

# MySQL 5.6 Reference Manual

---

# MySQL 5.6 Reference Manual

## Abstract

This is the MySQL™ Reference Manual. It documents MySQL 5.6 through 5.6.12.

MySQL Cluster is currently not supported in MySQL 5.6. For information about MySQL Cluster, please see [MySQL Cluster NDB 7.2](#).

**MySQL 5.6 features.** This manual describes features that are not included in every edition of MySQL 5.6; such features may not be included in the edition of MySQL 5.6 licensed to you. If you have any questions about the features included in your edition of MySQL 5.6, refer to your MySQL 5.6 license agreement or contact your Oracle sales representative.

For release notes detailing the changes in each release, see the [MySQL 5.6 Release Notes](#).

For legal information, see the [Legal Notices](#).

Document generated on: 2013-04-16 (revision: 34905)

General	Administrators	MySQL Enterprise	Developers & Functionality	Connectors & APIs	HA/Scalability
<a href="#">Tutorial</a>	<a href="#">Installation &amp; Upgrades</a>	<a href="#">MySQL Enterprise Edition</a>	<a href="#">MySQL Workbench</a>	<a href="#">Connectors and APIs</a>	<a href="#">» HA/Scalability Guide</a>
<a href="#">Server Administration</a>	<a href="#">» MySQL Installer</a>	<a href="#">MySQL Enterprise Monitor</a>	<a href="#">Globalization</a>	<a href="#">Connector/J</a>	<a href="#">MySQL and DRBD</a>
<a href="#">SQL Syntax</a>	<a href="#">» Security</a>	<a href="#">MySQL Enterprise Backup</a>	<a href="#">Optimization</a>	<a href="#">Connector/ODBC</a>	<a href="#">memcached</a>
<a href="#">Storage Engines</a>	<a href="#">» Startup / Shutdown</a>	<a href="#">MySQL Enterprise Security</a>	<a href="#">Functions and Operators</a>	<a href="#">Connector/Net</a>	<a href="#">memcached with InnoDB</a>
<a href="#">Server Option / Variable Reference</a>	<a href="#">» Backup and Recovery Overview</a>	<a href="#">MySQL Enterprise Audit</a>	<a href="#">Views and Stored Programs</a>	<a href="#">Connector/Python</a>	<a href="#">MySQL and Virtualization</a>
<a href="#">» Release Notes</a>	<a href="#">» MySQL Utilities</a>	<a href="#">MySQL Thread Pool</a>	<a href="#">Partitioning</a>	<a href="#">PHP</a>	<a href="#">MySQL Proxy</a>
<a href="#">» MySQL Version Reference</a>	<a href="#">» Linux/Unix Platform Guide</a>		<a href="#">Precision Math</a>	<a href="#">C API</a>	<a href="#">Replication</a>
<a href="#">FAQs</a>	<a href="#">» Windows Platform Guide</a>		<a href="#">Information Schema</a>	<a href="#">Connector/C</a>	<a href="#">Semisynchronous Replication</a>
	<a href="#">» Mac OS X Platform Guide</a>		<a href="#">Performance Schema</a>	<a href="#">Connector/C++</a>	
	<a href="#">» Solaris Platform Guide</a>		<a href="#">Spatial Extensions</a>	<a href="#">» MySQL for Excel</a>	
	<a href="#">» Building from Source</a>		<a href="#">Restrictions and Limitations</a>		

---



---

# Table of Contents

Preface and Legal Notices .....	xxi
1. General Information .....	1
1.1. About This Manual .....	2
1.2. Typographical and Syntax Conventions .....	2
1.3. Overview of the MySQL Database Management System .....	4
1.3.1. What is MySQL? .....	4
1.3.2. The Main Features of MySQL .....	6
1.3.3. History of MySQL .....	8
1.4. What Is New in MySQL 5.6 .....	9
1.5. MySQL Development History .....	17
1.6. MySQL Information Sources .....	18
1.6.1. MySQL Mailing Lists .....	18
1.6.2. MySQL Community Support at the MySQL Forums .....	20
1.6.3. MySQL Community Support on Internet Relay Chat (IRC) .....	21
1.6.4. MySQL Enterprise .....	21
1.7. How to Report Bugs or Problems .....	21
1.8. MySQL Standards Compliance .....	26
1.8.1. What Standards MySQL Follows .....	26
1.8.2. Selecting SQL Modes .....	26
1.8.3. Running MySQL in ANSI Mode .....	27
1.8.4. MySQL Extensions to Standard SQL .....	27
1.8.5. MySQL Differences from Standard SQL .....	30
1.8.6. How MySQL Deals with Constraints .....	34
1.9. Credits .....	38
1.9.1. Contributors to MySQL .....	38
1.9.2. Documenters and translators .....	42
1.9.3. Packages that support MySQL .....	44
1.9.4. Tools that were used to create MySQL .....	44
1.9.5. Supporters of MySQL .....	45
2. Installing and Upgrading MySQL .....	47
2.1. General Installation Guidance .....	49
2.1.1. Operating Systems Supported by MySQL Community Server .....	49
2.1.2. Choosing Which MySQL Distribution to Install .....	50
2.1.3. How to Get MySQL .....	53
2.1.4. Verifying Package Integrity Using MD5 Checksums or <a href="#">GnuPG</a> .....	54
2.1.5. Installation Layouts .....	62
2.1.6. Compiler-Specific Build Characteristics .....	62
2.2. Installing MySQL from Generic Binaries on Unix/Linux .....	62
2.3. Installing MySQL on Microsoft Windows .....	65
2.3.1. MySQL Installation Layout on Microsoft Windows .....	66
2.3.2. Choosing An Installation Package .....	67
2.3.3. Installing MySQL on Microsoft Windows Using MySQL Installer .....	67
2.3.4. Installing MySQL on Microsoft Windows Using a <a href="#">noinstall</a> Zip Archive .....	84
2.3.5. Troubleshooting a Microsoft Windows MySQL Server Installation .....	96
2.3.6. Upgrading MySQL on Windows .....	97
2.3.7. Windows Postinstallation Procedures .....	98
2.4. Installing MySQL on Mac OS X .....	100
2.4.1. General Notes on Installing MySQL on Mac OS X .....	100
2.4.2. Installing MySQL on Mac OS X Using Native Packages .....	101
2.4.3. Installing the MySQL Startup Item .....	104
2.4.4. Installing and Using the MySQL Preference Pane .....	107
2.4.5. Using the Bundled MySQL on Mac OS X Server .....	109
2.5. Installing MySQL on Linux .....	109
2.5.1. Installing MySQL from RPM Packages on Linux .....	110
2.5.2. Installing MySQL on Linux using Native Package Manager .....	114

2.6. Installing MySQL on Solaris and OpenSolaris .....	118
2.6.1. Installing MySQL on Solaris using a Solaris <code>PKG</code> .....	118
2.6.2. Installing MySQL on OpenSolaris using <code>IPS</code> .....	119
2.7. Installing MySQL on HP-UX .....	120
2.7.1. General Notes on Installing MySQL on HP-UX .....	120
2.7.2. Installing MySQL on HP-UX using <code>DEPOT</code> .....	121
2.8. Installing MySQL on FreeBSD .....	121
2.9. Installing MySQL from Source .....	122
2.9.1. MySQL Layout for Source Installation .....	124
2.9.2. Installing MySQL from a Standard Source Distribution .....	124
2.9.3. Installing MySQL from a Development Source Tree .....	127
2.9.4. MySQL Source-Configuration Options .....	129
2.9.5. Dealing with Problems Compiling MySQL .....	137
2.9.6. MySQL Configuration and Third-Party Tools .....	138
2.10. Postinstallation Setup and Testing .....	139
2.10.1. Unix Postinstallation Procedures .....	139
2.10.2. Securing the Initial MySQL Accounts .....	150
2.11. Upgrading or Downgrading MySQL .....	153
2.11.1. Upgrading MySQL .....	154
2.11.2. Downgrading MySQL .....	159
2.11.3. Checking Whether Tables or Indexes Must Be Rebuilt .....	161
2.11.4. Rebuilding or Repairing Tables or Indexes .....	163
2.11.5. Copying MySQL Databases to Another Machine .....	164
2.12. Environment Variables .....	165
2.13. Perl Installation Notes .....	167
2.13.1. Installing Perl on Unix .....	167
2.13.2. Installing ActiveState Perl on Windows .....	168
2.13.3. Problems Using the Perl <code>DBI/DBD</code> Interface .....	168
3. Tutorial .....	171
3.1. Connecting to and Disconnecting from the Server .....	171
3.2. Entering Queries .....	172
3.3. Creating and Using a Database .....	175
3.3.1. Creating and Selecting a Database .....	176
3.3.2. Creating a Table .....	177
3.3.3. Loading Data into a Table .....	178
3.3.4. Retrieving Information from a Table .....	179
3.4. Getting Information About Databases and Tables .....	191
3.5. Using <code>mysql</code> in Batch Mode .....	192
3.6. Examples of Common Queries .....	193
3.6.1. The Maximum Value for a Column .....	194
3.6.2. The Row Holding the Maximum of a Certain Column .....	194
3.6.3. Maximum of Column per Group .....	195
3.6.4. The Rows Holding the Group-wise Maximum of a Certain Column .....	195
3.6.5. Using User-Defined Variables .....	196
3.6.6. Using Foreign Keys .....	196
3.6.7. Searching on Two Keys .....	197
3.6.8. Calculating Visits Per Day .....	198
3.6.9. Using <code>AUTO_INCREMENT</code> .....	198
3.7. Using MySQL with Apache .....	200
4. MySQL Programs .....	203
4.1. Overview of MySQL Programs .....	204
4.2. Using MySQL Programs .....	208
4.2.1. Invoking MySQL Programs .....	208
4.2.2. Connecting to the MySQL Server .....	209
4.2.3. Specifying Program Options .....	212
4.2.4. Setting Environment Variables .....	224
4.3. MySQL Server and Server-Startup Programs .....	225
4.3.1. <code>mysqld</code> — The MySQL Server .....	225

4.3.2. <code>mysqld_safe</code> — MySQL Server Startup Script .....	225
4.3.3. <code>mysql.server</code> — MySQL Server Startup Script .....	230
4.3.4. <code>mysqld_multi</code> — Manage Multiple MySQL Servers .....	231
4.4. MySQL Installation-Related Programs .....	235
4.4.1. <code>comp_err</code> — Compile MySQL Error Message File .....	235
4.4.2. <code>mysqlbug</code> — Generate Bug Report .....	236
4.4.3. <code>mysql_install_db</code> — Initialize MySQL Data Directory .....	236
4.4.4. <code>mysql_plugin</code> — Configure MySQL Server Plugins .....	238
4.4.5. <code>mysql_secure_installation</code> — Improve MySQL Installation Security .....	240
4.4.6. <code>mysql_tzinfo_to_sql</code> — Load the Time Zone Tables .....	241
4.4.7. <code>mysql_upgrade</code> — Check and Upgrade MySQL Tables .....	241
4.5. MySQL Client Programs .....	244
4.5.1. <code>mysql</code> — The MySQL Command-Line Tool .....	244
4.5.2. <code>mysqladmin</code> — Client for Administering a MySQL Server .....	266
4.5.3. <code>mysqlcheck</code> — A Table Maintenance Program .....	272
4.5.4. <code>mysqldump</code> — A Database Backup Program .....	279
4.5.5. <code>mysqlimport</code> — A Data Import Program .....	297
4.5.6. <code>mysqlshow</code> — Display Database, Table, and Column Information .....	302
4.5.7. <code>mysqlslap</code> — Load Emulation Client .....	306
4.6. MySQL Administrative and Utility Programs .....	313
4.6.1. <code>innochecksum</code> — Offline InnoDB File Checksum Utility .....	313
4.6.2. <code>myisam_ftdump</code> — Display Full-Text Index information .....	314
4.6.3. <code>myisamchk</code> — MyISAM Table-Maintenance Utility .....	315
4.6.4. <code>myisamlog</code> — Display MyISAM Log File Contents .....	331
4.6.5. <code>myisampack</code> — Generate Compressed, Read-Only MyISAM Tables .....	332
4.6.6. <code>mysql_config_editor</code> — MySQL Configuration Utility .....	337
4.6.7. <code>mysqlaccess</code> — Client for Checking Access Privileges .....	343
4.6.8. <code>mysqlbinlog</code> — Utility for Processing Binary Log Files .....	345
4.6.9. <code>mysqldumpslow</code> — Summarize Slow Query Log Files .....	364
4.6.10. <code>mysqlhotcopy</code> — A Database Backup Program .....	365
4.6.11. <code>mysql_convert_table_format</code> — Convert Tables to Use a Given Storage Engine .....	368
4.6.12. <code>mysql_find_rows</code> — Extract SQL Statements from Files .....	369
4.6.13. <code>mysql_fix_extensions</code> — Normalize Table File Name Extensions .....	370
4.6.14. <code>mysql_setpermission</code> — Interactively Set Permissions in Grant Tables .....	370
4.6.15. <code>mysql_waitpid</code> — Kill Process and Wait for Its Termination .....	371
4.6.16. <code>mysql_zap</code> — Kill Processes That Match a Pattern .....	371
4.7. MySQL Program Development Utilities .....	372
4.7.1. <code>msql2mysql</code> — Convert mSQL Programs for Use with MySQL .....	372
4.7.2. <code>mysql_config</code> — Display Options for Compiling Clients .....	373
4.7.3. <code>my_print_defaults</code> — Display Options from Option Files .....	374
4.7.4. <code>resolve_stack_dump</code> — Resolve Numeric Stack Trace Dump to Symbols .....	375
4.8. Miscellaneous Programs .....	375
4.8.1. <code>perror</code> — Explain Error Codes .....	375
4.8.2. <code>replace</code> — A String-Replacement Utility .....	376
4.8.3. <code>resolveip</code> — Resolve Host name to IP Address or Vice Versa .....	377
5. MySQL Server Administration .....	379
5.1. The MySQL Server .....	380
5.1.1. Server Option and Variable Reference .....	380
5.1.2. Server Configuration Defaults .....	410
5.1.3. Server Command Options .....	412
5.1.4. Server System Variables .....	447
5.1.5. Using System Variables .....	564
5.1.6. Server Status Variables .....	576
5.1.7. Server SQL Modes .....	602
5.1.8. Server Plugins .....	609
5.1.9. IPv6 Support .....	613
5.1.10. Server-Side Help .....	617

5.1.11. Server Response to Signals .....	617
5.1.12. The Shutdown Process .....	618
5.2. MySQL Server Logs .....	619
5.2.1. Selecting General Query and Slow Query Log Output Destinations .....	620
5.2.2. The Error Log .....	622
5.2.3. The General Query Log .....	623
5.2.4. The Binary Log .....	624
5.2.5. The Slow Query Log .....	635
5.2.6. Server Log Maintenance .....	637
5.3. Managing Disk I/O and File Space for <a href="#">InnoDB</a> Tables .....	638
5.3.1. <a href="#">InnoDB</a> Disk I/O .....	638
5.3.2. File Space Management .....	639
5.3.3. <a href="#">InnoDB</a> Checkpoints .....	640
5.3.4. Defragmenting a Table .....	640
5.4. Creating and Using <a href="#">InnoDB</a> Tables and Indexes .....	640
5.4.1. Managing <a href="#">InnoDB</a> Tablespace .....	641
5.4.2. Grouping DML Operations with Transactions .....	645
5.4.3. Converting Tables from <a href="#">MyISAM</a> to <a href="#">InnoDB</a> .....	645
5.4.4. <a href="#">AUTO_INCREMENT</a> Handling in <a href="#">InnoDB</a> .....	650
5.4.5. <a href="#">InnoDB</a> and <a href="#">FOREIGN KEY</a> Constraints .....	655
5.4.6. Working with <a href="#">InnoDB</a> Compressed Tables .....	656
5.4.7. <a href="#">InnoDB</a> File-Format Management .....	667
5.4.8. How <a href="#">InnoDB</a> Stores Variable-Length Columns .....	672
5.5. Online DDL for <a href="#">InnoDB</a> Tables .....	674
5.5.1. Overview of Online DDL .....	674
5.5.2. Performance and Concurrency Considerations for Online DDL .....	679
5.5.3. SQL Syntax for Online DDL .....	681
5.5.4. Combining or Separating DDL Statements .....	682
5.5.5. Examples of Online DDL .....	682
5.5.6. Implementation Details of Online DDL .....	702
5.5.7. How Crash Recovery Works with Online DDL .....	703
5.5.8. DDL for Partitioned Tables .....	704
5.5.9. Limitations of Online DDL .....	704
5.6. Running Multiple MySQL Instances on One Machine .....	705
5.6.1. Setting Up Multiple Data Directories .....	706
5.6.2. Running Multiple MySQL Instances on Windows .....	707
5.6.3. Running Multiple MySQL Instances on Unix .....	710
5.6.4. Using Client Programs in a Multiple-Server Environment .....	711
5.7. Tracing <a href="#">mysqld</a> Using DTrace .....	711
5.7.1. <a href="#">mysqld</a> DTrace Probe Reference .....	713
6. Security .....	731
6.1. General Security Issues .....	732
6.1.1. Security Guidelines .....	732
6.1.2. Keeping Passwords Secure .....	733
6.1.3. Making MySQL Secure Against Attackers .....	746
6.1.4. Security-Related <a href="#">mysqld</a> Options and Variables .....	748
6.1.5. How to Run MySQL as a Normal User .....	748
6.1.6. Security Issues with <a href="#">LOAD DATA LOCAL</a> .....	749
6.1.7. Client Programming Security Guidelines .....	750
6.2. The MySQL Access Privilege System .....	751
6.2.1. Privileges Provided by MySQL .....	752
6.2.2. Privilege System Grant Tables .....	756
6.2.3. Specifying Account Names .....	762
6.2.4. Access Control, Stage 1: Connection Verification .....	763
6.2.5. Access Control, Stage 2: Request Verification .....	766
6.2.6. When Privilege Changes Take Effect .....	768
6.2.7. Causes of Access-Denied Errors .....	768
6.3. MySQL User Account Management .....	773

6.3.1. User Names and Passwords .....	773
6.3.2. Adding User Accounts .....	774
6.3.3. Removing User Accounts .....	778
6.3.4. Setting Account Resource Limits .....	778
6.3.5. Assigning Account Passwords .....	780
6.3.6. Password Expiration and Sandbox Mode .....	781
6.3.7. Pluggable Authentication .....	784
6.3.8. Proxy Users .....	803
6.3.9. Using SSL for Secure Connections .....	806
6.3.10. Connecting to MySQL Remotely from Windows with SSH .....	817
6.3.11. MySQL Enterprise Audit Log Plugin .....	817
6.3.12. SQL-Based MySQL Account Activity Auditing .....	827
7. Backup and Recovery .....	831
7.1. Backup and Recovery Types .....	832
7.2. Database Backup Methods .....	835
7.3. Example Backup and Recovery Strategy .....	837
7.3.1. Establishing a Backup Policy .....	838
7.3.2. Using Backups for Recovery .....	839
7.3.3. Backup Strategy Summary .....	840
7.4. Using <code>mysqldump</code> for Backups .....	840
7.4.1. Dumping Data in SQL Format with <code>mysqldump</code> .....	841
7.4.2. Reloading SQL-Format Backups .....	841
7.4.3. Dumping Data in Delimited-Text Format with <code>mysqldump</code> .....	842
7.4.4. Reloading Delimited-Text Format Backups .....	843
7.4.5. <code>mysqldump</code> Tips .....	844
7.5. Point-in-Time (Incremental) Recovery Using the Binary Log .....	845
7.5.1. Point-in-Time Recovery Using Event Times .....	847
7.5.2. Point-in-Time Recovery Using Event Positions .....	847
7.6. <code>MyISAM</code> Table Maintenance and Crash Recovery .....	848
7.6.1. Using <code>myisamchk</code> for Crash Recovery .....	848
7.6.2. How to Check <code>MyISAM</code> Tables for Errors .....	849
7.6.3. How to Repair <code>MyISAM</code> Tables .....	850
7.6.4. <code>MyISAM</code> Table Optimization .....	852
7.6.5. Setting Up a <code>MyISAM</code> Table Maintenance Schedule .....	853
8. Optimization .....	855
8.1. Optimization Overview .....	856
8.2. Optimizing SQL Statements .....	858
8.2.1. Optimizing <code>SELECT</code> Statements .....	858
8.2.2. Optimizing DML Statements .....	862
8.2.3. Optimizing Database Privileges .....	863
8.2.4. Optimizing <code>INFORMATION_SCHEMA</code> Queries .....	864
8.2.5. Other Optimization Tips .....	868
8.3. Optimization and Indexes .....	871
8.3.1. How MySQL Uses Indexes .....	871
8.3.2. Using Primary Keys .....	872
8.3.3. Using Foreign Keys .....	872
8.3.4. Column Indexes .....	872
8.3.5. Multiple-Column Indexes .....	873
8.3.6. Verifying Index Usage .....	875
8.3.7. <code>InnoDB</code> and <code>MyISAM</code> Index Statistics Collection .....	875
8.3.8. Comparison of B-Tree and Hash Indexes .....	876
8.4. Optimizing Database Structure .....	877
8.4.1. Optimizing Data Size .....	878
8.4.2. Optimizing MySQL Data Types .....	879
8.4.3. Optimizing for Many Tables .....	881
8.5. Optimizing for <code>InnoDB</code> Tables .....	883
8.5.1. Optimizing Storage Layout for <code>InnoDB</code> Tables .....	883
8.5.2. Optimizing <code>InnoDB</code> Transaction Management .....	884

8.5.3. Optimizing <code>InnoDB</code> Logging .....	885
8.5.4. Bulk Data Loading for <code>InnoDB</code> Tables .....	885
8.5.5. Optimizing <code>InnoDB</code> Queries .....	886
8.5.6. Optimizing <code>InnoDB</code> DDL Operations .....	886
8.5.7. Optimizing <code>InnoDB</code> Disk I/O .....	887
8.5.8. Optimizing <code>InnoDB</code> Configuration Variables .....	888
8.5.9. Optimizing <code>InnoDB</code> for Systems with Many Tables .....	889
8.6. Optimizing for <code>MyISAM</code> Tables .....	889
8.6.1. Optimizing <code>MyISAM</code> Queries .....	889
8.6.2. Bulk Data Loading for <code>MyISAM</code> Tables .....	891
8.6.3. Speed of <code>REPAIR TABLE</code> Statements .....	892
8.7. Optimizing for <code>MEMORY</code> Tables .....	893
8.8. Understanding the Query Execution Plan .....	894
8.8.1. Optimizing Queries with <code>EXPLAIN</code> .....	894
8.8.2. <code>EXPLAIN</code> Output Format .....	895
8.8.3. <code>EXPLAIN EXTENDED</code> Output Format .....	906
8.8.4. Estimating Query Performance .....	908
8.8.5. Controlling the Query Optimizer .....	908
8.9. Buffering and Caching .....	911
8.9.1. The <code>InnoDB</code> Buffer Pool .....	911
8.9.2. The <code>MyISAM</code> Key Cache .....	914
8.9.3. The MySQL Query Cache .....	918
8.9.4. Caching of Prepared Statements and Stored Programs .....	924
8.10. Optimizing Locking Operations .....	925
8.10.1. Internal Locking Methods .....	925
8.10.2. Table Locking Issues .....	927
8.10.3. Concurrent Inserts .....	929
8.10.4. Metadata Locking .....	929
8.10.5. External Locking .....	930
8.11. Optimizing the MySQL Server .....	931
8.11.1. System Factors and Startup Parameter Tuning .....	931
8.11.2. Tuning Server Parameters .....	932
8.11.3. Optimizing Disk I/O .....	939
8.11.4. Optimizing Memory Use .....	943
8.11.5. Optimizing Network Use .....	946
8.11.6. The Thread Pool Plugin .....	949
8.12. Measuring Performance (Benchmarking) .....	954
8.12.1. Measuring the Speed of Expressions and Functions .....	955
8.12.2. The MySQL Benchmark Suite .....	955
8.12.3. Using Your Own Benchmarks .....	956
8.12.4. Measuring Performance with <code>performance_schema</code> .....	956
8.12.5. Examining Thread Information .....	956
8.13. Internal Details of MySQL Optimizations .....	969
8.13.1. Range Optimization .....	970
8.13.2. Index Merge Optimization .....	973
8.13.3. Engine Condition Pushdown Optimization .....	976
8.13.4. Index Condition Pushdown Optimization .....	977
8.13.5. Use of Index Extensions .....	978
8.13.6. <code>IS NULL</code> Optimization .....	981
8.13.7. <code>LEFT JOIN</code> and <code>RIGHT JOIN</code> Optimization .....	981
8.13.8. Nested-Loop Join Algorithms .....	982
8.13.9. Nested Join Optimization .....	984
8.13.10. Outer Join Simplification .....	989
8.13.11. Multi-Range Read Optimization .....	991
8.13.12. Block Nested-Loop and Batched Key Access Joins .....	993
8.13.13. <code>ORDER BY</code> Optimization .....	995
8.13.14. <code>GROUP BY</code> Optimization .....	998
8.13.15. <code>DISTINCT</code> Optimization .....	1000



---

8.13.16. Subquery Optimization .....	1001
9. Language Structure .....	1011
9.1. Literal Values .....	1011
9.1.1. String Literals .....	1011
9.1.2. Number Literals .....	1013
9.1.3. Date and Time Literals .....	1014
9.1.4. Hexadecimal Literals .....	1016
9.1.5. Boolean Literals .....	1016
9.1.6. Bit-Field Literals .....	1016
9.1.7. <code>NULL</code> Values .....	1017
9.2. Schema Object Names .....	1017
9.2.1. Identifier Qualifiers .....	1019
9.2.2. Identifier Case Sensitivity .....	1020
9.2.3. Mapping of Identifiers to File Names .....	1021
9.2.4. Function Name Parsing and Resolution .....	1023
9.3. Reserved Words .....	1026
9.4. User-Defined Variables .....	1029
9.5. Expression Syntax .....	1032
9.6. Comment Syntax .....	1034
10. Globalization .....	1037
10.1. Character Set Support .....	1037
10.1.1. Character Sets and Collations in General .....	1038
10.1.2. Character Sets and Collations in MySQL .....	1039
10.1.3. Specifying Character Sets and Collations .....	1040
10.1.4. Connection Character Sets and Collations .....	1046
10.1.5. Configuring the Character Set and Collation for Applications .....	1049
10.1.6. Character Set for Error Messages .....	1050
10.1.7. Collation Issues .....	1051
10.1.8. String Repertoire .....	1060
10.1.9. Operations Affected by Character Set Support .....	1061
10.1.10. Unicode Support .....	1064
10.1.11. Upgrading from Previous to Current Unicode Support .....	1068
10.1.12. UTF-8 for Metadata .....	1070
10.1.13. Column Character Set Conversion .....	1071
10.1.14. Character Sets and Collations That MySQL Supports .....	1072
10.2. Setting the Error Message Language .....	1085
10.3. Adding a Character Set .....	1086
10.3.1. Character Definition Arrays .....	1088
10.3.2. String Collating Support for Complex Character Sets .....	1089
10.3.3. Multi-Byte Character Support for Complex Character Sets .....	1089
10.4. Adding a Collation to a Character Set .....	1089
10.4.1. Collation Implementation Types .....	1090
10.4.2. Choosing a Collation ID .....	1093
10.4.3. Adding a Simple Collation to an 8-Bit Character Set .....	1094
10.4.4. Adding a UCA Collation to a Unicode Character Set .....	1095
10.5. Character Set Configuration .....	1101
10.6. MySQL Server Time Zone Support .....	1102
10.6.1. Staying Current with Time Zone Changes .....	1104
10.6.2. Time Zone Leap Second Support .....	1105
10.7. MySQL Server Locale Support .....	1107
11. Data Types .....	1111
11.1. Data Type Overview .....	1112
11.1.1. Numeric Type Overview .....	1112
11.1.2. Date and Time Type Overview .....	1115
11.1.3. String Type Overview .....	1117
11.2. Numeric Types .....	1120
11.2.1. Integer Types (Exact Value) - <code>INTEGER</code> , <code>INT</code> , <code>SMALLINT</code> , <code>TINYINT</code> , <code>MEDIUMINT</code> , <code>BIGINT</code> .....	1121

---

11.2.2. Fixed-Point Types (Exact Value) - <code>DECIMAL</code> , <code>NUMERIC</code> .....	1121
11.2.3. Floating-Point Types (Approximate Value) - <code>FLOAT</code> , <code>DOUBLE</code> .....	1122
11.2.4. Bit-Value Type - <code>BIT</code> .....	1122
11.2.5. Numeric Type Attributes .....	1122
11.2.6. Out-of-Range and Overflow Handling .....	1123
11.3. Date and Time Types .....	1124
11.3.1. The <code>DATE</code> , <code>DATETIME</code> , and <code>TIMESTAMP</code> Types .....	1126
11.3.2. The <code>TIME</code> Type .....	1127
11.3.3. The <code>YEAR</code> Type .....	1128
11.3.4. <code>YEAR(2)</code> Limitations and Migrating to <code>YEAR(4)</code> .....	1128
11.3.5. Automatic Initialization and Updating for <code>TIMESTAMP</code> and <code>DATETIME</code> .....	1131
11.3.6. Fractional Seconds in Time Values .....	1134
11.3.7. Conversion Between Date and Time Types .....	1135
11.3.8. Two-Digit Years in Dates .....	1137
11.4. String Types .....	1137
11.4.1. The <code>CHAR</code> and <code>VARCHAR</code> Types .....	1137
11.4.2. The <code>BINARY</code> and <code>VARBINARY</code> Types .....	1139
11.4.3. The <code>BLOB</code> and <code>TEXT</code> Types .....	1140
11.4.4. The <code>ENUM</code> Type .....	1141
11.4.5. The <code>SET</code> Type .....	1144
11.5. Data Type Default Values .....	1146
11.6. Data Type Storage Requirements .....	1147
11.7. Choosing the Right Type for a Column .....	1151
11.8. Using Data Types from Other Database Engines .....	1151
12. Functions and Operators .....	1153
12.1. Function and Operator Reference .....	1154
12.2. Type Conversion in Expression Evaluation .....	1161
12.3. Operators .....	1164
12.3.1. Operator Precedence .....	1165
12.3.2. Comparison Functions and Operators .....	1165
12.3.3. Logical Operators .....	1171
12.3.4. Assignment Operators .....	1172
12.4. Control Flow Functions .....	1173
12.5. String Functions .....	1175
12.5.1. String Comparison Functions .....	1190
12.5.2. Regular Expressions .....	1193
12.6. Numeric Functions and Operators .....	1198
12.6.1. Arithmetic Operators .....	1199
12.6.2. Mathematical Functions .....	1201
12.7. Date and Time Functions .....	1209
12.8. What Calendar Is Used By MySQL? .....	1229
12.9. Full-Text Search Functions .....	1229
12.9.1. Natural Language Full-Text Searches .....	1231
12.9.2. Boolean Full-Text Searches .....	1234
12.9.3. Full-Text Searches with Query Expansion .....	1236
12.9.4. Full-Text Stopwords .....	1237
12.9.5. Full-Text Restrictions .....	1240
12.9.6. Fine-Tuning MySQL Full-Text Search .....	1241
12.9.7. Adding a Collation for Full-Text Indexing .....	1243
12.10. Cast Functions and Operators .....	1244
12.11. XML Functions .....	1247
12.12. Bit Functions .....	1257
12.13. Encryption and Compression Functions .....	1259
12.14. Information Functions .....	1265
12.15. GTID Functions .....	1273
12.16. Miscellaneous Functions .....	1274
12.17. Functions and Modifiers for Use with <code>GROUP BY</code> Clauses .....	1281
12.17.1. <code>GROUP BY</code> (Aggregate) Functions .....	1281



12.17.2. <code>GROUP BY</code> Modifiers .....	1285
12.17.3. MySQL Extensions to <code>GROUP BY</code> .....	1288
12.18. Spatial Extensions .....	1289
12.18.1. Introduction to MySQL Spatial Support .....	1289
12.18.2. The OpenGIS Geometry Model .....	1290
12.18.3. Supported Spatial Data Formats .....	1296
12.18.4. Creating a Spatially Enabled MySQL Database .....	1297
12.18.5. Spatial Analysis Functions .....	1302
12.18.6. Optimizing Spatial Analysis .....	1312
12.18.7. MySQL Conformance and Compatibility .....	1315
12.19. Precision Math .....	1315
12.19.1. Types of Numeric Values .....	1316
12.19.2. <code>DECIMAL</code> Data Type Changes .....	1316
12.19.3. Expression Handling .....	1317
12.19.4. Rounding Behavior .....	1319
12.19.5. Precision Math Examples .....	1319
13. SQL Statement Syntax .....	1325
13.1. Data Definition Statements .....	1326
13.1.1. <code>ALTER DATABASE</code> Syntax .....	1326
13.1.2. <code>ALTER EVENT</code> Syntax .....	1327
13.1.3. <code>ALTER FUNCTION</code> Syntax .....	1329
13.1.4. <code>ALTER PROCEDURE</code> Syntax .....	1329
13.1.5. <code>ALTER SERVER</code> Syntax .....	1329
13.1.6. <code>ALTER TABLE</code> Syntax .....	1330
13.1.7. <code>ALTER VIEW</code> Syntax .....	1343
13.1.8. <code>CREATE DATABASE</code> Syntax .....	1344
13.1.9. <code>CREATE EVENT</code> Syntax .....	1344
13.1.10. <code>CREATE FUNCTION</code> Syntax .....	1349
13.1.11. <code>CREATE INDEX</code> Syntax .....	1349
13.1.12. <code>CREATE PROCEDURE</code> and <code>CREATE FUNCTION</code> Syntax .....	1352
13.1.13. <code>CREATE SERVER</code> Syntax .....	1357
13.1.14. <code>CREATE TABLE</code> Syntax .....	1358
13.1.15. <code>CREATE TRIGGER</code> Syntax .....	1384
13.1.16. <code>CREATE VIEW</code> Syntax .....	1388
13.1.17. <code>DROP DATABASE</code> Syntax .....	1391
13.1.18. <code>DROP EVENT</code> Syntax .....	1392
13.1.19. <code>DROP FUNCTION</code> Syntax .....	1392
13.1.20. <code>DROP INDEX</code> Syntax .....	1392
13.1.21. <code>DROP PROCEDURE</code> and <code>DROP FUNCTION</code> Syntax .....	1393
13.1.22. <code>DROP SERVER</code> Syntax .....	1393
13.1.23. <code>DROP TABLE</code> Syntax .....	1393
13.1.24. <code>DROP TRIGGER</code> Syntax .....	1394
13.1.25. <code>DROP VIEW</code> Syntax .....	1394
13.1.26. <code>RENAME TABLE</code> Syntax .....	1394
13.1.27. <code>TRUNCATE TABLE</code> Syntax .....	1395
13.2. Data Manipulation Statements .....	1396
13.2.1. <code>CALL</code> Syntax .....	1396
13.2.2. <code>DELETE</code> Syntax .....	1398
13.2.3. <code>DO</code> Syntax .....	1402
13.2.4. <code>HANDLER</code> Syntax .....	1402
13.2.5. <code>INSERT</code> Syntax .....	1403
13.2.6. <code>LOAD DATA INFILE</code> Syntax .....	1412
13.2.7. <code>LOAD XML</code> Syntax .....	1420
13.2.8. <code>REPLACE</code> Syntax .....	1426
13.2.9. <code>SELECT</code> Syntax .....	1427
13.2.10. Subquery Syntax .....	1445
13.2.11. <code>UPDATE</code> Syntax .....	1455
13.3. MySQL Transactional and Locking Statements .....	1458

13.3.1. <code>START TRANSACTION</code> , <code>COMMIT</code> , and <code>ROLLBACK</code> Syntax .....	1458
13.3.2. Statements That Cannot Be Rolled Back .....	1460
13.3.3. Statements That Cause an Implicit Commit .....	1461
13.3.4. <code>SAVEPOINT</code> , <code>ROLLBACK TO SAVEPOINT</code> , and <code>RELEASE SAVEPOINT</code> Syntax .....	1462
13.3.5. <code>LOCK TABLES</code> and <code>UNLOCK TABLES</code> Syntax .....	1462
13.3.6. <code>SET TRANSACTION</code> Syntax .....	1468
13.3.7. XA Transactions .....	1470
13.4. Replication Statements .....	1474
13.4.1. SQL Statements for Controlling Master Servers .....	1474
13.4.2. SQL Statements for Controlling Slave Servers .....	1476
13.5. SQL Syntax for Prepared Statements .....	1486
13.5.1. <code>PREPARE</code> Syntax .....	1489
13.5.2. <code>EXECUTE</code> Syntax .....	1490
13.5.3. <code>DEALLOCATE PREPARE</code> Syntax .....	1490
13.6. MySQL Compound-Statement Syntax .....	1490
13.6.1. <code>BEGIN . . . END</code> Compound-Statement Syntax .....	1490
13.6.2. Statement Label Syntax .....	1491
13.6.3. <code>DECLARE</code> Syntax .....	1492
13.6.4. Variables in Stored Programs .....	1492
13.6.5. Flow Control Statements .....	1493
13.6.6. Cursors .....	1498
13.6.7. Condition Handling .....	1499
13.7. Database Administration Statements .....	1521
13.7.1. Account Management Statements .....	1521
13.7.2. Table Maintenance Statements .....	1536
13.7.3. Plugin and User-Defined Function Statements .....	1544
13.7.4. <code>SET</code> Syntax .....	1547
13.7.5. <code>SHOW</code> Syntax .....	1549
13.7.6. Other Administrative Statements .....	1588
13.8. MySQL Utility Statements .....	1596
13.8.1. <code>DESCRIBE</code> Syntax .....	1596
13.8.2. <code>EXPLAIN</code> Syntax .....	1596
13.8.3. <code>HELP</code> Syntax .....	1598
13.8.4. <code>USE</code> Syntax .....	1599
14. Storage Engines .....	1601
14.1. Setting the Storage Engine .....	1604
14.2. The <code>InnoDB</code> Storage Engine .....	1605
14.2.1. Getting Started with <code>InnoDB</code> Tables .....	1606
14.2.2. Administering <code>InnoDB</code> .....	1613
14.2.3. <code>InnoDB</code> Concepts and Architecture .....	1622
14.2.4. <code>InnoDB</code> Performance Tuning and Troubleshooting .....	1644
14.2.5. <code>InnoDB</code> Features for Flexibility, Ease of Use and Reliability .....	1686
14.2.6. <code>InnoDB</code> Startup Options and System Variables .....	1692
14.2.7. Limits on <code>InnoDB</code> Tables .....	1756
14.2.8. MySQL and the ACID Model .....	1759
14.2.9. <code>InnoDB</code> Integration with memcached .....	1761
14.3. The <code>MyISAM</code> Storage Engine .....	1787
14.3.1. <code>MyISAM</code> Startup Options .....	1789
14.3.2. Space Needed for Keys .....	1791
14.3.3. <code>MyISAM</code> Table Storage Formats .....	1791
14.3.4. <code>MyISAM</code> Table Problems .....	1794
14.4. The <code>MEMORY</code> Storage Engine .....	1795
14.5. The <code>CSV</code> Storage Engine .....	1799
14.5.1. Repairing and Checking CSV Tables .....	1800
14.5.2. CSV Limitations .....	1801
14.6. The <code>ARCHIVE</code> Storage Engine .....	1801
14.7. The <code>BLACKHOLE</code> Storage Engine .....	1802

14.8. The <code>MERGE</code> Storage Engine .....	1804
14.8.1. <code>MERGE</code> Table Advantages and Disadvantages .....	1806
14.8.2. <code>MERGE</code> Table Problems .....	1807
14.9. The <code>FEDERATED</code> Storage Engine .....	1809
14.9.1. <code>FEDERATED</code> Storage Engine Overview .....	1809
14.9.2. How to Create <code>FEDERATED</code> Tables .....	1810
14.9.3. <code>FEDERATED</code> Storage Engine Notes and Tips .....	1813
14.9.4. <code>FEDERATED</code> Storage Engine Resources .....	1814
14.10. The <code>EXAMPLE</code> Storage Engine .....	1814
14.11. Other Storage Engines .....	1815
14.12. Overview of MySQL Storage Engine Architecture .....	1815
14.12.1. Pluggable Storage Engine Architecture .....	1815
14.12.2. The Common Database Server Layer .....	1816
15. High Availability and Scalability .....	1819
15.1. Oracle VM Template for MySQL Enterprise Edition .....	1822
15.2. Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux .....	1822
15.3. Overview of MySQL with Windows Failover Clustering .....	1825
15.4. Using MySQL within an Amazon EC2 Instance .....	1827
15.4.1. Setting Up MySQL on an EC2 AMI .....	1828
15.4.2. EC2 Instance Limitations .....	1829
15.4.3. Deploying a MySQL Database Using EC2 .....	1829
15.5. Using ZFS Replication .....	1832
15.5.1. Using ZFS for File System Replication .....	1833
15.5.2. Configuring MySQL for ZFS Replication .....	1834
15.5.3. Handling MySQL Recovery with ZFS .....	1835
15.6. Using MySQL with <code>memcached</code> .....	1835
15.6.1. Installing <code>memcached</code> .....	1837
15.6.2. Using <code>memcached</code> .....	1838
15.6.3. Developing a <code>memcached</code> Application .....	1856
15.6.4. Getting <code>memcached</code> Statistics .....	1880
15.6.5. <code>memcached</code> FAQ .....	1888
15.7. MySQL Proxy .....	1891
15.7.1. MySQL Proxy Supported Platforms .....	1892
15.7.2. Installing MySQL Proxy .....	1892
15.7.3. MySQL Proxy Command Options .....	1895
15.7.4. MySQL Proxy Scripting .....	1904
15.7.5. Using MySQL Proxy .....	1918
15.7.6. MySQL Proxy FAQ .....	1923
16. Replication .....	1929
16.1. Replication Configuration .....	1930
16.1.1. How to Set Up Replication .....	1931
16.1.2. Replication Formats .....	1940
16.1.3. Replication with Global Transaction Identifiers .....	1947
16.1.4. Replication and Binary Logging Options and Variables .....	1955
16.1.5. Common Replication Administration Tasks .....	2024
16.2. Replication Implementation .....	2026
16.2.1. Replication Implementation Details .....	2027
16.2.2. Replication Relay and Status Logs .....	2028
16.2.3. How Servers Evaluate Replication Filtering Rules .....	2034
16.3. Replication Solutions .....	2040
16.3.1. Using Replication for Backups .....	2041
16.3.2. Using Replication with Different Master and Slave Storage Engines .....	2044
16.3.3. Using Replication for Scale-Out .....	2045
16.3.4. Replicating Different Databases to Different Slaves .....	2046
16.3.5. Improving Replication Performance .....	2047
16.3.6. Switching Masters During Failover .....	2048
16.3.7. Setting Up Replication Using SSL .....	2051
16.3.8. Semisynchronous Replication .....	2052

16.3.9. Delayed Replication .....	2057
16.4. Replication Notes and Tips .....	2057
16.4.1. Replication Features and Issues .....	2057
16.4.2. Replication Compatibility Between MySQL Versions .....	2080
16.4.3. Upgrading a Replication Setup .....	2081
16.4.4. Troubleshooting Replication .....	2082
16.4.5. How to Report Replication Bugs or Problems .....	2083
17. Partitioning .....	2085
17.1. Overview of Partitioning in MySQL .....	2086
17.2. Partitioning Types .....	2089
17.2.1. RANGE Partitioning .....	2091
17.2.2. LIST Partitioning .....	2094
17.2.3. COLUMNS Partitioning .....	2097
17.2.4. HASH Partitioning .....	2104
17.2.5. KEY Partitioning .....	2107
17.2.6. Subpartitioning .....	2108
17.2.7. How MySQL Partitioning Handles NULL .....	2111
17.3. Partition Management .....	2115
17.3.1. Management of RANGE and LIST Partitions .....	2116
17.3.2. Management of HASH and KEY Partitions .....	2121
17.3.3. Exchanging Partitions and Subpartitions with Tables .....	2122
17.3.4. Maintenance of Partitions .....	2127
17.3.5. Obtaining Information About Partitions .....	2128
17.4. Partition Pruning .....	2130
17.5. Partition Selection .....	2133
17.6. Restrictions and Limitations on Partitioning .....	2139
17.6.1. Partitioning Keys, Primary Keys, and Unique Keys .....	2144
17.6.2. Partitioning Limitations Relating to Storage Engines .....	2147
17.6.3. Partitioning Limitations Relating to Functions .....	2148
17.6.4. Partitioning and Locking .....	2149
18. Stored Programs and Views .....	2151
18.1. Defining Stored Programs .....	2152
18.2. Using Stored Routines (Procedures and Functions) .....	2153
18.2.1. Stored Routine Syntax .....	2153
18.2.2. Stored Routines and MySQL Privileges .....	2154
18.2.3. Stored Routine Metadata .....	2154
18.2.4. Stored Procedures, Functions, Triggers, and LAST_INSERT_ID() .....	2155
18.3. Using Triggers .....	2155
18.3.1. Trigger Syntax .....	2156
18.3.2. Trigger Metadata .....	2158
18.4. Using the Event Scheduler .....	2158
18.4.1. Event Scheduler Overview .....	2158
18.4.2. Event Scheduler Configuration .....	2159
18.4.3. Event Syntax .....	2161
18.4.4. Event Metadata .....	2161
18.4.5. Event Scheduler Status .....	2162
18.4.6. The Event Scheduler and MySQL Privileges .....	2163
18.5. Using Views .....	2165
18.5.1. View Syntax .....	2166
18.5.2. View Processing Algorithms .....	2166
18.5.3. Updatable and Insertable Views .....	2167
18.5.4. View Metadata .....	2169
18.6. Access Control for Stored Programs and Views .....	2169
18.7. Binary Logging of Stored Programs .....	2171
19. INFORMATION_SCHEMA Tables .....	2177
19.1. The INFORMATION_SCHEMA CHARACTER_SETS Table .....	2180
19.2. The INFORMATION_SCHEMA COLLATIONS Table .....	2180

19.3. The <code>INFORMATION_SCHEMA</code> <code>COLLATION_CHARACTER_SET_APPLICABILITY</code> Table .....	2181
19.4. The <code>INFORMATION_SCHEMA</code> <code>COLUMNS</code> Table .....	2181
19.5. The <code>INFORMATION_SCHEMA</code> <code>COLUMN_PRIVILEGES</code> Table .....	2182
19.6. The <code>INFORMATION_SCHEMA</code> <code>ENGINES</code> Table .....	2183
19.7. The <code>INFORMATION_SCHEMA</code> <code>EVENTS</code> Table .....	2183
19.8. The <code>INFORMATION_SCHEMA</code> <code>FILES</code> Table .....	2186
19.9. The <code>INFORMATION_SCHEMA</code> <code>GLOBAL_STATUS</code> and <code>SESSION_STATUS</code> Tables .....	2188
19.10. The <code>INFORMATION_SCHEMA</code> <code>GLOBAL_VARIABLES</code> and <code>SESSION_VARIABLES</code> Tables .....	2188
19.11. The <code>INFORMATION_SCHEMA</code> <code>KEY_COLUMN_USAGE</code> Table .....	2188
19.12. The <code>INFORMATION_SCHEMA</code> <code>OPTIMIZER_TRACE</code> Table .....	2189
19.13. The <code>INFORMATION_SCHEMA</code> <code>PARAMETERS</code> Table .....	2189
19.14. The <code>INFORMATION_SCHEMA</code> <code>PARTITIONS</code> Table .....	2190
19.15. The <code>INFORMATION_SCHEMA</code> <code>PLUGINS</code> Table .....	2193
19.16. The <code>INFORMATION_SCHEMA</code> <code>PROCESSLIST</code> Table .....	2194
19.17. The <code>INFORMATION_SCHEMA</code> <code>PROFILING</code> Table .....	2195
19.18. The <code>INFORMATION_SCHEMA</code> <code>REFERENTIAL_CONSTRAINTS</code> Table .....	2196
19.19. The <code>INFORMATION_SCHEMA</code> <code>ROUTINES</code> Table .....	2196
19.20. The <code>INFORMATION_SCHEMA</code> <code>SCHEMATA</code> Table .....	2198
19.21. The <code>INFORMATION_SCHEMA</code> <code>SCHEMA_PRIVILEGES</code> Table .....	2198
19.22. The <code>INFORMATION_SCHEMA</code> <code>STATISTICS</code> Table .....	2198
19.23. The <code>INFORMATION_SCHEMA</code> <code>TABLES</code> Table .....	2199
19.24. The <code>INFORMATION_SCHEMA</code> <code>TABLESPACES</code> Table .....	2200
19.25. The <code>INFORMATION_SCHEMA</code> <code>TABLE_CONSTRAINTS</code> Table .....	2201
19.26. The <code>INFORMATION_SCHEMA</code> <code>TABLE_PRIVILEGES</code> Table .....	2201
19.27. The <code>INFORMATION_SCHEMA</code> <code>TRIGGERS</code> Table .....	2202
19.28. The <code>INFORMATION_SCHEMA</code> <code>USER_PRIVILEGES</code> Table .....	2203
19.29. The <code>INFORMATION_SCHEMA</code> <code>VIEWS</code> Table .....	2204
19.30. <code>INFORMATION_SCHEMA</code> Tables for InnoDB .....	2205
19.30.1. The <code>INFORMATION_SCHEMA</code> <code>INNODB_CMP</code> and <code>INNODB_CMP_RESET</code> Tables .....	2205
19.30.2. The <code>INFORMATION_SCHEMA</code> <code>INNODB_CMP_PER_INDEX</code> and <code>INNODB_CMP_PER_INDEX_RESET</code> Tables .....	2206
19.30.3. The <code>INFORMATION_SCHEMA</code> <code>INNODB_CMPMEM</code> and <code>INNODB_CMPMEM_RESET</code> Tables .....	2207
19.30.4. The <code>INFORMATION_SCHEMA</code> <code>INNODB_TRX</code> Table .....	2207
19.30.5. The <code>INFORMATION_SCHEMA</code> <code>INNODB_LOCKS</code> Table .....	2209
19.30.6. The <code>INFORMATION_SCHEMA</code> <code>INNODB_LOCK_WAITS</code> Table .....	2210
19.30.7. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_TABLES</code> Table .....	2210
19.30.8. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_INDEXES</code> Table .....	2211
19.30.9. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_COLUMNS</code> Table .....	2212
19.30.10. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_FIELDS</code> Table .....	2212
19.30.11. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_FOREIGN</code> Table .....	2213
19.30.12. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_FOREIGN_COLS</code> Table .....	2213
19.30.13. The <code>INFORMATION_SCHEMA</code> <code>INNODB_SYS_TABLESTATS</code> View .....	2214
19.30.14. The <code>INFORMATION_SCHEMA</code> <code>INNODB_BUFFER_PAGE</code> Table .....	2214
19.30.15. The <code>INFORMATION_SCHEMA</code> <code>INNODB_BUFFER_PAGE_LRU</code> Table .....	2215
19.30.16. The <code>INFORMATION_SCHEMA</code> <code>INNODB_BUFFER_POOL_STATS</code> Table .....	2216
19.30.17. The <code>INFORMATION_SCHEMA</code> <code>INNODB_METRICS</code> Table .....	2217
19.30.18. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_CONFIG</code> Table .....	2222
19.30.19. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_DEFAULT_STOPWORD</code> Table .....	2223
19.30.20. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_INDEX_TABLE</code> Table .....	2223
19.30.21. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_INDEX_CACHE</code> Table .....	2224
19.30.22. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_DELETED</code> Table .....	2224
19.30.23. The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_BEING_DELETED</code> Table .....	2225
19.31. Thread Pool <code>INFORMATION_SCHEMA</code> Tables .....	2225
19.31.1. The <code>INFORMATION_SCHEMA</code> <code>TP_THREAD_STATE</code> Table .....	2226

19.31.2. The <code>INFORMATION_SCHEMA TP_THREAD_GROUP_STATE</code> Table .....	2226
19.31.3. The <code>INFORMATION_SCHEMA TP_THREAD_GROUP_STATS</code> Table .....	2228
19.32. Extensions to <code>SHOW</code> Statements .....	2229
20. MySQL Performance Schema .....	2233
20.1. Performance Schema Quick Start .....	2234
20.2. Performance Schema Configuration .....	2240
20.2.1. Performance Schema Build Configuration .....	2240
20.2.2. Performance Schema Startup Configuration .....	2241
20.2.3. Performance Schema Runtime Configuration .....	2243
20.3. Performance Schema Queries .....	2258
20.4. Performance Schema Instrument Naming Conventions .....	2259
20.5. Performance Schema Status Monitoring .....	2261
20.6. Performance Schema Atom and Molecule Events .....	2264
20.7. Performance Schema Statement Digests .....	2264
20.8. Performance Schema General Table Characteristics .....	2266
20.9. Performance Schema Table Descriptions .....	2266
20.9.1. Performance Schema Setup Tables .....	2266
20.9.2. Performance Schema Instance Tables .....	2270
20.9.3. Performance Schema Wait Event Tables .....	2275
20.9.4. Performance Schema Stage Event Tables .....	2278
20.9.5. Performance Schema Statement Event Tables .....	2279
20.9.6. Performance Schema Connection Tables .....	2284
20.9.7. Performance Schema Connection Attribute Tables .....	2286
20.9.8. Performance Schema Summary Tables .....	2286
20.9.9. Performance Schema Miscellaneous Tables .....	2299
20.10. Performance Schema Option and Variable Reference .....	2305
20.11. Performance Schema Command Options .....	2306
20.12. Performance Schema System Variables .....	2307
20.13. Performance Schema Status Variables .....	2318
20.14. Performance Schema and Plugins .....	2319
20.15. Using the Performance Schema to Diagnose Problems .....	2319
21. Connectors and APIs .....	2321
21.1. MySQL Connector/ODBC .....	2326
21.1.1. Connector/ODBC Versions .....	2327
21.1.2. Connector/ODBC Introduction .....	2327
21.1.3. Connector/ODBC Installation .....	2330
21.1.4. Configuring Connector/ODBC .....	2341
21.1.5. Connector/ODBC Examples .....	2360
21.1.6. Connector/ODBC Reference .....	2386
21.1.7. Connector/ODBC Notes and Tips .....	2392
21.1.8. Connector/ODBC Support .....	2402
21.2. MySQL Connector/Net .....	2403
21.2.1. Connector/Net Versions .....	2404
21.2.2. Connector/Net Installation .....	2406
21.2.3. Connector/Net Visual Studio Integration .....	2414
21.2.4. Connector/Net Tutorials .....	2444
21.2.5. Connector/Net Programming .....	2484
21.2.6. Connector/Net Connection String Options Reference .....	2523
21.2.7. Connector/Net API Reference .....	2530
21.2.8. Connector/Net Support .....	2633
21.2.9. Connector/Net FAQ .....	2634
21.3. MySQL Connector/J .....	2635
21.3.1. Overview of MySQL Connector/J .....	2635
21.3.2. Connector/J Versions .....	2635
21.3.3. Connector/J Installation .....	2637
21.3.4. Connector/J Examples .....	2641
21.3.5. Connector/J (JDBC) Reference .....	2642
21.3.6. JDBC Concepts .....	2675



21.3.7. Connection Pooling with Connector/J .....	2684
21.3.8. Load Balancing with Connector/J .....	2687
21.3.9. Failover with Connector/J .....	2689
21.3.10. Using the Connector/J Interceptor Classes .....	2691
21.3.11. Using Connector/J with Tomcat .....	2691
21.3.12. Using Connector/J with JBoss .....	2693
21.3.13. Using Connector/J with Spring .....	2694
21.3.14. Using Connector/J with GlassFish .....	2698
21.3.15. Troubleshooting Connector/J Applications .....	2704
21.3.16. Connector/J Support .....	2712
21.4. MySQL Connector/C++ .....	2713
21.4.1. How to Get MySQL Connector/C++ .....	2715
21.4.2. Installing MySQL Connector/C++ from a Binary Distribution .....	2716
21.4.3. Installing MySQL Connector/C++ from Source .....	2718
21.4.4. Building MySQL Connector/C++ Windows Applications with Microsoft Visual Studio .....	2723
21.4.5. Building MySQL Connector/C++ Linux Applications with NetBeans .....	2731
21.4.6. MySQL Connector/C++ Getting Started: Usage Examples .....	2735
21.4.7. MySQL Connector/C++ Tutorials .....	2741
21.4.8. MySQL Connector/C++ Debug Tracing .....	2749
21.4.9. MySQL Connector/C++ Usage Notes .....	2750
21.4.10. MySQL Connector/C++ Known Bugs and Issues .....	2754
21.4.11. MySQL Connector/C++ Support .....	2755
21.5. MySQL Connector/C .....	2755
21.5.1. MySQL Connector/C Supported Platforms .....	2755
21.5.2. MySQL Connector/C Distribution Contents .....	2756
21.5.3. Obtaining MySQL Connector/C .....	2756
21.5.4. Installing MySQL Connector/C from a Binary Distribution .....	2756
21.5.5. Installing MySQL Connector/C from Source .....	2756
21.5.6. Testing MySQL Connector/C .....	2758
21.5.7. Building MySQL Connector/C Applications .....	2758
21.6. MySQL Connector/Python .....	2759
21.6.1. Guidelines for Python Developers .....	2759
21.6.2. Connector/Python Versions .....	2761
21.6.3. Connector/Python Installation .....	2761
21.6.4. Connector/Python Coding Examples .....	2762
21.6.5. Connector/Python Tutorials .....	2768
21.6.6. Connector/Python Connection Arguments .....	2769
21.6.7. Connector/Python API Reference .....	2772
21.6.8. Connector/Python FAQ .....	2790
21.7. libmysqld, the Embedded MySQL Server Library .....	2790
21.7.1. Compiling Programs with <a href="#">libmysqld</a> .....	2790
21.7.2. Restrictions When Using the Embedded MySQL Server .....	2791
21.7.3. Options with the Embedded Server .....	2791
21.7.4. Embedded Server Examples .....	2792
21.8. MySQL C API .....	2795
21.8.1. MySQL C API Implementations .....	2796
21.8.2. Example C API Client Programs .....	2797
21.8.3. Building and Running C API Client Programs .....	2797
21.8.4. C API Data Structures .....	2800
21.8.5. C API Function Overview .....	2805
21.8.6. C API Function Descriptions .....	2809
21.8.7. C API Prepared Statements .....	2859
21.8.8. C API Prepared Statement Data Structures .....	2859
21.8.9. C API Prepared Statement Function Overview .....	2865
21.8.10. C API Prepared Statement Function Descriptions .....	2868
21.8.11. C API Threaded Function Descriptions .....	2890
21.8.12. C API Embedded Server Function Descriptions .....	2892

21.8.13. C API Client Plugin Functions .....	2892
21.8.14. Common Questions and Problems When Using the C API .....	2895
21.8.15. Controlling Automatic Reconnection Behavior .....	2897
21.8.16. C API Support for Multiple Statement Execution .....	2898
21.8.17. C API Prepared Statement Problems .....	2900
21.8.18. C API Prepared Statement Handling of Date and Time Values .....	2900
21.8.19. C API Support for Prepared <code>CALL</code> Statements .....	2902
21.9. MySQL PHP API .....	2905
21.9.1. Overview of the MySQL PHP drivers .....	2906
21.9.2. Original MySQL API ( <code>Mysql</code> ) .....	2912
21.9.3. MySQL Improved Extension ( <code>Mysqli</code> ) .....	2982
21.9.4. MySQL Functions (PDO_MYSQL) ( <code>MySQL (PDO)</code> ) .....	3196
21.9.5. MySQL Native Driver ( <code>Mysqlnd</code> ) .....	3200
21.9.6. Mysqlnd replication and load balancing plugin ( <code>mysqlnd_ms</code> ) .....	3232
21.9.7. Mysqlnd query result cache plugin ( <code>mysqlnd_qc</code> ) .....	3333
21.9.8. Mysqlnd user handler plugin ( <code>mysqlnd_uh</code> ) .....	3384
21.9.9. Mysqlnd connection multiplexing plugin ( <code>mysqlnd_mux</code> ) .....	3467
21.9.10. Mysqlnd Memcache plugin ( <code>mysqlnd_memcache</code> ) .....	3471
21.9.11. Connector/PHP .....	3480
21.9.12. Common Problems with MySQL and PHP .....	3480
21.9.13. Enabling Both <code>mysql</code> and <code>mysqli</code> in PHP .....	3480
21.10. MySQL Perl API .....	3481
21.11. MySQL Python API .....	3482
21.12. MySQL Ruby APIs .....	3482
21.12.1. The MySQL/Ruby API .....	3482
21.12.2. The Ruby/MySQL API .....	3482
21.13. MySQL Tcl API .....	3482
21.14. MySQL Eiffel Wrapper .....	3482
22. Extending MySQL .....	3483
22.1. MySQL Internals .....	3483
22.1.1. MySQL Threads .....	3483
22.1.2. The MySQL Test Suite .....	3484
22.2. The MySQL Plugin API .....	3485
22.2.1. Plugin API Characteristics .....	3485
22.2.2. Plugin API Components .....	3486
22.2.3. Types of Plugins .....	3487
22.2.4. Writing Plugins .....	3490
22.2.5. MySQL Services for Plugins .....	3530
22.3. Adding New Functions to MySQL .....	3532
22.3.1. Features of the User-Defined Function Interface .....	3532
22.3.2. Adding a New User-Defined Function .....	3532
22.3.3. Adding a New Native Function .....	3542
22.4. Debugging and Porting MySQL .....	3543
22.4.1. Debugging a MySQL Server .....	3544
22.4.2. Debugging a MySQL Client .....	3550
22.4.3. The DEBUG Package .....	3550
23. MySQL Enterprise Edition .....	3553
23.1. MySQL Enterprise Monitor .....	3553
23.2. MySQL Enterprise Backup .....	3554
23.3. MySQL Enterprise Security .....	3554
23.4. MySQL Enterprise Audit .....	3555
23.5. MySQL Thread Pool .....	3555
24. MySQL Workbench .....	3557
25. MySQL for Excel .....	3559
25.1. Installing and Configuring .....	3559
25.2. Edit MySQL Data .....	3560
25.3. Import MySQL Data .....	3561
25.4. Append MySQL Data .....	3562



25.5. Export MySQL Data .....	3565
25.6. MySQL for Excel FAQ .....	3568
A. Licenses for Third-Party Components .....	3569
A.1. Ant-Contrib License .....	3572
A.2. ANTLR 3.3 License .....	3573
A.3. Boost Library License .....	3573
A.4. c3p0 JDBC Library License .....	3574
A.5. <a href="#">dtoa.c</a> License .....	3574
A.6. Editline Library ( <a href="#">libedit</a> ) License .....	3574
A.7. Facebook Fast Checksum Patch License .....	3577
A.8. Facebook Patches License .....	3578
A.9. <a href="#">FindGTest.cmake</a> License .....	3578
A.10. Fred Fish's Dbug Library License .....	3579
A.11. <a href="#">getarg</a> License .....	3580
A.12. GLib License (for MySQL Proxy) .....	3580
A.13. GNU General Public License Version 2.0, June 1991 .....	3581
A.14. GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009 .....	3586
A.15. GNU Lesser General Public License Version 2.1, February 1999 .....	3597
A.16. GNU Libtool License .....	3604
A.17. GNU Readline License .....	3605
A.18. GNU Standard C++ Library (libstdc++) License .....	3605
A.19. Google Controlling Master Thread I/O Rate Patch License .....	3606
A.20. Google Perftools (TCMalloc utility) License .....	3607
A.21. Google SMP Patch License .....	3607
A.22. jboss-common-jdbc-wrapper.jar License .....	3608
A.23. <a href="#">lib_sql.cc</a> License .....	3608
A.24. Libaio License .....	3608
A.25. <a href="#">libevent</a> License .....	3609
A.26. Libiconv License .....	3610
A.27. <a href="#">libintl</a> License .....	3611
A.28. Linux-PAM License .....	3611
A.29. <a href="#">LPeg</a> Library License .....	3612
A.30. Lua (liblua) License .....	3612
A.31. <a href="#">LuaFileSystem</a> Library License .....	3613
A.32. md5 (Message-Digest Algorithm 5) License .....	3613
A.33. <a href="#">memcached</a> License .....	3613
A.34. <a href="#">mkpasswd.pl</a> License .....	3614
A.35. nt_servc (Windows NT Service class library) License .....	3617
A.36. OpenPAM License .....	3618
A.37. OpenSSL v1.0 License .....	3618
A.38. PCRE License .....	3620
A.39. Percona Multiple I/O Threads Patch License .....	3621
A.40. Red HAT RPM Spec File License .....	3621
A.41. RegEX-Spencer Library License .....	3622
A.42. RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License .....	3622
A.43. Richard A. O'Keefe String Library License .....	3622
A.44. SHA-1 in C License .....	3623
A.45. Simple Logging Facade for Java (SLF4J) License .....	3623
A.46. Unicode Data Files .....	3623
A.47. <a href="#">zlib</a> License .....	3624
A.48. ZLIB.NET License .....	3624
B. MySQL 5.6 Frequently Asked Questions .....	3627
B.1. MySQL 5.6 FAQ: General .....	3627
B.2. MySQL 5.6 FAQ: Storage Engines .....	3629
B.3. MySQL 5.6 FAQ: Server SQL Mode .....	3630
B.4. MySQL 5.6 FAQ: Stored Procedures and Functions .....	3631
B.5. MySQL 5.6 FAQ: Triggers .....	3634

B.6. MySQL 5.6 FAQ: Views .....	3637
B.7. MySQL 5.6 FAQ: <a href="#">INFORMATION_SCHEMA</a> .....	3637
B.8. MySQL 5.6 FAQ: Migration .....	3638
B.9. MySQL 5.6 FAQ: Security .....	3639
B.10. MySQL 5.6 FAQ: MySQL Cluster .....	3640
B.11. MySQL 5.6 FAQ: MySQL Chinese, Japanese, and Korean Character Sets .....	3640
B.12. MySQL 5.6 FAQ: Connectors & APIs .....	3652
B.13. MySQL 5.6 FAQ: Replication .....	3652
B.14. MySQL 5.6 FAQ: MySQL Enterprise Scalability Thread Pool .....	3656
C. Errors, Error Codes, and Common Problems .....	3659
C.1. Sources of Error Information .....	3659
C.2. Types of Error Values .....	3659
C.3. Server Error Codes and Messages .....	3660
C.4. Client Error Codes and Messages .....	3718
C.5. Problems and Common Errors .....	3722
C.5.1. How to Determine What Is Causing a Problem .....	3722
C.5.2. Common Errors When Using MySQL Programs .....	3723
C.5.3. Installation-Related Issues .....	3735
C.5.4. Administration-Related Issues .....	3736
C.5.5. Query-Related Issues .....	3743
C.5.6. Optimizer-Related Issues .....	3749
C.5.7. Table Definition-Related Issues .....	3750
C.5.8. Known Issues in MySQL .....	3751
D. MySQL Release Notes .....	3755
E. Restrictions and Limits .....	3757
E.1. Restrictions on Stored Programs .....	3757
E.2. Restrictions on Condition Handling .....	3760
E.3. Restrictions on Server-Side Cursors .....	3760
E.4. Restrictions on Subqueries .....	3761
E.5. Restrictions on Views .....	3763
E.6. Restrictions on XA Transactions .....	3764
E.7. Restrictions on Character Sets .....	3765
E.8. Restrictions on Performance Schema .....	3765
E.9. Restrictions on Pluggable Authentication .....	3766
E.10. Limits in MySQL .....	3767
E.10.1. Limits of Joins .....	3767
E.10.2. Limits on Number of Databases and Tables .....	3768
E.10.3. Limits on Table Size .....	3768
E.10.4. Limits on Table Column Count and Row Size .....	3769
E.10.5. Limits Imposed by <a href="#">.frm</a> File Structure .....	3770
E.10.6. Windows Platform Limitations .....	3771
MySQL Glossary .....	3775
General Index .....	3829
C Function Index .....	3901
Command Index .....	3911
Function Index .....	3931
<a href="#">INFORMATION_SCHEMA</a> Index .....	3949
Join Types Index .....	3953
Operator Index .....	3955
Option Index .....	3961
Privileges Index .....	4005
SQL Modes Index .....	4011
Statement/Syntax Index .....	4013
System Variable Index .....	4055
Status Variable Index .....	4083
Transaction Isolation Level Index .....	4091

---

# Preface and Legal Notices

This is the Reference Manual for the MySQL Database System, version 5.6, through release 5.6.12. Differences between minor versions of MySQL 5.6 are noted in the present text with reference to release numbers (5.6.x). For license information, see the [Legal Notices](#). This product may contain third-party code. For license information on third-party code, see [Appendix A, Licenses for Third-Party Components](#).

This manual is not intended for use with older versions of the MySQL software due to the many functional and other differences between MySQL 5.6 and previous versions. If you are using an earlier release of the MySQL software, please refer to the appropriate manual. For example, [MySQL 5.5 Reference Manual](#) covers the 5.5 series of MySQL software releases.

## Legal Notices

Copyright © 1997, 2013, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. MySQL is a trademark of Oracle Corporation and/or its affiliates, and shall not be used without Oracle's express written authorization. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this material is subject to the terms and conditions of your Oracle Software License and Service Agreement, which has been executed and with

which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle or as specifically provided below. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

For more information on the terms of this license, or for details on how the MySQL documentation is built and produced, please visit [MySQL Contact & Questions](#).

For additional licensing information, including licenses for third-party libraries used by MySQL products, see [Preface and Legal Notices](#).

For help with using MySQL, please visit either the [MySQL Forums](#) or [MySQL Mailing Lists](#) where you can discuss your issues with other MySQL users.

For additional documentation on MySQL products, including translations of the documentation into other languages, and downloadable versions in variety of formats, including HTML and PDF formats, see the [MySQL Documentation Library](#).

---

# Chapter 1. General Information

## Table of Contents

1.1. About This Manual .....	2
1.2. Typographical and Syntax Conventions .....	2
1.3. Overview of the MySQL Database Management System .....	4
1.3.1. What is MySQL? .....	4
1.3.2. The Main Features of MySQL .....	6
1.3.3. History of MySQL .....	8
1.4. What Is New in MySQL 5.6 .....	9
1.5. MySQL Development History .....	17
1.6. MySQL Information Sources .....	18
1.6.1. MySQL Mailing Lists .....	18
1.6.2. MySQL Community Support at the MySQL Forums .....	20
1.6.3. MySQL Community Support on Internet Relay Chat (IRC) .....	21
1.6.4. MySQL Enterprise .....	21
1.7. How to Report Bugs or Problems .....	21
1.8. MySQL Standards Compliance .....	26
1.8.1. What Standards MySQL Follows .....	26
1.8.2. Selecting SQL Modes .....	26
1.8.3. Running MySQL in ANSI Mode .....	27
1.8.4. MySQL Extensions to Standard SQL .....	27
1.8.5. MySQL Differences from Standard SQL .....	30
1.8.6. How MySQL Deals with Constraints .....	34
1.9. Credits .....	38
1.9.1. Contributors to MySQL .....	38
1.9.2. Documenters and translators .....	42
1.9.3. Packages that support MySQL .....	44
1.9.4. Tools that were used to create MySQL .....	44
1.9.5. Supporters of MySQL .....	45

The MySQL™ software delivers a very fast, multi-threaded, multi-user, and robust SQL (Structured Query Language) database server. MySQL Server is intended for mission-critical, heavy-load production systems as well as for embedding into mass-deployed software. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. MySQL is a trademark of Oracle Corporation and/or its affiliates, and shall not be used by Customer without Oracle's express written authorization. Other names may be trademarks of their respective owners.

The MySQL software is Dual Licensed. Users can choose to use the MySQL software as an Open Source product under the terms of the GNU General Public License (<http://www.fsf.org/licenses/>) or can purchase a standard commercial license from Oracle. See <http://www.mysql.com/company/legal/licensing/> for more information on our licensing policies.

The following list describes some sections of particular interest in this manual:

- For a discussion of MySQL Database Server capabilities, see [Section 1.3.2, “The Main Features of MySQL”](#).
- For an overview of new MySQL features, see [Section 1.4, “What Is New in MySQL 5.6”](#). For information about the changes in each version, see the [Release Notes](#).
- For installation instructions, see [Chapter 2, \*Installing and Upgrading MySQL\*](#). For information about upgrading MySQL, see [Section 2.11.1, “Upgrading MySQL”](#).
- For a tutorial introduction to the MySQL Database Server, see [Chapter 3, \*Tutorial\*](#).

- For information about configuring and administering MySQL Server, see [Chapter 5, MySQL Server Administration](#).
- For information about security in MySQL, see [Chapter 6, Security](#).
- For information about setting up replication servers, see [Chapter 16, Replication](#).
- For information about MySQL Enterprise, the commercial MySQL release with advanced features and management tools, see [Chapter 23, MySQL Enterprise Edition](#).
- For answers to a number of questions that are often asked concerning the MySQL Database Server and its capabilities, see [Appendix B, MySQL 5.6 Frequently Asked Questions](#).
- For a history of new features and bugfixes, see the [Release Notes](#).



### Important

To report problems or bugs, please use the instructions at [Section 1.7, “How to Report Bugs or Problems”](#). If you find a sensitive security bug in MySQL Server, please let us know immediately by sending an email message to [<secalert\\_us@oracle.com>](mailto:secalert_us@oracle.com). Exception: Support customers should report all problems, including security bugs, to Oracle Support.

## 1.1. About This Manual

This is the Reference Manual for the MySQL Database System, version 5.6, through release 5.6.12. Differences between minor versions of MySQL 5.6 are noted in the present text with reference to release numbers (5.6.x). For license information, see the [Legal Notices](#). This product may contain third-party code. For license information on third-party code, see [Appendix A, Licenses for Third-Party Components](#).

This manual is not intended for use with older versions of the MySQL software due to the many functional and other differences between MySQL 5.6 and previous versions. If you are using an earlier release of the MySQL software, please refer to the appropriate manual. For example, [MySQL 5.5 Reference Manual](#) covers the 5.5 series of MySQL software releases.

Because this manual serves as a reference, it does not provide general instruction on SQL or relational database concepts. It also does not teach you how to use your operating system or command-line interpreter.

The MySQL Database Software is under constant development, and the Reference Manual is updated frequently as well. The most recent version of the manual is available online in searchable form at <http://dev.mysql.com/doc/>. Other formats also are available there, including HTML, PDF, and EPUB versions.

The Reference Manual source files are written in DocBook XML format. The HTML version and other formats are produced automatically, primarily using the DocBook XSL stylesheets. For information about DocBook, see <http://docbook.org/>

If you have questions about using MySQL, you can ask them using our mailing lists or forums. See [Section 1.6.1, “MySQL Mailing Lists”](#), and [Section 1.6.2, “MySQL Community Support at the MySQL Forums”](#). If you have suggestions concerning additions or corrections to the manual itself, please send them to the <http://www.mysql.com/company/contact/>.

This manual was originally written by David Axmark and Michael “Monty” Widenius. It is maintained by the MySQL Documentation Team, consisting of Paul DuBois, Stefan Hinz, Philip Olson, John Russell, and Jon Stephens.

## 1.2. Typographical and Syntax Conventions

This manual uses certain typographical conventions:



- *Text in this style* is used for SQL statements; database, table, and column names; program listings and source code; and environment variables. Example: “To reload the grant tables, use the `FLUSH PRIVILEGES` statement.”
- *Text in this style* indicates input that you type in examples.
- *Text in this style* indicates the names of executable programs and scripts, examples being `mysql` (the MySQL command-line client program) and `mysqld` (the MySQL server executable).
- *Text in this style* is used for variable input for which you should substitute a value of your own choosing.
- *Text in this style* is used for emphasis.
- **Text in this style** is used in table headings and to convey especially strong emphasis.
- *Text in this style* is used to indicate a program option that affects how the program is executed, or that supplies information that is needed for the program to function in a certain way. *Example:* “The `--host` option (short form `-h`) tells the `mysql` client program the hostname or IP address of the MySQL server that it should connect to”.
- File names and directory names are written like this: “The global `my.cnf` file is located in the `/etc` directory.”
- Character sequences are written like this: “To specify a wildcard, use the ‘%’ character.”

When commands are shown that are meant to be executed from within a particular program, the prompt shown preceding the command indicates which command to use. For example, `shell>` indicates a command that you execute from your login shell, `root-shell>` is similar but should be executed as `root`, and `mysql>` indicates a statement that you execute from the `mysql` client program:

```
shell> type a shell command here
root-shell> type a shell command as root here
mysql> type a mysql statement here
```

In some areas different systems may be distinguished from each other to show that commands should be executed in two different environments. For example, while working with replication the commands might be prefixed with `master` and `slave`:

```
master> type a mysql command on the replication master here
slave> type a mysql command on the replication slave here
```

The “shell” is your command interpreter. On Unix, this is typically a program such as `sh`, `csch`, or `bash`. On Windows, the equivalent program is `command.com` or `cmd.exe`, typically run in a console window.

When you enter a command or statement shown in an example, do not type the prompt shown in the example.

Database, table, and column names must often be substituted into statements. To indicate that such substitution is necessary, this manual uses `db_name`, `tbl_name`, and `col_name`. For example, you might see a statement like this:

```
mysql> SELECT col_name FROM db_name.tbl_name;
```

This means that if you were to enter a similar statement, you would supply your own database, table, and column names, perhaps like this:

```
mysql> SELECT author_name FROM biblio_db.author_list;
```

SQL keywords are not case sensitive and may be written in any lettercase. This manual uses uppercase.

In syntax descriptions, square brackets (“[” and “]”) indicate optional words or clauses. For example, in the following statement, `IF EXISTS` is optional:

```
DROP TABLE [IF EXISTS] tbl_name
```

When a syntax element consists of a number of alternatives, the alternatives are separated by vertical bars (“|”). When one member from a set of choices *may* be chosen, the alternatives are listed within square brackets (“[” and “]”):

```
TRIM([[BOTH | LEADING | TRAILING] [remstr] FROM] str)
```

When one member from a set of choices *must* be chosen, the alternatives are listed within braces (“{” and “}”):

```
{DESCRIBE | DESC} tbl_name [col_name | wild]
```

An ellipsis (...) indicates the omission of a section of a statement, typically to provide a shorter version of more complex syntax. For example, `SELECT ... INTO OUTFILE` is shorthand for the form of `SELECT` statement that has an `INTO OUTFILE` clause following other parts of the statement.

An ellipsis can also indicate that the preceding syntax element of a statement may be repeated. In the following example, multiple `reset_option` values may be given, with each of those after the first preceded by commas:

```
RESET reset_option [,reset_option] ...
```

Commands for setting shell variables are shown using Bourne shell syntax. For example, the sequence to set the `CC` environment variable and run the `configure` command looks like this in Bourne shell syntax:

```
shell> CC=gcc ./configure
```

If you are using `csh` or `tcsh`, you must issue commands somewhat differently:

```
shell> setenv CC gcc
shell> ./configure
```

## 1.3. Overview of the MySQL Database Management System

### 1.3.1. What is MySQL?

MySQL, the most popular Open Source SQL database management system, is developed, distributed, and supported by Oracle Corporation.

The MySQL Web site (<http://www.mysql.com/>) provides the latest information about MySQL software.

- **MySQL is a database management system.**

A database is a structured collection of data. It may be anything from a simple shopping list to a picture gallery or the vast amounts of information in a corporate network. To add, access, and process data stored in a computer database, you need a database management system such as MySQL Server. Since computers are very good at handling large amounts of data, database management systems play a central role in computing, as standalone utilities, or as parts of other applications.

- **MySQL databases are relational.**

A relational database stores data in separate tables rather than putting all the data in one big storeroom. The database structures are organized into physical files optimized for speed. The logical model, with objects such as databases, tables, views, rows, and columns, offers a flexible



programming environment. You set up rules governing the relationships between different data fields, such as one-to-one, one-to-many, unique, required or optional, and “pointers” between different tables. The database enforces these rules, so that with a well-designed database, your application never sees inconsistent, duplicate, orphan, out-of-date, or missing data.

The SQL part of “MySQL” stands for “Structured Query Language”. SQL is the most common standardized language used to access databases. Depending on your programming environment, you might enter SQL directly (for example, to generate reports), embed SQL statements into code written in another language, or use a language-specific API that hides the SQL syntax.

SQL is defined by the ANSI/ISO SQL Standard. The SQL standard has been evolving since 1986 and several versions exist. In this manual, “SQL-92” refers to the standard released in 1992, “SQL:1999” refers to the standard released in 1999, and “SQL:2003” refers to the current version of the standard. We use the phrase “the SQL standard” to mean the current version of the SQL Standard at any time.

- **MySQL software is Open Source.**

Open Source means that it is possible for anyone to use and modify the software. Anybody can download the MySQL software from the Internet and use it without paying anything. If you wish, you may study the source code and change it to suit your needs. The MySQL software uses the GPL (GNU General Public License), <http://www.fsf.org/licenses/>, to define what you may and may not do with the software in different situations. If you feel uncomfortable with the GPL or need to embed MySQL code into a commercial application, you can buy a commercially licensed version from us. See the MySQL Licensing Overview for more information (<http://www.mysql.com/company/legal/licensing/>).

- **The MySQL Database Server is very fast, reliable, scalable, and easy to use.**

If that is what you are looking for, you should give it a try. MySQL Server can run comfortably on a desktop or laptop, alongside your other applications, web servers, and so on, requiring little or no attention. If you dedicate an entire machine to MySQL, you can adjust the settings to take advantage of all the memory, CPU power, and I/O capacity available. MySQL can also scale up to clusters of machines, networked together.

You can find a performance comparison of MySQL Server with other database managers on our benchmark page. See [Section 8.12.2, “The MySQL Benchmark Suite”](#).

MySQL Server was originally developed to handle large databases much faster than existing solutions and has been successfully used in highly demanding production environments for several years. Although under constant development, MySQL Server today offers a rich and useful set of functions. Its connectivity, speed, and security make MySQL Server highly suited for accessing databases on the Internet.

- **MySQL Server works in client/server or embedded systems.**

The MySQL Database Software is a client/server system that consists of a multi-threaded SQL server that supports different backends, several different client programs and libraries, administrative tools, and a wide range of application programming interfaces (APIs).

We also provide MySQL Server as an embedded multi-threaded library that you can link into your application to get a smaller, faster, easier-to-manage standalone product.

- **A large amount of contributed MySQL software is available.**

MySQL Server has a practical set of features developed in close cooperation with our users. It is very likely that your favorite application or language supports the MySQL Database Server.

The official way to pronounce “MySQL” is “My Ess Que Ell” (not “my sequel”), but we do not mind if you pronounce it as “my sequel” or in some other localized way.

## 1.3.2. The Main Features of MySQL

This section describes some of the important characteristics of the MySQL Database Software. See also [Section 1.5, “MySQL Development History”](#). In most respects, the roadmap applies to all versions of MySQL. For information about features as they are introduced into MySQL on a series-specific basis, see the “In a Nutshell” section of the appropriate Manual:

- MySQL 5.6: [MySQL 5.6 in a Nutshell](#)
- MySQL 5.5: [MySQL 5.5 in a Nutshell](#)
- MySQL 5.1: [MySQL 5.1 in a Nutshell](#)
- MySQL 5.0: [MySQL 5.0 in a Nutshell](#)

### Internals and Portability:

- Written in C and C++.
- Tested with a broad range of different compilers.
- Works on many different platforms. See [Section 2.1.1, “Operating Systems Supported by MySQL Community Server”](#).
- For portability, uses [CMake](#) in MySQL 5.5 and up. Previous series use GNU Automake, Autoconf, and Libtool.
- Tested with Purify (a commercial memory leakage detector) as well as with Valgrind, a GPL tool (<http://developer.kde.org/~sewardj/>).
- Uses multi-layered server design with independent modules.
- Designed to be fully multi-threaded using kernel threads, to easily use multiple CPUs if they are available.
- Provides transactional and nontransactional storage engines.
- Uses very fast B-tree disk tables ([MyISAM](#)) with index compression.
- Designed to make it relatively easy to add other storage engines. This is useful if you want to provide an SQL interface for an in-house database.
- Uses a very fast thread-based memory allocation system.
- Executes very fast joins using an optimized nested-loop join.
- Implements in-memory hash tables, which are used as temporary tables.
- Implements SQL functions using a highly optimized class library that should be as fast as possible. Usually there is no memory allocation at all after query initialization.
- Provides the server as a separate program for use in a client/server networked environment, and as a library that can be embedded (linked) into standalone applications. Such applications can be used in isolation or in environments where no network is available.

### Data Types:

- Many data types: signed/unsigned integers 1, 2, 3, 4, and 8 bytes long, [FLOAT](#), [DOUBLE](#), [CHAR](#), [VARCHAR](#), [BINARY](#), [VARBINARY](#), [TEXT](#), [BLOB](#), [DATE](#), [TIME](#), [DATETIME](#), [TIMESTAMP](#), [YEAR](#), [SET](#), [ENUM](#), and OpenGIS spatial types. See [Chapter 11, Data Types](#).
- Fixed-length and variable-length string types.

## Statements and Functions:

- Full operator and function support in the `SELECT` list and `WHERE` clause of queries. For example:

```
mysql> SELECT CONCAT(first_name, ' ', last_name)
-> FROM citizen
-> WHERE income/dependents > 10000 AND age > 30;
```

- Full support for SQL `GROUP BY` and `ORDER BY` clauses. Support for group functions (`COUNT()` [1283], `AVG()` [1282], `STD()` [1284], `SUM()` [1285], `MAX()` [1284], `MIN()` [1284], and `GROUP_CONCAT()` [1283]).
- Support for `LEFT OUTER JOIN` and `RIGHT OUTER JOIN` with both standard SQL and ODBC syntax.
- Support for aliases on tables and columns as required by standard SQL.
- Support for `DELETE`, `INSERT`, `REPLACE`, and `UPDATE` to return the number of rows that were changed (affected), or to return the number of rows matched instead by setting a flag when connecting to the server.
- Support for MySQL-specific `SHOW` statements that retrieve information about databases, storage engines, tables, and indexes. MySQL 5.0 adds support for the `INFORMATION_SCHEMA` database, implemented according to standard SQL.
- An `EXPLAIN` statement to show how the optimizer resolves a query.
- Independence of function names from table or column names. For example, `ABS` is a valid column name. The only restriction is that for a function call, no spaces are permitted between the function name and the "(" that follows it. See [Section 9.3, "Reserved Words"](#).
- You can refer to tables from different databases in the same statement.

## Security:

- A privilege and password system that is very flexible and secure, and that enables host-based verification.
- Password security by encryption of all password traffic when you connect to a server.

## Scalability and Limits:

- Support for large databases. We use MySQL Server with databases that contain 50 million records. We also know of users who use MySQL Server with 200,000 tables and about 5,000,000,000 rows.
- Support for up to 64 indexes per table (32 before MySQL 4.1.2). Each index may consist of 1 to 16 columns or parts of columns. The maximum index width is 767 bytes for `InnoDB` tables, or 1000 for `MyISAM`; before MySQL 4.1.2, the limit is 500 bytes. An index may use a prefix of a column for `CHAR`, `VARCHAR`, `BLOB`, or `TEXT` column types.

## Connectivity:

- Clients can connect to MySQL Server using several protocols:
  - Clients can connect using TCP/IP sockets on any platform.
  - On Windows systems in the NT family (NT, 2000, XP, 2003, or Vista), clients can connect using named pipes if the server is started with the `--enable-named-pipe` [421] option. In MySQL 4.1 and higher, Windows servers also support shared-memory connections if started with the `--shared-memory` [439] option. Clients can connect through shared memory by using the `--protocol=memory` option.
  - On Unix systems, clients can connect using Unix domain socket files.

- MySQL client programs can be written in many languages. A client library written in C is available for clients written in C or C++, or for any language that provides C bindings.
- APIs for C, C++, Eiffel, Java, Perl, PHP, Python, Ruby, and Tcl are available, enabling MySQL clients to be written in many languages. See [Chapter 21, Connectors and APIs](#).
- The Connector/ODBC (MyODBC) interface provides MySQL support for client programs that use ODBC (Open Database Connectivity) connections. For example, you can use MS Access to connect to your MySQL server. Clients can be run on Windows or Unix. Connector/ODBC source is available. All ODBC 2.5 functions are supported, as are many others. See [Section 21.1, “MySQL Connector/ODBC”](#).
- The Connector/J interface provides MySQL support for Java client programs that use JDBC connections. Clients can be run on Windows or Unix. Connector/J source is available. See [Section 21.3, “MySQL Connector/J”](#).
- MySQL Connector/Net enables developers to easily create .NET applications that require secure, high-performance data connectivity with MySQL. It implements the required ADO.NET interfaces and integrates into ADO.NET aware tools. Developers can build applications using their choice of .NET languages. MySQL Connector/Net is a fully managed ADO.NET driver written in 100% pure C#. See [Section 21.2, “MySQL Connector/Net”](#).

### Localization:

- The server can provide error messages to clients in many languages. See [Section 10.2, “Setting the Error Message Language”](#).
- Full support for several different character sets, including `latin1` (cp1252), `german`, `big5`, `ujis`, and more. For example, the Scandinavian characters “å”, “ä” and “ö” are permitted in table and column names. Unicode support is available as of MySQL 4.1.
- All data is saved in the chosen character set.
- Sorting and comparisons are done according to the chosen character set and collation (using `latin1` and Swedish collation by default). It is possible to change this when the MySQL server is started. To see an example of very advanced sorting, look at the Czech sorting code. MySQL Server supports many different character sets that can be specified at compile time and runtime.
- As of MySQL 4.1, the server time zone can be changed dynamically, and individual clients can specify their own time zone. [Section 10.6, “MySQL Server Time Zone Support”](#).

### Clients and Tools:

- MySQL includes several client and utility programs. These include both command-line programs such as `mysqldump` and `mysqladmin`, and graphical programs such as [MySQL Workbench](#).
- MySQL Server has built-in support for SQL statements to check, optimize, and repair tables. These statements are available from the command line through the `mysqlcheck` client. MySQL also includes `myisamchk`, a very fast command-line utility for performing these operations on `MyISAM` tables. See [Chapter 4, MySQL Programs](#).
- MySQL programs can be invoked with the `--help` or `-?` option to obtain online assistance.

### 1.3.3. History of MySQL

We started out with the intention of using the `mSQL` database system to connect to our tables using our own fast low-level (ISAM) routines. However, after some testing, we came to the conclusion that `mSQL` was not fast enough or flexible enough for our needs. This resulted in a new SQL interface to our database but with almost the same API interface as `mSQL`. This API was designed to enable third-party code that was written for use with `mSQL` to be ported easily for use with MySQL.

MySQL is named after co-founder Monty Widenius's daughter, My.

The name of the MySQL Dolphin (our logo) is “Sakila,” which was chosen from a huge list of names suggested by users in our “Name the Dolphin” contest. The winning name was submitted by Ambrose Twebaze, an Open Source software developer from Swaziland, Africa. According to Ambrose, the feminine name Sakila has its roots in SiSwati, the local language of Swaziland. Sakila is also the name of a town in Arusha, Tanzania, near Ambrose’s country of origin, Uganda.

## 1.4. What Is New in MySQL 5.6

This section summarizes what has been added to and removed from MySQL 5.6.

### Added Features

The following features have been added to MySQL 5.6:

- **Security improvements.** These security improvements were made:
  - MySQL now provides a method for storing authentication credentials encrypted in an option file named `.mylogin.cnf`. To create the file, use the `mysql_config_editor` utility. The file can be read later by MySQL client programs to obtain authentication credentials for connecting to a MySQL server. `mysql_config_editor` writes the `.mylogin.cnf` file using encryption so the credentials are not stored as clear text, and its contents when decrypted by client programs are used only in memory. In this way, passwords can be stored in a file in non-clear text format and used later without ever needing to be exposed on the command line or in an environment variable. For more information, see [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#).
  - MySQL now supports stronger encryption for user account passwords, available through an authentication plugin named `sha256_password` that implements SHA-256 password hashing. This plugin is built in, so it is always available and need not be loaded explicitly. For more information, including instructions for creating accounts that use SHA-256 passwords, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).
  - The `mysql.user` table now has a `password_expired` column. Its default value is 'N', but can be set to 'Y' with the new `ALTER USER` statement. After an account’s password has been expired, all operations performed in subsequent connections to the server using the account result in an error until the user issues a `SET PASSWORD` statement to establish a new account password. For more information, see [Section 13.7.1.1, “ALTER USER Syntax”](#).
  - MySQL now has provision for checking password security:
    - In statements that assign a password supplied as a cleartext value, the value is checked against the current password policy and rejected if it is weak (the statement returns an `ER_NOT_VALID_PASSWORD [3714]` error). This affects the `CREATE USER`, `GRANT`, and `SET PASSWORD` statements. Passwords given as arguments to the `PASSWORD()` [\[1263\]](#) and `OLD_PASSWORD()` [\[1263\]](#) functions are checked as well.
    - The strength of potential passwords can be assessed using the new `VALIDATE_PASSWORD_STRENGTH()` [\[1265\]](#) SQL function, which takes a password argument and returns an integer from 0 (weak) to 100 (strong).

Both capabilities are implemented by the `validate_password` plugin. For more information, see [Section 6.1.2.6, “The Password Validation Plugin”](#).

- `mysql_upgrade` now produces a warning if it finds user accounts with passwords hashed with the older pre-4.1 hashing method. Such accounts should be updated to use more secure password hashing. See [Section 6.1.2.4, “Password Hashing in MySQL”](#)
- On Unix platforms, `mysql_install_db` supports a new option, `--random-passwords` [\[237\]](#), that provides for more secure MySQL installation. Invoking `mysql_install_db` with `--random-`

`passwords` [237] causes it to assign a random password to the MySQL `root` accounts, set the “password expired” flag for those accounts, and remove the anonymous-user MySQL accounts. For additional details, see [Section 4.4.3, “mysql\\_install\\_db — Initialize MySQL Data Directory”](#).

- Logging has been modified so that passwords do not appear in plain text in statements written to the general query log, slow query log, and binary log. See [Section 6.1.2.3, “Passwords and Logging”](#).

The `mysql` client no longer logs to its history file statements that refer to passwords. See [Section 4.5.1.3, “mysql Logging”](#).

- `START SLAVE` syntax has been modified to permit connection parameters to be specified for connecting to the master. This provides an alternative to storing the password in the `master.info` file. See [Section 13.4.2.5, “START SLAVE Syntax”](#).
- **Changes to server defaults.** Beginning with MySQL 5.6.6, several MySQL Server parameter defaults differ from the defaults in previous releases. The motivation for these changes is to provide better out-of-box performance and to reduce the need for database administrators to change settings manually. For more information, see [Section 5.1.2.1, “Changes to Server Defaults”](#).
- **InnoDB enhancements.** These InnoDB enhancements were added:
  - You can create `FULLTEXT` indexes on InnoDB tables, and query them using the `MATCH( ) . . . AGAINST [1229]` syntax. This feature includes a new proximity search operator (`@`) and several new configuration options and `INFORMATION_SCHEMA` tables: See [Section 14.2.3.12.3, “FULLTEXT Indexes”](#) for more information.
  - Several `ALTER TABLE` operations can be performed without copying the table, without blocking inserts, updates, and deletes to the table, or both. These enhancements are known collectively as [online DDL](#). See [Section 5.5, “Online DDL for InnoDB Tables”](#) for details.
  - You have more flexibility to move around the `.ibd` files created in `file-per-table` mode to suit your storage devices and database servers. When creating a table, you can designate a location outside the MySQL data directory to hold the `.ibd` file, for example to put a busy table on an `SSD` device or a huge table on a high-capacity `HDD` device. You can export a table from one MySQL instance and import it in a different instance, without inconsistencies or mismatches caused by buffered data, in-progress transactions, and internal bookkeeping details such as the `space ID` and `LSN`. See [Section 14.2.4.2.34, “Improved Tablespace Management”](#) for details.
  - You can now set the InnoDB page size for uncompressed tables to 8KB or 4KB, as an alternative to the default 16KB. This setting is controlled by the `innodb_page_size [1740]` configuration option. You specify the size when creating the MySQL instance. All InnoDB tablespaces within an instance share the same page size. Smaller page sizes can help to avoid redundant or inefficient I/O for certain combinations of workload and storage devices, particularly `SSD` devices with small block sizes.
  - Improvements to the algorithms for [adaptive flushing](#) make I/O operations more efficient and consistent under a variety of `workloads`. The new algorithm and default configuration values are expected to improve performance and concurrency for most users. Advanced users can fine-tune their I/O responsiveness through several configuration options. See [Section 14.2.4.2.9, “Improvements to Buffer Pool Flushing”](#) for details.
  - You can code MySQL applications that access InnoDB tables through a NoSQL-style API. This feature uses the popular `memcached` daemon to relay requests such as `ADD`, `SET`, and `GET` for key-value pairs. These simple operations to store and retrieve data avoid the SQL overhead such as parsing and constructing a [query execution plan](#). You can access the same data through the NoSQL API and SQL. For example, you might use the NoSQL API for fast updates and lookups, and SQL for complex queries and compatibility with existing applications. See [Section 14.2.9, “InnoDB Integration with memcached”](#) for details.



- Optimizer statistics for [InnoDB](#) tables are gathered at more predictable intervals and can persist across server restarts, for improved [plan stability](#). You can also control the amount of sampling done for [InnoDB](#) indexes, to make the optimizer statistics more accurate and improve the query execution plan. See [Section 14.2.4.2.10, “Persistent Optimizer Statistics for InnoDB Tables”](#) for details.
- New optimizations apply to read-only [transactions](#), improving performance and concurrency for ad-hoc queries and report-generating applications. These optimizations are applied automatically when practical, or you can specify `START TRANSACTION READ ONLY` to ensure the transaction is read-only. See [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#) for details.
- You can move the [InnoDB undo log](#) out of the [system tablespace](#) into one or more separate [tablespaces](#). The I/O patterns for the undo log make these new tablespaces good candidates to move to SSD storage, while keeping the system tablespace on hard disk storage. For details, see [Section 14.2.4.2.4, “Separate Tablespaces for InnoDB Undo Logs”](#).
- The [InnoDB redo log](#) files now have a maximum combined size of 512GB, increased from 4GB. You can specify the larger values through the `innodb_log_file_size` [\[1732\]](#) option. The startup behavior now automatically handles the situation where the size of the existing redo log files does not match the size specified by `innodb_log_file_size` and `innodb_log_files_in_group` [\[1733\]](#).
- The `--innodb-read-only` [\[1743\]](#) option lets you run a MySQL server in read-only mode. You can access [InnoDB](#) tables on read-only media such as a DVD or CD, or set up a data warehouse with multiple instances all sharing the same data directory. See [Section 14.2.5.1, “Support for Read-Only Media”](#) for usage details.
- Several new [InnoDB](#)-related `INFORMATION_SCHEMA` tables provide information about the [InnoDB](#) buffer pool, metadata about tables, indexes, and foreign keys from the [InnoDB](#) data dictionary, and low-level information about performance metrics that complements the information from the Performance Schema tables.
- [InnoDB](#) now limits the memory used to hold table information when many tables are opened.
- [InnoDB](#) has several internal performance enhancements, including reducing contention by splitting the kernel mutex, moving flushing operations from the main thread to a separate thread, enabling multiple purge threads, and reducing contention for the buffer pool on large-memory systems.
- [InnoDB](#) uses a new, faster algorithm to detect [deadlocks](#). Information about all [InnoDB](#) deadlocks can be written to the MySQL server error log, to help diagnose application issues.
- To avoid a lengthy [warmup](#) period after restarting the server, particularly for instances with large [InnoDB buffer pools](#), you can reload pages into the buffer pool immediately after a restart. MySQL can dump a compact data file at shutdown, then consult that data file to find the [pages](#) to reload on the next restart. You can also manually dump or reload the buffer pool at any time, for example during benchmarking or after complex report-generation queries. See [Section 14.2.4.2.8, “Faster Restart by Preloading the InnoDB Buffer Pool”](#) for details.
- **Partitioning.** These table-partitioning enhancements were added:
  - The maximum number of partitions is increased to 8192. This number includes all partitions and all subpartitions of the table.
  - It is now possible to exchange a partition of a partitioned table or a subpartition of a subpartitioned table with a nonpartitioned table that otherwise has the same structure using the `ALTER TABLE ... EXCHANGE PARTITION` statement. This can be used, for example, to import and export partitions. For more information and examples, see [Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”](#).

- Explicit selection of one or more partitions or subpartitions is now supported for queries, as well as for many data modification statements, that act on partitioned tables. For example, assume a table `t` with some integer column `c` has 4 partitions named `p0`, `p1`, `p2`, and `p3`. Then the query `SELECT * FROM t PARTITION (p0, p1) WHERE c < 5` returns only those rows from partitions `p0` and `p1` for which `c` is less than 5.

The following statements support explicit partition selection:

- `SELECT`
- `DELETE`
- `INSERT`
- `REPLACE`
- `UPDATE`
- `LOAD DATA`.
- `LOAD XML`.

For syntax, see the descriptions of the individual statements. For additional information and examples, see [Section 17.5, “Partition Selection”](#).

- Partition lock pruning greatly improves performance of many DML and DDL statements acting on tables with many partitions by helping to eliminate locks on partitions that are not affected by these statements. Such statements include many `SELECT`, `SELECT ... PARTITION`, `UPDATE`, `REPLACE`, `INSERT`, as well as many other statements. For more information, including a complete listing of the statements whose performance has thus been improved, see [Section 17.6.4, “Partitioning and Locking”](#).
- **Performance Schema.** The Performance Schema includes several new features:
  - Instrumentation for table input and output. Instrumented operations include row-level accesses to persistent base tables or temporary tables. Operations that affect rows are fetch, insert, update, and delete.
  - Event filtering by table, based on schema and/or table names.
  - Event filtering by thread. More information is collected for threads.
  - Summary tables for table and index I/O, and for table locks.
  - Instrumentation for statements and stages within statements.
  - Configuration of instruments and consumers at server startup, which previously was possible only at runtime.
- **Replication and logging.** These replication enhancements were added:
  - MySQL now supports transaction-based replication using *global transaction identifiers* (also known as “GTIDs”). This makes it possible to identify and track each transaction when it is committed on the originating server and as it is applied by any slaves.

Enabling of GTIDs in a replication setup is done primarily using the new `--gtid-mode` [2019] and `--enforce-gtid-consistency` [2018] server options. For information about additional options and variables introduced in support of GTIDs, see [Section 16.1.4.5, “Global Transaction ID Options and Variables”](#).



When using GTIDs it is not necessary to refer to log files or positions within those files when starting a new slave or failing over to a new master, which greatly simplifies these tasks. For more information about provisioning servers for GTID replication with or without referring to binary log files, see [Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”](#).

GTID-based replication is completely transaction-based, which makes it simple to check the consistency of masters and slaves. If all transactions committed on a given master are also committed on a given slave, consistency between the two servers is guaranteed.

For more complete information about the implementation and use of GTIDs in MySQL Replication, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

- MySQL row-based replication now supports row image control. By logging only those columns required for uniquely identifying and executing changes on each row (as opposed to all columns) for each row change, it is possible to save disk space, network resources, and memory usage. You can determine whether full or minimal rows are logged by setting the `binlog_row_image` [2012] server system variable to one of the values `minimal` (log required columns only), `full` (log all columns), or `noblob` (log all columns except for unneeded `BLOB` or `TEXT` columns). See [System variables used with the binary log](#), for more information.
- Binary logs written and read by the MySQL Server are now crash-safe, because only complete events (or transactions) are logged or read back. By default, the server logs the length of the event as well as the event itself and uses this information to verify that the event was written correctly. You can also cause the server to write checksums for the events using CRC32 checksums by setting the `binlog_checksum` [2009] system variable. To cause the server to read checksums from the binary log, use the `master_verify_checksum` system variable. The `--slave-sql-verify-checksum` system variable causes the slave SQL thread to read checksums from the relay log.
- MySQL now supports logging of master connection information and of slave relay log information to tables as well as files. Use of these tables can be controlled independently, by the `--master-info-repository` [2000] and `--relay-log-info-repository` [2000] server options. Setting `--master-info-repository` [2000] to `TABLE` causes connection information to be logged in the `slave_master_info` table; setting `--relay-log-info-repository` [2000] to `TABLE` causes relay log information to be logged to the `slave_relay_log_info` table. Both of these tables are created automatically, in the `mysql` system database.

In order for replication to be crash-safe, the `slave_master_info` and `slave_relay_log_info` tables must each use a transactional storage engine, and beginning with MySQL 5.6.6, these tables are created using `InnoDB` for this reason. (Bug #13538891) If you are using a previous MySQL 5.6 release in which both of these tables use `MyISAM`, this means that, prior to starting replication, you must convert both of them to a transactional storage engine (such as `InnoDB`) if you wish for replication to be crash-safe. You can do this in such cases by means of the appropriate `ALTER TABLE ... ENGINE=...` statements. You should *not* attempt to change the storage engine used by either of these tables while replication is actually running.



**Note**

In order to guarantee crash safety on the slave, you must also run the slave with the `--relay-log-recovery` [1975] option enabled.

- `mysqlbinlog` now has the capability to back up a binary log in its original binary format. When invoked with the `--read-from-remote-server` [352] and `--raw` [352] options, `mysqlbinlog` connects to a server, requests the log files, and writes output files in the same format as the originals. See [Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”](#).

- MySQL now supports delayed replication such that a slave server deliberately lags behind the master by at least a specified amount of time. The default delay is 0 seconds. Use the new `MASTER_DELAY` option for `CHANGE MASTER TO` to set the delay.

Delayed replication can be used for purposes such as protecting against user mistakes on the master (a DBA can roll back a delayed slave to the time just before the disaster) or testing how the system behaves when there is a lag. See [Section 16.3.9, “Delayed Replication”](#).

- A replication slave having multiple network interfaces can now be caused to use only one of these (to the exclusion of the others) by using the `MASTER_BIND` option when issuing a `CHANGE MASTER TO` statement.
- The `log_bin_basename` [2015] system variable has been added. This variable contains the complete filename and path to the binary log file. Whereas the `log_bin` [2008] system variable shows only whether or not binary logging is enabled, `log_bin_basename` [2015] reflects the name set with the `--log-bin` [2002] server option.

Similarly, the `relay_log_basename` [1989] system variable shows the filename and complete path to the relay log file.

- MySQL Replication now supports parallel execution of transactions with multi-threading on the slave. When parallel execution is enabled, the slave SQL thread acts as the coordinator for a number of slave worker threads as determined by the value of the `slave_parallel_workers` [1993] server system variable. The current implementation of multi-threading on the slave assumes that data and updates are partitioned on a per-database basis, and that updates within a given database occur in the same relative order as they do on the master. However, it is not necessary to coordinate transactions between different databases. Transactions can then also be distributed per database, which means that a worker thread on the slave can process successive transactions on a given database without waiting for updates to other databases to complete.

Since transactions on different databases can occur in a different order on the slave than on the master, simply checking for the most recently executed transaction is not a guarantee that all previous transactions on the master have been executed on the slave. This has implications for logging and recovery when using a multi-threaded slave. For information about how to interpret binary logging information when using multi-threading on the slave, see [Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”](#).

- **Optimizer enhancements.** These query optimizer improvements were implemented:
  - The optimizer now more efficiently handles queries (and subqueries) of the following form:

```
SELECT ... FROM single_table ... ORDER BY non_index_column [DESC] LIMIT [M,]N;
```

That type of query is common in web applications that display only a few rows from a larger result set. For example:

```
SELECT col1, ... FROM t1 ... ORDER BY name LIMIT 10;
SELECT col1, ... FROM t1 ... ORDER BY RAND() LIMIT 15;
```

The sort buffer has a size of `sort_buffer_size` [542]. If the sort elements for  $N$  rows are small enough to fit in the sort buffer ( $M+N$  rows if  $M$  was specified), the server can avoid using a merge file and perform the sort entirely in memory. For details, see [Section 8.2.1.3, “Optimizing LIMIT Queries”](#).

- The optimizer implements Disk-Sweep Multi-Range Read. Reading rows using a range scan on a secondary index can result in many random disk accesses to the base table when the table is large and not stored in the storage engine's cache. With the Disk-Sweep Multi-Range Read (MRR) optimization, MySQL tries to reduce the number of random disk access for range scans by first scanning the index only and collecting the keys for the relevant rows. Then the keys are

sorted and finally the rows are retrieved from the base table using the order of the primary key. The motivation for Disk-sweep MRR is to reduce the number of random disk accesses and instead achieve a more sequential scan of the base table data. For more information, see [Section 8.13.11](#), “Multi-Range Read Optimization”.

- The optimizer implements Index Condition Pushdown (ICP), an optimization for the case where MySQL retrieves rows from a table using an index. Without ICP, the storage engine traverses the index to locate rows in the base table and returns them to the MySQL server which evaluates the `WHERE` condition for the rows. With ICP enabled, and if parts of the `WHERE` condition can be evaluated by using only fields from the index, the MySQL server pushes this part of the `WHERE` condition down to the storage engine. The storage engine then evaluates the pushed index condition by using the index entry and only if this is satisfied is base row be read. ICP can reduce the number of accesses the storage engine has to do against the base table and the number of accesses the MySQL server has to do against the storage engine. For more information, see [Section 8.13.4](#), “Index Condition Pushdown Optimization”.
- The `EXPLAIN` statement now provides execution plan information for `DELETE`, `INSERT`, `REPLACE`, and `UPDATE` statements. Previously, `EXPLAIN` provided information only for `SELECT` statements. In addition, the `EXPLAIN` statement now can produce output in JSON format. See [Section 13.8.2](#), “`EXPLAIN` Syntax”.
- The optimizer more efficiently handles subqueries in the `FROM` clause (that is, derived tables). Materialization of subqueries in the `FROM` clause is postponed until their contents are needed during query execution, which improves performance. In addition, during query execution, the optimizer may add an index to a derived table to speed up row retrieval from it. For more information, see [Section 8.13.16.3](#), “Optimizing Subqueries in the `FROM` Clause (Derived Tables)”.
- The optimizer uses semi-join and materialization strategies to optimize subquery execution. See [Section 8.13.16.1](#), “Optimizing Subqueries with Semi-Join Transformations”, and [Section 8.13.16.2](#), “Optimizing Subqueries with Subquery Materialization”.
- A Batched Key Access (BKA) join algorithm is now available that uses both index access to the joined table and a join buffer. The BKA algorithm supports inner join, outer join, and semi-join operations, including nested outer joins and nested semi-joins. Benefits of BKA include improved join performance due to more efficient table scanning. For more information, see [Section 8.13.12](#), “Block Nested-Loop and Batched Key Access Joins”.
- The optimizer now has a tracing capability, primarily for use by developers. The interface is provided by a set of `optimizer_trace_xxx` system variables and the `INFORMATION_SCHEMA.OPTIMIZER_TRACE` table. For details, see [MySQL Internals: Tracing the Optimizer](#).
- **Condition handling.** MySQL now supports the `GET DIAGNOSTICS` statement. `GET DIAGNOSTICS` provides applications a standardized way to obtain information from the diagnostics area, such as whether the previous SQL statement produced an exception and what it was. For more information, see [Section 13.6.7.3](#), “`GET DIAGNOSTICS` Syntax”.

In addition, several deficiencies in condition handler processing rules were corrected so that MySQL behavior is more like standard SQL:

- Block scope is used in determining which handler to select. Previously, a stored program was treated as having a single scope for handler selection.
- Condition precedence is more accurately resolved.
- Diagnostics area clearing has changed. Bug #55843 caused handled conditions to be cleared from the diagnostics area before activating the handler. This made condition information unavailable within the handler. Now condition information is available to the handler, which can inspect it with the `GET DIAGNOSTICS` statement. The condition information is cleared when the handler exits, if it has not already been cleared during handler execution.

- Previously, handlers were activated as soon as a condition occurred. Now they are not activated until the statement in which the condition occurred finishes execution, at which point the most appropriate handler is chosen. This can make a difference for statements that raise multiple conditions, if a condition raised later during statement execution has higher precedence than an earlier condition and there are handlers in the same scope for both conditions. Previously, the handler for the first condition raised would be chosen, even if it had a lower precedence than other handlers. Now the handler for the condition with highest precedence is chosen, even if it is not the first condition raised by the statement.

For more information, see [Section 13.6.7.6, “Scope Rules for Handlers”](#).

- **Data types.** These data type changes have been implemented:
  - MySQL now permits fractional seconds for `TIME`, `DATETIME`, and `TIMESTAMP` values, with up to microseconds (6 digits) precision. See [Section 11.3.6, “Fractional Seconds in Time Values”](#).
  - Previously, at most one `TIMESTAMP` column per table could be automatically initialized or updated to the current date and time. This restriction has been lifted. Any `TIMESTAMP` column definition can have any combination of `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` clauses. In addition, these clauses now can be used with `DATETIME` column definitions. For more information, see [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#).
  - In MySQL, the `TIMESTAMP` data type differs in nonstandard ways from other data types in terms of default value and assignment of automatic initialization and update attributes. These behaviors remain the default but now are deprecated, and can be turned off by enabling the `explicit_defaults_for_timestamp` [481] system variable at server startup. See [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#), and [Section 5.1.4, “Server System Variables”](#).
  - The `YEAR(2)` data type is now deprecated. `YEAR(2)` columns in existing tables are treated as before, but `YEAR(2)` in new or altered tables are converted to `YEAR(4)`. For more information, see [Section 11.3.4, “`YEAR\(2\)` Limitations and Migrating to `YEAR\(4\)`”](#).
- **Host cache.** MySQL now provides more information about the causes of errors that occur when clients connect to the server, as well as improved access to the host cache, which contains client IP address and host name information and is used to avoid DNS lookups. These changes have been implemented:
  - New `Connection_errors_xxx` [587] status variables provide information about connection errors that do not apply to specific client IP addresses.
  - Counters have been added to the host cache to track errors that do apply to specific IP addresses, and a new `host_cache` Performance Schema table exposes the contents of the host cache so that it can be examined using `SELECT` statements. Access to host cache contents makes it possible to answer questions such as how many hosts are cached, what kinds of connection errors are occurring for which hosts, or how close host error counts are to reaching the `max_connect_errors` [502] system variable limit.
  - The host cache size now is configurable using the `host_cache_size` [487] system variable.

For more information, see [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#), and [Section 20.9.9.1, “The `host\_cache` Table”](#).

- **OpenGIS.** The OpenGIS specification defines functions that test the relationship between two geometry values. MySQL originally implemented these functions such that they used object bounding rectangles and returned the same result as the corresponding MBR-based functions. Corresponding versions are now available that use precise object shapes. These versions are named with an `ST_` prefix. For example, `Contains()` [1311] uses object bounding

rectangles, whereas `ST_Contains()` [1310] uses object shapes. For more information, see [Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”](#).

## Removed Features

The following constructs are obsolete and have been removed in MySQL 5.6. Where alternatives are shown, applications should be updated to use them.

- The `--log` server option and the `log` system variable. Instead, use the `--general_log` [485] option to enable the general query log and the `--general_log_file=file_name` [486] option to set the general query log file name.
- The `--log-slow-queries` server option and the `log_slow_queries` system variable. Instead, use the `--slow_query_log` [541] option to enable the slow query log and the `--slow_query_log_file=file_name` [541] option to set the slow query log file name.
- The `--one-thread` server option. Use `--thread_handling=no-threads` [553] instead.
- The `--safe-mode` server option.
- The `--skip-thread-priority` server option.
- The `--table-cache` server option. Use the `table_open_cache` [551] system variable instead.
- The `--init-rpl-role` and `--rpl-recovery-rank` options, the `rpl_recovery_rank` system variable, and the `Rpl_status` status variable.
- The `engine_condition_pushdown` system variable. Use the `engine_condition_pushdown` flag of the `optimizer_switch` [520] variable instead.
- The `have_csv`, `have_innodb`, `have_ndbcluster`, and `have_partitioning` system variables. Use `SHOW ENGINES` instead.
- The `sql_big_tables` system variable. Use `big_tables` [466] instead.
- The `sql_low_priority_updates` system variable. Use `low_priority_updates` [500] instead.
- The `sql_max_join_size` system variable. Use `max_join_size` [505] instead.
- The `max_long_data_size` system variable. Use `max_allowed_packet` [502] instead.
- The `FLUSH MASTER` and `FLUSH SLAVE` statements. Use the `RESET MASTER` and `RESET SLAVE` statements instead.
- The `SLAVE START` and `SLAVE STOP` statements. Use The `START SLAVE` and `STOP SLAVE` statements.
- The `SHOW AUTHORS` and `SHOW CONTRIBUTORS` statements.
- The `OPTION` and `ONE_SHOT` modifiers for the `SET` statement.
- It is explicitly disallowed to assign the value `DEFAULT` to stored procedure or function parameters or stored program local variables (for example with a `SET var_name = DEFAULT` statement). It remains permissible to assign `DEFAULT` to system variables, as before.

## 1.5. MySQL Development History

This section describes the general MySQL development history, provides an overview about features that have been implemented in previous series and that are new in MySQL 5.6, the release series covered in this manual. The maturity level this release series is general availability. Information about maturity levels can be found in [Section 2.1.2.1, “Choosing Which Version of MySQL to Install”](#).

Before upgrading from one release series to the next, please see the notes in [Section 2.11.1](#), “Upgrading MySQL”.

The most requested features and the versions in which they were implemented are summarized in the following table.

Feature	MySQL Series
Unions	4.0
Subqueries	4.1
R-trees	4.1 (for the <a href="#">MyISAM</a> storage engine)
Stored procedures and functions	5.0
Views	5.0
Cursors	5.0
XA transactions	5.0
Triggers	5.0 and 5.1
Event scheduler	5.1
Partitioning	5.1
Pluggable storage engine API	5.1
Plugin API	5.1
Row-based replication	5.1
Server log tables	5.1
Scalability and performance improvements	5.1 (with InnoDB Plugin)
DTrace support	5.5
Semisynchronous replication	5.5
SIGNAL/RESIGNAL support in stored routines	5.5
Performance Schema	5.5
Supplementary Unicode characters	5.5

## 1.6. MySQL Information Sources

This section lists sources of additional information that you may find helpful, such as the MySQL mailing lists and user forums, and Internet Relay Chat.

### 1.6.1. MySQL Mailing Lists

This section introduces the MySQL mailing lists and provides guidelines as to how the lists should be used. When you subscribe to a mailing list, you receive all postings to the list as email messages. You can also send your own questions and answers to the list.

To subscribe to or unsubscribe from any of the mailing lists described in this section, visit <http://lists.mysql.com/>. For most of them, you can select the regular version of the list where you get individual messages, or a digest version where you get one large message per day.

Please *do not* send messages about subscribing or unsubscribing to any of the mailing lists, because such messages are distributed automatically to thousands of other users.

Your local site may have many subscribers to a MySQL mailing list. If so, the site may have a local mailing list, so that messages sent from [lists.mysql.com](http://lists.mysql.com) to your site are propagated to the local



list. In such cases, please contact your system administrator to be added to or dropped from the local MySQL list.

To have traffic for a mailing list go to a separate mailbox in your mail program, set up a filter based on the message headers. You can use either the `List-ID:` or `Delivered-To:` headers to identify list messages.

The MySQL mailing lists are as follows:

- [announce](#)

The list for announcements of new versions of MySQL and related programs. This is a low-volume list to which all MySQL users should subscribe.

- [mysql](#)

The main list for general MySQL discussion. Please note that some topics are better discussed on the more-specialized lists. If you post to the wrong list, you may not get an answer.

- [bugs](#)

The list for people who want to stay informed about issues reported since the last release of MySQL or who want to be actively involved in the process of bug hunting and fixing. See [Section 1.7, “How to Report Bugs or Problems”](#).

- [internals](#)

The list for people who work on the MySQL code. This is also the forum for discussions on MySQL development and for posting patches.

- [mysqldoc](#)

The list for people who work on the MySQL documentation.

- [benchmarks](#)

The list for anyone interested in performance issues. Discussions concentrate on database performance (not limited to MySQL), but also include broader categories such as performance of the kernel, file system, disk system, and so on.

- [packagers](#)

The list for discussions on packaging and distributing MySQL. This is the forum used by distribution maintainers to exchange ideas on packaging MySQL and on ensuring that MySQL looks and feels as similar as possible on all supported platforms and operating systems.

- [java](#)

The list for discussions about the MySQL server and Java. It is mostly used to discuss JDBC drivers such as MySQL Connector/J.

- [win32](#)

The list for all topics concerning the MySQL software on Microsoft operating systems, such as Windows 9x, Me, NT, 2000, XP, and 2003.

- [myodbc](#)

The list for all topics concerning connecting to the MySQL server with ODBC.

- [gui-tools](#)

The list for all topics concerning MySQL graphical user interface tools such as MySQL Workbench.

- [cluster](#)

The list for discussion of MySQL Cluster.

- [dotnet](#)

The list for discussion of the MySQL server and the .NET platform. It is mostly related to MySQL Connector/Net.

- [plusplus](#)

The list for all topics concerning programming with the C++ API for MySQL.

- [perl](#)

The list for all topics concerning Perl support for MySQL with `DBD: :mysql`.

If you're unable to get an answer to your questions from a MySQL mailing list or forum, one option is to purchase support from Oracle. This puts you in direct contact with MySQL developers.

The following MySQL mailing lists are in languages other than English. These lists are not operated by Oracle.

- [<mysql-france-subscribe@yahoogroups.com>](mailto:mysql-france-subscribe@yahoogroups.com)

A French mailing list.

- [<list@tinc.net>](mailto:list@tinc.net)

A Korean mailing list. To subscribe, email `subscribe mysql your@email.address` to this list.

- [<mysql-de-request@lists.4t2.com>](mailto:mysql-de-request@lists.4t2.com)

A German mailing list. To subscribe, email `subscribe mysql-de your@email.address` to this list. You can find information about this mailing list at <http://www.4t2.com/mysql/>.

- [<mysql-br-request@listas.linkway.com.br>](mailto:mysql-br-request@listas.linkway.com.br)

A Portuguese mailing list. To subscribe, email `subscribe mysql-br your@email.address` to this list.

- [<mysql-alta@elistas.net>](mailto:mysql-alta@elistas.net)

A Spanish mailing list. To subscribe, email `subscribe mysql your@email.address` to this list.

### 1.6.1.1. Guidelines for Using the Mailing Lists

Please do not post mail messages from your browser with HTML mode turned on. Many users do not read mail with a browser.

When you answer a question sent to a mailing list, if you consider your answer to have broad interest, you may want to post it to the list instead of replying directly to the individual who asked. Try to make your answer general enough that people other than the original poster may benefit from it. When you post to the list, please make sure that your answer is not a duplication of a previous answer.

Try to summarize the essential part of the question in your reply. Do not feel obliged to quote the entire original message.

When answers are sent to you individually and not to the mailing list, it is considered good etiquette to summarize the answers and send the summary to the mailing list so that others may have the benefit of responses you received that helped you solve your problem.

### 1.6.2. MySQL Community Support at the MySQL Forums

The forums at <http://forums.mysql.com> are an important community resource. Many forums are available, grouped into these general categories:

- Migration
- MySQL Usage
- MySQL Connectors
- Programming Languages
- Tools
- 3rd-Party Applications
- Storage Engines
- MySQL Technology
- SQL Standards
- Business

### 1.6.3. MySQL Community Support on Internet Relay Chat (IRC)

In addition to the various MySQL mailing lists and forums, you can find experienced community people on Internet Relay Chat (IRC). These are the best networks/channels currently known to us:

**freenode** (see <http://www.freenode.net/> for servers)

- `#mysql` is primarily for MySQL questions, but other database and general SQL questions are welcome. Questions about PHP, Perl, or C in combination with MySQL are also common.

If you are looking for IRC client software to connect to an IRC network, take a look at **xChat** (<http://www.xchat.org/>). X-Chat (GPL licensed) is available for Unix as well as for Windows platforms (a free Windows build of X-Chat is available at <http://www.silverex.org/download/>).

### 1.6.4. MySQL Enterprise

Oracle offers technical support in the form of MySQL Enterprise. For organizations that rely on the MySQL DBMS for business-critical production applications, MySQL Enterprise is a commercial subscription offering which includes:

- MySQL Enterprise Server
- MySQL Enterprise Monitor
- Monthly Rapid Updates and Quarterly Service Packs
- MySQL Knowledge Base
- 24x7 Technical and Consultative Support

MySQL Enterprise is available in multiple tiers, giving you the flexibility to choose the level of service that best matches your needs. For more information, see [MySQL Enterprise](#).

## 1.7. How to Report Bugs or Problems

Before posting a bug report about a problem, please try to verify that it is a bug and that it has not been reported already:

- Start by searching the MySQL online manual at <http://dev.mysql.com/doc/>. We try to keep the manual up to date by updating it frequently with solutions to newly found problems. In addition, the

release notes accompanying the manual can be particularly useful since it is quite possible that a newer version contains a solution to your problem. The release notes are available at the location just given for the manual.

- If you get a parse error for an SQL statement, please check your syntax closely. If you cannot find something wrong with it, it is extremely likely that your current version of MySQL Server doesn't support the syntax you are using. If you are using the current version and the manual doesn't cover the syntax that you are using, MySQL Server doesn't support your statement.

If the manual covers the syntax you are using, but you have an older version of MySQL Server, you should check the MySQL change history to see when the syntax was implemented. In this case, you have the option of upgrading to a newer version of MySQL Server.

- For solutions to some common problems, see [Section C.5, "Problems and Common Errors"](#).
- Search the bugs database at <http://bugs.mysql.com/> to see whether the bug has been reported and fixed.
- Search the MySQL mailing list archives at <http://lists.mysql.com/>. See [Section 1.6.1, "MySQL Mailing Lists"](#).
- You can also use <http://www.mysql.com/search/> to search all the Web pages (including the manual) that are located at the MySQL Web site.

If you cannot find an answer in the manual, the bugs database, or the mailing list archives, check with your local MySQL expert. If you still cannot find an answer to your question, please use the following guidelines for reporting the bug.

The normal way to report bugs is to visit <http://bugs.mysql.com/>, which is the address for our bugs database. This database is public and can be browsed and searched by anyone. If you log in to the system, you can enter new reports.

Bugs posted in the bugs database at <http://bugs.mysql.com/> that are corrected for a given release are noted in the release notes.

If you find a sensitive security bug in MySQL Server, please let us know immediately by sending an email message to [<secalert\\_us@oracle.com>](mailto:secalert_us@oracle.com). Exception: Support customers should report all problems, including security bugs, to Oracle Support at <http://support.oracle.com/>.

To discuss problems with other users, you can use one of the MySQL mailing lists. [Section 1.6.1, "MySQL Mailing Lists"](#).

Writing a good bug report takes patience, but doing it right the first time saves time both for us and for yourself. A good bug report, containing a full test case for the bug, makes it very likely that we will fix the bug in the next release. This section helps you write your report correctly so that you do not waste your time doing things that may not help us much or at all. Please read this section carefully and make sure that all the information described here is included in your report.

Preferably, you should test the problem using the latest production or development version of MySQL Server before posting. Anyone should be able to repeat the bug by just using `mysql test <script_file` on your test case or by running the shell or Perl script that you include in the bug report. Any bug that we are able to repeat has a high chance of being fixed in the next MySQL release.

It is most helpful when a good description of the problem is included in the bug report. That is, give a good example of everything you did that led to the problem and describe, in exact detail, the problem itself. The best reports are those that include a full example showing how to reproduce the bug or problem. See [MySQL Internals: Porting to Other Systems](#).

Remember that it is possible for us to respond to a report containing too much information, but not to one containing too little. People often omit facts because they think they know the cause of a problem

and assume that some details do not matter. A good principle to follow is that if you are in doubt about stating something, state it. It is faster and less troublesome to write a couple more lines in your report than to wait longer for the answer if we must ask you to provide information that was missing from the initial report.

The most common errors made in bug reports are (a) not including the version number of the MySQL distribution that you use, and (b) not fully describing the platform on which the MySQL server is installed (including the platform type and version number). These are highly relevant pieces of information, and in 99 cases out of 100, the bug report is useless without them. Very often we get questions like, “Why doesn't this work for me?” Then we find that the feature requested wasn't implemented in that MySQL version, or that a bug described in a report has been fixed in newer MySQL versions. Errors often are platform-dependent. In such cases, it is next to impossible for us to fix anything without knowing the operating system and the version number of the platform.

If you compiled MySQL from source, remember also to provide information about your compiler if it is related to the problem. Often people find bugs in compilers and think the problem is MySQL-related. Most compilers are under development all the time and become better version by version. To determine whether your problem depends on your compiler, we need to know what compiler you used. Note that every compiling problem should be regarded as a bug and reported accordingly.

If a program produces an error message, it is very important to include the message in your report. If we try to search for something from the archives, it is better that the error message reported exactly matches the one that the program produces. (Even the lettercase should be observed.) It is best to copy and paste the entire error message into your report. You should never try to reproduce the message from memory.

If you have a problem with Connector/ODBC (MyODBC), please try to generate a trace file and send it with your report. See [Section 21.1.8.2, “How to Report Connector/ODBC Problems or Bugs”](#).

If your report includes long query output lines from test cases that you run with the `mysql` command-line tool, you can make the output more readable by using the `--vertical` [254] option or the `\G` statement terminator. The `EXPLAIN SELECT` example later in this section demonstrates the use of `\G`.

Please include the following information in your report:

- The version number of the MySQL distribution you are using (for example, MySQL 5.0.19). You can find out which version you are running by executing `mysqladmin version`. The `mysqladmin` program can be found in the `bin` directory under your MySQL installation directory.
- The manufacturer and model of the machine on which you experience the problem.
- The operating system name and version. If you work with Windows, you can usually get the name and version number by double-clicking your My Computer icon and pulling down the “Help/About Windows” menu. For most Unix-like operating systems, you can get this information by executing the command `uname -a`.
- Sometimes the amount of memory (real and virtual) is relevant. If in doubt, include these values.
- If you are using a source distribution of the MySQL software, include the name and version number of the compiler that you used. If you have a binary distribution, include the distribution name.
- If the problem occurs during compilation, include the exact error messages and also a few lines of context around the offending code in the file where the error occurs.
- If `mysqld` died, you should also report the statement that crashed `mysqld`. You can usually get this information by running `mysqld` with query logging enabled, and then looking in the log after `mysqld` crashes. See [MySQL Internals: Porting to Other Systems](#).
- If a database table is related to the problem, include the output from the `SHOW CREATE TABLE db_name.tbl_name` statement in the bug report. This is a very easy way to get the definition of

any table in a database. The information helps us create a situation matching the one that you have experienced.

- The SQL mode in effect when the problem occurred can be significant, so please report the value of the `sql_mode` [545] system variable. For stored procedure, stored function, and trigger objects, the relevant `sql_mode` [545] value is the one in effect when the object was created. For a stored procedure or function, the `SHOW CREATE PROCEDURE` or `SHOW CREATE FUNCTION` statement shows the relevant SQL mode, or you can query `INFORMATION_SCHEMA` for the information:

```
SELECT ROUTINE_SCHEMA, ROUTINE_NAME, SQL_MODE
FROM INFORMATION_SCHEMA.ROUTINES;
```

For triggers, you can use this statement:

```
SELECT EVENT_OBJECT_SCHEMA, EVENT_OBJECT_TABLE, TRIGGER_NAME, SQL_MODE
FROM INFORMATION_SCHEMA.TRIGGERS;
```

- For performance-related bugs or problems with `SELECT` statements, you should always include the output of `EXPLAIN SELECT ...`, and at least the number of rows that the `SELECT` statement produces. You should also include the output from `SHOW CREATE TABLE tbl_name` for each table that is involved. The more information you provide about your situation, the more likely it is that someone can help you.

The following is an example of a very good bug report. The statements are run using the `mysql` command-line tool. Note the use of the `\G` statement terminator for statements that would otherwise provide very long output lines that are difficult to read.

```
mysql> SHOW VARIABLES;
mysql> SHOW COLUMNS FROM ...\G
<output from SHOW COLUMNS>
mysql> EXPLAIN SELECT ...\G
<output from EXPLAIN>
mysql> FLUSH STATUS;
mysql> SELECT ...;
<A short version of the output from SELECT,
including the time taken to run the query>
mysql> SHOW STATUS;
<output from SHOW STATUS>
```

- If a bug or problem occurs while running `mysqld`, try to provide an input script that reproduces the anomaly. This script should include any necessary source files. The more closely the script can reproduce your situation, the better. If you can make a reproducible test case, you should upload it to be attached to the bug report.

If you cannot provide a script, you should at least include the output from `mysqladmin variables extended-status processlist` in your report to provide some information on how your system is performing.

- If you cannot produce a test case with only a few rows, or if the test table is too big to be included in the bug report (more than 10 rows), you should dump your tables using `mysqldump` and create a `README` file that describes your problem. Create a compressed archive of your files using `tar` and `gzip` or `zip`. After you initiate a bug report for our bugs database at <http://bugs.mysql.com/>, click the Files tab in the bug report for instructions on uploading the archive to the bugs database.
- If you believe that the MySQL server produces a strange result from a statement, include not only the result, but also your opinion of what the result should be, and an explanation describing the basis for your opinion.
- When you provide an example of the problem, it is better to use the table names, variable names, and so forth that exist in your actual situation than to come up with new names. The problem could be related to the name of a table or variable. These cases are rare, perhaps, but it is better to be safe than sorry. After all, it should be easier for you to provide an example that uses your actual situation, and it is by all means better for us. If you have data that you do not want to be visible to others in the bug report, you can upload it using the Files tab as previously described. If the



information is really top secret and you do not want to show it even to us, go ahead and provide an example using other names, but please regard this as the last choice.

- Include all the options given to the relevant programs, if possible. For example, indicate the options that you use when you start the `mysqld` server, as well as the options that you use to run any MySQL client programs. The options to programs such as `mysqld` and `mysql`, and to the `configure` script, are often key to resolving problems and are very relevant. It is never a bad idea to include them. If your problem involves a program written in a language such as Perl or PHP, please include the language processor's version number, as well as the version for any modules that the program uses. For example, if you have a Perl script that uses the `DBI` and `DBD: :mysql` modules, include the version numbers for Perl, `DBI`, and `DBD: :mysql`.
- If your question is related to the privilege system, please include the output of `mysqlaccess`, the output of `mysqladmin reload`, and all the error messages you get when trying to connect. When you test your privileges, you should first run `mysqlaccess`. After this, execute `mysqladmin reload version` and try to connect with the program that gives you trouble. `mysqlaccess` can be found in the `bin` directory under your MySQL installation directory.
- If you have a patch for a bug, do include it. But do not assume that the patch is all we need, or that we can use it, if you do not provide some necessary information such as test cases showing the bug that your patch fixes. We might find problems with your patch or we might not understand it at all. If so, we cannot use it.

If we cannot verify the exact purpose of the patch, we will not use it. Test cases help us here. Show that the patch handles all the situations that may occur. If we find a borderline case (even a rare one) where the patch will not work, it may be useless.

- Guesses about what the bug is, why it occurs, or what it depends on are usually wrong. Even the MySQL team cannot guess such things without first using a debugger to determine the real cause of a bug.
- Indicate in your bug report that you have checked the reference manual and mail archive so that others know you have tried to solve the problem yourself.
- If your data appears corrupt or you get errors when you access a particular table, first check your tables with `CHECK TABLE`. If that statement reports any errors:
  - The `InnoDB` crash recovery mechanism handles cleanup when the server is restarted after being killed, so in typical operation there is no need to “repair” tables. If you encounter an error with `InnoDB` tables, restart the server and see whether the problem persists, or whether the error affected only cached data in memory. If data is corrupted on disk, consider restarting with the `innodb_force_recovery [1722]` option enabled so that you can dump the affected tables.
  - For non-transactional tables, try to repair them with `REPAIR TABLE` or with `myisamchk`. See [Chapter 5, MySQL Server Administration](#).

If you are running Windows, please verify the value of `lower_case_table_names [501]` using the `SHOW VARIABLES LIKE 'lower_case_table_names'` statement. This variable affects how the server handles lettercase of database and table names. Its effect for a given value should be as described in [Section 9.2.2, “Identifier Case Sensitivity”](#).

- If you often get corrupted tables, you should try to find out when and why this happens. In this case, the error log in the MySQL data directory may contain some information about what happened. (This is the file with the `.err` suffix in the name.) See [Section 5.2.2, “The Error Log”](#). Please include any relevant information from this file in your bug report. Normally `mysqld` should *never* crash a table if nothing killed it in the middle of an update. If you can find the cause of `mysqld` dying, it is much easier for us to provide you with a fix for the problem. See [Section C.5.1, “How to Determine What Is Causing a Problem”](#).
- If possible, download and install the most recent version of MySQL Server and check whether it solves your problem. All versions of the MySQL software thoroughly tested and should work without

problems. We believe in making everything as backward-compatible as possible, and you should be able to switch MySQL versions without difficulty. See [Section 2.1.2, “Choosing Which MySQL Distribution to Install”](#).

## 1.8. MySQL Standards Compliance

This section describes how MySQL relates to the ANSI/ISO SQL standards. MySQL Server has many extensions to the SQL standard, and here you can find out what they are and how to use them. You can also find information about functionality missing from MySQL Server, and how to work around some of the differences.

The SQL standard has been evolving since 1986 and several versions exist. In this manual, “SQL-92” refers to the standard released in 1992, “SQL:1999” refers to the standard released in 1999, “SQL:2003” refers to the standard released in 2003, and “SQL:2008” refers to the most recent version of the standard, released in 2008. We use the phrase “the SQL standard” or “standard SQL” to mean the current version of the SQL Standard at any time.

One of our main goals with the product is to continue to work toward compliance with the SQL standard, but without sacrificing speed or reliability. We are not afraid to add extensions to SQL or support for non-SQL features if this greatly increases the usability of MySQL Server for a large segment of our user base. The `HANDLER` interface is an example of this strategy. See [Section 13.2.4, “HANDLER Syntax”](#).

We continue to support transactional and nontransactional databases to satisfy both mission-critical 24/7 usage and heavy Web or logging usage.

MySQL Server was originally designed to work with medium-sized databases (10-100 million rows, or about 100MB per table) on small computer systems. Today MySQL Server handles terabyte-sized databases, but the code can also be compiled in a reduced version suitable for hand-held and embedded devices. The compact design of the MySQL server makes development in both directions possible without any conflicts in the source tree.

Currently, we are not targeting real-time support, although MySQL replication capabilities offer significant functionality.

MySQL supports high-availability database clustering using the `NDBCLUSTER` storage engine. See [MySQL Cluster NDB 7.2](#).

We implement XML functionality which supports most of the W3C XPath standard. See [Section 12.11, “XML Functions”](#).

### 1.8.1. What Standards MySQL Follows

Our aim is to support the full ANSI/ISO SQL standard, but without making concessions to speed and quality of the code.

ODBC levels 0 to 3.51.

### 1.8.2. Selecting SQL Modes

The MySQL server can operate in different SQL modes, and can apply these modes differentially for different clients. This capability enables each application to tailor the server's operating mode to its own requirements.

SQL modes control aspects of server operation such as what SQL syntax MySQL should support and what kind of data validation checks it should perform. This makes it easier to use MySQL in different environments and to use MySQL together with other database servers.

You can set the default SQL mode by starting `mysqld` with the `--sql-mode="mode_value"` [\[442\]](#) option. You can also change the mode at runtime by setting the `sql_mode` [\[545\]](#) system variable with a `SET [GLOBAL|SESSION] sql_mode='mode_value'` statement.

For more information on setting the SQL mode, see [Section 5.1.7, “Server SQL Modes”](#).

### 1.8.3. Running MySQL in ANSI Mode

You can tell `mysqld` to run in ANSI mode with the `--ansi` [413] startup option. Running the server in ANSI mode is the same as starting it with the following options:

```
--transaction-isolation=SERIALIZABLE --sql-mode=ANSI
```

You can achieve the same effect at runtime by executing these two statements:

```
SET GLOBAL TRANSACTION ISOLATION LEVEL SERIALIZABLE;
SET GLOBAL sql_mode = 'ANSI';
```

You can see that setting the `sql_mode` [545] system variable to `'ANSI'` enables all SQL mode options that are relevant for ANSI mode as follows:

```
mysql> SET GLOBAL sql_mode='ANSI';
mysql> SELECT @@global.sql_mode;
-> 'REAL_AS_FLOAT,PIPES_AS_CONCAT,ANSI_QUOTES,IGNORE_SPACE,ANSI'
```

Running the server in ANSI mode with `--ansi` [413] is not quite the same as setting the SQL mode to `'ANSI'`. The `--ansi` [413] option affects the SQL mode and also sets the transaction isolation level. Setting the SQL mode to `'ANSI'` has no effect on the isolation level.

See [Section 5.1.3, “Server Command Options”](#), and [Section 1.8.2, “Selecting SQL Modes”](#).

### 1.8.4. MySQL Extensions to Standard SQL

MySQL Server supports some extensions that you probably won't find in other SQL DBMSs. Be warned that if you use them, your code won't be portable to other SQL servers. In some cases, you can write code that includes MySQL extensions, but is still portable, by using comments of the following form:

```
/*! MySQL-specific code */
```

In this case, MySQL Server parses and executes the code within the comment as it would any other SQL statement, but other SQL servers will ignore the extensions. For example, MySQL Server recognizes the `STRAIGHT_JOIN` keyword in the following statement, but other servers will not:

```
SELECT /*! STRAIGHT_JOIN */ col1 FROM table1,table2 WHERE ...
```

If you add a version number after the `!` character, the syntax within the comment is executed only if the MySQL version is greater than or equal to the specified version number. The `TEMPORARY` keyword in the following comment is executed only by servers from MySQL 3.23.02 or higher:

```
CREATE /*!32302 TEMPORARY */ TABLE t (a INT);
```

The following descriptions list MySQL extensions, organized by category.

- Organization of data on disk

MySQL Server maps each database to a directory under the MySQL data directory, and maps tables within a database to file names in the database directory. This has a few implications:

- Database and table names are case sensitive in MySQL Server on operating systems that have case-sensitive file names (such as most Unix systems). See [Section 9.2.2, “Identifier Case Sensitivity”](#).
- You can use standard system commands to back up, rename, move, delete, and copy tables that are managed by the `MyISAM` storage engine. For example, it is possible to rename a `MyISAM` table

by renaming the `.MYD`, `.MYI`, and `.frm` files to which the table corresponds. (Nevertheless, it is preferable to use `RENAME TABLE` or `ALTER TABLE ... RENAME` and let the server rename the files.)

- General language syntax
  - By default, strings can be enclosed by either “`”` or “`'`”, not just by “`'`”. (If the [ANSI\\_QUOTES \[603\]](#) SQL mode is enabled, strings can be enclosed only by “`'`” and the server interprets strings enclosed by “`”` as identifiers.)
  - “`\`” is the escape character in strings.
  - In SQL statements, you can access tables from different databases with the `db_name.tbl_name` syntax. Some SQL servers provide the same functionality but call this `User space`. MySQL Server doesn't support tablespaces such as used in statements like this: `CREATE TABLE ralph.my_table ... IN my_tablespace`.
- SQL statement syntax
  - The `ANALYZE TABLE`, `CHECK TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements.
  - The `CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE` statements. See [Section 13.1.8, “CREATE DATABASE Syntax”](#), [Section 13.1.17, “DROP DATABASE Syntax”](#), and [Section 13.1.1, “ALTER DATABASE Syntax”](#).
  - The `DO` statement.
  - `EXPLAIN SELECT` to obtain a description of how tables are processed by the query optimizer.
  - The `FLUSH` and `RESET` statements.
  - The `SET` statement. See [Section 13.7.4, “SET Syntax”](#).
  - The `SHOW` statement. See [Section 13.7.5, “SHOW Syntax”](#). The information produced by many of the MySQL-specific `SHOW` statements can be obtained in more standard fashion by using `SELECT` to query `INFORMATION_SCHEMA`. See [Chapter 19, \*INFORMATION\\_SCHEMA Tables\*](#).
  - Use of `LOAD DATA INFILE`. In many cases, this syntax is compatible with Oracle's `LOAD DATA INFILE`. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).
  - Use of `RENAME TABLE`. See [Section 13.1.26, “RENAME TABLE Syntax”](#).
  - Use of `REPLACE` instead of `DELETE` plus `INSERT`. See [Section 13.2.8, “REPLACE Syntax”](#).
  - Use of `CHANGE col_name`, `DROP col_name`, or `DROP INDEX, IGNORE` or `RENAME` in `ALTER TABLE` statements. Use of multiple `ADD`, `ALTER`, `DROP`, or `CHANGE` clauses in an `ALTER TABLE` statement. See [Section 13.1.6, “ALTER TABLE Syntax”](#).
  - Use of index names, indexes on a prefix of a column, and use of `INDEX` or `KEY` in `CREATE TABLE` statements. See [Section 13.1.14, “CREATE TABLE Syntax”](#).
  - Use of `TEMPORARY` or `IF NOT EXISTS` with `CREATE TABLE`.
  - Use of `IF EXISTS` with `DROP TABLE` and `DROP DATABASE`.
  - The capability of dropping multiple tables with a single `DROP TABLE` statement.
  - The `ORDER BY` and `LIMIT` clauses of the `UPDATE` and `DELETE` statements.
  - `INSERT INTO tbl_name SET col_name = ...` syntax.
  - The `DELAYED` clause of the `INSERT` and `REPLACE` statements.

- The `LOW_PRIORITY` clause of the `INSERT`, `REPLACE`, `DELETE`, and `UPDATE` statements.
- Use of `INTO OUTFILE` or `INTO DUMPFILE` in `SELECT` statements. See [Section 13.2.9](#), “`SELECT` Syntax”.
- Options such as `STRAIGHT_JOIN` or `SQL_SMALL_RESULT` in `SELECT` statements.
- You don't need to name all selected columns in the `GROUP BY` clause. This gives better performance for some very specific, but quite normal queries. See [Section 12.17](#), “`Functions and Modifiers for Use with GROUP BY Clauses`”.
- You can specify `ASC` and `DESC` with `GROUP BY`, not just with `ORDER BY`.
- The ability to set variables in a statement with the `:=` assignment operator. See [Section 9.4](#), “`User-Defined Variables`”.
- Data types
  - The `MEDIUMINT`, `SET`, and `ENUM` data types, and the various `BLOB` and `TEXT` data types.
  - The `AUTO_INCREMENT`, `BINARY`, `NULL`, `UNSIGNED`, and `ZEROFILL` data type attributes.
- Functions and operators
  - To make it easier for users who migrate from other SQL environments, MySQL Server supports aliases for many functions. For example, all string functions support both standard SQL syntax and ODBC syntax.
  - MySQL Server understands the `||` [1172] and `&&` [1171] operators to mean logical OR and AND, as in the C programming language. In MySQL Server, `||` [1172] and `OR` [1172] are synonyms, as are `&&` [1171] and `AND` [1171]. Because of this nice syntax, MySQL Server doesn't support the standard SQL `||` [1172] operator for string concatenation; use `CONCAT()` [1178] instead. Because `CONCAT()` [1178] takes any number of arguments, it is easy to convert use of the `||` [1172] operator to MySQL Server.
  - Use of `COUNT(DISTINCT value_list)` [1283] where `value_list` has more than one element.
  - String comparisons are case-insensitive by default, with sort ordering determined by the collation of the current character set, which is `latin1` (cp1252 West European) by default. If you don't like this, you should declare your columns with the `BINARY` attribute or use the `BINARY` cast, which causes comparisons to be done using the underlying character code values rather than a lexical ordering.
  - The `%` [1205] operator is a synonym for `MOD()` [1205]. That is, `N % M` is equivalent to `MOD(N,M)` [1205]. `%` [1205] is supported for C programmers and for compatibility with PostgreSQL.
  - The `=` [1166], `<>` [1167], `<=` [1167], `<` [1167], `>=` [1167], `>` [1167], `<<` [1258], `>>` [1258], `<=>` [1167], `AND` [1171], `OR` [1172], or `LIKE` [1190] operators may be used in expressions in the output column list (to the left of the `FROM`) in `SELECT` statements. For example:

```
mysql> SELECT col1=1 AND col2=2 FROM my_table;
```
- The `LAST_INSERT_ID()` [1269] function returns the most recent `AUTO_INCREMENT` value. See [Section 12.14](#), “`Information Functions`”.
- `LIKE` [1190] is permitted on numeric values.
- The `REGEXP` [1193] and `NOT REGEXP` [1193] extended regular expression operators.

- `CONCAT()` [1178] or `CHAR()` [1177] with one argument or more than two arguments. (In MySQL Server, these functions can take a variable number of arguments.)
- The `BIT_COUNT()` [1259], `CASE` [1174], `ELT()` [1179], `FROM_DAYS()` [1218], `FORMAT()` [1180], `IF()` [1174], `PASSWORD()` [1263], `ENCRYPT()` [1262], `MD5()` [1263], `ENCODE()` [1262], `DECODE()` [1261], `PERIOD_ADD()` [1221], `PERIOD_DIFF()` [1221], `TO_DAYS()` [1225], and `WEEKDAY()` [1228] functions.
- Use of `TRIM()` [1186] to trim substrings. Standard SQL supports removal of single characters only.
- The `GROUP BY` functions `STD()` [1284], `BIT_OR()` [1282], `BIT_AND()` [1282], `BIT_XOR()` [1283], and `GROUP_CONCAT()` [1283]. See [Section 12.17, “Functions and Modifiers for Use with GROUP BY Clauses”](#).

## 1.8.5. MySQL Differences from Standard SQL

We try to make MySQL Server follow the ANSI SQL standard and the ODBC SQL standard, but MySQL Server performs operations differently in some cases:

- There are several differences between the MySQL and standard SQL privilege systems. For example, in MySQL, privileges for a table are not automatically revoked when you delete a table. You must explicitly issue a `REVOKE` statement to revoke privileges for a table. For more information, see [Section 13.7.1.6, “REVOKE Syntax”](#).
- The `CAST()` [1245] function does not support cast to `REAL` or `BIGINT`. See [Section 12.10, “Cast Functions and Operators”](#).

### 1.8.5.1. SELECT INTO TABLE Differences

MySQL Server doesn't support the `SELECT ... INTO TABLE` Sybase SQL extension. Instead, MySQL Server supports the `INSERT INTO ... SELECT` standard SQL syntax, which is basically the same thing. See [Section 13.2.5.1, “INSERT ... SELECT Syntax”](#). For example:

```
INSERT INTO tbl_temp2 (fld_id)
  SELECT tbl_temp1.fld_order_id
  FROM tbl_temp1 WHERE tbl_temp1.fld_order_id > 100;
```

Alternatively, you can use `SELECT ... INTO OUTFILE` or `CREATE TABLE ... SELECT`.

You can use `SELECT ... INTO` with user-defined variables. The same syntax can also be used inside stored routines using cursors and local variables. See [Section 13.2.9.1, “SELECT ... INTO Syntax”](#).

### 1.8.5.2. UPDATE Differences

If you access a column from the table to be updated in an expression, `UPDATE` uses the current value of the column. The second assignment in the following statement sets `col2` to the current (updated) `col1` value, not the original `col1` value. The result is that `col1` and `col2` have the same value. This behavior differs from standard SQL.

```
UPDATE t1 SET col1 = col1 + 1, col2 = col1;
```

### 1.8.5.3. Transaction and Atomic Operation Differences

MySQL Server (version 3.23-max and all versions 4.0 and above) supports transactions with the `InnoDB` transactional storage engine. In MySQL 5.5 and up, newly created tables use `InnoDB` by default, as explained in [Section 14.2.1.1, “InnoDB as the Default MySQL Storage Engine”](#). By default, `InnoDB` provides *full* ACID compliance; see [Section 14.2.8, “MySQL and the ACID Model”](#) for ways



that you can adjust settings to balance ACID compliance with raw performance. For information about [InnoDB](#) differences from standard SQL with regard to treatment of transaction errors, see [Section 14.2.3.14, “InnoDB Error Handling”](#).

The nontransactional storage engines in MySQL Server (such as [MyISAM](#)) follow a different paradigm for data integrity called “atomic operations”. [MyISAM](#) tables effectively always operate in `autocommit = 1` [464] mode. Because changed data is written to disk one statement at a time, it is harder to guarantee the consistency of a sequence of related DML operations, which could be interrupted partway through. Thus, this mode is suitable for read-mostly workloads. In transactional terms, while each specific update is running, no other user can interfere with it, there can never be an automatic rollback, and there are no dirty reads. However, these features apply to single operations, not related updates that succeed or fail as a unit. Workarounds such as the `LOCK TABLES` statement limit concurrent write access to nontransactional tables.

You can choose which paradigm to use, even for different tables within the same application: transactional features for reliability combined with high performance, or atomic operations for non-critical, read-mostly data (for example, on replication slave servers).

Transactional storage engines such as [InnoDB](#) offer many significant features to support high reliability for heavy read/write workloads. As a result, transactional tables can have higher memory and disk space requirements, and more CPU overhead. MySQL Server's modular design enables the concurrent use of different storage engines to suit different requirements and deliver optimum performance in all situations.

## Workarounds for Reliability with Non-Transactional Tables

But how do you use the features of MySQL Server to maintain integrity even with the nontransactional [MyISAM](#) tables, and how do these features compare with the transactional storage engines?

- If your applications are written in a way that is dependent on being able to call `ROLLBACK` rather than `COMMIT` in critical situations, transactions are more convenient. Transactions also ensure that unfinished updates or corrupting activities are not committed to the database; the server is given the opportunity to do an automatic rollback and your database is saved.

If you use nontransactional tables, you must resolve potential problems at the application level by including checks before updates and by running scripts that check the databases for inconsistencies and automatically repair or warn if such an inconsistency occurs. You can normally fix tables with no data integrity loss by using the MySQL log or even adding one extra log.

- Sometimes, critical transactional updates can be rewritten to be atomic. Multiple DML operations can be done with `LOCK TABLES` or atomic updates, ensuring that there are no deadlocks by limiting concurrent write access. If you obtain a `READ LOCAL` lock (as opposed to a write lock) for a table that enables concurrent inserts at the end of the table, reads are permitted, as are inserts by other clients. The newly inserted records are not be seen by the client that has the read lock until it releases the lock. With `INSERT DELAYED`, you can write inserts that go into a local queue until the locks are released, without having the client wait for the insert to complete. See [Section 8.10.3, “Concurrent Inserts”](#), and [Section 13.2.5.2, “INSERT DELAYED Syntax”](#).
- To be safe with MySQL Server, regardless of what kinds of tables you use, make regular backups and have binary logging turned on. It is always good to have backups, regardless of which database system you use.

Following are some techniques for working with nontransactional tables:

- Loops that need transactions normally can be coded with the help of `LOCK TABLES`, and you don't need cursors to update records on the fly.
- To avoid using `ROLLBACK`, you can employ the following strategy:
  1. Use `LOCK TABLES` to lock all the tables you want to access.

2. Test the conditions that must be true before performing the update.
3. Update if the conditions are satisfied.
4. Use `UNLOCK TABLES` to release your locks.



#### Note

This solution does not handle the situation when someone kills the threads in the middle of an update. In that case, all locks are released but some of the updates may not have been executed.

- You can also use functions to update records in a single operation, using the following techniques:
  - Modify columns relative to their current value. This makes the update correct even if another client has changed the column values in the meantime.
  - Update only those columns that actually have changed. This is a good database practice in general.
- When managing unique identifiers, you can avoid statements such as `LOCK TABLES` or `ROLLBACK` by using an `AUTO_INCREMENT` column and either the `LAST_INSERT_ID()` [1269] SQL function or the `mysql_insert_id()` C API function. See [Section 12.14, “Information Functions”](#), and [Section 21.8.6.37, “mysql\\_insert\\_id\(\)”](#).

For situations that require row-level locking, use `InnoDB` tables. Otherwise, with `MyISAM` tables, you can use a flag column in the table and do something like the following:

```
UPDATE tbl_name SET row_flag=1 WHERE id=ID;
```

MySQL returns 1 for the number of affected rows if the row was found and `row_flag` wasn't 1 in the original row. You can think of this as though MySQL Server changed the preceding statement to:

```
UPDATE tbl_name SET row_flag=1 WHERE id=ID AND row_flag <> 1;
```

### 1.8.5.4. Foreign Key Differences

The `InnoDB` storage engine supports checking of foreign key constraints, including `CASCADE`, `ON DELETE`, and `ON UPDATE`. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

For storage engines other than `InnoDB`, MySQL Server parses the `FOREIGN KEY` syntax in `CREATE TABLE` statements, but does not use or store it. In the future, the implementation will be extended to store this information in the table specification file so that it may be retrieved by `mysqldump` and ODBC. At a later stage, foreign key constraints will be implemented for `MyISAM` tables as well.

Foreign key enforcement offers several benefits to database developers:

- Assuming proper design of the relationships, foreign key constraints make it more difficult for a programmer to introduce an inconsistency into the database.
- Centralized checking of constraints by the database server makes it unnecessary to perform these checks on the application side. This eliminates the possibility that different applications may not all check the constraints in the same way.
- Using cascading updates and deletes can simplify the application code.
- Properly designed foreign key rules aid in documenting relationships between tables.

Do keep in mind that these benefits come at the cost of additional overhead for the database server to perform the necessary checks. Additional checking by the server affects performance, which for some applications may be sufficiently undesirable as to be avoided if possible. (Some major commercial applications have coded the foreign key logic at the application level for this reason.)

MySQL gives database developers the choice of which approach to use. If you don't need foreign keys and want to avoid the overhead associated with enforcing referential integrity, you can choose another storage engine instead, such as [MyISAM](#). (For example, the [MyISAM](#) storage engine offers very fast performance for applications that perform only [INSERT](#) and [SELECT](#) operations. In this case, the table has no holes in the middle and the inserts can be performed concurrently with retrievals. See [Section 8.10.3, "Concurrent Inserts"](#).)

If you choose not to take advantage of referential integrity checks, keep the following considerations in mind:

- In the absence of server-side foreign key relationship checking, the application itself must handle relationship issues. For example, it must take care to insert rows into tables in the proper order, and to avoid creating orphaned child records. It must also be able to recover from errors that occur in the middle of multiple-record insert operations.
- If [ON DELETE](#) is the only referential integrity capability an application needs, you can achieve a similar effect as of MySQL Server 4.0 by using multiple-table [DELETE](#) statements to delete rows from many tables with a single statement. See [Section 13.2.2, "DELETE Syntax"](#).
- A workaround for the lack of [ON DELETE](#) is to add the appropriate [DELETE](#) statements to your application when you delete records from a table that has a foreign key. In practice, this is often as quick as using foreign keys and is more portable.

Be aware that the use of foreign keys can sometimes lead to problems:

- Foreign key support addresses many referential integrity issues, but it is still necessary to design key relationships carefully to avoid circular rules or incorrect combinations of cascading deletes.
- It is not uncommon for a DBA to create a topology of relationships that makes it difficult to restore individual tables from a backup. (MySQL alleviates this difficulty by enabling you to temporarily disable foreign key checks when reloading a table that depends on other tables. See [Section 5.4.5, "InnoDB and FOREIGN KEY Constraints"](#). As of MySQL 4.1.1, [mysqldump](#) generates dump files that take advantage of this capability automatically when they are reloaded.)

Foreign keys in SQL are used to check and enforce referential integrity, not to join tables. If you want to get results from multiple tables from a [SELECT](#) statement, you do this by performing a join between them:

```
SELECT * FROM t1 INNER JOIN t2 ON t1.id = t2.id;
```

See [Section 13.2.9.2, "JOIN Syntax"](#), and [Section 3.6.6, "Using Foreign Keys"](#).

The [FOREIGN KEY](#) syntax without [ON DELETE . . .](#) is often used by ODBC applications to produce automatic [WHERE](#) clauses.

### 1.8.5.5. '--' as the Start of a Comment

Standard SQL uses the C syntax `/* this is a comment */` for comments, and MySQL Server supports this syntax as well. MySQL also support extensions to this syntax that enable MySQL-specific SQL to be embedded in the comment, as described in [Section 9.6, "Comment Syntax"](#).

Standard SQL uses `--` as a start-comment sequence. MySQL Server uses `#` as the start comment character. MySQL Server 3.23.3 and up also supports a variant of the `--` comment style. That is, the `--` start-comment sequence must be followed by a space (or by a control character such as a newline). The space is required to prevent problems with automatically generated SQL queries that use constructs such as the following, where we automatically insert the value of the payment for `payment`:

```
UPDATE account SET credit=credit-payment
```

Consider about what happens if `payment` has a negative value such as `-1`:

```
UPDATE account SET credit=credit--1
```

`credit--1` is a valid expression in SQL, but “--” is interpreted as the start of a comment, part of the expression is discarded. The result is a statement that has a completely different meaning than intended:

```
UPDATE account SET credit=credit
```

The statement produces no change in value at all. This illustrates that permitting comments to start with “--” can have serious consequences.

Using our implementation requires a space following the “--” for it to be recognized as a start-comment sequence in MySQL Server 3.23.3 and newer. Therefore, `credit--1` is safe to use.

Another safe feature is that the `mysql` command-line client ignores lines that start with “--”.

The following information is relevant only if you are running a MySQL version earlier than 3.23.3:

If you have an SQL script in a text file that contains “--” comments, you should use the `replace` utility as follows to convert the comments to use “#” characters before executing the script:

```
shell> replace " --" " #" < text-file-with-funny-comments.sql \
| mysql db_name
```

That is safer than executing the script in the usual way:

```
shell> mysql db_name < text-file-with-funny-comments.sql
```

You can also edit the script file “in place” to change the “--” comments to “#” comments:

```
shell> replace " --" " #" -- text-file-with-funny-comments.sql
```

Change them back with this command:

```
shell> replace " #" " --" -- text-file-with-funny-comments.sql
```

See [Section 4.8.2, “replace — A String-Replacement Utility”](#).

## 1.8.6. How MySQL Deals with Constraints

MySQL enables you to work both with transactional tables that permit rollback and with nontransactional tables that do not. Because of this, constraint handling is a bit different in MySQL than in other DBMSs. We must handle the case when you have inserted or updated a lot of rows in a nontransactional table for which changes cannot be rolled back when an error occurs.

The basic philosophy is that MySQL Server tries to produce an error for anything that it can detect while parsing a statement to be executed, and tries to recover from any errors that occur while executing the statement. We do this in most cases, but not yet for all.

The options MySQL has when an error occurs are to stop the statement in the middle or to recover as well as possible from the problem and continue. By default, the server follows the latter course. This means, for example, that the server may coerce invalid values to the closest valid values.

Several SQL mode options are available to provide greater control over handling of bad data values and whether to continue statement execution or abort when errors occur. Using these options, you can configure MySQL Server to act in a more traditional fashion that is like other DBMSs that reject improper input. The SQL mode can be set globally at server startup to affect all clients. Individual clients can set the SQL mode at runtime, which enables each client to select the behavior most appropriate for its requirements. See [Section 5.1.7, “Server SQL Modes”](#).

The following sections describe how MySQL Server handles different types of constraints.

### 1.8.6.1. PRIMARY KEY and UNIQUE Index Constraints

Normally, errors occurs for data-change statements (such as `INSERT` or `UPDATE`) that would violate primary-key, unique-key, or foreign-key constraints. If you are using a transactional storage engine such as `InnoDB`, MySQL automatically rolls back the statement. If you are using a nontransactional storage engine, MySQL stops processing the statement at the row for which the error occurred and leaves any remaining rows unprocessed.

MySQL supports an `IGNORE` keyword for `INSERT`, `UPDATE`, and so forth. If you use it, MySQL ignores primary-key or unique-key violations and continues processing with the next row. See the section for the statement that you are using ([Section 13.2.5, “INSERT Syntax”](#), [Section 13.2.11, “UPDATE Syntax”](#), and so forth).

You can get information about the number of rows actually inserted or updated with the `mysql_info()` C API function. You can also use the `SHOW WARNINGS` statement. See [Section 21.8.6.35, “mysql\\_info\(\)”](#), and [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

Currently, only `InnoDB` tables support foreign keys. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

### 1.8.6.2. FOREIGN KEY Constraints

Foreign keys let you cross-reference related data across tables, and [foreign key constraints](#) help keep this spread-out data consistent.

MySQL supports `ON UPDATE` and `ON DELETE` foreign key references in `CREATE TABLE` and `ALTER TABLE` statements. The available referential actions are `RESTRICT` (the default), `CASCADE`, `SET NULL`, and `NO ACTION`.

`SET DEFAULT` is also supported by the MySQL Server but is currently rejected as invalid by `InnoDB`. Since MySQL does not support deferred constraint checking, `NO ACTION` is treated as `RESTRICT`. For the exact syntax supported by MySQL for foreign keys, see [Section 13.1.14.2, “Using FOREIGN KEY Constraints”](#).

`MATCH FULL`, `MATCH PARTIAL`, and `MATCH SIMPLE` are allowed, but their use should be avoided, as they cause the MySQL Server to ignore any `ON DELETE` or `ON UPDATE` clause used in the same statement. `MATCH` options do not have any other effect in MySQL, which in effect enforces `MATCH SIMPLE` semantics full-time.

MySQL requires that foreign key columns be indexed; if you create a table with a foreign key constraint but no index on a given column, an index is created.

You can obtain information about foreign keys from the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table. An example of a query against this table is shown here:

```
mysql> SELECT TABLE_SCHEMA, TABLE_NAME, COLUMN_NAME, CONSTRAINT_NAME
> FROM INFORMATION_SCHEMA.KEY_COLUMN_USAGE
> WHERE REFERENCED_TABLE_SCHEMA IS NOT NULL;
```

TABLE_SCHEMA	TABLE_NAME	COLUMN_NAME	CONSTRAINT_NAME
fk1	myuser	myuser_id	f
fk1	product_order	customer_id	f2
fk1	product_order	product_id	f1

```
3 rows in set (0.01 sec)
```

Information about foreign keys on `InnoDB` tables can also be found in the `INNODB_SYS_FOREIGN` and `INNODB_SYS_FOREIGN_COLS` tables, in the `INFORMATION_SCHEMA` database.

Currently, only `InnoDB` tables support foreign keys. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#), for information specific to foreign key support in `InnoDB`.

## Deviations from SQL Standards

MySQL's implementation of foreign keys differs from the SQL standard in the following key respects:

- If there are several rows in the parent table that have the same referenced key value, [InnoDB](#) acts in foreign key checks as if the other parent rows with the same key value do not exist. For example, if you have defined a [RESTRICT](#) type constraint, and there is a child row with several parent rows, [InnoDB](#) does not permit the deletion of any of those parent rows.

[InnoDB](#) performs cascading operations through a depth-first algorithm, based on records in the indexes corresponding to the foreign key constraints.

- A [FOREIGN KEY](#) constraint that references a non-[UNIQUE](#) key is not standard SQL but rather an [InnoDB](#) extension.
- If [ON UPDATE CASCADE](#) or [ON UPDATE SET NULL](#) recurses to update the *same table* it has previously updated during the same cascade, it acts like [RESTRICT](#). This means that you cannot use self-referential [ON UPDATE CASCADE](#) or [ON UPDATE SET NULL](#) operations. This is to prevent infinite loops resulting from cascaded updates. A self-referential [ON DELETE SET NULL](#), on the other hand, is possible, as is a self-referential [ON DELETE CASCADE](#). Cascading operations may not be nested more than 15 levels deep.
- In an SQL statement that inserts, deletes, or updates many rows, foreign key constraints (like unique constraints) are checked row-by-row. When performing foreign key checks, [InnoDB](#) sets shared row-level locks on child or parent records that it must examine. MySQL checks foreign key constraints immediately; the check is not deferred to transaction commit. According to the SQL standard, the default behavior should be deferred checking. That is, constraints are only checked after the *entire SQL statement* has been processed. This means that it is not possible to delete a row that refers to itself using a foreign key.

For information how [InnoDB](#) foreign keys differ from the SQL standard, see [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

### 1.8.6.3. Constraints on Invalid Data

By default, MySQL is forgiving of invalid or improper data values and coerces them to valid values for data entry. However, you can change the server SQL mode to select more traditional treatment of bad values such that the server rejects them and aborts the statement in which they occur. See [Section 5.1.7, “Server SQL Modes”](#).

This section describes the default (forgiving) behavior of MySQL, as well as the strict SQL mode and how it differs.

If you are not using strict mode, then whenever you insert an “incorrect” value into a column, such as a [NULL](#) into a [NOT NULL](#) column or a too-large numeric value into a numeric column, MySQL sets the column to the “best possible value” instead of producing an error: The following rules describe in more detail how this works:

- If you try to store an out of range value into a numeric column, MySQL Server instead stores zero, the smallest possible value, or the largest possible value, whichever is closest to the invalid value.
- For strings, MySQL stores either the empty string or as much of the string as can be stored in the column.
- If you try to store a string that doesn't start with a number into a numeric column, MySQL Server stores 0.
- Invalid values for [ENUM](#) and [SET](#) columns are handled as described in [Section 1.8.6.4, “ENUM and SET Constraints”](#).
- MySQL enables you to store certain incorrect date values into [DATE](#) and [DATETIME](#) columns (such as '2000-02-31' or '2000-02-00'). The idea is that it is not the job of the SQL server to validate



dates. If MySQL can store a date value and retrieve exactly the same value, MySQL stores it as given. If the date is totally wrong (outside the server's ability to store it), the special “zero” date value '0000-00-00' is stored in the column instead.

- If you try to store `NULL` into a column that doesn't take `NULL` values, an error occurs for single-row `INSERT` statements. For multiple-row `INSERT` statements or for `INSERT INTO ... SELECT` statements, MySQL Server stores the implicit default value for the column data type. In general, this is 0 for numeric types, the empty string ( ' ' ) for string types, and the “zero” value for date and time types. Implicit default values are discussed in [Section 11.5, “Data Type Default Values”](#).
- If an `INSERT` statement specifies no value for a column, MySQL inserts its default value if the column definition includes an explicit `DEFAULT` clause. If the definition has no such `DEFAULT` clause, MySQL inserts the implicit default value for the column data type.

The reason for using the preceding rules in nonstrict mode is that we can't check these conditions until the statement has begun executing. We can't just roll back if we encounter a problem after updating a few rows, because the storage engine may not support rollback. The option of terminating the statement is not that good; in this case, the update would be “half done,” which is probably the worst possible scenario. In this case, it is better to “do the best you can” and then continue as if nothing happened.

In MySQL 5.0.2 and up, you can select stricter treatment of input values by using the `STRICT_TRANS_TABLES` [\[607\]](#) or `STRICT_ALL_TABLES` [\[607\]](#) SQL modes:

```
SET sql_mode = 'STRICT_TRANS_TABLES';
SET sql_mode = 'STRICT_ALL_TABLES';
```

`STRICT_TRANS_TABLES` [\[607\]](#) enables strict mode for transactional storage engines, and also to some extent for nontransactional engines. It works like this:

- For transactional storage engines, bad data values occurring anywhere in a statement cause the statement to abort and roll back.
- For nontransactional storage engines, a statement aborts if the error occurs in the first row to be inserted or updated. (When the error occurs in the first row, the statement can be aborted to leave the table unchanged, just as for a transactional table.) Errors in rows after the first do not abort the statement, because the table has already been changed by the first row. Instead, bad data values are adjusted and result in warnings rather than errors. In other words, with `STRICT_TRANS_TABLES` [\[607\]](#), a wrong value causes MySQL to roll back all updates done so far, if that can be done without changing the table. But once the table has been changed, further errors result in adjustments and warnings.

For even stricter checking, enable `STRICT_ALL_TABLES` [\[607\]](#). This is the same as `STRICT_TRANS_TABLES` [\[607\]](#) except that for nontransactional storage engines, errors abort the statement even for bad data in rows following the first row. This means that if an error occurs partway through a multiple-row insert or update for a nontransactional table, a partial update results. Earlier rows are inserted or updated, but those from the point of the error on are not. To avoid this for nontransactional tables, either use single-row statements or else use `STRICT_TRANS_TABLES` [\[607\]](#) if conversion warnings rather than errors are acceptable. To avoid problems in the first place, do not use MySQL to check column content. It is safest (and often faster) to let the application ensure that it passes only valid values to the database.

With either of the strict mode options, you can cause errors to be treated as warnings by using `INSERT IGNORE` or `UPDATE IGNORE` rather than `INSERT` or `UPDATE` without `IGNORE`.

#### 1.8.6.4. `ENUM` and `SET` Constraints

`ENUM` and `SET` columns provide an efficient way to define columns that can contain only a given set of values. See [Section 11.4.4, “The `ENUM` Type”](#), and [Section 11.4.5, “The `SET` Type”](#). However, before MySQL 5.0.2, `ENUM` and `SET` columns do not provide true constraints on entry of invalid data:

- [ENUM](#) columns always have a default value. If you specify no default value, then it is [NULL](#) for columns that can have [NULL](#), otherwise it is the first enumeration value in the column definition.
- If you insert an incorrect value into an [ENUM](#) column or if you force a value into an [ENUM](#) column with [IGNORE](#), it is set to the reserved enumeration value of 0, which is displayed as an empty string in string context.
- If you insert an incorrect value into a [SET](#) column, the incorrect value is ignored. For example, if the column can contain the values 'a', 'b', and 'c', an attempt to assign 'a,x,b,y' results in a value of 'a,b'.

As of MySQL 5.0.2, you can configure the server to use strict SQL mode. See [Section 5.1.7, “Server SQL Modes”](#). With strict mode enabled, the definition of a [ENUM](#) or [SET](#) column does act as a constraint on values entered into the column. An error occurs for values that do not satisfy these conditions:

- An [ENUM](#) value must be one of those listed in the column definition, or the internal numeric equivalent thereof. The value cannot be the error value (that is, 0 or the empty string). For a column defined as [ENUM\('a','b','c'\)](#), values such as '', 'd', or 'ax' are invalid and are rejected.
- A [SET](#) value must be the empty string or a value consisting only of the values listed in the column definition separated by commas. For a column defined as [SET\('a','b','c'\)](#), values such as 'd' or 'a,b,c,d' are invalid and are rejected.

Errors for invalid values can be suppressed in strict mode if you use [INSERT IGNORE](#) or [UPDATE IGNORE](#). In this case, a warning is generated rather than an error. For [ENUM](#), the value is inserted as the error member (0). For [SET](#), the value is inserted as given except that any invalid substrings are deleted. For example, 'a,x,b,y' results in a value of 'a,b'.

## 1.9. Credits

The following sections list developers, contributors, and supporters that have helped to make MySQL what it is today.

### 1.9.1. Contributors to MySQL

Although Oracle Corporation and/or its affiliates own all copyrights in the [MySQL server](#) and the [MySQL manual](#), we wish to recognize those who have made contributions of one kind or another to the [MySQL distribution](#). Contributors are listed here, in somewhat random order:

- Gianmassimo Vigazzola <[qwerg@mbbox.vol.it](mailto:qwerg@mbbox.vol.it)> or <[qwerg@tin.it](mailto:qwerg@tin.it)>

The initial port to Win32/NT.

- Per Eric Olsson

For constructive criticism and real testing of the dynamic record format.

- Irena Pancirov <[irena@mail.yacc.it](mailto:irena@mail.yacc.it)>

Win32 port with Borland compiler. [mysqlshutdown.exe](#) and [mysqlwatch.exe](#).

- David J. Hughes

For the effort to make a shareware SQL database. At TcX, the predecessor of MySQL AB, we started with [mSQL](#), but found that it couldn't satisfy our purposes so instead we wrote an SQL interface to our application builder Unireg. [mysqladmin](#) and [mysql](#) client are programs that were largely influenced by their [mSQL](#) counterparts. We have put a lot of effort into making the MySQL syntax a superset of [mSQL](#). Many of the API's ideas are borrowed from [mSQL](#) to make it easy to port free [mSQL](#) programs to the MySQL API. The MySQL software doesn't contain any code from [mSQL](#). Two files in the distribution ([client/insert\\_test.c](#) and [client/select\\_test.c](#)) are based

on the corresponding (noncopyrighted) files in the `mSQL` distribution, but are modified as examples showing the changes necessary to convert code from `mSQL` to MySQL Server. (`mSQL` is copyrighted David J. Hughes.)

- Patrick Lynch

For helping us acquire <http://www.mysql.com/>.

- Fred Lindberg

For setting up qmail to handle the MySQL mailing list and for the incredible help we got in managing the MySQL mailing lists.

- Igor Romanenko <[igor@frog.kiev.ua](mailto:igor@frog.kiev.ua)>

`mysqldump` (previously `msqldump`, but ported and enhanced by Monty).

- Yuri Dario

For keeping up and extending the MySQL OS/2 port.

- Tim Bunce

Author of `mysqlhotcopy`.

- Zarko Mocnik <[zarko.mocnik@dem.si](mailto:zarko.mocnik@dem.si)>

Sorting for Slovenian language.

- "TAMITO" <[tommy@valley.ne.jp](mailto:tommy@valley.ne.jp)>

The `_MB` character set macros and the `ujis` and `sjis` character sets.

- Joshua Chamas <[joshua@chamas.com](mailto:joshua@chamas.com)>

Base for concurrent insert, extended date syntax, debugging on NT, and answering on the MySQL mailing list.

- Yves Carlier <[Yves.Carlier@rug.ac.be](mailto:Yves.Carlier@rug.ac.be)>

`mysqlaccess`, a program to show the access rights for a user.

- Rhys Jones <[rhys@wales.com](mailto:rhys@wales.com)> (And GWE Technologies Limited)

For one of the early JDBC drivers.

- Dr Xiaokun Kelvin ZHU <[X.Zhu@brad.ac.uk](mailto:X.Zhu@brad.ac.uk)>

Further development of one of the early JDBC drivers and other MySQL-related Java tools.

- James Cooper <[pixel@organic.com](mailto:pixel@organic.com)>

For setting up a searchable mailing list archive at his site.

- Rick Mehalick <[Rick\\_Mehalick@i-o.com](mailto:Rick_Mehalick@i-o.com)>

For `xmysql`, a graphical X client for MySQL Server.

- Doug Sisk <[sisk@wix.com](mailto:sisk@wix.com)>

For providing RPM packages of MySQL for Red Hat Linux.

- Diemand Alexander V. <[axeld@vial.ethz.ch](mailto:axeld@vial.ethz.ch)>

For providing RPM packages of MySQL for Red Hat Linux-Alpha.

- Antoni Pamies Olive <[toni@readysoft.es](mailto:toni@readysoft.es)>

For providing RPM versions of a lot of MySQL clients for Intel and SPARC.

- Jay Bloodworth <[jay@pathways.sde.state.sc.us](mailto:jay@pathways.sde.state.sc.us)>

For providing RPM versions for MySQL 3.21.

- David Sacerdote <[davids@secnet.com](mailto:davids@secnet.com)>

Ideas for secure checking of DNS host names.

- Wei-Jou Chen <[jou@nematic.ieo.nctu.edu.tw](mailto:jou@nematic.ieo.nctu.edu.tw)>

Some support for Chinese(BIG5) characters.

- Wei He <[hewei@mail.ied.ac.cn](mailto:hewei@mail.ied.ac.cn)>

A lot of functionality for the Chinese(GBK) character set.

- Jan Pazdziora <[adelton@fi.muni.cz](mailto:adelton@fi.muni.cz)>

Czech sorting order.

- Zeev Suraski <[bourbon@netvision.net.il](mailto:bourbon@netvision.net.il)>

[FROM\\_UNIXTIME\(\)](#) [1218] time formatting, [ENCRYPT\(\)](#) [1262] functions, and [bison](#) advisor.  
Active mailing list member.

- Luuk de Boer <[luuk@wxs.nl](mailto:luuk@wxs.nl)>

Ported (and extended) the benchmark suite to [DBI/DBD](#). Have been of great help with [crash-me](#) and running benchmarks. Some new date functions. The [mysql\\_setpermission](#) script.

- Alexis Mikhailov <[root@medinf.chuvashia.su](mailto:root@medinf.chuvashia.su)>

User-defined functions (UDFs); [CREATE FUNCTION](#) and [DROP FUNCTION](#).

- Andreas F. Bobak <[bobak@relog.ch](mailto:bobak@relog.ch)>

The [AGGREGATE](#) extension to user-defined functions.

- Ross Wakelin <[R.Wakelin@march.co.uk](mailto:R.Wakelin@march.co.uk)>

Help to set up InstallShield for MySQL-Win32.

- Jethro Wright III <[jetman@li.net](mailto:jetman@li.net)>

The [libmysql.dll](#) library.

- James Pereria <[jpereira@iafrica.com](mailto:jpereira@iafrica.com)>

MySQLmanager, a Win32 GUI tool for administering MySQL Servers.

- Curt Sampson <[cjs@portal.ca](mailto:cjs@portal.ca)>

Porting of MIT-pthreads to NetBSD/Alpha and NetBSD 1.3/i386.

- Martin Ramsch <[m.ramsch@computer.org](mailto:m.ramsch@computer.org)>

Examples in the MySQL Tutorial.

- Steve Harvey  
For making `mysqlaccess` more secure.
- Konark IA-64 Centre of Persistent Systems Private Limited  
<http://www.pspl.co.in/konark/>. Help with the Win64 port of the MySQL server.
- Albert Chin-A-Young.  
Configure updates for Tru64, large file support and better TCP wrappers support.
- John Birrell  
Emulation of `pthread_mutex()` for OS/2.
- Benjamin Pflugmann  
Extended `MERGE` tables to handle `INSERTS`. Active member on the MySQL mailing lists.
- Jocelyn Fournier  
Excellent spotting and reporting innumerable bugs (especially in the MySQL 4.1 subquery code).
- Marc Liyanage  
Maintaining the Mac OS X packages and providing invaluable feedback on how to create Mac OS X packages.
- Robert Rutherford  
Providing invaluable information and feedback about the QNX port.
- Previous developers of NDB Cluster  
Lots of people were involved in various ways summer students, master thesis students, employees. In total more than 100 people so too many to mention here. Notable name is Ataulah Dabaghi who up until 1999 contributed around a third of the code base. A special thanks also to developers of the AXE system which provided much of the architectural foundations for NDB Cluster with blocks, signals and crash tracing functionality. Also credit should be given to those who believed in the ideas enough to allocate of their budgets for its development from 1992 to present time.
- Google Inc.  
We wish to recognize Google Inc. for contributions to the MySQL distribution: Mark Callaghan's SMP Performance patches and other patches.

Other contributors, bugfinders, and testers: James H. Thompson, Maurizio Menghini, Wojciech Tryc, Luca Berra, Zarko Mocnik, Wim Bonis, Elmar Haneke, <[jehamby@lightside](mailto:jehamby@lightside)>, <[psmith@BayNetworks.com](mailto:psmith@BayNetworks.com)>, <[duane@connect.com.au](mailto:duane@connect.com.au)>, Ted Deppner <[ted@psyber.com](mailto:ted@psyber.com)>, Mike Simons, Jaakko Hyvatti.

And lots of bug report/patches from the folks on the mailing list.

A big tribute goes to those that help us answer questions on the MySQL mailing lists:

- Daniel Koch <[dkoch@amcity.com](mailto:dkoch@amcity.com)>  
Irix setup.
- Luuk de Boer <[luuk@wxs.nl](mailto:luuk@wxs.nl)>  
Benchmark questions.

- Tim Sailer <[tps@users.buoy.com](mailto:tps@users.buoy.com)>  
DBD: :mysql questions.
- Boyd Lynn Gerber <[gerberb@zenez.com](mailto:gerberb@zenez.com)>  
SCO-related questions.
- Richard Mehalick <[RM186061@shellus.com](mailto:RM186061@shellus.com)>  
xmysql-related questions and basic installation questions.
- Zeev Suraski <[bourbon@netvision.net.il](mailto:bourbon@netvision.net.il)>  
Apache module configuration questions (log & auth), PHP-related questions, SQL syntax-related questions and other general questions.
- Francesc Guasch <[frankie@citel.upc.es](mailto:frankie@citel.upc.es)>  
General questions.
- Jonathan J Smith <[jsmith@wtp.net](mailto:jsmith@wtp.net)>  
Questions pertaining to OS-specifics with Linux, SQL syntax, and other things that might need some work.
- David Sklar <[sklar@student.net](mailto:sklar@student.net)>  
Using MySQL from PHP and Perl.
- Alistair MacDonald <[A.MacDonald@uel.ac.uk](mailto:A.MacDonald@uel.ac.uk)>  
Is flexible and can handle Linux and perhaps HP-UX.
- John Lyon <[jlyon@imag.net](mailto:jlyon@imag.net)>  
Questions about installing MySQL on Linux systems, using either `.rpm` files or compiling from source.
- Lorvid Ltd. <[lorvid@WOLFENET.com](mailto:lorvid@WOLFENET.com)>  
Simple billing/license/support/copyright issues.
- Patrick Sherrill <[patrick@coconet.com](mailto:patrick@coconet.com)>  
ODBC and VisualC++ interface questions.
- Randy Harmon <[rjharmon@uptimecomputers.com](mailto:rjharmon@uptimecomputers.com)>  
DBD, Linux, some SQL syntax questions.

## 1.9.2. Documenters and translators

The following people have helped us with writing the MySQL documentation and translating the documentation or error messages in MySQL.

- Paul DuBois  
Ongoing help with making this manual correct and understandable. That includes rewriting Monty's and David's attempts at English into English as other people know it.
- Kim Aldale  
Helped to rewrite Monty's and David's early attempts at English into English.



- Michael J. Miller Jr. <[mke@terrapin.turbolift.com](mailto:mke@terrapin.turbolift.com)>

For the first MySQL manual. And a lot of spelling/language fixes for the FAQ (that turned into the MySQL manual a long time ago).

- Yan Cailin

First translator of the MySQL Reference Manual into simplified Chinese in early 2000 on which the Big5 and HK coded (<http://mysql.hitstar.com/>) versions were based. [Personal home page at linuxdb.yeah.net](http://linuxdb.yeah.net).

- Jay Flaherty <[fty@mediapulse.com](mailto:fty@mediapulse.com)>

Big parts of the Perl [DBI/DBD](#) section in the manual.

- Paul Southworth <[pauls@etext.org](mailto:pauls@etext.org)>, Ray Loyzaga <[yar@cs.su.oz.au](mailto:yar@cs.su.oz.au)>

Proof-reading of the Reference Manual.

- Therrien Gilbert <[gilbert@ican.net](mailto:gilbert@ican.net)>, Jean-Marc Pouyot <[jmp@scaltaire.fr](mailto:jmp@scaltaire.fr)>

French error messages.

- Petr Snajdr, <[snajdr@pvt.net](mailto:snajdr@pvt.net)>

Czech error messages.

- Jaroslaw Lewandowski <[jotel@itnet.com.pl](mailto:jotel@itnet.com.pl)>

Polish error messages.

- Miguel Angel Fernandez Roiz

Spanish error messages.

- Roy-Magne Mo <[rmo@www.hivolda.no](mailto:rmo@www.hivolda.no)>

Norwegian error messages and testing of MySQL 3.21.xx.

- Timur I. Bakeyev <[root@timur.tatarstan.ru](mailto:root@timur.tatarstan.ru)>

Russian error messages.

- <[brenno@dewinter.com](mailto:brenno@dewinter.com)> & Filippo Grassilli <[phil@hyppo.com](mailto:phil@hyppo.com)>

Italian error messages.

- Dirk Munzinger <[dirk@trinity.saar.de](mailto:dirk@trinity.saar.de)>

German error messages.

- Billik Stefan <[billik@sun.uniag.sk](mailto:billik@sun.uniag.sk)>

Slovak error messages.

- Stefan Saroiu <[tzoompy@cs.washington.edu](mailto:tzoompy@cs.washington.edu)>

Romanian error messages.

- Peter Feher

Hungarian error messages.

- Roberto M. Serqueira

Portuguese error messages.

- Carsten H. Pedersen

Danish error messages.

- Arjen Lentz

Dutch error messages, completing earlier partial translation (also work on consistency and spelling).

### 1.9.3. Packages that support MySQL

The following is a list of creators/maintainers of some of the most important API/packages/applications that a lot of people use with MySQL.

We cannot list every possible package here because the list would then be way to hard to maintain. For other packages, please refer to the software portal at <http://solutions.mysql.com/software/>.

- Tim Bunce, Alligator Descartes

For the `DBD` (Perl) interface.

- Andreas Koenig <[a.koenig@mind.de](mailto:a.koenig@mind.de)>

For the Perl interface for MySQL Server.

- Jochen Wiedmann <[wiedmann@neckar-alb.de](mailto:wiedmann@neckar-alb.de)>

For maintaining the Perl `DBD::mysql` module.

- Eugene Chan <[eugene@acenet.com.sg](mailto:eugene@acenet.com.sg)>

For porting PHP for MySQL Server.

- Georg Richter

MySQL 4.1 testing and bug hunting. New PHP 5.0 `mysqli` extension (API) for use with MySQL 4.1 and up.

- Giovanni Maruzzelli <[maruzz@matrice.it](mailto:maruzz@matrice.it)>

For porting iODBC (Unix ODBC).

- Xavier Leroy <[Xavier.Leroy@inria.fr](mailto:Xavier.Leroy@inria.fr)>

The author of LinuxThreads (used by the MySQL Server on Linux).

### 1.9.4. Tools that were used to create MySQL

The following is a list of some of the tools we have used to create MySQL. We use this to express our thanks to those that has created them as without these we could not have made MySQL what it is today.

- Free Software Foundation

From whom we got an excellent compiler (`gcc`), an excellent debugger (`gdb`) and the `libc` library (from which we have borrowed `strto.c` to get some code working in Linux).

- Free Software Foundation & The XEmacs development team

For a really great editor/environment.

- Julian Seward

Author of [valgrind](#), an excellent memory checker tool that has helped us find a lot of otherwise hard to find bugs in MySQL.

- Dorothea Lütkehaus and Andreas Zeller

For [DDD](#) (The Data Display Debugger) which is an excellent graphical front end to [gdb](#).

## 1.9.5. Supporters of MySQL

Although Oracle Corporation and/or its affiliates own all copyrights in the [MySQL server](#) and the [MySQL manual](#), we wish to recognize the following companies, which helped us finance the development of the [MySQL server](#), such as by paying us for developing a new feature or giving us hardware for development of the [MySQL server](#).

- VA Linux / Andover.net

Funded replication.

- NuSphere

Editing of the MySQL manual.

- Stork Design studio

The MySQL Web site in use between 1998-2000.

- Intel

Contributed to development on Windows and Linux platforms.

- Compaq

Contributed to Development on Linux/Alpha.

- SWSOft

Development on the embedded [mysqld](#) version.

- FutureQuest

The `--skip-show-database` [\[441\]](#) option.



---

# Chapter 2. Installing and Upgrading MySQL

## Table of Contents

2.1. General Installation Guidance .....	49
2.1.1. Operating Systems Supported by MySQL Community Server .....	49
2.1.2. Choosing Which MySQL Distribution to Install .....	50
2.1.3. How to Get MySQL .....	53
2.1.4. Verifying Package Integrity Using MD5 Checksums or <a href="#">GnuPG</a> .....	54
2.1.5. Installation Layouts .....	62
2.1.6. Compiler-Specific Build Characteristics .....	62
2.2. Installing MySQL from Generic Binaries on Unix/Linux .....	62
2.3. Installing MySQL on Microsoft Windows .....	65
2.3.1. MySQL Installation Layout on Microsoft Windows .....	66
2.3.2. Choosing An Installation Package .....	67
2.3.3. Installing MySQL on Microsoft Windows Using MySQL Installer .....	67
2.3.4. Installing MySQL on Microsoft Windows Using a <a href="#">noinstall</a> Zip Archive .....	84
2.3.5. Troubleshooting a Microsoft Windows MySQL Server Installation .....	96
2.3.6. Upgrading MySQL on Windows .....	97
2.3.7. Windows Postinstallation Procedures .....	98
2.4. Installing MySQL on Mac OS X .....	100
2.4.1. General Notes on Installing MySQL on Mac OS X .....	100
2.4.2. Installing MySQL on Mac OS X Using Native Packages .....	101
2.4.3. Installing the MySQL Startup Item .....	104
2.4.4. Installing and Using the MySQL Preference Pane .....	107
2.4.5. Using the Bundled MySQL on Mac OS X Server .....	109
2.5. Installing MySQL on Linux .....	109
2.5.1. Installing MySQL from RPM Packages on Linux .....	110
2.5.2. Installing MySQL on Linux using Native Package Manager .....	114
2.6. Installing MySQL on Solaris and OpenSolaris .....	118
2.6.1. Installing MySQL on Solaris using a Solaris <a href="#">PKG</a> .....	118
2.6.2. Installing MySQL on OpenSolaris using IPS .....	119
2.7. Installing MySQL on HP-UX .....	120
2.7.1. General Notes on Installing MySQL on HP-UX .....	120
2.7.2. Installing MySQL on HP-UX using DEPOT .....	121
2.8. Installing MySQL on FreeBSD .....	121
2.9. Installing MySQL from Source .....	122
2.9.1. MySQL Layout for Source Installation .....	124
2.9.2. Installing MySQL from a Standard Source Distribution .....	124
2.9.3. Installing MySQL from a Development Source Tree .....	127
2.9.4. MySQL Source-Configuration Options .....	129
2.9.5. Dealing with Problems Compiling MySQL .....	137
2.9.6. MySQL Configuration and Third-Party Tools .....	138
2.10. Postinstallation Setup and Testing .....	139
2.10.1. Unix Postinstallation Procedures .....	139
2.10.2. Securing the Initial MySQL Accounts .....	150
2.11. Upgrading or Downgrading MySQL .....	153
2.11.1. Upgrading MySQL .....	154
2.11.2. Downgrading MySQL .....	159
2.11.3. Checking Whether Tables or Indexes Must Be Rebuilt .....	161
2.11.4. Rebuilding or Repairing Tables or Indexes .....	163
2.11.5. Copying MySQL Databases to Another Machine .....	164
2.12. Environment Variables .....	165
2.13. Perl Installation Notes .....	167
2.13.1. Installing Perl on Unix .....	167
2.13.2. Installing ActiveState Perl on Windows .....	168

This chapter describes how to obtain and install MySQL. A summary of the procedure follows and later sections provide the details. If you plan to upgrade an existing version of MySQL to a newer version rather than install MySQL for the first time, see [Section 2.11.1, "Upgrading MySQL"](#), for information about upgrade procedures and about issues that you should consider before upgrading.

If you are interested in migrating to MySQL from another database system, you may wish to read [Section B.8, "MySQL 5.6 FAQ: Migration"](#), which contains answers to some common questions concerning migration issues.

Installation of MySQL generally follows the steps outlined here:

**1. Determine whether MySQL runs and is supported on your platform.**

Please note that not all platforms are equally suitable for running MySQL, and that not all platforms on which MySQL is known to run are officially supported by Oracle Corporation:

**2. Choose which distribution to install.**

Several versions of MySQL are available, and most are available in several distribution formats. You can choose from pre-packaged distributions containing binary (precompiled) programs or source code. When in doubt, use a binary distribution. We also provide public access to our current source tree for those who want to see our most recent developments and help us test new code. To determine which version and type of distribution you should use, see [Section 2.1.2, "Choosing Which MySQL Distribution to Install"](#).

**3. Download the distribution that you want to install.**

For instructions, see [Section 2.1.3, "How to Get MySQL"](#). To verify the integrity of the distribution, use the instructions in [Section 2.1.4, "Verifying Package Integrity Using MD5 Checksums or GnuPG"](#).

**4. Install the distribution.**

To install MySQL from a binary distribution, use the instructions in [Section 2.2, "Installing MySQL from Generic Binaries on Unix/Linux"](#).

To install MySQL from a source distribution or from the current development source tree, use the instructions in [Section 2.9, "Installing MySQL from Source"](#).

**5. Perform any necessary postinstallation setup.**

After installing MySQL, read [Section 2.10, "Postinstallation Setup and Testing"](#). This section contains important information about making sure the MySQL server is working properly. It also describes how to secure the initial MySQL user accounts, *which have no passwords* until you assign passwords. The section applies whether you install MySQL using a binary or source distribution.

**6. If you want to run the MySQL benchmark scripts, Perl support for MySQL must be available. See [Section 2.13, "Perl Installation Notes"](#).**

Instructions for installing MySQL on different platforms and environments is available on a platform by platform basis:

- **Unix, Linux, FreeBSD**

For instructions on installing MySQL on most Linux and Unix platforms using a generic binary (for example, a `.tar.gz` package), see [Section 2.2, "Installing MySQL from Generic Binaries on Unix/Linux"](#).



For information on building MySQL entirely from the source code distributions or the source code repositories, see [Section 2.9, “Installing MySQL from Source”](#)

For specific platform help on installation, configuration, and building from source see the corresponding platform section:

- Linux, including notes on distribution specific methods, see [Section 2.5, “Installing MySQL on Linux”](#).
- Solaris and OpenSolaris, including PKG and IPS formats, see [Section 2.6, “Installing MySQL on Solaris and OpenSolaris”](#).
- IBM AIX, see [Section 2.6, “Installing MySQL on Solaris and OpenSolaris”](#).
- Hewlett-Packard HP-UX, including the DEPOT package format, see [Section 2.7, “Installing MySQL on HP-UX”](#).
- FreeBSD, see [Section 2.8, “Installing MySQL on FreeBSD”](#).

- **Microsoft Windows**

For instructions on installing MySQL on Microsoft Windows, using either a Zipped binary or an MSI package, see [Section 2.3, “Installing MySQL on Microsoft Windows”](#).

For details and instructions on building MySQL from source code using Microsoft Visual Studio, see [Section 2.9, “Installing MySQL from Source”](#).

- **Mac OS X**

For installation on Mac OS X, including using both the binary package and native PKG formats, see [Section 2.4, “Installing MySQL on Mac OS X”](#).

For information on making use of the MySQL Startup Item to automatically start and stop MySQL, see [Section 2.4.3, “Installing the MySQL Startup Item”](#).

For information on the MySQL Preference Pane, see [Section 2.4.4, “Installing and Using the MySQL Preference Pane”](#).

## 2.1. General Installation Guidance

The immediately following sections contain the information necessary to choose, download, and verify your distribution. The instructions in later sections of the chapter describe how to install the distribution that you choose. For binary distributions, see the instructions at [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#) or the corresponding section for your platform if available. To build MySQL from source, use the instructions in [Section 2.9, “Installing MySQL from Source”](#).

### 2.1.1. Operating Systems Supported by MySQL Community Server

This section lists the operating systems on which MySQL Community Server is known to run.



#### Important

Oracle Corporation does not necessarily provide official support for all the platforms listed in this section. For information about those platforms that are officially supported, see <http://www.mysql.com/support/supportedplatforms.html> on the MySQL Web site.

We use `CMake`, so it is possible to port MySQL to all modern systems that have a C++ compiler and a working implementation of POSIX threads. (Thread support is needed for the server. To compile only the client code, the only requirement is a C++ compiler.)

MySQL has been reported to compile successfully on the following combinations of operating system and thread package.

- FreeBSD 5.x and up with native threads. See [Section 2.8, “Installing MySQL on FreeBSD”](#).
- HP-UX 11.x with the native threads. See [Section 2.7, “Installing MySQL on HP-UX”](#).
- Linux. Builds on all recent Linux distributions based on the 2.6 kernel. See [Section 2.5, “Installing MySQL on Linux”](#).
- Mac OS X. See [Section 2.4, “Installing MySQL on Mac OS X”](#).
- Solaris 2.8 on SPARC and x86, including support for native threads. See [Section 2.6, “Installing MySQL on Solaris and OpenSolaris”](#).
- Windows XP, Windows Vista, Windows Server 2003, and Windows Server 2008. See [Section 2.3, “Installing MySQL on Microsoft Windows”](#).

MySQL has also been known to run on other systems in the past. See [Section 2.1, “General Installation Guidance”](#). Some porting effort might be required for current versions of MySQL on these systems.

Not all platforms are equally well-suited for running MySQL. How well a certain platform is suited for a high-load mission-critical MySQL server is determined by the following factors:

- General stability of the thread library. A platform may have an excellent reputation otherwise, but MySQL is only as stable as the thread library it calls, even if everything else is perfect.
- The capability of the kernel and the thread library to take advantage of symmetric multi-processor (SMP) systems. In other words, when a process creates a thread, it should be possible for that thread to run on a CPU different from the original process.
- The capability of the kernel and the thread library to run many threads that acquire and release a mutex over a short critical region frequently without excessive context switches. If the implementation of `pthread_mutex_lock()` is too anxious to yield CPU time, this hurts MySQL tremendously. If this issue is not taken care of, adding extra CPUs actually makes MySQL slower.
- General file system stability and performance.
- Table size. If your tables are large, performance is affected by the ability of the file system to deal with large files and dealing with them efficiently.
- Our level of expertise here at Oracle Corporation with the platform. If we know a platform well, we enable platform-specific optimizations and fixes at compile time. We can also provide advice on configuring your system optimally for MySQL.
- The amount of testing we have done internally for similar configurations.
- The number of users that have run MySQL successfully on the platform in similar configurations. If this number is high, the likelihood of encountering platform-specific surprises is much smaller.

### 2.1.2. Choosing Which MySQL Distribution to Install

When preparing to install MySQL, you should decide which version to use. MySQL development occurs in several release series, and you can pick the one that best fits your needs. After deciding which version to install, you can choose a distribution format. Releases are available in binary or source format.

#### 2.1.2.1. Choosing Which Version of MySQL to Install

The first decision to make is whether you want to use a production (stable) release or a development release. In the MySQL development process, multiple release series co-exist, each at a different stage of maturity.

## Production Releases

- MySQL 5.6: Latest General Availability (Production) release
- MySQL 5.5: Previous General Availability (Production) release
- MySQL 5.1: Older General Availability (Production) release
- MySQL 5.0: Older Production release nearing the end of the product lifecycle

MySQL 4.1, 4.0, and 3.23 are old releases that are no longer supported.

See <http://www.mysql.com/about/legal/lifecycle/> for information about support policies and schedules.

Normally, if you are beginning to use MySQL for the first time or trying to port it to some system for which there is no binary distribution, use the most recent General Availability series listed in the preceding descriptions. All MySQL releases, even those from development series, are checked with the MySQL benchmarks and an extensive test suite before being issued.

If you are running an older system and want to upgrade, but do not want to take the chance of having a nonseamless upgrade, you should upgrade to the latest version in the same release series you are using (where only the last part of the version number is newer than yours). We have tried to fix only fatal bugs and make only small, relatively “safe” changes to that version.

If you want to use new features not present in the production release series, you can use a version from a development series. Be aware that development releases are not as stable as production releases.

We do not use a complete code freeze because this prevents us from making bugfixes and other fixes that must be done. We may add small things that should not affect anything that currently works in a production release. Naturally, relevant bugfixes from an earlier series propagate to later series.

If you want to use the very latest sources containing all current patches and bugfixes, you can use one of our source code repositories (see [Section 2.9.3, “Installing MySQL from a Development Source Tree”](#)). These are not “releases” as such, but are available as previews of the code on which future releases are to be based.

The naming scheme in MySQL 5.6 uses release names that consist of three numbers and a suffix; for example, **mysql-5.6.1-m1**. The numbers within the release name are interpreted as follows:

- The first number (**5**) is the major version and describes the file format. All MySQL 5 releases have the same file format.
- The second number (**6**) is the release level. Taken together, the major version and release level constitute the release series number.
- The third number (**1**) is the version number within the release series. This is incremented for each new release. Usually you want the latest version for the series you have chosen.

For each minor update, the last number in the version string is incremented. When there are major new features or minor incompatibilities with previous versions, the second number in the version string is incremented. When the file format changes, the first number is increased.

Release names also include a suffix to indicate the stability level of the release. Releases within a series progress through a set of suffixes to indicate how the stability level improves. The possible suffixes are:

- **mN** (for example, **m1**, **m2**, **m3**, ...) indicate a milestone number. MySQL development uses a milestone model, in which each milestone proceeds through a small number of versions with a tight focus on a small subset of thoroughly tested features. Following the releases for one milestone, development proceeds with another small number of releases that focuses on the next small set of

features, also thoroughly tested. Features within milestone releases may be considered to be of pre-production quality.

- **rc** indicates a Release Candidate. Release candidates are believed to be stable, having passed all of MySQL's internal testing, and with all known fatal runtime bugs fixed. However, the release has not been in widespread use long enough to know for sure that all bugs have been identified. Only minor fixes are added.
- If there is no suffix, it indicates that the release is a General Availability (GA) or Production release. GA releases are stable, having successfully passed through all earlier release stages and are believed to be reliable, free of serious bugs, and suitable for use in production systems. Only critical bugfixes are applied to the release.

All releases of MySQL are run through our standard tests and benchmarks to ensure that they are relatively safe to use. Because the standard tests are extended over time to check for all previously found bugs, the test suite keeps getting better.

All releases have been tested at least with these tools:

- **An internal test suite.** The `mysql-test` directory contains an extensive set of test cases. We run these tests for every server binary. See [Section 22.1.2, "The MySQL Test Suite"](#), for more information about this test suite.
- **The MySQL benchmark suite.** This suite runs a range of common queries. It is also a test to determine whether the latest batch of optimizations actually made the code faster. See [Section 8.12.2, "The MySQL Benchmark Suite"](#).

We also perform additional integration and nonfunctional testing of the latest MySQL version in our internal production environment. Integration testing is done with different connectors, storage engines, replication modes, backup, partitioning, stored programs, and so forth in various combinations. Additional nonfunctional testing is done in areas of performance, concurrency, stress, high volume, upgrade and downgrade.

### 2.1.2.2. Choosing a Distribution Format

After choosing which version of MySQL to install, you should decide whether to use a binary distribution or a source distribution. In most cases, you should probably use a binary distribution, if one exists for your platform. Binary distributions are available in native format for many platforms, such as RPM files for Linux or PKG package installers for Mac OS X or Solaris. Distributions also are available as Zip archives or compressed `tar` files.

Reasons to choose a binary distribution include the following:

- Binary distributions generally are easier to install than source distributions.
- To satisfy different user requirements, we provide several servers in binary distributions. `mysqld` is an optimized server that is a smaller, faster binary. `mysqld-debug` is compiled with debugging support.

Each of these servers is compiled from the same source distribution, though with different configuration options. All native MySQL clients can connect to servers from either MySQL version.

Under some circumstances, you may be better off installing MySQL from a source distribution:

- You want to install MySQL at some explicit location. The standard binary distributions are ready to run at any installation location, but you might require even more flexibility to place MySQL components where you want.
- You want to configure `mysqld` to ensure that features are available that might not be included in the standard binary distributions. Here is a list of the most common extra options that you may want to use to ensure feature availability:

- `-DWITH_LIBWRAP=1` [136] for TCP wrappers support.
- `-DWITH_ZLIB={system|bundled}` [137] for features that depend on compression
- `-DWITH_DEBUG=1` [136] for debugging support
- You want to configure `mysqld` without some features that are included in the standard binary distributions. For example, distributions normally are compiled with support for all character sets. If you want a smaller MySQL server, you can recompile it with support for only the character sets you need.
- You want to use the latest sources from one of the Bazaar repositories to have access to all current bugfixes. For example, if you have found a bug and reported it to the MySQL development team, the bugfix is committed to the source repository and you can access it there. The bugfix does not appear in a release until a release actually is issued.
- You want to read (or modify) the C and C++ code that makes up MySQL. For this purpose, you should get a source distribution, because the source code is always the ultimate manual.
- Source distributions contain more tests and examples than binary distributions.

### 2.1.2.3. How and When Updates Are Released

MySQL is evolving quite rapidly and we want to share new developments with other MySQL users. We try to produce a new release whenever we have new and useful features that others also seem to have a need for.

We also try to help users who request features that are easy to implement. We take note of what our licensed users want, and we especially take note of what our support customers want and try to help them in this regard.

No one is *required* to download a new release. The [Release Notes](#) help you determine whether the new release has something you really want.

We use the following policy when updating MySQL:

- Enterprise Server releases are meant to appear every 18 months, supplemented by quarterly service packs and monthly rapid updates. Community Server releases are meant to appear 2 to 3 times per year.
- Releases are issued within each series. For each release, the last number in the version is one more than the previous release within the same series.
- Binary distributions for some platforms are made by us for major releases. Other people may make binary distributions for other systems, but probably less frequently.
- We make fixes available as soon as we have identified and corrected small or noncritical but annoying bugs. The fixes are available in source form immediately from our public Bazaar repositories, and are included in the next release.
- If by any chance a security vulnerability or critical bug is found in a release, our policy is to fix it in a new release as soon as possible. (We would like other companies to do this, too!)

### 2.1.3. How to Get MySQL

Check our downloads page at <http://dev.mysql.com/downloads/> for information about the current version of MySQL and for downloading instructions. For a complete up-to-date list of MySQL download mirror sites, see <http://dev.mysql.com/downloads/mirrors.html>. You can also find information there about becoming a MySQL mirror site and how to report a bad or out-of-date mirror.

To obtain the latest development source, see [Section 2.9.3, “Installing MySQL from a Development Source Tree”](#).

## 2.1.4. Verifying Package Integrity Using MD5 Checksums or GnuPG

After you have downloaded the MySQL package that suits your needs and before you attempt to install it, you should make sure that it is intact and has not been tampered with. There are three means of integrity checking:

- MD5 checksums
- Cryptographic signatures using [GnuPG](#), the GNU Privacy Guard
- For RPM packages, the built-in RPM integrity verification mechanism

The following sections describe how to use these methods.

If you notice that the MD5 checksum or GPG signatures do not match, first try to download the respective package one more time, perhaps from another mirror site.

### 2.1.4.1. Verifying the MD5 Checksum

After you have downloaded a MySQL package, you should make sure that its MD5 checksum matches the one provided on the MySQL download pages. Each package has an individual checksum that you can verify against the package that you downloaded. The correct MD5 checksum is listed on the downloads page for each MySQL product, and you will compare it against the MD5 checksum of the file (product) that you download.

Each operating system and setup offers its own version of tools for checking the MD5 checksum. Typically the command is named `md5sum`, or it may be named `md5`, and some operating systems do not ship it at all. On Linux, it is part of the **GNU Text Utilities** package, which is available for a wide range of platforms. You can also download the source code from <http://www.gnu.org/software/textutils/>. If you have OpenSSL installed, you can use the command `openssl md5 package_name` instead. A Windows implementation of the `md5` command line utility is available from <http://www.fourmilab.ch/md5/>. `winMd5Sum` is a graphical MD5 checking tool that can be obtained from <http://www.nullriver.com/index/products/winmd5sum>. Our Microsoft Windows examples will assume the name `md5.exe`.

Linux and Microsoft Windows examples:

```
shell> md5sum mysql-standard-5.6.12-linux-i686.tar.gz
aaab65abbec64d5e907dcd41b8699945  mysql-standard-5.6.12-linux-i686.tar.gz
```

```
shell> md5.exe mysql-installer-community-5.6.12.msi
aaab65abbec64d5e907dcd41b8699945  mysql-installer-community-5.6.12.msi
```

You should verify that the resulting checksum (the string of hexadecimal digits) matches the one displayed on the download page immediately below the respective package.



#### Note

Make sure to verify the checksum of the *archive file* (for example, the `.zip`, `.tar.gz`, or `.msi` file) and not of the files that are contained inside of the archive. In other words, verify the file before extracting its contents.

### 2.1.4.2. Signature Checking Using GnuPG

Another method of verifying the integrity and authenticity of a package is to use cryptographic signatures. This is more reliable than using [MD5 checksums](#), but requires more work.

We sign MySQL downloadable packages with [GnuPG](#) (GNU Privacy Guard). [GnuPG](#) is an Open Source alternative to the well-known Pretty Good Privacy ([PGP](#)) by Phil Zimmermann. See <http://www.gnupg.org/> for more information about [GnuPG](#) and how to obtain and install it on your system.



Most Linux distributions ship with [GnuPG](#) installed by default. For more information about [GnuPG](#), see <http://www.openpgp.org/>.

To verify the signature for a specific package, you first need to obtain a copy of our public GPG build key, which you can download from <http://pgp.mit.edu/>. The key that you want to obtain is named [mysql-build@oss.oracle.com](mailto:mysql-build@oss.oracle.com). Alternatively, you can cut and paste the key directly from the following text:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: PGP Universal 2.9.1 (Build 347)

mQGIBD4+owwRBAC14GIUFuFCyEDSIEPvEW3SAFUdJBtoQHH/nJKZyQT7h9bP1UWC3
RODjQReyCITRrdwyrKUGku2FmeVGwn2u2WmDMNABLnpPrWPkBDck96+OmSLN9brz
fw2v0UgCmYv2hW0hyDHuvYlQA/BThQoADgJ8AW6/0Lo7V1W9/8VuHP0gQwCgvzV3
BqOxRznNCRcRxAuAuVzthRcEAJooQK1+iSiunZMYD1WufeXfshc57S/+yeJkegNW
hxwR9pRWWArNYJdDRT+rf2RUe3vpquKNQU/hnEIHJRQqYHo8gTxvxXNQC7fJYLV
K2HtkrPbP72vwsEKMYhhr0eKCbTlGfLs9krjJ6sBgACyP/Vb7hiPwxh6rDZ7ITnE
kYpXBACmWpP8NJTkamEnPCia2ZoOHODANwpUkP43I7jsDmgtobZX9qnrAXw+uNDI
QJEXM6FSbi0LLtZciNLYsafwAPEOMDKpMqAK6IyisNtPvaLd8lH0bPAnWqcyefep
rv0sxxquEMcM3o7wggfN83POkDasDbs3pjwPhxvzh6//62zQJ7Q2TXlTUUwgUmVs
ZWFzZSBFbmdpbmVlcm1uZyA8bXlzcWwtYnVpbGRAb3NzLm9yYWNsZS5jb20+iGYE
ExECACyCGyMGcwkIBwMCBBUCCAMEFgIDAQIeAQIXgAUCtnc+KgUJE/sCFQAKCRCM
cY07UHLh9SbMAJ411+qBz2BZNSGCzwwA6YbhGPC7FwCgpp8z5TzIw4YQuL5NGJ/sy
0oSazqmJASIEEAECAAwFAk53QS4FAwASdQAACgkQlxc4m8pXrXwJ8Qf/be/UO9mq
foc2sMyhwMpn4/fdBWwflkA12FXQDOQMvW9HsmEjnfUgYKXschZRi+DuHXe1P7l
8G2aQLubhBsQf9eJkVrFTzuWMQkdIq+6Koulxv6ofkCcv3d1xtO2W7nb5yxcPVBp
rRfGFGebJvZa58DymCNgyGtAU6AOz4veavNmI2+GIDQsY66+tYDvZ+CxwzdYu+HD
V9HmrJfc6deM0mnBn7SRjqzxJPgoTQihTav6q/R5/2p5NvQ/H84OgS6GjosfGc2
duUDzCP/kheMRkfzuyKCOHQPtJuIj8++gfpHtEU7IDUX1So3c9n0PdpeBvclSDBp
RnCNxQWU4mBot7kCDQQ+PqMdEAgA7+GJfxbmDY4wslPnjH9rF4N2qfWsEN/lxaZo
JYc3a6M02WCnH16aht2/tBK2w1QI4YFteR47gCvtgb6O1JHffOo2HfLmRDRiRjd1
DTCHqeyX7CHhcgj/dNR1W2Z015QFEcmV9U0Vhp3aFfWC4Ujfs3LU+hkAWzE7zaD
5cH9J7yv/6xuZVw411x0h4UqsTcWMu0im1BzELqX1DY7LwoPEb/O9Rkbf4fmLe11
EzIaCa4PqARXQZc4dhSinMt6K3X4BrRsKTFozBu74F47D8I1bf5vSYHbuE5p/1oI
Dznkg/p8kW+3FwuWrycciqFTcNz215yyX39LXFnlLzKUB/F5GwADBQf+Lwqqa8CG
rRfs0AJxim63CHft5mUc5rUSnTslGYEIOCR1BeQauyPZbPDsDD9MZ1ZaSafanFv
wFG6Llx9xkuU7tzq+vKLoWkm4u5xf3vn55VjnSdlaQ9eQnUcXiL4cnBGoTbOWI39E
cyzgs1zBdC++mpjCQTcA7p6JUvS6oAB3FQWg54tuUo0Ec8bsM8b3Ev42LmuQT5N
dKHGwHsXTptl0klk4bQk4OajHsiy1BMahpT27jWjJlMiJc+IwJ0mghkKHt926s/y
mfd5HkdQ1cyvz5tryVI3F78XeSYfQvuuwqp2H139pXGekg0n6KdUOetdZWhe7
0YGNPwlyjWJT1IhUBBGRAGAMBQJodz3tBQkT+wG4ABIHZUdQRwABAQkQjHGN01By
4fUUmwCbByr2+bBEN/L2B0cnw9Z/QFWuhRMAoKVgCFm5fadQ3Afi+UQ1AcOphrnJ
=Eto8
-----END PGP PUBLIC KEY BLOCK-----
```

To import the build key into your personal public GPG keyring, use `gpg --import`. For example, if you have saved the key in a file named `mysql_pubkey.asc`, the import command looks like this:

```
shell> gpg --import mysql_pubkey.asc
gpg: key 5072E1F5: public key "MySQL Release Engineering
<mysql-build@oss.oracle.com>" imported
gpg: Total number processed: 1
gpg: imported: 1
gpg: no ultimately trusted keys found
```

You can also download the key from the public keyserver using the public key id, `5072E1F5`:

```
shell> gpg --recv-keys 5072E1F5
gpg: requesting key 5072E1F5 from hkp server keys.gnupg.net
gpg: key 5072E1F5: "MySQL Release Engineering <mysql-build@oss.oracle.com>"
1 new user ID
gpg: key 5072E1F5: "MySQL Release Engineering <mysql-build@oss.oracle.com>"
53 new signatures
gpg: no ultimately trusted keys found
gpg: Total number processed: 1
gpg: new user IDs: 1
gpg: new signatures: 53
```

If you want to import the key into your RPM configuration to validate RPM install packages, you should be able to import the key directly:

```
shell> rpm --import mysql_pubkey.asc
```

If you experience problems or require RPM specific information, see [Section 2.1.4.4, "Signature Checking Using RPM"](#).

After you have downloaded and imported the public build key, download your desired MySQL package and the corresponding signature, which also is available from the download page. The signature file has the same name as the distribution file with an `.asc` extension, as shown by the examples in the following table.

**Table 2.1. MySQL Package and Signature Files for Source files**

File Type	File Name
Distribution file	<code>mysql-standard-5.6.12-linux-i686.tar.gz</code>
Signature file	<code>mysql-standard-5.6.12-linux-i686.tar.gz.asc</code>

Make sure that both files are stored in the same directory and then run the following command to verify the signature for the distribution file:

```
shell> gpg --verify package_name.asc
```

If the downloaded package is valid, you will see a "Good signature" similar to:

```
shell> gpg --verify mysql-standard-5.6.12-linux-i686.tar.gz.asc
gpg: Signature made Tue 01 Feb 2011 02:38:30 AM CST using DSA key ID 5072E1F5
gpg: Good signature from "MySQL Release Engineering <mysql-build@oss.oracle.com>"
```

The `Good signature` message indicates that the file signature is valid, when compared to the signature listed on our site. But you might also see warnings, like so:

```
shell> gpg --verify mysql-standard-5.6.12-linux-i686.tar.gz.asc
gpg: Signature made Wed 23 Jan 2013 02:25:45 AM PST using DSA key ID 5072E1F5
gpg: checking the trustdb
gpg: no ultimately trusted keys found
gpg: Good signature from "MySQL Release Engineering <mysql-build@oss.oracle.com>"
gpg: WARNING: This key is not certified with a trusted signature!
gpg:          There is no indication that the signature belongs to the owner.
Primary key fingerprint: A4A9 4068 76FC BD3C 4567 70C8 8C71 8D3B 5072 E1F5
```

That is normal, as they depend on your setup and configuration. Here are explanations for these warnings:

- *gpg: no ultimately trusted keys found*: This means that the specific key is not "ultimately trusted" by you or your web of trust, which is okay for the purposes of verifying file signatures.
- *WARNING: This key is not certified with a trusted signature! There is no indication that the signature belongs to the owner.*: This refers to your level of trust in your belief that you possess our real public key. This is a personal decision. Ideally, a MySQL developer would hand you the key in person, but more commonly, you downloaded it. Was the download tampered with? Probably not, but this decision is up to you. Setting up a web of trust is one method for trusting them.


See the GPG documentation for more information on how to work with public keys.

### 2.1.4.3. Signature Checking Using [Gpg4win](#) for Windows

The [Section 2.1.4.2, "Signature Checking Using GnuPG"](#) section describes how to verify MySQL downloads using GPG. That guide also applies to Microsoft Windows, but another option is to use a GUI tool like [Gpg4win](#). You may use a different tool but our examples are based on Gpg4win, and utilize its bundled [Kleopatra](#) GUI.


Download and install Gpg4win, and then load Kleopatra. The dialog should look similar to:

Figure 2.1. Initial screen after loading Kleopatra



Next, add the MySQL Release Engineering certificate. Do this by clicking **File**, **Lookup Certificates on Server**. Type "Mysql Release Engineering" into the search box and press **Search**.

Figure 2.2. Finding the MySQL Release Engineering certificate

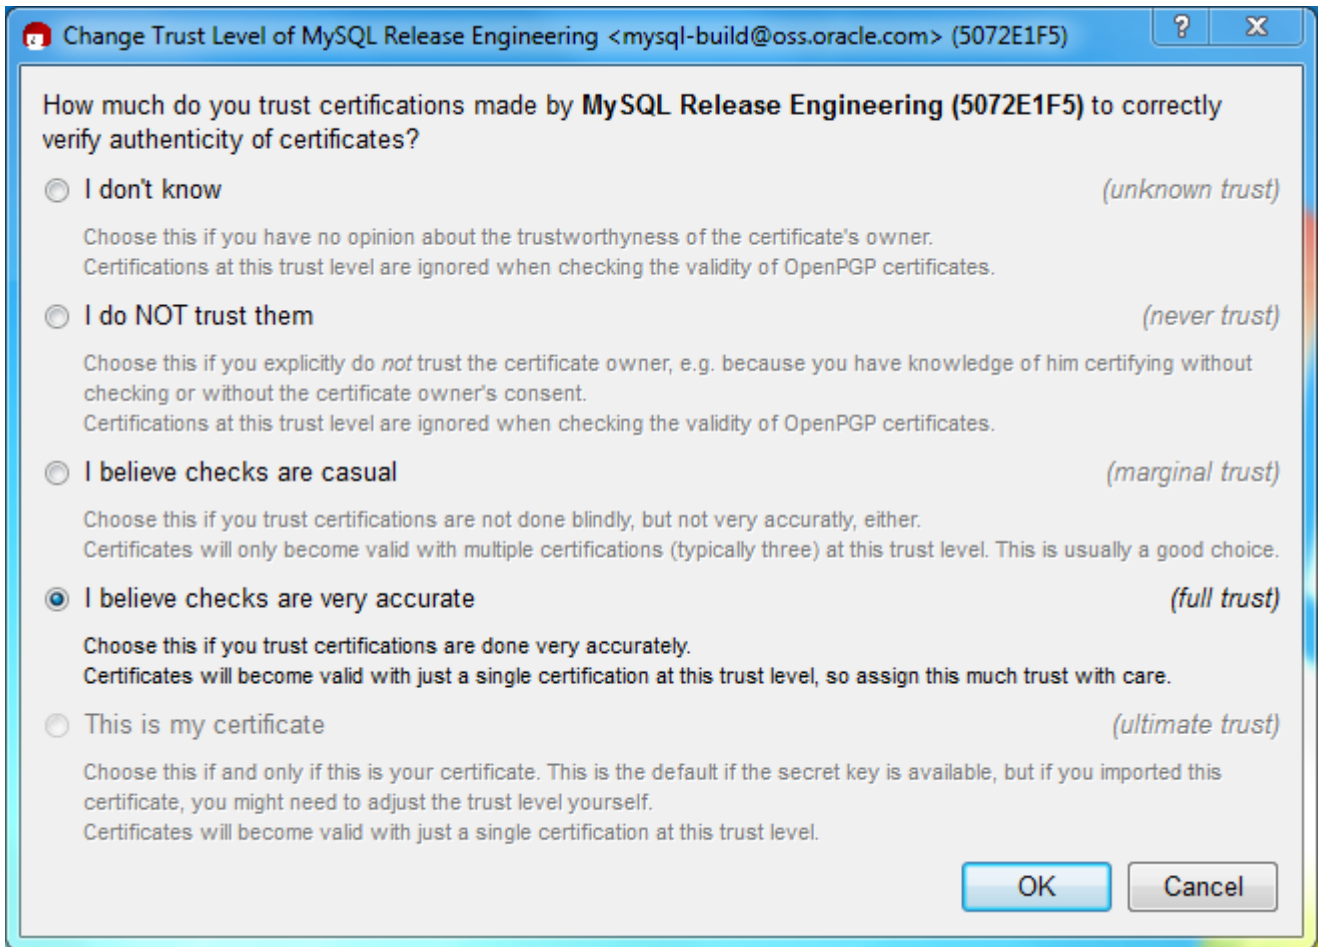


Select the "MySQL Release Engineering" certificate. The Fingerprint and Key-ID must be "5072E1F5", or choose **Details...** to confirm the certificate is valid. Now, import it by clicking **Import**. An import dialog will be displayed, choose **Okay**, and this certificate will now be listed under the **Imported Certificates** tab.

Next, configure the trust level for our certificate. Select our certificate, then from the main menu select **Certificates**, **Change Owner Trust...**. We suggest choosing **I believe checks are very accurate** for our

certificate, as otherwise you might not be able to verify our signature. Select **I believe checks are very accurate** and then press **OK**.

**Figure 2.3. Changing the Trust level**

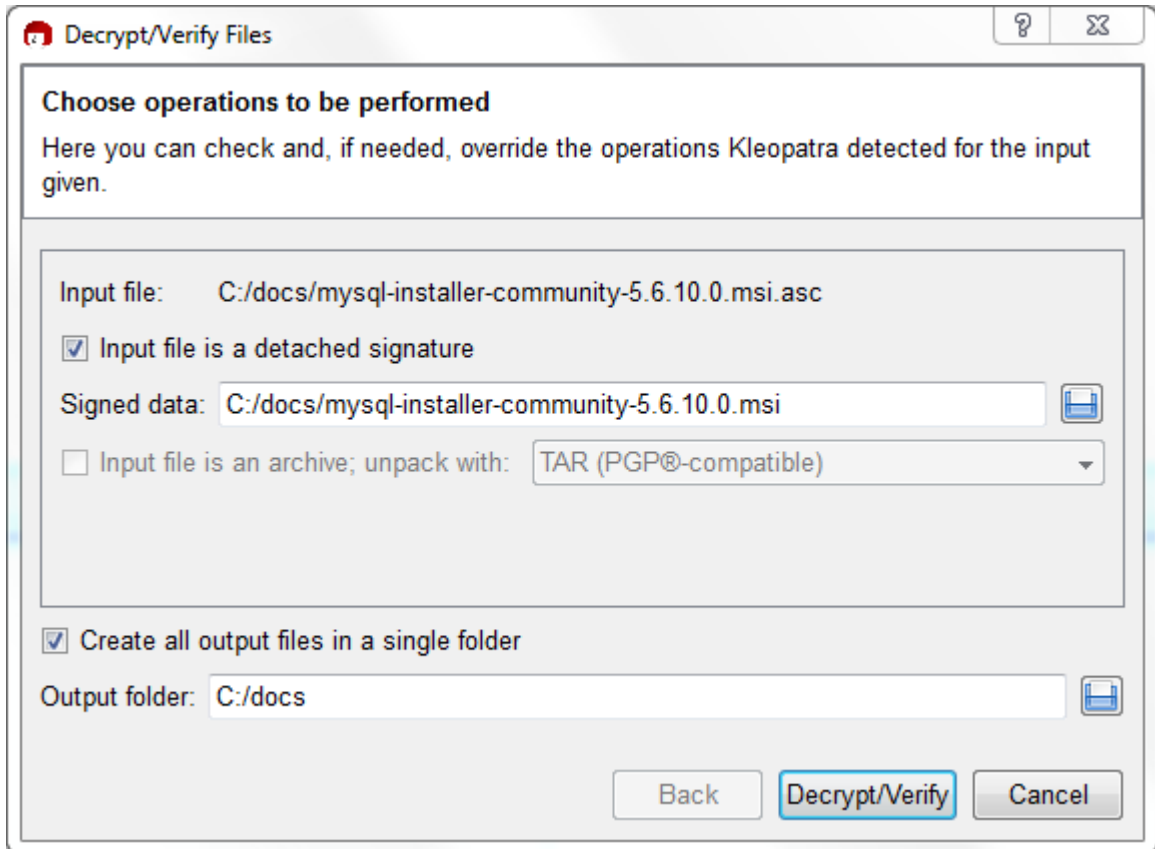


Next, verify the downloaded MySQL package file. This requires files for both the packaged file, and the signature. The signature file must have the same name as the packaged file but with an appended `.asc` extension, as shown by the example in the following table. The signature is linked to on the downloads page for each MySQL product. You must create the `.asc` file with this signature.

**Table 2.2. MySQL Package and Signature Files for MySQL Installer for Microsoft Windows**


File Type	File Name
Distribution file	<code>mysql-installer-community-5.6.12.msi</code>
Signature file	<code>mysql-installer-community-5.6.12.msi.asc</code>

Make sure that both files are stored in the same directory and then run the following command to verify the signature for the distribution file. Either drag and drop the signature (`.asc`) file into Kleopatra, or load the dialog from **File**, **Decrypt/Verify Files...**, and then choose either the `.msi` or `.asc` file.

**Figure 2.4. The Decrypt/Verify Files dialog**

Click **Decrypt/Verify** to check the file. The two most common results will look like the following, and although the yellow warning looks problematic, the following means that the file check passed with success. You may now run this installer.

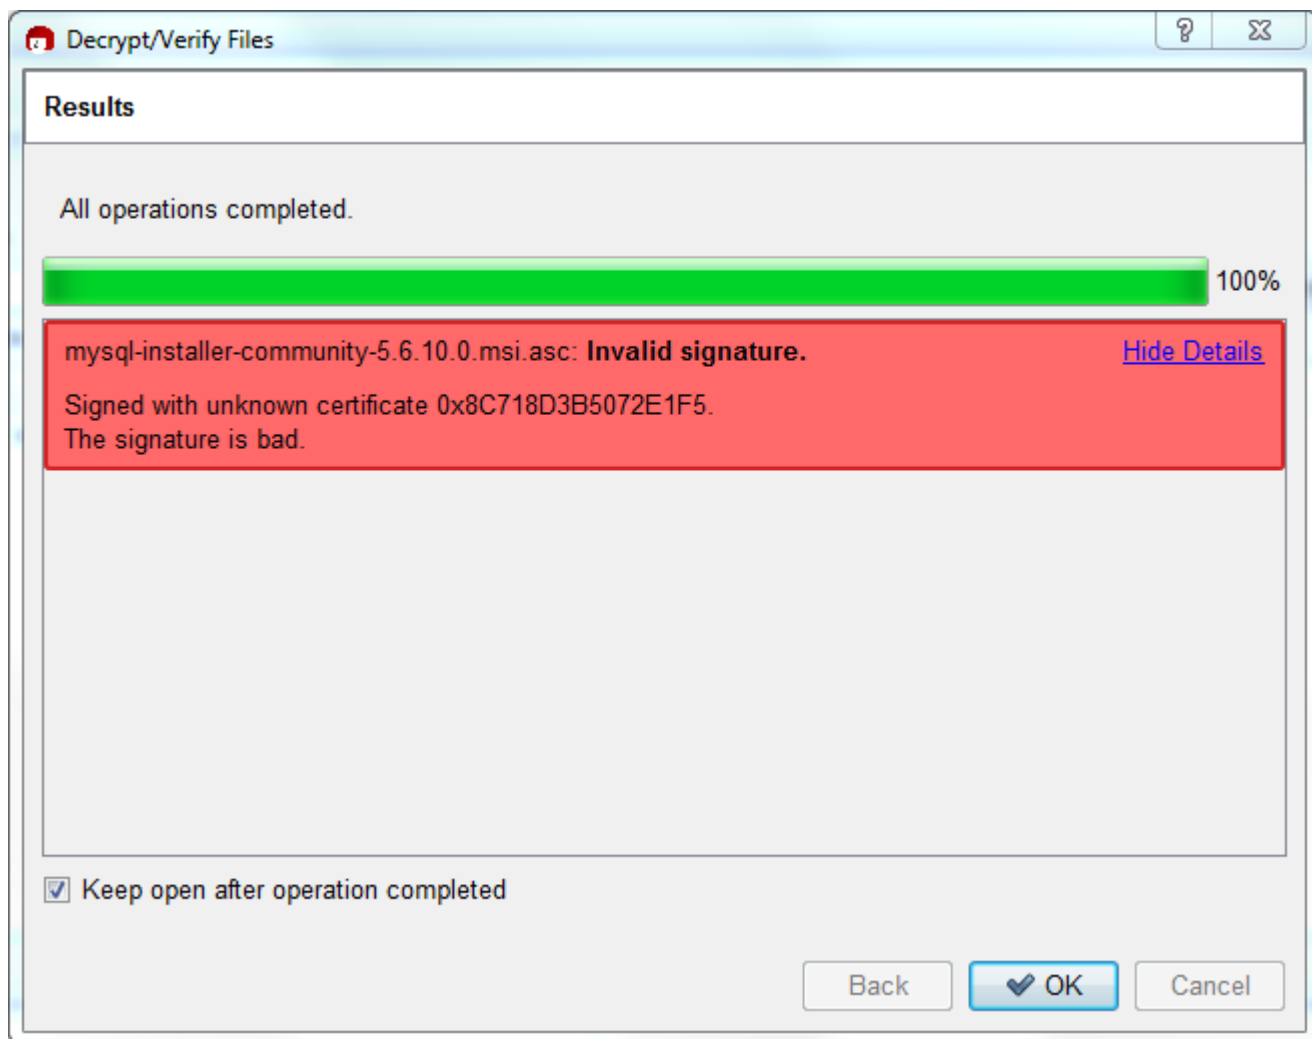
Figure 2.5. The Decrypt/Verify Results: Good



Seeing a red "The signature is bad" error means the file is invalid. Do not execute the MSI file if you see this error.



Figure 2.6. The Decrypt/Verify Results: Bad



The [Section 2.1.4.2, “Signature Checking Using GnuPG”](#) section explains why you probably don't see a green `Good signature` result.

#### 2.1.4.4. Signature Checking Using RPM

For RPM packages, there is no separate signature. RPM packages have a built-in GPG signature and MD5 checksum. You can verify a package by running the following command:

```
shell> rpm --checksig package_name.rpm
```

Example:

```
shell> rpm --checksig MySQL-server-5.6.12-0.glibc23.i386.rpm
MySQL-server-5.6.12-0.glibc23.i386.rpm: md5 gpg OK
```



#### Note

If you are using RPM 4.1 and it complains about `(GPG) NOT OK (MISSING KEYS: GPG#5072e1f5)`, even though you have imported the MySQL public build key into your own GPG keyring, you need to import the key into the RPM keyring first. RPM 4.1 no longer uses your personal GPG keyring (or GPG itself). Rather, RPM maintains a separate keyring because it is a system-wide application and a user's GPG public keyring is a user-specific file. To import the MySQL public key into the RPM keyring, first obtain the key, then use `rpm --import` to import the key. For example:

```
shell> gpg --export -a 5072elf5 > 5072elf5.asc
shell> rpm --import 5072elf5.asc
```

Alternatively, `rpm` also supports loading the key directly from a URL, and you can use this manual page:

```
shell> rpm --import http://dev.mysql.com/doc/refman/5.6/en/checking-gpg-signature.html
```

If you need to obtain the MySQL public key, see [Section 2.1.4.2, “Signature Checking Using GnuPG”](#).

## 2.1.5. Installation Layouts

The installation layout differs for different installation types (for example, native packages, binary tarballs, and source tarballs), which can lead to confusion when managing different systems or using different installation sources. The individual layouts are given in the corresponding installation type or platform chapter, as described following. Note that the layout of installations from vendors other than Oracle may differ from these layouts.

- [Section 2.3.1, “MySQL Installation Layout on Microsoft Windows”](#)
- [Section 2.9.1, “MySQL Layout for Source Installation”](#)
- [Table 2.3, “MySQL Installation Layout for Generic Unix/Linux Binary Package”](#)
- [Table 2.9, “MySQL Installation Layout for Linux RPM”](#)
- [Table 2.6, “MySQL Installation Layout on Mac OS X”](#)

## 2.1.6. Compiler-Specific Build Characteristics

In some cases, the compiler used to build MySQL affects the features available for use. The notes in this section apply for binary distributions provided by Oracle Corporation or that you compile yourself from source.

### `icc` (Intel C++ Compiler) Builds

A server built with `icc` has these characteristics:

- SSL support is not included.

## 2.2. Installing MySQL from Generic Binaries on Unix/Linux

Oracle provides a set of binary distributions of MySQL. These include binary distributions in the form of compressed `tar` files (files with a `.tar.gz` extension) for a number of platforms, as well as binaries in platform-specific package formats for selected platforms.

This section covers the installation of MySQL from a compressed `tar` file binary distribution. For other platform-specific package formats, see the other platform-specific sections. For example, for Windows distributions, see [Section 2.3, “Installing MySQL on Microsoft Windows”](#).

To obtain MySQL, see [Section 2.1.3, “How to Get MySQL”](#).

MySQL compressed `tar` file binary distributions have names of the form `mysql-VERSION-OS.tar.gz`, where `VERSION` is a number (for example, `5.6.12`), and `OS` indicates the type of operating system for which the distribution is intended (for example, `pc-linux-i686` or `winx64`).

To install MySQL from a compressed `tar` file binary distribution, your system must have GNU `gunzip` to uncompress the distribution and a reasonable `tar` to unpack it. If your `tar` program supports the `z` option, it can both uncompress and unpack the file.

GNU `tar` is known to work. The standard `tar` provided with some operating systems is not able to unpack the long file names in the MySQL distribution. You should download and install GNU `tar`, or if available, use a preinstalled version of GNU `tar`. Usually this is available as `gnutar`, `gtar`, or as `tar` within a GNU or Free Software directory, such as `/usr/sfw/bin` or `/usr/local/bin`. GNU `tar` is available from <http://www.gnu.org/software/tar/>.



### Warning

If you have previously installed MySQL using your operating system native package management system, such as `yum` or `apt-get`, you may experience problems installing using a native binary. Make sure your previous MySQL previous installation has been removed entirely (using your package management system), and that any additional files, such as old versions of your data files, have also been removed. You should also check the existence of configuration files such as `/etc/my.cnf` or the `/etc/mysql` directory have been deleted.

If you run into problems and need to file a bug report, please use the instructions in [Section 1.7, “How to Report Bugs or Problems”](#).

On Unix, to install a compressed `tar` file binary distribution, unpack it at the installation location you choose (typically `/usr/local/mysql`). This creates the directories shown in the following table.

**Table 2.3. MySQL Installation Layout for Generic Unix/Linux Binary Package**

Directory	Contents of Directory
<code>bin</code>	Client programs and the <code>mysqld</code> server
<code>data</code>	Log files, databases
<code>docs</code>	Manual in Info format
<code>man</code>	Unix manual pages
<code>include</code>	Include (header) files
<code>lib</code>	Libraries
<code>scripts</code>	<code>mysql_install_db</code>
<code>share</code>	Miscellaneous support files, including error messages, sample configuration files, SQL for database installation
<code>sql-bench</code>	Benchmarks

Debug versions of the `mysqld` binary are available as `mysqld-debug`. To compile your own debug version of MySQL from a source distribution, use the appropriate configuration options to enable debugging support. For more information on compiling from source, see [Section 2.9, “Installing MySQL from Source”](#).

To install and use a MySQL binary distribution, the basic command sequence looks like this:

```
shell> groupadd mysql
shell> useradd -r -g mysql mysql
shell> cd /usr/local
shell> tar zxvf /path/to/mysql-VERSION-OS.tar.gz
shell> ln -s full-path-to-mysql-VERSION-OS mysql
shell> cd mysql
shell> chown -R mysql .
shell> chgrp -R mysql .
shell> scripts/mysql_install_db --user=mysql
shell> chown -R root .
shell> chown -R mysql data
# Next command is optional
shell> cp support-files/my-medium.cnf /etc/my.cnf
shell> bin/mysqld_safe --user=mysql &
```

```
# Next command is optional
shell> cp support-files/mysql.server /etc/init.d/mysql.server
```

A more detailed version of the preceding description for installing a binary distribution follows.



### Note

This procedure assumes that you have `root` (administrator) access to your system. Alternatively, you can prefix each command using the `sudo` (Linux) or `pfexec` (OpenSolaris) command.

The procedure does not set up any passwords for MySQL accounts. After following the procedure, proceed to [Section 2.10, “Postinstallation Setup and Testing”](#).

## Create a `mysql` User and Group

If your system does not already have a user and group for `mysqld` to run as, you may need to create one. The following commands add the `mysql` group and the `mysql` user. You might want to call the user and group something else instead of `mysql`. If so, substitute the appropriate name in the following instructions. The syntax for `useradd` and `groupadd` may differ slightly on different versions of Unix, or they may have different names such as `adduser` and `addgroup`.

```
shell> groupadd mysql
shell> useradd -r -g mysql mysql
```



### Note

Because the user is required only for ownership purposes, not login purposes, the `useradd` command uses the `-r` option to create a user that does not have login permissions to your server host. Omit this option to permit logins for the user (or if your `useradd` does not support the option).

## Obtain and Unpack the Distribution

Pick the directory under which you want to unpack the distribution and change location into it. The example here unpacks the distribution under `/usr/local`. The instructions, therefore, assume that you have permission to create files and directories in `/usr/local`. If that directory is protected, you must perform the installation as `root`.

```
shell> cd /usr/local
```

Obtain a distribution file using the instructions in [Section 2.1.3, “How to Get MySQL”](#). For a given release, binary distributions for all platforms are built from the same MySQL source distribution.

Unpack the distribution, which creates the installation directory. Then create a symbolic link to that directory. `tar` can uncompress and unpack the distribution if it has `z` option support:

```
shell> tar zxvf /path/to/mysql-VERSION-OS.tar.gz
shell> ln -s full-path-to-mysql-VERSION-OS mysql
```

The `tar` command creates a directory named `mysql-VERSION-OS`. The `ln` command makes a symbolic link to that directory. This enables you to refer more easily to the installation directory as `/usr/local/mysql`.

If your `tar` does not have `z` option support, use `gunzip` to unpack the distribution and `tar` to unpack it. Replace the preceding `tar` command with the following alternative command to uncompress and extract the distribution:

```
shell> gunzip < /path/to/mysql-VERSION-OS.tar.gz | tar xvf -
```

## Perform Postinstallation Setup

The remainder of the installation process involves setting up the configuration file, creating the core databases, and starting the MySQL server. For instructions, see [Section 2.10, “Postinstallation Setup and Testing”](#).



### Note

The accounts that are listed in the MySQL grant tables initially have no passwords. After starting the server, you should set up passwords for them using the instructions in [Section 2.10, “Postinstallation Setup and Testing”](#).

## 2.3. Installing MySQL on Microsoft Windows

MySQL for Microsoft Windows is available in a number of different forms. A Microsoft Windows operating system such as Windows Vista, Windows 7, Windows 8, Windows Server 2008, or Windows Server 2012. Both 32-bit and 64-bit versions are supported. For supported platform information, see <http://www.mysql.com/support/supportedplatforms/database.html>.

In addition to running MySQL as a standard application, you can also run the MySQL server as a Windows service. By using a service you can monitor and control the operation of the server through the standard Windows service management tools. For more information, see [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

Generally, you should install MySQL on Windows using an account that has administrator rights. Otherwise, you may encounter problems with certain operations such as editing the `PATH` environment variable or accessing the `Service Control Manager`. Once installed, MySQL does not need to be executed using a user with Administrator privileges.

For a list of limitations within the Windows version of MySQL, see [Section E.10.6, “Windows Platform Limitations”](#).

In addition to the MySQL Server package, you may need or want additional components to use MySQL with your application or development environment. These include, but are not limited to:

- If you plan to connect to the MySQL server using ODBC, you need a Connector/ODBC driver. For more information, including installation and configuration instructions, see [Section 21.1, “MySQL Connector/ODBC”](#). But note that MySQL Installer will install and configure Connector/ODBC for you.
- If you plan to use MySQL server with .NET applications, you need the Connector/Net driver. For more information, including installation and configuration instructions, see [Section 21.2, “MySQL Connector/Net”](#). But note that MySQL Installer will install and configure Connector/NET for you.

MySQL distributions for Windows can be downloaded from <http://dev.mysql.com/downloads/>. See [Section 2.1.3, “How to Get MySQL”](#).

MySQL for Windows is available in several distribution formats, detailed below. Generally speaking, you should use MySQL Installer. It is simpler to use than the ZIP file, and you need no additional tools to get MySQL up and running. MySQL Installer will automatically install MySQL Server, install additional MySQL products, create an options file, starts the server, and allow you to create default user accounts. For more information on choosing a package, see [Section 2.3.2, “Choosing An Installation Package”](#).

- MySQL Installer. It includes MySQL Server and additional MySQL products including MySQL Workbench, MySQL Notifier, and MySQL for Excel. And MySQL Installer can also be used to upgrade these product in the future.

For instructions on installing MySQL using MySQL Installer, see [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#).

- Standard binary distribution format packaged as a Zip file containing all of the necessary files that you unpack into your chosen location. This package contains all of the files in the full Windows MSI Installer package, but does not include an installation program.

For instructions on installing MySQL using the Zip file, see [Section 2.3.4, “Installing MySQL on Microsoft Windows Using a `noinstall` Zip Archive”](#).

- The source distribution contains all the code and support files for building the executables using the Visual Studio compiler system.

For instructions on building MySQL from source on Windows, see [Section 2.9, “Installing MySQL from Source”](#).

**MySQL on Windows considerations:**

- **Large Table Support**

If you need tables with a size larger than 4GB, install MySQL on an NTFS or newer file system. Do not forget to use `MAX_ROWS` and `AVG_ROW_LENGTH` when you create tables. See [Section 13.1.14, “CREATE TABLE Syntax”](#).

- **MySQL and Virus Checking Software**

Using virus scanning software such as Norton/Symantec Anti-Virus on directories containing MySQL data and temporary tables can cause issues, both in terms of the performance of MySQL and the virus-scanning software mis-identifying the contents of the files as containing spam. This is because of the fingerprinting mechanism used by the virus scanning software, and the way in which MySQL rapidly updates different files, which may be identified as a potential security risk.

After installing MySQL Server, it is recommended that you disable virus scanning on the main directory (`datadir` [472]) being used to store your MySQL table data. There is usually a system built into the virus scanning software to permit certain directories to be specifically ignored during virus scanning.

In addition, by default, MySQL creates temporary files in the standard Windows temporary directory. To prevent the temporary files also being scanned, you should configure a separate temporary directory for MySQL temporary files and add this to the virus scanning exclusion list. To do this, add a configuration option for the `tmpdir` [446] parameter to your `my.ini` configuration file. For more information, see [Section 2.3.4.2, “Creating an Option File”](#).

## 2.3.1. MySQL Installation Layout on Microsoft Windows

For MySQL 5.6 on Windows, the default installation directory is `C:\Program Files\MySQL\MySQL Server 5.6`. Some Windows users prefer to install in `C:\mysql`, the directory that formerly was used as the default. However, the layout of the subdirectories remains the same.

All of the files are located within this parent directory, using the structure shown in the following table.

**Table 2.4. MySQL Installation Layout for Windows**

Directory	Contents of Directory
<code>bin</code>	Client programs and the <code>mysqld</code> server
<code>C:\Documents and Settings\All Users\Application Data\MySQL</code>	Log files, databases (Windows XP, Windows Server 2003)
<code>C:\ProgramData\MySQL</code>	Log files, databases (Windows 7, Windows Server 2008, and newer)
<code>examples</code>	Example programs and scripts
<code>include</code>	Include (header) files



Directory	Contents of Directory
<code>lib</code>	Libraries
<code>scripts</code>	Utility scripts
<code>share</code>	Miscellaneous support files, including error messages, character set files, sample configuration files, SQL for database installation

If you install MySQL using the MySQL Installer, this package creates and sets up the data directory that the installed server will use, but also creates a pristine “template” data directory named `data` under the installation directory. After an installation has been performed using this package, the template data directory can be copied to set up additional MySQL instances. See [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).

## 2.3.2. Choosing An Installation Package

For MySQL 5.6, there are installation package formats to choose from when installing MySQL on Windows:

- **MySQL Installer:** This package has a file name similar to `mysql-installer-community-5.6.12.0.msi` or `mysql-installer-commercial-5.6.12.0.msi`, and utilizes MSI's to automatically install MySQL server and other products. It will download and apply updates to itself, and for each of the installed products. It also configures the additional non-server products.

The installed products are configurable, and this includes: documentation with samples and examples, connectors (such as C, C++, J, NET, and ODBC), MySQL Workbench, MySQL Notifier for Microsoft Windows, MySQL for Excel, and the MySQL Server with its components.

MySQL Installer will run on all Windows platforms that are [supported by MySQL](#). Note: because MySQL is not a native component of Windows and depends on .NET, it will not work on minimal installation options like the "Server Core" version of Windows Server 2008.

For instructions on installing MySQL using MySQL Installer, see [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#).

- **The Noinstall Archive:** This package has a file name similar to `mysql-5.6.12-win32.zip` or `mysql-5.6.12-winx64.zip`, and contains all the files found in the Complete install package, with the exception of the GUI. This package does not include an automated installer, and must be manually installed and configured.

Your choice of install package affects the installation process you must follow. If you choose to use MySQL Installer, see [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#). If you choose to install a Noinstall archive, see [Section 2.3.4, “Installing MySQL on Microsoft Windows Using a noinstall Zip Archive”](#).

## 2.3.3. Installing MySQL on Microsoft Windows Using MySQL Installer

MySQL Installer is an application that simplifies the installation and updating process for a wide range of MySQL products, including [MySQL Notifier for Microsoft Windows](#), [MySQL Workbench](#), and [MySQL for Excel](#). From this central application, you can see which MySQL products are already installed, configure them, and update or remove them if necessary. The installer can also install plugins, documentation, tutorials, and example databases. The MySQL Installer is only available for Microsoft Windows, and includes both a GUI and command-line interface.

### Community edition

Download the community edition of MySQL Installer at <http://dev.mysql.com/downloads/installer/>. It will install the community version of all MySQL products.

## Commercial edition

Download the commercial edition of MySQL Installer at either [My Oracle Support \(MOS\)](#) or <https://edelivery.oracle.com/>. It will install the commercial version of all MySQL products, including Workbench SE. It also integrates with your MOS account, so enter in your MOS credentials to automatically receive updates for your commercial MySQL products.

For release notes detailing the changes in each release of MySQL Installer, see [MySQL Installer Release Notes](#).

MySQL Installer is compatible with pre-existing installations, and adds them to its list of installed components. And while the MySQL Installer is bundled with a specific version of MySQL Server, a single MySQL Installer instance can install and manage multiple MySQL Server versions. For example, a single MySQL Installer instance can install versions 5.1, 5.5, and 5.6. It can also manage either commercial or community editions of the MySQL Server.



### Note

A single host can not have both community and commercial editions of MySQL Server installed. For example, if you want both MySQL Server 5.5 and 5.6 installed on a single host, then both must be the same commercial or community edition.

MySQL Installer handles the initial configuration and setup of the applications. For example:

1. It will create MySQL Server connections in MySQL Workbench.
2. It creates the configuration file (`my.ini`) that is used to configure the MySQL Server. The values written to this file are influenced by choices you make during the installation process.
3. It imports example databases.
4. It creates MySQL Server user accounts with configurable permissions based on general roles, such as DB Administrator, DB Designer, and Backup Admin. It optionally creates a Windows user named `MysqISys` with limited privileges, which would then run the MySQL Server.

This feature is only available during the initial installation of the MySQL Server, and not during future updates. User accounts may also be added with MySQL Workbench.

5. If the "Advanced Configuration" option is checked, then the **Logging Options** are also configured. This includes defining file paths for the error log, general log, slow query log (including the configuration of seconds it requires to execute a query), and the binary log.

MySQL Installer can optionally check for updated components and download them for you automatically.

### 2.3.3.1. MySQL Installer GUI

After installation of the GUI version, the installer will have add its own Start Menu item under MySQL.

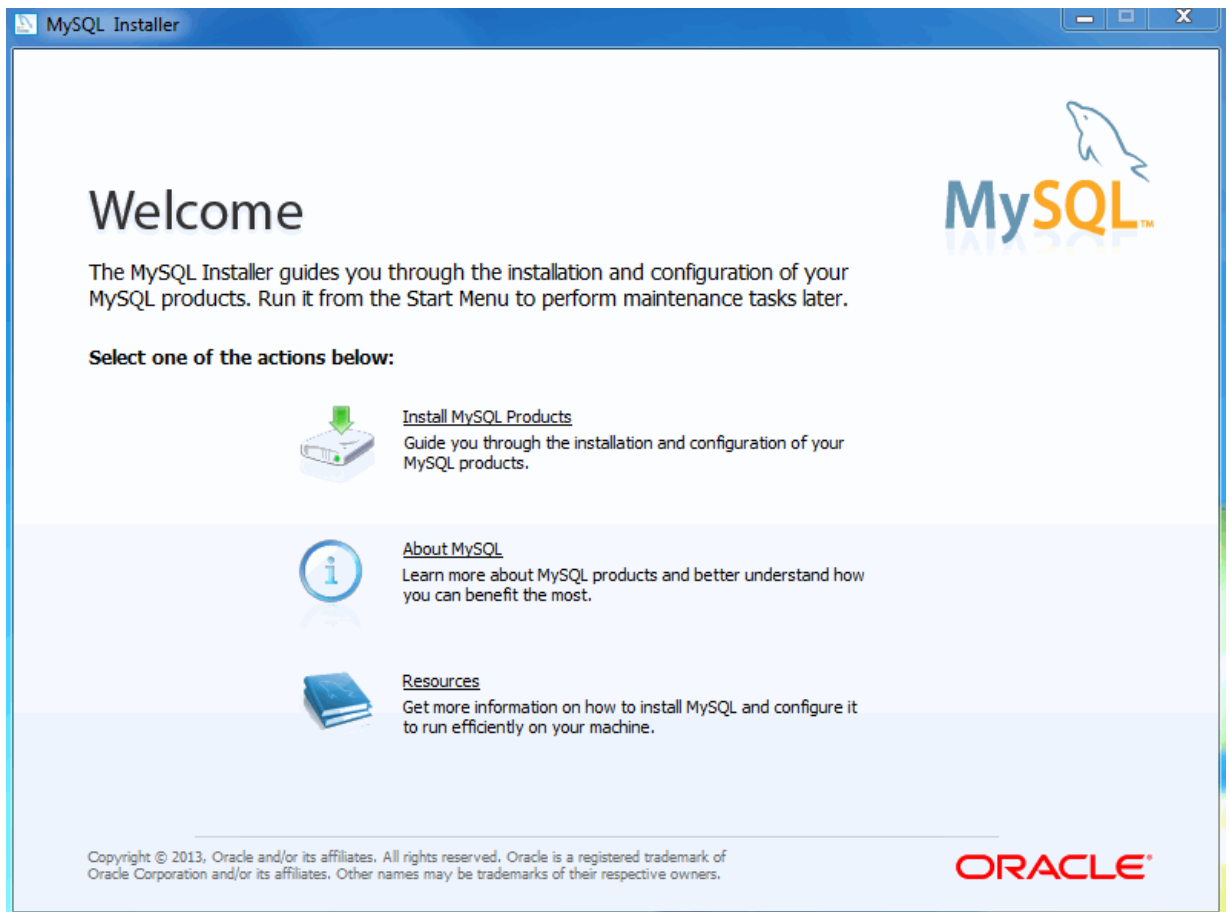


### Note

Files that are generated by MySQL Installer grant full permissions to the user that executes MySQL Installer, including `my.ini`. This does not apply to files and directories for specific products such as the MySQL Server data directory in `ProgramData`, that is owned by `SYSTEM`.

After the installer itself has been installed and started, the following screen is displayed:

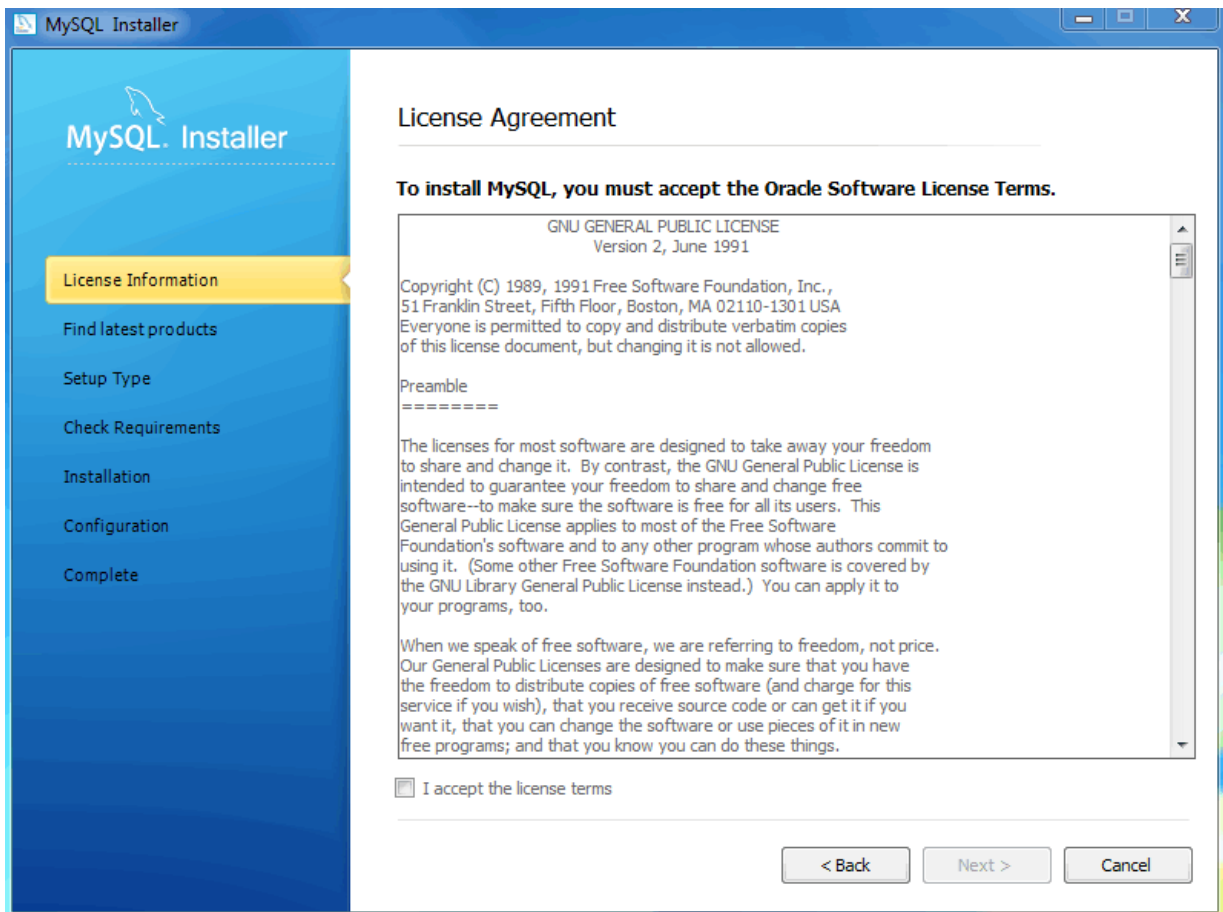
Figure 2.7. MySQL Installer - Welcome Screen



There are three main options:

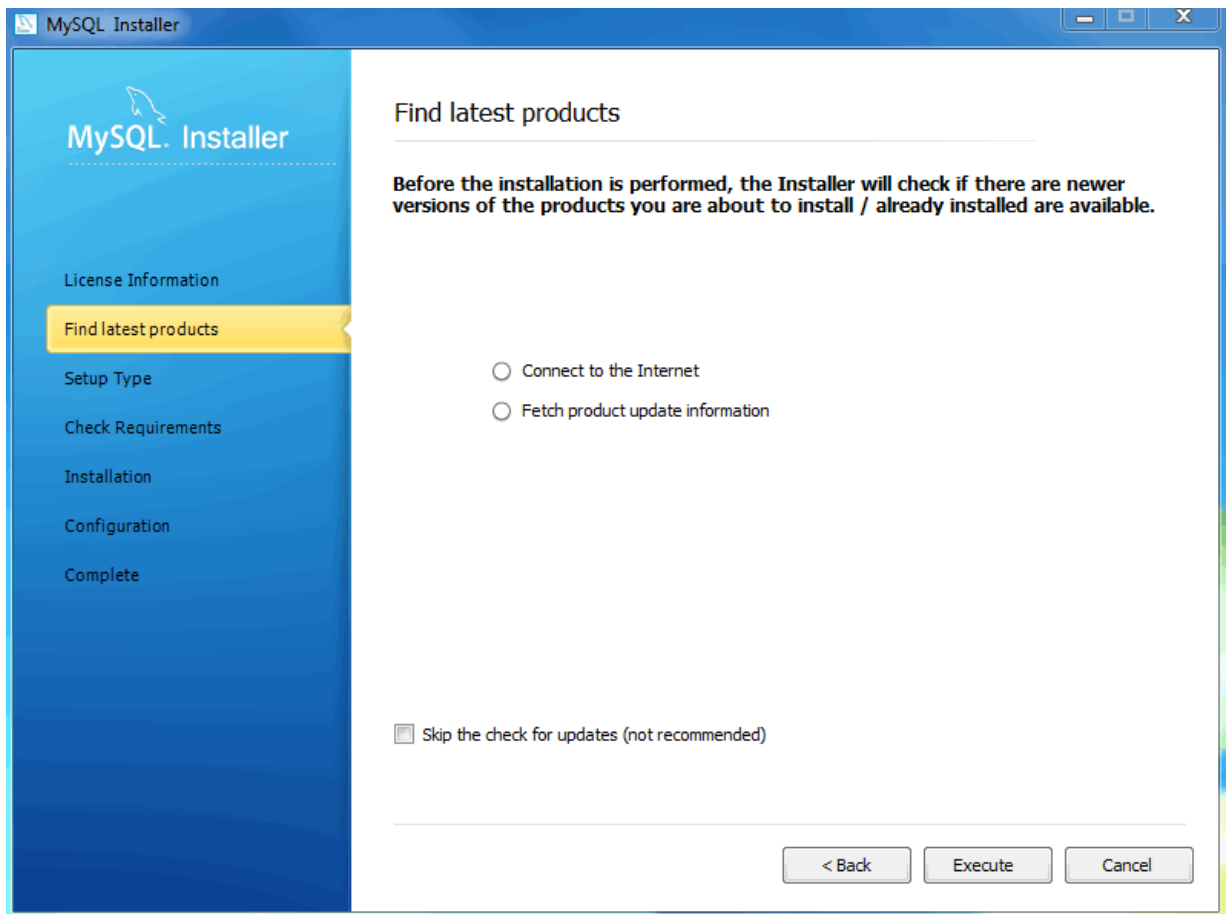
1. Install MySQL Products - The Installation Wizard.
2. About MySQL - Learn about MySQL products and features.
3. Resources - Information to help install and configure MySQL.

To **Install MySQL Products** after executing MySQL Installer for the first time, you must accept the license agreement before proceeding with the installation process.

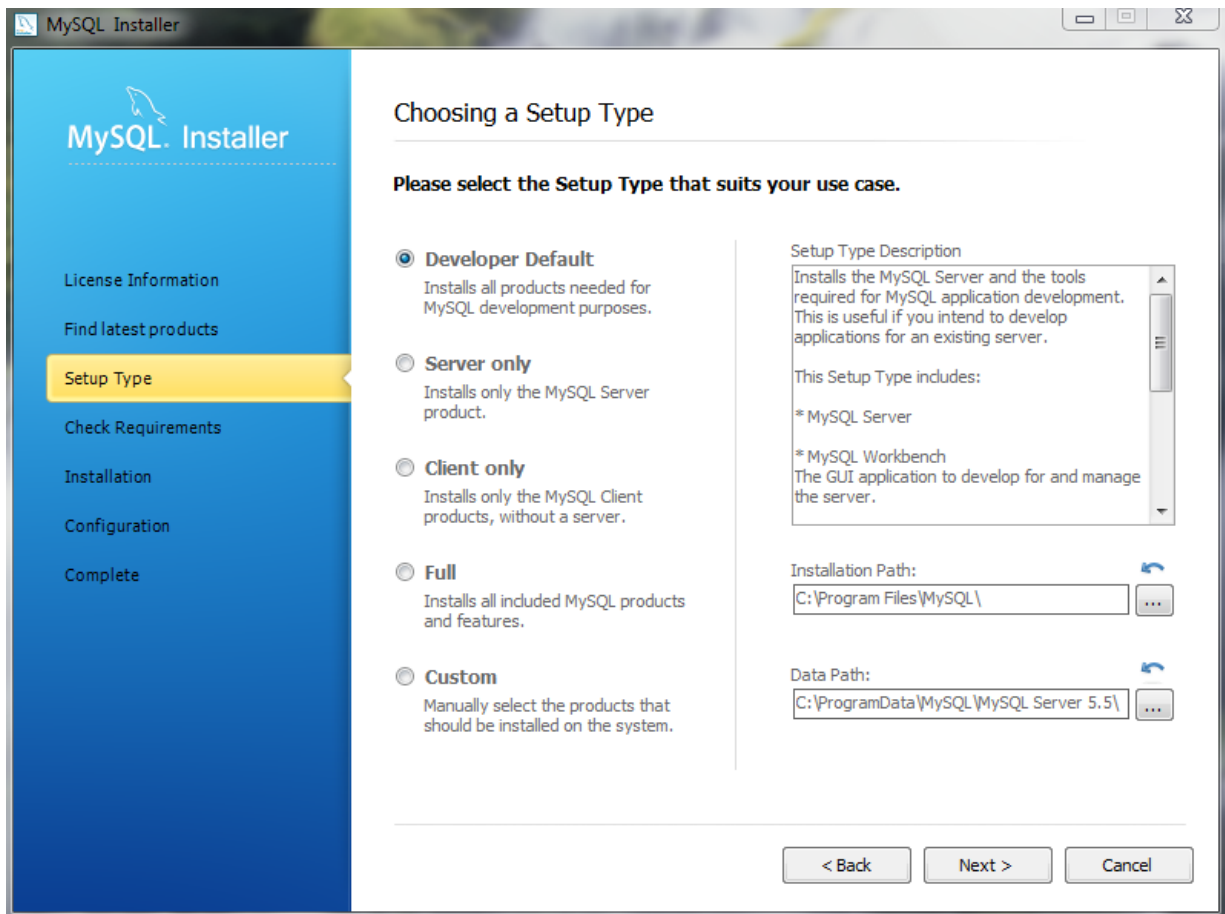
**Figure 2.8. MySQL Installer - License Agreement**

If you are connected to the Internet, then the Installer will search for the latest MySQL components and add them to the installation bundle. Click **Connect to the Internet** to complete this step, or otherwise check the **Skip** checkbox and then **Continue**.

**Figure 2.9. MySQL Installer - Find latest products**



If you chose "Connect to the Internet," the next page will show the progress of MySQL Installer's search for available updates. When the search is complete (or if you opted to skip the search), you will be taken to the **Choose Setup Type** page:

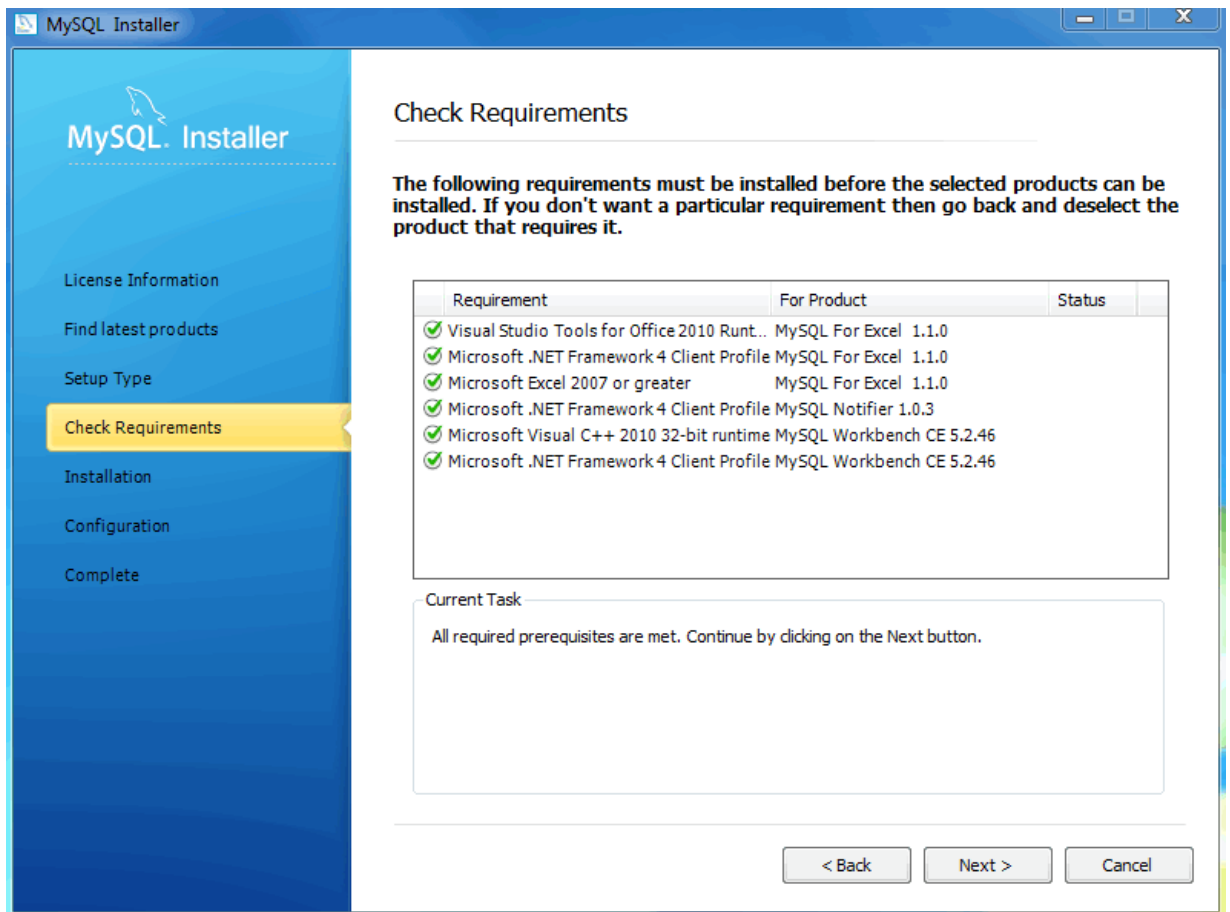
**Figure 2.10. MySQL Installer - Choosing a Setup Type**

Determine the option most compatible with your preferences by reading the **Setup Type Description** descriptions.

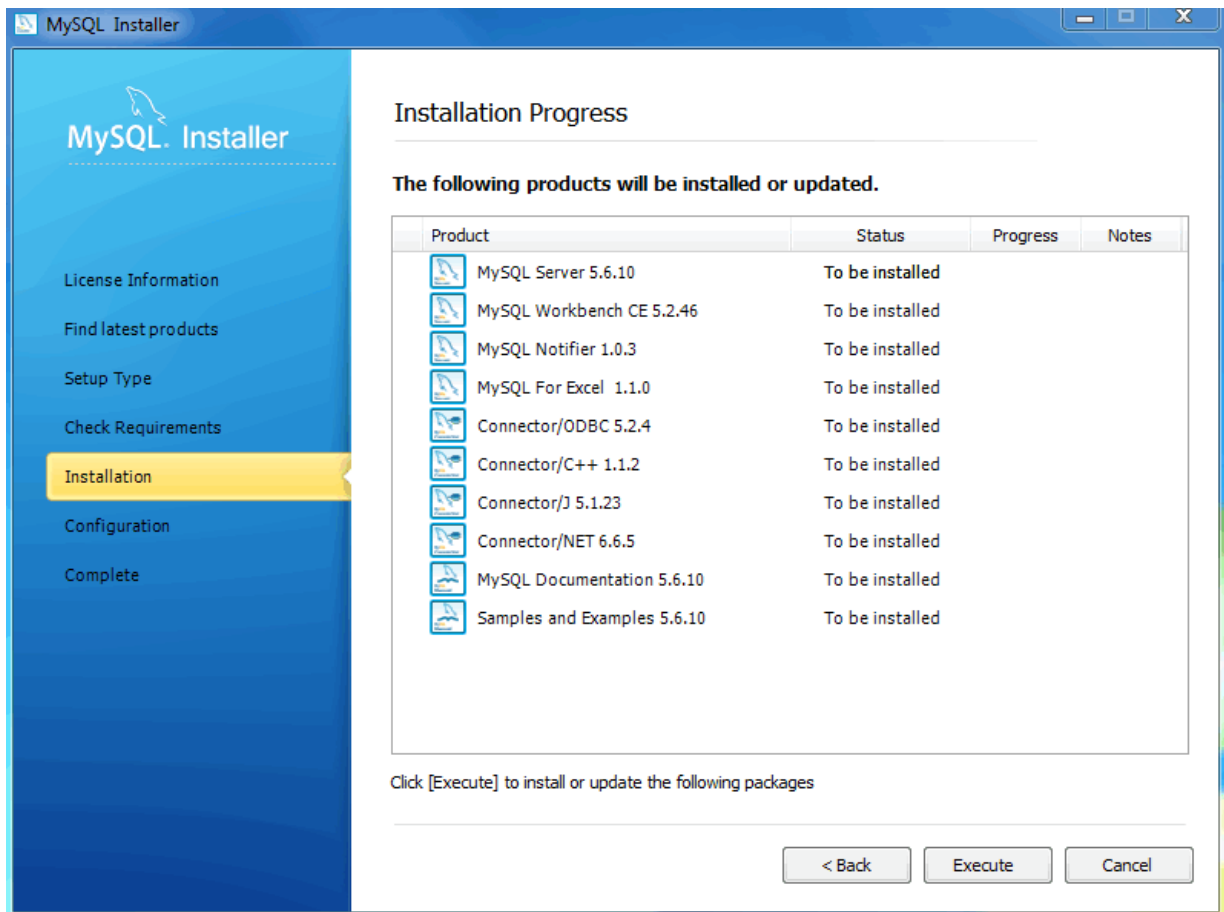
The Installation and Data paths are also defined here, and a caution flag will notify you if the data path you define already exists.

After you select a setup type, the MySQL Installer will check your system for the necessary external requirements and download then install missing components onto your system.

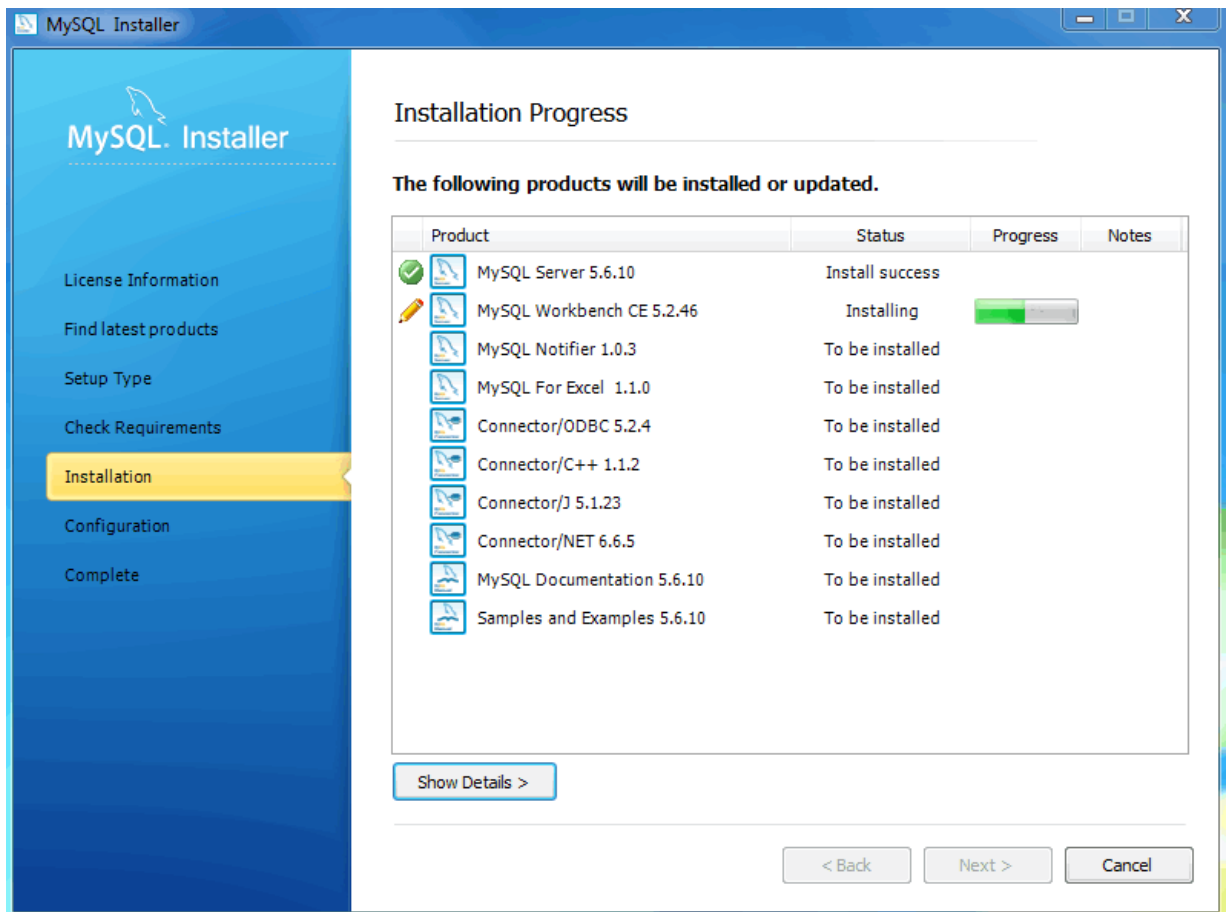


**Figure 2.11. MySQL Installer - Check Requirements**

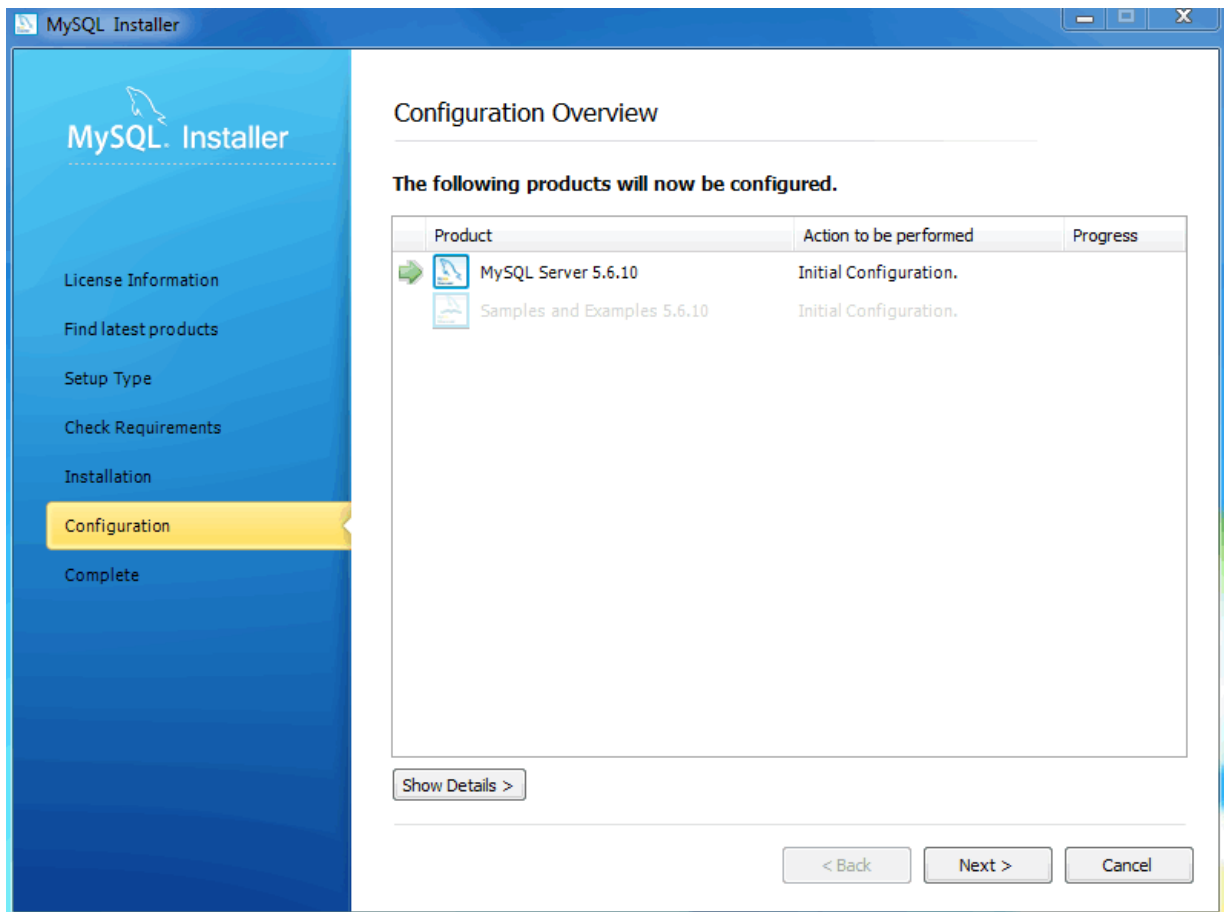
The next window lists the MySQL products that are scheduled to be installed:

**Figure 2.12. MySQL Installer - Installation Progress**

As components are installed, you'll see their status change from "to be installed" to "install success."

**Figure 2.13. MySQL Installer - Installation Progress status**

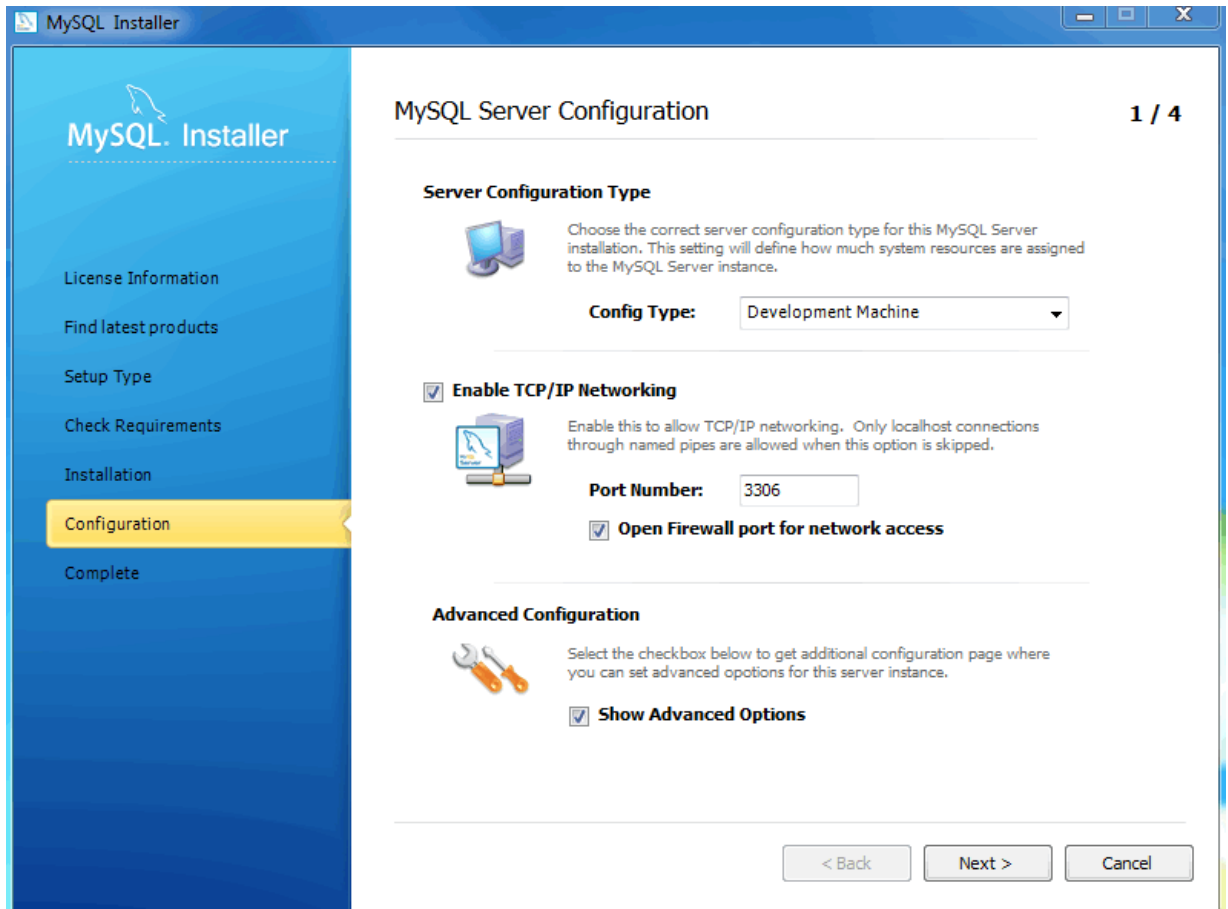
After all components are installed, the next step involves configuring the products. The [Configuration Overview](#) window displays the progress and then loads a configuration window if it is required.

**Figure 2.14. MySQL Installer - Configuration Overview**

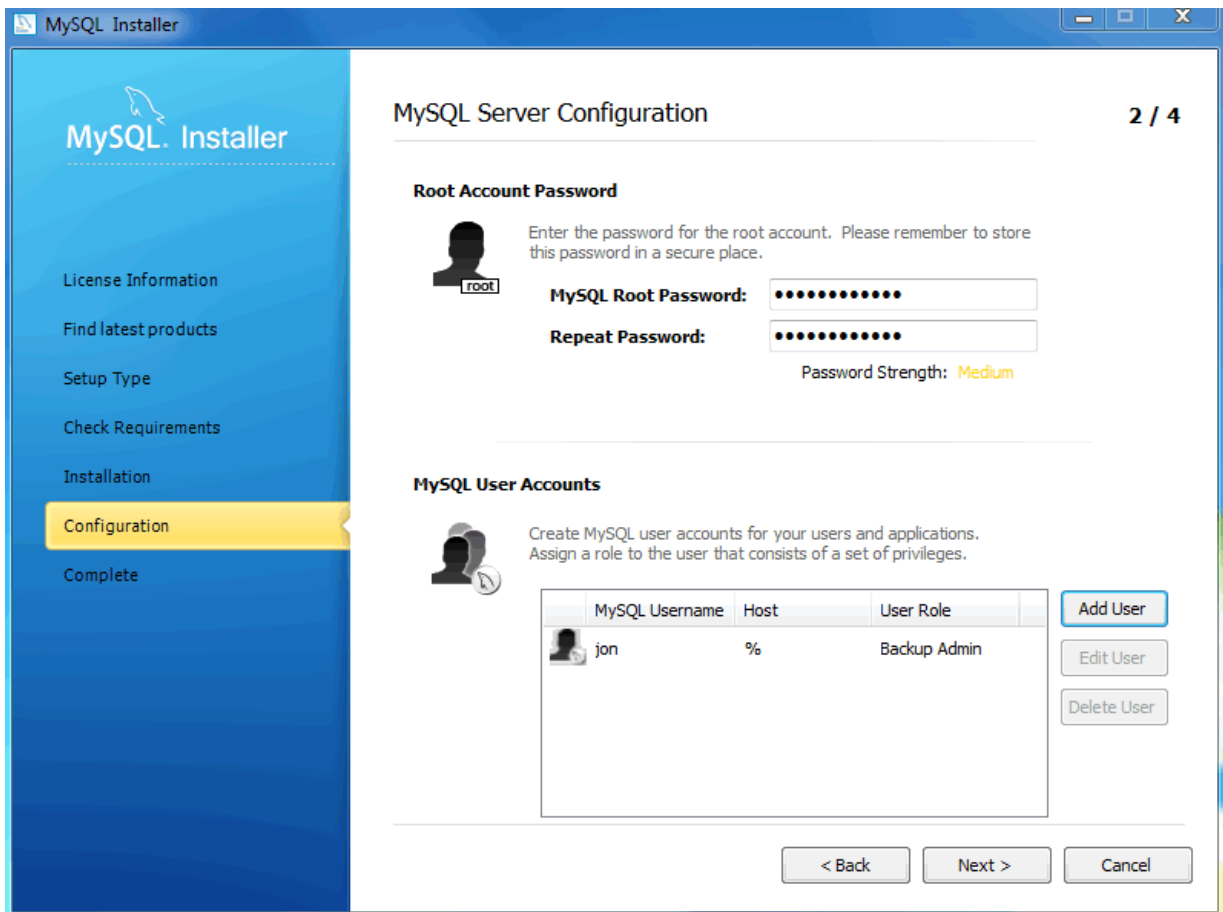
The ideal MySQL Server configuration depends on your intended use, as explained in the next window. Choose the description that most closely applies to your machine.

You may enable TCP/IP Networking here as otherwise only localhost connections are allowed.

Checking the "Advanced Configuration" option provides additional **Logging Options** to configure. This includes defining file paths for the error log, general log, slow query log (including the configuration of seconds it requires to execute a query), and the binary log.

**Figure 2.15. MySQL Installer - MySQL Server Configuration: Define platform, networking, and logging options**

Next, choose your account information. Defining a root password is required, whereas it's optional to create additional users. There are several different predefined user roles that each have different permission levels. For example, a "DB Admin" will have more privileges than a "DB Designer."

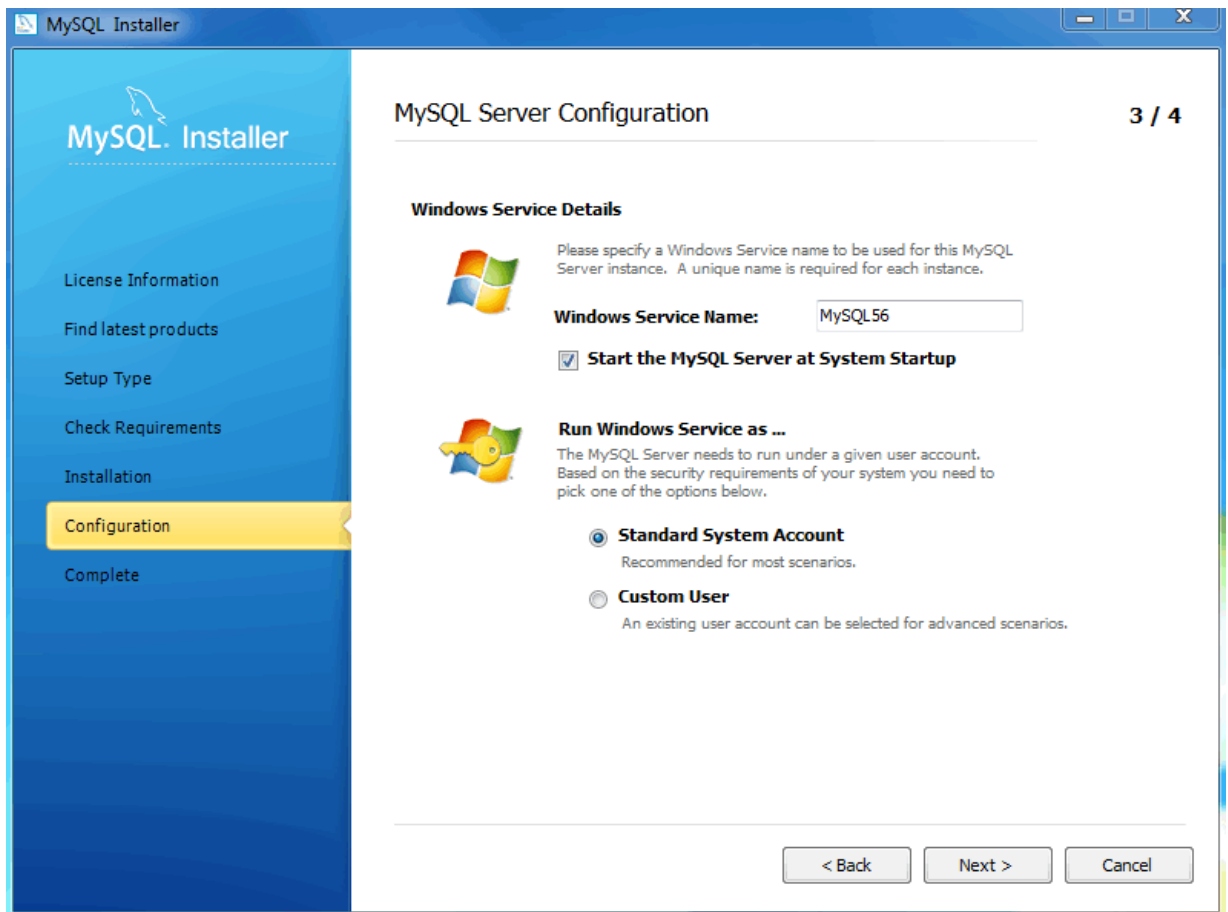
**Figure 2.16. MySQL Installer - MySQL Server Configuration: User accounts****Note**

If the MySQL Server is already installed, then the [Current Root Password](#) will also be needed.

Next, configure the **Windows Service Details**. This includes the service name, how the MySQL Server should be loaded at startup, and how the Windows Service for MySQL Server will be run.



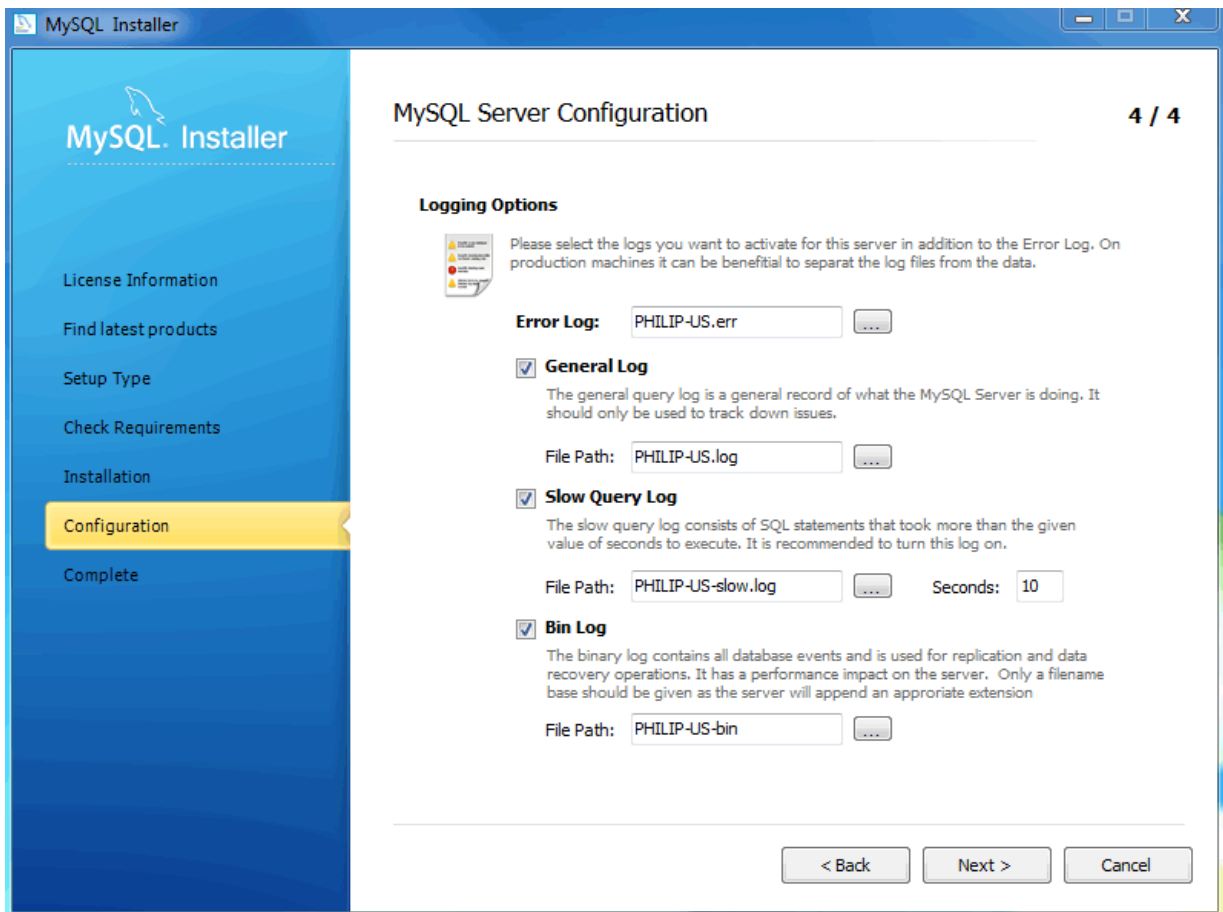
Figure 2.17. MySQL Installer - MySQL Server Configuration: Windows service details

**Note**

When configuring **Run Windows Services as ...** using a **Custom User**, the custom user must have privileges to log on to Windows as a service. And the **Next** button will be disabled until this user is given these user rights.

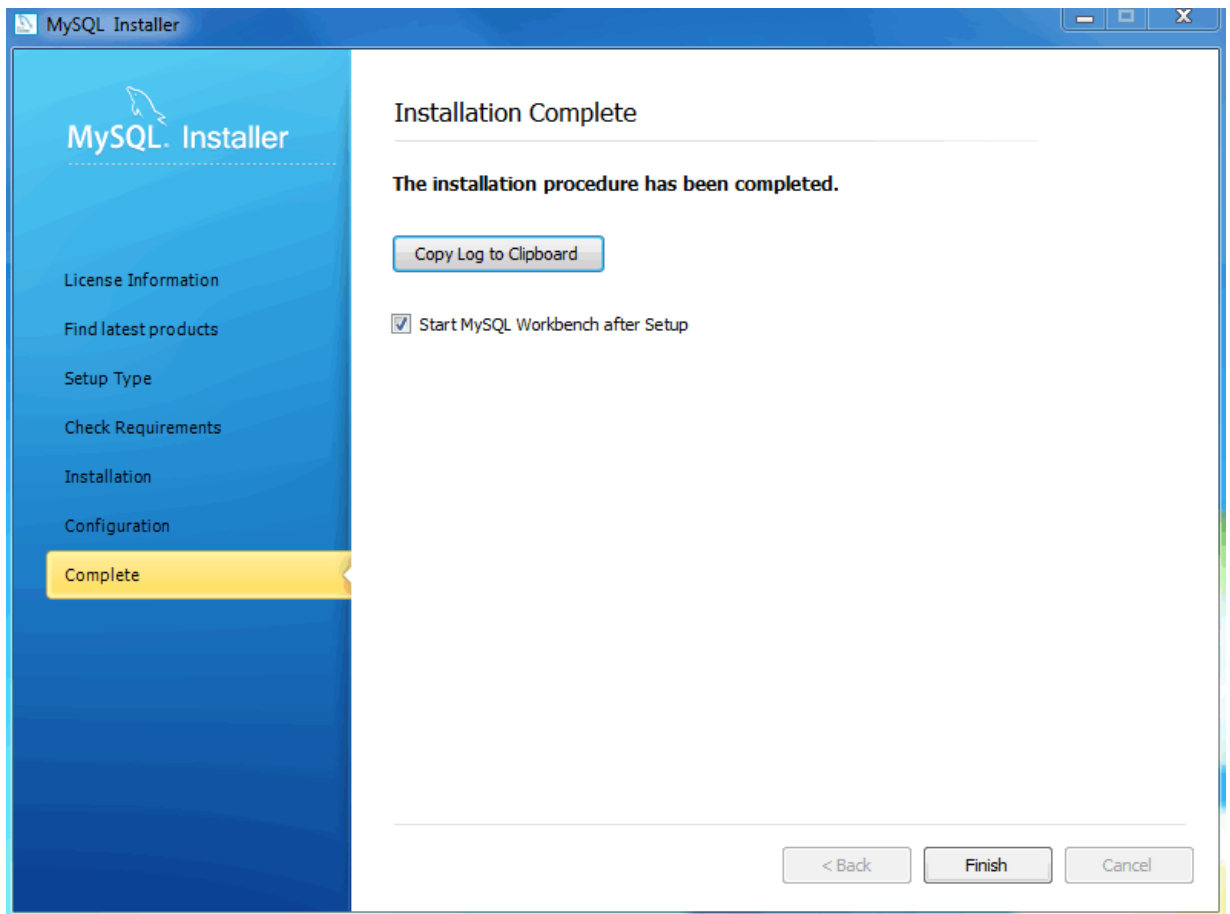
On Microsoft Windows 7, this is configured by loading the [Start Menu](#), [Control Panel](#), [Administrative Tools](#), [Local Security Policy](#), [Local Policies](#), [User Rights Assignment](#), then [Log On As A Service](#). Choose [Add User or Group](#) here to add the custom user, and then **OK**, **OK** to save.

The final configuration step is available if the **Advanced Configuration** option was checked, and it includes configuration options related to log file names:

**Figure 2.18. MySQL Installer - MySQL Server Configuration: Logging options**

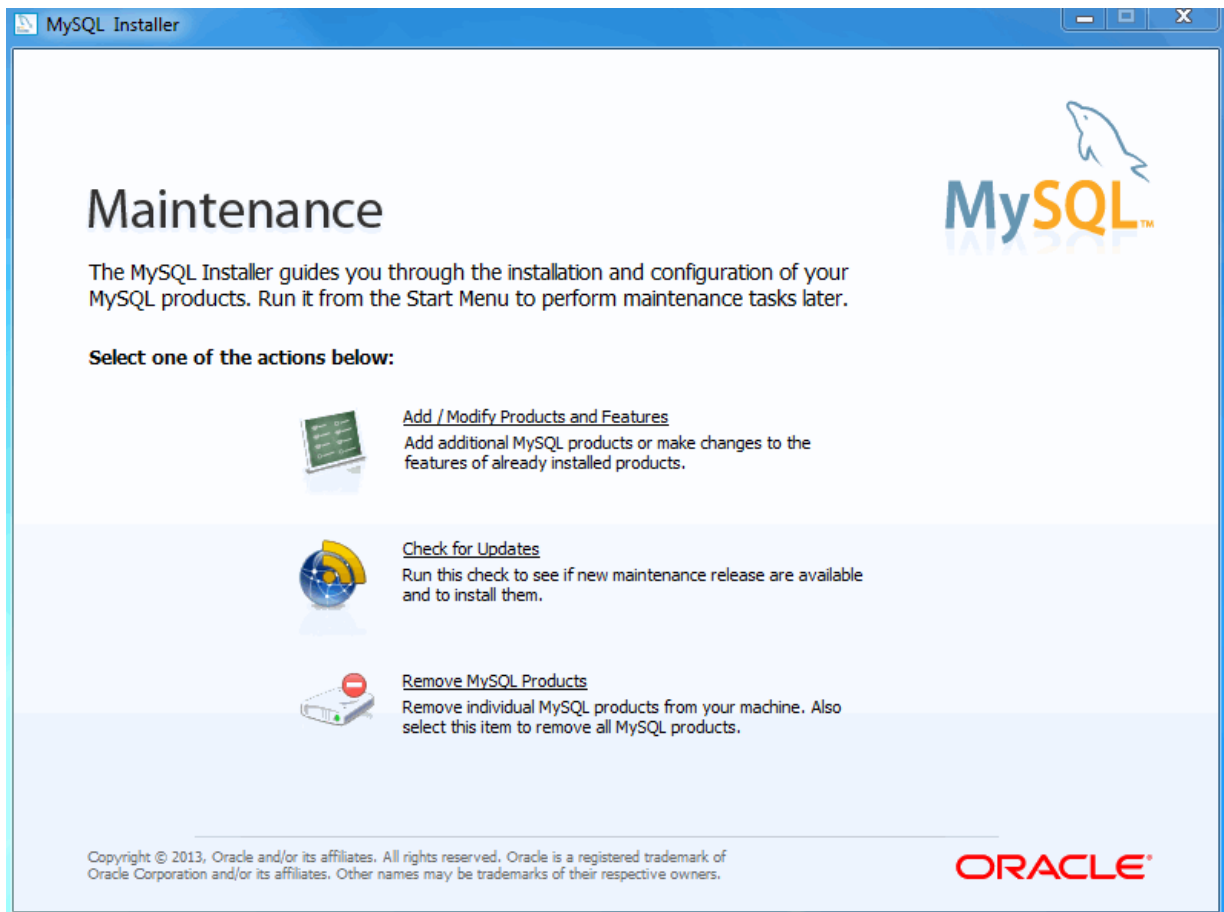
After the MySQL Installer configuration process is completed, you may save the installation log, and then load MySQL Workbench if the **Start MySQL Workbench after Setup** option is checked:

**Figure 2.19. MySQL Installer - Installation Complete**



You can now open MySQL Installer from the Microsoft Windows Start menu under the [MySQL](#) group, which will load the MySQL Installer [Maintenance Screen](#). This is used to add, update, and remove features.

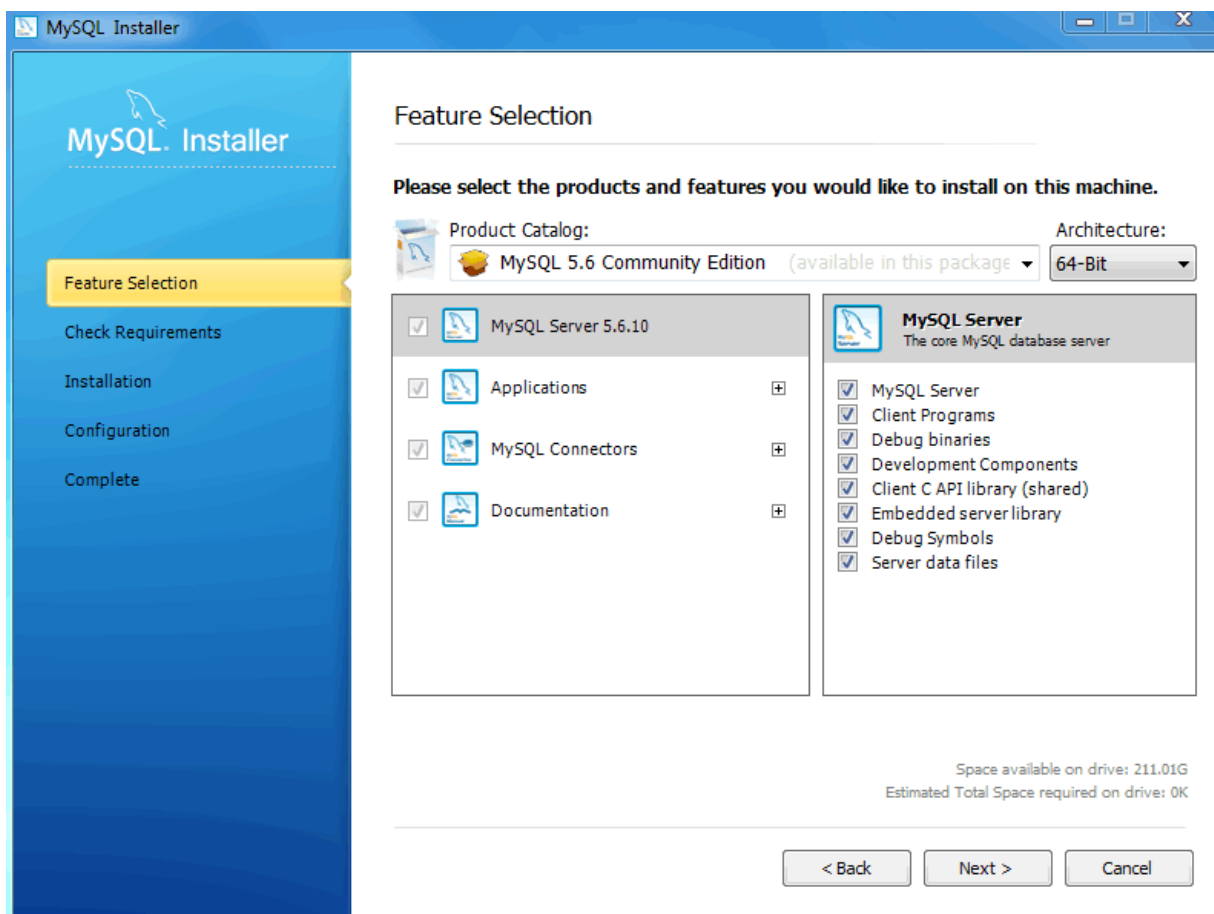
Figure 2.20. MySQL Installer - Maintenance Screen



**Note**

An [Update Screen](#) screen is shown if MySQL Installer is used on a machine with older products installed, as opposed to the [Maintenance Screen](#) shown above. However, the functionality remains the same.

**Add/Modify Products and Features** will list all installed and available MySQL products.

**Figure 2.21. MySQL Installer - Add/Modify Products and Features**

The installation is now complete. MySQL Server should be running, and most MySQL products installed and available for use.

See also the MySQL Workbench documentation (<http://dev.mysql.com/doc/workbench/en/>).

### 2.3.3.2. MySQL Installer Console

`MySQLInstallerConsole` provides functionality similar to the GUI version of MySQL Installer, but from the command-line. It is installed when MySQL Installer is initially executed, and then available within the `MySQL Installer` directory. Typically that is in `C:\Program Files (x86)\MySQL\MySQL Installer\`, and the console must be executed with administrative privileges.

To use, invoke the Command Prompt with administrative privileges by choosing Start, Accessories, then right-click on Command Prompt and choose Run as administrator. And from the command-line, optionally change the directory to where `MySQLInstallerConsole` is located:

```
C:\> cd "C:\Program Files (x86)\MySQL\MySQL Installer"
```

`MySQLInstallerConsole` supports the following options, which are specified on the command line:

- `--help`, `-h`, or `-?`

Displays a help message with usage examples, and then exits.

```
C:\> MySQLInstallerConsole --help
```

- `--updates` (or `-u`)

Checks for new products before any further action is taken. Disabled by default.

- `--nowait`

Skips the final pause when the program finishes. Otherwise, a "Press Enter to continue." dialogue is generated. It is used in conjunction with other options.

- `--catalog=catalog_name` (or `-c`)

Sets the default catalog. Use `--list` to view a list of available catalogs.

- `--type=installation_type` (or `-t`)

Sets the installation type.

The possible values for `installation_type` are: `developer`, `server`, `client`, `full`, and `custom`.

- `--action=action_name`

The action being performed.

The possible values are: `install`, `remove`, `upgrade`, `list`, and `status`.

- **install**: Installs a product or products, as defined by `--products`
- **upgrade**: Upgrades a product or products, as defined by `--products`.
- **remove**: Removes a product or products, as defined by `--products`.
- **list**: Lists the product manifest, both installed and available products.
- **status**: Shows the status after another action is performed.

- `--product=product_name[:feature1],[feature2], [...]` (or `-p`)

Set the feature list of a product. Use `--list` to view available products, or pass in `--product=*` (an asterisk) to install all available products.

- `--config=product_name:passwd=root_password[;parameter1=value],[;parameter2=value], ...`

The configuration parameters for the most recently listed products.

- `--user=product_name:name=username,host:hostname,role=rolename,password=password` or `--user=product_name:name=username,host:hostname,role=rolename,tokens=tokens`

Creates a new user.

Requires: `name`, `host`, `role`, and the `password` or `tokens`. Tokens are separated by pipe ("|") characters.

### 2.3.4. Installing MySQL on Microsoft Windows Using a `noinstall` Zip Archive

Users who are installing from the `noinstall` package can use the instructions in this section to manually install MySQL. The process for installing MySQL from a Zip archive is as follows:

1. Extract the archive to the desired install directory
2. Create an option file

3. Choose a MySQL server type
4. Start the MySQL server
5. Secure the default user accounts

This process is described in the sections that follow.

### 2.3.4.1. Extracting the Install Archive

To install MySQL manually, do the following:

1. If you are upgrading from a previous version please refer to [Section 2.3.6, “Upgrading MySQL on Windows”](#), before beginning the upgrade process.
2. Make sure that you are logged in as a user with administrator privileges.
3. Choose an installation location. Traditionally, the MySQL server is installed in `C:\mysql`. The MySQL Installer installs MySQL under `C:\Program Files\MySQL`. If you do not install MySQL at `C:\mysql`, you must specify the path to the install directory during startup or in an option file. See [Section 2.3.4.2, “Creating an Option File”](#).
4. Extract the install archive to the chosen installation location using your preferred Zip archive tool. Some tools may extract the archive to a folder within your chosen installation location. If this occurs, you can move the contents of the subfolder into the chosen installation location.

### 2.3.4.2. Creating an Option File

If you need to specify startup options when you run the server, you can indicate them on the command line or place them in an option file. For options that are used every time the server starts, you may find it most convenient to use an option file to specify your MySQL configuration. This is particularly true under the following circumstances:

- The installation or data directory locations are different from the default locations (`C:\Program Files\MySQL\MySQL Server 5.6` and `C:\Program Files\MySQL\MySQL Server 5.6\data`).
- You need to tune the server settings, such as memory, cache, or InnoDB configuration information.

When the MySQL server starts on Windows, it looks for option files in several locations, such as the Windows directory, `C:\`, and the MySQL installation directory (for the full list of locations, see [Section 4.2.3.3, “Using Option Files”](#)). The Windows directory typically is named something like `C:\WINDOWS`. You can determine its exact location from the value of the `WINDIR` environment variable using the following command:

```
C:\> echo %WINDIR%
```

MySQL looks for options in each location first in the `my.ini` file, and then in the `my.cnf` file. However, to avoid confusion, it is best if you use only one file. If your PC uses a boot loader where `C:` is not the boot drive, your only option is to use the `my.ini` file. Whichever option file you use, it must be a plain text file.



#### Note

When using the MySQL Installer to install MySQL Server, it will create the `my.ini` at the default location. And as of MySQL Server 5.5.27, the user running MySQL Installer is granted full permissions to this new `my.ini`.

In other words, be sure that the MySQL Server user has permission to read the `my.ini` file.



You can also make use of the example option files included with your MySQL distribution; see [Section 5.1.2, “Server Configuration Defaults”](#).

An option file can be created and modified with any text editor, such as Notepad. For example, if MySQL is installed in `E:\mysql` and the data directory is in `E:\mydata\data`, you can create an option file containing a `[mysqld]` section to specify values for the `basedir` and `datadir` options:

```
[mysqld]
# set basedir to your installation path
basedir=E:/mysql
# set datadir to the location of your data directory
datadir=E:/mydata/data
```

Note that Windows path names are specified in option files using (forward) slashes rather than backslashes. If you do use backslashes, double them:

```
[mysqld]
# set basedir to your installation path
basedir=E:\\mysql
# set datadir to the location of your data directory
datadir=E:\\mydata\\data
```

The rules for use of backslash in option file values are given in [Section 4.2.3.3, “Using Option Files”](#).

The data directory is located within the `AppData` directory for the user running MySQL.

If you would like to use a data directory in a different location, you should copy the entire contents of the `data` directory to the new location. For example, if you want to use `E:\mydata` as the data directory instead, you must do two things:

1. Move the entire `data` directory and all of its contents from the default location (for example `C:\Program Files\MySQL\MySQL Server 5.6\data`) to `E:\mydata`.
2. Use a `--datadir [418]` option to specify the new data directory location each time you start the server.

### 2.3.4.3. Selecting a MySQL Server Type

The following table shows the available servers for Windows in MySQL 5.6.

Binary	Description
<code>mysqld</code>	Optimized binary with named-pipe support
<code>mysqld-debug</code>	Like <code>mysqld</code> , but compiled with full debugging and automatic memory allocation checking

All of the preceding binaries are optimized for modern Intel processors, but should work on any Intel i386-class or higher processor.

Each of the servers in a distribution support the same set of storage engines. The `SHOW ENGINES` statement displays which engines a given server supports.

All Windows MySQL 5.6 servers have support for symbolic linking of database directories.

MySQL supports TCP/IP on all Windows platforms. MySQL servers on Windows support named pipes as indicated in the following list. However, the default is to use TCP/IP regardless of platform. (Named pipes are slower than TCP/IP in many Windows configurations.)

Named pipes are enabled only if you start the server with the `--enable-named-pipe [421]` option. It is necessary to use this option explicitly because some users have experienced problems with shutting down the MySQL server when named pipes were used.

### 2.3.4.4. Starting the Server for the First Time

This section gives a general overview of starting the MySQL server. The following sections provide more specific information for starting the MySQL server from the command line or as a Windows service.

The information here applies primarily if you installed MySQL using the `Noinstall` version, or if you wish to configure and test MySQL manually rather than with the GUI tools.



#### Note

The MySQL server will automatically start after using the MySQL Installer, and the [MySQL Notifier for Microsoft Windows](#) GUI can be used to start/stop/restart at any time.

The examples in these sections assume that MySQL is installed under the default location of `C:\Program Files\MySQL\MySQL Server 5.6`. Adjust the path names shown in the examples if you have MySQL installed in a different location.

Clients have two options. They can use TCP/IP, or they can use a named pipe if the server supports named-pipe connections.

MySQL for Windows also supports shared-memory connections if the server is started with the `--shared-memory` [439] option. Clients can connect through shared memory by using the `--protocol=MEMORY` [211] option.

For information about which server binary to run, see [Section 2.3.4.3, “Selecting a MySQL Server Type”](#).

Testing is best done from a command prompt in a console window (or “DOS window”). In this way you can have the server display status messages in the window where they are easy to see. If something is wrong with your configuration, these messages make it easier for you to identify and fix any problems.

To start the server, enter this command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld" --console
```

For a server that includes `InnoDB` support, you should see the messages similar to those following as it starts (the path names and sizes may differ):

```
InnoDB: The first specified datafile c:\ibdata\ibdata1 did not exist:
InnoDB: a new database to be created!
InnoDB: Setting file c:\ibdata\ibdata1 size to 209715200
InnoDB: Database physically writes the file full: wait...
InnoDB: Log file c:\iblogs\ib_logfile0 did not exist: new to be created
InnoDB: Setting log file c:\iblogs\ib_logfile0 size to 31457280
InnoDB: Log file c:\iblogs\ib_logfile1 did not exist: new to be created
InnoDB: Setting log file c:\iblogs\ib_logfile1 size to 31457280
InnoDB: Log file c:\iblogs\ib_logfile2 did not exist: new to be created
InnoDB: Setting log file c:\iblogs\ib_logfile2 size to 31457280
InnoDB: Doublewrite buffer not found: creating new
InnoDB: Doublewrite buffer created
InnoDB: creating foreign key constraint system tables
InnoDB: foreign key constraint system tables created
011024 10:58:25 InnoDB: Started
```

When the server finishes its startup sequence, you should see something like this, which indicates that the server is ready to service client connections:

```
mysqld: ready for connections
Version: '5.6.12' socket: '' port: 3306
```

The server continues to write to the console any further diagnostic output it produces. You can open a new console window in which to run client programs.

If you omit the `--console` [417] option, the server writes diagnostic output to the error log in the data directory (`C:\Program Files\MySQL\MySQL Server 5.6\data` by default). The error log is the file with the `.err` extension, and may be set using the `--log-error` [427] option.

**Note**

The accounts that are listed in the MySQL grant tables initially have no passwords. After starting the server, you should set up passwords for them using the instructions in [Section 2.10, “Postinstallation Setup and Testing”](#).

### 2.3.4.5. Starting MySQL from the Windows Command Line

The MySQL server can be started manually from the command line. This can be done on any version of Windows.

**Note**

The [MySQL Notifier for Microsoft Windows](#) GUI can also be used to start/stop/restart the MySQL server.

To start the `mysqld` server from the command line, you should start a console window (or “DOS window”) and enter this command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld"
```

The path to `mysqld` may vary depending on the install location of MySQL on your system.

You can stop the MySQL server by executing this command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqladmin" -u root shutdown
```

**Note**

If the MySQL `root` user account has a password, you need to invoke `mysqladmin` with the `-p` option and supply the password when prompted.

This command invokes the MySQL administrative utility `mysqladmin` to connect to the server and tell it to shut down. The command connects as the MySQL `root` user, which is the default administrative account in the MySQL grant system. Note that users in the MySQL grant system are wholly independent from any login users under Windows.

If `mysqld` doesn't start, check the error log to see whether the server wrote any messages there to indicate the cause of the problem. By default, the error log is located in the `C:\Program Files\MySQL\MySQL Server 5.6\data` directory. It is the file with a suffix of `.err`, or may be specified by passing in the `--log-error` [427] option. Alternatively, you can try to start the server as `mysqld --console`; in this case, you may get some useful information on the screen that may help solve the problem.

The last option is to start `mysqld` with the `--standalone` [440] and `--debug` [418] options. In this case, `mysqld` writes a log file `C:\mysqld.trace` that should contain the reason why `mysqld` doesn't start. See [MySQL Internals: Porting to Other Systems](#).

Use `mysqld --verbose --help` to display all the options that `mysqld` supports.

### 2.3.4.6. MySQL Notifier for Microsoft Windows

The MySQL Notifier for Microsoft Windows is a tool that allows you to monitor and adjust the status of your local MySQL Server instances through an indicator that resides in the system tray. The MySQL Notifier for Microsoft Windows also gives quick access to several MySQL GUI tools (such as MySQL Workbench) through its context menu.

The MySQL Notifier for Microsoft Windows is installed by MySQL Installer, and (by default) will start-up when Microsoft Windows is started.



#### Note

To install, download and execute the [MySQL Installer](#), be sure the MySQL Notifier for Microsoft Windows product is selected, then proceed with the installation. See the [MySQL Installer manual](#) for additional details.

For release notes detailing the changes in each release of MySQL Notifier for Microsoft Windows, see [MySQL Notifier Release Notes](#).

Visit the [MySQL Notifier forum](#) for additional MySQL Notifier for Microsoft Windows help and support.

Features include:

- Start, Stop, and Restart instances of the MySQL Server.
- Automatically detects (and adds) new MySQL Server services. These are listed under [Manage Services](#), and may also be configured.
- The Tray icon changes, depending on the status. It's green if all monitored MySQL Server instances are running, or red if at least one service is stopped.
- Links to other applications like MySQL Workbench, Windows Installer, and the MySQL Utilities. For example, choosing [Configure Instance](#) will load the MySQL Workbench Server Administration window for that particular instance.
- If MySQL Workbench is also installed, then the [Configure Instance](#) and [SQL Editor](#) options are available.


The MySQL Notifier for Microsoft Windows resides in the system tray and provides visual status information for your MySQL Server instances. A green icon is displayed at the top left corner of the tray icon if the current MySQL Server is running, or a red icon if the service is stopped.

The MySQL Notifier for Microsoft Windows automatically adds discovered MySQL Services on the local machine, and each service is saved and configurable. By default, the **Automatically add new services that match this pattern** option is enabled and set to `.*mysql\d.*`. Related **Notifications Options** include being notified when new services are either discovered or experience status changes, and are also enabled by default. And uninstalling a service will also remove the service from the MySQL Notifier for Microsoft Windows.

Clicking the system tray icon will reveal several options, as seen in the screenshots below:

The Service Instance menu is the main MySQL Notifier for Microsoft Windows window, and allows you to Stop, Start, and Restart the MySQL Server.

**Figure 2.22. MySQL Notifier for Microsoft Windows Service Instance menu**



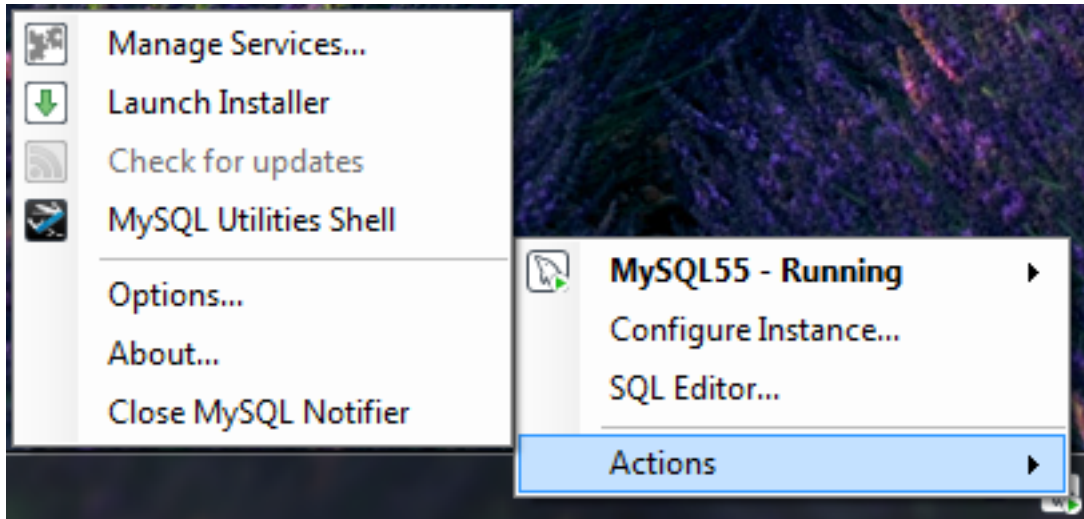
The Actions menu actions include several links to external applications (if they are installed).



**Note**

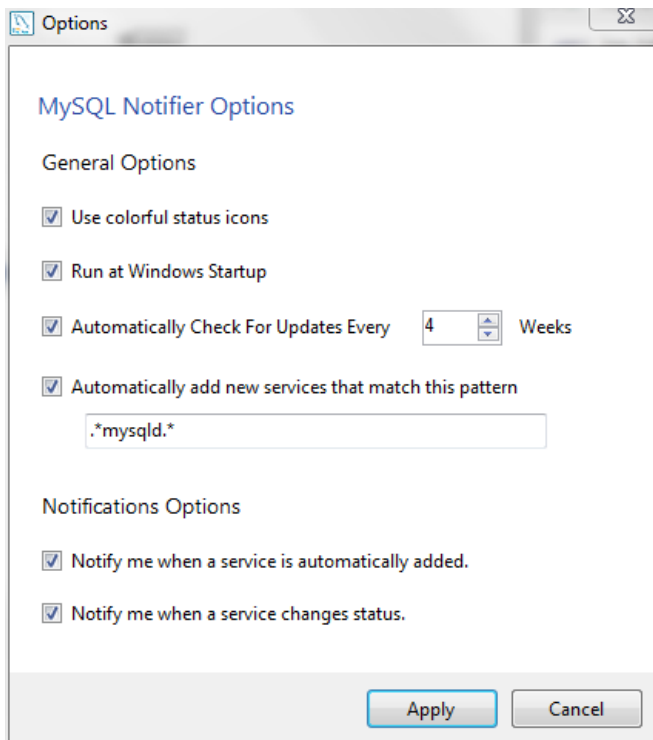
The main menu will not show the Actions menu when there are no services being monitored by MySQL Notifier for Microsoft Windows.

**Figure 2.23. MySQL Notifier for Microsoft Windows Actions menu**

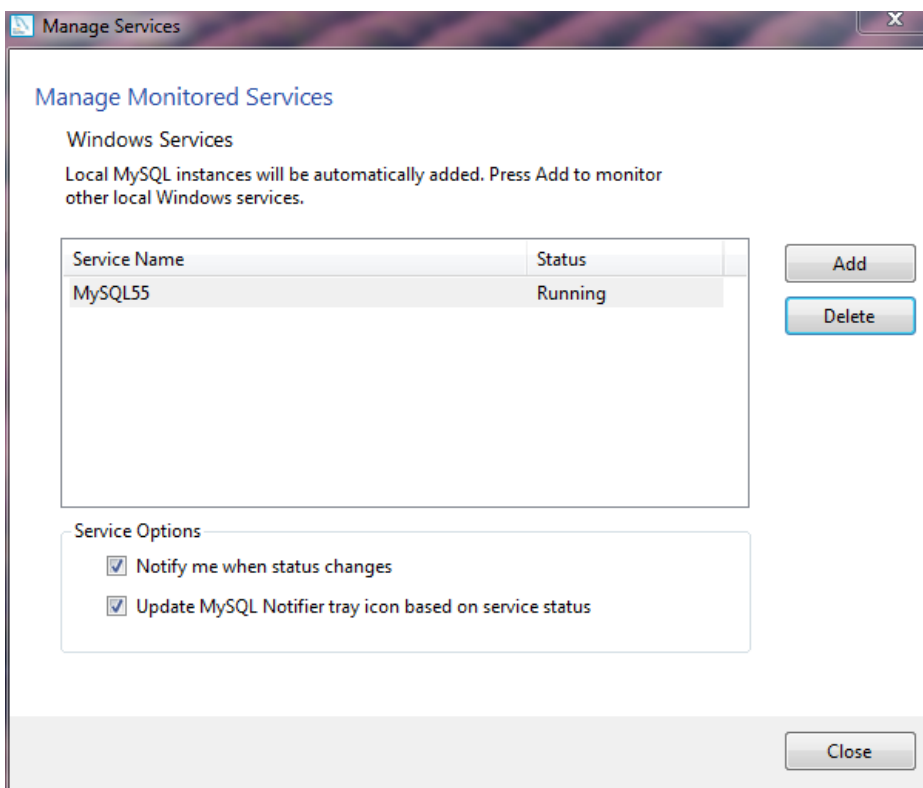


The Actions, Options menu configures MySQL Notifier for Microsoft Windows and includes options to:

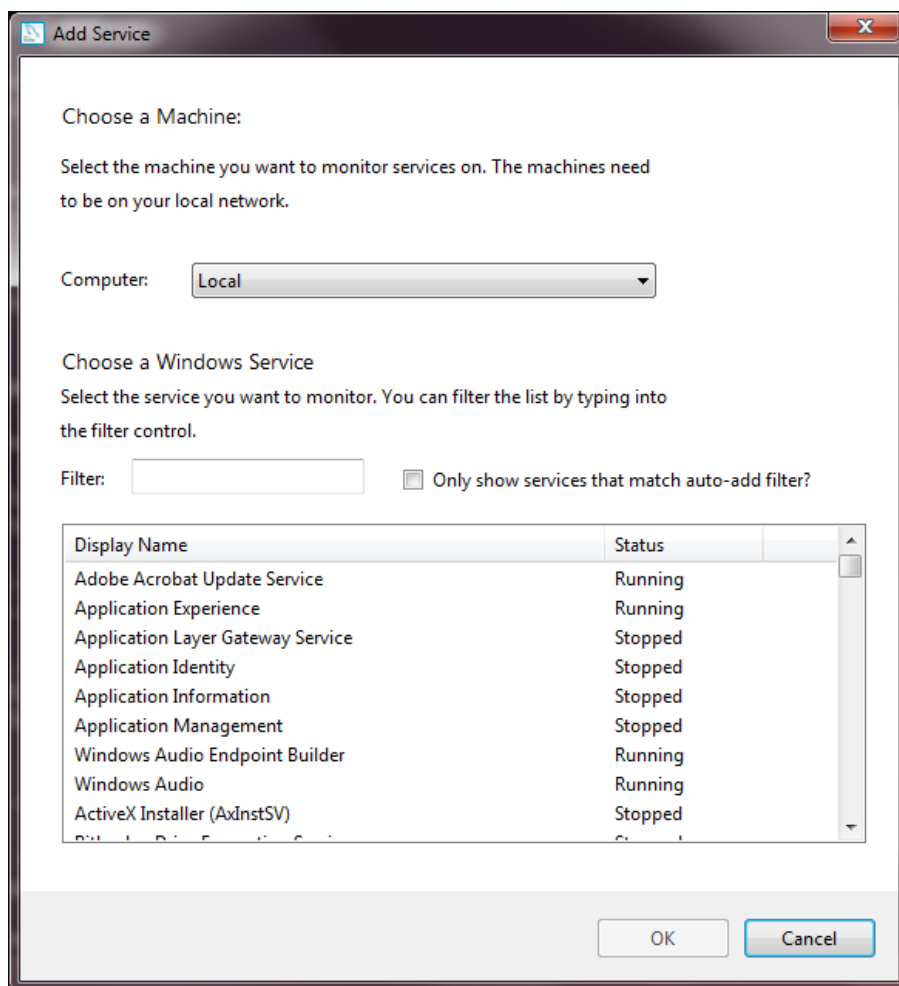
- **Use colorful status icons:** Enables a colorful style of icons for the tray of the MySQL Notifier for Microsoft Windows.
- **Run at Windows Startup:** Allows the application to be loaded when Microsoft Windows starts.
- **Automatically Check For Updates Every # Weeks:** Checks for a new version of MySQL Notifier for Microsoft Windows, and runs this check every # weeks.
- **Automatically add new services that match this pattern:** A custom regular expression that filters the monitored services, and also filters the list of the Microsoft Windows services for the **Add New Service** dialog.
- **Notify me when a service is automatically added:** Will display a notification from the taskbar when a newly discovered service is added to the monitored services list.
- **Notify me when a service changes status:** Will display a notification from the taskbar when a monitored service changes its status.

**Figure 2.24. MySQL Notifier for Microsoft Windows Options menu**

The **Actions**, **Manage Services** menu allows you to configure the monitored services.

**Figure 2.25. MySQL Notifier for Microsoft Windows Service Instance menu**

Adding a service (after clicking **Add** in the **Manage Services** window) allows you to select a running Microsoft Windows service, and configure MySQL Notifier for Microsoft Windows to monitor it. Add a new service by clicking service name from the list, then **OK** to accept. Multiple services may be selected.

**Figure 2.26. MySQL Notifier for Microsoft Windows Adding a new services**

### 2.3.4.7. Customizing the PATH for MySQL Tools

To make it easier to invoke MySQL programs, you can add the path name of the MySQL `bin` directory to your Windows system `PATH` environment variable:

- On the Windows desktop, right-click the My Computer icon, and select Properties.
- Next select the Advanced tab from the System Properties menu that appears, and click the Environment Variables button.
- Under **System Variables**, select Path, and then click the Edit button. The Edit System Variable dialogue should appear.
- Place your cursor at the end of the text appearing in the space marked **Variable Value**. (Use the **End** key to ensure that your cursor is positioned at the very end of the text in this space.) Then enter the complete path name of your MySQL `bin` directory (for example, `C:\Program Files\MySQL\MySQL Server 5.6\bin`)



#### Note

There must be a semicolon separating this path from any values present in this field.

Dismiss this dialogue, and each dialogue in turn, by clicking **OK** until all of the dialogues that were opened have been dismissed. You should now be able to invoke any MySQL executable program by typing its name at the DOS prompt from any directory on the system, without having to supply



the path. This includes the servers, the `mysql` client, and all MySQL command-line utilities such as `mysqladmin` and `mysqldump`.

You should not add the MySQL `bin` directory to your Windows `PATH` if you are running multiple MySQL servers on the same machine.



#### Warning

You must exercise great care when editing your system `PATH` by hand; accidental deletion or modification of any portion of the existing `PATH` value can leave you with a malfunctioning or even unusable system.

### 2.3.4.8. Starting MySQL as a Windows Service

On Windows, the recommended way to run MySQL is to install it as a Windows service, so that MySQL starts and stops automatically when Windows starts and stops. A MySQL server installed as a service can also be controlled from the command line using `NET` commands, or with the graphical `Services` utility. Generally, to install MySQL as a Windows service you should be logged in using an account that has administrator rights.



#### Note

The [MySQL Notifier for Microsoft Windows](#) GUI can also be used to monitor the status of the MySQL service.

The `Services` utility (the Windows [Service Control Manager](#)) can be found in the Windows Control Panel (under Administrative Tools on Windows 2000, XP, Vista, and Server 2003). To avoid conflicts, it is advisable to close the `Services` utility while performing server installation or removal operations from the command line.

Before installing MySQL as a Windows service, you should first stop the current server if it is running by using the following command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqladmin"  
-u root shutdown
```



#### Note

If the MySQL `root` user account has a password, you need to invoke `mysqladmin` with the `-p` option and supply the password when prompted.

This command invokes the MySQL administrative utility `mysqladmin` to connect to the server and tell it to shut down. The command connects as the MySQL `root` user, which is the default administrative account in the MySQL grant system. Note that users in the MySQL grant system are wholly independent from any login users under Windows.

Install the server as a service using this command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld" --install
```

The service-installation command does not start the server. Instructions for that are given later in this section.

To make it easier to invoke MySQL programs, you can add the path name of the MySQL `bin` directory to your Windows system `PATH` environment variable:

- On the Windows desktop, right-click the My Computer icon, and select Properties.
- Next select the Advanced tab from the [System Properties](#) menu that appears, and click the [Environment Variables](#) button.

- Under **System Variables**, select Path, and then click the **Edit** button. The **Edit System Variable** dialogue should appear.
- Place your cursor at the end of the text appearing in the space marked **Variable Value**. (Use the **End** key to ensure that your cursor is positioned at the very end of the text in this space.) Then enter the complete path name of your MySQL `bin` directory (for example, `C:\Program Files\MySQL\MySQL Server 5.6\bin`), Note that there should be a semicolon separating this path from any values present in this field. Dismiss this dialogue, and each dialogue in turn, by clicking **OK** until all of the dialogues that were opened have been dismissed. You should now be able to invoke any MySQL executable program by typing its name at the DOS prompt from any directory on the system, without having to supply the path. This includes the servers, the `mysql` client, and all MySQL command-line utilities such as `mysqladmin` and `mysqldump`.

You should not add the MySQL `bin` directory to your Windows `PATH` if you are running multiple MySQL servers on the same machine.



### Warning

You must exercise great care when editing your system `PATH` by hand; accidental deletion or modification of any portion of the existing `PATH` value can leave you with a malfunctioning or even unusable system.

The following additional arguments can be used when installing the service:

- You can specify a service name immediately following the `--install` option. The default service name is `MySQL`.
- If a service name is given, it can be followed by a single option. By convention, this should be `--defaults-file=file_name [219]` to specify the name of an option file from which the server should read options when it starts.

The use of a single option other than `--defaults-file [219]` is possible but discouraged. `--defaults-file [219]` is more flexible because it enables you to specify multiple startup options for the server by placing them in the named option file.

- You can also specify a `--local-service` option following the service name. This causes the server to run using the `LocalService` Windows account that has limited system privileges. This account is available only for Windows XP or newer. If both `--defaults-file [219]` and `--local-service` are given following the service name, they can be in any order.

For a MySQL server that is installed as a Windows service, the following rules determine the service name and option files that the server uses:

- If the service-installation command specifies no service name or the default service name (`MySQL`) following the `--install` option, the server uses the a service name of `MySQL` and reads options from the `[mysqld]` group in the standard option files.
- If the service-installation command specifies a service name other than `MySQL` following the `--install` option, the server uses that service name. It reads options from the `[mysqld]` group and the group that has the same name as the service in the standard option files. This enables you to use the `[mysqld]` group for options that should be used by all MySQL services, and an option group with the service name for use by the server installed with that service name.
- If the service-installation command specifies a `--defaults-file [219]` option after the service name, the server reads options the same way as described in the previous item, except that it reads options only from the the named file and ignores the standard option files.

As a more complex example, consider the following command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld"  
--install MySQL --defaults-file=C:\my-opts.cnf
```

Here, the default service name (`MySQL`) is given after the `--install` option. If no `--defaults-file` [219] option had been given, this command would have the effect of causing the server to read the `[mysqld]` group from the standard option files. However, because the `--defaults-file` [219] option is present, the server reads options from the `[mysqld]` option group, and only from the named file.

You can also specify options as Start parameters in the Windows `Services` utility before you start the MySQL service.

Once a MySQL server has been installed as a service, Windows starts the service automatically whenever Windows starts. The service also can be started immediately from the `Services` utility, or by using a `NET START MySQL` command. The `NET` command is not case sensitive.

When run as a service, `mysqld` has no access to a console window, so no messages can be seen there. If `mysqld` does not start, check the error log to see whether the server wrote any messages there to indicate the cause of the problem. The error log is located in the MySQL data directory (for example, `C:\Program Files\MySQL\MySQL Server 5.6\data`). It is the file with a suffix of `.err`.

When a MySQL server has been installed as a service, and the service is running, Windows stops the service automatically when Windows shuts down. The server also can be stopped manually by using the `Services` utility, the `NET STOP MySQL` command, or the `mysqladmin shutdown` command.

You also have the choice of installing the server as a manual service if you do not wish for the service to be started automatically during the boot process. To do this, use the `--install-manual` option rather than the `--install` option:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld" --install-manual
```

To remove a server that is installed as a service, first stop it if it is running by executing `NET STOP MySQL`. Then use the `--remove` [437] option to remove it:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld" --remove
```

If `mysqld` is not running as a service, you can start it from the command line. For instructions, see [Section 2.3.4.5, "Starting MySQL from the Windows Command Line"](#).

If you encounter difficulties during installation, see [Section 2.3.5, "Troubleshooting a Microsoft Windows MySQL Server Installation"](#).

### 2.3.4.9. Testing The MySQL Installation

You can test whether the MySQL server is working by executing any of the following commands:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqlshow"  
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqlshow" -u root mysql  
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqladmin" version status proc  
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysql" test
```

If `mysqld` is slow to respond to TCP/IP connections from client programs, there is probably a problem with your DNS. In this case, start `mysqld` with the `--skip-name-resolve` [440] option and use only `localhost` and IP addresses in the `Host` column of the MySQL grant tables.

You can force a MySQL client to use a named-pipe connection rather than TCP/IP by specifying the `--pipe` [211] or `--protocol=PIPE` [211] option, or by specifying `.` (period) as the host name. Use the `--socket` [211] option to specify the name of the pipe if you do not want to use the default pipe name.

Note that if you have set a password for the `root` account, deleted the anonymous account, or created a new user account, then you must use the appropriate `-u` and `-p` options with the commands shown above to connect with the MySQL Server. See [Section 4.2.2, "Connecting to the MySQL Server"](#).

For more information about `mysqlshow`, see [Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”](#).

## 2.3.5. Troubleshooting a Microsoft Windows MySQL Server Installation

When installing and running MySQL for the first time, you may encounter certain errors that prevent the MySQL server from starting. This section helps you diagnose and correct some of these errors.

Your first resource when troubleshooting server issues is the [error log](#). The MySQL server uses the error log to record information relevant to the error that prevents the server from starting. The error log is located in the [data directory](#) specified in your `my.ini` file. The default data directory location is `C:\Program Files\MySQL\MySQL Server 5.6\data`, or `C:\ProgramData\Mysql` on Windows 7 and Windows Server 2008. The `C:\ProgramData` directory is hidden by default. You need to change your folder options to see the directory and contents. For more information on the error log and understanding the content, see [Section 5.2.2, “The Error Log”](#).

For information regarding possible errors, also consult the console messages displayed when the MySQL service is starting. Use the `NET START MySQL` command from the command line after installing `mysqld` as a service to see any error messages regarding the starting of the MySQL server as a service. See [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

The following examples show other common error messages you might encounter when installing MySQL and starting the server for the first time:

- If the MySQL server cannot find the `mysql` privileges database or other critical files, it displays these messages:

```
System error 1067 has occurred.
Fatal error: Can't open and lock privilege tables:
Table 'mysql.user' doesn't exist
```

These messages often occur when the MySQL base or data directories are installed in different locations than the default locations (`C:\Program Files\MySQL\MySQL Server 5.6` and `C:\Program Files\MySQL\MySQL Server 5.6\data`, respectively).

This situation can occur when MySQL is upgraded and installed to a new location, but the configuration file is not updated to reflect the new location. In addition, old and new configuration files might conflict. Be sure to delete or rename any old configuration files when upgrading MySQL.

If you have installed MySQL to a directory other than `C:\Program Files\MySQL\MySQL Server 5.6`, ensure that the MySQL server is aware of this through the use of a configuration (`my.ini`) file. Put the `my.ini` file in your Windows directory, typically `C:\WINDOWS`. To determine its exact location from the value of the `WINDIR` environment variable, issue the following command from the command prompt:

```
C:\> echo %WINDIR%
```

You can create or modify an option file with any text editor, such as Notepad. For example, if MySQL is installed in `E:\mysql` and the data directory is `D:\MySQLdata`, you can create the option file and set up a `[mysqld]` section to specify values for the `basedir` and `datadir` options:

```
[mysqld]
# set basedir to your installation path
basedir=E:/mysql
# set datadir to the location of your data directory
datadir=D:/MySQLdata
```

Note that Windows path names are specified in option files using (forward) slashes rather than backslashes. If you do use backslashes, double them:

```
[mysqld]
# set basedir to your installation path
basedir=C:\\Program Files\\MySQL\\MySQL Server 5.6
```

```
# set datadir to the location of your data directory
datadir=D:\\MySQLdata
```

The rules for use of backslash in option file values are given in [Section 4.2.3.3, “Using Option Files”](#).

If you change the `datadir` value in your MySQL configuration file, you must move the contents of the existing MySQL data directory before restarting the MySQL server.

See [Section 2.3.4.2, “Creating an Option File”](#).

- If you reinstall or upgrade MySQL without first stopping and removing the existing MySQL service and install MySQL using the MySQL Installer, you might see this error:

```
Error: Cannot create Windows service for MySQL. Error: 0
```

This occurs when the Configuration Wizard tries to install the service and finds an existing service with the same name.

One solution to this problem is to choose a service name other than `mysql` when using the configuration wizard. This enables the new service to be installed correctly, but leaves the outdated service in place. Although this is harmless, it is best to remove old services that are no longer in use.

To permanently remove the old `mysql` service, execute the following command as a user with administrative privileges, on the command line:

```
C:\> sc delete mysql
[SC] DeleteService SUCCESS
```

If the `sc` utility is not available for your version of Windows, download the `delsrv` utility from <http://www.microsoft.com/windows2000/techinfo/reskit/tools/existing/delsrv-o.asp> and use the `delsrv mysql` syntax.

## 2.3.6. Upgrading MySQL on Windows

This section lists some of the steps to take when upgrading MySQL on Windows.

1. Review [Section 2.11.1, “Upgrading MySQL”](#), for additional information on upgrading MySQL that is not specific to Windows.
2. Always back up your current MySQL installation before performing an upgrade. See [Section 7.2, “Database Backup Methods”](#).
3. Download the latest Windows distribution of MySQL from <http://dev.mysql.com/downloads/>.
4. Before upgrading MySQL, stop the server. If the server is installed as a service, stop the service with the following command from the command prompt:

```
C:\> NET STOP MySQL
```

If you are not running the MySQL server as a service, use `mysqladmin` to stop it. For example, before upgrading from MySQL 5.5 to 5.6, use `mysqladmin` from MySQL 5.5 as follows:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.5\bin\mysqladmin" -u root shutdown
```



### Note

If the MySQL `root` user account has a password, invoke `mysqladmin` with the `-p` option and supply the password when prompted.

5. Before upgrading to MySQL 5.6 from a version previous to 4.1.5, or from a version of MySQL installed from a Zip archive to a version of MySQL installed with the MySQL Installation Wizard, you must first manually remove the previous installation and MySQL service (if the server is installed as a service).

To remove the MySQL service, use the following command:

```
C:\> C:\mysql\bin\mysqld --remove
```

**If you do not remove the existing service, the MySQL Installation Wizard may fail to properly install the new MySQL service.**

6. If you are using the MySQL Installer, start it as described in [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#).
7. If you are installing MySQL from a Zip archive, extract the archive. You may either overwrite your existing MySQL installation (usually located at `C:\mysql`), or install it into a different directory, such as `C:\mysql5`. Overwriting the existing installation is recommended.
8. If you were running MySQL as a Windows service and you had to remove the service earlier in this procedure, reinstall the service. (See [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).)
9. Restart the server. For example, use `NET START MySQL` if you run MySQL as a service, or invoke `mysqld` directly otherwise.
10. As Administrator, run `mysql_upgrade` to check your tables, attempt to repair them if necessary, and update your grant tables if they have changed so that you can take advantage of any new capabilities. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).
11. If you encounter errors, see [Section 2.3.5, “Troubleshooting a Microsoft Windows MySQL Server Installation”](#).

## 2.3.7. Windows Postinstallation Procedures

GUI tools exist that perform most of the tasks described below, including:

- [MySQL Installer](#): Used to install and upgrade MySQL products.
- [MySQL Workbench](#): Manages the MySQL server and edits SQL queries.
- [MySQL Notifier](#): Starts, stops, or restarts the MySQL server, and monitors its status.
- [MySQL for Excel](#): Edits MySQL data with Microsoft Excel.

On Windows, you need not create the data directory and the grant tables. MySQL Windows distributions include the grant tables with a set of preinitialized accounts in the `mysql` database under the data directory. Regarding passwords, if you installed MySQL using the MySQL Installer, you may have already assigned passwords to the accounts. (See [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#).) Otherwise, use the password-assignment procedure given in [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

Before setting up passwords, you might want to try running some client programs to make sure that you can connect to the server and that it is operating properly. Make sure that the server is running (see [Section 2.3.4.4, “Starting the Server for the First Time”](#)), and then issue the following commands to verify that you can retrieve information from the server. You may need to specify directory different from `C:\mysql\bin` on the command line. If you used the MySQL Installer, the default directory is `C:\Program Files\MySQL\MySQL Server 5.6`, and the `mysql` and `mysqlshow` client programs are in `C:\Program Files\MySQL\MySQL Server 5.6\bin`. See [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#), for more information.

Use `mysqlshow` to see what databases exist:

```
C:\> C:\mysql\bin\mysqlshow
+-----+
|          |
| Databases |
|          |
```



```
+-----+
| information_schema |
| mysql              |
| test               |
+-----+
```

The list of installed databases may vary, but will always include the minimum of `mysql` and `information_schema`. In most cases, the `test` database will also be installed automatically.

The preceding command (and commands for other MySQL programs such as `mysql`) may not work if the correct MySQL account does not exist. For example, the program may fail with an error, or you may not be able to view all databases. If you installed using MySQL Installer, then the `root` user will have been created automatically with the password you supplied. In this case, you should use the `-u root` and `-p` options. (You will also need to use the `-u root` and `-p` options if you have already secured the initial MySQL accounts.) With `-p`, you will be prompted for the `root` password. For example:

```
C:\> C:\mysql\bin\mysqlshow -u root -p
Enter password: (enter root password here)
+-----+
|      Databases      |
+-----+
| information_schema  |
| mysql              |
| test               |
+-----+
```

If you specify a database name, `mysqlshow` displays a list of the tables within the database:

```
C:\> C:\mysql\bin\mysqlshow mysql
Database: mysql
+-----+
|      Tables        |
+-----+
| columns_priv      |
| db                 |
| event              |
| func               |
| help_category     |
| help_keyword      |
| help_relation     |
| help_topic        |
| host               |
| plugin             |
| proc               |
| procs_priv        |
| servers            |
| tables_priv       |
| time_zone          |
| time_zone_leap_second |
| time_zone_name    |
| time_zone_transition |
| time_zone_transition_type |
| user               |
+-----+
```

Use the `mysql` program to select information from a table in the `mysql` database:

```
C:\> C:\mysql\bin\mysql -e "SELECT Host,Db,User FROM mysql.db"
+-----+-----+-----+
| host | db   | user |
+-----+-----+-----+
| %    | test |      |
| %    | test_% |    |
+-----+-----+-----+
```

For more information about `mysqlshow` and `mysql`, see [Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”](#), and [Section 4.5.1, “mysql — The MySQL Command-Line Tool”](#).



If you are running a version of Windows that supports services, you can set up the MySQL server to run automatically when Windows starts. See [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

## 2.4. Installing MySQL on Mac OS X

MySQL for Mac OS X is available in a number of different forms:

- Native Package Installer format, which uses the native Mac OS X installer to walk you through the installation of MySQL. For more information, see [Section 2.4.2, “Installing MySQL on Mac OS X Using Native Packages”](#). You can use the package installer with Mac OS X 10.3 and later, and the package is available for both PowerPC and Intel architectures, and 32-bit and 64-bit architectures. There is no Universal Binary available using the package installation method. The user you use to perform the installation must have administrator privileges.
- Tar package format, which uses a file packaged using the Unix `tar` and `gzip` commands. To use this method, you will need to open a `Terminal` window. You do not need administrator privileges using this method, as you can install the MySQL server anywhere using this method. For more information on using this method, you can use the generic instructions for using a tarball, [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#). You can use the package installer with Mac OS X 10.3 and later, and available for both PowerPC and Intel architectures, and both 32-bit and 64-bit architectures. A Universal Binary, incorporating both Power PC and Intel architectures and 32-bit and 64-bit binaries is available.

In addition to the core installation, the Package Installer also includes [Section 2.4.3, “Installing the MySQL Startup Item”](#) and [Section 2.4.4, “Installing and Using the MySQL Preference Pane”](#), both of which simplify the management of your installation.

- Mac OS X server includes a version of MySQL as standard. If you want to use a more recent version than that supplied with the Mac OS X server release, you can make use of the package or tar formats. For more information on using the MySQL bundled with Mac OS X, see [Section 2.4.5, “Using the Bundled MySQL on Mac OS X Server”](#).

For additional information on using MySQL on Mac OS X, see [Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”](#).

### 2.4.1. General Notes on Installing MySQL on Mac OS X

You should keep the following issues and notes in mind:

- The default location for the MySQL Unix socket is different on Mac OS X and Mac OS X Server depending on the installation type you chose. The following table shows the default locations by installation type.

**Table 2.5. MySQL Unix Socket Locations on Mac OS X by Installation Type**

Installation Type	Socket Location
Package Installer from MySQL	<code>/tmp/mysql.sock</code>
Tarball from MySQL	<code>/tmp/mysql.sock</code>
MySQL Bundled with Mac OS X Server	<code>/var/mysql/mysql.sock</code>

To prevent issues, you should either change the configuration of the socket used within your application (for example, changing `php.ini`), or you should configure the socket location using a MySQL configuration file and the `socket [442]` option. For more information, see [Section 5.1.3, “Server Command Options”](#).

- You may need (or want) to create a specific `mysql` user to own the MySQL directory and data. On Mac OS X 10.4 and lower you can do this by using the `Netinfo Manager` application, located within the `Utilities` folder within the `Applications` folder. On Mac OS X 10.5 and later you

can do this through the [Directory Utility](#). From Mac OS X 10.5 and later (including Mac OS X Server 10.5) the `mysql` should already exist. For use in single user mode, an entry for `_mysql` (note the underscore prefix) should already exist within the system `/etc/passwd` file.

- Due to a bug in the Mac OS X package installer, you may see this error message in the destination disk selection dialog:

```
You cannot install this software on this disk. (null)
```

If this error occurs, click the [Go Back](#) button once to return to the previous screen. Then click [Continue](#) to advance to the destination disk selection again, and you should be able to choose the destination disk correctly. We have reported this bug to Apple and it is investigating this problem.

- If you get an “insecure startup item disabled” error when MySQL launches, use the following procedure. Adjust the pathnames appropriately for your system.

1. Modify the `mysql.script` using this command (enter it on a single line):

```
shell> sudo /Applications/TextEdit.app/Contents/MacOS/TextEdit  
/usr/local/mysql/support-files/mysql.server
```

2. Locate the option file that defines the `basedir` value and modify it to contain these lines:

```
basedir=/usr/local/mysql  
datadir=/usr/local/mysql/data
```

In the `/Library/StartupItems/MySQLCOM/` directory, make the following group ID changes from `staff` to `wheel`:

```
shell> sudo chgrp wheel MySQLCOM StartupParameters.plist
```

3. Start the server from System Preferences or Terminal.app.

- Because the MySQL package installer installs the MySQL contents into a version and platform specific directory, you can use this to upgrade and migrate your database between versions. You will need to either copy the `data` directory from the old version to the new version, or alternatively specify an alternative `datadir` value to set location of the data directory.
- You might want to add aliases to your shell's resource file to make it easier to access commonly used programs such as `mysql` and `mysqladmin` from the command line. The syntax for `bash` is:

```
alias mysql=/usr/local/mysql/bin/mysql  
alias mysqladmin=/usr/local/mysql/bin/mysqladmin
```

For `tcsh`, use:

```
alias mysql /usr/local/mysql/bin/mysql  
alias mysqladmin /usr/local/mysql/bin/mysqladmin
```

Even better, add `/usr/local/mysql/bin` to your `PATH` environment variable. You can do this by modifying the appropriate startup file for your shell. For more information, see [Section 4.2.1](#), “[Invoking MySQL Programs](#)”.

- After you have copied over the MySQL database files from the previous installation and have successfully started the new server, you should consider removing the old installation files to save disk space. Additionally, you should also remove older versions of the Package Receipt directories located in `/Library/Receipts/mysql-VERSION.pkg`.

## 2.4.2. Installing MySQL on Mac OS X Using Native Packages

You can install MySQL on Mac OS X 10.3.x (“Panther”) or newer using a Mac OS X binary package in PKG format instead of the binary tarball distribution. Please note that older versions of Mac OS X (for example, 10.1.x or 10.2.x) are *not* supported by this package.

The package is located inside a disk image (`.dmg`) file that you first need to mount by double-clicking its icon in the Finder. It should then mount the image and display its contents.



#### Note

Before proceeding with the installation, be sure to stop all running MySQL server instances by using either the MySQL Manager Application (on Mac OS X Server) or `mysqladmin shutdown` on the command line.

When installing from the package version, you should also install the MySQL Preference Pane, which will enable you to control the startup and execution of your MySQL server from System Preferences. For more information, see [Section 2.4.4, “Installing and Using the MySQL Preference Pane”](#).

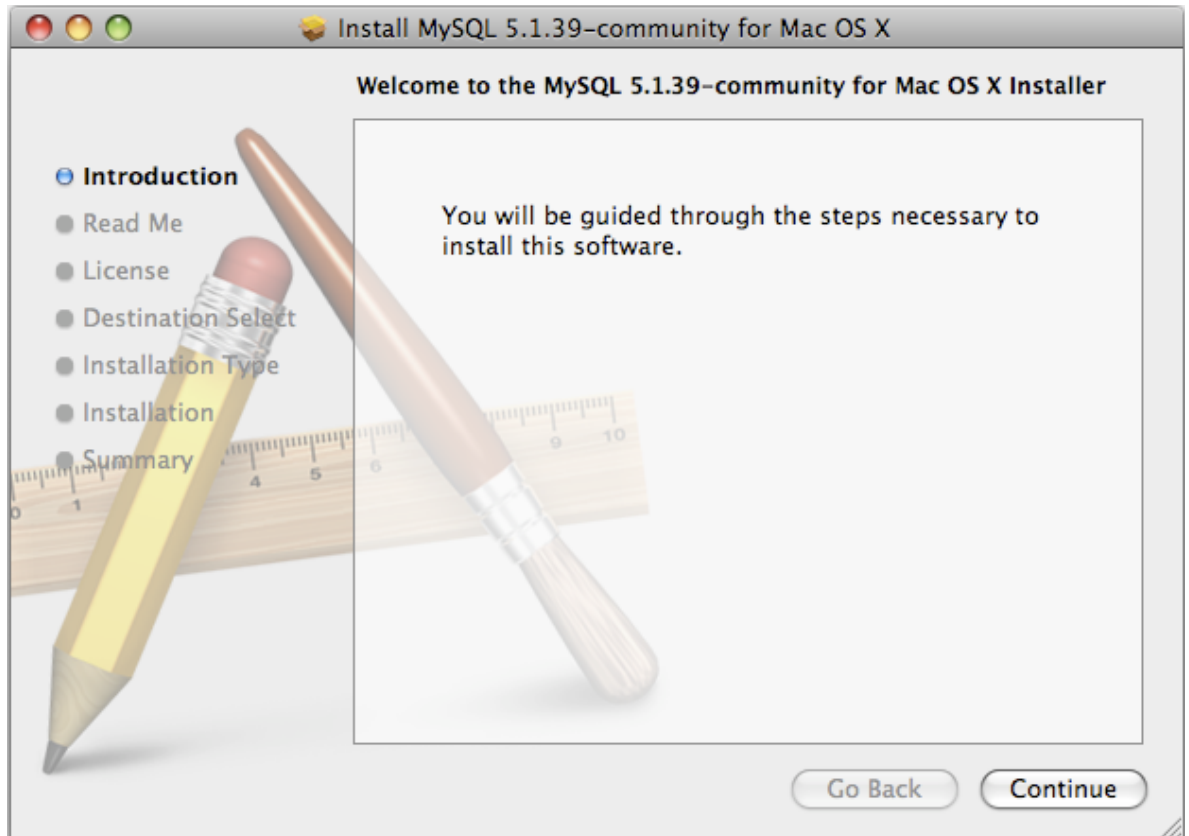
When installing using the package installer, the files are installed into a directory within `/usr/local` matching the name of the installation version and platform. For example, the installer file `mysql-5.1.39-osx10.5-x86_64.pkg` installs MySQL into `/usr/local/mysql-5.1.39-osx10.5-x86_64`. The following table shows the layout of the installation directory.

**Table 2.6. MySQL Installation Layout on Mac OS X**

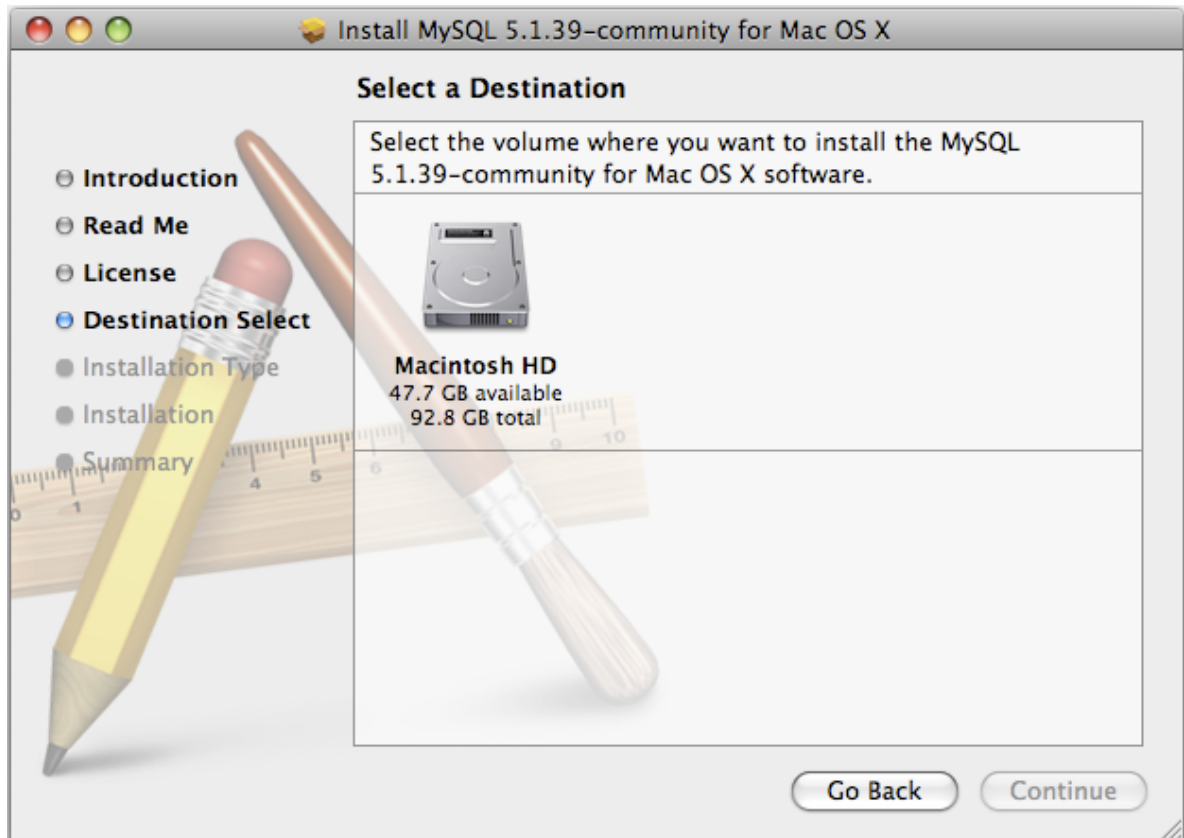
Directory	Contents of Directory
<code>bin</code>	Client programs and the <code>mysqld</code> server
<code>data</code>	Log files, databases
<code>docs</code>	Manual in Info format
<code>include</code>	Include (header) files
<code>lib</code>	Libraries
<code>man</code>	Unix manual pages
<code>mysql-test</code>	MySQL test suite
<code>scripts</code>	<code>mysql_install_db</code>
<code>share</code>	Miscellaneous support files, including error messages, sample configuration files, SQL for database installation
<code>sql-bench</code>	Benchmarks
<code>support-files</code>	Scripts and sample configuration files
<code>/tmp/mysql.sock</code>	Location of the MySQL Unix socket

During the package installer process, a symbolic link from `/usr/local/mysql` to the version/platform specific directory created during installation will be created automatically.

1. Download and open the MySQL package installer, which is provided on a disk image (`.dmg`) that includes the main MySQL installation package, the `MySQLStartupItem.pkg` installation package, and the `MySQL.prefPane`. Double-click the disk image to open it.
2. Double-click the MySQL installer package. It will be named according to the version of MySQL you have downloaded. For example, if you have downloaded MySQL 5.1.39, double-click `mysql-5.1.39-osx10.5-x86.pkg`.
3. You will be presented with the opening installer dialog. Click **Continue** to begin installation.



4. A copy of the installation instructions and other important information relevant to this installation are displayed. Click **Continue** .
5. If you have downloaded the community version of MySQL, you will be shown a copy of the relevant GNU General Public License. Click **Continue** .
6. Select the drive you want to use to install the MySQL Startup Item. The drive must have a valid, bootable, Mac OS X operating system installed. Click **Continue**.



7. You will be asked to confirm the details of the installation, including the space required for the installation. To change the drive on which the startup item is installed, click either **Go Back** or **Change Install Location...**. To install the startup item, click **Install**.
8. Once the installation has been completed successfully, you will be shown an **Install Succeeded** message.

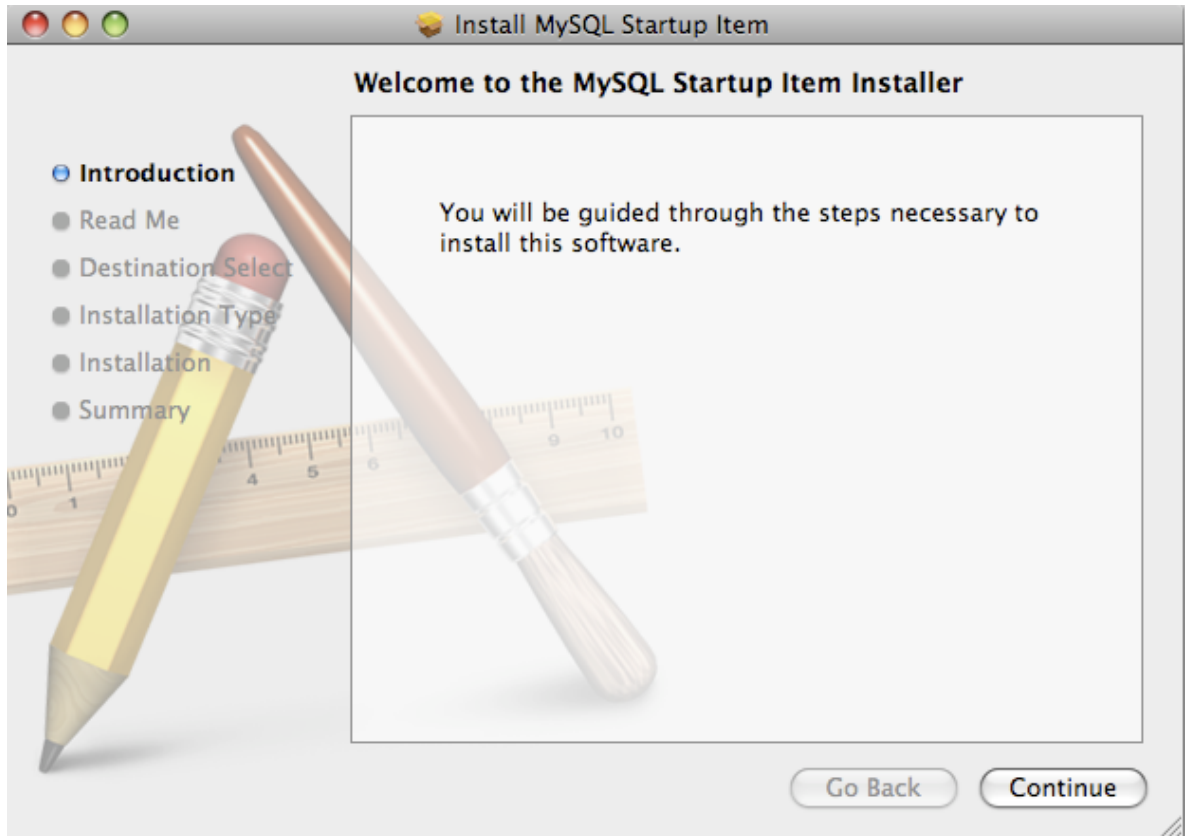
For convenience, you may also want to install the startup item and preference pane. See [Section 2.4.3, “Installing the MySQL Startup Item”](#), and [Section 2.4.4, “Installing and Using the MySQL Preference Pane”](#).

### 2.4.3. Installing the MySQL Startup Item

The MySQL Installation Package includes a startup item that can be used to automatically start and stop MySQL.

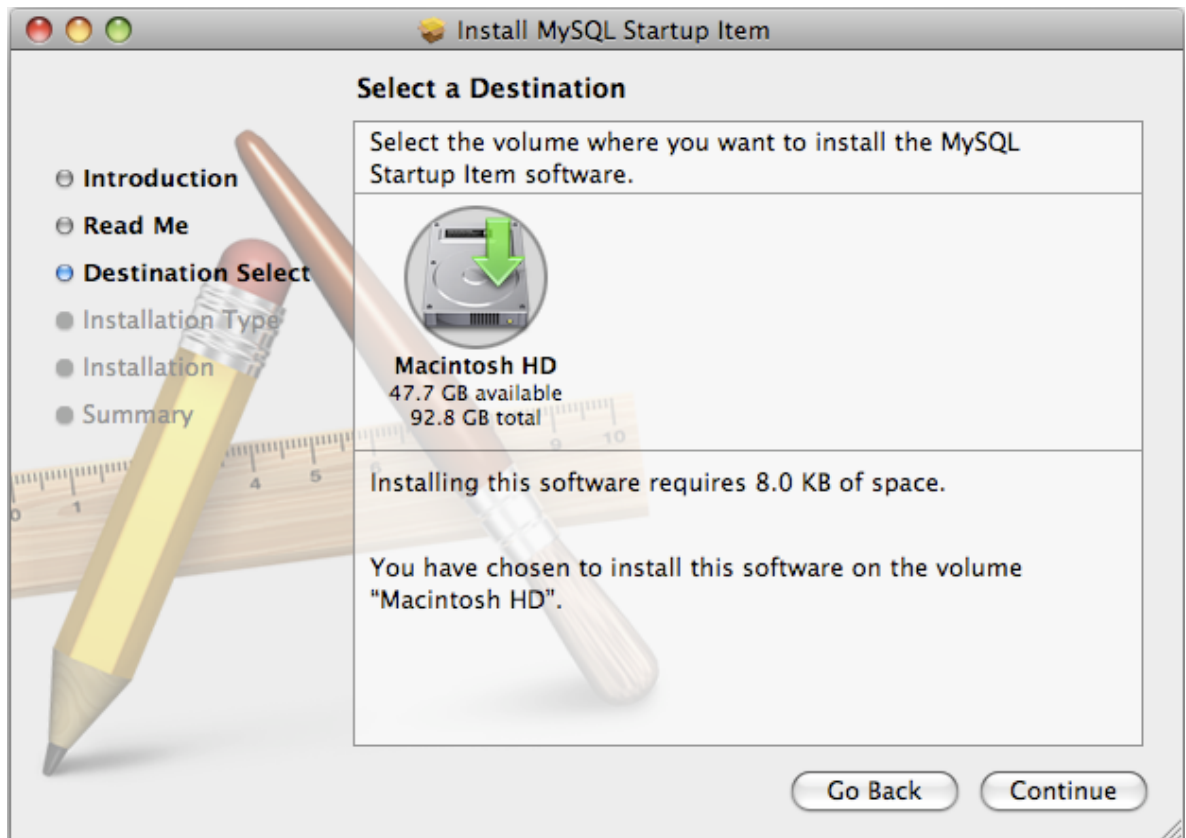
To install the MySQL Startup Item:

1. Download and open the MySQL package installer, which is provided on a disk image (`.dmg`) that includes the main MySQL installation package, the `MySQLStartupItem.pkg` installation package, and the `MySQL.prefPane`. Double-click the disk image to open it.
2. Double-click the `MySQLStartItem.pkg` file to start the installation process.
3. You will be presented with the **Install MySQL Startup Item** dialog.



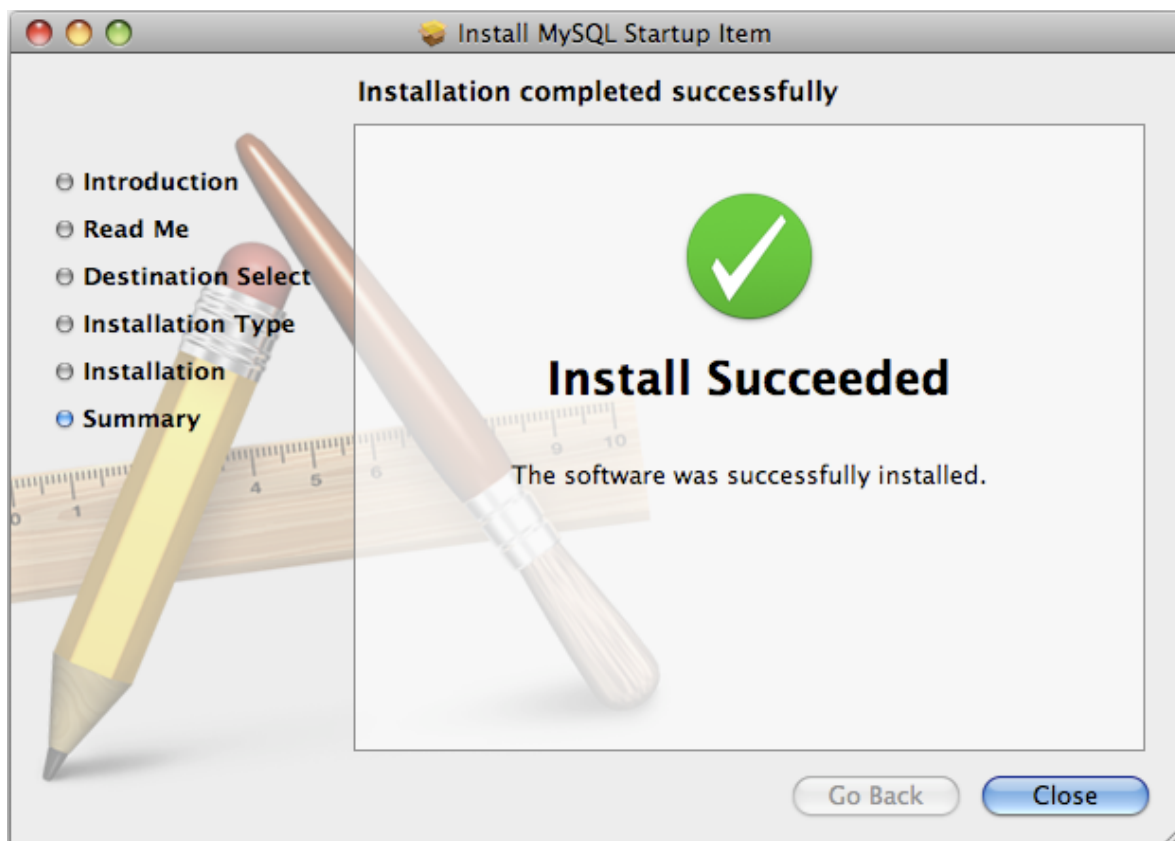
Click **Continue** to continue the installation process.

4. A copy of the installation instructions and other important information relevant to this installation are displayed. Click **Continue** .
5. Select the drive you want to use to install the MySQL Startup Item. The drive must have a valid, bootable, Mac OS X operating system installed. Click **Continue**.



6. You will be asked to confirm the details of the installation. To change the drive on which the startup item is installed, click either **Go Back** or **Change Install Location...**. To install the startup item, click **Install**.
7. Once the installation has been completed successfully, you will be shown an **Install Succeeded** message.





The Startup Item for MySQL is installed into `/Library/StartupItems/MySQLCOM`. The Startup Item installation adds a variable `MYSQLCOM=-YES-` to the system configuration file `/etc/hostconfig`. If you want to disable the automatic startup of MySQL, change this variable to `MYSQLCOM=-NO-`.

After the installation, you can start and stop MySQL by running the following commands in a terminal window. You must have administrator privileges to perform these tasks, and you may be prompted for your password.

If you have installed the Startup Item, use this command to start the server:

```
shell> sudo /Library/StartupItems/MySQLCOM/MySQLCOM start
```

If you have installed the Startup Item, use this command to stop the server:

```
shell> sudo /Library/StartupItems/MySQLCOM/MySQLCOM stop
```

## 2.4.4. Installing and Using the MySQL Preference Pane

The MySQL Package installer disk image also includes a custom MySQL Preference Pane that enables you to start, stop, and control automated startup during boot of your MySQL installation.

To install the MySQL Preference Pane:

1. Download and open the MySQL package installer package, which is provided on a disk image (`.dmg`) that includes the main MySQL installation package, the `MySQLStartupItem.pkg` installation package, and the `MySQL.prefPane`. Double-click the disk image to open it.
2. Double-click the `MySQL.prefPane`. The MySQL System Preferences will open.
3. If this is the first time you have installed the preference pane, you will be asked to confirm installation and whether you want to install the preference pane for all users, or only the current user. To install the preference pane for all users you will need administrator privileges. If necessary, you will be prompted for the username and password for a user with administrator privileges.


4. If you already have the MySQL Preference Pane installed, you will be asked to confirm whether you want to overwrite the existing MySQL Preference Pane.



#### Note

The MySQL Preference Pane only starts and stops MySQL installation installed from the MySQL package installation that have been installed in the default location.

Once the MySQL Preference Pane has been installed, you can control your MySQL server instance using the preference pane. To use the preference pane, open the **System Preferences...** from the Apple menu. Select the MySQL preference pane by clicking the MySQL logo within the **Other** section of the preference panes list.



The MySQL Preference Pane shows the current status of the MySQL server, showing **stopped** (in red) if the server is not running and **running** (in green) if the server has already been started. The preference pane also shows the current setting for whether the MySQL server has been set to start automatically.

- **To start MySQL using the preference pane:**

Click **Start MySQL Server**. You may be prompted for the username and password of a user with administrator privileges to start the MySQL server.

- **To stop MySQL using the preference pane:**

Click **Stop MySQL Server**. You may be prompted for the username and password of a user with administrator privileges to stop the MySQL server.

- **To automatically start the MySQL server when the system boots:**

Check the check box next to **Automatically Start MySQL Server on Startup**.

- **To disable automatic MySQL server startup when the system boots:**

Uncheck the check box next to **Automatically Start MySQL Server on Startup**.

You can close the **System Preferences...** window once you have completed your settings.

## 2.4.5. Using the Bundled MySQL on Mac OS X Server

If you are running Mac OS X Server, a version of MySQL should already be installed. The following table shows the versions of MySQL that ship with Mac OS X Server versions.

**Table 2.7. MySQL Versions Preinstalled with Mac OS X Server**

Mac OS X Server Version	MySQL Version
10.2-10.2.2	3.23.51
10.2.3-10.2.6	3.23.53
10.3	4.0.14
10.3.2	4.0.16
10.4.0	4.1.10a
10.5.0	5.0.45
10.6.0	5.0.82

The following table shows the installation layout of MySQL on Mac OS X Server.

**Table 2.8. MySQL Directory Layout for Preinstalled MySQL Installations on Mac OS X Server**

Directory	Contents of Directory
<code>/usr/bin</code>	Client programs
<code>/var/mysql</code>	Log files, databases
<code>/usr/libexec</code>	The <code>mysqld</code> server
<code>/usr/share/man</code>	Unix manual pages
<code>/usr/share/mysql/mysql-test</code>	MySQL test suite
<code>/usr/share/mysql</code>	Miscellaneous support files, including error messages, character set files, sample configuration files, SQL for database installation
<code>/var/mysql/mysql.sock</code>	Location of the MySQL Unix socket

### Additional Resources

- For more information on managing the bundled MySQL instance in Mac OS X Server 10.5, see [Mac OS X Server: Web Technologies Administration For Version 10.5 Leopard](#).
- For more information on managing the bundled MySQL instance in Mac OS X Server 10.6, see [Mac OS X Server: Web Technologies Administration Version 10.6 Snow Leopard](#).
- The MySQL server bundled with Mac OS X Server does not include the MySQL client libraries and header files required to access and use MySQL from a third-party driver, such as Perl DBI or PHP. For more information on obtaining and installing MySQL libraries, see [Mac OS X Server version 10.5: MySQL libraries available for download](#). Alternatively, you can ignore the bundled MySQL server and install MySQL from the package or tarball installation.

## 2.5. Installing MySQL on Linux

Linux supports a number of different solutions for installing MySQL. The recommended method is to use one of the distributions from Oracle. If you choose this method, there are three options available:

- Installing from a generic binary package in `.tar.gz` format. See [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#) for more information.

- Extracting and compiling MySQL from a source distribution. For detailed instructions, see [Section 2.9, “Installing MySQL from Source”](#).
- Installing using a pre-compiled RPM package. For more information on using the RPM solution, see [Section 2.5.1, “Installing MySQL from RPM Packages on Linux”](#).

As an alternative, you can use the native package manager within your Linux distribution to automatically download and install MySQL for you. Native package installations can take of the download and dependencies required to run MySQL, but the MySQL version will often be some way behind the currently available release. You will also normally be unable to install developmental releases, as these are not usually made available in the native repository. For more information on using the native package installers, see [Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”](#).



#### Note

For many Linux installations, you will want to set up MySQL to be started automatically when your machine starts. Many of the native package installations perform this operation for you, but for source, binary and RPM solutions you may need to set this up separately. The required script, `mysql.server`, can be found in the `support-files` directory under the MySQL installation directory or in a MySQL source tree. You can install it as `/etc/init.d/mysql` for automatic MySQL startup and shutdown. See [Section 2.10.1.2, “Starting and Stopping MySQL Automatically”](#).

## 2.5.1. Installing MySQL from RPM Packages on Linux



#### Note

To install or upgrade to MySQL 5.6.11, be sure to read the special instructions at the end of this section.

The recommended way to install MySQL on RPM-based Linux distributions is by using the RPM packages. The RPMs that we provide to the community should work on all versions of Linux that support RPM packages and use `glibc` 2.3. To obtain RPM packages, see [Section 2.1.3, “How to Get MySQL”](#).

For non-RPM Linux distributions, you can install MySQL using a `.tar.gz` package. See [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#).

Installations created from our Linux RPM distributions result in files under the following system directories.

**Table 2.9. MySQL Installation Layout for Linux RPM**

Directory	Contents of Directory
<code>/usr/bin</code>	Client programs and scripts
<code>/usr/sbin</code>	The <code>mysqld</code> server
<code>/var/lib/mysql</code>	Log files, databases
<code>/usr/share/info</code>	Manual in Info format
<code>/usr/share/man</code>	Unix manual pages
<code>/usr/include/mysql</code>	Include (header) files
<code>/usr/lib/mysql</code>	Libraries
<code>/usr/share/mysql</code>	Miscellaneous support files, including error messages, character set files, sample configuration files, SQL for database installation
<code>/usr/share/sql-bench</code>	Benchmarks

**Note**

RPM distributions of MySQL are also provided by other vendors. Be aware that they may differ from those built by us in features, capabilities, and conventions (including communication setup), and that the instructions in this manual do not necessarily apply to installing them. The vendor's instructions should be consulted instead. Because of these differences, RPM packages built by us check whether such RPMs built by other vendors are installed. If so, the RPM does not install and produces a message explaining this.

In most cases, you need to install only the `MySQL-server` and `MySQL-client` packages to get a functional MySQL installation. The other packages are not required for a standard installation.

As of MySQL 5.6.8, new RPM install operations (not upgrades) invoke `mysql_install_db` with the `--random-passwords` option that provides for more secure MySQL installation. Invoking `mysql_install_db` with `--random-passwords` [237] causes it to assign a random password to the MySQL `root` accounts, set the “password expired” flag for those accounts, and remove the anonymous-user MySQL accounts. It will be necessary after installation to start the server, connect as `root` using the password written to the `$HOME/.mysql_secret` file, and assign a new `root` password. Until this is done, `root` cannot do anything else. This must be done for each `root` account you intend to use. To change the password, you can use the `SET PASSWORD` statement (for example, with the `mysql` client). You can also use `mysqladmin` or `mysql_secure_installation`. For additional details, see [Section 4.4.3, “mysql\\_install\\_db — Initialize MySQL Data Directory”](#). (Install operations using RPMs for Unbreakable Linux Network are unaffected because they do not use `mysql_install_db`.)

**RPMs for MySQL Cluster.** Standard MySQL server RPMs built by MySQL do not provide support for the `NDBCLUSTER` storage engine.

**Important**

When upgrading a MySQL Cluster RPM installation, you must upgrade *all* installed RPMs, including the `Server` and `Client` RPMs.

For more information about installing MySQL Cluster from RPMs, see [MySQL Cluster Installation and Upgrades](#).

For upgrades, if your installation was originally produced by installing multiple RPM packages, it is best to upgrade all the packages, not just some. For example, if you previously installed the server and client RPMs, do not upgrade just the server RPM.

If you get a dependency failure when trying to install MySQL packages (for example, `error: removing these packages would break dependencies: libmysqlclient.so.10 is needed by ...`), you should also install the `MySQL-shared-compat` package, which includes the shared libraries for older releases for backward compatibility.

The RPM packages shown in the following list are available. The names shown here use a suffix of `.glibc23.i386.rpm`, but particular packages can have different suffixes, described later.

- `MySQL-server-VERSION.glibc23.i386.rpm`

The MySQL server. You need this unless you only want to connect to a MySQL server running on another machine.

- `MySQL-client-VERSION.glibc23.i386.rpm`

The standard MySQL client programs. You probably always want to install this package.

- `MySQL-devel-VERSION.glibc23.i386.rpm`

The libraries and include files that are needed if you want to compile other MySQL clients, such as the Perl modules.

- [MySQL-shared-VERSION.glibc23.i386.rpm](#)

This package contains the shared libraries (`libmysqlclient.so*`) that certain languages and applications need to dynamically load and use MySQL. It contains single-threaded and thread-safe libraries.

- [MySQL-shared-compat-VERSION.glibc23.i386.rpm](#)

This package includes the shared libraries for older releases, but not the libraries for the current release. It contains single-threaded and thread-safe libraries. Install this package if you have applications installed that are dynamically linked against older versions of MySQL but you want to upgrade to the current version without breaking the library dependencies.

As of MySQL 5.6.5, the [MySQL-shared-compat](#) RPM package enables users of Red Hat-provided `mysql-*5.1` RPM packages to migrate to Oracle-provided `MySQL-*5.5` packages. [MySQL-shared-compat](#) replaces the Red Hat `mysql-libs` package by replacing `libmysqlclient.so` files of the latter package, thus satisfying dependencies of other packages on `mysql-libs`. This change affects only users of Red Hat (or Red Hat-compatible) RPM packages. Nothing is different for users of Oracle RPM packages.

- [MySQL-embedded-VERSION.glibc23.i386.rpm](#)

The embedded MySQL server library.

- [MySQL-test-VERSION.glibc23.i386.rpm](#)

This package includes the MySQL test suite.

- [MySQL-VERSION.src.rpm](#)

This contains the source code for all of the previous packages. It can also be used to rebuild the RPMs on other architectures (for example, Alpha or SPARC).

The suffix of RPM package names (following the `VERSION` value) has the following syntax:

```
.PLATFORM.CPU.rpm
```

The `PLATFORM` and `CPU` values indicate the type of system for which the package is built. `PLATFORM` indicates the platform and `CPU` indicates the processor type or family.

All packages are dynamically linked against `glibc` 2.3. The `PLATFORM` value indicates whether the package is platform independent or intended for a specific platform, as shown in the following table.

**Table 2.10. MySQL Linux Installation Packages**

<code>PLATFORM</code> Value	Intended Use
<code>glibc23</code>	Platform independent, should run on any Linux distribution that supports <code>glibc</code> 2.3
<code>rhel4, rhel5</code>	Red Hat Enterprise Linux 4 or 5
<code>e16</code>	Enterprise Linux 6
<code>sles10, sles11</code>	SuSE Linux Enterprise Server 10 or 11

In MySQL 5.6, only `glibc23` packages are available currently.

The `CPU` value indicates the processor type or family for which the package is built.

**Table 2.11. MySQL Installation Packages for Linux CPU Identifier**

<code>CPU</code> Value	Intended Processor Type or Family
<code>i386, i586, i686</code>	Pentium processor or better, 32 bit
<code>x86_64</code>	64-bit x86 processor



CPU Value	Intended Processor Type or Family
ia64	Itanium (IA-64) processor

To see all files in an RPM package (for example, a `MySQL-server` RPM), run a command like this:

```
shell> rpm -qpl MySQL-server-VERSION.glibc23.i386.rpm
```

To perform a standard minimal installation, install the server and client RPMs:

```
shell> rpm -i MySQL-server-VERSION.glibc23.i386.rpm
shell> rpm -i MySQL-client-VERSION.glibc23.i386.rpm
```

To install only the client programs, install just the client RPM:

```
shell> rpm -i MySQL-client-VERSION.glibc23.i386.rpm
```

RPM provides a feature to verify the integrity and authenticity of packages before installing them. If you would like to learn more about this feature, see [Section 2.1.4, “Verifying Package Integrity Using MD5 Checksums or GnuPG”](#).

The server RPM places data under the `/var/lib/mysql` directory. The RPM also creates a login account for a user named `mysql` (if one does not exist) to use for running the MySQL server, and creates the appropriate entries in `/etc/init.d/` to start the server automatically at boot time. (This means that if you have performed a previous installation and have made changes to its startup script, you may want to make a copy of the script so that you do not lose it when you install a newer RPM.) See [Section 2.10.1.2, “Starting and Stopping MySQL Automatically”](#), for more information on how MySQL can be started automatically on system startup.

In MySQL 5.6, during a new installation, the server boot scripts are installed, but the MySQL server is not started at the end of the installation, since the status of the server during an unattended installation is not known.

In MySQL 5.6, during an upgrade installation using the RPM packages, if the MySQL server is running when the upgrade occurs, the MySQL server is stopped, the upgrade occurs, and the MySQL server is restarted. If the MySQL server is not already running when the RPM upgrade occurs, the MySQL server is not started at the end of the installation.

If something goes wrong, you can find more information in the binary installation section. See [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#).



#### Note

The accounts that are listed in the MySQL grant tables initially have no passwords. After starting the server, you should set up passwords for them using the instructions in [Section 2.10, “Postinstallation Setup and Testing”](#).

During RPM installation, a user named `mysql` and a group named `mysql` are created on the system. This is done using the `useradd`, `groupadd`, and `usermod` commands. Those commands require appropriate administrative privileges, which is required for locally managed users and groups (as listed in the `/etc/passwd` and `/etc/group` files) by the RPM installation process being run by `root`.

If you log in as the `mysql` user, you may find that MySQL displays “Invalid (old?) table or database name” errors that mention `.mysqlgui`, `lost+found`, `.mysqlgui`, `.bash_history`, `.fonts.cache-1`, `.lesshtst`, `.mysql_history`, `.profile`, `.viminfo`, and similar files created by MySQL or operating system utilities. You can safely ignore these error messages or remove the files or directories that cause them if you do not need them.

For nonlocal user management (LDAP, NIS, and so forth), the administrative tools may require additional authentication (such as a password), and will fail if the installing user does not provide this authentication. Even if they fail, the RPM installation will not abort but succeed, and this is intentional. If they failed, some of the intended transfer of ownership may be missing, and it is recommended that the system administrator then manually ensures some appropriate user and group exists and manually transfers ownership following the actions in the RPM spec file.



In MySQL 5.6.11, the RPM spec file has been updated, which has the following consequences:

- For a non-upgrade installation (no existing MySQL version installed), it is possible to install MySQL using `yum`.
- For upgrades, it is necessary to clean up any earlier MySQL installations. In effect, the update is performed by removing the old installations and installing the new one.

Additional details follow.

For a non-upgrade installation of MySQL 5.6.11, it is possible to install using `yum`:

```
shell> yum install MySQL-server-NEWVERSION.glibc23.i386.rpm
```

For upgrades to MySQL 5.6.11, the upgrade is performed by removing the old installation and installing the new one. To do this, use the following procedure:

1. Remove the existing 5.6.X installation. `OLDVERSION` is the version to remove.

```
shell> rpm -e MySQL-server-OLDVERSION.glibc23.i386.rpm
```

Repeat this step for all installed MySQL RPMs.

2. Install the new version. `NEWVERSION` is the version to install.

```
shell> rpm -ivh MySQL-server-NEWVERSION.glibc23.i386.rpm
```

Alternatively, the removal and installation can be done using `yum`:

```
shell> yum remove MySQL-server-OLDVERSION.glibc23.i386.rpm
shell> yum install MySQL-server-NEWVERSION.glibc23.i386.rpm
```

## 2.5.2. Installing MySQL on Linux using Native Package Manager

Many Linux distributions include a version of the MySQL server, client tools, and development components into the standard package management system built into distributions such as Fedora, Debian, Ubuntu, and Gentoo. This section provides basic instructions for installing MySQL using these systems.



### Important

Native package installations can take care of the download and dependencies required to run MySQL, but the MySQL version will often be some way behind the currently available release. You will also normally be unable to install developmental releases, as these are not usually made available in the native repository.

Distribution specific instructions are shown below:

- **Red Hat Linux, Fedora, CentOS**

For Red Hat and similar distributions, the MySQL distribution is divided into a number of separate packages, `mysql` for the client tools, `mysql-server` for the server and associated tools, and `mysql-libs` for the libraries. The libraries are required if you want to provide connectivity from different languages and environments such as Perl, Python and others.

To install, use the `yum` command to specify the packages that you want to install. For example:

```
root-shell> yum install mysql mysql-server mysql-libs mysql-server
Loaded plugins: presto, refresh-packagekit
Setting up Install Process
Resolving Dependencies
--> Running transaction check
----> Package mysql.x86_64 0:5.1.48-2.fc13 set to be updated
----> Package mysql-libs.x86_64 0:5.1.48-2.fc13 set to be updated
----> Package mysql-server.x86_64 0:5.1.48-2.fc13 set to be updated
```

```
--> Processing Dependency: perl-DBD-MySQL for package: mysql-server-5.1.48-2.fc13.x86_64
--> Running transaction check
---> Package perl-DBD-MySQL.x86_64 0:4.017-1.fc13 set to be updated
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch             Version          Repository       Size
=====
Installing:
mysql                  x86_64           5.1.48-2.fc13   updates         889 k
mysql-libs             x86_64           5.1.48-2.fc13   updates         1.2 M
mysql-server          x86_64           5.1.48-2.fc13   updates         8.1 M
Installing for dependencies:
perl-DBD-MySQL        x86_64           4.017-1.fc13   updates         136 k

Transaction Summary
=====
Install                4 Package(s)
Upgrade                0 Package(s)

Total download size: 10 M
Installed size: 30 M
Is this ok [y/N]: y
Downloading Packages:
Setting up and reading Presto delta metadata
Processing delta metadata
Package(s) data still to download: 10 M
(1/4): mysql-5.1.48-2.fc13.x86_64.rpm           | 889 kB    00:04
(2/4): mysql-libs-5.1.48-2.fc13.x86_64.rpm     | 1.2 MB    00:06
(3/4): mysql-server-5.1.48-2.fc13.x86_64.rpm   | 8.1 MB    00:40
(4/4): perl-DBD-MySQL-4.017-1.fc13.x86_64.rpm  | 136 kB    00:00
-----
Total                                           201 kB/s | 10 MB    00:52
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing      : mysql-libs-5.1.48-2.fc13.x86_64          1/4
  Installing      : mysql-5.1.48-2.fc13.x86_64             2/4
  Installing      : perl-DBD-MySQL-4.017-1.fc13.x86_64     3/4
  Installing      : mysql-server-5.1.48-2.fc13.x86_64     4/4

Installed:
mysql.x86_64 0:5.1.48-2.fc13          mysql-libs.x86_64 0:5.1.48-2.fc13
mysql-server.x86_64 0:5.1.48-2.fc13

Dependency Installed:
perl-DBD-MySQL.x86_64 0:4.017-1.fc13

Complete!
```

MySQL and the MySQL server should now be installed. A sample configuration file is installed into `/etc/my.cnf`. An init script, to start and stop the server, will have been installed into `/etc/init.d/mysql`. To start the MySQL server use `service`:

```
root-shell> service mysqld start
```

To enable the server to be started and stopped automatically during boot, use `chkconfig`:

```
root-shell> chkconfig --levels 235 mysqld on
```

Which enables the MySQL server to be started (and stopped) automatically at the specified the run levels.

The database tables will have been automatically created for you, if they do not already exist. You should, however, run `mysql_secure_installation` to set the root passwords on your server.

- **Debian, Ubuntu, Kubuntu**

On Debian and related distributions, there are two packages, `mysql-client` and `mysql-server`, for the client and server components respectively. You should specify an explicit version, for example `mysql-client-5.1`, to ensure that you install the version of MySQL that you want.

To download and install, including any dependencies, use the `apt-get` command, specifying the packages that you want to install.



### Note

Before installing, make sure that you update your `apt-get` index files to ensure you are downloading the latest available version.

A sample installation of the MySQL packages might look like this (some sections trimmed for clarity):

```
root-shell> apt-get install mysql-client-5.1 mysql-server-5.1
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following packages were automatically installed and are no longer required:
  linux-headers-2.6.28-11 linux-headers-2.6.28-11-generic
Use 'apt-get autoremove' to remove them.
The following extra packages will be installed:
  bsd-mailx libdbd-mysql-perl libdbi-perl libhtml-template-perl
  libmysqlclient15off libmysqlclient16 libnet-daemon-perl libplrpc-perl mailx
  mysql-common postfix
Suggested packages:
  dbshell libipc-sharedcache-perl tinyca procmail postfix-mysql postfix-pgsql
  postfix-ldap postfix-pcre sasl2-bin resolvconf postfix-cdb
The following NEW packages will be installed:
  bsd-mailx libdbd-mysql-perl libdbi-perl libhtml-template-perl
  libmysqlclient15off libmysqlclient16 libnet-daemon-perl libplrpc-perl mailx
  mysql-client-5.1 mysql-common mysql-server-5.1 postfix
0 upgraded, 13 newly installed, 0 to remove and 182 not upgraded.
Need to get 1907kB/25.3MB of archives.
After this operation, 59.5MB of additional disk space will be used.
Do you want to continue [Y/n]? Y
Get: 1 http://gb.archive.ubuntu.com jaunty-updates/main mysql-common 5.1.30really5.0.75-0ubuntu10.5 [63.0kB]
Get: 2 http://gb.archive.ubuntu.com jaunty-updates/main libmysqlclient15off 5.1.30really5.0.75-0ubuntu10.5 [1872kB]
Fetched 1907kB in 9s (205kB/s)
Preconfiguring packages ...
Selecting previously deselected package mysql-common.
(Reading database ... 121260 files and directories currently installed.)
...
Processing 1 added doc-base file(s)...
Registering documents with scrollkeeper...
Setting up libnet-daemon-perl (0.43-1) ...
Setting up libplrpc-perl (0.2020-1) ...
Setting up libdbi-perl (1.607-1) ...
Setting up libmysqlclient15off(5.1.30really5.0.75-0ubuntu10.5) ...

Setting up libdbd-mysql-perl (4.008-1) ...
Setting up libmysqlclient16 (5.1.31-lubuntu2) ...

Setting up mysql-client-5.1 (5.1.31-lubuntu2) ...

Setting up mysql-server-5.1 (5.1.31-lubuntu2) ...
* Stopping MySQL database server mysqld
  ...done.
100825 11:46:15 InnoDB: Started; log sequence number 0 46409
100825 11:46:15 InnoDB: Starting shutdown...
100825 11:46:17 InnoDB: Shutdown completed; log sequence number 0 46409
100825 11:46:17 [Warning] Forcing shutdown of 1 plugins
* Starting MySQL database server mysqld
  ...done.
* Checking for corrupt, not cleanly closed and upgrade needing tables.
...
Processing triggers for libc6 ...
ldconfig deferred processing now taking place
```



### Note

The `apt-get` command will install a number of packages, including the MySQL server, in order to provide the typical tools and application environment. This can mean that you install a large number of packages in addition to the main MySQL package.

During installation, the initial database will be created, and you will be prompted for the MySQL root password (and confirmation). A configuration file will have been created in `/etc/mysql/my.cnf`. An init script will have been created in `/etc/init.d/mysql`.

The server will already be started. You can manually start and stop the server using:

```
root-shell> service mysql [start|stop]
```

The service will automatically be added to the 2, 3 and 4 run levels, with stop scripts in the single, shutdown and restart levels.

### • Gentoo Linux

As a source-based distribution, installing MySQL on Gentoo involves downloading the source, patching the Gentoo specifics, and then compiling the MySQL server and installing it. This process is handled automatically by the `emerge` command. Depending on the version of MySQL that you want to install, you may need to unmask the specific version that you want for your chosen platform.

The MySQL server and client tools are provided within a single package, `dev-db/mysql`. You can obtain a list of the versions available to install by looking at the portage directory for the package:

```
root-shell> ls /usr/portage/dev-db/mysql/mysql-5.1*
mysql-5.1.39-r1.ebuild
mysql-5.1.44-r1.ebuild
mysql-5.1.44-r2.ebuild
mysql-5.1.44-r3.ebuild
mysql-5.1.44.ebuild
mysql-5.1.45-r1.ebuild
mysql-5.1.45.ebuild
mysql-5.1.46.ebuild
```

To install a specific MySQL version, you must specify the entire atom. For example:

```
root-shell> emerge =dev-db/mysql-5.1.46
```

A simpler alternative is to use the `virtual/mysql-5.1` package, which will install the latest version:

```
root-shell> emerge =virtual/mysql-5.1
```

If the package is masked (because it is not tested or certified for the current platform), use the `ACCEPT_KEYWORDS` environment variable. For example:

```
root-shell> ACCEPT_KEYWORDS="~x86" emerge =virtual/mysql-5.1
```

After installation, you should create a new database using `mysql_install_db`, and set the password for the root user on MySQL. You can use the configuration interface to set the password and create the initial database:

```
root-shell> emerge --config =dev-db/mysql-5.1.46
```

A sample configuration file will have been created for you in `/etc/mysql/my.cnf`, and an init script will have been created in `/etc/init.d/mysql`.

To enable MySQL to start automatically at the normal (default) run levels, you can use:

```
root-shell> rc-update add mysql default
```

## 2.6. Installing MySQL on Solaris and OpenSolaris

MySQL on Solaris and OpenSolaris is available in a number of different formats.

- For information on installing using the native Solaris `PKG` format, see [Section 2.6.1, “Installing MySQL on Solaris using a Solaris `PKG`”](#).
- On OpenSolaris, the standard package repositories include MySQL packages specially built for OpenSolaris that include entries for the Service Management Framework (SMF) to enable control of the installation using the SMF administration commands. For more information, see [Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”](#).
- To use a standard `tar` binary installation, use the notes provided in [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#). Check the notes and hints at the end of this section for Solaris specific notes that you may need before or after installation.

To obtain a binary MySQL distribution for Solaris in tarball or PKG format, <http://dev.mysql.com/downloads/mysql/5.6.html>.

Additional notes to be aware of when installing and using MySQL on Solaris:

- If you want to use MySQL with the `mysql` user and group, use the `groupadd` and `useradd` commands:

```
groupadd mysql
useradd -g mysql mysql
```

- If you install MySQL using a binary tarball distribution on Solaris, you may run into trouble even before you get the MySQL distribution unpacked, as the Solaris `tar` cannot handle long file names. This means that you may see errors when you try to unpack MySQL.

If this occurs, you must use GNU `tar` (`gtar`) to unpack the distribution. In Solaris 10 and OpenSolaris `gtar` is normally located in `/usr/sfw/bin/gtar`, but may not be included in the default path definition.

- When using Solaris 10 for x86\_64, you should mount any file systems on which you intend to store `InnoDB` files with the `forcedirectio` option. (By default mounting is done without this option.) Failing to do so will cause a significant drop in performance when using the `InnoDB` storage engine on this platform.
- If you would like MySQL to start automatically, you can copy `support-files/mysql.server` to `/etc/init.d` and create a symbolic link to it named `/etc/rc3.d/S99mysql.server`.
- If too many processes try to connect very rapidly to `mysqld`, you should see this error in the MySQL log:

```
Error in accept: Protocol error
```

You might try starting the server with the `--back_log=50` [\[465\]](#) option as a workaround for this.

- To configure the generation of core files on Solaris you should use the `coreadm` command. Because of the security implications of generating a core on a `setuid()` application, by default, Solaris does not support core files on `setuid()` programs. However, you can modify this behavior using `coreadm`. If you enable `setuid()` core files for the current user, they will be generated using the mode 600 and owned by the superuser.

### 2.6.1. Installing MySQL on Solaris using a Solaris `PKG`

You can install MySQL on Solaris and OpenSolaris using a binary package using the native Solaris `PKG` format instead of the binary tarball distribution.

To use this package, download the corresponding `mysql-VERSION-solaris10-PLATFORM.pkg.gz` file, then decompress it. For example:

```
shell> gunzip mysql-5.6.12-solaris10-x86_64.pkg.gz
```

To install a new package, use `pkgadd` and follow the onscreen prompts. You must have root privileges to perform this operation:

```
shell> pkgadd -d mysql-5.6.12-solaris10-x86_64.pkg

The following packages are available:
 1  mysql      MySQL Community Server (GPL)
      (i86pc) 5.6.12

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]:
```

The `PKG` installer installs all of the files and tools needed, and then initializes your database if one does not exist. To complete the installation, you should set the root password for MySQL as provided in the instructions at the end of the installation. Alternatively, you can run the `mysql_secure_installation` script that comes with the installation.

The default installation directory is `/opt/mysql`. You can only change the root path of the installation when using `pkgadd`, which can be used to install MySQL in a different Solaris zone. If you need to install in a specific directory, use the binary `tar` file.

The `pkg` installer copies a suitable startup script for MySQL into `/etc/init.d/mysql`. To enable MySQL to startup and shutdown automatically, you should create a link between this file and the init script directories. For example, to ensure safe startup and shutdown of MySQL you could use the following commands to add the right links:

```
shell> ln /etc/init.d/mysql /etc/rc3.d/S91mysql
shell> ln /etc/init.d/mysql /etc/rc0.d/K02mysql
```

To remove MySQL, the installed package name is `mysql`. You can use this in combination with the `pkgrm` command to remove the installation.

To upgrade when using the Solaris package file format, you must remove the existing installation before installing the updated package. Removal of the package does not delete the existing database information, only the server, binaries and support files. The typical upgrade sequence is therefore:

```
shell> mysqladmin shutdown
shell> pkgrm mysql
shell> pkgadd -d mysql-5.6.12-solaris10-x86_64.pkg
shell> mysqld_safe &
shell> mysql_upgrade
```

You should check the notes in [Section 2.11, “Upgrading or Downgrading MySQL”](#) before performing any upgrade.

## 2.6.2. Installing MySQL on OpenSolaris using IPS

OpenSolaris includes standard packages for MySQL in the core repository. The MySQL packages are based on a specific release of MySQL and updated periodically. For the latest release you must use either the native Solaris `PKG`, `tar`, or source installations. The native OpenSolaris packages include SMF files so that you can easily control your MySQL installation, including automatic startup and recovery, using the native service management tools.

To install MySQL on OpenSolaris, use the `pkg` command. You will need to be logged in as root, or use the `pfexec` tool, as shown in the example below:

```
shell> pfexec pkg install SUNWmysql156
```

The package set installs three individual packages, `SUNWmysql156lib`, which contains the MySQL client libraries; `SUNWmysql156r` which contains the root components, including SMF and configuration

files; and `SUNWmysql56u` which contains the scripts, binary tools and other files. You can install these packages individually if you only need the corresponding components.

The MySQL files are installed into `/usr/mysql` which symbolic links for the sub directories (`bin`, `lib`, etc.) to a version specific directory. For MySQL 5.6, the full installation is located in `/usr/mysql/5.6`. The default data directory is `/var/mysql/5.6/data`. The configuration file is installed in `/etc/mysql/5.6/my.cnf`. This layout permits multiple versions of MySQL to be installed, without overwriting the data and binaries from other versions.

Once installed, you must run `mysql_install_db` to initialize the database, and use the `mysql_secure_installation` to secure your installation.

## Using SMF to manage your MySQL installation

Once installed, you can start and stop your MySQL server using the installed SMF configuration. The service name is `mysql`, or if you have multiple versions installed, you should use the full version name, for example `mysql:version_56`. To start and enable MySQL to be started at boot time:

```
shell> svcadm enable mysql
```

To disable MySQL from starting during boot time, and shut the MySQL server down if it is running, use:

```
shell> svcadm disable mysql
```

To restart MySQL, for example after a configuration file changes, use the `restart` option:

```
shell> svcadm restart mysql
```

You can also use SMF to configure the data directory and enable full 64-bit mode. For example, to set the data directory used by MySQL:

```
shell> svccfg
svc:> select mysql:version_56
svc:/application/database/mysql:version_56> setprop mysql/data=/data0/mysql
```

By default, the 32-bit binaries are used. To enable the 64-bit server on 64-bit platforms, set the `enable_64bit` parameter. For example:

```
svc:/application/database/mysql:version_56> setprop mysql/enable_64bit=1
```

You need to refresh the SMF after settings these options:

```
shell> svcadm refresh mysql
```

## 2.7. Installing MySQL on HP-UX

MySQL for HP-UX is available in a number of different forms:

- Using a DEPOT distribution provided at <http://dev.mysql.com/downloads/>. Please read the [general notes on HP-UX installation](#) before continuing. For more information on DEPOT installations, see [Section 2.7.2, "Installing MySQL on HP-UX using DEPOT"](#).
- Using a binary tarball distribution provided at <http://dev.mysql.com/downloads/>. Please read the [general notes on HP-UX installation](#) before continuing. For more information on binary installations, see [Section 2.2, "Installing MySQL from Generic Binaries on Unix/Linux"](#).

### 2.7.1. General Notes on Installing MySQL on HP-UX

Some additional notes on installing and using MySQL on HP-UX:

- If you install MySQL using a binary tarball distribution on HP-UX, you may run into trouble even before you get the MySQL distribution unpacked, as the HP-UX `tar` cannot handle long file names. This means that you may see errors when you try to unpack MySQL.

If this occurs, you must use GNU `tar` (`gtar`) to unpack the distribution.



- Because of some critical bugs in the standard HP-UX libraries, you should install the following patches before trying to run MySQL on HP-UX 11.0:

```
PHKL_22840 Streams cumulative
PHNE_22397 ARPA cumulative
```

This solves the problem of getting `EWOULDBLOCK` from `recv()` and `EBADF` from `accept()` in threaded applications.

## 2.7.2. Installing MySQL on HP-UX using DEPOT

The HP-UX DEPOT format packages can be installed using the `swinstall` command. You should install the `ncurses` and `zlib` libraries before installing the MySQL DEPOT package. You can use the free software `depothelper` tool to install these packages and any dependencies for you automatically.

To install using the MySQL DEPOT packages, follow this guide:

1. Download the MySQL DEPOT package from <http://dev.mysql.com/downloads/>. You must decompress the package before installation:

```
root-shell> gunzip mysql-5.1.48-hpux11.31-ia64-64bit.depot.gz
```

2. Install the DEPOT package using `swinstall`:

```
root-shell> swinstall -s mysql-5.1.49-hpux11.31-ia64-64bit.depot
```

MySQL will be installed into a directory matching the depot package name, within `/usr/local`. For convenience, you may want to create a symbolic link to the installed directory, for example:

```
root-shell> ln -s mysql-5.1.49-hpux11.31-ia64-64bit mysql
```

3. Your package is now installed. You should complete the configuration of MySQL by creating a user and group:

```
root-shell> /usr/sbin/groupadd mysql
root-shell> /usr/sbin/useradd -g mysql -d /var/lib/mysql/ -s /bin/false mysql
```

4. Create the standard database using the new user/group you have created, and set the permissions:

```
root-shell> cd /usr/local/
root-shell> scripts/mysql_install_db --user=mysql
root-shell> chown -R root .
root-shell> chown -R mysql data
```

5. Finally, secure your new installation by setting the root passwords, and then start your MySQL server using the `mysql` user:

```
root-shell> mysql_secure_installation
root-shell> mysqld_safe --user=mysql &
```

## 2.8. Installing MySQL on FreeBSD

This section provides information about installing MySQL on variants of FreeBSD Unix.

You can install MySQL on FreeBSD by using the binary distribution provided by Oracle. For more information, see [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#).

The easiest (and preferred) way to install MySQL is to use the `mysql-server` and `mysql-client` ports available at <http://www.freebsd.org/>. Using these ports gives you the following benefits:

- A working MySQL with all optimizations enabled that are known to work on your version of FreeBSD.
- Automatic configuration and build.
- Startup scripts installed in `/usr/local/etc/rc.d`.

- The ability to use `pkg_info -L` to see which files are installed.
- The ability to use `pkg_delete` to remove MySQL if you no longer want it on your machine.

The MySQL build process requires GNU make (`gmake`) to work. If GNU `make` is not available, you must install it first before compiling MySQL.

To install using the ports system:

```
# cd /usr/ports/databases/mysql51-server
# make
...
# cd /usr/ports/databases/mysql51-client
# make
...
```

The standard port installation places the server into `/usr/local/libexec/mysqld`, with the startup script for the MySQL server placed in `/usr/local/etc/rc.d/mysql-server`.

Some additional notes on the BSD implementation:

- To remove MySQL after installation using the ports system:

```
# cd /usr/ports/databases/mysql51-server
# make deinstall
...
# cd /usr/ports/databases/mysql51-client
# make deinstall
...
```

- If you get problems with the current date in MySQL, setting the `TZ` variable should help. See [Section 2.12, “Environment Variables”](#).

## 2.9. Installing MySQL from Source

Building MySQL from the source code enables you to customize build parameters, compiler optimizations, and installation location. For a list of systems on which MySQL is known to run, see [Section 2.1.1, “Operating Systems Supported by MySQL Community Server”](#).

Before you proceed with an installation from source, check whether Oracle produces a precompiled binary distribution for your platform and whether it works for you. We put a great deal of effort into ensuring that our binaries are built with the best possible options for optimal performance. Instructions for installing binary distributions are available in [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#).

### Source Installation Methods

There are two methods for installing MySQL from source:

- Use a standard MySQL source distribution. To obtain a standard distribution, see [Section 2.1.3, “How to Get MySQL”](#). For instructions on building from a standard distribution, see [Section 2.9.2, “Installing MySQL from a Standard Source Distribution”](#).

Standard distributions are available as compressed `tar` files, Zip archives, or RPM packages. Distribution files have names of the form `mysql-VERSION.tar.gz`, `mysql-VERSION.zip`, or `mysql-VERSION.rpm`, where `VERSION` is a number like `5.6.12`. File names for source distributions can be distinguished from those for precompiled binary distributions in that source distribution names are generic and include no platform name, whereas binary distribution names include a platform name indicating the type of system for which the distribution is intended (for example, `pc-linux-i686` or `winx64`).

- Use a MySQL development tree. Development trees have not necessarily received the same level of testing as standard release distributions, so this installation method is usually required only if you

need the most recent code changes. For information on building from one of the development trees, see [Section 2.9.3, “Installing MySQL from a Development Source Tree”](#).

## Source Installation System Requirements

Installation of MySQL from source requires several development tools. Some of these tools are needed no matter whether you use a standard source distribution or a development source tree. Other tool requirements depend on which installation method you use.

To install MySQL from source, your system must have the following tools, regardless of installation method:

- [CMake](#), which is used as the build framework on all platforms. [CMake](#) can be downloaded from <http://www.cmake.org>.
- A good [make](#) program. Although some platforms come with their own [make](#) implementations, it is highly recommended that you use GNU [make](#) 3.75 or newer. It may already be available on your system as [gmake](#). GNU [make](#) is available from <http://www.gnu.org/software/make/>.
- A working ANSI C++ compiler. GCC 4.2.1 or later, Sun Studio 10 or later, Visual Studio 2008 or later, and many current vendor-supplied compilers are known to work.
- Perl is needed if you intend to run test scripts. Most Unix-like systems include Perl. On Windows, you can use a version such as ActiveState Perl.

To install MySQL from a standard source distribution, one of the following tools is required to unpack the distribution file:

- For a [.tar.gz](#) compressed [tar](#) file: GNU [gunzip](#) to uncompress the distribution and a reasonable [tar](#) to unpack it. If your [tar](#) program supports the [z](#) option, it can both uncompress and unpack the file.

GNU [tar](#) is known to work. The standard [tar](#) provided with some operating systems is not able to unpack the long file names in the MySQL distribution. You should download and install GNU [tar](#), or if available, use a preinstalled version of GNU [tar](#). Usually this is available as [gnutar](#), [gtar](#), or as [tar](#) within a GNU or Free Software directory, such as [/usr/sfw/bin](#) or [/usr/local/bin](#). GNU [tar](#) is available from <http://www.gnu.org/software/tar/>.

- For a [.zip](#) Zip archive: [WinZip](#) or another tool that can read [.zip](#) files.
- For an [.rpm](#) RPM package: The [rpmbuild](#) program used to build the distribution unpacks it.

To install MySQL from a development source tree, the following additional tools are required:

- To obtain the source tree, you must have Bazaar installed. The [Bazaar VCS Web site](#) has instructions for downloading and installing Bazaar on different platforms. Bazaar is supported on any platform that supports Python, and is therefore compatible with any Linux, Unix, Windows, or Mac OS X host.
- [bison](#) is needed to generate [sql\\_yacc.cc](#) from [sql\\_yacc.yy](#). You should use the latest version of [bison](#) where possible. Versions 1.75 and 2.1 are known to work. There have been reported problems with [bison](#) 1.875. If you experience problems, upgrade to a later, rather than earlier, version.

[bison](#) is available from <http://www.gnu.org/software/bison/>. [bison](#) for Windows can be downloaded from <http://gnuwin32.sourceforge.net/packages/bison.htm>. Download the package labeled “Complete package, excluding sources”. On Windows, the default location for [bison](#) is the [C:\Program Files\GnuWin32](#) directory. Some utilities may fail to find [bison](#) because of the space in the directory name. Also, Visual Studio may simply hang if there are spaces in the path. You can resolve these problems by installing into a directory that does not contain a space; for example [C:\GnuWin32](#).

- On OpenSolaris and Solaris Express, `m4` must be installed in addition to `bison`. `m4` is available from <http://www.gnu.org/software/m4/>.



#### Note

If you have to install any programs, modify your `PATH` environment variable to include any directories in which the programs are located. See [Section 4.2.4, “Setting Environment Variables”](#).

If you run into problems and need to file a bug report, please use the instructions in [Section 1.7, “How to Report Bugs or Problems”](#).

## 2.9.1. MySQL Layout for Source Installation

By default, when you install MySQL after compiling it from source, the installation step installs files under `/usr/local/mysql`. The component locations under the installation directory are the same as for binary distributions. See [Table 2.3, “MySQL Installation Layout for Generic Unix/Linux Binary Package”](#), and [Section 2.3.1, “MySQL Installation Layout on Microsoft Windows”](#). To configure installation locations different from the defaults, use the options described at [Section 2.9.4, “MySQL Source-Configuration Options”](#).

## 2.9.2. Installing MySQL from a Standard Source Distribution

To install MySQL from a standard source distribution:

1. Verify that your system satisfies the tool requirements listed at [Section 2.9, “Installing MySQL from Source”](#).
2. Obtain a distribution file using the instructions in [Section 2.1.3, “How to Get MySQL”](#).
3. Configure, build, and install the distribution using the instructions in this section.
4. Perform postinstallation procedures using the instructions in [Section 2.10, “Postinstallation Setup and Testing”](#).

In MySQL 5.6, `CMake` is used as the build framework on all platforms. The instructions given here should enable you to produce a working installation. For additional information on using `CMake` to build MySQL, see [How to Build MySQL Server with CMake](#).

If you start from a source RPM, use the following command to make a binary RPM that you can install. If you do not have `rpmbuild`, use `rpm` instead.

```
shell> rpmbuild --rebuild --clean MySQL-VERSION.src.rpm
```

The result is one or more binary RPM packages that you install as indicated in [Section 2.5.1, “Installing MySQL from RPM Packages on Linux”](#).

The sequence for installation from a compressed `tar` file or Zip archive source distribution is similar to the process for installing from a generic binary distribution (see [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#)), except that it is used on all platforms and includes steps to configure and compile the distribution. For example, with a compressed `tar` file source distribution on Unix, the basic installation command sequence looks like this:

```
# Preconfiguration setup
shell> groupadd mysql
shell> useradd -r -g mysql mysql
# Beginning of source-build specific instructions
shell> tar zxvf mysql-VERSION.tar.gz
shell> cd mysql-VERSION
shell> cmake .
shell> make
shell> make install
# End of source-build specific instructions
# Postinstallation setup
shell> cd /usr/local/mysql
```

```

shell> chown -R mysql .
shell> chgrp -R mysql .
shell> scripts/mysql_install_db --user=mysql
shell> chown -R root .
shell> chown -R mysql data
# Next command is optional
shell> cp support-files/my-medium.cnf /etc/my.cnf
shell> bin/mysqld_safe --user=mysql &
# Next command is optional
shell> cp support-files/mysql.server /etc/init.d/mysql.server

```

A more detailed version of the source-build specific instructions is shown following.



#### Note

The procedure shown here does not set up any passwords for MySQL accounts. After following the procedure, proceed to [Section 2.10](#), “Postinstallation Setup and Testing”, for postinstallation setup and testing.

## Perform Preconfiguration Setup

On Unix, set up the `mysql` user and group that will be used to run and execute the MySQL server and own the database directory. For details, see [Creating a `mysql` System User and Group](#), in [Section 2.2](#), “Installing MySQL from Generic Binaries on Unix/Linux”. Then perform the following steps as the `mysql` user, except as noted.

## Obtain and Unpack the Distribution

Pick the directory under which you want to unpack the distribution and change location into it.

Obtain a distribution file using the instructions in [Section 2.1.3](#), “How to Get MySQL”.

Unpack the distribution into the current directory:

- To unpack a compressed `tar` file, `tar` can uncompress and unpack the distribution if it has `z` option support:

```
shell> tar zxvf mysql-VERSION.tar.gz
```

If your `tar` does not have `z` option support, use `gunzip` to unpack the distribution and `tar` to unpack it:

```
shell> gunzip < mysql-VERSION.tar.gz | tar xvf -
```

Alternatively, `CMake` can uncompress and unpack the distribution:

```
shell> cmake -E tar zxvf mysql-VERSION.tar.gz
```

- To unpack a Zip archive, use `WinZip` or another tool that can read `.zip` files.

Unpacking the distribution file creates a directory named `mysql-VERSION`.

## Configure the Distribution

Change location into the top-level directory of the unpacked distribution:

```
shell> cd mysql-VERSION
```

Configure the source directory. The minimum configuration command includes no options to override configuration defaults:

```
shell> cmake .
```

On Windows, specify the development environment. For example, the following commands configure MySQL for 32-bit or 64-bit builds, respectively:

```
shell> cmake . -G "Visual Studio 9 2008"
```

```
shell> cmake . -G "Visual Studio 9 2008 Win64"
```

On Mac OS X, to use the Xcode IDE:

```
shell> cmake . -G Xcode
```

When you run `cmake`, you might want to add options to the command line. Here are some examples:

- `-DBUILD_CONFIG=mysql_release` [131]: Configure the source with the same build options used by Oracle to produce binary distributions for official MySQL releases.
- `-DCMAKE_INSTALL_PREFIX=dir_name` [132]: Configure the distribution for installation under a particular location.
- `-DCPACK_MONOLITHIC_INSTALL=1` [131]: Cause `make package` to generate a single installation file rather than multiple files.
- `-DWITH_DEBUG=1` [136]: Build the distribution with debugging support.

For a more extensive list of options, see [Section 2.9.4, “MySQL Source-Configuration Options”](#).

To list the configuration options, use one of the following commands:

```
shell> cmake . -L # overview
shell> cmake . -LH # overview with help text
shell> cmake . -LAH # all params with help text
shell> ccmake . # interactive display
```

If `CMake` fails, you might need to reconfigure by running it again with different options. If you do reconfigure, take note of the following:

- If `CMake` is run after it has previously been run, it may use information that was gathered during its previous invocation. This information is stored in `CMakeCache.txt`. When `CMake` starts up, it looks for that file and reads its contents if it exists, on the assumption that the information is still correct. That assumption is invalid when you reconfigure.
- Each time you run `CMake`, you must run `make` again to recompile. However, you may want to remove old object files from previous builds first because they were compiled using different configuration options.

To prevent old object files or configuration information from being used, run these commands on Unix before re-running `CMake`:

```
shell> make clean
shell> rm CMakeCache.txt
```

Or, on Windows:

```
shell> devenv MySQL.sln /clean
shell> del CMakeCache.txt
```

If you build out of the source tree (as described later), the `CMakeCache.txt` file and all built files are in the build directory, so you can remove that directory to object files and cached configuration information.

If you are going to send mail to a MySQL mailing list to ask for configuration assistance, first check the files in the `CMakeFiles` directory for useful information about the failure. To file a bug report, please use the instructions in [Section 1.7, “How to Report Bugs or Problems”](#).

## Build the Distribution

On Unix:

```
shell> make
shell> make VERBOSE=1
```

The second command sets `VERBOSE` to show the commands for each compiled source.

Use `gmake` instead on systems where you are using GNU `make` and it has been installed as `gmake`.

On Windows:

```
shell> devenv MySQL.sln /build RelWithDebInfo
```

It is possible to build out of the source tree to keep the tree clean. If the top-level source directory is named `mysql-src` under your current working directory, you can build in a directory named `build` at the same level like this:

```
shell> mkdir build
shell> cd build
shell> cmake ../mysql-src
```

If you have gotten to the compilation stage, but the distribution does not build, see [Section 2.9.5, “Dealing with Problems Compiling MySQL”](#), for help. If that does not solve the problem, please enter it into our bugs database using the instructions given in [Section 1.7, “How to Report Bugs or Problems”](#). If you have installed the latest versions of the required tools, and they crash trying to process our configuration files, please report that also. However, if you get a `command not found` error or a similar problem for required tools, do not report it. Instead, make sure that all the required tools are installed and that your `PATH` variable is set correctly so that your shell can find them.

## Install the Distribution

On Unix:

```
shell> make install
```

This installs the files under the configured installation directory (by default, `/usr/local/mysql`). You might need to run the command as `root`.

To install in a specific directory, add a `DESTDIR` parameter to the command line:

```
shell> make install DESTDIR="/opt/mysql"
```

Alternatively, generate installation package files that you can install where you like:

```
shell> make package
```

This operation produces one or more `.tar.gz` files that can be installed like generic binary distribution packages. See [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#). If you run `CMake` with `-DCPACK_MONOLITHIC_INSTALL=1` [131], the operation produces a single file. Otherwise, it produces multiple files.

On Windows, generate the data directory, then create a `.zip` archive installation package:

```
shell> devenv MySQL.sln /build RelWithDebInfo /project initial_database
shell> devenv MySQL.sln /build RelWithDebInfo /project package
```

You can install the resulting `.zip` archive where you like. See [Section 2.3.4, “Installing MySQL on Microsoft Windows Using a noinstall Zip Archive”](#).

## Perform Postinstallation Setup

The remainder of the installation process involves setting up the configuration file, creating the core databases, and starting the MySQL server. For instructions, see [Section 2.10, “Postinstallation Setup and Testing”](#).



### Note

The accounts that are listed in the MySQL grant tables initially have no passwords. After starting the server, you should set up passwords for them using the instructions in [Section 2.10, “Postinstallation Setup and Testing”](#).

## 2.9.3. Installing MySQL from a Development Source Tree



This section discusses how to install MySQL from the latest development source code. Development trees have not necessarily received the same level of testing as standard release distributions, so this installation method is usually required only if you need the most recent code changes. Do not use a development tree for production systems. If your goal is simply to get MySQL up and running on your system, you should use a standard release distribution (either a binary or source distribution). See [Section 2.1.3, “How to Get MySQL”](#).

MySQL development projects are hosted on [Launchpad](#). MySQL projects, including MySQL Server, MySQL Workbench, and others are available from the [Oracle/MySQL Engineering](#) page. For the repositories related only to MySQL Server, see the [MySQL Server](#) page.

To install MySQL from a development source tree, your system must satisfy the tool requirements listed at [Section 2.9, “Installing MySQL from Source”](#), including the requirements for Bazaar and `bison`.

To create a local branch of the MySQL development tree on your machine, use this procedure:

1. To obtain a copy of the MySQL source code, you must create a new Bazaar branch. If you do not already have a Bazaar repository directory set up, you must initialize a new directory:

```
shell> mkdir mysql-server
shell> bzz init-repo --trees mysql-server
```

This is a one-time operation.

2. Assuming that you have an initialized repository directory, you can branch from the public MySQL server repositories to create a local source tree. To create a branch of a specific version:

```
shell> cd mysql-server
shell> bzz branch lp:mysql-server/5.6 mysql-5.6
```

This is a one-time operation per source tree. You can branch the source trees for several versions of MySQL under the `mysql-server` directory.

3. The initial download will take some time to complete, depending on the speed of your connection. Please be patient. Once you have downloaded the first tree, additional trees should take significantly less time to download.
4. When building from the Bazaar branch, you may want to create a copy of your active branch so that you can make configuration and other changes without affecting the original branch contents. You can achieve this by branching from the original branch:

```
shell> bzz branch mysql-5.6 mysql-5.6-build
```

5. To obtain changes made after you have set up the branch initially, update it using the `pull` option periodically. Use this command in the top-level directory of the local copy:

```
shell> bzz pull
```

To examine the changeset comments for the tree, use the `log` option to `bzz`:

```
shell> bzz log
```

You can also browse changesets, comments, and source code online at the Launchpad [MySQL Server](#) page.

If you see diffs (changes) or code that you have a question about, do not hesitate to send email to the MySQL [internals](#) mailing list. See [Section 1.6.1, “MySQL Mailing Lists”](#). If you think you have a better idea on how to do something, send an email message to the list with a patch.

After you have the local branch, you can build MySQL server from the source code. For information, see [Section 2.9.2, “Installing MySQL from a Standard Source Distribution”](#), except that you skip the part about obtaining and unpacking the distribution.

Be careful about installing a build from a distribution source tree on a production machine. The installation command may overwrite your live release installation. If you already have MySQL installed and do not want to overwrite it, run `CMake` with values for the `CMAKE_INSTALL_PREFIX` [132], `MYSQL_TCP_PORT` [135], and `MYSQL_UNIX_ADDR` [135] options different from those used by your production server. For additional information about preventing multiple servers from interfering with each other, see [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).

Play hard with your new installation. For example, try to make new features crash. Start by running `make test`. See [Section 22.1.2, “The MySQL Test Suite”](#).

## 2.9.4. MySQL Source-Configuration Options

The `CMake` program provides a great deal of control over how you configure a MySQL source distribution. Typically, you do this using options on the `CMake` command line. For information about options supported by `CMake`, run either of these commands in the top-level source directory:

```
shell> cmake . -LH
shell> ccmake .
```

You can also affect `CMake` using certain environment variables. See [Section 2.12, “Environment Variables”](#).

The following table shows the available `CMake` options. In the `Default` column, `PREFIX` stands for the value of the `CMAKE_INSTALL_PREFIX` [132] option, which specifies the installation base directory. This value is used as the parent location for several of the installation subdirectories.

**Table 2.12. MySQL Source-Configuration Option Reference (CMake)**

Formats	Description	Default	Introduced	Removed
<code>BUILD_CONFIG</code> [131]	Use same build options as official releases			
<code>CMAKE_BUILD_TYPE</code> [131]	Type of build to produce	<code>RelWithDebInfo</code>		
<code>CMAKE_INSTALL_PREFIX</code> [132]	Installation base directory	<code>/usr/local/mysql</code>		
<code>COMPILATION_COMMENT</code> [134]	Comment about compilation environment			
<code>CPACK_MONOLITHIC_INSTALLATION</code> [133]	Whether package build produces single file	<code>OFF</code>		
<code>DEFAULT_CHARSET</code> [134]	The default server character set	<code>latin1</code>		
<code>DEFAULT_COLLATION</code> [134]	The default server collation	<code>latin1_swedish_ci</code>		
<code>ENABLE_DEBUG_SYNC</code> [134]	Whether to enable Debug Sync support	<code>ON</code>		
<code>ENABLE_DOWNLOADS</code> [135]	Whether to download optional files	<code>OFF</code>		
<code>ENABLE_DTRACE</code> [135]	Whether to include DTrace support			
<code>ENABLE_GCOV</code> [135]	Whether to include gcov support		5.6.3	
<code>ENABLE_GPROF</code> [135]	Enable gprof (optimized Linux builds only)	<code>OFF</code>	5.6.6	
<code>ENABLED_LOCAL_INFILE</code> [135]	Whether to enable LOCAL for LOAD DATA INFILE	<code>OFF</code>		
<code>ENABLED_PROFILING</code> [135]	Whether to enable query profiling code	<code>ON</code>		

MySQL Source-Configuration Options

Formats	Description	Default	Introduced	Removed
<a href="#">IGNORE_AIO_CHECK [135]</a>	With -DBUILD_CONFIG=mysql_release, ignore libaio check	OFF	5.6.1	
<a href="#">INSTALL_BINDIR [132]</a>	User executables directory	PREFIX/bin		
<a href="#">INSTALL_DOCDIR [132]</a>	Documentation directory	PREFIX/docs		
<a href="#">INSTALL_DOCREADMEDIR [132]</a>	README file directory	PREFIX		
<a href="#">INSTALL_INCLUDEDIR [132]</a>	Header file directory	PREFIX/include		
<a href="#">INSTALL_INFODIR [132]</a>	Info file directory	PREFIX/docs		
<a href="#">INSTALL_LAYOUT [132]</a>	Select predefined installation layout	STANDALONE		
<a href="#">INSTALL_LIBDIR [132]</a>	Library file directory	PREFIX/lib		
<a href="#">INSTALL_MANDIR [132]</a>	Manual page directory	PREFIX/man		
<a href="#">INSTALL_MYSQLSHAREDIR [132]</a>	Shared data directory	PREFIX/share		
<a href="#">INSTALL_MYSQLTESTDIR [132]</a>	mysql-test directory	PREFIX/mysql-test		
<a href="#">INSTALL_PLUGINDIR [133]</a>	Plugin directory	PREFIX/lib/plugin		
<a href="#">INSTALL_SBINDIR [133]</a>	Server executable directory	PREFIX/bin		
<a href="#">INSTALL_SCRIPTDIR [133]</a>	Scripts directory	PREFIX/scripts		
<a href="#">INSTALL_SHAREDIR [133]</a>	aclocal/mysql.m4 installation directory	PREFIX/share		
<a href="#">INSTALL_SQLBENCHDIR [133]</a>	sql-bench directory	PREFIX		
<a href="#">INSTALL_SUPPORTFILESDIR [133]</a>	Extra support files directory	PREFIX/support-files		
<a href="#">MYSQL_DATADIR [133]</a>	Data directory			
<a href="#">MYSQL_MAINTAINER_MODE [133]</a>	Whether to enable MySQL maintainer-specific development environment	OFF		
<a href="#">MYSQL_PROJECT_NAME [135]</a>	Windows/Mac OS X project name	3306	5.6.5	
<a href="#">MYSQL_TCP_PORT [135]</a>	TCP/IP port number	3306		
<a href="#">MYSQL_UNIX_ADDR [135]</a>	Unix socket file	/tmp/mysql.sock		
<a href="#">ODBC_INCLUDES [133]</a>	ODBC includes directory			
<a href="#">ODBC_LIB_DIR [133]</a>	ODBC library directory			
<a href="#">OPTIMIZER_TRACE [133]</a>	Whether to support optimizer tracing		5.6.3	
<a href="#">SYSCONFDIR [133]</a>	Option file directory			
<a href="#">WITH_DEBUG [136]</a>	Whether to include debugging support	OFF		
<a href="#">WITH_DEFAULT_COMPILER_OPTIONS [136]</a>	Whether to use default compiler options	ON	5.6.6	
<a href="#">WITH_DEFAULT_FEATURE_SET [135]</a>	Whether to use default feature set	ON	5.6.6	

Formats	Description	Default	Introduced	Removed
<a href="#">WITH_EMBEDDED_SERVER [136]</a>	Whether to build embedded server	OFF		
<a href="#">WITH_xxx_STORAGE_ENGINE [136]</a>	Compile storage engine xxx statically into server			
<a href="#">WITH_EXTRA_CHARSETS [136]</a>	Which extra character sets to include	all		
<a href="#">WITH_LIBEDIT [136]</a>	Use bundled libedit library	ON		
<a href="#">WITH_LIBWRAP [136]</a>	Whether to include libwrap (TCP wrappers) support	OFF		
<a href="#">WITH_READLINE [136]</a>	Use bundled readline library	OFF		5.6.5
<a href="#">WITH_SSL [136]</a>	Type of SSL support	no		
<a href="#">WITH_UNIXODBC [137]</a>	Enable unixODBC support	OFF		
<a href="#">WITH_ZLIB [137]</a>	Type of zlib support	system		
<a href="#">WITHOUT_xxx_STORAGE_ENGINE [136]</a>	Exclude storage engine xxx from build			

The following sections provide more information about [CMake](#) options.

- [General Options](#)
- [Installation Layout Options](#)
- [Feature Options](#)
- [Compiler Flags](#)

For boolean options, the value may be specified as 1 or [ON](#) to enable the option, or as 0 or [OFF](#) to disable the option.

Many options configure compile-time defaults that can be overridden at server startup. For example, the [CMAKE\\_INSTALL\\_PREFIX \[132\]](#), [MYSQL\\_TCP\\_PORT \[135\]](#), and [MYSQL\\_UNIX\\_ADDR \[135\]](#) options that configure the default installation base directory location, TCP/IP port number, and Unix socket file can be changed at server startup with the [--basedir \[414\]](#), [--port \[437\]](#), and [--socket \[442\]](#) options for `mysqld`. Where applicable, configuration option descriptions indicate the corresponding `mysqld` startup option.

## General Options

- `-DBUILD_CONFIG=mysql_release`

This option configures a source distribution with the same build options used by Oracle to produce binary distributions for official MySQL releases.

- `-DCMAKE_BUILD_TYPE=type`

The type of build to produce:

- `RelWithDebInfo`: Enable optimizations and generate debugging information. This is the default MySQL build type.
- `Debug`: Disable optimizations and generate debugging information. This build type is also used if the [WITH\\_DEBUG \[136\]](#) option is enabled. That is, `-DWITH_DEBUG=1 [136]` has the same effect as `-DCMAKE_BUILD_TYPE=Debug [131]`.
- `-DCPACK_MONOLITHIC_INSTALL=bool`

This option affects whether the `make package` operation produces multiple installation package files or a single file. If disabled, the operation produces multiple installation package files, which may be useful if you want to install only a subset of a full MySQL installation. If enabled, it produces a single file for installing everything.

## Installation Layout Options

The `CMAKE_INSTALL_PREFIX` [132] option indicates the base installation directory. Other options with names of the form `INSTALL_XXX` that indicate component locations are interpreted relative to the prefix and their values are relative pathnames. Their values should not include the prefix.

- `-DCMAKE_INSTALL_PREFIX=dir_name`

The installation base directory.

This value can be set at server startup with the `--basedir` [414] option.

- `-DINSTALL_BINDIR=dir_name`

Where to install user programs.

- `-DINSTALL_DOCDIR=dir_name`

Where to install documentation.

- `-DINSTALL_DOCREADMEDIR=dir_name`

Where to install `README` files.

- `-DINSTALL_INCLUDEDIR=dir_name`

Where to install header files.

- `-DINSTALL_INFODIR=dir_name`

Where to install Info files.

- `-DINSTALL_LAYOUT=name`

Select a predefined installation layout:

- `STANDALONE`: Same layout as used for `.tar.gz` and `.zip` packages. This is the default.
- `RPM`: Layout similar to RPM packages.
- `SVR4`: Solaris package layout.
- `DEB`: DEB package layout (experimental).

You can select a predefined layout but modify individual component installation locations by specifying other options. For example:

```
shell> cmake . -DINSTALL_LAYOUT=SVR4 -DMYSQL_DATADIR=/var/mysql/data
```

- `-DINSTALL_LIBDIR=dir_name`

Where to install library files.

- `-DINSTALL_MANDIR=dir_name`

Where to install manual pages.

- `-DINSTALL_MYSQLSHAREDIR=dir_name`

Where to install shared data files.

- `-DINSTALL_MYSQLTESTDIR=dir_name`

Where to install the `mysql-test` directory. As of MySQL 5.6.12, to suppress installation of this directory, explicitly set the option to the empty value (`-DINSTALL_MYSQLTESTDIR= [133]`).

- `-DINSTALL_PLUGINDIR=dir_name`

The location of the plugin directory.

This value can be set at server startup with the `--plugin_dir [525]` option.

- `-DINSTALL_SBINDIR=dir_name`

Where to install the `mysqld` server.

- `-DINSTALL_SCRIPTDIR=dir_name`

Where to install `mysql_install_db`.

- `-DINSTALL_SHAREDIR=dir_name`

Where to install `aclocal/mysql.m4`.

- `-DINSTALL_SQLBENCHDIR=dir_name`

Where to install the `sql-bench` directory. To suppress installation of this directory, explicitly set the option to the empty value (`-DINSTALL_SQLBENCHDIR= [133]`).

- `-DINSTALL_SUPPORTFILESDIR=dir_name`

Where to install extra support files.

- `-DMYSQL_DATADIR=dir_name`

The location of the MySQL data directory.

This value can be set at server startup with the `--datadir [418]` option.

- `-DOPTIMIZER_TRACE=bool`

Whether to support optimizer tracing. See [MySQL Internals: Tracing the Optimizer](#). This option was added in MySQL 5.6.3.

- `-DSYSCONFDIR=dir_name`

The default `my.cnf` option file directory.

This location cannot be set at server startup, but you can start the server with a given option file using the `--defaults-file=file_name [219]` option, where `file_name` is the full path name to the file.

- `-DODBC_INCLUDES=dir_name`

The location of the ODBC includes directory, and may be used while configuring Connector/ODBC.

- `-DODBC_LIB_DIR=dir_name`

The location of the ODBC library directory, and may be used while configuring Connector/ODBC.

## Storage Engine Options

Storage engines are built as plugins. You can build a plugin as a static module (compiled into the server) or a dynamic module (built as a dynamic library that must be installed into the server using the `INSTALL PLUGIN` statement or the `--plugin-load` [435] option before it can be used). Some plugins might not support static or dynamic building.

The `MyISAM`, `MERGE`, `MEMORY`, and `CSV` engines are mandatory (always compiled into the server) and need not be installed explicitly.

To compile a storage engine statically into the server, use `-DWITH_engine_STORAGE_ENGINE=1`. Some permissible *engine* values are `ARCHIVE`, `BLACKHOLE`, `EXAMPLE`, `FEDERATED`, `INNOBASE` (InnoDB), `PARTITION` (partitioning support), and `PERFSCHEMA` (Performance Schema). Examples:

```
-DWITH_INNOBASE_STORAGE_ENGINE=1
-DWITH_ARCHIVE_STORAGE_ENGINE=1
-DWITH_BLACKHOLE_STORAGE_ENGINE=1
-DWITH_PERFSCHEMA_STORAGE_ENGINE=1
```

To exclude a storage engine from the build, use `-DWITHOUT_engine_STORAGE_ENGINE=1`. Examples:

```
-DWITHOUT_EXAMPLE_STORAGE_ENGINE=1
-DWITHOUT_FEDERATED_STORAGE_ENGINE=1
-DWITHOUT_PARTITION_STORAGE_ENGINE=1
```

If neither `-DWITH_engine_STORAGE_ENGINE` nor `-DWITHOUT_engine_STORAGE_ENGINE` are specified for a given storage engine, the engine is built as a shared module, or excluded if it cannot be built as a shared module.

## Feature Options

- `-DCOMPILATION_COMMENT=string`

A descriptive comment about the compilation environment.

- `-DDEFAULT_CHARSET=charset_name`

The server character set. By default, MySQL uses the `latin1` (cp1252 West European) character set.

*charset\_name* may be one of `binary`, `armscii8`, `ascii`, `big5`, `cp1250`, `cp1251`, `cp1256`, `cp1257`, `cp850`, `cp852`, `cp866`, `cp932`, `dec8`, `eucjms`, `euckr`, `gb2312`, `gbk`, `geostd8`, `greek`, `hebrew`, `hp8`, `keybcs2`, `koi8r`, `koi8u`, `latin1`, `latin2`, `latin5`, `latin7`, `macce`, `macroman`, `sjis`, `swe7`, `tis620`, `ucs2`, `ujis`, `utf8`, `utf8mb4`, `utf16`, `utf16le`, `utf32`. The permissible character sets are listed in the `cmake/character_sets.cmake` file as the value of `CHARSETS_AVAILABLE`.

This value can be set at server startup with the `--character_set_server` [468] option.

- `-DDEFAULT_COLLATION=collation_name`

The server collation. By default, MySQL uses `latin1_swedish_ci`. Use the `SHOW COLLATION` statement to determine which collations are available for each character set.

This value can be set at server startup with the `--collation_server` [470] option.

- `-DENABLE_DEBUG_SYNC=bool`

Whether to compile the Debug Sync facility into the server. This facility is used for testing and debugging. This option is enabled by default, but has no effect unless MySQL is configured with debugging enabled. If debugging is enabled and you want to disable Debug Sync, use `-DENABLE_DEBUG_SYNC=0` [134].



When compiled in, Debug Sync is disabled by default at runtime. To enable it, start `mysqld` with the `--debug-sync-timeout=N` [419] option, where *N* is a timeout value greater than 0. (The default value is 0, which disables Debug Sync.) *N* becomes the default timeout for individual synchronization points.

For a description of the Debug Sync facility and how to use synchronization points, see [MySQL Internals: Test Synchronization](#).

- `-DENABLE_DOWNLOADS=bool`

Whether to download optional files. For example, with this option enabled, `CMake` downloads the Google Test distribution that is used by the test suite to run unit tests.

- `-DENABLE_DTRACE=bool`

Whether to include support for DTrace probes. For information about DTrace, see [Section 5.7, “Tracing `mysqld` Using DTrace”](#)

- `-DENABLE_GCOV=bool`

Whether to include gcov support (Linux only).

- `-DENABLE_GPROF=bool`

Whether to enable `gprof` (optimized Linux builds only). This option was added in MySQL 5.6.6.

- `-DENABLED_LOCAL_INFILE=bool`

Whether to enable `LOCAL` capability in the client library for `LOAD DATA INFILE`.

This option controls client-side `LOCAL` capability, but the capability can be set on the server side at server startup with the `--local-infile` [496] option. See [Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”](#).

- `-DENABLED_PROFILING=bool`

Whether to enable query profiling code (for the `SHOW PROFILE` and `SHOW PROFILES` statements).

- `-DIGNORE_AIO_CHECK=bool`

If the `-DBUILD_CONFIG=mysql_release` [131] option is given on Linux, the `libaio` library must be linked in by default. If you do not have `libaio` or do not want to install it, you can suppress the check for it by specifying `-DIGNORE_AIO_CHECK=1` [135]. This option was added in MySQL 5.6.1.

- `-DMYSQL_MAINTAINER_MODE=bool`

Whether to enable a MySQL maintainer-specific development environment. If enabled, this option causes compiler warnings to become errors.

- `-DMYSQL_PROJECT_NAME=name`

For Windows or Mac OS X, the project name to incorporate into the project file name. This option was added in MySQL 5.6.5.

- `-DMYSQL_TCP_PORT=port_num`

The port number on which the server listens for TCP/IP connections. The default is 3306.

This value can be set at server startup with the `--port` [437] option.

- `-DMYSQL_UNIX_ADDR=file_name`

The Unix socket file path on which the server listens for socket connections. This must be an absolute path name. The default is `/tmp/mysql.sock`.

This value can be set at server startup with the `--socket` [442] option.

- `-DWITH_DEBUG=bool`

Whether to include debugging support.

Configuring MySQL with debugging support enables you to use the `--debug="d,parser_debug"` [418] option when you start the server. This causes the Bison parser that is used to process SQL statements to dump a parser trace to the server's standard error output. Typically, this output is written to the error log.

- `-DWITH_DEFAULT_COMPILER_OPTIONS=bool`

Whether to use the flags from `cmake/build_configurations/compiler_options.cmake`. This option was added in MySQL 5.6.6.

- `-DWITH_DEFAULT_FEATURE_SET=bool`

Whether to use the flags from `cmake/build_configurations/feature_set.cmake`. This option was added in MySQL 5.6.6.

- `-DWITH_EMBEDDED_SERVER=bool`

Whether to build the `libmysqld` embedded server library.

- `-DWITH_EXTRA_CHARSETS=name`

Which extra character sets to include:

- `all`: All character sets. This is the default.
- `complex`: Complex character sets.
- `none`: No extra character sets.

- `-DWITH_LIBEDIT=bool`

Whether to use the `libedit` library bundled with the distribution.

- `-DWITH_LIBWRAP=bool`

Whether to include `libwrap` (TCP wrappers) support.

- `-DWITH_READLINE=bool`

Whether to use the `readline` library bundled with the distribution. This option was removed in MySQL 5.6.5 because `readline` is no longer bundled.

- `-DWITH_SSL=ssl_type`

The type of SSL support to include, if any:

- `no`: No SSL support. This is the default before MySQL 5.6.6. As of 5.6.6, this is no longer a permitted value and the default is `bundled`.
- `yes`: Use the system SSL library if present, else the library bundled with the distribution.
- `bundled`: Use the SSL library bundled with the distribution. This is the default as of MySQL 5.6.6.
- `system`: Use the system SSL library.

- `path_name`: As of MySQL 5.6.7, it is permitted to specify the path name to the OpenSSL installation to use. This can be useful instead of a value of `system` when CMake detects an older or incorrect installed OpenSSL version. (Another permitted way to do the same thing is to set the `CMAKE_PREFIX_PATH` option to `path_name`.)

For information about using SSL support, see [Section 6.3.9, “Using SSL for Secure Connections”](#).

- `-DWITH_ZLIB=zlib_type`

Some features require that the server be built with compression library support, such as the `COMPRESS()` [1261] and `UNCOMPRESS()` [1265] functions, and compression of the client/server protocol. The `WITH_ZLIB` [137] indicates the source of `zlib` support:

- `bundled`: Use the `zlib` library bundled with the distribution.
- `system`: Use the system `zlib` library. This is the default.
- `-DWITH_UNIXODBC=1`

Enables unixODBC support, for Connector/ODBC.

## Compiler Flags

To specify compiler flags, set the `CFLAGS` and `CXXFLAGS` environment variables before running CMake. Example:

```
shell> CFLAGS=-DHAVE_EMBEDDED_PRIVILEGE_CONTROL
shell> CXXFLAGS=-DHAVE_EMBEDDED_PRIVILEGE_CONTROL
shell> export CFLAGS CXXFLAGS
shell> cmake [options]
```

The following flags control configuration features:

- `HAVE_EMBEDDED_PRIVILEGE_CONTROL`

By default, authentication for connections to the embedded server is disabled. To enable connection authentication, define this flag.

## 2.9.5. Dealing with Problems Compiling MySQL

The solution to many problems involves reconfiguring. If you do reconfigure, take note of the following:

- If CMake is run after it has previously been run, it may use information that was gathered during its previous invocation. This information is stored in `CMakeCache.txt`. When CMake starts up, it looks for that file and reads its contents if it exists, on the assumption that the information is still correct. That assumption is invalid when you reconfigure.
- Each time you run CMake, you must run `make` again to recompile. However, you may want to remove old object files from previous builds first because they were compiled using different configuration options.

To prevent old object files or configuration information from being used, run these commands on Unix before re-running CMake:

```
shell> make clean
shell> rm CMakeCache.txt
```

Or, on Windows:

```
shell> devenv MySQL.sln /clean
shell> del CMakeCache.txt
```

If you build out of the source tree, remove and recreate your build directory before re-running CMake. For instructions on building outside of the source tree, see [How to Build MySQL Server with CMake](#).

On some systems, warnings may occur due to differences in system include files. The following list describes other problems that have been found to occur most often when compiling MySQL:

- To define flags to be used by your C or C++ compilers, specify them using the `CFLAGS` and `CXXFLAGS` environment variables. You can also specify the compiler names this way using `CC` and `CXX`. For example:

```
shell> CC=gcc
shell> CFLAGS=-O3
shell> CXX=g++
shell> CXXFLAGS=-O3
shell> export CC CFLAGS CXX CXXFLAGS
```

To see what flags you might need to specify, invoke `mysql_config` with the `--cflags` [373] and `--cxxflags` [373] options.

- If compilation fails, check whether the `MYSQL_MAINTAINER_MODE` [135] option is enabled. This mode causes compiler warnings to become errors, so disabling it may enable compilation to proceed.
- If your compile fails with errors such as any of the following, you must upgrade your version of `make` to GNU `make`:

```
make: Fatal error in reader: Makefile, line 18:
Badly formed macro assignment
```

Or:

```
make: file `Makefile' line 18: Must be a separator (:
```

Or:

```
pthread.h: No such file or directory
```

Solaris and FreeBSD are known to have troublesome `make` programs.

GNU `make` 3.75 is known to work.

- The `sql_yacc.cc` file is generated from `sql_yacc.yy`. Normally, the build process does not need to create `sql_yacc.cc` because MySQL comes with a pregenerated copy. However, if you do need to re-create it, you might encounter this error:

```
"sql_yacc.yy", line xxx fatal: default action causes potential...
```

This is a sign that your version of `yacc` is deficient. You probably need to install `bison` (the GNU version of `yacc`) and use that instead.

Versions of `bison` older than 1.75 may report this error:

```
sql_yacc.yy:#####: fatal error: maximum table size (32767) exceeded
```

The maximum table size is not actually exceeded; the error is caused by bugs in older versions of `bison`.

For information about acquiring or updating tools, see the system requirements in [Section 2.9, "Installing MySQL from Source"](#).

## 2.9.6. MySQL Configuration and Third-Party Tools

Third-party tools that need to determine the MySQL version from the MySQL source can read the `VERSION` file in the top-level source directory. The file lists the pieces of the version separately. For example, if the version is 5.5.8, the file looks like this:

```
MYSQL_VERSION_MAJOR=5
```

```
MYSQL_VERSION_MINOR=5
MYSQL_VERSION_PATCH=8
MYSQL_VERSION_EXTRA=
```

If the source is not for a General Availability (GA) release, the `MYSQL_VERSION_EXTRA` value will be nonempty. For example, the value for a Release Candidate release would look like this:

```
MYSQL_VERSION_EXTRA=rc
```

To construct a five-digit number from the version components, use this formula:

```
MYSQL_VERSION_MAJOR*10000 + MYSQL_VERSION_MINOR*100 + MYSQL_VERSION_PATCH
```

## 2.10. Postinstallation Setup and Testing

After installing MySQL, there are some issues that you should address. For example, on Unix, you should initialize the data directory and create the MySQL grant tables. On all platforms, an important security concern is that the initial accounts in the grant tables have no passwords. You should assign passwords to prevent unauthorized access to the MySQL server. Optionally, you can create time zone tables to enable recognition of named time zones.

The following sections include postinstallation procedures that are specific to Windows systems and to Unix systems. Another section, [Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”](#), applies to all platforms; it describes what to do if you have trouble getting the server to start. [Section 2.10.2, “Securing the Initial MySQL Accounts”](#), also applies to all platforms. You should follow its instructions to make sure that you have properly protected your MySQL accounts by assigning passwords to them.

When you are ready to create additional user accounts, you can find information on the MySQL access control system and account management in [Section 6.2, “The MySQL Access Privilege System”](#), and [Section 6.3, “MySQL User Account Management”](#).

### 2.10.1. Unix Postinstallation Procedures

After installing MySQL on Unix, you must initialize the grant tables, start the server, and make sure that the server works satisfactorily. You may also wish to arrange for the server to be started and stopped automatically when your system starts and stops. You should also assign passwords to the accounts in the grant tables.

On Unix, the grant tables are set up by the `mysql_install_db` program. For some installation methods, this program is run for you automatically if an existing database cannot be found.

- If you install MySQL on Linux using RPM distributions, the server RPM runs `mysql_install_db`.
- Using the native packaging system on many platforms, including Debian Linux, Ubuntu Linux, Gentoo Linux and others, the `mysql_install_db` command is run for you.
- If you install MySQL on Mac OS X using a PKG distribution, the installer runs `mysql_install_db`.

For other platforms and installation types, including generic binary and source installs, you will need to run `mysql_install_db` yourself.

The following procedure describes how to initialize the grant tables (if that has not previously been done) and start the server. It also suggests some commands that you can use to test whether the server is accessible and working properly. For information about starting and stopping the server automatically, see [Section 2.10.1.2, “Starting and Stopping MySQL Automatically”](#).

After you complete the procedure and have the server running, you should assign passwords to the accounts created by `mysql_install_db` and perhaps restrict access to test databases. For instructions, see [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

In the examples shown here, the server runs under the user ID of the `mysql` login account. This assumes that such an account exists. Either create the account if it does not exist, or substitute the name of a different existing login account that you plan to use for running the server. For information

about creating the account, see [Creating a `mysql` System User and Group](#), in [Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”](#).

1. Change location into the top-level directory of your MySQL installation, represented here by `BASEDIR`:

```
shell> cd BASEDIR
```

`BASEDIR` is the installation directory for your MySQL instance. It is likely to be something like `/usr/local/mysql` or `/usr/local`. The following steps assume that you have changed location to this directory.

You will find several files and subdirectories in the `BASEDIR` directory. The most important for installation purposes are the `bin` and `scripts` subdirectories:

- The `bin` directory contains client programs and the server. You should add the full path name of this directory to your `PATH` environment variable so that your shell finds the MySQL programs properly. See [Section 2.12, “Environment Variables”](#).
  - The `scripts` directory contains the `mysql_install_db` script used to initialize the `mysql` database containing the grant tables that store the server access permissions.
2. If necessary, ensure that the distribution contents are accessible to `mysql`. If you installed the distribution as `mysql`, no further action is required. If you installed the distribution as `root`, its contents will be owned by `root`. Change its ownership to `mysql` by executing the following commands as `root` in the installation directory. The first command changes the owner attribute of the files to the `mysql` user. The second changes the group attribute to the `mysql` group.

```
shell> chown -R mysql .
shell> chgrp -R mysql .
```

3. If necessary, run the `mysql_install_db` program to set up the initial MySQL grant tables containing the privileges that determine how users are permitted to connect to the server. You will need to do this if you used a distribution type for which the installation procedure does not run the program for you.

```
shell> scripts/mysql_install_db --user=mysql
```

Typically, `mysql_install_db` needs to be run only the first time you install MySQL, so you can skip this step if you are upgrading an existing installation. However, `mysql_install_db` does not overwrite any existing privilege tables, so it should be safe to run in any circumstances.

It might be necessary to specify other options such as `--basedir` [\[237\]](#) or `--datadir` [\[237\]](#) if `mysql_install_db` does not identify the correct locations for the installation directory or data directory. For example:

```
shell> scripts/mysql_install_db --user=mysql \
    --basedir=/opt/mysql/mysql \
    --datadir=/opt/mysql/mysql/data
```

The `mysql_install_db` script creates the server's data directory with `mysql` as the owner. Under the data directory, it creates directories for the `mysql` database that holds the grant tables and the `test` database that you can use to test MySQL. The script also creates privilege table entries for `root` and anonymous-user accounts. The accounts have no passwords initially. [Section 2.10.2, “Securing the Initial MySQL Accounts”](#), describes the initial privileges. Briefly, these privileges permit the MySQL `root` user to do anything, and permit anybody to create or use databases with a name of `test` or starting with `test_`. See [Section 6.2, “The MySQL Access Privilege System”](#), for a complete listing and description of the grant tables.

For a more secure installation, invoke `mysql_install_db` with the `--random-passwords` option. This causes it to assign a random password to the MySQL `root` accounts, set the “password expired” flag for those accounts, and remove the anonymous-user MySQL accounts. For

additional details, see [Section 4.4.3, “mysql\\_install\\_db — Initialize MySQL Data Directory”](#). (Install operations using RPMs for Unbreakable Linux Network are unaffected because they do not use `mysql_install_db`.)

It is important to make sure that the database directories and files are owned by the `mysql` login account so that the server has read and write access to them when you run it later. To ensure this if you run `mysql_install_db` as `root`, include the `--user [238]` option as shown. Otherwise, you should execute the script while logged in as `mysql`, in which case you can omit the `--user [238]` option from the command.

If you do not want to have the `test` database, you can remove it after starting the server, using the instructions in [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

If you have trouble with `mysql_install_db` at this point, see [Section 2.10.1.1, “Problems Running mysql\\_install\\_db”](#).

- Most of the MySQL installation can be owned by `root` if you like. The exception is that the data directory must be owned by `mysql`. To accomplish this, run the following commands as `root` in the installation directory:

```
shell> chown -R root .
shell> chown -R mysql data
```

- If the plugin directory (the directory named by the `plugin_dir [525]` system variable) is writable by the server, it may be possible for a user to write executable code to a file in the directory using `SELECT ... INTO DUMPFILE`. This can be prevented by making `plugin_dir [525]` read only to the server or by setting `--secure-file-priv [438]` to a directory where `SELECT` writes can be made safely.
- If you installed MySQL using a source distribution, you may want to optionally copy one of the provided configuration files from the `support-files` directory into your `/etc` directory. There are different sample configuration files for different use cases, server types, and CPU and RAM configurations. If you want to use one of these standard files, you should copy it to `/etc/my.cnf`, or `/etc/mysql/my.cnf` and edit and check the configuration before starting your MySQL server for the first time.

If you do not copy one of the standard configuration files, the MySQL server will be started with the default settings.

If you want MySQL to start automatically when you boot your machine, you can copy `support-files/mysql.server` to the location where your system has its startup files. More information can be found in the `mysql.server` script itself, and in [Section 2.10.1.2, “Starting and Stopping MySQL Automatically”](#).

- Start the MySQL server:

```
shell> bin/mysqld_safe --user=mysql &
```

It is important that the MySQL server be run using an unprivileged (non-`root`) login account. To ensure this if you run `mysqld_safe` as `root`, include the `--user [229]` option as shown. Otherwise, you should execute the script while logged in as `mysql`, in which case you can omit the `--user [229]` option from the command.

For further instructions for running MySQL as an unprivileged user, see [Section 6.1.5, “How to Run MySQL as a Normal User”](#).

If the command fails immediately and prints `mysqld ended`, look for information in the error log (which by default is the `host_name.err` file in the data directory).

If you neglected to create the grant tables by running `mysql_install_db` before proceeding to this step, the following message appears in the error log file when you start the server:



```
mysqld: Can't find file: 'host.frm'
```

This error also occurs if you run `mysql_install_db` as `root` without the `--user` [238] option. Remove the `data` directory and run `mysql_install_db` with the `--user` [238] option as described previously.

If you have other problems starting the server, see [Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”](#). For more information about `mysqld_safe`, see [Section 4.3.2, “mysqld\\_safe — MySQL Server Startup Script”](#).

- Use `mysqladmin` to verify that the server is running. The following commands provide simple tests to check whether the server is up and responding to connections:

```
shell> bin/mysqladmin version
shell> bin/mysqladmin variables
```

The output from `mysqladmin version` varies slightly depending on your platform and version of MySQL, but should be similar to that shown here:

```
shell> bin/mysqladmin version
mysqladmin Ver 14.12 Distrib 5.6.12, for pc-linux-gnu on i686
...

Server version          5.6.12
Protocol version        10
Connection              Localhost via UNIX socket
UNIX socket             /var/lib/mysql/mysql.sock
Uptime:                 14 days 5 hours 5 min 21 sec

Threads: 1  Questions: 366  Slow queries: 0
Opens: 0  Flush tables: 1  Open tables: 19
Queries per second avg: 0.000
```

To see what else you can do with `mysqladmin`, invoke it with the `--help` [270] option.

- Verify that you can shut down the server:

```
shell> bin/mysqladmin -u root shutdown
```

- Verify that you can start the server again. Do this by using `mysqld_safe` or by invoking `mysqld` directly. For example:

```
shell> bin/mysqld_safe --user=mysql &
```

If `mysqld_safe` fails, see [Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”](#).

- Run some simple tests to verify that you can retrieve information from the server. The output should be similar to what is shown here:

```
shell> bin/mysqlshow
+-----+
|          Databases          |
+-----+
| information_schema |
| mysql              |
| test               |
+-----+

shell> bin/mysqlshow mysql
Database: mysql
+-----+
|          Tables          |
+-----+
| columns_priv |
| db           |
| event       |
| func        |
```

```

| help_category
| help_keyword
| help_relation
| help_topic
| host
| plugin
| proc
| procs_priv
| servers
| tables_priv
| time_zone
| time_zone_leap_second
| time_zone_name
| time_zone_transition
| time_zone_transition_type
| user
+-----+

shell> bin/mysql -e "SELECT Host,Db,User FROM db" mysql
+-----+-----+-----+
| host | db      | user |
+-----+-----+-----+
| %    | test   |      |
| %    | test_% |      |
+-----+-----+-----+

```

12. There is a benchmark suite in the `sql-bench` directory (under the MySQL installation directory) that you can use to compare how MySQL performs on different platforms. The benchmark suite is written in Perl. It requires the Perl DBI module that provides a database-independent interface to the various databases, and some other additional Perl modules:

```

DBI
DBD::mysql
Data::Dumper
Data::ShowTable

```

These modules can be obtained from CPAN (<http://www.cpan.org/>). See also [Section 2.13.1](#), “Installing Perl on Unix”.

The `sql-bench/Results` directory contains the results from many runs against different databases and platforms. To run all tests, execute these commands:

```

shell> cd sql-bench
shell> perl run-all-tests

```

If you do not have the `sql-bench` directory, you probably installed MySQL using RPM files other than the source RPM. (The source RPM includes the `sql-bench` benchmark directory.) In this case, you must first install the benchmark suite before you can use it. There are separate benchmark RPM files named `mysql-bench-VERSION.i386.rpm` that contain benchmark code and data.

If you have a source distribution, there are also tests in its `tests` subdirectory that you can run. For example, to run `auto_increment.tst`, execute this command from the top-level directory of your source distribution:

```

shell> mysql -vfv test < ./tests/auto_increment.tst

```

The expected result of the test can be found in the `./tests/auto_increment.res` file.

13. At this point, you should have the server running. However, none of the initial MySQL accounts have a password, and the server permits permissive access to test databases. To tighten security, follow the instructions in [Section 2.10.2](#), “Securing the Initial MySQL Accounts”.

The MySQL 5.6 installation procedure creates time zone tables in the `mysql` database but does not populate them. To do so, use the instructions in [Section 10.6](#), “MySQL Server Time Zone Support”.

To make it more convenient to invoke programs installed in the `bin` directory under the installation directory, you can add that directory to your `PATH` environment variable setting. That enables you to run a program by typing only its name, not its entire path name. See [Section 4.2.4, “Setting Environment Variables”](#).

You can set up new accounts using the `bin/mysql_setpermission` script if you install the `DBI` and `DBD::mysql` Perl modules. See [Section 4.6.14, “mysql\\_setpermission — Interactively Set Permissions in Grant Tables”](#). For Perl module installation instructions, see [Section 2.13, “Perl Installation Notes”](#).

If you would like to use `mysqlaccess` and have the MySQL distribution in some nonstandard location, you must change the location where `mysqlaccess` expects to find the `mysql` client. Edit the `bin/mysqlaccess` script at approximately line 18. Search for a line that looks like this:

```
$MYSQL = '/usr/local/bin/mysql'; # path to mysql executable
```

Change the path to reflect the location where `mysql` actually is stored on your system. If you do not do this, a [Broken pipe](#) error will occur when you run `mysqlaccess`.

### 2.10.1.1. Problems Running `mysql_install_db`

The purpose of the `mysql_install_db` script is to generate new MySQL privilege tables. It does not overwrite existing MySQL privilege tables, and it does not affect any other data.

If you want to re-create your privilege tables, first stop the `mysqld` server if it is running. Then rename the `mysql` directory under the data directory to save it, and then run `mysql_install_db`. Suppose that your current directory is the MySQL installation directory and that `mysql_install_db` is located in the `bin` directory and the data directory is named `data`. To rename the `mysql` database and re-run `mysql_install_db`, use these commands.

```
shell> mv data/mysql data/mysql.old
shell> scripts/mysql_install_db --user=mysql
```

When you run `mysql_install_db`, you might encounter the following problems:

- **`mysql_install_db` fails to install the grant tables**

You may find that `mysql_install_db` fails to install the grant tables and terminates after displaying the following messages:

```
Starting mysqld daemon with databases from XXXXXX
mysqld ended
```

In this case, you should examine the error log file very carefully. The log should be located in the directory `XXXXXX` named by the error message and should indicate why `mysqld` did not start. If you do not understand what happened, include the log when you post a bug report. See [Section 1.7, “How to Report Bugs or Problems”](#).

- **There is a `mysqld` process running**

This indicates that the server is running, in which case the grant tables have probably been created already. If so, there is no need to run `mysql_install_db` at all because it needs to be run only once (when you install MySQL the first time).

- **Installing a second `mysqld` server does not work when one server is running**

This can happen when you have an existing MySQL installation, but want to put a new installation in a different location. For example, you might have a production installation, but you want to create a second installation for testing purposes. Generally the problem that occurs when you try to run a second server is that it tries to use a network interface that is in use by the first server. In this case, you should see one of the following error messages:

```
Can't start server: Bind on TCP/IP port:
Address already in use
```

```
Can't start server: Bind on unix socket...
```

For instructions on setting up multiple servers, see [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).

- **You do not have write access to the `/tmp` directory**

If you do not have write access to create temporary files or a Unix socket file in the default location (the `/tmp` directory) or the `TMP_DIR` environment variable, if it has been set, an error occurs when you run `mysql_install_db` or the `mysqld` server.

You can specify different locations for the temporary directory and Unix socket file by executing these commands prior to starting `mysql_install_db` or `mysqld`, where `some_tmp_dir` is the full path name to some directory for which you have write permission:

```
shell> TMPDIR=/some_tmp_dir/
shell> MYSQL_UNIX_PORT=/some_tmp_dir/mysql.sock
shell> export TMPDIR MYSQL_UNIX_PORT
```

Then you should be able to run `mysql_install_db` and start the server with these commands:

```
shell> scripts/mysql_install_db --user=mysql
shell> bin/mysqld_safe --user=mysql &
```

If `mysql_install_db` is located in the `scripts` directory, modify the first command to `scripts/mysql_install_db`.

See [Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”](#), and [Section 2.12, “Environment Variables”](#).

There are some alternatives to running the `mysql_install_db` script provided in the MySQL distribution:

- If you want the initial privileges to be different from the standard defaults, you can modify `mysql_install_db` before you run it. However, it is preferable to use `GRANT` and `REVOKE` to change the privileges *after* the grant tables have been set up. In other words, you can run `mysql_install_db`, and then use `mysql -u root mysql` to connect to the server as the MySQL `root` user so that you can issue the necessary `GRANT` and `REVOKE` statements.

If you want to install MySQL on several machines with the same privileges, you can put the `GRANT` and `REVOKE` statements in a file and execute the file as a script using `mysql` after running `mysql_install_db`. For example:

```
shell> scripts/mysql_install_db --user=mysql
shell> bin/mysql -u root < your_script_file
```

By doing this, you can avoid having to issue the statements manually on each machine.

- It is possible to re-create the grant tables completely after they have previously been created. You might want to do this if you are just learning how to use `GRANT` and `REVOKE` and have made so many modifications after running `mysql_install_db` that you want to wipe out the tables and start over.

To re-create the grant tables, remove all the `.frm`, `.MYI`, and `.MYD` files in the `mysql` database directory. Then run the `mysql_install_db` script again.

- You can start `mysqld` manually using the `--skip-grant-tables` [\[439\]](#) option and add the privilege information yourself using `mysql`:

```
shell> bin/mysqld_safe --user=mysql --skip-grant-tables &
shell> bin/mysql mysql
```

From `mysql`, manually execute the SQL commands contained in `mysql_install_db`. Make sure that you run `mysqladmin flush-privileges` or `mysqladmin reload` afterward to tell the server to reload the grant tables.

Note that by not using `mysql_install_db`, you not only have to populate the grant tables manually, you also have to create them first.

### 2.10.1.2. Starting and Stopping MySQL Automatically

Generally, you start the `mysqld` server in one of these ways:

- Invoke `mysqld` directly. This works on any platform.
- Run the MySQL server as a Windows service. The service can be set to start the server automatically when Windows starts, or as a manual service that you start on request. For instructions, see [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).
- Invoke `mysqld_safe`, which tries to determine the proper options for `mysqld` and then runs it with those options. This script is used on Unix and Unix-like systems. See [Section 4.3.2, “mysqld\\_safe — MySQL Server Startup Script”](#).
- Invoke `mysql.server`. This script is used primarily at system startup and shutdown on systems that use System V-style run directories (that is, `/etc/init.d` and run-level specific directories), where it usually is installed under the name `mysql`. The `mysql.server` script starts the server by invoking `mysqld_safe`. See [Section 4.3.3, “mysql.server — MySQL Server Startup Script”](#).
- On Mac OS X, install a separate MySQL Startup Item package to enable the automatic startup of MySQL on system startup. The Startup Item starts the server by invoking `mysql.server`. See [Section 2.4.3, “Installing the MySQL Startup Item”](#), for details. A MySQL Preference Pane also provides control for starting and stopping MySQL through the System Preferences, see [Section 2.4.4, “Installing and Using the MySQL Preference Pane”](#).
- Use the Solaris/OpenSolaris service management framework (SMF) system to initiate and control MySQL startup. For more information, see [Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”](#).

The `mysqld_safe` and `mysql.server` scripts, Windows server, Solaris/OpenSolaris SMF, and the Mac OS X Startup Item (or MySQL Preference Pane) can be used to start the server manually, or automatically at system startup time. `mysql.server` and the Startup Item also can be used to stop the server.

To start or stop the server manually using the `mysql.server` script, invoke it with `start` or `stop` arguments:

```
shell> mysql.server start
shell> mysql.server stop
```

Before `mysql.server` starts the server, it changes location to the MySQL installation directory, and then invokes `mysqld_safe`. If you want the server to run as some specific user, add an appropriate `user` option to the `[mysqld]` group of the `/etc/my.cnf` option file, as shown later in this section. (It is possible that you will need to edit `mysql.server` if you've installed a binary distribution of MySQL in a nonstandard location. Modify it to change location into the proper directory before it runs `mysqld_safe`. If you do this, your modified version of `mysql.server` may be overwritten if you upgrade MySQL in the future, so you should make a copy of your edited version that you can reinstall.)

`mysql.server stop` stops the server by sending a signal to it. You can also stop the server manually by executing `mysqladmin shutdown`.

To start and stop MySQL automatically on your server, you need to add start and stop commands to the appropriate places in your `/etc/rc*` files.

If you use the Linux server RPM package (`MySQL-server-VERSION.rpm`), or a native Linux package installation, the `mysql.server` script may be installed in the `/etc/init.d` directory with the name `mysql`. See [Section 2.5.1, “Installing MySQL from RPM Packages on Linux”](#), for more information on the Linux RPM packages.

Some vendors provide RPM packages that install a startup script under a different name such as `mysqld`.

If you install MySQL from a source distribution or using a binary distribution format that does not install `mysql.server` automatically, you can install it manually. The script can be found in the `support-files` directory under the MySQL installation directory or in a MySQL source tree.

To install `mysql.server` manually, copy it to the `/etc/init.d` directory with the name `mysql`, and then make it executable. Do this by changing location into the appropriate directory where `mysql.server` is located and executing these commands:

```
shell> cp mysql.server /etc/init.d/mysql
shell> chmod +x /etc/init.d/mysql
```



#### Note

Older Red Hat systems use the `/etc/rc.d/init.d` directory rather than `/etc/init.d`. Adjust the preceding commands accordingly. Alternatively, first create `/etc/init.d` as a symbolic link that points to `/etc/rc.d/init.d`:

```
shell> cd /etc
shell> ln -s rc.d/init.d .
```

After installing the script, the commands needed to activate it to run at system startup depend on your operating system. On Linux, you can use `chkconfig`:

```
shell> chkconfig --add mysql
```

On some Linux systems, the following command also seems to be necessary to fully enable the `mysql` script:

```
shell> chkconfig --level 345 mysql on
```

On FreeBSD, startup scripts generally should go in `/usr/local/etc/rc.d/`. The `rc(8)` manual page states that scripts in this directory are executed only if their basename matches the `*.sh` shell file name pattern. Any other files or directories present within the directory are silently ignored. In other words, on FreeBSD, you should install the `mysql.server` script as `/usr/local/etc/rc.d/mysql.server.sh` to enable automatic startup.

As an alternative to the preceding setup, some operating systems also use `/etc/rc.local` or `/etc/init.d/boot.local` to start additional services on startup. To start up MySQL using this method, you could append a command like the one following to the appropriate startup file:

```
/bin/sh -c 'cd /usr/local/mysql; ./bin/mysqld_safe --user=mysql &'
```

For other systems, consult your operating system documentation to see how to install startup scripts.

You can add options for `mysql.server` in a global `/etc/my.cnf` file. A typical `/etc/my.cnf` file might look like this:

```
[mysqld]
datadir=/usr/local/mysql/var
socket=/var/tmp/mysql.sock
port=3306
user=mysql

[mysql.server]
basedir=/usr/local/mysql
```

The `mysql.server` script supports the following options: `basedir`, `datadir`, and `pid-file`. If specified, they *must* be placed in an option file, not on the command line. `mysql.server` supports only `start` and `stop` as command-line arguments.

The following table shows which option groups the server and each startup script read from option files.

**Table 2.13. MySQL Startup scripts and supported server option groups**

Script	Option Groups
<code>mysqld</code>	<code>[mysqld]</code> , <code>[server]</code> , <code>[mysqld-major_version]</code>
<code>mysqld_safe</code>	<code>[mysqld]</code> , <code>[server]</code> , <code>[mysqld_safe]</code>
<code>mysql.server</code>	<code>[mysqld]</code> , <code>[mysql.server]</code> , <code>[server]</code>

`[mysqld-major_version]` means that groups with names like `[mysqld-5.5]` and `[mysqld-5.6]` are read by servers having versions 5.5.x, 5.6.x, and so forth. This feature can be used to specify options that can be read only by servers within a given release series.

For backward compatibility, `mysql.server` also reads the `[mysql_server]` group and `mysqld_safe` also reads the `[safe_mysqld]` group. However, you should update your option files to use the `[mysql.server]` and `[mysqld_safe]` groups instead when using MySQL 5.6.

For more information on MySQL configuration files and their structure and contents, see [Section 4.2.3.3, “Using Option Files”](#).

### 2.10.1.3. Starting and Troubleshooting the MySQL Server

This section provides troubleshooting suggestions for problems starting the server on Unix. If you are using Windows, see [Section 2.3.5, “Troubleshooting a Microsoft Windows MySQL Server Installation”](#).

If you have problems starting the server, here are some things to try:

- Check the [error log](#) to see why the server does not start. Log files are located in the [data directory](#) (typically `C:\Program Files\MySQL\MySQL Server 5.6\data` on Windows, `/usr/local/mysql/data` for a Unix binary distribution, and `/usr/local/var` for a Unix source distribution). Look in the data directory for files with names of the form `host_name.err` and `host_name.log`, where `host_name` is the name of your server host. Then examine the last few lines of these files. On Unix, use `tail` to display them:

```
shell> tail host_name.err
shell> tail host_name.log
```

- Specify any special options needed by the storage engines you are using. You can create a `my.cnf` file and specify startup options for the engines that you plan to use. If you are going to use storage engines that support transactional tables ([InnoDB](#), [NDB](#)), be sure that you have them configured the way you want before starting the server. If you are using [InnoDB](#) tables, see [Section 14.2.1.2, “Configuring InnoDB”](#) for guidelines and [Section 14.2.6, “InnoDB Startup Options and System Variables”](#) for option syntax.

Although storage engines use default values for options that you omit, Oracle recommends that you review the available options and specify explicit values for any options whose defaults are not appropriate for your installation.

- Make sure that the server knows where to find the [data directory](#). The `mysqld` server uses this directory as its current directory. This is where it expects to find databases and where it expects to write log files. The server also writes the pid (process ID) file in the data directory.

The default data directory location is hardcoded when the server is compiled. To determine what the default path settings are, invoke `mysqld` with the `--verbose` [\[447\]](#) and `--help` [\[413\]](#) options. If the data directory is located somewhere else on your system, specify that location with the `--datadir` [\[418\]](#) option to `mysqld` or `mysqld_safe`, on the command line or in an option file. Otherwise, the server will not work properly. As an alternative to the `--datadir` [\[418\]](#) option, you can specify `mysqld` the location of the base directory under which MySQL is installed with the `--basedir` [\[414\]](#), and `mysqld` looks for the `data` directory there.

To check the effect of specifying path options, invoke `mysqld` with those options followed by the `--verbose` [\[447\]](#) and `--help` [\[413\]](#) options. For example, if you change location into the directory



where `mysqld` is installed and then run the following command, it shows the effect of starting the server with a base directory of `/usr/local`:

```
shell> ./mysqld --basedir=/usr/local --verbose --help
```

You can specify other options such as `--datadir` [418] as well, but `--verbose` [447] and `--help` [413] must be the last options.

Once you determine the path settings you want, start the server without `--verbose` [447] and `--help` [413].

If `mysqld` is currently running, you can find out what path settings it is using by executing this command:

```
shell> mysqladmin variables
```

Or:

```
shell> mysqladmin -h host_name variables
```

`host_name` is the name of the MySQL server host.

- Make sure that the server can access the [data directory](#). The ownership and permissions of the data directory and its contents must allow the server to read and modify them.

If you get [Errcode 13](#) (which means [Permission denied](#)) when starting `mysqld`, this means that the privileges of the data directory or its contents do not permit server access. In this case, you change the permissions for the involved files and directories so that the server has the right to use them. You can also start the server as `root`, but this raises security issues and should be avoided.

On Unix, change location into the data directory and check the ownership of the data directory and its contents to make sure the server has access. For example, if the data directory is `/usr/local/mysql/var`, use this command:

```
shell> ls -la /usr/local/mysql/var
```

If the data directory or its files or subdirectories are not owned by the login account that you use for running the server, change their ownership to that account. If the account is named `mysql`, use these commands:

```
shell> chown -R mysql /usr/local/mysql/var
shell> chgrp -R mysql /usr/local/mysql/var
```

Even with correct ownership, MySQL might fail to start up if there is other security software running on your system that manages application access to various parts of the file system. In this case, reconfigure that software to enable `mysqld` to access the directories it uses during normal operation.

- Verify that the network interfaces the server wants to use are available.

If either of the following errors occur, it means that some other program (perhaps another `mysqld` server) is using the TCP/IP port or Unix socket file that `mysqld` is trying to use:

```
Can't start server: Bind on TCP/IP port: Address already in use
Can't start server: Bind on unix socket...
```

Use `ps` to determine whether you have another `mysqld` server running. If so, shut down the server before starting `mysqld` again. (If another server is running, and you really want to run multiple servers, you can find information about how to do so in [Section 5.6, "Running Multiple MySQL Instances on One Machine"](#).)

If no other server is running, execute the command `telnet your_host_name tcp_ip_port_number`. (The default MySQL port number is 3306.) Then press Enter a couple of times. If you do not get an error message like `telnet: Unable to connect to remote`

`host: Connection refused`, some other program is using the TCP/IP port that `mysqld` is trying to use. Track down what program this is and disable it, or tell `mysqld` to listen to a different port with the `--port [437]` option. In this case, specify the same non-default port number for client programs when connecting to the server using TCP/IP.

Another reason the port might be inaccessible is that you have a firewall running that blocks connections to it. If so, modify the firewall settings to permit access to the port.

If the server starts but you cannot connect to it, make sure that you have an entry in `/etc/hosts` that looks like this:

```
127.0.0.1 localhost
```

- If you cannot get `mysqld` to start, try to make a trace file to find the problem by using the `--debug [418]` option. See [MySQL Internals: Porting to Other Systems](#).

## 2.10.2. Securing the Initial MySQL Accounts

Part of the MySQL installation process is to set up the `mysql` database that contains the grant tables:

- Windows distributions contain preinitialized grant tables.
- On Unix, the `mysql_install_db` program populates the grant tables. Some installation methods run this program for you. Others require that you execute it manually. For details, see [Section 2.10.1, “Unix Postinstallation Procedures”](#).

The `mysql.user` grant table defines the initial MySQL user accounts and their access privileges:

- Some accounts have the user name `root`. These are superuser accounts that have all privileges and can do anything. The initial `root` account passwords are empty, so anyone can connect to the MySQL server as `root` *without a password* and be granted all privileges.
- On Windows, `root` accounts are created that permit connections from the local host only. Connections can be made by specifying the host name `localhost`, the IP address `127.0.0.1`, or the IPv6 address `::1`. If the user selects the **Enable root access from remote machines** option during installation, the Windows installer creates another `root` account that permits connections from any host.
- On Unix, each `root` account permits connections from the local host. Connections can be made by specifying the host name `localhost`, the IP address `127.0.0.1`, the IPv6 address `::1`, or the actual host name or IP address.

An attempt to connect to the host `127.0.0.1` normally resolves to the `localhost` account. However, this fails if the server is run with the `--skip-name-resolve [440]` option, so the `127.0.0.1` account is useful in that case. The `::1` account is used for IPv6 connections.

- Some accounts are for anonymous users. These have an empty user name. The anonymous accounts have no password, so anyone can use them to connect to the MySQL server.
- On Windows, there is one anonymous account that permits connections from the local host. Connections can be made by specifying a host name of `localhost`.
- On Unix, each anonymous account permits connections from the local host. Connections can be made by specifying a host name of `localhost` for one of the accounts, or the actual host name or IP address for the other.

To display which accounts exist in the `mysql.user` table and check whether their passwords are empty, use the following statement:

```
mysql> SELECT User, Host, Password FROM mysql.user;
+-----+-----+-----+
| User | Host | Password |
+-----+-----+-----+
```

root	localhost		
root	myhost.example.com		
root	127.0.0.1		
root	::1		
	localhost		
	myhost.example.com		

This output indicates that there are several `root` and anonymous-user accounts, none of which have passwords. The output might differ on your system, but the presence of accounts with empty passwords means that your MySQL installation is unprotected until you do something about it:

- You should assign a password to each MySQL `root` account.
- If you want to prevent clients from connecting as anonymous users without a password, you should either assign a password to each anonymous account or else remove the accounts.

In addition, the `mysql.db` table contains rows that permit all accounts to access the `test` database and other databases with names that start with `test_`. This is true even for accounts that otherwise have no special privileges such as the default anonymous accounts. This is convenient for testing but inadvisable on production servers. Administrators who want database access restricted only to accounts that have permissions granted explicitly for that purpose should remove these `mysql.db` table rows.

The following instructions describe how to set up passwords for the initial MySQL accounts, first for the `root` accounts, then for the anonymous accounts. The instructions also cover how to remove the anonymous accounts, should you prefer not to permit anonymous access at all, and describe how to remove permissive access to test databases. Replace `newpwd` in the examples with the password that you want to use. Replace `host_name` with the name of the server host. You can determine this name from the output of the preceding `SELECT` statement. For the output shown, `host_name` is `myhost.example.com`.



#### Note

For additional information about setting passwords, see [Section 6.3.5, “Assigning Account Passwords”](#). If you forget your `root` password after setting it, see [Section C.5.4.1, “How to Reset the Root Password”](#).

You might want to defer setting the passwords until later, to avoid the need to specify them while you perform additional setup or testing. However, be sure to set them before using your installation for production purposes.

To set up additional accounts, see [Section 6.3.2, “Adding User Accounts”](#).

## Assigning `root` Account Passwords

The `root` account passwords can be set several ways. The following discussion demonstrates three methods:

- Use the `SET PASSWORD` statement
- Use the `UPDATE` statement
- Use the `mysqladmin` command-line client program

To assign passwords using `SET PASSWORD`, connect to the server as `root` and issue a `SET PASSWORD` statement for each `root` account listed in the `mysql.user` table. Be sure to encrypt the password using the `PASSWORD()` [1263] function.

For Windows, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR 'root'@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'127.0.0.1' = PASSWORD('newpwd');
```

```
mysql> SET PASSWORD FOR 'root'@'::1' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@ '%' = PASSWORD('newpwd');
```

The last statement is unnecessary if the `mysql.user` table has no `root` account with a host value of `%`.

For Unix, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR 'root'@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'127.0.0.1' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'::1' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'host_name' = PASSWORD('newpwd');
```

You can also use a single statement that assigns a password to all `root` accounts by using `UPDATE` to modify the `mysql.user` table directly. This method works on any platform:

```
shell> mysql -u root
mysql> UPDATE mysql.user SET Password = PASSWORD('newpwd')
-> WHERE User = 'root';
mysql> FLUSH PRIVILEGES;
```

The `FLUSH` statement causes the server to reread the grant tables. Without it, the password change remains unnoticed by the server until you restart it.

To assign passwords to the `root` accounts using `mysqladmin`, execute the following commands:

```
shell> mysqladmin -u root password "newpwd"
shell> mysqladmin -u root -h host_name password "newpwd"
```

Those commands apply both to Windows and to Unix. The double quotation marks around the password are not always necessary, but you should use them if the password contains spaces or other characters that are special to your command interpreter.

The `mysqladmin` method of setting the `root` account passwords does not work for the `'root'@'127.0.0.1'` or `'root'@'::1'` account. Use the `SET PASSWORD` method shown earlier.

After the `root` passwords have been set, you must supply the appropriate password whenever you connect as `root` to the server. For example, to shut down the server with `mysqladmin`, use this command:

```
shell> mysqladmin -u root -p shutdown
Enter password: (enter root password here)
```

## Assigning Anonymous Account Passwords

The `mysql` commands in the following instructions include a `-p` option based on the assumption that you have set the `root` account passwords using the preceding instructions and must specify that password when connecting to the server.

To assign passwords to the anonymous accounts, connect to the server as `root`, then use either `SET PASSWORD` or `UPDATE`. Be sure to encrypt the password using the `PASSWORD()` [1263] function.

To use `SET PASSWORD` on Windows, do this:

```
shell> mysql -u root -p
Enter password: (enter root password here)
mysql> SET PASSWORD FOR ''@'localhost' = PASSWORD('newpwd');
```

To use `SET PASSWORD` on Unix, do this:

```
shell> mysql -u root -p
Enter password: (enter root password here)
mysql> SET PASSWORD FOR ''@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR ''@'host_name' = PASSWORD('newpwd');
```

To set the anonymous-user account passwords with a single `UPDATE` statement, do this (on any platform):

```

shell> mysql -u root -p
Enter password: (enter root password here)
mysql> UPDATE mysql.user SET Password = PASSWORD('newpwd')
-> WHERE User = '';
mysql> FLUSH PRIVILEGES;

```

The `FLUSH` statement causes the server to reread the grant tables. Without it, the password change remains unnoticed by the server until you restart it.

## Removing Anonymous Accounts

If you prefer to remove any anonymous accounts rather than assigning them passwords, do so as follows on Windows:

```

shell> mysql -u root -p
Enter password: (enter root password here)
mysql> DROP USER '@localhost';

```

On Unix, remove the anonymous accounts like this:

```

shell> mysql -u root -p
Enter password: (enter root password here)
mysql> DROP USER '@localhost';
mysql> DROP USER '@host_name';

```

## Securing Test Databases

By default, the `mysql.db` table contains rows that permit access by any user to the `test` database and other databases with names that start with `test_`. (These rows have an empty `User` column value, which for access-checking purposes matches any user name.) This means that such databases can be used even by accounts that otherwise possess no privileges. If you want to remove any-user access to test databases, do so as follows:

```

shell> mysql -u root -p
Enter password: (enter root password here)
mysql> DELETE FROM mysql.db WHERE Db LIKE 'test%';
mysql> FLUSH PRIVILEGES;

```

The `FLUSH` statement causes the server to reread the grant tables. Without it, the privilege change remains unnoticed by the server until you restart it.

With the preceding change, only users who have global database privileges or privileges granted explicitly for the `test` database can use it. However, if you do not want the database to exist at all, drop it:

```

mysql> DROP DATABASE test;

```



### Note

On Windows, you can also perform the process described in this section during installation with MySQL Installer (see [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#)). On other platforms, the MySQL distribution includes `mysql_secure_installation`, a command-line utility that automates much of the process of securing a MySQL installation. And MySQL Workbench is available on all platforms, and also offers the ability to manage user accounts (see [Chapter 24, MySQL Workbench](#)).

## 2.11. Upgrading or Downgrading MySQL

This section describes the steps to upgrade or downgrade a MySQL installation.

Upgrading is a common procedure, as you pick up bug fixes within the same MySQL release series or significant features between major MySQL releases. You perform this procedure first on some test systems to make sure everything works smoothly, and then on the production systems.

Downgrading is less common. Typically, you undo an upgrade because of some compatibility or performance issue that occurs on a production system, and was not uncovered during initial upgrade verification on the test systems. As with the upgrade procedure, perform and verify the downgrade procedure on some test systems first, before using it on a production system.

### 2.11.1. Upgrading MySQL

As a general rule, to upgrade from one release series to another, go to the next series rather than skipping a series. To upgrade from a release series previous to MySQL 5.5, upgrade to each successive release series in turn until you have reached MySQL 5.5, and then proceed with the upgrade to MySQL 5.6. For example, if you currently are running MySQL 5.1 and wish to upgrade to a newer series, upgrade to MySQL 5.5 first before upgrading to 5.6, and so forth. For information on upgrading to MySQL 5.5, see the *MySQL 5.5 Reference Manual*.

To upgrade to MySQL 5.6, use the items in the following checklist as a guide:

- Before any upgrade, back up your databases, including the `mysql` database that contains the grant tables. See [Section 7.2, “Database Backup Methods”](#).
- Read *all* the notes in [Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”](#). These notes enable you to identify upgrade issues that apply to your current MySQL installation. Some incompatibilities discussed in that section require your attention *before* upgrading. Others require some action *after* upgrading.
- Read the [Release Notes](#) as well, which provide information about features that are new in MySQL 5.6 or differ from those found in earlier MySQL releases.
- After upgrading to a new version of MySQL, run `mysql_upgrade` (see [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#)). This program checks your tables, and attempts to repair them if necessary. It also updates your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. (Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features.)

`mysql_upgrade` does not upgrade the contents of the help tables. For upgrade instructions, see [Section 5.1.10, “Server-Side Help”](#).

`mysql_upgrade` should not be used when the server is running with `--gtid-mode=ON` [2019], since it may make changes in nontransactional system tables in the `mysql` database, many of which are `MyISAM` and cannot be changed to use a different storage engine. See [GTID mode and mysql\\_upgrade](#).

- If you run MySQL Server on Windows, see [Section 2.3.6, “Upgrading MySQL on Windows”](#).
- If you use replication, see [Section 16.4.3, “Upgrading a Replication Setup”](#), for information on upgrading your replication setup.
- If you upgrade an installation originally produced by installing multiple RPM packages, it is best to upgrade all the packages, not just some. For example, if you previously installed the server and client RPMs, do not upgrade just the server RPM.
- If you have created a user-defined function (UDF) with a given name and upgrade MySQL to a version that implements a new built-in function with the same name, the UDF becomes inaccessible. To correct this, use `DROP FUNCTION` to drop the UDF, and then use `CREATE FUNCTION` to re-create the UDF with a different nonconflicting name. The same is true if the new version of MySQL implements a built-in function with the same name as an existing stored function. See [Section 9.2.4, “Function Name Parsing and Resolution”](#), for the rules describing how the server interprets references to different kinds of functions.

For upgrades between versions of a MySQL release series that has reached General Availability status, you can move the MySQL format files and data files between different versions on systems with the same architecture. For upgrades to a version of a MySQL release series that is in development status, that is not necessarily true. Use of development releases is at your own risk.



If you are cautious about using new versions, you can always rename your old `mysqld` before installing a newer one. For example, if you are using a version of MySQL 5.5 and want to upgrade to 5.6, rename your current server from `mysqld` to `mysqld-5.5`. If your new `mysqld` then does something unexpected, you can simply shut it down and restart with your old `mysqld`.

If problems occur, such as that the new `mysqld` server does not start or that you cannot connect without a password, verify that you do not have an old `my.cnf` file from your previous installation. You can check this with the `--print-defaults` [220] option (for example, `mysqld --print-defaults`). If this command displays anything other than the program name, you have an active `my.cnf` file that affects server or client operation.

If, after an upgrade, you experience problems with compiled client programs, such as `Commands out of sync` or unexpected core dumps, you probably have used old header or library files when compiling your programs. In this case, check the date for your `mysql.h` file and `libmysqlclient.a` library to verify that they are from the new MySQL distribution. If not, recompile your programs with the new headers and libraries. Recompile might also be necessary for programs compiled against the shared client library if the library major version number has changed (for example from `libmysqlclient.so.15` to `libmysqlclient.so.16`).

If your MySQL installation contains a large amount of data that might take a long time to convert after an in-place upgrade, you might find it useful to create a “dummy” database instance for assessing what conversions might be needed and the work involved to perform them. Make a copy of your MySQL instance that contains a full copy of the `mysql` database, plus all other databases without data. Run your upgrade procedure on this dummy instance to see what actions might be needed so that you can better evaluate the work involved when performing actual data conversion on your original database instance.

It is a good idea to rebuild and reinstall the Perl `DBD: :mysql` module whenever you install a new release of MySQL. The same applies to other MySQL interfaces as well, such as PHP `mysql` extensions and the Python `MySQLdb` module.

### 2.11.1.1. Upgrading from MySQL 5.5 to 5.6



#### Note

Beginning with MySQL 5.6.6, several MySQL Server parameters have defaults that differ from previous releases. See the notes regarding these changes later in this section, particularly regarding overriding them to preserve backward compatibility if that is a concern.



#### Note

It is good practice to back up your data before installing any new version of software. Although MySQL works very hard to ensure a high level of quality, protect your data by making a backup.

To upgrade to 5.6 from any previous version, MySQL recommends that you dump your tables with `mysqldump` before upgrading and reload the dump file after upgrading. Use the `--all-databases` [292] option to include all databases in the dump. If your databases include stored programs, use the `--routines` [293] and `--events` [293] options as well.

In general, do the following when upgrading from MySQL 5.5 to 5.6:

- Read *all* the items in these sections to see whether any of them might affect your applications:
  - [Section 2.11.1, “Upgrading MySQL”](#), has general update information.
  - The items in the change lists provided later in this section enable you to identify upgrade issues that apply to your current MySQL installation. Some incompatibilities discussed there require your attention *before* upgrading. Others should be dealt with *after* upgrading.



- The MySQL 5.6 [Release Notes](#) describe significant new features you can use in 5.6 or that differ from those found in earlier MySQL releases. Some of these changes may result in incompatibilities.

Note particularly any changes that are marked **Known issue** or **Incompatible change**. These incompatibilities with earlier versions of MySQL may require your attention *before you upgrade*. Our aim is to avoid these changes, but occasionally they are necessary to correct problems that would be worse than an incompatibility between releases. If any upgrade issue applicable to your installation involves an incompatibility that requires special handling, follow the instructions given in the incompatibility description. Sometimes this involves dumping and reloading tables, or use of a statement such as `CHECK TABLE` or `REPAIR TABLE`.

For dump and reload instructions, see [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#). Any procedure that involves `REPAIR TABLE` with the `USE_FRM` option *must* be done before upgrading. Use of this statement with a version of MySQL different from the one used to create the table (that is, using it after upgrading) may damage the table. See [Section 13.7.2.5, “REPAIR TABLE Syntax”](#).

- Before upgrading to a new version of MySQL, [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#), to see whether changes to table formats or to character sets or collations were made between your current version of MySQL and the version to which you are upgrading. If so and these changes result in an incompatibility between MySQL versions, you will need to upgrade the affected tables using the instructions in [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).
- After upgrading to a new version of MySQL, run `mysql_upgrade` (see [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#)). This program checks your tables, and attempts to repair them if necessary. It also updates your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. (Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features.)

`mysql_upgrade` does not upgrade the contents of the help tables. For upgrade instructions, see [Section 5.1.10, “Server-Side Help”](#).

- If you run MySQL Server on Windows, see [Section 2.3.6, “Upgrading MySQL on Windows”](#).
- If you use replication, see [Section 16.4.3, “Upgrading a Replication Setup”](#), for information on upgrading your replication setup.

If your MySQL installation contains a large amount of data that might take a long time to convert after an in-place upgrade, you might find it useful to create a “dummy” database instance for assessing what conversions might be needed and the work involved to perform them. Make a copy of your MySQL instance that contains a full copy of the `mysql` database, plus all other databases without data. Run your upgrade procedure on this dummy instance to see what actions might be needed so that you can better evaluate the work involved when performing actual data conversion on your original database instance.

Read *all* the items in the following sections to see whether any of them might affect your applications:

### Configuration Changes

- Beginning with MySQL 5.6.6, several MySQL Server parameters have defaults that differ from previous releases. The motivation for these changes is to provide better out-of-box performance and to reduce the need for the database administrator to change settings manually. These changes are subject to possible revision in future releases as we gain feedback.

In some cases, a parameter has a different static default value. In other cases, the server autosizes a parameter at startup using a formula based on other related parameters or server host configuration, rather than using a static value. For example, the setting for `back_log` [465] now is its previous default of 50, adjusted up by an amount proportional to the value of

`max_connections` [503]. The idea behind autosizing is that when the server has information available to make a decision about a parameter setting likely to be better than a fixed default, it will.

The following table summarizes changes to defaults. Any of these can be overridden by specifying an explicit value at server startup.

Parameter	Old Default	New Default
<code>back_log</code> [465]	50	Autosized using <code>max_connections</code> [503]
<code>binlog_checksum</code> [2009]	NONE	CRC32
<code>--binlog-row-event-max-size</code> [2001]	1024	8192
<code>flush_time</code> [482]	1800 (on Windows)	0
<code>innodb_autoextend_increment</code> [1784]	64	64
<code>innodb_buffer_pool_instances</code> [1706]	1	8 (platform dependent)
<code>innodb_checksum_algorithm</code> [1710]	INNODB	CRC32
<code>innodb_concurrency_tickets</code> [1715]	500	5000
<code>innodb_file_per_table</code> [1718]	0	1
<code>innodb_old_blocks_time</code> [1738]	0	1000
<code>innodb_open_files</code> [1739]	300	Autosized using <code>innodb_file_per_table</code> [1718], <code>table_open_cache</code> [551]
<code>innodb_stats_on_metadata</code> [1747]	ON	OFF
<code>join_buffer_size</code> [490]	128KB	256KB
<code>max_allowed_packet</code> [502]	1MB	4MB
<code>max_connect_errors</code> [502]	10	100
<code>sync_master_info</code> [1997]	0	10000
<code>sync_relay_log</code> [1998]	0	10000
<code>sync_relay_log_info</code> [1999]	0	10000

With regard to compatibility with previous releases, the most important changes are:

- `innodb_file_per_table` [1718] is enabled (previously disabled).
- `innodb_checksum_algorithm` [1710] is CRC32 (previously INNODB).
- `binlog_checksum` [2009] is CRC32 (previously NONE).

Therefore, if you are upgrading an existing MySQL installation, have not already changed the values of these parameters from their previous defaults, and backward compatibility is a concern, you may want to explicitly set these parameters to their previous defaults. For example, put these lines in the server option file:

```
[mysqld]
innodb_file_per_table=0
innodb_checksum_algorithm=INNODB
binlog_checksum=NONE
```

Those settings preserve compatibility as follows:

- With the new default of `innodb_file_per_table` [1718] enabled, `ALTER TABLE` operations following an upgrade will move InnoDB tables that are in the system tablespace to individual `.ibd` files. Using `innodb_file_per_table=0` [1718] will prevent this from happening.

- Setting `innodb_checksum_algorithm=INNODB` [1710] permits binary downgrades after upgrading to this release. With a setting of `CRC32`, InnoDB would use checksumming that older MySQL versions cannot use.
- With `binlog_checksum=NONE` [2009], the server can be used as a replication master without causing failure of older slaves that do not understand binary log checksums.

## Server Changes

- **Incompatible change:** As of MySQL 5.6, the full-text stopword file is loaded and searched using `latin1` if `character_set_server` is `ucs2`, `utf16`, `utf16le`, or `utf32`. If any table was created with `FULLTEXT` indexes while the server character set was `ucs2`, `utf16`, `utf16le`, or `utf32`, repair it using this statement:

```
REPAIR TABLE tbl_name QUICK;
```

- **Incompatible change:** MySQL 5.6.11 and later supports `CREATE TABLE ... [SUB]PARTITION BY ALGORITHM=n [LINEAR] KEY (...)`, which can be used to create a table whose `KEY` partitioning is compatible with a MySQL 5.1 server ( $n=1$ ). (Bug #14521864, Bug #66462) This syntax is not accepted by MySQL 5.6.10 and earlier, although it is supported in MySQL 5.5 beginning with MySQL 5.5.31. `mysqldump` in MySQL 5.5.31 and later MySQL 5.5 releases includes the `ALGORITHM` option when dumping tables using this option, but surrounds it with conditional comments, like this:

```
CREATE TABLE t1 (a INT)
/*!50100 PARTITION BY KEY */ /*!50531 ALGORITHM = 1 */ /*!50100 ( )
PARTITIONS 3 */
```

When importing a dump containing such `CREATE TABLE` statements into a MySQL 5.6.10 or earlier MySQL 5.6 server, the versioned comment is not ignored, which causes a syntax error. Therefore, prior to importing such a dump file, you must either change the comments so that the MySQL 5.6 server ignores them (by removing the string `!50531` or replacing it with `!50611`, wherever it occurs), or remove them.

This is not an issue with dump files made using MySQL 5.6.11 or later, where the `ALGORITHM` option is written using `/*!50611 ... */`.

## SQL Changes

- Some keywords may be reserved in MySQL 5.6 that were not reserved in MySQL 5.5. See [Section 9.3, “Reserved Words”](#).
- The `YEAR(2)` data type has certain issues that you should consider before choosing to use it. As of MySQL 5.6.6, `YEAR(2)` is deprecated. `YEAR(2)` columns in existing tables are treated as before, but `YEAR(2)` in new or altered tables are converted to `YEAR(4)`. For more information, see [Section 11.3.4, “YEAR\(2\) Limitations and Migrating to YEAR\(4\)”](#).
- As of MySQL 5.6.6, it is explicitly disallowed to assign the value `DEFAULT` to stored procedure or function parameters or stored program local variables (for example with a `SET var_name = DEFAULT` statement). This was not previously supported, or documented as permitted, but is flagged as an incompatible change in case existing code inadvertently used this construct. It remains permissible to assign `DEFAULT` to system variables, as before, but assigning `DEFAULT` to parameters or local variables now results in a syntax error.

After an upgrade to MySQL 5.6.6 or later, existing stored programs that use this construct produce a syntax error when invoked. If a `mysqldump` file from 5.6.5 or earlier is loaded into 5.6.6 or later, the load operation fails and affected stored program definitions must be changed.

- In MySQL, the `TIMESTAMP` data type differs in nonstandard ways from other data types:

- `TIMESTAMP` columns not explicitly declared with the `NULL` attribute are assigned the `NOT NULL` attribute. (Columns of other data types, if not explicitly declared as `NOT NULL`, permit `NULL` values.) Setting such a column to `NULL` sets it to the current timestamp.
- The first `TIMESTAMP` column in a table, if not declared with the `NULL` attribute or an explicit `DEFAULT` or `ON UPDATE` clause, is automatically assigned the `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` attributes.
- `TIMESTAMP` columns following the first one, if not declared with the `NULL` attribute or an explicit `DEFAULT` clause, are automatically assigned `DEFAULT '0000-00-00 00:00:00'` (the “zero” timestamp). For inserted rows that specify no explicit value for such a column, the column is assigned `'0000-00-00 00:00:00'` and no warning occurs.

Those nonstandard behaviors remain the default for `TIMESTAMP` but as of MySQL 5.6.6 are deprecated and this warning appears at startup:

```
[Warning] TIMESTAMP with implicit DEFAULT value is deprecated.
Please use --explicit_defaults_for_timestamp server option (see
documentation for more details).
```

As indicated by the warning, to turn off the nonstandard behaviors, enable the new `explicit_defaults_for_timestamp` [481] system variable at server startup. With this variable enabled, the server handles `TIMESTAMP` as follows instead:

- `TIMESTAMP` columns not explicitly declared as `NOT NULL` permit `NULL` values. Setting such a column to `NULL` sets it to `NULL`, not the current timestamp.
- No `TIMESTAMP` column is assigned the `DEFAULT CURRENT_TIMESTAMP` or `ON UPDATE CURRENT_TIMESTAMP` attributes automatically. Those attributes must be explicitly specified.
- `TIMESTAMP` columns declared as `NOT NULL` and without an explicit `DEFAULT` clause are treated as having no default value. For inserted rows that specify no explicit value for such a column, the result depends on the SQL mode. If strict SQL mode is enabled, an error occurs. If strict SQL mode is not enabled, the column is assigned the implicit default of `'0000-00-00 00:00:00'` and a warning occurs. This is similar to how MySQL treats other temporal types such as `DATETIME`.

To upgrade servers used for replication, upgrade the slaves first, then the master. Replication between the master and its slaves should work provided that all use the same value of `explicit_defaults_for_timestamp` [481]:

1. Bring down the slaves, upgrade them, configure them with the desired value of `explicit_defaults_for_timestamp` [481], and bring them back up.

The slaves will recognize from the format of the binary logs received from the master that the master is older (predates the introduction of `explicit_defaults_for_timestamp` [481]) and that operations on `TIMESTAMP` columns coming from the master use the old `TIMESTAMP` behavior.

2. Bring down the master, upgrade it, and configure it with the same `explicit_defaults_for_timestamp` [481] value used on the slaves, and bring it back up.

## 2.11.2. Downgrading MySQL

This section describes what to do to downgrade to an older MySQL version, in the unlikely case that the previous version worked better than the new one.

It is always a good idea to make a backup beforehand, in case a downgrade fails and leaves the instance in an unusable state.

To downgrade between General Availability (GA) status versions within the same release series, typically you just install the new binaries on top of the old ones and do not make any changes to the databases.

Downgrades between milestone releases (or from a GA release to a milestone release) within the same release series are not supported and you may encounter issues.

The following items form a checklist of things to do whenever you perform a downgrade:

- Read the upgrading section for the release series from which you are downgrading to be sure that it does not have any features you really need. See [Section 2.11.1, “Upgrading MySQL”](#).
- If there is a downgrading section for that version, read that as well.
- To see which new features were added between the version to which you are downgrading and your current version, see the [Release Notes](#).
- Check [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#), to see whether changes to table formats or to character sets or collations were made between your current version of MySQL and the version to which you are downgrading. If so and these changes result in an incompatibility between MySQL versions, you will need to downgrade the affected tables using the instructions in [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).

In most cases, you can move the MySQL format files and data files between different GA versions on the same architecture as long as you stay within versions for the same release series of MySQL.

If you downgrade from one release series to another, there may be incompatibilities in table storage formats. In this case, use `mysqldump` to dump your tables before downgrading. After downgrading, reload the dump file using `mysql` or `mysqlimport` to re-create your tables. For examples, see [Section 2.11.5, “Copying MySQL Databases to Another Machine”](#).

A typical symptom of a downward-incompatible table format change when you downgrade is that you cannot open tables. In that case, use the following procedure:

1. Stop the older MySQL server that you are downgrading to.
2. Restart the newer MySQL server you are downgrading from.
3. Dump any tables that were inaccessible to the older server by using `mysqldump` to create a dump file.
4. Stop the newer MySQL server and restart the older one.
5. Reload the dump file into the older server. Your tables should be accessible.

If system tables in the `mysql` database changed, downgrading might introduce some loss of functionality or require some adjustments. Here are some examples:

- Trigger creation requires the `TRIGGER` privilege as of MySQL 5.1. In MySQL 5.0, there is no `TRIGGER` privilege and `SUPER` is required instead. If you downgrade from MySQL 5.1 to 5.0, you will need to give the `SUPER` privilege to those accounts that had the `TRIGGER` privilege in 5.1.
- Triggers were added in MySQL 5.0, so if you downgrade from 5.0 to 4.1, you cannot use triggers at all.
- The `mysql.proc.comment` column definition changed between MySQL 5.1 and 5.5. After a downgrade from 5.5 to 5.1, this table is seen as corrupt and in need of repair. To workaround this problem, execute `mysql_upgrade` from the version of MySQL to which you downgraded.

### 2.11.2.1. Downgrading to MySQL 5.5

When downgrading to MySQL 5.5 from MySQL 5.6, keep in mind the following issues relating to features found in MySQL 5.6, but not in MySQL 5.5:

## System Tables

- The `mysql.user` table in MySQL 5.6 has a `password_expired` column. The `mysql.user` table in MySQL 5.5 does not. This means that an account with an expired password in MySQL 5.6 will work normally in MySQL 5.5.

## InnoDB

- InnoDB search indexes (with a type of `FULLTEXT`), introduced in MySQL 5.6.4, are not compatible with earlier versions of MySQL, including earlier releases in the 5.6 series. Drop such indexes before performing a downgrade.
- InnoDB small page sizes specified by the `innodb_page_size [1740]` configuration option, introduced in MySQL 5.6.4, are not compatible with earlier versions of MySQL, including earlier releases in the 5.6 series. Dump all InnoDB tables in instances that use a smaller InnoDB page size, drop the tables, and re-create and reload them after the downgrade.

## Replication

- As of MySQL 5.6, the `relay-log.info` file contains a line count and a replication delay value, so the file format differs from that in older versions. See [Section 16.2.2.2, “Slave Status Logs”](#). If you downgrade a slave server to a version older than MySQL 5.6, the older server will not read the file correctly. To address this, modify the file in a text editor to delete the initial line containing the number of lines.
- Beginning with MySQL 5.6.6, the MySQL Server employs Version 2 binary log events when writing the binary log. Binary logs written using Version 2 log events cannot be read by earlier versions of MySQL Server. To generate a binary log that is written using Version 1 log events readable by older servers, start the MySQL 5.6.6 or later server using `--log-bin-use-v1-row-events=1 [2003]`, which forces the server to employ Version 1 events when writing the binary log.

### 2.11.3. Checking Whether Tables or Indexes Must Be Rebuilt

A binary upgrade or downgrade is one that installs one version of MySQL “in place” over an existing version, without dumping and reloading tables:

1. Stop the server for the existing version if it is running.
2. Install a different version of MySQL. This is an upgrade if the new version is higher than the original version, a downgrade if the version is lower.
3. Start the server for the new version.

In many cases, the tables from the previous version of MySQL can be used without problem by the new version. However, sometimes changes occur that require tables or table indexes to be rebuilt, as described in this section. If you have tables that are affected by any of the issues described here, rebuild the tables or indexes as necessary using the instructions given in [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).

## Table Incompatibilities

After a binary upgrade to MySQL 5.1 from a MySQL 5.0 installation that contains `ARCHIVE` tables, accessing those tables causes the server to crash, even if you have run `mysql_upgrade` or `CHECK TABLE ... FOR UPGRADE`. To work around this problem, use `mysqldump` to dump all `ARCHIVE` tables before upgrading, and reload them into MySQL 5.1 after upgrading. The same problem occurs for binary downgrades from MySQL 5.1 to 5.0.

The upgrade problem is fixed in MySQL 5.6.4: The server can open `ARCHIVE` tables created in MySQL 5.0. However, it remains the recommended upgrade procedure to dump 5.0 `ARCHIVE` tables before upgrading and reload them after upgrading.



## Index Incompatibilities

In MySQL 5.6.3, the length limit for index prefix keys is increased from 767 bytes to 3072 bytes, for InnoDB tables using `ROW_FORMAT=DYNAMIC` or `ROW_FORMAT=COMPRESSED`. See [Section 14.2.7, “Limits on InnoDB Tables”](#) for details. This change is also backported to MySQL 5.5.14. If you downgrade from one of these releases or higher, to an earlier release with a lower length limit, the index prefix keys could be truncated at 767 bytes or the downgrade could fail. This issue could only occur if the configuration option `innodb_large_prefix [1728]` was enabled on the server being downgraded.

If you perform a binary upgrade without dumping and reloading tables, you cannot upgrade directly from MySQL 4.1 to 5.1 or higher. This occurs due to an incompatible change in the `MyISAM` table index format in MySQL 5.0. Upgrade from MySQL 4.1 to 5.0 and repair all `MyISAM` tables. Then upgrade from MySQL 5.0 to 5.1 and check and repair your tables.

Modifications to the handling of character sets or collations might change the character sort order, which causes the ordering of entries in any index that uses an affected character set or collation to be incorrect. Such changes result in several possible problems:

- Comparison results that differ from previous results
- Inability to find some index values due to misordered index entries
- Misordered `ORDER BY` results
- Tables that `CHECK TABLE` reports as being in need of repair

The solution to these problems is to rebuild any indexes that use an affected character set or collation, either by dropping and re-creating the indexes, or by dumping and reloading the entire table. In some cases, it is possible to alter affected columns to use a different collation. For information about rebuilding indexes, see [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).

To check whether a table has indexes that must be rebuilt, consult the following list. It indicates which versions of MySQL introduced character set or collation changes that require indexes to be rebuilt. Each entry indicates the version in which the change occurred and the character sets or collations that the change affects. If the change is associated with a particular bug report, the bug number is given.

The list applies both for binary upgrades and downgrades. For example, Bug #27877 was fixed in MySQL 5.1.24 and 5.4.0, so it applies to upgrades from versions older than 5.1.24 to 5.1.24 or newer, and to downgrades from 5.1.24 or newer to versions older than 5.1.24.

In many cases, you can use `CHECK TABLE ... FOR UPGRADE` to identify tables for which index rebuilding is required. It will report this message:

```
Table upgrade required.
Please do "REPAIR TABLE `tbl_name`" or dump/reload to fix it!
```

In these cases, you can also use `mysqlcheck --check-upgrade` or `mysql_upgrade`, which execute `CHECK TABLE`. However, the use of `CHECK TABLE` applies only after upgrades, not downgrades. Also, `CHECK TABLE` is not applicable to all storage engines. For details about which storage engines `CHECK TABLE` supports, see [Section 13.7.2.2, “CHECK TABLE Syntax”](#).

These changes cause index rebuilding to be necessary:

- MySQL 5.1.24, 5.4.0 (Bug #27877)

Affects indexes that use the `utf8_general_ci` or `ucs2_general_ci` collation for columns that contain 'ß' LATIN SMALL LETTER SHARP S (German). The bug fix corrected an error in the original collations but introduced an incompatibility such that 'ß' compares equal to characters with which it previously compared different.



Affected tables can be detected by `CHECK TABLE ... FOR UPGRADE` as of MySQL 5.1.30, 5.4.0 (see Bug #40053).

A workaround for this issue is implemented as of MySQL 5.1.62, 5.5.21, and 5.6.5. The workaround involves altering affected columns to use the `utf8_general_mysql500_ci` and `ucs2_general_mysql500_ci` collations, which preserve the original pre-5.1.24 ordering of `utf8_general_ci` and `ucs2_general_ci`.

- MySQL 5.0.48, 5.1.23 (Bug #27562)

Affects indexes that use the `ascii_general_ci` collation for columns that contain any of these characters: ``` GRAVE ACCENT, `[` LEFT SQUARE BRACKET, `\` REVERSE SOLIDUS, `]` RIGHT SQUARE BRACKET, `~` TILDE

Affected tables can be detected by `CHECK TABLE ... FOR UPGRADE` as of MySQL 5.1.29, 5.4.0 (see Bug #39585).

- MySQL 5.0.48, 5.1.21 (Bug #29461)

Affects indexes for columns that use any of these character sets: `eurjpm`, `eurkr`, `gb2312`, `latin7`, `macce`, `ujis`

Affected tables can be detected by `CHECK TABLE ... FOR UPGRADE` as of MySQL 5.1.29, 5.4.0 (see Bug #39585).

## 2.11.4. Rebuilding or Repairing Tables or Indexes

This section describes how to rebuild a table, following changes to MySQL such as how data types or character sets are handled. For example, an error in a collation might have been corrected, requiring a table rebuild to update the indexes for character columns that use the collation. (For examples, see [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#).) You might also need to repair or upgrade a table, as indicated by a table check operation such as that performed by `CHECK TABLE`, `mysqlcheck`, or `mysql_upgrade`.

Methods for rebuilding a table include dumping and reloading it, or using `ALTER TABLE` or `REPAIR TABLE`.



### Note

If you are rebuilding tables because a different version of MySQL will not handle them after a binary (in-place) upgrade or downgrade, you must use the dump-and-reload method. Dump the tables *before* upgrading or downgrading using your original version of MySQL. Then reload the tables *after* upgrading or downgrading.

If you use the dump-and-reload method of rebuilding tables only for the purpose of rebuilding indexes, you can perform the dump either before or after upgrading or downgrading. Reloading still must be done afterward.

To rebuild a table by dumping and reloading it, use `mysqldump` to create a dump file and `mysql` to reload the file:

```
shell> mysqldump db_name t1 > dump.sql
shell> mysql db_name < dump.sql
```

To rebuild all the tables in a single database, specify the database name without any following table name:

```
shell> mysqldump db_name > dump.sql
shell> mysql db_name < dump.sql
```

To rebuild all tables in all databases, use the `--all-databases` [292] option:

```
shell> mysqldump --all-databases > dump.sql
shell> mysql < dump.sql
```

To rebuild a table with `ALTER TABLE`, use a “null” alteration; that is, an `ALTER TABLE` statement that “changes” the table to use the storage engine that it already has. For example, if `t1` is a `MyISAM` table, use this statement:

```
mysql> ALTER TABLE t1 ENGINE = MyISAM;
```

If you are not sure which storage engine to specify in the `ALTER TABLE` statement, use `SHOW CREATE TABLE` to display the table definition.

If you must rebuild a table because a table checking operation indicates that the table is corrupt or needs an upgrade, you can use `REPAIR TABLE` if that statement supports the table's storage engine. For example, to repair a `MyISAM` table, use this statement:

```
mysql> REPAIR TABLE t1;
```

For storage engines such as `InnoDB` that `REPAIR TABLE` does not support, use `mysqldump` to create a dump file and `mysql` to reload the file, as described earlier.

For specifics about which storage engines `REPAIR TABLE` supports, see [Section 13.7.2.5, “REPAIR TABLE Syntax”](#).

`mysqlcheck --repair` provides command-line access to the `REPAIR TABLE` statement. This can be a more convenient means of repairing tables because you can use the `--databases` [276] or `--all-databases` [276] option to repair all tables in specific databases or all databases, respectively:

```
shell> mysqlcheck --repair --databases db_name ...
shell> mysqlcheck --repair --all-databases
```

For incompatibilities introduced in MySQL 5.1.24 by the fix for Bug #27877 that corrected the `utf8_general_ci` and `ucs2_general_ci` collations, a workaround is implemented as of MySQL 5.1.62, 5.5.21, and 5.6.5. Upgrade to one of those versions, then convert each affected table using one of the following methods. In each case, the workaround altering affected columns to use the `utf8_general_mysql500_ci` and `ucs2_general_mysql500_ci` collations, which preserve the original pre-5.1.24 ordering of `utf8_general_ci` and `ucs2_general_ci`.

- To convert an affected table after a binary upgrade that leaves the table files in place, alter the table to use the new collation. Suppose that the table `t1` contains one or more problematic `utf8` columns. To convert the table at the table level, use a statement like this:

```
ALTER TABLE t1
CONVERT TO CHARACTER SET utf8 COLLATE utf8_general_mysql500_ci;
```

To apply the change on a column-specific basis, use a statement like this (be sure to repeat the column definition as originally specified except for the `COLLATE` clause):

```
ALTER TABLE t1
MODIFY c1 CHAR(N) CHARACTER SET utf8 COLLATE utf8_general_mysql500_ci;
```

- To upgrade the table using a dump and reload procedure, dump the table using `mysqldump`, modify the `CREATE TABLE` statement in the dump file to use the new collation, and reload the table.

After making the appropriate changes, `CHECK TABLE` should report no error.

## 2.11.5. Copying MySQL Databases to Another Machine

You can copy the `.frm`, `.MYI`, and `.MYD` files for `MyISAM` tables between different architectures that support the same floating-point format. (MySQL takes care of any byte-swapping issues.) See [Section 14.3, “The `MyISAM` Storage Engine”](#).

In cases where you need to transfer databases between different architectures, you can use `mysqldump` to create a file containing SQL statements. You can then transfer the file to the other machine and feed it as input to the `mysql` client.

Use `mysqldump --help` to see what options are available.

The easiest (although not the fastest) way to move a database between two machines is to run the following commands on the machine on which the database is located:

```
shell> mysqladmin -h 'other_hostname' create db_name
shell> mysqldump db_name | mysql -h 'other_hostname' db_name
```

If you want to copy a database from a remote machine over a slow network, you can use these commands:

```
shell> mysqladmin create db_name
shell> mysqldump -h 'other_hostname' --compress db_name | mysql db_name
```

You can also store the dump in a file, transfer the file to the target machine, and then load the file into the database there. For example, you can dump a database to a compressed file on the source machine like this:

```
shell> mysqldump --quick db_name | gzip > db_name.gz
```

Transfer the file containing the database contents to the target machine and run these commands there:

```
shell> mysqladmin create db_name
shell> gunzip < db_name.gz | mysql db_name
```

You can also use `mysqldump` and `mysqlimport` to transfer the database. For large tables, this is much faster than simply using `mysqldump`. In the following commands, `DUMPDIR` represents the full path name of the directory you use to store the output from `mysqldump`.

First, create the directory for the output files and dump the database:

```
shell> mkdir DUMPDIR
shell> mysqldump --tab=DUMPDIR db_name
```

Then transfer the files in the `DUMPDIR` directory to some corresponding directory on the target machine and load the files into MySQL there:

```
shell> mysqladmin create db_name           # create database
shell> cat DUMPDIR/*.sql | mysql db_name   # create tables in database
shell> mysqlimport db_name DUMPDIR/*.txt  # load data into tables
```

Do not forget to copy the `mysql` database because that is where the grant tables are stored. You might have to run commands as the MySQL `root` user on the new machine until you have the `mysql` database in place.

After you import the `mysql` database on the new machine, execute `mysqladmin flush-privileges` so that the server reloads the grant table information.

## 2.12. Environment Variables

This section lists all the environment variables that are used directly or indirectly by MySQL. Most of these can also be found in other places in this manual.

Note that any options on the command line take precedence over values specified in option files and environment variables, and values in option files take precedence over values in environment variables.

In many cases, it is preferable to use an option file instead of environment variables to modify the behavior of MySQL. See [Section 4.2.3.3, “Using Option Files”](#).

Variable	Description
CXX	The name of your C++ compiler (for running <code>CMake</code> ).
CC	The name of your C compiler (for running <code>CMake</code> ).
CFLAGS	Flags for your C compiler (for running <code>CMake</code> ).
CXXFLAGS	Flags for your C++ compiler (for running <code>CMake</code> ).
DBI_USER	The default user name for Perl DBI.
DBI_TRACE	Trace options for Perl DBI.
HOME	The default path for the <code>mysql</code> history file is <code>\$HOME/.mysql_history</code> .
LD_RUN_PATH	Used to specify the location of <code>libmysqlclient.so</code> .
LIBMYSQL_ENABLE_CLEARTEXT_PLUGIN	Enable <code>mysql_clear_password</code> authentication plugin; see <a href="#">Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”</a> .
LIBMYSQL_PLUGIN_DIR	Directory in which to look for client plugins.
LIBMYSQL_PLUGINS	Client plugins to preload.
MYSQL_DEBUG	Debug trace options when debugging.
MYSQL_GROUP_SUFFIX	Option group suffix value (like specifying <code>--defaults-group-suffix [219]</code> ).
MYSQL_HISTFILE	The path to the <code>mysql</code> history file. If this variable is set, its value overrides the default for <code>\$HOME/.mysql_history</code> .
MYSQL_HISTIGNORE	Patterns specifying statements not to log to <code>\$HOME/.mysql_history</code> .
MYSQL_HOME	The path to the directory in which the server-specific <code>my.cnf</code> file resides.
MYSQL_HOST	The default host name used by the <code>mysql</code> command-line client.
MYSQL_PS1	The command prompt to use in the <code>mysql</code> command-line client.
MYSQL_PWD	The default password when connecting to <code>mysqld</code> . Note that using this is insecure. See <a href="#">Section 6.1.2.1, “End-User Guidelines for Password Security”</a> .
MYSQL_TCP_PORT	The default TCP/IP port number.
MYSQL_TEST_LOGIN_FILE	The name of the <code>.mylogin.cnf</code> login file.
MYSQL_UNIX_PORT	The default Unix socket file name; used for connections to <code>localhost</code> .
PATH	Used by the shell to find MySQL programs.
TMPDIR	The directory where temporary files are created.
TZ	This should be set to your local time zone. See <a href="#">Section C.5.4.6, “Time Zone Problems”</a> .
UMASK	The user-file creation mode when creating files. See note following table.
UMASK_DIR	The user-directory creation mode when creating directories. See note following table.
USER	The default user name on Windows when connecting to <code>mysqld</code> .

For information about the `mysql` history file, see [Section 4.5.1.3, “mysql Logging”](#).

`MYSQL_TEST_LOGIN_FILE` is the path name of the login file (the file created by `mysql_config_editor`). If not set, the default value is `%APPDATA%\MySQL\mylogin.cnf` directory on Windows and `$HOME/.mylogin.cnf` on non-Windows systems. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#).

The `UMASK` and `UMASK_DIR` variables, despite their names, are used as modes, not masks:

- If `UMASK` is set, `mysqld` uses `( $UMASK | 0600 )` as the mode for file creation, so that newly created files have a mode in the range from `0600` to `0666` (all values octal).

- If `UMASK_DIR` is set, `mysqld` uses `( $\$UMASK\_DIR \mid 0700$ )` as the base mode for directory creation, which then is AND-ed with `~( $\sim\$UMASK \& 0666$ )`, so that newly created directories have a mode in the range from 0700 to 0777 (all values octal). The AND operation may remove read and write permissions from the directory mode, but not execute permissions.

MySQL assumes that the value for `UMASK` or `UMASK_DIR` is in octal if it starts with a zero.

## 2.13. Perl Installation Notes

The Perl `DBI` module provides a generic interface for database access. You can write a `DBI` script that works with many different database engines without change. To use `DBI`, you must install the `DBI` module, as well as a DataBase Driver (DBD) module for each type of database server you want to access. For MySQL, this driver is the `DBD:mysql` module.

Perl, and the `DBD:mysql` module for `DBI` must be installed if you want to run the MySQL benchmark scripts; see [Section 8.12.2, “The MySQL Benchmark Suite”](#).



### Note

Perl support is not included with MySQL distributions. You can obtain the necessary modules from <http://search.cpan.org> for Unix, or by using the ActiveState `ppm` program on Windows. The following sections describe how to do this.

The `DBI/DBD` interface requires Perl 5.6.0, and 5.6.1 or later is preferred. `DBI` *does not work* if you have an older version of Perl. You should use `DBD:mysql` 4.009 or higher. Although earlier versions are available, they do not support the full functionality of MySQL 5.6.

### 2.13.1. Installing Perl on Unix

MySQL Perl support requires that you have installed MySQL client programming support (libraries and header files). Most installation methods install the necessary files. If you install MySQL from RPM files on Linux, be sure to install the developer RPM as well. The client programs are in the client RPM, but client programming support is in the developer RPM.

The files you need for Perl support can be obtained from the CPAN (Comprehensive Perl Archive Network) at <http://search.cpan.org>.

The easiest way to install Perl modules on Unix is to use the `CPAN` module. For example:

```
shell> perl -MCPAN -e shell
cpan> install DBI
cpan> install DBD:mysql
```

The `DBD:mysql` installation runs a number of tests. These tests attempt to connect to the local MySQL server using the default user name and password. (The default user name is your login name on Unix, and `ODBC` on Windows. The default password is “no password.”) If you cannot connect to the server with those values (for example, if your account has a password), the tests fail. You can use `force install DBD:mysql` to ignore the failed tests.

`DBI` requires the `Data:Dumper` module. It may be installed; if not, you should install it before installing `DBI`.

It is also possible to download the module distributions in the form of compressed `tar` archives and build the modules manually. For example, to unpack and build a `DBI` distribution, use a procedure such as this:

1. Unpack the distribution into the current directory:

```
shell> gunzip < DBI-VERSION.tar.gz | tar xvf -
```

This command creates a directory named `DBI-VERSION`.

2. Change location into the top-level directory of the unpacked distribution:

```
shell> cd DBI-VERSION
```

3. Build the distribution and compile everything:

```
shell> perl Makefile.PL
shell> make
shell> make test
shell> make install
```

The `make test` command is important because it verifies that the module is working. Note that when you run that command during the `DBD:mysql` installation to exercise the interface code, the MySQL server must be running or the test fails.

It is a good idea to rebuild and reinstall the `DBD:mysql` distribution whenever you install a new release of MySQL. This ensures that the latest versions of the MySQL client libraries are installed correctly.

If you do not have access rights to install Perl modules in the system directory or if you want to install local Perl modules, the following reference may be useful: <http://servers.digitaldaze.com/extensions/perl/modules.html#modules>

Look under the heading “Installing New Modules that Require Locally Installed Modules.”

## 2.13.2. Installing ActiveState Perl on Windows

On Windows, you should do the following to install the MySQL `DBD` module with ActiveState Perl:

1. Get ActiveState Perl from <http://www.activestate.com/Products/ActivePerl/> and install it.
2. Open a console window.
3. If necessary, set the `HTTP_proxy` variable. For example, you might try a setting like this:

```
C:\> set HTTP_proxy=my.proxy.com:3128
```

4. Start the PPM program:

```
C:\> C:\perl\bin\ppm.pl
```

5. If you have not previously done so, install `DBI`:

```
ppm> install DBI
```

6. If this succeeds, run the following command:

```
ppm> install DBD-mysql
```

This procedure should work with ActiveState Perl 5.6 or newer.

If you cannot get the procedure to work, you should install the ODBC driver instead and connect to the MySQL server through ODBC:

```
use DBI;
$dbh= DBI->connect("DBI:ODBC:$dsn", $user, $password) ||
die "Got error $DBI::errstr when connecting to $dsn\n";
```

## 2.13.3. Problems Using the Perl `DBI/DBD` Interface

If Perl reports that it cannot find the `./mysql/mysql.so` module, the problem is probably that Perl cannot locate the `libmysqlclient.so` shared library. You should be able to fix this problem by one of the following methods:

- Copy `libmysqlclient.so` to the directory where your other shared libraries are located (probably `/usr/lib` or `/lib`).
- Modify the `-L` options used to compile `DBD:mysql` to reflect the actual location of `libmysqlclient.so`.
- On Linux, you can add the path name of the directory where `libmysqlclient.so` is located to the `/etc/ld.so.conf` file.
- Add the path name of the directory where `libmysqlclient.so` is located to the `LD_RUN_PATH` environment variable. Some systems use `LD_LIBRARY_PATH` instead.

Note that you may also need to modify the `-L` options if there are other libraries that the linker fails to find. For example, if the linker cannot find `libc` because it is in `/lib` and the link command specifies `-L/usr/lib`, change the `-L` option to `-L/lib` or add `-L/lib` to the existing link command.

If you get the following errors from `DBD:mysql`, you are probably using `gcc` (or using an old binary compiled with `gcc`):

```
/usr/bin/perl: can't resolve symbol '__moddi3'  
/usr/bin/perl: can't resolve symbol '__divdi3'
```

Add `-L/usr/lib/gcc-lib/... -lgcc` to the link command when the `mysql.so` library gets built (check the output from `make` for `mysql.so` when you compile the Perl client). The `-L` option should specify the path name of the directory where `libgcc.a` is located on your system.

Another cause of this problem may be that Perl and MySQL are not both compiled with `gcc`. In this case, you can solve the mismatch by compiling both with `gcc`.





---

## Chapter 3. Tutorial

### Table of Contents

3.1. Connecting to and Disconnecting from the Server .....	171
3.2. Entering Queries .....	172
3.3. Creating and Using a Database .....	175
3.3.1. Creating and Selecting a Database .....	176
3.3.2. Creating a Table .....	177
3.3.3. Loading Data into a Table .....	178
3.3.4. Retrieving Information from a Table .....	179
3.4. Getting Information About Databases and Tables .....	191
3.5. Using <code>mysql</code> in Batch Mode .....	192
3.6. Examples of Common Queries .....	193
3.6.1. The Maximum Value for a Column .....	194
3.6.2. The Row Holding the Maximum of a Certain Column .....	194
3.6.3. Maximum of Column per Group .....	195
3.6.4. The Rows Holding the Group-wise Maximum of a Certain Column .....	195
3.6.5. Using User-Defined Variables .....	196
3.6.6. Using Foreign Keys .....	196
3.6.7. Searching on Two Keys .....	197
3.6.8. Calculating Visits Per Day .....	198
3.6.9. Using <code>AUTO_INCREMENT</code> .....	198
3.7. Using MySQL with Apache .....	200

This chapter provides a tutorial introduction to MySQL by showing how to use the `mysql` client program to create and use a simple database. `mysql` (sometimes referred to as the “terminal monitor” or just “monitor”) is an interactive program that enables you to connect to a MySQL server, run queries, and view the results. `mysql` may also be used in batch mode: you place your queries in a file beforehand, then tell `mysql` to execute the contents of the file. Both ways of using `mysql` are covered here.

To see a list of options provided by `mysql`, invoke it with the `--help` [248] option:

```
shell> mysql --help
```

This chapter assumes that `mysql` is installed on your machine and that a MySQL server is available to which you can connect. If this is not true, contact your MySQL administrator. (If *you* are the administrator, you need to consult the relevant portions of this manual, such as [Chapter 5, MySQL Server Administration](#).)

This chapter describes the entire process of setting up and using a database. If you are interested only in accessing an existing database, you may want to skip over the sections that describe how to create the database and the tables it contains.

Because this chapter is tutorial in nature, many details are necessarily omitted. Consult the relevant sections of the manual for more information on the topics covered here.

### 3.1. Connecting to and Disconnecting from the Server

To connect to the server, you will usually need to provide a MySQL user name when you invoke `mysql` and, most likely, a password. If the server runs on a machine other than the one where you log in, you will also need to specify a host name. Contact your administrator to find out what connection parameters you should use to connect (that is, what host, user name, and password to use). Once you know the proper parameters, you should be able to connect like this:

```
shell> mysql -h host -u user -p
```

```
Enter password: *****
```

`host` and `user` represent the host name where your MySQL server is running and the user name of your MySQL account. Substitute appropriate values for your setup. The `*****` represents your password; enter it when `mysql` displays the `Enter password:` prompt.

If that works, you should see some introductory information followed by a `mysql>` prompt:

```
shell> mysql -h host -u user -p
Enter password: *****
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 25338 to server version: 5.6.12-standard

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>
```

The `mysql>` prompt tells you that `mysql` is ready for you to enter commands.

If you are logging in on the same machine that MySQL is running on, you can omit the host, and simply use the following:

```
shell> mysql -u user -p
```

If, when you attempt to log in, you get an error message such as `ERROR 2002 (HY000): Can't connect to local MySQL server through socket '/tmp/mysql.sock' (2)`, it means that the MySQL server daemon (Unix) or service (Windows) is not running. Consult the administrator or see the section of [Chapter 2, Installing and Upgrading MySQL](#) that is appropriate to your operating system.

For help with other problems often encountered when trying to log in, see [Section C.5.2, "Common Errors When Using MySQL Programs"](#).

Some MySQL installations permit users to connect as the anonymous (unnamed) user to the server running on the local host. If this is the case on your machine, you should be able to connect to that server by invoking `mysql` without any options:

```
shell> mysql
```

After you have connected successfully, you can disconnect any time by typing `QUIT` (or `\q`) at the `mysql>` prompt:

```
mysql> QUIT
Bye
```

On Unix, you can also disconnect by pressing Control+D.

Most examples in the following sections assume that you are connected to the server. They indicate this by the `mysql>` prompt.

## 3.2. Entering Queries

Make sure that you are connected to the server, as discussed in the previous section. Doing so does not in itself select any database to work with, but that is okay. At this point, it is more important to find out a little about how to issue queries than to jump right in creating tables, loading data into them, and retrieving data from them. This section describes the basic principles of entering commands, using several queries you can try out to familiarize yourself with how `mysql` works.

Here is a simple command that asks the server to tell you its version number and the current date. Type it in as shown here following the `mysql>` prompt and press Enter:

```
mysql> SELECT VERSION(), CURRENT_DATE;
+-----+-----+
```

```
| VERSION() | CURRENT_DATE |
+-----+-----+
| 5.6.1-m4-log | 2010-08-06 |
+-----+-----+
1 row in set (0.01 sec)
mysql>
```

This query illustrates several things about `mysql`:

- A command normally consists of an SQL statement followed by a semicolon. (There are some exceptions where a semicolon may be omitted. `QUIT`, mentioned earlier, is one of them. We'll get to others later.)
- When you issue a command, `mysql` sends it to the server for execution and displays the results, then prints another `mysql>` prompt to indicate that it is ready for another command.
- `mysql` displays query output in tabular form (rows and columns). The first row contains labels for the columns. The rows following are the query results. Normally, column labels are the names of the columns you fetch from database tables. If you're retrieving the value of an expression rather than a table column (as in the example just shown), `mysql` labels the column using the expression itself.
- `mysql` shows how many rows were returned and how long the query took to execute, which gives you a rough idea of server performance. These values are imprecise because they represent wall clock time (not CPU or machine time), and because they are affected by factors such as server load and network latency. (For brevity, the "rows in set" line is sometimes not shown in the remaining examples in this chapter.)

Keywords may be entered in any lettercase. The following queries are equivalent:

```
mysql> SELECT VERSION(), CURRENT_DATE;
mysql> select version(), current_date;
mysql> SeLeCt vErSiOn(), current_DATE;
```

Here is another query. It demonstrates that you can use `mysql` as a simple calculator:

```
mysql> SELECT SIN(PI()/4), (4+1)*5;
+-----+-----+
| SIN(PI()/4) | (4+1)*5 |
+-----+-----+
| 0.70710678118655 | 25 |
+-----+-----+
1 row in set (0.02 sec)
```

The queries shown thus far have been relatively short, single-line statements. You can even enter multiple statements on a single line. Just end each one with a semicolon:

```
mysql> SELECT VERSION(); SELECT NOW();
+-----+
| VERSION() |
+-----+
| 5.6.1-m4-log |
+-----+
1 row in set (0.00 sec)
+-----+
| NOW() |
+-----+
| 2010-08-06 12:17:13 |
+-----+
1 row in set (0.00 sec)
```

A command need not be given all on a single line, so lengthy commands that require several lines are not a problem. `mysql` determines where your statement ends by looking for the terminating semicolon, not by looking for the end of the input line. (In other words, `mysql` accepts free-format input: it collects input lines but does not execute them until it sees the semicolon.)

Here is a simple multiple-line statement:

```
mysql> SELECT
-> USER()
-> ,
-> CURRENT_DATE;
+-----+-----+
| USER()      | CURRENT_DATE |
+-----+-----+
| jon@localhost | 2010-08-06   |
+-----+-----+
```

In this example, notice how the prompt changes from `mysql>` to `->` after you enter the first line of a multiple-line query. This is how `mysql` indicates that it has not yet seen a complete statement and is waiting for the rest. The prompt is your friend, because it provides valuable feedback. If you use that feedback, you can always be aware of what `mysql` is waiting for.

If you decide you do not want to execute a command that you are in the process of entering, cancel it by typing `\c`:

```
mysql> SELECT
-> USER()
-> \c
mysql>
```

Here, too, notice the prompt. It switches back to `mysql>` after you type `\c`, providing feedback to indicate that `mysql` is ready for a new command.

The following table shows each of the prompts you may see and summarizes what they mean about the state that `mysql` is in.

Prompt	Meaning
<code>mysql&gt;</code>	Ready for new command.
<code>-&gt;</code>	Waiting for next line of multiple-line command.
<code>'&gt;</code>	Waiting for next line, waiting for completion of a string that began with a single quote (“'”).
<code>"&gt;</code>	Waiting for next line, waiting for completion of a string that began with a double quote (“”).
<code>`&gt;</code>	Waiting for next line, waiting for completion of an identifier that began with a backtick (“`”).
<code>/*&gt;</code>	Waiting for next line, waiting for completion of a comment that began with <code>/*</code> .

Multiple-line statements commonly occur by accident when you intend to issue a command on a single line, but forget the terminating semicolon. In this case, `mysql` waits for more input:

```
mysql> SELECT USER()
->
```

If this happens to you (you think you've entered a statement but the only response is a `->` prompt), most likely `mysql` is waiting for the semicolon. If you don't notice what the prompt is telling you, you might sit there for a while before realizing what you need to do. Enter a semicolon to complete the statement, and `mysql` executes it:

```
mysql> SELECT USER()
-> ;
+-----+-----+
| USER()      |
+-----+-----+
| jon@localhost |
+-----+-----+
```

The `'>` and `">` prompts occur during string collection (another way of saying that MySQL is waiting for completion of a string). In MySQL, you can write strings surrounded by either “'” or “”” characters (for example, `'hello'` or `"goodbye"`), and `mysql` lets you enter strings that span multiple lines. When you see a `'>` or `">` prompt, it means that you have entered a line containing a string that begins with a “'” or “”” quote character, but have not yet entered the matching quote that terminates the string. This often indicates that you have inadvertently left out a quote character. For example:

```
mysql> SELECT * FROM my_table WHERE name = 'Smith AND age < 30;
'>
```

If you enter this `SELECT` statement, then press **Enter** and wait for the result, nothing happens. Instead of wondering why this query takes so long, notice the clue provided by the `'>` prompt. It tells you that `mysql` expects to see the rest of an unterminated string. (Do you see the error in the statement? The string `'Smith` is missing the second single quotation mark.)

At this point, what do you do? The simplest thing is to cancel the command. However, you cannot just type `\c` in this case, because `mysql` interprets it as part of the string that it is collecting. Instead, enter the closing quote character (so `mysql` knows you've finished the string), then type `\c`:

```
mysql> SELECT * FROM my_table WHERE name = 'Smith AND age < 30;
'> '\c
mysql>
```

The prompt changes back to `mysql>`, indicating that `mysql` is ready for a new command.

The ``>` prompt is similar to the `'>` and `">` prompts, but indicates that you have begun but not completed a backtick-quoted identifier.

It is important to know what the `'>`, `">`, and ``>` prompts signify, because if you mistakenly enter an unterminated string, any further lines you type appear to be ignored by `mysql`—including a line containing `QUIT`. This can be quite confusing, especially if you do not know that you need to supply the terminating quote before you can cancel the current command.

### 3.3. Creating and Using a Database

Once you know how to enter commands, you are ready to access a database.

Suppose that you have several pets in your home (your menagerie) and you would like to keep track of various types of information about them. You can do so by creating tables to hold your data and loading them with the desired information. Then you can answer different sorts of questions about your animals by retrieving data from the tables. This section shows you how to perform the following operations:

- Create a database
- Create a table
- Load data into the table
- Retrieve data from the table in various ways
- Use multiple tables

The menagerie database is simple (deliberately), but it is not difficult to think of real-world situations in which a similar type of database might be used. For example, a database like this could be used by a farmer to keep track of livestock, or by a veterinarian to keep track of patient records. A menagerie distribution containing some of the queries and sample data used in the following sections can be obtained from the MySQL Web site. It is available in both compressed `tar` file and Zip formats at <http://dev.mysql.com/doc/>.

Use the `SHOW` statement to find out what databases currently exist on the server:

```
mysql> SHOW DATABASES;
+-----+
| Database |
+-----+
| mysql    |
| test     |
| tmp      |
```

```
+-----+
```

The `mysql` database describes user access privileges. The `test` database often is available as a workspace for users to try things out.

The list of databases displayed by the statement may be different on your machine; `SHOW DATABASES` does not show databases that you have no privileges for if you do not have the `SHOW DATABASES` privilege. See [Section 13.7.5.15, “SHOW DATABASES Syntax”](#).

If the `test` database exists, try to access it:

```
mysql> USE test
Database changed
```

`USE`, like `QUIT`, does not require a semicolon. (You can terminate such statements with a semicolon if you like; it does no harm.) The `USE` statement is special in another way, too: it must be given on a single line.

You can use the `test` database (if you have access to it) for the examples that follow, but anything you create in that database can be removed by anyone else with access to it. For this reason, you should probably ask your MySQL administrator for permission to use a database of your own. Suppose that you want to call yours `menagerie`. The administrator needs to execute a command like this:

```
mysql> GRANT ALL ON menagerie.* TO 'your_mysql_name'@'your_client_host';
```

where `your_mysql_name` is the MySQL user name assigned to you and `your_client_host` is the host from which you connect to the server.

### 3.3.1. Creating and Selecting a Database

If the administrator creates your database for you when setting up your permissions, you can begin using it. Otherwise, you need to create it yourself:

```
mysql> CREATE DATABASE menagerie;
```

Under Unix, database names are case sensitive (unlike SQL keywords), so you must always refer to your database as `menagerie`, not as `Menagerie`, `MENAGERIE`, or some other variant. This is also true for table names. (Under Windows, this restriction does not apply, although you must refer to databases and tables using the same lettercase throughout a given query. However, for a variety of reasons, the recommended best practice is always to use the same lettercase that was used when the database was created.)



#### Note

If you get an error such as `ERROR 1044 (42000): Access denied for user 'monty'@'localhost' to database 'menagerie' when attempting to create a database`, this means that your user account does not have the necessary privileges to do so. Discuss this with the administrator or see [Section 6.2, “The MySQL Access Privilege System”](#).

Creating a database does not select it for use; you must do that explicitly. To make `menagerie` the current database, use this command:

```
mysql> USE menagerie
Database changed
```

Your database needs to be created only once, but you must select it for use each time you begin a `mysql` session. You can do this by issuing a `USE` statement as shown in the example. Alternatively, you can select the database on the command line when you invoke `mysql`. Just specify its name after any connection parameters that you might need to provide. For example:

```
shell> mysql -h host -u user -p menagerie
Enter password: *****
```



**Important**

`menagerie` in the command just shown is **not** your password. If you want to supply your password on the command line after the `-p` option, you must do so with no intervening space (for example, as `-pmypassword`, not as `-p mypassword`). However, putting your password on the command line is not recommended, because doing so exposes it to snooping by other users logged in on your machine.

**Note**

You can see at any time which database is currently selected using `SELECT DATABASE() [1268]`.

### 3.3.2. Creating a Table

Creating the database is the easy part, but at this point it is empty, as `SHOW TABLES` tells you:

```
mysql> SHOW TABLES;
Empty set (0.00 sec)
```

The harder part is deciding what the structure of your database should be: what tables you need and what columns should be in each of them.

You want a table that contains a record for each of your pets. This can be called the `pet` table, and it should contain, as a bare minimum, each animal's name. Because the name by itself is not very interesting, the table should contain other information. For example, if more than one person in your family keeps pets, you might want to list each animal's owner. You might also want to record some basic descriptive information such as species and sex.

How about age? That might be of interest, but it is not a good thing to store in a database. Age changes as time passes, which means you'd have to update your records often. Instead, it is better to store a fixed value such as date of birth. Then, whenever you need age, you can calculate it as the difference between the current date and the birth date. MySQL provides functions for doing date arithmetic, so this is not difficult. Storing birth date rather than age has other advantages, too:

- You can use the database for tasks such as generating reminders for upcoming pet birthdays. (If you think this type of query is somewhat silly, note that it is the same question you might ask in the context of a business database to identify clients to whom you need to send out birthday greetings in the current week or month, for that computer-assisted personal touch.)
- You can calculate age in relation to dates other than the current date. For example, if you store death date in the database, you can easily calculate how old a pet was when it died.

You can probably think of other types of information that would be useful in the `pet` table, but the ones identified so far are sufficient: name, owner, species, sex, birth, and death.

Use a `CREATE TABLE` statement to specify the layout of your table:

```
mysql> CREATE TABLE pet (name VARCHAR(20), owner VARCHAR(20),
-> species VARCHAR(20), sex CHAR(1), birth DATE, death DATE);
```

`VARCHAR` is a good choice for the `name`, `owner`, and `species` columns because the column values vary in length. The lengths in those column definitions need not all be the same, and need not be `20`. You can normally pick any length from `1` to `65535`, whatever seems most reasonable to you. If you make a poor choice and it turns out later that you need a longer field, MySQL provides an `ALTER TABLE` statement.

Several types of values can be chosen to represent sex in animal records, such as `'m'` and `'f'`, or perhaps `'male'` and `'female'`. It is simplest to use the single characters `'m'` and `'f'`.

The use of the `DATE` data type for the `birth` and `death` columns is a fairly obvious choice.

Once you have created a table, `SHOW TABLES` should produce some output:

```
mysql> SHOW TABLES;
+-----+
| Tables in menagerie |
+-----+
| pet                |
+-----+
```

To verify that your table was created the way you expected, use a `DESCRIBE` statement:

```
mysql> DESCRIBE pet;
+-----+-----+-----+-----+-----+-----+
| Field | Type   | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| name  | varchar(20) | YES |     | NULL    |       |
| owner | varchar(20) | YES |     | NULL    |       |
| species | varchar(20) | YES |     | NULL    |       |
| sex   | char(1)   | YES |     | NULL    |       |
| birth | date      | YES |     | NULL    |       |
| death | date      | YES |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
```

You can use `DESCRIBE` any time, for example, if you forget the names of the columns in your table or what types they have.

For more information about MySQL data types, see [Chapter 11, Data Types](#).

### 3.3.3. Loading Data into a Table

After creating your table, you need to populate it. The `LOAD DATA` and `INSERT` statements are useful for this.

Suppose that your pet records can be described as shown here. (Observe that MySQL expects dates in `'YYYY-MM-DD'` format; this may be different from what you are used to.)

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	
Claws	Gwen	cat	m	1994-03-17	
Buffy	Harold	dog	f	1989-05-13	
Fang	Benny	dog	m	1990-08-27	
Bowser	Diane	dog	m	1979-08-31	1995-07-29
Chirpy	Gwen	bird	f	1998-09-11	
Whistler	Gwen	bird		1997-12-09	
Slim	Benny	snake	m	1996-04-29	

Because you are beginning with an empty table, an easy way to populate it is to create a text file containing a row for each of your animals, then load the contents of the file into the table with a single statement.

You could create a text file `pet.txt` containing one record per line, with values separated by tabs, and given in the order in which the columns were listed in the `CREATE TABLE` statement. For missing values (such as unknown sexes or death dates for animals that are still living), you can use `NULL` values. To represent these in your text file, use `\N` (backslash, capital-N). For example, the record for Whistler the bird would look like this (where the whitespace between values is a single tab character):

```
Whistler      Gwen      bird      \N      1997-12-09      \N
```

To load the text file `pet.txt` into the `pet` table, use this statement:

```
mysql> LOAD DATA LOCAL INFILE '/path/pet.txt' INTO TABLE pet;
```

If you created the file on Windows with an editor that uses `\r\n` as a line terminator, you should use this statement instead:

```
mysql> LOAD DATA LOCAL INFILE '/path/pet.txt' INTO TABLE pet
-> LINES TERMINATED BY '\r\n';
```

(On an Apple machine running OS X, you would likely want to use `LINES TERMINATED BY '\r'`.)

You can specify the column value separator and end of line marker explicitly in the `LOAD DATA` statement if you wish, but the defaults are tab and linefeed. These are sufficient for the statement to read the file `pet.txt` properly.

If the statement fails, it is likely that your MySQL installation does not have local file capability enabled by default. See [Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”](#), for information on how to change this.

When you want to add new records one at a time, the `INSERT` statement is useful. In its simplest form, you supply values for each column, in the order in which the columns were listed in the `CREATE TABLE` statement. Suppose that Diane gets a new hamster named “Puffball.” You could add a new record using an `INSERT` statement like this:

```
mysql> INSERT INTO pet
-> VALUES ('Puffball','Diane','hamster','f','1999-03-30',NULL);
```

String and date values are specified as quoted strings here. Also, with `INSERT`, you can insert `NULL` directly to represent a missing value. You do not use `\N` like you do with `LOAD DATA`.

From this example, you should be able to see that there would be a lot more typing involved to load your records initially using several `INSERT` statements rather than a single `LOAD DATA` statement.

### 3.3.4. Retrieving Information from a Table

The `SELECT` statement is used to pull information from a table. The general form of the statement is:

```
SELECT what_to_select
FROM which_table
WHERE conditions_to_satisfy;
```

*what\_to\_select* indicates what you want to see. This can be a list of columns, or `*` to indicate “all columns.” *which\_table* indicates the table from which you want to retrieve data. The `WHERE` clause is optional. If it is present, *conditions\_to\_satisfy* specifies one or more conditions that rows must satisfy to qualify for retrieval.

#### 3.3.4.1. Selecting All Data

The simplest form of `SELECT` retrieves everything from a table:

```
mysql> SELECT * FROM pet;
```

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	NULL
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL
Fang	Benny	dog	m	1990-08-27	NULL
Bowser	Diane	dog	m	1979-08-31	1995-07-29
Chirpy	Gwen	bird	f	1998-09-11	NULL
Whistler	Gwen	bird	NULL	1997-12-09	NULL
Slim	Benny	snake	m	1996-04-29	NULL
Puffball	Diane	hamster	f	1999-03-30	NULL

This form of `SELECT` is useful if you want to review your entire table, for example, after you've just loaded it with your initial data set. For example, you may happen to think that the birth date for Bowser

doesn't seem quite right. Consulting your original pedigree papers, you find that the correct birth year should be 1989, not 1979.

There are at least two ways to fix this:

- Edit the file `pet.txt` to correct the error, then empty the table and reload it using `DELETE` and `LOAD DATA`:

```
mysql> DELETE FROM pet;
mysql> LOAD DATA LOCAL INFILE 'pet.txt' INTO TABLE pet;
```

However, if you do this, you must also re-enter the record for Puffball.

- Fix only the erroneous record with an `UPDATE` statement:

```
mysql> UPDATE pet SET birth = '1989-08-31' WHERE name = 'Bowser';
```

The `UPDATE` changes only the record in question and does not require you to reload the table.

### 3.3.4.2. Selecting Particular Rows

As shown in the preceding section, it is easy to retrieve an entire table. Just omit the `WHERE` clause from the `SELECT` statement. But typically you don't want to see the entire table, particularly when it becomes large. Instead, you're usually more interested in answering a particular question, in which case you specify some constraints on the information you want. Let's look at some selection queries in terms of questions about your pets that they answer.

You can select only particular rows from your table. For example, if you want to verify the change that you made to Bowser's birth date, select Bowser's record like this:

```
mysql> SELECT * FROM pet WHERE name = 'Bowser';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
| Bowser | Diane | dog      | m   | 1989-08-31 | 1995-07-29 |
+-----+-----+-----+-----+-----+-----+
```

The output confirms that the year is correctly recorded as 1989, not 1979.

String comparisons normally are case-insensitive, so you can specify the name as `'bowser'`, `'BOWSER'`, and so forth. The query result is the same.

You can specify conditions on any column, not just `name`. For example, if you want to know which animals were born during or after 1998, test the `birth` column:

```
mysql> SELECT * FROM pet WHERE birth >= '1998-1-1';
+-----+-----+-----+-----+-----+-----+
| name      | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
| Chirpy    | Gwen  | bird    | f   | 1998-09-11 | NULL       |
| Puffball  | Diane | hamster | f   | 1999-03-30 | NULL       |
+-----+-----+-----+-----+-----+-----+
```

You can combine conditions, for example, to locate female dogs:

```
mysql> SELECT * FROM pet WHERE species = 'dog' AND sex = 'f';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
| Buffy | Harold | dog      | f   | 1989-05-13 | NULL       |
+-----+-----+-----+-----+-----+-----+
```

The preceding query uses the `AND` [1171] logical operator. There is also an `OR` [1172] operator:

```
mysql> SELECT * FROM pet WHERE species = 'snake' OR species = 'bird';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
```

Chirpy	Gwen	bird	f	1998-09-11	NULL
Whistler	Gwen	bird	NULL	1997-12-09	NULL
Slim	Benny	snake	m	1996-04-29	NULL

**AND** [1171] and **OR** [1172] may be intermixed, although **AND** [1171] has higher precedence than **OR** [1172]. If you use both operators, it is a good idea to use parentheses to indicate explicitly how conditions should be grouped:

```
mysql> SELECT * FROM pet WHERE (species = 'cat' AND sex = 'm')
-> OR (species = 'dog' AND sex = 'f');
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

### 3.3.4.3. Selecting Particular Columns

If you do not want to see entire rows from your table, just name the columns in which you are interested, separated by commas. For example, if you want to know when your animals were born, select the `name` and `birth` columns:

```
mysql> SELECT name, birth FROM pet;
```

name	birth
Fluffy	1993-02-04
Claws	1994-03-17
Buffy	1989-05-13
Fang	1990-08-27
Bowser	1989-08-31
Chirpy	1998-09-11
Whistler	1997-12-09
Slim	1996-04-29
Puffball	1999-03-30

To find out who owns pets, use this query:

```
mysql> SELECT owner FROM pet;
```

owner
Harold
Gwen
Harold
Benny
Diane
Gwen
Gwen
Benny
Diane

Notice that the query simply retrieves the `owner` column from each record, and some of them appear more than once. To minimize the output, retrieve each unique output record just once by adding the keyword **DISTINCT**:

```
mysql> SELECT DISTINCT owner FROM pet;
```

owner
Benny
Diane
Gwen
Harold

You can use a `WHERE` clause to combine row selection with column selection. For example, to get birth dates for dogs and cats only, use this query:

```
mysql> SELECT name, species, birth FROM pet
-> WHERE species = 'dog' OR species = 'cat';
```

name	species	birth
Fluffy	cat	1993-02-04
Claws	cat	1994-03-17
Buffy	dog	1989-05-13
Fang	dog	1990-08-27
Bowser	dog	1989-08-31

### 3.3.4.4. Sorting Rows

You may have noticed in the preceding examples that the result rows are displayed in no particular order. It is often easier to examine query output when the rows are sorted in some meaningful way. To sort a result, use an `ORDER BY` clause.

Here are animal birthdays, sorted by date:

```
mysql> SELECT name, birth FROM pet ORDER BY birth;
```

name	birth
Buffy	1989-05-13
Bowser	1989-08-31
Fang	1990-08-27
Fluffy	1993-02-04
Claws	1994-03-17
Slim	1996-04-29
Whistler	1997-12-09
Chirpy	1998-09-11
Puffball	1999-03-30

On character type columns, sorting—like all other comparison operations—is normally performed in a case-insensitive fashion. This means that the order is undefined for columns that are identical except for their case. You can force a case-sensitive sort for a column by using `BINARY [1244]` like so: `ORDER BY BINARY col_name`.

The default sort order is ascending, with smallest values first. To sort in reverse (descending) order, add the `DESC` keyword to the name of the column you are sorting by:

```
mysql> SELECT name, birth FROM pet ORDER BY birth DESC;
```

name	birth
Puffball	1999-03-30
Chirpy	1998-09-11
Whistler	1997-12-09
Slim	1996-04-29
Claws	1994-03-17
Fluffy	1993-02-04
Fang	1990-08-27
Bowser	1989-08-31
Buffy	1989-05-13

You can sort on multiple columns, and you can sort different columns in different directions. For example, to sort by type of animal in ascending order, then by birth date within animal type in descending order (youngest animals first), use the following query:

```
mysql> SELECT name, species, birth FROM pet
-> ORDER BY species, birth DESC;
```

name	species	birth
Chirpy	bird	1998-09-11
Whistler	bird	1997-12-09
Claws	cat	1994-03-17
Fluffy	cat	1993-02-04
Fang	dog	1990-08-27
Bowser	dog	1989-08-31
Buffy	dog	1989-05-13
Puffball	hamster	1999-03-30
Slim	snake	1996-04-29

The `DESC` keyword applies only to the column name immediately preceding it (`birth`); it does not affect the `species` column sort order.

### 3.3.4.5. Date Calculations

MySQL provides several functions that you can use to perform calculations on dates, for example, to calculate ages or extract parts of dates.

To determine how many years old each of your pets is, use the `TIMESTAMPDIFF()` [1224] function. Its arguments are the unit in which you want the result expressed, and the two date for which to take the difference. The following query shows, for each pet, the birth date, the current date, and the age in years. An *alias* (`age`) is used to make the final output column label more meaningful.

```
mysql> SELECT name, birth, CURDATE(),
->   TIMESTAMPDIFF(YEAR,birth,CURDATE()) AS age
->   FROM pet;
```

name	birth	CURDATE()	age
Fluffy	1993-02-04	2003-08-19	10
Claws	1994-03-17	2003-08-19	9
Buffy	1989-05-13	2003-08-19	14
Fang	1990-08-27	2003-08-19	12
Bowser	1989-08-31	2003-08-19	13
Chirpy	1998-09-11	2003-08-19	4
Whistler	1997-12-09	2003-08-19	5
Slim	1996-04-29	2003-08-19	7
Puffball	1999-03-30	2003-08-19	4

The query works, but the result could be scanned more easily if the rows were presented in some order. This can be done by adding an `ORDER BY name` clause to sort the output by name:

```
mysql> SELECT name, birth, CURDATE(),
->   TIMESTAMPDIFF(YEAR,birth,CURDATE()) AS age
->   FROM pet ORDER BY name;
```

name	birth	CURDATE()	age
Bowser	1989-08-31	2003-08-19	13
Buffy	1989-05-13	2003-08-19	14
Chirpy	1998-09-11	2003-08-19	4
Claws	1994-03-17	2003-08-19	9
Fang	1990-08-27	2003-08-19	12
Fluffy	1993-02-04	2003-08-19	10
Puffball	1999-03-30	2003-08-19	4
Slim	1996-04-29	2003-08-19	7
Whistler	1997-12-09	2003-08-19	5

To sort the output by `age` rather than `name`, just use a different `ORDER BY` clause:

```
mysql> SELECT name, birth, CURDATE(),
->   TIMESTAMPDIFF(YEAR,birth,CURDATE()) AS age
->   FROM pet ORDER BY age;
```



name	birth	CURDATE()	age
Chirpy	1998-09-11	2003-08-19	4
Puffball	1999-03-30	2003-08-19	4
Whistler	1997-12-09	2003-08-19	5
Slim	1996-04-29	2003-08-19	7
Claws	1994-03-17	2003-08-19	9
Fluffy	1993-02-04	2003-08-19	10
Fang	1990-08-27	2003-08-19	12
Bowser	1989-08-31	2003-08-19	13
Buffy	1989-05-13	2003-08-19	14

A similar query can be used to determine age at death for animals that have died. You determine which animals these are by checking whether the `death` value is `NULL`. Then, for those with non-`NULL` values, compute the difference between the `death` and `birth` values:

```
mysql> SELECT name, birth, death,
-> TIMESTAMPDIFF(YEAR,birth,death) AS age
-> FROM pet WHERE death IS NOT NULL ORDER BY age;
```

name	birth	death	age
Bowser	1989-08-31	1995-07-29	5

The query uses `death IS NOT NULL` rather than `death <> NULL` because `NULL` is a special value that cannot be compared using the usual comparison operators. This is discussed later. See [Section 3.3.4.6, “Working with NULL Values”](#).

What if you want to know which animals have birthdays next month? For this type of calculation, year and day are irrelevant; you simply want to extract the month part of the `birth` column. MySQL provides several functions for extracting parts of dates, such as `YEAR()` [1229], `MONTH()` [1220], and `DAYOFMONTH()` [1217]. `MONTH()` [1220] is the appropriate function here. To see how it works, run a simple query that displays the value of both `birth` and `MONTH(birth)` [1220]:

```
mysql> SELECT name, birth, MONTH(birth) FROM pet;
```

name	birth	MONTH(birth)
Fluffy	1993-02-04	2
Claws	1994-03-17	3
Buffy	1989-05-13	5
Fang	1990-08-27	8
Bowser	1989-08-31	8
Chirpy	1998-09-11	9
Whistler	1997-12-09	12
Slim	1996-04-29	4
Puffball	1999-03-30	3

Finding animals with birthdays in the upcoming month is also simple. Suppose that the current month is April. Then the month value is `4` and you can look for animals born in May (month `5`) like this:

```
mysql> SELECT name, birth FROM pet WHERE MONTH(birth) = 5;
```

name	birth
Buffy	1989-05-13

There is a small complication if the current month is December. You cannot merely add one to the month number (`12`) and look for animals born in month `13`, because there is no such month. Instead, you look for animals born in January (month `1`).

You can write the query so that it works no matter what the current month is, so that you do not have to use the number for a particular month. `DATE_ADD()` [1213] enables you to add a time interval to

a given date. If you add a month to the value of `CURDATE()` [1212], then extract the month part with `MONTH()` [1220], the result produces the month in which to look for birthdays:

```
mysql> SELECT name, birth FROM pet
  -> WHERE MONTH(birth) = MONTH(DATE_ADD(CURDATE(),INTERVAL 1 MONTH));
```

A different way to accomplish the same task is to add 1 to get the next month after the current one after using the modulo function (`MOD`) to wrap the month value to 0 if it is currently 12:

```
mysql> SELECT name, birth FROM pet
  -> WHERE MONTH(birth) = MOD(MONTH(CURDATE()), 12) + 1;
```

`MONTH()` [1220] returns a number between 1 and 12. And `MOD(something,12)` [1205] returns a number between 0 and 11. So the addition has to be after the `MOD()` [1205], otherwise we would go from November (11) to January (1).

### 3.3.4.6. Working with `NULL` Values

The `NULL` value can be surprising until you get used to it. Conceptually, `NULL` means “a missing unknown value” and it is treated somewhat differently from other values.

To test for `NULL`, use the `IS NULL` [1168] and `IS NOT NULL` [1168] operators, as shown here:

```
mysql> SELECT 1 IS NULL, 1 IS NOT NULL;
+-----+-----+
| 1 IS NULL | 1 IS NOT NULL |
+-----+-----+
|          0 |                1 |
+-----+-----+
```

You cannot use arithmetic comparison operators such as `=` [1166], `<` [1167], or `<>` [1167] to test for `NULL`. To demonstrate this for yourself, try the following query:

```
mysql> SELECT 1 = NULL, 1 <> NULL, 1 < NULL, 1 > NULL;
+-----+-----+-----+-----+
| 1 = NULL | 1 <> NULL | 1 < NULL | 1 > NULL |
+-----+-----+-----+-----+
|        NULL |        NULL |        NULL |        NULL |
+-----+-----+-----+-----+
```

Because the result of any arithmetic comparison with `NULL` is also `NULL`, you cannot obtain any meaningful results from such comparisons.

In MySQL, 0 or `NULL` means false and anything else means true. The default truth value from a boolean operation is 1.

This special treatment of `NULL` is why, in the previous section, it was necessary to determine which animals are no longer alive using `death IS NOT NULL` instead of `death <> NULL`.

Two `NULL` values are regarded as equal in a `GROUP BY`.

When doing an `ORDER BY`, `NULL` values are presented first if you do `ORDER BY ... ASC` and last if you do `ORDER BY ... DESC`.

A common error when working with `NULL` is to assume that it is not possible to insert a zero or an empty string into a column defined as `NOT NULL`, but this is not the case. These are in fact values, whereas `NULL` means “not having a value.” You can test this easily enough by using `IS [NOT] NULL` as shown:

```
mysql> SELECT 0 IS NULL, 0 IS NOT NULL, '' IS NULL, '' IS NOT NULL;
+-----+-----+-----+-----+
| 0 IS NULL | 0 IS NOT NULL | '' IS NULL | '' IS NOT NULL |
+-----+-----+-----+-----+
|          0 |                1 |          0 |                1 |
+-----+-----+-----+-----+
```

Thus it is entirely possible to insert a zero or empty string into a `NOT NULL` column, as these are in fact `NOT NULL`. See [Section C.5.5.3, “Problems with NULL Values”](#).

### 3.3.4.7. Pattern Matching

MySQL provides standard SQL pattern matching as well as a form of pattern matching based on extended regular expressions similar to those used by Unix utilities such as `vi`, `grep`, and `sed`.

SQL pattern matching enables you to use “`_`” to match any single character and “`%`” to match an arbitrary number of characters (including zero characters). In MySQL, SQL patterns are case-insensitive by default. Some examples are shown here. You do not use `=` or `<>` when you use SQL patterns; use the `LIKE [1190]` or `NOT LIKE [1192]` comparison operators instead.

To find names beginning with “`b`”:

```
mysql> SELECT * FROM pet WHERE name LIKE 'b%';
```

name	owner	species	sex	birth	death
Buffy	Harold	dog	f	1989-05-13	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29

To find names ending with “`fy`”:

```
mysql> SELECT * FROM pet WHERE name LIKE '%fy';
```

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

To find names containing a “`w`”:

```
mysql> SELECT * FROM pet WHERE name LIKE '%w%';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29
Whistler	Gwen	bird	NULL	1997-12-09	NULL

To find names containing exactly five characters, use five instances of the “`_`” pattern character:

```
mysql> SELECT * FROM pet WHERE name LIKE '_____';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

The other type of pattern matching provided by MySQL uses extended regular expressions. When you test for a match for this type of pattern, use the `REGEXP [1193]` and `NOT REGEXP [1193]` operators (or `RLIKE [1193]` and `NOT RLIKE [1193]`, which are synonyms).

The following list describes some characteristics of extended regular expressions:

- “`.`” matches any single character.
- A character class “[`...`]” matches any character within the brackets. For example, “[`abc`]” matches “`a`”, “`b`”, or “`c`”. To name a range of characters, use a dash. “[`a-z`]” matches any letter, whereas “[`0-9`]” matches any digit.

- “`*`” matches zero or more instances of the thing preceding it. For example, “`x*`” matches any number of “`x`” characters, “`[0-9]*`” matches any number of digits, and “`.*`” matches any number of anything.
- A [REGEXP \[1193\]](#) pattern match succeeds if the pattern matches anywhere in the value being tested. (This differs from a [LIKE \[1190\]](#) pattern match, which succeeds only if the pattern matches the entire value.)
- To anchor a pattern so that it must match the beginning or end of the value being tested, use “`^`” at the beginning or “`$`” at the end of the pattern.

To demonstrate how extended regular expressions work, the [LIKE \[1190\]](#) queries shown previously are rewritten here to use [REGEXP \[1193\]](#).

To find names beginning with “`b`”, use “`^`” to match the beginning of the name:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^b';
```

name	owner	species	sex	birth	death
Buffy	Harold	dog	f	1989-05-13	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29

If you really want to force a [REGEXP \[1193\]](#) comparison to be case sensitive, use the [BINARY \[1244\]](#) keyword to make one of the strings a binary string. This query matches only lowercase “`b`” at the beginning of a name:

```
mysql> SELECT * FROM pet WHERE name REGEXP BINARY '^b';
```

To find names ending with “`fy`”, use “`$`” to match the end of the name:

```
mysql> SELECT * FROM pet WHERE name REGEXP 'fy$';
```

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

To find names containing a “`w`”, use this query:

```
mysql> SELECT * FROM pet WHERE name REGEXP 'w';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29
Whistler	Gwen	bird	NULL	1997-12-09	NULL

Because a regular expression pattern matches if it occurs anywhere in the value, it is not necessary in the previous query to put a wildcard on either side of the pattern to get it to match the entire value like it would be if you used an SQL pattern.

To find names containing exactly five characters, use “`^`” and “`$`” to match the beginning and end of the name, and five instances of “`.`” in between:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^.....$';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

You could also write the previous query using the `{n}` (“repeat-*n*-times”) operator:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^.{5}$';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

Section 12.5.2, “Regular Expressions”, provides more information about the syntax for regular expressions.

### 3.3.4.8. Counting Rows

Databases are often used to answer the question, “How often does a certain type of data occur in a table?” For example, you might want to know how many pets you have, or how many pets each owner has, or you might want to perform various kinds of census operations on your animals.

Counting the total number of animals you have is the same question as “How many rows are in the `pet` table?” because there is one record per pet. `COUNT(*)` [1283] counts the number of rows, so the query to count your animals looks like this:

```
mysql> SELECT COUNT(*) FROM pet;
```

COUNT(*)
9

Earlier, you retrieved the names of the people who owned pets. You can use `COUNT()` [1283] if you want to find out how many pets each owner has:

```
mysql> SELECT owner, COUNT(*) FROM pet GROUP BY owner;
```

owner	COUNT(*)
Benny	2
Diane	2
Gwen	3
Harold	2

The preceding query uses `GROUP BY` to group all records for each `owner`. The use of `COUNT()` [1283] in conjunction with `GROUP BY` is useful for characterizing your data under various groupings. The following examples show different ways to perform animal census operations.

Number of animals per species:

```
mysql> SELECT species, COUNT(*) FROM pet GROUP BY species;
```

species	COUNT(*)
bird	2
cat	2
dog	3
hamster	1
snake	1

Number of animals per sex:

```
mysql> SELECT sex, COUNT(*) FROM pet GROUP BY sex;
```

sex	COUNT(*)
NULL	1
f	4
m	4

(In this output, `NULL` indicates that the sex is unknown.)

Number of animals per combination of species and sex:

```
mysql> SELECT species, sex, COUNT(*) FROM pet GROUP BY species, sex;
```

species	sex	COUNT(*)
bird	NULL	1
bird	f	1
cat	f	1
cat	m	1
dog	f	1
dog	m	2
hamster	f	1
snake	m	1

You need not retrieve an entire table when you use `COUNT()` [1283]. For example, the previous query, when performed just on dogs and cats, looks like this:

```
mysql> SELECT species, sex, COUNT(*) FROM pet
-> WHERE species = 'dog' OR species = 'cat'
-> GROUP BY species, sex;
```

species	sex	COUNT(*)
cat	f	1
cat	m	1
dog	f	1
dog	m	2

Or, if you wanted the number of animals per sex only for animals whose sex is known:

```
mysql> SELECT species, sex, COUNT(*) FROM pet
-> WHERE sex IS NOT NULL
-> GROUP BY species, sex;
```

species	sex	COUNT(*)
bird	f	1
cat	f	1
cat	m	1
dog	f	1
dog	m	2
hamster	f	1
snake	m	1

If you name columns to select in addition to the `COUNT()` [1283] value, a `GROUP BY` clause should be present that names those same columns. Otherwise, the following occurs:

- If the `ONLY_FULL_GROUP_BY` [606] SQL mode is enabled, an error occurs:

```
mysql> SET sql_mode = 'ONLY_FULL_GROUP_BY';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT owner, COUNT(*) FROM pet;
ERROR 1140 (42000): Mixing of GROUP columns (MIN(),MAX(),COUNT()...)
with no GROUP columns is illegal if there is no GROUP BY clause
```

- If `ONLY_FULL_GROUP_BY` [606] is not enabled, the query is processed by treating all rows as a single group, but the value selected for each named column is indeterminate. The server is free to select the value from any row:

```
mysql> SET sql_mode = '';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT owner, COUNT(*) FROM pet;
```

```

+-----+-----+
| owner | COUNT(*) |
+-----+-----+
| Harold |      8 |
+-----+-----+
1 row in set (0.00 sec)

```

See also [Section 12.17.3, “MySQL Extensions to GROUP BY”](#).

### 3.3.4.9. Using More Than one Table

The `pet` table keeps track of which pets you have. If you want to record other information about them, such as events in their lives like visits to the vet or when litters are born, you need another table. What should this table look like? It needs to contain the following information:

- The pet name so that you know which animal each event pertains to.
- A date so that you know when the event occurred.
- A field to describe the event.
- An event type field, if you want to be able to categorize events.

Given these considerations, the `CREATE TABLE` statement for the `event` table might look like this:

```

mysql> CREATE TABLE event (name VARCHAR(20), date DATE,
-> type VARCHAR(15), remark VARCHAR(255));

```

As with the `pet` table, it is easiest to load the initial records by creating a tab-delimited text file containing the following information.

name	date	type	remark
Fluffy	1995-05-15	litter	4 kittens, 3 female, 1 male
Buffy	1993-06-23	litter	5 puppies, 2 female, 3 male
Buffy	1994-06-19	litter	3 puppies, 3 female
Chirpy	1999-03-21	vet	needed beak straightened
Slim	1997-08-03	vet	broken rib
Bowser	1991-10-12	kennel	
Fang	1991-10-12	kennel	
Fang	1998-08-28	birthday	Gave him a new chew toy
Claws	1998-03-17	birthday	Gave him a new flea collar
Whistler	1998-12-09	birthday	First birthday

Load the records like this:

```

mysql> LOAD DATA LOCAL INFILE 'event.txt' INTO TABLE event;

```

Based on what you have learned from the queries that you have run on the `pet` table, you should be able to perform retrievals on the records in the `event` table; the principles are the same. But when is the `event` table by itself insufficient to answer questions you might ask?

Suppose that you want to find out the ages at which each pet had its litters. We saw earlier how to calculate ages from two dates. The litter date of the mother is in the `event` table, but to calculate her age on that date you need her birth date, which is stored in the `pet` table. This means the query requires both tables:

```

mysql> SELECT pet.name,
-> (YEAR(date)-YEAR(birth)) - (RIGHT(date,5)<RIGHT(birth,5)) AS age,

```



```

-> remark
-> FROM pet INNER JOIN event
->   ON pet.name = event.name
-> WHERE event.type = 'litter';
+-----+-----+-----+
| name  | age  | remark                                     |
+-----+-----+-----+
| Fluffy | 2    | 4 kittens, 3 female, 1 male             |
| Buffy  | 4    | 5 puppies, 2 female, 3 male             |
| Buffy  | 5    | 3 puppies, 3 female                       |
+-----+-----+-----+

```

There are several things to note about this query:

- The `FROM` clause joins two tables because the query needs to pull information from both of them.
- When combining (joining) information from multiple tables, you need to specify how records in one table can be matched to records in the other. This is easy because they both have a `name` column. The query uses an `ON` clause to match up records in the two tables based on the `name` values.

The query uses an `INNER JOIN` to combine the tables. An `INNER JOIN` permits rows from either table to appear in the result if and only if both tables meet the conditions specified in the `ON` clause. In this example, the `ON` clause specifies that the `name` column in the `pet` table must match the `name` column in the `event` table. If a name appears in one table but not the other, the row will not appear in the result because the condition in the `ON` clause fails.

- Because the `name` column occurs in both tables, you must be specific about which table you mean when referring to the column. This is done by prepending the table name to the column name.

You need not have two different tables to perform a join. Sometimes it is useful to join a table to itself, if you want to compare records in a table to other records in that same table. For example, to find breeding pairs among your pets, you can join the `pet` table with itself to produce candidate pairs of males and females of like species:

```

mysql> SELECT p1.name, p1.sex, p2.name, p2.sex, p1.species
-> FROM pet AS p1 INNER JOIN pet AS p2
->   ON p1.species = p2.species AND p1.sex = 'f' AND p2.sex = 'm';
+-----+-----+-----+-----+-----+
| name  | sex  | name  | sex  | species |
+-----+-----+-----+-----+-----+
| Fluffy | f    | Claws | m    | cat     |
| Buffy  | f    | Fang  | m    | dog     |
| Buffy  | f    | Bowser | m    | dog     |
+-----+-----+-----+-----+-----+

```

In this query, we specify aliases for the table name to refer to the columns and keep straight which instance of the table each column reference is associated with.

### 3.4. Getting Information About Databases and Tables

What if you forget the name of a database or table, or what the structure of a given table is (for example, what its columns are called)? MySQL addresses this problem through several statements that provide information about the databases and tables it supports.

You have previously seen `SHOW DATABASES`, which lists the databases managed by the server. To find out which database is currently selected, use the `DATABASE()` [1268] function:

```

mysql> SELECT DATABASE();
+-----+
| DATABASE() |
+-----+
| menagerie  |
+-----+

```

If you have not yet selected any database, the result is `NULL`.

To find out what tables the default database contains (for example, when you are not sure about the name of a table), use this command:

```
mysql> SHOW TABLES;
+-----+
| Tables_in_menagerie |
+-----+
| event                |
| pet                  |
+-----+
```

The name of the column in the output produced by this statement is always `Tables_in_db_name`, where `db_name` is the name of the database. See [Section 13.7.5.38, “SHOW TABLES Syntax”](#), for more information.

If you want to find out about the structure of a table, the `DESCRIBE` statement is useful; it displays information about each of a table's columns:

```
mysql> DESCRIBE pet;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| name  | varchar(20)   | YES  |     | NULL    |       |
| owner | varchar(20)   | YES  |     | NULL    |       |
| species | varchar(20)  | YES  |     | NULL    |       |
| sex   | char(1)       | YES  |     | NULL    |       |
| birth | date          | YES  |     | NULL    |       |
| death | date          | YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
```

`Field` indicates the column name, `Type` is the data type for the column, `NULL` indicates whether the column can contain `NULL` values, `Key` indicates whether the column is indexed, and `Default` specifies the column's default value. `Extra` displays special information about columns: If a column was created with the `AUTO_INCREMENT` option, the value will be `auto_increment` rather than empty.

`DESC` is a short form of `DESCRIBE`. See [Section 13.8.1, “DESCRIBE Syntax”](#), for more information.

You can obtain the `CREATE TABLE` statement necessary to create an existing table using the `SHOW CREATE TABLE` statement. See [Section 13.7.5.12, “SHOW CREATE TABLE Syntax”](#).

If you have indexes on a table, `SHOW INDEX FROM tbl_name` produces information about them. See [Section 13.7.5.23, “SHOW INDEX Syntax”](#), for more about this statement.

## 3.5. Using `mysql` in Batch Mode

In the previous sections, you used `mysql` interactively to enter queries and view the results. You can also run `mysql` in batch mode. To do this, put the commands you want to run in a file, then tell `mysql` to read its input from the file:

```
shell> mysql < batch-file
```

If you are running `mysql` under Windows and have some special characters in the file that cause problems, you can do this:

```
C:\> mysql -e "source batch-file"
```

If you need to specify connection parameters on the command line, the command might look like this:

```
shell> mysql -h host -u user -p < batch-file
Enter password: *****
```

When you use `mysql` this way, you are creating a script file, then executing the script.

If you want the script to continue even if some of the statements in it produce errors, you should use the `--force` [250] command-line option.

Why use a script? Here are a few reasons:

- If you run a query repeatedly (say, every day or every week), making it a script enables you to avoid retyping it each time you execute it.
- You can generate new queries from existing ones that are similar by copying and editing script files.
- Batch mode can also be useful while you're developing a query, particularly for multiple-line commands or multiple-statement sequences of commands. If you make a mistake, you don't have to retype everything. Just edit your script to correct the error, then tell `mysql` to execute it again.
- If you have a query that produces a lot of output, you can run the output through a pager rather than watching it scroll off the top of your screen:

```
shell> mysql < batch-file | more
```

- You can catch the output in a file for further processing:

```
shell> mysql < batch-file > mysql.out
```

- You can distribute your script to other people so that they can also run the commands.
- Some situations do not allow for interactive use, for example, when you run a query from a `cron` job. In this case, you must use batch mode.

The default output format is different (more concise) when you run `mysql` in batch mode than when you use it interactively. For example, the output of `SELECT DISTINCT species FROM pet` looks like this when `mysql` is run interactively:

```
+-----+
| species |
+-----+
| bird    |
| cat     |
| dog     |
| hamster |
| snake   |
+-----+
```

In batch mode, the output looks like this instead:

```
species
bird
cat
dog
hamster
snake
```

If you want to get the interactive output format in batch mode, use `mysql -t`. To echo to the output the commands that are executed, use `mysql -vvv`.

You can also use scripts from the `mysql` prompt by using the `source` command or `\.` command:

```
mysql> source filename;
mysql> \. filename
```

See [Section 4.5.1.5, "Executing SQL Statements from a Text File"](#), for more information.

## 3.6. Examples of Common Queries

Here are examples of how to solve some common problems with MySQL.

Some of the examples use the table `shop` to hold the price of each article (item number) for certain traders (dealers). Supposing that each trader has a single fixed price per article, then (`article`, `dealer`) is a primary key for the records.

Start the command-line tool `mysql` and select a database:

```
shell> mysql your-database-name
```

(In most MySQL installations, you can use the database named `test`).

You can create and populate the example table with these statements:

```
CREATE TABLE shop (
  article INT(4) UNSIGNED ZEROFILL DEFAULT '0000' NOT NULL,
  dealer CHAR(20) DEFAULT '' NOT NULL,
  price DOUBLE(16,2) DEFAULT '0.00' NOT NULL,
  PRIMARY KEY(article, dealer));
INSERT INTO shop VALUES
  (1, 'A', 3.45), (1, 'B', 3.99), (2, 'A', 10.99), (3, 'B', 1.45),
  (3, 'C', 1.69), (3, 'D', 1.25), (4, 'D', 19.95);
```

After issuing the statements, the table should have the following contents:

```
SELECT * FROM shop;
```

article	dealer	price
0001	A	3.45
0001	B	3.99
0002	A	10.99
0003	B	1.45
0003	C	1.69
0003	D	1.25
0004	D	19.95

### 3.6.1. The Maximum Value for a Column

“What is the highest item number?”

```
SELECT MAX(article) AS article FROM shop;
```

article
4

### 3.6.2. The Row Holding the Maximum of a Certain Column

*Task: Find the number, dealer, and price of the most expensive article.*

This is easily done with a subquery:

```
SELECT article, dealer, price
FROM shop
WHERE price=(SELECT MAX(price) FROM shop);
```

article	dealer	price
0004	D	19.95

Other solutions are to use a `LEFT JOIN` or to sort all rows descending by price and get only the first row using the MySQL-specific `LIMIT` clause:

```
SELECT s1.article, s1.dealer, s1.price
FROM shop s1
LEFT JOIN shop s2 ON s1.price < s2.price
WHERE s2.article IS NULL;
```

```
SELECT article, dealer, price
FROM shop
ORDER BY price DESC
LIMIT 1;
```



**Note**

If there were several most expensive articles, each with a price of 19.95, the `LIMIT` solution would show only one of them.

### 3.6.3. Maximum of Column per Group

*Task: Find the highest price per article.*

```
SELECT article, MAX(price) AS price
FROM shop
GROUP BY article;
```

```
+-----+-----+
| article | price |
+-----+-----+
|    0001 |   3.99 |
|    0002 |  10.99 |
|    0003 |   1.69 |
|    0004 |  19.95 |
+-----+-----+
```

### 3.6.4. The Rows Holding the Group-wise Maximum of a Certain Column

*Task: For each article, find the dealer or dealers with the most expensive price.*

This problem can be solved with a subquery like this one:

```
SELECT article, dealer, price
FROM shop s1
WHERE price=(SELECT MAX(s2.price)
              FROM shop s2
              WHERE s1.article = s2.article);
```

```
+-----+-----+-----+
| article | dealer | price |
+-----+-----+-----+
|    0001 | B      |   3.99 |
|    0002 | A      |  10.99 |
|    0003 | C      |   1.69 |
|    0004 | D      |  19.95 |
+-----+-----+-----+
```

The preceding example uses a correlated subquery, which can be inefficient (see [Section 13.2.10.7, "Correlated Subqueries"](#)). Other possibilities for solving the problem are to use an uncorrelated subquery in the `FROM` clause or a `LEFT JOIN`.

Uncorrelated subquery:

```
SELECT s1.article, dealer, s1.price
FROM shop s1
JOIN (
  SELECT article, MAX(price) AS price
  FROM shop
  GROUP BY article) AS s2
ON s1.article = s2.article AND s1.price = s2.price;
```

`LEFT JOIN`:

```
SELECT s1.article, s1.dealer, s1.price
FROM shop s1
LEFT JOIN shop s2 ON s1.article = s2.article AND s1.price < s2.price
WHERE s2.article IS NULL;
```

The `LEFT JOIN` works on the basis that when `s1.price` is at its maximum value, there is no `s2.price` with a greater value and the `s2` rows values will be `NULL`. See [Section 13.2.9.2, “JOIN Syntax”](#).

### 3.6.5. Using User-Defined Variables

You can employ MySQL user variables to remember results without having to store them in temporary variables in the client. (See [Section 9.4, “User-Defined Variables”](#).)

For example, to find the articles with the highest and lowest price you can do this:

```
mysql> SELECT @min_price:=MIN(price),@max_price:=MAX(price) FROM shop;
mysql> SELECT * FROM shop WHERE price=@min_price OR price=@max_price;
+-----+-----+-----+
| article | dealer | price |
+-----+-----+-----+
| 0003   | D     | 1.25  |
| 0004   | D     | 19.95 |
+-----+-----+-----+
```



#### Note

It is also possible to store the name of a database object such as a table or a column in a user variable and then to use this variable in an SQL statement; however, this requires the use of a prepared statement. See [Section 13.5, “SQL Syntax for Prepared Statements”](#), for more information.

### 3.6.6. Using Foreign Keys

In MySQL, `InnoDB` tables support checking of foreign key constraints. See [Section 14.2, “The InnoDB Storage Engine”](#), and [Section 1.8.5.4, “Foreign Key Differences”](#).

A foreign key constraint is not required merely to join two tables. For storage engines other than `InnoDB`, it is possible when defining a column to use a `REFERENCES tbl_name(col_name)` clause, which has no actual effect, and *serves only as a memo or comment to you that the column which you are currently defining is intended to refer to a column in another table*. It is extremely important to realize when using this syntax that:

- MySQL does not perform any sort of `CHECK` to make sure that `col_name` actually exists in `tbl_name` (or even that `tbl_name` itself exists).
- MySQL does not perform any sort of action on `tbl_name` such as deleting rows in response to actions taken on rows in the table which you are defining; in other words, this syntax induces no `ON DELETE` or `ON UPDATE` behavior whatsoever. (Although you can write an `ON DELETE` or `ON UPDATE` clause as part of the `REFERENCES` clause, it is also ignored.)
- This syntax creates a *column*; it does **not** create any sort of index or key.

You can use a column so created as a join column, as shown here:

```
CREATE TABLE person (
  id SMALLINT UNSIGNED NOT NULL AUTO_INCREMENT,
  name CHAR(60) NOT NULL,
  PRIMARY KEY (id)
);

CREATE TABLE shirt (
  id SMALLINT UNSIGNED NOT NULL AUTO_INCREMENT,
  style ENUM('t-shirt', 'polo', 'dress') NOT NULL,
  color ENUM('red', 'blue', 'orange', 'white', 'black') NOT NULL,
  owner SMALLINT UNSIGNED NOT NULL REFERENCES person(id),
  PRIMARY KEY (id)
);

INSERT INTO person VALUES (NULL, 'Antonio Paz');
```

```

SELECT @last := LAST_INSERT_ID();

INSERT INTO shirt VALUES
(NULL, 'polo', 'blue', @last),
(NULL, 'dress', 'white', @last),
(NULL, 't-shirt', 'blue', @last);

INSERT INTO person VALUES (NULL, 'Lilliana Angelovska');

SELECT @last := LAST_INSERT_ID();

INSERT INTO shirt VALUES
(NULL, 'dress', 'orange', @last),
(NULL, 'polo', 'red', @last),
(NULL, 'dress', 'blue', @last),
(NULL, 't-shirt', 'white', @last);

SELECT * FROM person;
+-----+-----+
| id | name                |
+-----+-----+
|  1 | Antonio Paz         |
|  2 | Lilliana Angelovska |
+-----+-----+

SELECT * FROM shirt;
+-----+-----+-----+-----+
| id | style | color | owner |
+-----+-----+-----+-----+
|  1 | polo  | blue  | 1     |
|  2 | dress | white | 1     |
|  3 | t-shirt | blue  | 1     |
|  4 | dress | orange | 2     |
|  5 | polo  | red   | 2     |
|  6 | dress | blue  | 2     |
|  7 | t-shirt | white | 2     |
+-----+-----+-----+-----+

SELECT s.* FROM person p INNER JOIN shirt s
ON s.owner = p.id
WHERE p.name LIKE 'Lilliana%'
AND s.color <> 'white';

+-----+-----+-----+-----+
| id | style | color | owner |
+-----+-----+-----+-----+
|  4 | dress | orange | 2     |
|  5 | polo  | red   | 2     |
|  6 | dress | blue  | 2     |
+-----+-----+-----+-----+

```

When used in this fashion, the [REFERENCES](#) clause is not displayed in the output of `SHOW CREATE TABLE` or `DESCRIBE`:

```

SHOW CREATE TABLE shirt\G
***** 1. row *****
Table: shirt
Create Table: CREATE TABLE `shirt` (
  `id` smallint(5) unsigned NOT NULL auto_increment,
  `style` enum('t-shirt','polo','dress') NOT NULL,
  `color` enum('red','blue','orange','white','black') NOT NULL,
  `owner` smallint(5) unsigned NOT NULL,
  PRIMARY KEY (`id`)
) ENGINE=MyISAM DEFAULT CHARSET=latin1

```

The use of [REFERENCES](#) in this way as a comment or “reminder” in a column definition works with [MyISAM](#) tables.

### 3.6.7. Searching on Two Keys



An `OR` [1172] using a single key is well optimized, as is the handling of `AND` [1171].

The one tricky case is that of searching on two different keys combined with `OR` [1172]:

```
SELECT field1_index, field2_index FROM test_table
WHERE field1_index = '1' OR field2_index = '1'
```

This case is optimized. See [Section 8.13.2, “Index Merge Optimization”](#).

You can also solve the problem efficiently by using a `UNION` that combines the output of two separate `SELECT` statements. See [Section 13.2.9.4, “UNION Syntax”](#).

Each `SELECT` searches only one key and can be optimized:

```
SELECT field1_index, field2_index
FROM test_table WHERE field1_index = '1'
UNION
SELECT field1_index, field2_index
FROM test_table WHERE field2_index = '1';
```

### 3.6.8. Calculating Visits Per Day

The following example shows how you can use the bit group functions to calculate the number of days per month a user has visited a Web page.

```
CREATE TABLE t1 (year YEAR(4), month INT(2) UNSIGNED ZEROFILL,
                 day INT(2) UNSIGNED ZEROFILL);
INSERT INTO t1 VALUES(2000,1,1),(2000,1,20),(2000,1,30),(2000,2,2),
                    (2000,2,23),(2000,2,23);
```

The example table contains year-month-day values representing visits by users to the page. To determine how many different days in each month these visits occur, use this query:

```
SELECT year,month,BIT_COUNT(BIT_OR(1<<day)) AS days FROM t1
GROUP BY year,month;
```

Which returns:

```
+-----+-----+-----+
| year | month | days |
+-----+-----+-----+
| 2000 |    01 |    3 |
| 2000 |    02 |    2 |
+-----+-----+-----+
```

The query calculates how many different days appear in the table for each year/month combination, with automatic removal of duplicate entries.

### 3.6.9. Using `AUTO_INCREMENT`

The `AUTO_INCREMENT` attribute can be used to generate a unique identity for new rows:

```
CREATE TABLE animals (
  id MEDIUMINT NOT NULL AUTO_INCREMENT,
  name CHAR(30) NOT NULL,
  PRIMARY KEY (id)
);

INSERT INTO animals (name) VALUES
  ('dog'),('cat'),('penguin'),
  ('lax'),('whale'),('ostrich');

SELECT * FROM animals;
```

Which returns:

```
+-----+-----+
| id | name |
+-----+-----+
| 1 | dog |
| 2 | cat |
| 3 | penguin |
| 4 | lax |
| 5 | whale |
| 6 | ostrich |
+-----+-----+
```

No value was specified for the `AUTO_INCREMENT` column, so MySQL assigned sequence numbers automatically. You can also explicitly assign `NULL` or 0 to the column to generate sequence numbers.

You can retrieve the most recent `AUTO_INCREMENT` value with the `LAST_INSERT_ID()` [1269] SQL function or the `mysql_insert_id()` C API function. These functions are connection-specific, so their return values are not affected by another connection which is also performing inserts.

Use the smallest integer data type for the `AUTO_INCREMENT` column that is large enough to hold the maximum sequence value you will need. When the column reaches the upper limit of the data type, the next attempt to generate a sequence number fails. Use the `UNSIGNED` attribute if possible to allow a greater range. For example, if you use `TINYINT`, the maximum permissible sequence number is 127. For `TINYINT UNSIGNED`, the maximum is 255. See Section 11.2.1, “Integer Types (Exact Value) - `INTEGER`, `INT`, `SMALLINT`, `TINYINT`, `MEDIUMINT`, `BIGINT`” for the ranges of all the integer types.



#### Note

For a multiple-row insert, `LAST_INSERT_ID()` [1269] and `mysql_insert_id()` actually return the `AUTO_INCREMENT` key from the *first* of the inserted rows. This enables multiple-row inserts to be reproduced correctly on other servers in a replication setup.

If the `AUTO_INCREMENT` column is part of multiple indexes, MySQL generates sequence values using the index that begins with the `AUTO_INCREMENT` column, if there is one. For example, if the `animals` table contained indexes `PRIMARY KEY (grp, id)` and `INDEX (id)`, MySQL would ignore the `PRIMARY KEY` for generating sequence values. As a result, the table would contain a single sequence, not a sequence per `grp` value.

To start with an `AUTO_INCREMENT` value other than 1, set that value with `CREATE TABLE` or `ALTER TABLE`, like this:

```
mysql> ALTER TABLE tbl AUTO_INCREMENT = 100;
```

## InnoDB Notes

For `InnoDB` tables, be careful if you modify the column containing the auto-increment value in the middle of a sequence of `INSERT` statements. For example, if you use an `UPDATE` statement to put a new, larger value in the auto-increment column, a subsequent `INSERT` could encounter a “Duplicate entry” error. The test whether an auto-increment value is already present occurs if you do a `DELETE` followed by more `INSERT` statements, or when you `COMMIT` the transaction, but not after an `UPDATE` statement.

## MyISAM Notes

For `MyISAM` tables, you can specify `AUTO_INCREMENT` on a secondary column in a multiple-column index. In this case, the generated value for the `AUTO_INCREMENT` column is calculated as `MAX(auto_increment_column) + 1 WHERE prefix=given-prefix` [1284]. This is useful when you want to put data into ordered groups.

```
CREATE TABLE animals (
  grp ENUM('fish','mammal','bird') NOT NULL,
  id MEDIUMINT NOT NULL AUTO_INCREMENT,
  name CHAR(30) NOT NULL,
```

```

PRIMARY KEY (grp,id)
) ENGINE=MyISAM;

INSERT INTO animals (grp,name) VALUES
  ('mammal','dog'),('mammal','cat'),
  ('bird','penguin'),('fish','lax'),('mammal','whale'),
  ('bird','ostrich');

SELECT * FROM animals ORDER BY grp,id;

```

Which returns:

```

+-----+-----+-----+
| grp   | id  | name  |
+-----+-----+-----+
| fish  | 1   | lax   |
| mammal| 1   | dog   |
| mammal| 2   | cat   |
| mammal| 3   | whale |
| bird  | 1   | penguin|
| bird  | 2   | ostrich|
+-----+-----+-----+

```

In this case (when the [AUTO\\_INCREMENT](#) column is part of a multiple-column index), [AUTO\\_INCREMENT](#) values are reused if you delete the row with the biggest [AUTO\\_INCREMENT](#) value in any group. This happens even for [MyISAM](#) tables, for which [AUTO\\_INCREMENT](#) values normally are not reused.

## Further Reading

More information about [AUTO\\_INCREMENT](#) is available here:

- How to assign the [AUTO\\_INCREMENT](#) attribute to a column: [Section 13.1.14](#), “[CREATE TABLE Syntax](#)”, and [Section 13.1.6](#), “[ALTER TABLE Syntax](#)”.
- How [AUTO\\_INCREMENT](#) behaves depending on the [NO\\_AUTO\\_VALUE\\_ON\\_ZERO \[604\]](#) SQL mode: [Section 5.1.7](#), “[Server SQL Modes](#)”.
- How to use the [LAST\\_INSERT\\_ID\(\) \[1269\]](#) function to find the row that contains the most recent [AUTO\\_INCREMENT](#) value: [Section 12.14](#), “[Information Functions](#)”.
- Setting the [AUTO\\_INCREMENT](#) value to be used: [Section 5.1.4](#), “[Server System Variables](#)”.
- [AUTO\\_INCREMENT](#) and replication: [Section 16.4.1.1](#), “[Replication and AUTO\\_INCREMENT](#)”.
- Server-system variables related to [AUTO\\_INCREMENT](#) ([auto\\_increment\\_increment \[1965\]](#) and [auto\\_increment\\_offset \[1968\]](#)) that can be used for replication: [Section 5.1.4](#), “[Server System Variables](#)”.

## 3.7. Using MySQL with Apache

There are programs that let you authenticate your users from a MySQL database and also let you write your log files into a MySQL table.

You can change the Apache logging format to be easily readable by MySQL by putting the following into the Apache configuration file:

```

LogFormat \
  "%h" "%Y%m%d%H%M%S"t,%>s,"%b",\ "%{Content-Type}o", \
  \"%U\", \"%{Referer}i\", \"%{User-Agent}i\"

```

To load a log file in that format into MySQL, you can use a statement something like this:

```

LOAD DATA INFILE '/local/access_log' INTO TABLE tbl_name
FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"' ESCAPED BY '\\

```

The named table should be created to have columns that correspond to those that the `LogFormat` line writes to the log file.



---

## Chapter 4. MySQL Programs

### Table of Contents

4.1. Overview of MySQL Programs .....	204
4.2. Using MySQL Programs .....	208
4.2.1. Invoking MySQL Programs .....	208
4.2.2. Connecting to the MySQL Server .....	209
4.2.3. Specifying Program Options .....	212
4.2.4. Setting Environment Variables .....	224
4.3. MySQL Server and Server-Startup Programs .....	225
4.3.1. <code>mysqld</code> — The MySQL Server .....	225
4.3.2. <code>mysqld_safe</code> — MySQL Server Startup Script .....	225
4.3.3. <code>mysql.server</code> — MySQL Server Startup Script .....	230
4.3.4. <code>mysqld_multi</code> — Manage Multiple MySQL Servers .....	231
4.4. MySQL Installation-Related Programs .....	235
4.4.1. <code>comp_err</code> — Compile MySQL Error Message File .....	235
4.4.2. <code>mysqlbug</code> — Generate Bug Report .....	236
4.4.3. <code>mysql_install_db</code> — Initialize MySQL Data Directory .....	236
4.4.4. <code>mysql_plugin</code> — Configure MySQL Server Plugins .....	238
4.4.5. <code>mysql_secure_installation</code> — Improve MySQL Installation Security .....	240
4.4.6. <code>mysql_tzinfo_to_sql</code> — Load the Time Zone Tables .....	241
4.4.7. <code>mysql_upgrade</code> — Check and Upgrade MySQL Tables .....	241
4.5. MySQL Client Programs .....	244
4.5.1. <code>mysql</code> — The MySQL Command-Line Tool .....	244
4.5.2. <code>mysqladmin</code> — Client for Administering a MySQL Server .....	266
4.5.3. <code>mysqlcheck</code> — A Table Maintenance Program .....	272
4.5.4. <code>mysqldump</code> — A Database Backup Program .....	279
4.5.5. <code>mysqlimport</code> — A Data Import Program .....	297
4.5.6. <code>mysqlshow</code> — Display Database, Table, and Column Information .....	302
4.5.7. <code>mysqlslap</code> — Load Emulation Client .....	306
4.6. MySQL Administrative and Utility Programs .....	313
4.6.1. <code>innochecksum</code> — Offline InnoDB File Checksum Utility .....	313
4.6.2. <code>myisam_ftdump</code> — Display Full-Text Index information .....	314
4.6.3. <code>myisamchk</code> — MyISAM Table-Maintenance Utility .....	315
4.6.4. <code>myisamlog</code> — Display MyISAM Log File Contents .....	331
4.6.5. <code>myisampack</code> — Generate Compressed, Read-Only MyISAM Tables .....	332
4.6.6. <code>mysql_config_editor</code> — MySQL Configuration Utility .....	337
4.6.7. <code>mysqlaccess</code> — Client for Checking Access Privileges .....	343
4.6.8. <code>mysqlbinlog</code> — Utility for Processing Binary Log Files .....	345
4.6.9. <code>mysqldumpslow</code> — Summarize Slow Query Log Files .....	364
4.6.10. <code>mysqlhotcopy</code> — A Database Backup Program .....	365
4.6.11. <code>mysql_convert_table_format</code> — Convert Tables to Use a Given Storage Engine .....	368
4.6.12. <code>mysql_find_rows</code> — Extract SQL Statements from Files .....	369
4.6.13. <code>mysql_fix_extensions</code> — Normalize Table File Name Extensions .....	370
4.6.14. <code>mysql_setpermission</code> — Interactively Set Permissions in Grant Tables .....	370
4.6.15. <code>mysql_waitpid</code> — Kill Process and Wait for Its Termination .....	371
4.6.16. <code>mysql_zap</code> — Kill Processes That Match a Pattern .....	371
4.7. MySQL Program Development Utilities .....	372
4.7.1. <code>msql2mysql</code> — Convert mSQL Programs for Use with MySQL .....	372
4.7.2. <code>mysql_config</code> — Display Options for Compiling Clients .....	373
4.7.3. <code>my_print_defaults</code> — Display Options from Option Files .....	374
4.7.4. <code>resolve_stack_dump</code> — Resolve Numeric Stack Trace Dump to Symbols .....	375
4.8. Miscellaneous Programs .....	375
4.8.1. <code>pererror</code> — Explain Error Codes .....	375

4.8.2. <code>replace</code> — A String-Replacement Utility .....	376
4.8.3. <code>resolveip</code> — Resolve Host name to IP Address or Vice Versa .....	377

This chapter provides a brief overview of the MySQL command-line programs provided by Oracle Corporation. It also discusses the general syntax for specifying options when you run these programs. Most programs have options that are specific to their own operation, but the option syntax is similar for all of them. Finally, the chapter provides more detailed descriptions of individual programs, including which options they recognize.

## 4.1. Overview of MySQL Programs

There are many different programs in a MySQL installation. This section provides a brief overview of them. Later sections provide a more detailed description of each one. Each program's description indicates its invocation syntax and the options that it supports.

Most MySQL distributions include all of these programs, except for those programs that are platform-specific. (For example, the server startup scripts are not used on Windows.) The exception is that RPM distributions are more specialized. There is one RPM for the server, another for client programs, and so forth. If you appear to be missing one or more programs, see [Chapter 2, \*Installing and Upgrading MySQL\*](#), for information on types of distributions and what they contain. It may be that you have a distribution that does not include all programs and you need to install an additional package.

Each MySQL program takes many different options. Most programs provide a `--help` option that you can use to get a description of the program's different options. For example, try `mysql --help`.

You can override default option values for MySQL programs by specifying options on the command line or in an option file. See [Section 4.2, "Using MySQL Programs"](#), for general information on invoking programs and specifying program options.

The MySQL server, `mysqld`, is the main program that does most of the work in a MySQL installation. The server is accompanied by several related scripts that assist you in starting and stopping the server:

- `mysqld`

The SQL daemon (that is, the MySQL server). To use client programs, `mysqld` must be running, because clients gain access to databases by connecting to the server. See [Section 4.3.1, "mysqld — The MySQL Server"](#).

- `mysqld_safe`

A server startup script. `mysqld_safe` attempts to start `mysqld`. See [Section 4.3.2, "mysqld\\_safe — MySQL Server Startup Script"](#).

- `mysql.server`

A server startup script. This script is used on systems that use System V-style run directories containing scripts that start system services for particular run levels. It invokes `mysqld_safe` to start the MySQL server. See [Section 4.3.3, "mysql.server — MySQL Server Startup Script"](#).

- `mysqld_multi`

A server startup script that can start or stop multiple servers installed on the system. See [Section 4.3.4, "mysqld\\_multi — Manage Multiple MySQL Servers"](#).

Several programs perform setup operations during MySQL installation or upgrading:

- `comp_err`

This program is used during the MySQL build/installation process. It compiles error message files from the error source files. See [Section 4.4.1, "comp\\_err — Compile MySQL Error Message File"](#).

- `mysql_install_db`



This script creates the MySQL database, initializes the grant tables with default privileges, and sets up the [InnoDB system tablespace](#). It is usually executed only once, when first installing MySQL on a system. See [Section 4.4.3, “mysql\\_install\\_db — Initialize MySQL Data Directory”](#), [Section 2.10.1, “Unix Postinstallation Procedures”](#), and [Section 4.4.3, “mysql\\_install\\_db — Initialize MySQL Data Directory”](#).

- `mysql_plugin`

This program configures MySQL server plugins. See [Section 4.4.4, “mysql\\_plugin — Configure MySQL Server Plugins”](#).

- `mysql_secure_installation`

This program enables you to improve the security of your MySQL installation. SQL. See [Section 4.4.5, “mysql\\_secure\\_installation — Improve MySQL Installation Security”](#).

- `mysql_tzinfo_to_sql`

This program loads the time zone tables in the `mysql` database using the contents of the host system `zoneinfo` database (the set of files describing time zones). SQL. See [Section 4.4.6, “mysql\\_tzinfo\\_to\\_sql — Load the Time Zone Tables”](#).

- `mysql_upgrade`

This program is used after a MySQL upgrade operation. It checks tables for incompatibilities and repairs them if necessary, and updates the grant tables with any changes that have been made in newer versions of MySQL. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

MySQL client programs that connect to the MySQL server:

- `mysql`

The command-line tool for interactively entering SQL statements or executing them from a file in batch mode. See [Section 4.5.1, “mysql — The MySQL Command-Line Tool”](#).

- `mysqladmin`

A client that performs administrative operations, such as creating or dropping databases, reloading the grant tables, flushing tables to disk, and reopening log files. `mysqladmin` can also be used to retrieve version, process, and status information from the server. See [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#).

- `mysqlcheck`

A table-maintenance client that checks, repairs, analyzes, and optimizes tables. See [Section 4.5.3, “mysqlcheck — A Table Maintenance Program”](#).

- `mysqldump`

A client that dumps a MySQL database into a file as SQL, text, or XML. See [Section 4.5.4, “mysqldump — A Database Backup Program”](#).

- `mysqlimport`

A client that imports text files into their respective tables using `LOAD DATA INFILE`. See [Section 4.5.5, “mysqlimport — A Data Import Program”](#).

- `mysqlshow`

A client that displays information about databases, tables, columns, and indexes. See [Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”](#).

- `mysqlslap`

A client that is designed to emulate client load for a MySQL server and report the timing of each stage. It works as if multiple clients are accessing the server. See [Section 4.5.7, “mysqlslap — Load Emulation Client”](#).

MySQL administrative and utility programs:

- `innochecksum`

An offline InnoDB offline file checksum utility. See [Section 4.6.1, “innochecksum — Offline InnoDB File Checksum Utility”](#).

- `myisam_ftdump`

A utility that displays information about full-text indexes in MyISAM tables. See [Section 4.6.2, “myisam\\_ftdump — Display Full-Text Index information”](#).

- `myisamchk`

A utility to describe, check, optimize, and repair MyISAM tables. See [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

- `myisamlog`

A utility that processes the contents of a MyISAM log file. See [Section 4.6.4, “myisamlog — Display MyISAM Log File Contents”](#).

- `myisampack`

A utility that compresses MyISAM tables to produce smaller read-only tables. See [Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).

- `mysql_config_editor`

A utility that enables you to store authentication credentials in a secure, encrypted login file named `.mylogin.cnf`. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#).

- `mysqlaccess`

A script that checks the access privileges for a host name, user name, and database combination. See [Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”](#).

- `mysqlbinlog`

A utility for reading statements from a binary log. The log of executed statements contained in the binary log files can be used to help recover from a crash. See [Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#).

- `mysqldumpslow`

A utility to read and summarize the contents of a slow query log. See [Section 4.6.9, “mysqldumpslow — Summarize Slow Query Log Files”](#).

- `mysqlhotcopy`

A utility that quickly makes backups of MyISAM tables while the server is running. See [Section 4.6.10, “mysqlhotcopy — A Database Backup Program”](#).

- `mysql_convert_table_format`

A utility that converts tables in a database to use a given storage engine. See [Section 4.6.11, “mysql\\_convert\\_table\\_format — Convert Tables to Use a Given Storage Engine”](#).

- `mysql_find_rows`

A utility that reads files containing SQL statements (such as update logs) and extracts statements that match a given regular expression. See [Section 4.6.12, “mysql\\_find\\_rows — Extract SQL Statements from Files”](#).

- `mysql_fix_extensions`

A utility that converts the extensions for `MyISAM` table files to lowercase. This can be useful after transferring the files from a system with case-insensitive file names to a system with case-sensitive file names. See [Section 4.6.13, “mysql\\_fix\\_extensions — Normalize Table File Name Extensions”](#).

- `mysql_setpermission`

A utility for interactively setting permissions in the MySQL grant tables. See [Section 4.6.14, “mysql\\_setpermission — Interactively Set Permissions in Grant Tables”](#).

- `mysql_waitpid`

A utility that kills the process with a given process ID. See [Section 4.6.15, “mysql\\_waitpid — Kill Process and Wait for Its Termination”](#).

- `mysql_zap`

A utility that kills processes that match a pattern. See [Section 4.6.16, “mysql\\_zap — Kill Processes That Match a Pattern”](#).

MySQL program-development utilities:

- `msql2mysql`

A shell script that converts `mSQL` programs to MySQL. It doesn't handle every case, but it gives a good start when converting. See [Section 4.7.1, “msql2mysql — Convert mSQL Programs for Use with MySQL”](#).

- `mysql_config`

A shell script that produces the option values needed when compiling MySQL programs. See [Section 4.7.2, “mysql\\_config — Display Options for Compiling Clients”](#).

- `my_print_defaults`

A utility that shows which options are present in option groups of option files. See [Section 4.7.3, “my\\_print\\_defaults — Display Options from Option Files”](#).

- `resolve_stack_dump`

A utility program that resolves a numeric stack trace dump to symbols. See [Section 4.7.4, “resolve\\_stack\\_dump — Resolve Numeric Stack Trace Dump to Symbols”](#).

Miscellaneous utilities:

- `perror`

A utility that displays the meaning of system or MySQL error codes. See [Section 4.8.1, “perror — Explain Error Codes”](#).

- `replace`

A utility program that performs string replacement in the input text. See [Section 4.8.2, “replace — A String-Replacement Utility”](#).

- `resolveip`

A utility program that resolves a host name to an IP address or vice versa. See [Section 4.8.3, “resolveip — Resolve Host name to IP Address or Vice Versa”](#).

Oracle Corporation also provides the [MySQL Workbench](#) GUI tool, which is used to administer MySQL servers and databases, to create, execute, and evaluate queries, and to migrate schemas and data from other relational database management systems for use with MySQL. Additional GUI tools include [MySQL Notifier for Microsoft Windows](#) and [MySQL for Excel](#).

MySQL client programs that communicate with the server using the MySQL client/server library use the following environment variables.

Environment Variable	Meaning
<code>MYSQL_UNIX_PORT</code>	The default Unix socket file; used for connections to <code>localhost</code>
<code>MYSQL_TCP_PORT</code>	The default port number; used for TCP/IP connections
<code>MYSQL_PWD</code>	The default password
<code>MYSQL_DEBUG</code>	Debug trace options when debugging
<code>TMPDIR</code>	The directory where temporary tables and files are created

For a full list of environment variables used by MySQL programs, see [Section 2.12, “Environment Variables”](#).

Use of `MYSQL_PWD` is insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#).

## 4.2. Using MySQL Programs

### 4.2.1. Invoking MySQL Programs

To invoke a MySQL program from the command line (that is, from your shell or command prompt), enter the program name followed by any options or other arguments needed to instruct the program what you want it to do. The following commands show some sample program invocations. “`shell>`” represents the prompt for your command interpreter; it is not part of what you type. The particular prompt you see depends on your command interpreter. Typical prompts are `$` for `sh`, `ksh`, or `bash`, `%` for `csch` or `tcsh`, and `C:\>` for the Windows `command.com` or `cmd.exe` command interpreters.

```
shell> mysql --user=root test
shell> mysqladmin extended-status variables
shell> mysqlshow --help
shell> mysqldump -u root personnel
```

Arguments that begin with a single or double dash (“-”, “--”) specify program options. Options typically indicate the type of connection a program should make to the server or affect its operational mode. Option syntax is described in [Section 4.2.3, “Specifying Program Options”](#).

Nonoption arguments (arguments with no leading dash) provide additional information to the program. For example, the `mysql` program interprets the first nonoption argument as a database name, so the command `mysql --user=root test` indicates that you want to use the `test` database.

Later sections that describe individual programs indicate which options a program supports and describe the meaning of any additional nonoption arguments.

Some options are common to a number of programs. The most frequently used of these are the `--host [211]` (or `-h`), `--user [212]` (or `-u`), and `--password [211]` (or `-p`) options that specify connection parameters. They indicate the host where the MySQL server is running, and the user name and password of your MySQL account. All MySQL client programs understand these options;

they enable you to specify which server to connect to and the account to use on that server. Other connection options are `--port [211]` (or `-P`) to specify a TCP/IP port number and `--socket [211]` (or `-S`) to specify a Unix socket file on Unix (or named pipe name on Windows). For more information on options that specify connection options, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

You may find it necessary to invoke MySQL programs using the path name to the `bin` directory in which they are installed. This is likely to be the case if you get a “program not found” error whenever you attempt to run a MySQL program from any directory other than the `bin` directory. To make it more convenient to use MySQL, you can add the path name of the `bin` directory to your `PATH` environment variable setting. That enables you to run a program by typing only its name, not its entire path name. For example, if `mysql` is installed in `/usr/local/mysql/bin`, you can run the program by invoking it as `mysql`, and it is not necessary to invoke it as `/usr/local/mysql/bin/mysql`.

Consult the documentation for your command interpreter for instructions on setting your `PATH` variable. The syntax for setting environment variables is interpreter-specific. (Some information is given in [Section 4.2.4, “Setting Environment Variables”](#).) After modifying your `PATH` setting, open a new console window on Windows or log in again on Unix so that the setting goes into effect.

## 4.2.2. Connecting to the MySQL Server

For a client program to be able to connect to the MySQL server, it must use the proper connection parameters, such as the name of the host where the server is running and the user name and password of your MySQL account. Each connection parameter has a default value, but you can override them as necessary using program options specified either on the command line or in an option file.

The examples here use the `mysql` client program, but the principles apply to other clients such as `mysqldump`, `mysqladmin`, or `mysqlshow`.

This command invokes `mysql` without specifying any connection parameters explicitly:

```
shell> mysql
```

Because there are no parameter options, the default values apply:

- The default host name is `localhost`. On Unix, this has a special meaning, as described later.
- The default user name is `ODBC` on Windows or your Unix login name on Unix.
- No password is sent if neither `-p` nor `--password [211]` is given.
- For `mysql`, the first nonoption argument is taken as the name of the default database. If there is no such option, `mysql` does not select a default database.

To specify the host name and user name explicitly, as well as a password, supply appropriate options on the command line:

```
shell> mysql --host=localhost --user=myname --password=mypass mydb
shell> mysql -h localhost -u myname -pmypass mydb
```

For password options, the password value is optional:

- If you use a `-p` or `--password [211]` option and specify the password value, there must be *no space* between `-p` or `--password= [211]` and the password following it.
- If you use a `-p` or `--password [211]` option but do not specify the password value, the client program prompts you to enter the password. The password is not displayed as you enter it. This is more secure than giving the password on the command line. Other users on your system may be able to see a password specified on the command line by executing a command such as `ps auxw`. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#).

As just mentioned, including the password value on the command line can be a security risk. To avoid this problem, specify the `--password` or `-p` option without any following password value:

```
shell> mysql --host=localhost --user=myname --password mydb
shell> mysql -h localhost -u myname -p mydb
```

When the password option has no password value, the client program prints a prompt and waits for you to enter the password. (In these examples, `mydb` is *not* interpreted as a password because it is separated from the preceding password option by a space.)

On some systems, the library routine that MySQL uses to prompt for a password automatically limits the password to eight characters. That is a problem with the system library, not with MySQL. Internally, MySQL does not have any limit for the length of the password. To work around the problem, change your MySQL password to a value that is eight or fewer characters long, or put your password in an option file.

On Unix, MySQL programs treat the host name `localhost` specially, in a way that is likely different from what you expect compared to other network-based programs. For connections to `localhost`, MySQL programs attempt to connect to the local server by using a Unix socket file. This occurs even if a `--port [211]` or `-P` option is given to specify a port number. To ensure that the client makes a TCP/IP connection to the local server, use `--host [211]` or `-h` to specify a host name value of `127.0.0.1`, or the IP address or name of the local server. You can also specify the connection protocol explicitly, even for `localhost`, by using the `--protocol=TCP [211]` option. For example:

```
shell> mysql --host=127.0.0.1
shell> mysql --protocol=TCP
```

The `--protocol [211]` option enables you to establish a particular type of connection even when the other options would normally default to some other protocol.

If the server is configured to accept IPv6 connections, client can connect over IPv6 using `--host>:::1 [211]`. See [Section 5.1.9, "IPv6 Support"](#).

On Windows, you can force a MySQL client to use a named-pipe connection by specifying the `--pipe [211]` or `--protocol=PIPE [211]` option, or by specifying `.` (period) as the host name. If named-pipe connections are not enabled, an error occurs. Use the `--socket [211]` option to specify the name of the pipe if you do not want to use the default pipe name.

Connections to remote servers always use TCP/IP. This command connects to the server running on `remote.example.com` using the default port number (3306):

```
shell> mysql --host=remote.example.com
```

To specify a port number explicitly, use the `--port [211]` or `-P` option:

```
shell> mysql --host=remote.example.com --port=13306
```

You can specify a port number for connections to a local server, too. However, as indicated previously, connections to `localhost` on Unix will use a socket file by default. You will need to force a TCP/IP connection as already described or any option that specifies a port number will be ignored.

For this command, the program uses a socket file on Unix and the `--port [211]` option is ignored:

```
shell> mysql --port=13306 --host=localhost
```

To cause the port number to be used, invoke the program in either of these ways:

```
shell> mysql --port=13306 --host=127.0.0.1
shell> mysql --port=13306 --protocol=TCP
```

The following list summarizes the options that can be used to control how client programs connect to the server:

- `--host=host_name [211], -h host_name`

The host where the server is running. The default value is `localhost`.

- `--password[=pass_val] [211], -p[pass_val]`

The password of the MySQL account. As described earlier, the password value is optional, but if given, there must be *no space* between `-p` or `--password=` [211] and the password following it. The default is to send no password.

- `--pipe [211], -W`

On Windows, connect to the server using a named pipe. The server must be started with the `--enable-named-pipe` [421] option to enable named-pipe connections.

- `--port=port_num [211], -P port_num`

The port number to use for the connection, for connections made using TCP/IP. The default port number is 3306.

- `--protocol={TCP|SOCKET|PIPE|MEMORY} [211]`

This option explicitly specifies a protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For example, connections on Unix to `localhost` are made using a Unix socket file by default:

```
shell> mysql --host=localhost
```

To force a TCP/IP connection to be used instead, specify a `--protocol` [211] option:

```
shell> mysql --host=localhost --protocol=TCP
```

The following table shows the permissible `--protocol` [211] option values and indicates the platforms on which each value may be used. The values are not case sensitive.

<code>--protocol</code> [211] Value	Connection Protocol	Permissible Operating Systems
TCP	TCP/IP connection to local or remote server	All
SOCKET	Unix socket file connection to local server	Unix only
PIPE	Named-pipe connection to local or remote server	Windows only
MEMORY	Shared-memory connection to local server	Windows only

- `--shared-memory-base-name=name [211]`

On Windows, the shared-memory name to use, for connections made using shared memory to a local server. The default value is `MYSQL`. The shared-memory name is case sensitive.

The server must be started with the `--shared-memory` [439] option to enable shared-memory connections.

- `--socket=file_name [211], -S file_name`

On Unix, the name of the Unix socket file to use, for connections made using a named pipe to a local server. The default Unix socket file name is `/tmp/mysql.sock`.

On Windows, the name of the named pipe to use, for connections to a local server. The default Windows pipe name is `MySQL`. The pipe name is not case sensitive.

The server must be started with the `--enable-named-pipe` [421] option to enable named-pipe connections.



- `--ssl*`

Options that begin with `--ssl` [810] are used for establishing a secure connection to the server using SSL, if the server is configured with SSL support. For details, see [Section 6.3.9.4, “SSL Command Options”](#).

- `--user=user_name` [212], `-u user_name`

The user name of the MySQL account you want to use. The default user name is `ODBC` on Windows or your Unix login name on Unix.

It is possible to specify different default values to be used when you make a connection so that you need not enter them on the command line each time you invoke a client program. This can be done in a couple of ways:

- You can specify connection parameters in the `[client]` section of an option file. The relevant section of the file might look like this:

```
[client]
host=host_name
user=user_name
password=your_pass
```

[Section 4.2.3.3, “Using Option Files”](#), discusses option files further.

- You can specify some connection parameters using environment variables. The host can be specified for `mysql` using `MYSQL_HOST`. The MySQL user name can be specified using `USER` (this is for Windows only). The password can be specified using `MYSQL_PWD`, although this is insecure; see [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). For a list of variables, see [Section 2.12, “Environment Variables”](#).

### 4.2.3. Specifying Program Options

There are several ways to specify options for MySQL programs:

- List the options on the command line following the program name. This is common for options that apply to a specific invocation of the program.
- List the options in an option file that the program reads when it starts. This is common for options that you want the program to use each time it runs.
- List the options in environment variables (see [Section 4.2.4, “Setting Environment Variables”](#)). This method is useful for options that you want to apply each time the program runs. In practice, option files are used more commonly for this purpose, but [Section 5.6.3, “Running Multiple MySQL Instances on Unix”](#), discusses one situation in which environment variables can be very helpful. It describes a handy technique that uses such variables to specify the TCP/IP port number and Unix socket file for the server and for client programs.

Options are processed in order, so if an option is specified multiple times, the last occurrence takes precedence. The following command causes `mysql` to connect to the server running on `localhost`:

```
shell> mysql -h example.com -h localhost
```

If conflicting or related options are given, later options take precedence over earlier options. The following command runs `mysql` in “no column names” mode:

```
shell> mysql --column-names --skip-column-names
```

MySQL programs determine which options are given first by examining environment variables, then by reading option files, and then by checking the command line. This means that environment variables have the lowest precedence and command-line options the highest.

You can take advantage of the way that MySQL programs process options by specifying default option values for a program in an option file. That enables you to avoid typing them each time you run the program while enabling you to override the defaults if necessary by using command-line options.

An option can be specified by writing it in full or as any unambiguous prefix. For example, the `--compress` [284] option can be given to `mysqldump` as `--compr`, but not as `--comp` because the latter is ambiguous:

```
shell> mysqldump --comp
mysqldump: ambiguous option '--comp' (compatible, compress)
```

Be aware that the use of option prefixes can cause problems in the event that new options are implemented for a program. A prefix that is unambiguous now might become ambiguous in the future.

### 4.2.3.1. Using Options on the Command Line

Program options specified on the command line follow these rules:

- Options are given after the command name.
- An option argument begins with one dash or two dashes, depending on whether it is a short form or long form of the option name. Many options have both short and long forms. For example, `-?` and `--help` are the short and long forms of the option that instructs a MySQL program to display its help message.
- Option names are case sensitive. `-v` and `-V` are both legal and have different meanings. (They are the corresponding short forms of the `--verbose` and `--version` options.)
- Some options take a value following the option name. For example, `-h localhost` or `--host=localhost` [211] indicate the MySQL server host to a client program. The option value tells the program the name of the host where the MySQL server is running.
- For a long option that takes a value, separate the option name and the value by an “=” sign. For a short option that takes a value, the option value can immediately follow the option letter, or there can be a space between: `-hlocalhost` and `-h localhost` are equivalent. An exception to this rule is the option for specifying your MySQL password. This option can be given in long form as `--password=pass_val` [211] or as `--password` [211]. In the latter case (with no password value given), the program prompts you for the password. The password option also may be given in short form as `-ppass_val` or as `-p`. However, for the short form, if the password value is given, it must follow the option letter with *no intervening space*. The reason for this is that if a space follows the option letter, the program has no way to tell whether a following argument is supposed to be the password value or some other kind of argument. Consequently, the following two commands have two completely different meanings:

```
shell> mysql -ptest
shell> mysql -p test
```

The first command instructs `mysql` to use a password value of `test`, but specifies no default database. The second instructs `mysql` to prompt for the password value and to use `test` as the default database.

- Within option names, dash (“-”) and underscore (“\_”) may be used interchangeably. For example, `--skip-grant-tables` [439] and `--skip_grant_tables` [439] are equivalent. (However, the leading dashes cannot be given as underscores.)
- For options that take a numeric value, the value can be given with a suffix of `K`, `M`, or `G` (either uppercase or lowercase) to indicate a multiplier of 1024, 1024<sup>2</sup> or 1024<sup>3</sup>. For example, the following command tells `mysqladmin` to ping the server 1024 times, sleeping 10 seconds between each ping:

```
mysql> mysqladmin --count=1K --sleep=10 ping
```

Option values that contain spaces must be quoted when given on the command line. For example, the `--execute` [250] (or `-e`) option can be used with `mysql` to pass SQL statements to the server.

When this option is used, `mysql` executes the statements in the option value and exits. The statements must be enclosed by quotation marks. For example, you can use the following command to obtain a list of user accounts:

```
mysql> mysql -u root -p --execute="SELECT User, Host FROM mysql.user"
Enter password: *****
+-----+-----+
| User | Host |
+-----+-----+
| root | gigan |
| root | gigan |
| jon  | localhost |
| root | localhost |
+-----+-----+
shell>
```

Note that the long form (`--execute [250]`) is followed by an equals sign (=).

If you wish to use quoted values within a statement, you will either need to escape the inner quotation marks, or use a different type of quotation marks within the statement from those used to quote the statement itself. The capabilities of your command processor dictate your choices for whether you can use single or double quotation marks and the syntax for escaping quote characters. For example, if your command processor supports quoting with single or double quotation marks, you can use double quotation marks around the statement, and single quotation marks for any quoted values within the statement.

Multiple SQL statements may be passed in the option value on the command line, separated by semicolons:

```
shell> mysql -u root -p -e "SELECT VERSION();SELECT NOW()"
Enter password: *****
+-----+
| VERSION() |
+-----+
| 5.1.5-alpha-log |
+-----+
+-----+
| NOW() |
+-----+
| 2006-01-05 21:19:04 |
+-----+
```

### 4.2.3.2. Program Option Modifiers

Some options are “boolean” and control behavior that can be turned on or off. For example, the `mysql` client supports a `--column-names [249]` option that determines whether or not to display a row of column names at the beginning of query results. By default, this option is enabled. However, you may want to disable it in some instances, such as when sending the output of `mysql` into another program that expects to see only data and not an initial header line.

To disable column names, you can specify the option using any of these forms:

```
--disable-column-names
--skip-column-names
--column-names=0
```

The `--disable` and `--skip` prefixes and the `=0` suffix all have the same effect: They turn the option off.

The “enabled” form of the option may be specified in any of these ways:

```
--column-names
--enable-column-names
--column-names=1
```

As of MySQL 5.6.2, the values `ON`, `TRUE`, `OFF`, and `FALSE` are also recognized for boolean options (not case sensitive).

If an option is prefixed by `--loose`, a program does not exit with an error if it does not recognize the option, but instead issues only a warning:

```
shell> mysql --loose-no-such-option
mysql: WARNING: unknown option '--no-such-option'
```

The `--loose` prefix can be useful when you run programs from multiple installations of MySQL on the same machine and list options in an option file. An option that may not be recognized by all versions of a program can be given using the `--loose` prefix (or `loose` in an option file). Versions of the program that recognize the option process it normally, and versions that do not recognize it issue a warning and ignore it.

`mysqld` enables a limit to be placed on how large client programs can set dynamic system variables. To do this, use a `--maximum` prefix with the variable name. For example, `--maximum-query_cache_size=4M` prevents any client from making the query cache size larger than 4MB.

### 4.2.3.3. Using Option Files

Most MySQL programs can read startup options from option files (also sometimes called configuration files). Option files provide a convenient way to specify commonly used options so that they need not be entered on the command line each time you run a program. For the MySQL server, MySQL provides a number of [preconfigured option files](#).

To determine whether a program reads option files, invoke it with the `--help` option. (For `mysqld`, use `--verbose` [447] and `--help` [413].) If the program reads option files, the help message indicates which files it looks for and which option groups it recognizes.

The `.mylogin.cnf` file that contains login path options is created by the `mysql_config_editor` utility. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#). A “login path” is an option group that permits only a limited set of options: `host`, `user`, and `password`. Client programs specify which login path to read from `.mylogin.cnf` using the `--login-path` [219] option.

On Windows, MySQL programs read startup options from the following files, in the specified order (top items are used first).

File Name	Purpose
<code>%WINDIR%\my.ini</code> , <code>%WINDIR%\my.cnf</code>	Global options
<code>C:\my.ini</code> , <code>C:\my.cnf</code>	Global options
<code>INSTALLDIR\my.ini</code> , <code>INSTALLDIR\my.cnf</code>	Global options
<code>defaults-extra-file</code>	The file specified with <code>--defaults-extra-file=path</code> [219], if any
<code>%APPDATA%\MySQL</code> <code>\.mylogin.cnf</code>	Login path options

`%WINDIR%` represents the location of your Windows directory. This is commonly `C:\WINDOWS`. You can determine its exact location from the value of the `WINDIR` environment variable using the following command:

```
C:\> echo %WINDIR%
```

`INSTALLDIR` represents the MySQL installation directory. This is typically `C:\PROGRAMDIR\MySQL\MySQL 5.6 Server` where `PROGRAMDIR` represents the programs directory (usually `Program Files` on English-language versions of Windows), when MySQL 5.6 has been installed using the

installation and configuration wizards. See [Section 2.3.3, “Installing MySQL on Microsoft Windows Using MySQL Installer”](#).

`%APPDATA%` represents the value of the Windows application data directory. You can determine its exact location from the value of the `APPDATA` environment variable using the following command:

```
C:\> echo %APPDATA%
```

On Unix, Linux and Mac OS X, MySQL programs read startup options from the following files, in the specified order (top items are used first).

File Name	Purpose
<code>/etc/my.cnf</code>	Global options
<code>/etc/mysql/my.cnf</code>	Global options
<code>SYSCONFDIR/my.cnf</code>	Global options
<code>\$MYSQL_HOME/my.cnf</code>	Server-specific options
<code>defaults-extra-file</code>	The file specified with <code>--defaults-extra-file=path</code> [219], if any
<code>~/.my.cnf</code>	User-specific options
<code>~/.mylogin.cnf</code>	Login path options

`~` represents the current user's home directory (the value of `$HOME`).

`SYSCONFDIR` represents the directory specified with the `SYSCONFDIR` [133] option to `CMake` when MySQL was built. By default, this is the `etc` directory located under the compiled-in installation directory.

`MYSQL_HOME` is an environment variable containing the path to the directory in which the server-specific `my.cnf` file resides. If `MYSQL_HOME` is not set and you start the server using the `mysqld_safe` program, `mysqld_safe` attempts to set `MYSQL_HOME` as follows:

- Let `BASEDIR` and `DATADIR` represent the path names of the MySQL base directory and data directory, respectively.
- If there is a `my.cnf` file in `DATADIR` but not in `BASEDIR`, `mysqld_safe` sets `MYSQL_HOME` to `DATADIR`.
- Otherwise, if `MYSQL_HOME` is not set and there is no `my.cnf` file in `DATADIR`, `mysqld_safe` sets `MYSQL_HOME` to `BASEDIR`.

In MySQL 5.6, use of `DATADIR` as the location for `my.cnf` is deprecated.

Typically, `DATADIR` is `/usr/local/mysql/data` for a binary installation or `/usr/local/var` for a source installation. Note that this is the data directory location that was specified at configuration time, not the one specified with the `--datadir` [418] option when `mysqld` starts. Use of `--datadir` [418] at runtime has no effect on where the server looks for option files, because it looks for them before processing any options.

MySQL looks for option files in the order just described and reads any that exist. If an option file that you want to use does not exist, create it with a plain text editor.

If multiple instances of a given option are found, the last instance takes precedence. There is one exception: For `mysqld`, the *first* instance of the `--user` [446] option is used as a security precaution, to prevent a user specified in an option file from being overridden on the command line.



**Note**

On Unix platforms, MySQL ignores configuration files that are world-writable. This is intentional as a security measure.

Any long option that may be given on the command line when running a MySQL program can be given in an option file as well. To get the list of available options for a program, run it with the `--help` option.

The syntax for specifying options in an option file is similar to command-line syntax (see [Section 4.2.3.1, “Using Options on the Command Line”](#)). However, in an option file, you omit the leading two dashes from the option name and you specify only one option per line. For example, `--quick` and `--host=localhost` on the command line should be specified as `quick` and `host=localhost` on separate lines in an option file. To specify an option of the form `--loose-opt_name` in an option file, write it as `loose-opt_name`.

Empty lines in option files are ignored. Nonempty lines can take any of the following forms:

- `#comment, ;comment`

Comment lines start with “#” or “;”. A “#” comment can start in the middle of a line as well.

- `[group]`

`group` is the name of the program or group for which you want to set options. After a group line, any option-setting lines apply to the named group until the end of the option file or another group line is given. Option group names are not case sensitive.

- `opt_name`

This is equivalent to `--opt_name` on the command line.

- `opt_name=value`

This is equivalent to `--opt_name=value` on the command line. In an option file, you can have spaces around the “=” character, something that is not true on the command line. You can optionally enclose the value within single quotation marks or double quotation marks, which is useful if the value contains a “#” comment character.

Leading and trailing spaces are automatically deleted from option names and values.

You can use the escape sequences “\b”, “\t”, “\n”, “\r”, “\\”, and “\s” in option values to represent the backspace, tab, newline, carriage return, backslash, and space characters. The escaping rules in option files are:

- If a backslash is followed by a valid escape sequence character, the sequence is converted to the character represented by the sequence. For example, “\s” is converted to a space.
- If a backslash is not followed by a valid escape sequence character, it remains unchanged. For example, “\S” is retained as is.

The preceding rules mean that a literal backslash can be given as “\\”, or as “\” if it is not followed by a valid escape sequence character.

The rules for escape sequences in option files differ slightly from the rules for escape sequences in string literals in SQL statements. In the latter context, if “x” is not a value escape sequence character, “\x” becomes “x” rather than “\x”. See [Section 9.1.1, “String Literals”](#).

The escaping rules for option file values are especially pertinent for Windows path names, which use “\” as a path name separator. A separator in a Windows path name must be written as “\\” if it is followed by an escape sequence character. It can be written as “\” or “\” if it is not. Alternatively, “/” may be used in Windows path names and will be treated as “\”. Suppose that you want to specify a base directory of `C:\Program Files\MySQL\MySQL Server 5.6` in an option file. This can be done several ways. Some examples:

```
basedir="C:\Program Files\MySQL\MySQL Server 5.6"
basedir="C:\\Program Files\\MySQL\\MySQL Server 5.6"
basedir="C:/Program Files/MySQL/MySQL Server 5.6"
basedir=C:\\Program\sFiles\\MySQL\\MySQL\sServer\s5.6
```

If an option group name is the same as a program name, options in the group apply specifically to that program. For example, the `[mysqld]` and `[mysql]` groups apply to the `mysqld` server and the `mysql` client program, respectively.

The `[client]` option group is read by all client programs (but *not* by `mysqld`). This enables you to specify options that apply to all clients. For example, `[client]` is the perfect group to use to specify the password that you use to connect to the server. (But make sure that the option file is readable and writable only by yourself, so that other people cannot find out your password.) Be sure not to put an option in the `[client]` group unless it is recognized by *all* client programs that you use. Programs that do not understand the option quit after displaying an error message if you try to run them.

Here is a typical global option file:

```
[client]
port=3306
socket=/tmp/mysql.sock

[mysqld]
port=3306
socket=/tmp/mysql.sock
key_buffer_size=16M
max_allowed_packet=8M

[mysqldump]
quick
```

The preceding option file uses `var_name=value` syntax for the lines that set the `key_buffer_size` [492] and `max_allowed_packet` [502] variables.

Here is a typical user option file:

```
[client]
# The following password will be sent to all standard MySQL clients
password="my_password"

[mysql]
no-auto-rehash
connect_timeout=2

[mysqlhotcopy]
interactive_timeout
```

If you want to create option groups that should be read by `mysqld` servers from a specific MySQL release series only, you can do this by using groups with names of `[mysqld-5.5]`, `[mysqld-5.6]`, and so forth. The following group indicates that the `--new` option should be used only by MySQL servers with 5.6.x version numbers:

```
[mysqld-5.6]
new
```

It is possible to use `!include` directives in option files to include other option files and `!includedir` to search specific directories for option files. For example, to include the `/home/mydir/myopt.cnf` file, use the following directive:

```
!include /home/mydir/myopt.cnf
```

To search the `/home/mydir` directory and read option files found there, use this directive:

```
!includedir /home/mydir
```

There is no guarantee about the order in which the option files in the directory will be read.



**Note**

Currently, any files to be found and included using the `!includedir` directive on Unix operating systems *must* have file names ending in `.cnf`. On Windows, this directive checks for files with the `.ini` or `.cnf` extension.



Write the contents of an included option file like any other option file. That is, it should contain groups of options, each preceded by a `[group]` line that indicates the program to which the options apply.

While an included file is being processed, only those options in groups that the current program is looking for are used. Other groups are ignored. Suppose that a `my.cnf` file contains this line:

```
!include /home/mydir/myopt.cnf
```

And suppose that `/home/mydir/myopt.cnf` looks like this:

```
[mysqladmin]
force

[mysqld]
key_buffer_size=16M
```

If `my.cnf` is processed by `mysqld`, only the `[mysqld]` group in `/home/mydir/myopt.cnf` is used. If the file is processed by `mysqladmin`, only the `[mysqladmin]` group is used. If the file is processed by any other program, no options in `/home/mydir/myopt.cnf` are used.

The `!includedir` directive is processed similarly except that all option files in the named directory are read.

#### 4.2.3.4. Command-Line Options that Affect Option-File Handling

Most MySQL programs that support option files handle the following options. They affect option-file handling, so they must be given on the command line and not in an option file. To work properly, each of these options must be given before other options, with these exceptions:

- `--print-defaults` [220] may be used immediately after `--defaults-file` [219] or `--defaults-extra-file` [219].
- On Windows, if the server is started with the `--defaults-file` [219] and `--install` [424] options, `--install` [424] must be first. See Section 2.3.4.8, “Starting MySQL as a Windows Service”.

When specifying file names, you should avoid the use of the “~” shell metacharacter because it might not be interpreted as you expect.

- `--defaults-extra-file=file_name` [219]

Read this option file after the global option file but (on Unix) before the user option file. If the file does not exist or is otherwise inaccessible, the program exits with an error. `file_name` is interpreted relative to the current directory if given as a relative path name rather than a full path name.

- `--defaults-file=file_name` [219]

Use only the given option file. If the file does not exist or is otherwise inaccessible, the program exits with an error. `file_name` is interpreted relative to the current directory if given as a relative path name rather than a full path name.

- `--defaults-group-suffix=str` [219]

If this option is given, the program reads not only its usual option groups, but also groups with the usual names and a suffix of `str`. For example, the `mysql` client normally reads the `[client]` and `[mysql]` groups. If the `--defaults-group-suffix=_other` [219] option is given, `mysql` also reads the `[client_other]` and `[mysql_other]` groups.

- `--login-path=name` [219]

Read options from the named login path in the `.mylogin.cnf` login file. A “login path” is an option group that permits only a limited set of options: `host`, `user`, and `password`. Think of a

login path as a set of values that indicate the server host and the credentials for authenticating with the server. To create the login file, use the `mysql_config_editor` utility. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#). This option was added in MySQL 5.6.6.

- `--no-defaults` [\[220\]](#)

Do not read any option files. If a program does not start because it is reading unknown options from an option file, `--no-defaults` [\[220\]](#) can be used to prevent the program from reading them.

The exception is that the `.mylogin.cnf` file, if it exists, is read in all cases. This permits passwords to be specified in a safer way than on the command line even when `--no-defaults` [\[220\]](#) is used. (`.mylogin.cnf` is created by the `mysql_config_editor` utility. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#).)

- `--print-defaults` [\[220\]](#)

Print the program name and all options that it gets from option files.

### 4.2.3.5. Using Options to Set Program Variables

Many MySQL programs have internal variables that can be set at runtime using the `SET` statement. See [Section 13.7.4, “SET Syntax”](#), and [Section 5.1.5, “Using System Variables”](#).

Most of these program variables also can be set at server startup by using the same syntax that applies to specifying program options. For example, `mysql` has a `max_allowed_packet` variable that controls the maximum size of its communication buffer. To set the `max_allowed_packet` variable for `mysql` to a value of 16MB, use either of the following commands:

```
shell> mysql --max_allowed_packet=16777216
shell> mysql --max_allowed_packet=16M
```

The first command specifies the value in bytes. The second specifies the value in megabytes. For variables that take a numeric value, the value can be given with a suffix of `K`, `M`, or `G` (either uppercase or lowercase) to indicate a multiplier of  $1024$ ,  $1024^2$  or  $1024^3$ . (For example, when used to set `max_allowed_packet`, the suffixes indicate units of kilobytes, megabytes, or gigabytes.)

In an option file, variable settings are given without the leading dashes:

```
[mysql]
max_allowed_packet=16777216
```

Or:

```
[mysql]
max_allowed_packet=16M
```

If you like, underscores in a variable name can be specified as dashes. The following option groups are equivalent. Both set the size of the server's key buffer to 512MB:

```
[mysqld]
key_buffer_size=512M

[mysqld]
key-buffer-size=512M
```

A variable can be specified by writing it in full or as any unambiguous prefix. For example, the `max_allowed_packet` variable can be set for `mysql` as `--max_a`, but not as `--max` because the latter is ambiguous:

```
shell> mysql --max=1000000
mysql: ambiguous option '--max=1000000' (max_allowed_packet, max_join_size)
```

Be aware that the use of variable prefixes can cause problems in the event that new variables are implemented for a program. A prefix that is unambiguous now might become ambiguous in the future.

Suffixes for specifying a value multiplier can be used when setting a variable at server startup, but not to set the value with `SET` at runtime. On the other hand, with `SET` you can assign a variable's value using an expression, which is not true when you set a variable at server startup. For example, the first of the following lines is legal at server startup, but the second is not:

```
shell> mysql --max_allowed_packet=16M
shell> mysql --max_allowed_packet=16*1024*1024
```

Conversely, the second of the following lines is legal at runtime, but the first is not:

```
mysql> SET GLOBAL max_allowed_packet=16M;
mysql> SET GLOBAL max_allowed_packet=16*1024*1024;
```

#### 4.2.3.6. Option Defaults, Options Expecting Values, and the = Sign

By convention, long forms of options that assign a value are written with an equals (=) sign, like this:

```
shell> mysql --host=tonfisk --user=jon
```

For options that require a value (that is, not having a default value), the equals sign is not required, and so the following is also valid:

```
shell> mysql --host tonfisk --user jon
```

In both cases, the `mysql` client attempts to connect to a MySQL server running on the host named “tonfisk” using an account with the user name “jon”.

Due to this behavior, problems can occasionally arise when no value is provided for an option that expects one. Consider the following example, where a user connects to a MySQL server running on host `tonfisk` as user `jon`:

```
shell> mysql --host 85.224.35.45 --user jon
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 3
Server version: 5.6.12 Source distribution

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SELECT CURRENT_USER();
+-----+
| CURRENT_USER() |
+-----+
| jon@%          |
+-----+
1 row in set (0.00 sec)
```

Omitting the required value for one of these option yields an error, such as the one shown here:

```
shell> mysql --host 85.224.35.45 --user
mysql: option '--user' requires an argument
```

In this case, `mysql` was unable to find a value following the `--user` [212] option because nothing came after it on the command line. However, if you omit the value for an option that is *not* the last option to be used, you obtain a different error that you may not be expecting:

```
shell> mysql --host --user jon
ERROR 2005 (HY000): Unknown MySQL server host '--user' (1)
```

Because `mysql` assumes that any string following `--host` [211] on the command line is a host name, `--host` [211] `--user` [212] is interpreted as `--host=--user` [211], and the client attempts to connect to a MySQL server running on a host named “--user”.

Options having default values always require an equals sign when assigning a value; failing to do so causes an error. For example, the MySQL server `--log-error` [427] option has the default value

`host_name.err`, where `host_name` is the name of the host on which MySQL is running. Assume that you are running MySQL on a computer whose host name is “tonfisk”, and consider the following invocation of `mysqld_safe`:

```
shell> mysqld_safe &
[1] 11699
shell> 080112 12:53:40 mysqld_safe Logging to '/usr/local/mysql/var/tonfisk.err'.
080112 12:53:40 mysqld_safe Starting mysqld daemon with databases from /usr/local/mysql/var
shell>
```

After shutting down the server, restart it as follows:

```
shell> mysqld_safe --log-error &
[1] 11699
shell> 080112 12:53:40 mysqld_safe Logging to '/usr/local/mysql/var/tonfisk.err'.
080112 12:53:40 mysqld_safe Starting mysqld daemon with databases from /usr/local/mysql/var
shell>
```

The result is the same, since `--log-error` [227] is not followed by anything else on the command line, and it supplies its own default value. (The `&` character tells the operating system to run MySQL in the background; it is ignored by MySQL itself.) Now suppose that you wish to log errors to a file named `my-errors.err`. You might try starting the server with `--log-error my-errors`, but this does not have the intended effect, as shown here:

```
shell> mysqld_safe --log-error my-errors &
[1] 31357
shell> 080111 22:53:31 mysqld_safe Logging to '/usr/local/mysql/var/tonfisk.err'.
080111 22:53:32 mysqld_safe Starting mysqld daemon with databases from /usr/local/mysql/var
080111 22:53:34 mysqld_safe mysqld from pid file /usr/local/mysql/var/tonfisk.pid ended

[1]+  Done                  ./mysqld_safe --log-error my-errors
```

The server attempted to start using `/usr/local/mysql/var/tonfisk.err` as the error log, but then shut down. Examining the last few lines of this file shows the reason:

```
shell> tail /usr/local/mysql/var/tonfisk.err
080111 22:53:32 InnoDB: Started; log sequence number 0 46409
/usr/local/mysql/libexec/mysqld: Too many arguments (first extra is 'my-errors').
Use --verbose --help to get a list of available options
080111 22:53:32 [ERROR] Aborting

080111 22:53:32 InnoDB: Starting shutdown...
080111 22:53:34 InnoDB: Shutdown completed; log sequence number 0 46409
080111 22:53:34 [Note] /usr/local/mysql/libexec/mysqld: Shutdown complete

080111 22:53:34 mysqld_safe mysqld from pid file /usr/local/mysql/var/tonfisk.pid ended
```

Because the `--log-error` [227] option supplies a default value, you must use an equals sign to assign a different value to it, as shown here:

```
shell> mysqld_safe --log-error=my-errors &
[1] 31437
shell> 080111 22:54:15 mysqld_safe Logging to '/usr/local/mysql/var/my-errors.err'.
080111 22:54:15 mysqld_safe Starting mysqld daemon with databases from /usr/local/mysql/var
shell>
```

Now the server has been started successfully, and is logging errors to the file `/usr/local/mysql/var/my-errors.err`.

Similar issues can arise when specifying option values in option files. For example, consider a `my.cnf` file that contains the following:

```
[mysql]

host
user
```

When the `mysql` client reads this file, these entries are parsed as `--host [250] --user [254]` or `--host=--user [250]`, with the result shown here:

```
shell> mysql
ERROR 2005 (HY000): Unknown MySQL server host '--user' (1)
```

However, in option files, an equals sign is not assumed. Suppose the `my.cnf` file is as shown here:

```
[mysql]
user jon
```

Trying to start `mysql` in this case causes a different error:

```
shell> mysql
mysql: unknown option '--user jon'
```

A similar error would occur if you were to write `host tonfisk` in the option file rather than `host=tonfisk`. Instead, you must use the equals sign:

```
[mysql]
user=jon
```

Now the login attempt succeeds:

```
shell> mysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 5
Server version: 5.6.12 Source distribution

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SELECT USER();
+-----+
| USER() |
+-----+
| jon@localhost |
+-----+
1 row in set (0.00 sec)
```

This is not the same behavior as with the command line, where the equals sign is not required:

```
shell> mysql --user jon --host tonfisk
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 6
Server version: 5.6.12 Source distribution

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SELECT USER();
+-----+
| USER() |
+-----+
| jon@tonfisk |
+-----+
1 row in set (0.00 sec)
```

In MySQL 5.6, specifying an option requiring a value without a value in an option file causes the server to abort with an error. Suppose that `my.cnf` contains the following:

```
[mysqld]
log_error
relay_log
relay_log_index
```

This causes the server to fail on startup, as shown here:

```
shell> mysqld_safe &
```

```
090514 09:48:39 mysqld_safe Logging to '/home/jon/bin/mysql-5.5/var/tonfisk.err'.
090514 09:48:39 mysqld_safe Starting mysqld daemon with databases from /home/jon/bin/mysql-5.5/var
090514 09:48:39 mysqld_safe mysqld from pid file /home/jon/bin/mysql-5.5/var/tonfisk.pid ended
```

The `--log-error` [427] option does not require an argument; however, the `--relay-log` [1974] option requires one, as shown in the error log (which in the absence of a specified value, defaults to `datadir/hostname.err`):

```
shell> tail -n 3 ../var/tonfisk.err

090514 09:48:39 mysqld_safe Starting mysqld daemon with databases from /home/jon/bin/mysql-5.5/var
090514 9:48:39 [ERROR] /home/jon/bin/mysql-5.5/libexec/mysqld: option '--relay-log' requires an argument
090514 9:48:39 [ERROR] Aborting
```

This is a change from previous behavior, where the server would have interpreted the last two lines in the example `my.cnf` file as `--relay-log=relay_log_index` and created a relay log file using “`relay_log_index`” as the basename. (Bug #25192)

## 4.2.4. Setting Environment Variables

Environment variables can be set at the command prompt to affect the current invocation of your command processor, or set permanently to affect future invocations. To set a variable permanently, you can set it in a startup file or by using the interface provided by your system for this purpose. Consult the documentation for your command interpreter for specific details. [Section 2.12, “Environment Variables”](#), lists all environment variables that affect MySQL program operation.

To specify a value for an environment variable, use the syntax appropriate for your command processor. For example, on Windows, you can set the `USER` variable to specify your MySQL account name. To do so, use this syntax:

```
SET USER=your_name
```

The syntax on Unix depends on your shell. Suppose that you want to specify the TCP/IP port number using the `MYSQL_TCP_PORT` variable. Typical syntax (such as for `sh`, `ksh`, `bash`, `zsh`, and so on) is as follows:

```
MYSQL_TCP_PORT=3306
export MYSQL_TCP_PORT
```

The first command sets the variable, and the `export` command exports the variable to the shell environment so that its value becomes accessible to MySQL and other processes.

For `csh` and `tcsh`, use `setenv` to make the shell variable available to the environment:

```
setenv MYSQL_TCP_PORT 3306
```

The commands to set environment variables can be executed at your command prompt to take effect immediately, but the settings persist only until you log out. To have the settings take effect each time you log in, use the interface provided by your system or place the appropriate command or commands in a startup file that your command interpreter reads each time it starts.

On Windows, you can set environment variables using the System Control Panel (under Advanced).

On Unix, typical shell startup files are `.bashrc` or `.bash_profile` for `bash`, or `.tcshrc` for `tcsh`.

Suppose that your MySQL programs are installed in `/usr/local/mysql/bin` and that you want to make it easy to invoke these programs. To do this, set the value of the `PATH` environment variable to include that directory. For example, if your shell is `bash`, add the following line to your `.bashrc` file:

```
PATH=${PATH}:/usr/local/mysql/bin
```

`bash` uses different startup files for login and nonlogin shells, so you might want to add the setting to `.bashrc` for login shells and to `.bash_profile` for nonlogin shells to make sure that `PATH` is set regardless.

If your shell is `tcsh`, add the following line to your `.tcshrc` file:

```
setenv PATH ${PATH}:/usr/local/mysql/bin
```

If the appropriate startup file does not exist in your home directory, create it with a text editor.

After modifying your `PATH` setting, open a new console window on Windows or log in again on Unix so that the setting goes into effect.

## 4.3. MySQL Server and Server-Startup Programs

This section describes `mysqld`, the MySQL server, and several programs that are used to start the server.

### 4.3.1. `mysqld` — The MySQL Server

`mysqld`, also known as MySQL Server, is the main program that does most of the work in a MySQL installation. MySQL Server manages access to the MySQL data directory that contains databases and tables. The data directory is also the default location for other information such as log files and status files.

When MySQL server starts, it listens for network connections from client programs and manages access to databases on behalf of those clients.

The `mysqld` program has many options that can be specified at startup. For a complete list of options, run this command:

```
shell> mysqld --verbose --help
```

MySQL Server also has a set of system variables that affect its operation as it runs. System variables can be set at server startup, and many of them can be changed at runtime to effect dynamic server reconfiguration. MySQL Server also has a set of status variables that provide information about its operation. You can monitor these status variables to access runtime performance characteristics.

For a full description of MySQL Server command options, system variables, and status variables, see [Section 5.1, “The MySQL Server”](#). For information about installing MySQL and setting up the initial configuration, see [Chapter 2, \*Installing and Upgrading MySQL\*](#).

### 4.3.2. `mysqld_safe` — MySQL Server Startup Script

`mysqld_safe` is the recommended way to start a `mysqld` server on Unix. `mysqld_safe` adds some safety features such as restarting the server when an error occurs and logging runtime information to an error log file. A description of error logging is given later in this section.

`mysqld_safe` tries to start an executable named `mysqld`. To override the default behavior and specify explicitly the name of the server you want to run, specify a `--mysqld [228]` or `--mysqld-version [228]` option to `mysqld_safe`. You can also use `--ledir [227]` to indicate the directory where `mysqld_safe` should look for the server.

Many of the options to `mysqld_safe` are the same as the options to `mysqld`. See [Section 5.1.3, “Server Command Options”](#).

Options unknown to `mysqld_safe` are passed to `mysqld` if they are specified on the command line, but ignored if they are specified in the `[mysqld_safe]` group of an option file. See [Section 4.2.3.3, “Using Option Files”](#).



`mysqld_safe` reads all options from the `[mysqld]`, `[server]`, and `[mysqld_safe]` sections in option files. For example, if you specify a `[mysqld]` section like this, `mysqld_safe` will find and use the `--log-error` [227] option:

```
[mysqld]
log-error=error.log
```

For backward compatibility, `mysqld_safe` also reads `[safe_mysqld]` sections, although you should rename such sections to `[mysqld_safe]` in MySQL 5.6 installations.

`mysqld_safe` supports the following options. It also reads option files and supports the options for processing them described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.1. `mysqld_safe` Options**

Format	Option File	Description
<code>--basedir=path</code> [227]	<code>basedir</code> [227]	The path to the MySQL installation directory
<code>--core-file-size=size</code> [227]	<code>core-file-size</code> [227]	The size of the core file that <code>mysqld</code> should be able to create
<code>--datadir=path</code> [227]	<code>datadir</code> [227]	The path to the data directory
<code>--defaults-extra-file=path</code> [227]	<code>defaults-extra-file</code> [227]	The name of an option file to be read in addition to the usual option files
<code>--defaults-file=file_name</code> [227]	<code>defaults-file</code> [227]	The name of an option file to be read instead of the usual option files
<code>--help</code> [227]		Display a help message and exit
<code>--ledir=path</code> [227]	<code>ledir</code> [227]	Use this option to indicate the path name to the directory where the server is located
<code>--log-error=file_name</code> [227]	<code>log-error</code> [227]	Write the error log to the given file
<code>--malloc-lib=[lib-name]</code> [227]	<code>malloc-lib</code> [227]	Alternative malloc library to use for <code>mysqld</code>
<code>--mysqld=prog_name</code> [228]	<code>mysqld</code> [228]	The name of the server program (in the <code>ledir</code> directory) that you want to start
<code>--mysqld-version=suffix</code> [228]	<code>mysqld-version</code> [228]	This option is similar to the <code>--mysqld</code> option, but you specify only the suffix for the server program name
<code>--nice=priority</code> [228]	<code>nice</code> [228]	Use the <code>nice</code> program to set the server's scheduling priority to the given value
<code>--no-defaults</code> [228]	<code>no-defaults</code> [228]	Do not read any option files
<code>--open-files-limit=count</code> [228]	<code>open-files-limit</code> [228]	The number of files that <code>mysqld</code> should be able to open
<code>--pid-file=file_name</code> [228]	<code>pid-file=file_name</code> [228]	The path name of the process ID file
<code>--plugin-dir=path</code>	<code>plugin-dir=path</code>	The directory where plugins are located
<code>--port=number</code> [228]	<code>port</code> [228]	The port number that the server should use when listening for TCP/IP connections
<code>--skip-kill-mysqld</code> [229]	<code>skip-kill-mysqld</code> [229]	Do not try to kill stray <code>mysqld</code> processes
<code>--skip-syslog</code> [229]	<code>skip-syslog</code> [229]	Do not write error messages to syslog; use error log file

Format	Option File	Description
<code>--socket=path [229]</code>	<code>socket [229]</code>	The Unix socket file that the server should use when listening for local connections
<code>--syslog [229]</code>	<code>syslog [229]</code>	Tag suffix for messages written to syslog
<code>--timezone=timezone [229]</code>	<code>timezone [229]</code>	Set the TZ time zone environment variable to the given option value
<code>--user={user_name user_id} [229]</code>	<code>user [229]</code>	Run the mysqld server as the user having the name <code>user_name</code> or the numeric user ID <code>user_id</code>

- `--help [227]`  
Display a help message and exit.
- `--basedir=path [227]`  
The path to the MySQL installation directory.
- `--core-file-size=size [227]`  
The size of the core file that `mysqld` should be able to create. The option value is passed to `ulimit -c`.
- `--datadir=path [227]`  
The path to the data directory.
- `--defaults-extra-file=path [227]`  
The name of an option file to be read in addition to the usual option files. This must be the first option on the command line if it is used. If the file does not exist or is otherwise inaccessible, the server will exit with an error.
- `--defaults-file=file_name [227]`  
The name of an option file to be read instead of the usual option files. This must be the first option on the command line if it is used.
- `--ledir=path [227]`  
If `mysqld_safe` cannot find the server, use this option to indicate the path name to the directory where the server is located.
- `--log-error=file_name [227]`  
Write the error log to the given file. See [Section 5.2.2, “The Error Log”](#).
- `--malloc-lib=[lib_name] [227]`  
The name of the library to use for memory allocation instead of the system `malloc()` library. Any library can be used by specifying its path name, but there is a shortcut form to enable use of the `tcmalloc` library that is shipped with binary MySQL distributions for Linux in MySQL 5.6.  
The `--malloc-lib [227]` option works by modifying the `LD_PRELOAD` environment value to affect dynamic linking to enable the loader to find the memory-allocation library when `mysqld` runs:
  - If the option is not given, or is given without a value (`--malloc-lib= [227]`), `LD_PRELOAD` is not modified and no attempt is made to use `tcmalloc`.
  - If the option is given as `--malloc-lib=tcmalloc [227]`, `mysqld_safe` looks for a `tcmalloc` library in `/usr/lib` and then in the MySQL `pkglibdir` location (for example, `/usr/local/`

`mysql/lib` or whatever is appropriate). If `tmalloc` is found, its path name is added to the beginning of the `LD_PRELOAD` value for `mysqld`. If `tcmalloc` is not found, `mysqld_safe` aborts with an error.

- If the option is given as `--malloc-lib=/path/to/some/library` [227], that full path is added to the beginning of the `LD_PRELOAD` value. If the full path points to a nonexistent or unreadable file, `mysqld_safe` aborts with an error.
- For cases where `mysqld_safe` adds a path name to `LD_PRELOAD`, it adds the path to the beginning of any existing value the variable already has.

Linux users can use the `libtcmalloc_minimal.so` included in binary packages by adding these lines to the `my.cnf` file:

```
[mysqld_safe]
malloc-lib=tcmalloc
```

Those lines also suffice for users on any platform who have installed a `tcmalloc` package in `/usr/lib`. To use a specific `tcmalloc` library, specify its full path name. Example:

```
[mysqld_safe]
malloc-lib=/opt/lib/libtcmalloc_minimal.so
```

- `--mysqld=prog_name` [228]

The name of the server program (in the `ledir` directory) that you want to start. This option is needed if you use the MySQL binary distribution but have the data directory outside of the binary distribution. If `mysqld_safe` cannot find the server, use the `--ledir` [227] option to indicate the path name to the directory where the server is located.

- `--mysqld-version=suffix` [228]

This option is similar to the `--mysqld` [228] option, but you specify only the suffix for the server program name. The basename is assumed to be `mysqld`. For example, if you use `--mysqld-version=debug` [228], `mysqld_safe` starts the `mysqld-debug` program in the `ledir` directory. If the argument to `--mysqld-version` [228] is empty, `mysqld_safe` uses `mysqld` in the `ledir` directory.

- `--nice=priority` [228]

Use the `nice` program to set the server's scheduling priority to the given value.

- `--no-defaults` [228]

Do not read any option files. This must be the first option on the command line if it is used.

- `--open-files-limit=count` [228]

The number of files that `mysqld` should be able to open. The option value is passed to `ulimit -n`. Note that you need to start `mysqld_safe` as `root` for this to work properly!

- `--pid-file=file_name` [228]

The path name of the process ID file.

- `--plugin-dir=path` [228]

The path name of the plugin directory.

- `--port=port_num` [228]

The port number that the server should use when listening for TCP/IP connections. The port number must be 1024 or higher unless the server is started by the `root` system user.

- `--skip-kill-mysqld` [229]

Do not try to kill stray `mysqld` processes at startup. This option works only on Linux.

- `--socket=path` [229]

The Unix socket file that the server should use when listening for local connections.

- `--syslog` [229], `--skip-syslog` [229]

`--syslog` [229] causes error messages to be sent to `syslog` on systems that support the `logger` program. `--skip-syslog` suppresses the use of `syslog`; messages are written to an error log file.

When `syslog` is used, the `daemon.err` syslog priority/facility is used for all log messages.

- `--syslog-tag=tag` [229]

For logging to `syslog`, messages from `mysqld_safe` and `mysqld` are written with a tag of `mysqld_safe` and `mysqld`, respectively. To specify a suffix for the tag, use `--syslog-tag=tag` [229], which modifies the tags to be `mysqld_safe-tag` and `mysqld-tag`.

- `--timezone=timezone` [229]

Set the `TZ` time zone environment variable to the given option value. Consult your operating system documentation for legal time zone specification formats.

- `--user={user_name|user_id}` [229]

Run the `mysqld` server as the user having the name `user_name` or the numeric user ID `user_id`. (“User” in this context refers to a system login account, not a MySQL user listed in the grant tables.)

If you execute `mysqld_safe` with the `--defaults-file` [227] or `--defaults-extra-file` [227] option to name an option file, the option must be the first one given on the command line or the option file will not be used. For example, this command will not use the named option file:

```
mysql> mysqld_safe --port=port_num --defaults-file=file_name
```

Instead, use the following command:

```
mysql> mysqld_safe --defaults-file=file_name --port=port_num
```

The `mysqld_safe` script is written so that it normally can start a server that was installed from either a source or a binary distribution of MySQL, even though these types of distributions typically install the server in slightly different locations. (See [Section 2.1.5, “Installation Layouts”](#).) `mysqld_safe` expects one of the following conditions to be true:

- The server and databases can be found relative to the working directory (the directory from which `mysqld_safe` is invoked). For binary distributions, `mysqld_safe` looks under its working directory for `bin` and `data` directories. For source distributions, it looks for `libexec` and `var` directories. This condition should be met if you execute `mysqld_safe` from your MySQL installation directory (for example, `/usr/local/mysql` for a binary distribution).
- If the server and databases cannot be found relative to the working directory, `mysqld_safe` attempts to locate them by absolute path names. Typical locations are `/usr/local/libexec` and `/usr/local/var`. The actual locations are determined from the values configured into the distribution at the time it was built. They should be correct if MySQL is installed in the location specified at configuration time.

Because `mysqld_safe` tries to find the server and databases relative to its own working directory, you can install a binary distribution of MySQL anywhere, as long as you run `mysqld_safe` from the MySQL installation directory:

```
shell> cd mysql_installation_directory
shell> bin/mysqld_safe &
```

If `mysqld_safe` fails, even when invoked from the MySQL installation directory, you can specify the `--ledir` [227] and `--datadir` [227] options to indicate the directories in which the server and databases are located on your system.

In MySQL 5.6.5 and later, `mysqld_safe` tries to use the `sleep` and `date` system utilities to determine how many times it has attempted to start this second, and—if these are present and this is greater than 5 times—is forced to wait 1 full second before starting again. This is intended to prevent excessive CPU usage in the event of repeated failures. (Bug #11761530, Bug #54035)

When you use `mysqld_safe` to start `mysqld`, `mysqld_safe` arranges for error (and notice) messages from itself and from `mysqld` to go to the same destination.

There are several `mysqld_safe` options for controlling the destination of these messages:

- `--syslog` [229]: Write error messages to `syslog` on systems that support the `logger` program.
- `--skip-syslog` [229]: Do not write error messages to `syslog`. Messages are written to the default error log file (`host_name.err` in the data directory), or to a named file if the `--log-error` [227] option is given.
- `--log-error=file_name` [227]: Write error messages to the named error file.

If none of these options is given, the default is `--skip-syslog` [229].

If `--syslog` [229] and `--log-error` [227] are both given, a warning is issued and `--log-error` [227] takes precedence.

When `mysqld_safe` writes a message, notices go to the logging destination (`syslog` or the error log file) and `stdout`. Errors go to the logging destination and `stderr`.

Normally, you should not edit the `mysqld_safe` script. Instead, configure `mysqld_safe` by using command-line options or options in the `[mysqld_safe]` section of a `my.cnf` option file. In rare cases, it might be necessary to edit `mysqld_safe` to get it to start the server properly. However, if you do this, your modified version of `mysqld_safe` might be overwritten if you upgrade MySQL in the future, so you should make a copy of your edited version that you can reinstall.

### 4.3.3. mysql.server — MySQL Server Startup Script

MySQL distributions on Unix include a script named `mysql.server`. It can be used on systems such as Linux and Solaris that use System V-style run directories to start and stop system services. It is also used by the Mac OS X Startup Item for MySQL.

`mysql.server` can be found in the `support-files` directory under your MySQL installation directory or in a MySQL source distribution.

If you use the Linux server RPM package (`MySQL-server-VERSION.rpm`), the `mysql.server` script will be installed in the `/etc/init.d` directory with the name `mysql`. You need not install it manually. See Section 2.5.1, “Installing MySQL from RPM Packages on Linux”, for more information on the Linux RPM packages.

Some vendors provide RPM packages that install a startup script under a different name such as `mysqld`.

If you install MySQL from a source distribution or using a binary distribution format that does not install `mysql.server` automatically, you can install it manually. Instructions are provided in Section 2.10.1.2, “Starting and Stopping MySQL Automatically”.

`mysql.server` reads options from the `[mysql.server]` and `[mysqld]` sections of option files. For backward compatibility, it also reads `[mysql_server]` sections, although you should rename such sections to `[mysql.server]` when using MySQL 5.6.

`mysql.server` supports the following options.

- `--basedir=path` [231]

The path to the MySQL installation directory.

- `--datadir=path` [231]

The path to the MySQL data directory.

- `--pid-file=file_name` [231]

The path name of the file in which the server should write its process ID.

- `--service-startup-timeout=file_name` [231]

How long in seconds to wait for confirmation of server startup. If the server does not start within this time, `mysql.server` exits with an error. The default value is 900. A value of 0 means not to wait at all for startup. Negative values mean to wait forever (no timeout).

- `--use-mysqld_safe` [231]

Use `mysqld_safe` to start the server. This is the default.

- `--user=user_name` [231]

The login user name to use for running `mysqld`.

#### 4.3.4. `mysqld_multi` — Manage Multiple MySQL Servers

`mysqld_multi` is designed to manage several `mysqld` processes that listen for connections on different Unix socket files and TCP/IP ports. It can start or stop servers, or report their current status.

`mysqld_multi` searches for groups named `[mysqldN]` in `my.cnf` (or in the file named by the `--config-file` option). `N` can be any positive integer. This number is referred to in the following discussion as the option group number, or *GNR*. Group numbers distinguish option groups from one another and are used as arguments to `mysqld_multi` to specify which servers you want to start, stop, or obtain a status report for. Options listed in these groups are the same that you would use in the `[mysqld]` group used for starting `mysqld`. (See, for example, [Section 2.10.1.2, “Starting and Stopping MySQL Automatically”](#).) However, when using multiple servers, it is necessary that each one use its own value for options such as the Unix socket file and TCP/IP port number. For more information on which options must be unique per server in a multiple-server environment, see [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).

To invoke `mysqld_multi`, use the following syntax:

```
shell> mysqld_multi [options] {start|stop|reload|report} [GNR[,GNR] ...]
```

`start`, `stop`, `reload` (stop and restart), and `report` indicate which operation to perform. (`reload` is available as of MySQL 5.6.3.) You can perform the designated operation for a single server or multiple servers, depending on the *GNR* list that follows the option name. If there is no list, `mysqld_multi` performs the operation for all servers in the option file.

Each *GNR* value represents an option group number or range of group numbers. The value should be the number at the end of the group name in the option file. For example, the *GNR* for a group named `[mysqld17]` is 17. To specify a range of numbers, separate the first and last numbers by a dash. The *GNR* value `10-13` represents groups `[mysqld10]` through `[mysqld13]`. Multiple groups or group ranges can be specified on the command line, separated by commas. There must be no whitespace characters (spaces or tabs) in the *GNR* list; anything after a whitespace character is ignored.

This command starts a single server using option group `[mysqld17]`:

```
shell> mysqld_multi start 17
```

This command stops several servers, using option groups `[mysqld8]` and `[mysqld10]` through `[mysqld13]`:

```
shell> mysqld_multi stop 8,10-13
```

For an example of how you might set up an option file, use this command:

```
shell> mysqld_multi --example
```

`mysqld_multi` searches for option files as follows:

- With `--no-defaults` [232], no option files are read.
- With `--defaults-file=file_name` [232], only the named file is read.
- Otherwise, option files in the standard list of locations are read, including any file named by the `--defaults-extra-file=file_name` [232] option, if one is given. (If the option is given multiple times, the last value is used.)

Option files read are searched for `[mysqld_multi]` and `[mysqldN]` option groups. The `[mysqld_multi]` group can be used for options to `mysqld_multi` itself. `[mysqldN]` groups can be used for options passed to specific `mysqld` instances.

The `[mysqld]` or `[mysqld_safe]` groups can be used for common options read by all instances of `mysqld` or `mysqld_safe`. You can specify a `--defaults-file=file_name` [219] option to use a different configuration file for that instance, in which case the `[mysqld]` or `[mysqld_safe]` groups from that file will be used for that instance.

`mysqld_multi` supports the following options.

- `--help` [232]

Display a help message and exit.

- `--example` [232]

Display a sample option file.

- `--log=file_name` [232]

Specify the name of the log file. If the file exists, log output is appended to it.

- `--mysqladmin=prog_name` [232]

The `mysqladmin` binary to be used to stop servers.

- `--mysqld=prog_name` [232]

The `mysqld` binary to be used. Note that you can specify `mysqld_safe` as the value for this option also. If you use `mysqld_safe` to start the server, you can include the `mysqld` or `ledir` options in the corresponding `[mysqldN]` option group. These options indicate the name of the server that `mysqld_safe` should start and the path name of the directory where the server is located. (See the descriptions for these options in [Section 4.3.2, “mysqld\\_safe — MySQL Server Startup Script”](#).)  
Example:

```
[mysqld38]
mysqld = mysqld-debug
ledir  = /opt/local/mysql/libexec
```

- `--no-log` [232]



Print log information to `stdout` rather than to the log file. By default, output goes to the log file.

- `--password=password` [233]

The password of the MySQL account to use when invoking `mysqladmin`. Note that the password value is not optional for this option, unlike for other MySQL programs.

- `--silent` [233]

Silent mode; disable warnings.

- `--tcp-ip` [233]

Connect to each MySQL server through the TCP/IP port instead of the Unix socket file. (If a socket file is missing, the server might still be running, but accessible only through the TCP/IP port.) By default, connections are made using the Unix socket file. This option affects `stop` and `report` operations.

- `--user=user_name` [233]

The user name of the MySQL account to use when invoking `mysqladmin`.

- `--verbose` [233]

Be more verbose.

- `--version` [233]

Display version information and exit.

Some notes about `mysqld_multi`:

- **Most important:** Before using `mysqld_multi` be sure that you understand the meanings of the options that are passed to the `mysqld` servers and *why* you would want to have separate `mysqld` processes. Beware of the dangers of using multiple `mysqld` servers with the same data directory. Use separate data directories, unless you *know* what you are doing. Starting multiple servers with the same data directory does *not* give you extra performance in a threaded system. See [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).



#### Important

Make sure that the data directory for each server is fully accessible to the Unix account that the specific `mysqld` process is started as. *Do not* use the Unix `root` account for this, unless you *know* what you are doing. See [Section 6.1.5, “How to Run MySQL as a Normal User”](#).

- Make sure that the MySQL account used for stopping the `mysqld` servers (with the `mysqladmin` program) has the same user name and password for each server. Also, make sure that the account has the `SHUTDOWN` [756] privilege. If the servers that you want to manage have different user names or passwords for the administrative accounts, you might want to create an account on each server that has the same user name and password. For example, you might set up a common `multi_admin` account by executing the following commands for each server:

```
shell> mysql -u root -S /tmp/mysql.sock -p
Enter password:
mysql> GRANT SHUTDOWN ON *.*
-> TO 'multi_admin'@'localhost' IDENTIFIED BY 'multipass';
```

See [Section 6.2, “The MySQL Access Privilege System”](#). You have to do this for each `mysqld` server. Change the connection parameters appropriately when connecting to each one. Note that the host name part of the account name must permit you to connect as `multi_admin` from the host where you want to run `mysqld_multi`.

- The Unix socket file and the TCP/IP port number must be different for every `mysqld`. (Alternatively, if the host has multiple network addresses, you can use `--bind-address` [414] to cause different servers to listen to different interfaces.)
- The `--pid-file` [228] option is very important if you are using `mysqld_safe` to start `mysqld` (for example, `--mysqld=mysqld_safe` [228]). Every `mysqld` should have its own process ID file. The advantage of using `mysqld_safe` instead of `mysqld` is that `mysqld_safe` monitors its `mysqld` process and restarts it if the process terminates due to a signal sent using `kill -9` or for other reasons, such as a segmentation fault. Please note that the `mysqld_safe` script might require that you start it from a certain place. This means that you might have to change location to a certain directory before running `mysqld_multi`. If you have problems starting, please see the `mysqld_safe` script. Check especially the lines:

```
-----
MY_PWD=`pwd`
# Check if we are starting this relative (for the binary release)
if test -d $MY_PWD/data/mysql -a \
  -f ./share/mysql/english/errmsg.sys -a \
  -x ./bin/mysqld
-----
```

The test performed by these lines should be successful, or you might encounter problems. See [Section 4.3.2, “mysqld\\_safe — MySQL Server Startup Script”](#).

- You might want to use the `--user` [446] option for `mysqld`, but to do this you need to run the `mysqld_multi` script as the Unix superuser (`root`). Having the option in the option file doesn't matter; you just get a warning if you are not the superuser and the `mysqld` processes are started under your own Unix account.

The following example shows how you might set up an option file for use with `mysqld_multi`. The order in which the `mysqld` programs are started or stopped depends on the order in which they appear in the option file. Group numbers need not form an unbroken sequence. The first and fifth [`mysqldN`] groups were intentionally omitted from the example to illustrate that you can have “gaps” in the option file. This gives you more flexibility.

```
# This file should probably be in your home dir (~/.my.cnf)
# or /etc/my.cnf
# Version 2.1 by Jani Tolonen

[mysqld_multi]
mysqld      = /usr/local/bin/mysqld_safe
mysqladmin  = /usr/local/bin/mysqladmin
user        = multi_admin
password    = multipass

[mysqld2]
socket      = /tmp/mysql.sock2
port        = 3307
pid-file    = /usr/local/mysql/var2/hostname.pid2
datadir     = /usr/local/mysql/var2
language    = /usr/local/share/mysql/english
user        = john

[mysqld3]
socket      = /tmp/mysql.sock3
port        = 3308
pid-file    = /usr/local/mysql/var3/hostname.pid3
datadir     = /usr/local/mysql/var3
language    = /usr/local/share/mysql/swedish
user        = monty

[mysqld4]
socket      = /tmp/mysql.sock4
port        = 3309
pid-file    = /usr/local/mysql/var4/hostname.pid4
datadir     = /usr/local/mysql/var4
language    = /usr/local/share/mysql/estonia
```

```

user      = tonu

[mysqld6]
socket    = /tmp/mysql.sock6
port      = 3311
pid-file  = /usr/local/mysql/var6/hostname.pid6
datadir   = /usr/local/mysql/var6
language  = /usr/local/share/mysql/japanese
user      = jani

```

See [Section 4.2.3.3, “Using Option Files”](#).

## 4.4. MySQL Installation-Related Programs

The programs in this section are used when installing or upgrading MySQL.

### 4.4.1. `comp_err` — Compile MySQL Error Message File

`comp_err` creates the `errmsg.sys` file that is used by `mysqld` to determine the error messages to display for different error codes. `comp_err` normally is run automatically when MySQL is built. It compiles the `errmsg.sys` file from the plaintext file located at `sql/share/errmsg.txt` in MySQL source distributions.

`comp_err` also generates `mysqld_error.h`, `mysqld_ename.h`, and `sql_state.h` header files.

For more information about how error messages are defined, see the [MySQL Internals Manual](#).

Invoke `comp_err` like this:

```
shell> comp_err [options]
```

`comp_err` supports the following options.

- `--help [235], -?`

Display a help message and exit.

- `--charset=path [235], -C path`

The character set directory. The default is `../sql/share/charsets`.

- `--debug=debug_options [235], -# debug_options`

Write a debugging log. A typical `debug_options` string is `'d:t:O,file_name'`. The default is `'d:t:O,/tmp/comp_err.trace'`.

- `--debug-info [235], -T`

Print some debugging information when the program exits.

- `--header_file=file_name [235], -H file_name`

The name of the error header file. The default is `mysqld_error.h`.

- `--in_file=file_name [235], -F file_name`

The name of the input file. The default is `../sql/share/errmsg.txt`.

- `--name_file=file_name [235], -N file_name`

The name of the error name file. The default is `mysqld_ename.h`.

- `--out_dir=path [235], -D path`

The name of the output base directory. The default is `../sql/share/`.

- `--out_file=file_name` [236], `-O file_name`

The name of the output file. The default is `errmsg.sys`.

- `--statefile=file_name` [236], `-S file_name`

The name for the SQLSTATE header file. The default is `sql_state.h`.

- `--version` [236], `-V`

Display version information and exit.

## 4.4.2. mysqlbug — Generate Bug Report

This program is obsolete.

The normal way to report bugs is to visit <http://bugs.mysql.com/>, which is the address for our bugs database. This database is public and can be browsed and searched by anyone. If you log in to the system, you can enter new reports.

## 4.4.3. mysql\_install\_db — Initialize MySQL Data Directory

`mysql_install_db` initializes the MySQL data directory and creates the system tables that it contains, if they do not exist. It also initializes the [system tablespace](#) and related data structures needed to manage [InnoDB](#) tables. As of MySQL 5.6.8, `mysql_install_db` is a Perl script and can be used on any system with Perl installed. Before 5.6.8, it is a shell script and is available only on Unix platforms.

As of MySQL 5.6.8, on Unix platforms, `mysql_install_db` creates a default option file named `my.cnf` in the base installation directory. This file is created from a template included in the distribution package named `my-default.cnf`. You can find the template in or under the base installation directory. When started using `mysqld_safe`, the server uses `my.cnf` file by default. If `my.cnf` already exists, `mysql_install_db` assumes it to be in use and writes a new file named `my-new.cnf` instead.

With one exception, the settings in the default option file are commented and have no effect. The exception is that the file changes the `sql_mode` [545] system variable from its default of `NO_ENGINE_SUBSTITUTION` [605] to also include `STRICT_TRANS_TABLES` [607]. This setting produces a server configuration that results in errors rather than warnings for bad data in operations that modify transactional tables. See [Section 5.1.7, “Server SQL Modes”](#).

To invoke `mysql_install_db`, use the following syntax:

```
shell> mysql_install_db [options]
```

Because the MySQL server, `mysqld`, needs to access the data directory when it runs later, you should either run `mysql_install_db` from the same system account that will be used for running `mysqld` or run it as `root` and use the `--user` [238] option to indicate the user name that `mysqld` will run as. It might be necessary to specify other options such as `--basedir` [237] or `--datadir` [237] if `mysql_install_db` does not use the correct locations for the installation directory or data directory. For example:

```
shell> scripts/mysql_install_db --user=mysql \
      --basedir=/opt/mysql/mysql \
      --datadir=/opt/mysql/mysql/data
```

`mysql_install_db` needs to invoke `mysqld` with the `--bootstrap` [416] and `--skip-grant-tables` [439] options.

**Note**

If you have set a custom `TMPDIR` environment variable when performing the installation, and the specified directory is not accessible, `mysql_install_db` may fail. If so, unset `TMPDIR` or set `TMPDIR` to point to the system temporary directory (usually `/tmp`).

After `mysql_install_db` sets up the InnoDB system tablespace, changes to some of tablespace characteristics require setting up a whole new instance. This includes the file name of the first file in the system tablespace and the number of undo logs. If you do not want to use the default values, make sure that the settings for the `innodb_data_file_path` [1714] and `innodb_log_file_size` [1732] configuration options are in place in the MySQL configuration file before running `mysql_install_db`.

`mysql_install_db` supports the following options, which can be specified on the command line or in the `[mysql_install_db]` group of an option file. (Options that are common to `mysqld` can also be specified in the `[mysqld]` group.) Other options are passed to `mysqld`. For information about option files, see Section 4.2.3.3, “Using Option Files”. `mysql_install_db` also supports the options for processing option files described at Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”.

- `--basedir=path` [237]

The path to the MySQL installation directory.

- `--datadir=path` [237], `--ldata=path` [237]

The path to the MySQL data directory. Beginning with MySQL 5.6.8, `mysql_install_db` is more strict about the option value. Only the last component of the path name is created if it does not exist; the parent directory must already exist or an error occurs.

- `--force` [237]

Cause `mysql_install_db` to run even if DNS does not work. Grant table entries that normally use host names will use IP addresses.

- `--random-passwords` [237]

On Unix platforms, this option provides for more secure MySQL installation. Invoking `mysql_install_db` with `--random-passwords` [237] causes it to perform the following actions in addition to its normal operation:

- Create a random password, assign it to the initial MySQL `root` accounts, and set the “password expired” flag for those accounts.
- Write the initial password file to the `.mysql_secret` file in the directory named by the `HOME` environment variable. Depending on operating system, using a command such as `sudo` may cause the value of `HOME` to refer to the home directory of the `root` system user.

If `.mysql_secret` already exists, the new password information is appended to it. Each password entry includes a timestamp so that in the event of multiple install operations it is possible to determine the password associated with each one.

`.mysql_secret` is created with mode 600 to be accessible only to the system user for whom it is created.

- Remove the anonymous-user MySQL accounts.

As a result of these actions, it is necessary after installation to start the server, connect as `root` using the password written to the `.mysql_secret` file, and to assign a new `root` password. Until

this is done, `root` cannot do anything else. This must be done for each `root` account you intend to use. To change the password, you can use the `SET PASSWORD` statement (for example, with the `mysql` client). You can also use `mysqladmin` or `mysql_secure_installation`.

New RPM install operations (not upgrades) invoke `mysql_install_db` with the `--random-passwords` option. (Install operations using RPMs for Unbreakable Linux Network are unaffected because they do not use `mysql_install_db`.)

As of MySQL 5.6.9, new Solaris PKG install operations (not upgrades) invoke `mysql_install_db` with the `--random-passwords` option.

For install operations using a binary `.tar.gz` distribution or a source distribution, you can invoke `mysql_install_db` with the `--random-passwords` option manually to make your MySQL installation more secure. This is recommended, particularly for sites with sensitive data.

This option was added in MySQL 5.6.8.

- `--rpm` [238]

For internal use. This option is used during the MySQL installation process for install operations performed using RPM packages.

- `--skip-name-resolve` [238]

Use IP addresses rather than host names when creating grant table entries. This option can be useful if your DNS does not work.

- `--srcdir=path` [238]

For internal use. This option specifies the directory under which `mysql_install_db` looks for support files such as the error message file and the file for populating the help tables.

- `--user=user_name` [238]

The system (login) user name to use for running `mysqld`. Files and directories created by `mysqld` will be owned by this user. You must be `root` to use this option. By default, `mysqld` runs using your current login name and files and directories that it creates will be owned by you.

- `--verbose` [238]

Verbose mode. Print more information about what the program does.

- `--windows` [238]

For internal use. This option is used for creating Windows distributions.

#### 4.4.4. `mysql_plugin` — Configure MySQL Server Plugins

The `mysql_plugin` utility enables MySQL administrators to manage which plugins a MySQL server loads. It provides an alternative to manually specifying the `--plugin-load` [435] option at server startup or using the `INSTALL PLUGIN` and `UNINSTALL PLUGIN` statements at runtime. `mysql_plugin` is available as of MySQL 5.6.3.

Depending on whether `mysql_plugin` is invoked to enable or disable plugins, it inserts or deletes rows in the `mysql.plugin` table that serves as a plugin registry. (To perform this operation, `mysql_plugin` invokes the MySQL server in bootstrap mode. This means that the server must not already be running.) For normal server startups, the server loads and enables plugins listed in `mysql.plugin` automatically. For additional control over plugin activation, use `--plugin_name` options named for specific plugins, as described in [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

Each invocation of `mysql_plugin` reads a configuration file to determine how to configure the plugins contained in a single plugin library object file. To invoke `mysql_plugin`, use this syntax:

```
mysql_plugin [options] plugin {ENABLE|DISABLE}
```

*plugin* is the name of the plugin to configure. `ENABLE` or `DISABLE` (not case sensitive) specify whether to enable or disable components of the plugin library named in the configuration file. The order of the *plugin* and `ENABLE` or `DISABLE` arguments does not matter.

For example, to configure components of a plugin library file named `myplugins.so` on Linux or `myplugins.dll` on Windows, specify a *plugin* value of `myplugins`. Suppose that this plugin library contains three plugins, `plugin1`, `plugin2`, and `plugin3`, all of which should be configured under `mysql_plugin` control. By convention, configuration files have a suffix of `.ini` and the same basename as the plugin library, so the default configuration file name for this plugin library is `myplugins.ini`. The configuration file contents look like this:

```
myplugins
plugin1
plugin2
plugin3
```

The first line in the `myplugins.ini` file is the name of the library object file, without any extension such as `.so` or `.dll`. The remaining lines are the names of the components to be enabled or disabled. Each value in the file should be on a separate line. Lines on which the first character is `'#'` are taken as comments and ignored.

To enable the plugins listed in the configuration file, invoke `mysql_plugin` this way:

```
shell> mysql_plugin myplugins ENABLE
```

To disable the plugins, use `DISABLE` rather than `ENABLE`.

An error occurs if `mysql_plugin` cannot find the configuration file or plugin library file, or if `mysql_plugin` cannot start the MySQL server.

`mysql_plugin` supports the following options, which can be specified on the command line or in the `[mysqld]` group of any option file. For options specified in a `[mysqld]` group, `mysql_plugin` recognizes the `--basedir` [240], `--datadir` [240], and `--plugin-dir` [240] options and ignores others. For information about option files, see [Section 4.2.3.3, “Using Option Files”](#).

**Table 4.2. `mysql_plugin` Options**

Format	Option File	Description
<code>--basedir=path</code> [240]	<code>basedir=path</code> [240]	The server base directory
<code>--datadir=path</code> [240]	<code>datadir=path</code> [240]	The server data directory
<code>--help</code> [240]		Display help message and exit
<code>--my-print-defaults=path</code> [240]	<code>my-print-defaults=path</code> [240]	The path to <code>my_print_defaults</code>
<code>--mysqld=path</code> [240]	<code>mysqld=path</code> [240]	The path to the server
<code>--no-defaults</code> [240]	<code>no-defaults</code> [240]	Do not read configuration file
<code>--plugin-dir=path</code> [240]	<code>plugin-dir=path</code> [240]	The directory where plugins are located
<code>--plugin-ini=file_name</code> [240]	<code>plugin-ini=file_name</code> [240]	The plugin configuration file
<code>--print-defaults</code> [240]	<code>print-defaults</code> [240]	Show configuration file defaults
<code>--verbose</code> [240]		Verbose mode



Format	Option File	Description
<code>--version [240]</code>		Display version information and exit

- `--help [240], -?`  
Display a help message and exit.
- `--basedir=path [240], -b path`  
The server base directory.
- `--datadir=path [240], -d path`  
The server data directory.
- `--my-print-defaults=path [240], -b path`  
The path to the `my_print_defaults` program.
- `--mysqld=path [240], -b path`  
The path to the `mysqld` server.
- `--no-defaults [240], -p`  
Do not read values from the configuration file. This option enables an administrator to skip reading defaults from the configuration file.  
  
With `mysql_plugin`, this option need not be given first on the command line, unlike most other MySQL programs that support `--no-defaults [240]`.
- `--plugin-dir=path [240], -p path`  
The server plugin directory.
- `--plugin-ini=file_name [240], -i file_name`  
The `mysql_plugin` configuration file. Relative path names are interpreted relative to the current directory. If this option is not given, the default is `plugin.ini` in the plugin directory, where `plugin` is the `plugin` argument on the command line.
- `--print-defaults [240], -P`  
Display the default values from the configuration file. This option causes `mysql_plugin` to print the defaults for `--basedir [240]`, `--datadir [240]`, and `--plugin-dir [240]` if they are found in the configuration file. If no value for a variable is found, nothing is shown.  
  
With `mysql_plugin`, this option need not be given first on the command line, unlike most other MySQL programs that support `--print-defaults [240]`.
- `--verbose [240], -v`  
Verbose mode. Print more information about what the program does. This option can be used multiple times to increase the amount of information.
- `--version [240], -V`  
Display version information and exit.

#### 4.4.5. mysql\_secure\_installation — Improve MySQL Installation Security

This program enables you to improve the security of your MySQL installation in the following ways:

- You can set a password for `root` accounts.
- You can remove `root` accounts that are accessible from outside the local host.
- You can remove anonymous-user accounts.
- You can remove the `test` database (which by default can be accessed by all users, even anonymous users), and privileges that permit anyone to access databases with names that start with `test_`.

`mysql_secure_installation` helps you implement security recommendations similar to those described at [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

Invoke `mysql_secure_installation` without arguments:

```
shell> mysql_secure_installation
```

The script will prompt you to determine which actions to perform.

`mysql_secure_installation` is not available on Windows.

#### 4.4.6. `mysql_tzinfo_to_sql` — Load the Time Zone Tables

The `mysql_tzinfo_to_sql` program loads the time zone tables in the `mysql` database. It is used on systems that have a `zoneinfo` database (the set of files describing time zones). Examples of such systems are Linux, FreeBSD, Solaris, and Mac OS X. One likely location for these files is the `/usr/share/zoneinfo` directory (`/usr/share/lib/zoneinfo` on Solaris). If your system does not have a `zoneinfo` database, you can use the downloadable package described in [Section 10.6, “MySQL Server Time Zone Support”](#).

`mysql_tzinfo_to_sql` can be invoked several ways:

```
shell> mysql_tzinfo_to_sql tz_dir
shell> mysql_tzinfo_to_sql tz_file tz_name
shell> mysql_tzinfo_to_sql --leap tz_file
```

For the first invocation syntax, pass the `zoneinfo` directory path name to `mysql_tzinfo_to_sql` and send the output into the `mysql` program. For example:

```
shell> mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root mysql
```

`mysql_tzinfo_to_sql` reads your system's time zone files and generates SQL statements from them. `mysql` processes those statements to load the time zone tables.

The second syntax causes `mysql_tzinfo_to_sql` to load a single time zone file `tz_file` that corresponds to a time zone name `tz_name`:

```
shell> mysql_tzinfo_to_sql tz_file tz_name | mysql -u root mysql
```

If your time zone needs to account for leap seconds, invoke `mysql_tzinfo_to_sql` using the third syntax, which initializes the leap second information. `tz_file` is the name of your time zone file:

```
shell> mysql_tzinfo_to_sql --leap tz_file | mysql -u root mysql
```

After running `mysql_tzinfo_to_sql`, it is best to restart the server so that it does not continue to use any previously cached time zone data.

#### 4.4.7. `mysql_upgrade` — Check and Upgrade MySQL Tables

`mysql_upgrade` examines all tables in all databases for incompatibilities with the current version of MySQL Server. `mysql_upgrade` also upgrades the system tables so that you can take advantage of new privileges or capabilities that might have been added.

`mysql_upgrade` should be executed each time you upgrade MySQL.

If `mysql_upgrade` finds that a table has a possible incompatibility, it performs a table check and, if problems are found, attempts a table repair. If the table cannot be repaired, see [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#) for manual table repair strategies.



#### Note

On Windows Server 2008, Vista, and newer, you must run `mysql_upgrade` with administrator privileges. You can do this by running a Command Prompt as Administrator and running the command. Failure to do so may result in the upgrade failing to execute correctly.



#### Caution

You should always back up your current MySQL installation *before* performing an upgrade. See [Section 7.2, “Database Backup Methods”](#).

Some upgrade incompatibilities may require special handling before you upgrade your MySQL installation and run `mysql_upgrade`. See [Section 2.11.1, “Upgrading MySQL”](#), for instructions on determining whether any such incompatibilities apply to your installation and how to handle them.

To use `mysql_upgrade`, make sure that the server is running, and then invoke it like this:

```
shell> mysql_upgrade [options]
```

After running `mysql_upgrade`, stop the server and restart it so that any changes made to the system tables take effect.

`mysql_upgrade` executes the following commands to check and repair tables and to upgrade the system tables:

```
mysqlcheck --all-databases --check-upgrade --auto-repair
mysql < fix_priv_tables
mysqlcheck --all-databases --check-upgrade --fix-db-names --fix-table-names
```

Notes about the preceding commands:

- Because `mysql_upgrade` invokes `mysqlcheck` with the `--all-databases` [276] option, it processes all tables in all databases, which might take a long time to complete. Each table is locked and therefore unavailable to other sessions while it is being processed. Check and repair operations can be time-consuming, particularly for large tables.
- For details about what checks the `--check-upgrade` [276] option entails, see the description of the `FOR UPGRADE` option of the `CHECK TABLE` statement (see [Section 13.7.2.2, “CHECK TABLE Syntax”](#)).
- `fix_priv_tables` represents a script generated internally by `mysql_upgrade` that contains SQL statements to upgrade the tables in the `mysql` database.

All checked and repaired tables are marked with the current MySQL version number. This ensures that next time you run `mysql_upgrade` with the same version of the server, it can tell whether there is any need to check or repair the table again.

`mysql_upgrade` also saves the MySQL version number in a file named `mysql_upgrade_info` in the data directory. This is used to quickly check whether all tables have been checked for this release so that table-checking can be skipped. To ignore this file and perform the check regardless, use the `--force` [243] option.

If you install MySQL from RPM packages on Linux, you must install the server and client RPMs. `mysql_upgrade` is included in the server RPM but requires the client RPM because the latter includes `mysqlcheck`. (See [Section 2.5.1, “Installing MySQL from RPM Packages on Linux”](#).)

`mysql_upgrade` does not upgrade the contents of the help tables. For upgrade instructions, see [Section 5.1.10, “Server-Side Help”](#).

`mysql_upgrade` supports the following options, which can be specified on the command line or in the `[mysql_upgrade]` and `[client]` groups of an option file. Other options are passed to `mysqlcheck`. For example, it might be necessary to specify the `--password[=password]` [\[277\]](#) option. `mysql_upgrade` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

- `--help` [\[243\]](#)  
Display a short help message and exit.
- `--basedir=path` [\[243\]](#)  
The path to the MySQL installation directory. This option is accepted for backward compatibility but ignored.
- `--datadir=path` [\[243\]](#)  
The path to the data directory. This option is accepted for backward compatibility but ignored.
- `--debug-check` [\[243\]](#)  
Print some debugging information when the program exits.
- `--debug-info` [\[243\]](#), `-T`  
Print debugging information and memory and CPU usage statistics when the program exits.
- `--default-auth=plugin` [\[243\]](#)  
The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).  
This option was added in MySQL 5.6.2.
- `--force` [\[243\]](#)  
Ignore the `mysql_upgrade_info` file and force execution of `mysqlcheck` even if `mysql_upgrade` has already been executed for the current version of MySQL.
- `--plugin-dir=path` [\[243\]](#)  
The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [\[243\]](#) option is used to specify an authentication plugin but `mysql_upgrade` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).  
This option was added in MySQL 5.6.2.
- `--tmpdir=path` [\[243\]](#), `-t path`  
The path name of the directory to use for creating temporary files.
- `--upgrade-system-tables` [\[243\]](#), `-s`  
Upgrade only the system tables, do not upgrade data.
- `--user=user_name` [\[243\]](#), `-u user_name`  
The MySQL user name to use when connecting to the server. The default user name is `root`.
- `--verbose` [\[243\]](#)  
Verbose mode. Print more information about what the program does.

- `--version-check` [244], `-k`

Check the version of the server to which `mysql_upgrade` is connecting to verify that it is the same as the version for which `mysql_upgrade` was built. If not, `mysql_upgrade` exits. This option is enabled by default; to disable the check, use `--skip-version-check`. This option was added in MySQL 5.6.12.

- `--write-binlog` [244]

Cause binary logging to be enabled while `mysql_upgrade` runs. In MySQL 5.6.6 and earlier, this was the default behavior. (To disable binary logging during the upgrade, it was necessary to use the inverse of this option, by starting the program with `--skip-write-binlog`.) Beginning with MySQL 5.6.7, binary logging by `mysql_upgrade` is disabled by default (Bug #14221043), and you must invoke the program explicitly with `--write-binlog` if you want its actions to be written to the binary log. (Also beginning with MySQL 5.6.7, the `--skip-write-binlog` option effectively does nothing.)

Running `mysql_upgrade` is not recommended with a MySQL Server that is running with global transaction identifiers enabled (Bug #13833710). This is because enabling GTIDs means that any updates which `mysql_upgrade` might need to perform on system tables using a nontransactional storage engine such as `MyISAM` to fail. See [Section 16.1.3.4, “Restrictions on Replication with GTIDs”](#), for more information.

## 4.5. MySQL Client Programs

This section describes client programs that connect to the MySQL server.

### 4.5.1. `mysql` — The MySQL Command-Line Tool

`mysql` is a simple SQL shell with input line editing capabilities. It supports interactive and noninteractive use. When used interactively, query results are presented in an ASCII-table format. When used noninteractively (for example, as a filter), the result is presented in tab-separated format. The output format can be changed using command options.

If you have problems due to insufficient memory for large result sets, use the `--quick` [252] option. This forces `mysql` to retrieve results from the server a row at a time rather than retrieving the entire result set and buffering it in memory before displaying it. This is done by returning the result set using the `mysql_use_result()` C API function in the client/server library rather than `mysql_store_result()`.

Using `mysql` is very easy. Invoke it from the prompt of your command interpreter as follows:

```
shell> mysql db_name
```

Or:

```
shell> mysql --user=user_name --password=your_password db_name
```

Then type an SQL statement, end it with “;”, `\g`, or `\G` and press Enter.

Typing **Control+C** causes `mysql` to attempt to kill the current statement. If this cannot be done, or **Control+C** is typed again before the statement is killed, `mysql` exits.

You can execute SQL statements in a script file (batch file) like this:

```
shell> mysql db_name < script.sql > output.tab
```

On Unix, the `mysql` client logs statements executed interactively to a history file. See [Section 4.5.1.3, “mysql Logging”](#).

#### 4.5.1.1. `mysql` Options

mysql supports the following options, which can be specified on the command line or in the [mysql] and [client] groups of an option file. mysql also supports the options for processing option files described at Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”.

**Table 4.3. mysql Options**

Format	Option File	Description	Introduced
--auto-rehash [248]	auto-rehash [248]	Enable automatic rehashing	
--auto-vertical-output [248]	auto-vertical-output [248]	Enable automatic vertical result set display	
--batch [248]	batch [248]	Don't use history file	
--binary-mode [248]	binary-mode [248]	Disable \r\n - to - \n translation and treatment of \0 as end-of-query	
--bind-address=ip_address [248]	bind-address [248]	Use the specified network interface to connect to the MySQL Server	
--character-sets-dir=path [248]	character-sets-dir [248]	Set the default character set	
--column-names [249]	column-names [249]	Write column names in results	
--column-type-info [249]	column-type-info [249]	Display result set metadata	
--comments [249]	comments [249]	Whether to retain or strip comments in statements sent to the server	
--compress [249]	compress [249]	Compress all information sent between the client and the server	
--connect-expired-password [249]		Indicate to server that client can handle expired-password sandbox mode.	5.6.12
--connect_timeout=value [255]	connect_timeout [255]	Set the number of seconds before connection timeout	
--database=dbname [249]	database [249]	The database to use	
--debug[=debug_options] [249]	debug [249]	Write a debugging log	
--debug-check [249]	debug-check [249]	Print debugging information when the program exits	
--debug-info [249]	debug-info [249]	Print debugging information, memory and CPU statistics when the program exits	
--default-auth=plugin [249]	default-auth=plugin [249]	The authentication plugin to use	
--default-character-set=charset_name [249]	default-character-set [249]	Use charset_name as the default character set	
--delimiter=str [249]	delimiter [249]	Set the statement delimiter	
--enable-cleartext-plugin [250]	enable-cleartext-plugin [250]	Enable cleartext authentication plugin	5.6.7
--execute=statement [250]	execute [250]	Execute the statement and quit	
--force [250]	force [250]	Continue even if an SQL error occurs	

Format	Option File	Description	Introduced
<code>--help</code> [248]		Display help message and exit	
<code>--histignore=pattern_list</code> [250]	<code>histignore</code> [250] <code>list</code> [250]	Patterns specifying which statements to ignore for logging	5.6.8
<code>--host=host_name</code> [250]	<code>host</code> [250]	Connect to the MySQL server on the given host	
<code>--html</code> [250]	<code>html</code> [250]	Produce HTML output	
<code>--ignore-spaces</code> [250]	<code>ignore-spaces</code> [250]	Ignore spaces after function names	
<code>--init-command=str</code> [250]	<code>init-command</code> [250]	SQL statement to execute after connecting	
<code>--line-numbers</code> [250]	<code>line-numbers</code> [250]	Write line numbers for errors	
<code>--local-infile[={0 1}]</code> [250]	<code>local-infile</code> [250]	Enable or disable for LOCAL capability for LOAD DATA INFILE	
<code>--login-path=name</code> [219]		Read login path options from <code>.mylogin.cnf</code>	5.6.6
<code>--max_allowed_packet=value</code> [255]	<code>max_allowed_packet</code> [255]	The maximum packet length to send to or receive from the server	
<code>--max_join_size=value</code> [255]	<code>max_join_size</code> [255]	The automatic limit for rows in a join when using <code>--safe-updates</code>	
<code>--named-commands</code> [250]	<code>named-commands</code> [250]	Enable named mysql commands	
<code>--net_buffer_length=value</code>	<code>net_buffer_length</code>	The buffer size for TCP/IP and socket communication	
<code>--no-auto-rehash</code> [251]		Disable automatic rehashing	
<code>--no-beep</code> [251]	<code>no-beep</code> [251]	Do not beep when errors occur	
<code>--one-database</code> [251]	<code>one-database</code> [251]	Ignore statements except those for the default database named on the command line	
<code>--pager[=command]</code> [251]	<code>pager</code> [251]	Use the given command for paging query output	
<code>--password[=password]</code> [251]	<code>password</code> [251]	The password to use when connecting to the server	
<code>--plugin-dir=path</code> [252]	<code>plugin-dir=path</code> [252]	The directory where plugins are located	
<code>--port=port_num</code> [252]	<code>port</code> [252]	The TCP/IP port number to use for the connection	
<code>--prompt=format_str</code> [252]	<code>prompt</code> [252]	Set the prompt to the specified format	
<code>--protocol=type</code> [252]	<code>protocol</code> [252]	The connection protocol to use	
<code>--quick</code> [252]	<code>quick</code> [252]	Do not cache each query result	
<code>--raw</code> [252]	<code>raw</code> [252]	Write column values without escape conversion	
<code>--reconnect</code> [252]	<code>reconnect</code> [252]	If the connection to the server is lost, automatically try to reconnect	
<code>--safe-updates</code> [252]	<code>safe-updates</code> [252]	Allow only UPDATE and DELETE statements that specify key values	



Format	Option File	Description	Introduced
<code>--secure-auth</code> [253]	<code>secure-auth</code> [253]	Do not send passwords to the server in old (pre-4.1.1) format	
<code>--select_limit=value</code>	<code>select_limit</code> [255]	The automatic limit for SELECT statements when using <code>--safe-updates</code>	
<code>--server-public-key-path=file_name</code> [253]	<code>server-public-key-path=file_name</code> [253]	Display help message and exit	5.6.6
<code>--show-warnings</code> [253]	<code>show-warnings</code> [253]	Show warnings after each statement if there are any	
<code>--sigint-ignore</code> [253]	<code>sigint-ignore</code> [253]	Ignore SIGINT signals (typically the result of typing Control+C)	
<code>--silent</code> [253]	<code>silent</code> [253]	Silent mode	
<code>--skip-auto-rehash</code> [248]	<code>skip-auto-rehash</code> [248]	Disable automatic rehashing	
<code>--skip-column-names</code> [253]	<code>skip-column-names</code> [253]	Do not write column names in results	
<code>--skip-line-numbers</code> [253]	<code>skip-line-numbers</code> [253]	Skip line numbers for errors	
<code>--skip-named-commands</code> [250]	<code>skip-named-commands</code> [250]	Disable named mysql commands	
<code>--skip-pager</code> [251]	<code>skip-pager</code> [251]	Disable paging	
<code>--skip-reconnect</code> [252]	<code>skip-reconnect</code> [252]	Disable reconnecting	
<code>--socket=path</code> [253]	<code>socket</code> [253]	For connections to localhost	
<code>--ssl-ca=file_name</code> [253]	<code>ssl-ca</code> [253]	The path to a file that contains a list of trusted SSL CAs	
<code>--ssl-capath=dir_name</code> [253]	<code>ssl-capath</code> [253]	The path to a directory that contains trusted SSL CA certificates in PEM format	
<code>--ssl-cert=file_name</code> [253]	<code>ssl-cert</code> [253]	The name of the SSL certificate file to use for establishing a secure connection	
<code>--ssl-cipher=cipher_list</code> [253]	<code>ssl-cipher</code> [253]	A list of allowable ciphers to use for SSL encryption	
<code>--ssl-crl=file_name</code> [253]	<code>ssl-crl</code> [253]	The path to a file that contains certificate revocation lists	5.6.3
<code>--ssl-crlpath=dir_name</code> [253]	<code>ssl-crlpath</code> [253]	The path to a directory that contains certificate revocation list files	5.6.3
<code>--ssl-key=file_name</code> [253]	<code>ssl-key</code> [253]	The name of the SSL key file to use for establishing a secure connection	
<code>--ssl-verify-server-cert</code> [253]	<code>ssl-verify-server-cert</code> [253]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<code>--table</code> [254]	<code>table</code> [254]	Display output in tabular format	
<code>--tee=file_name</code> [254]	<code>tee</code> [254]	Append a copy of output to the given file	

Format	Option File	Description	Introduced
<code>--unbuffered</code> [254]	<code>unbuffered</code> [254]	Flush the buffer after each query	
<code>--user=user_name</code> [254]	<code>user</code> [254]	MySQL user name to use when connecting to server	
<code>--verbose</code> [254]		Verbose mode	
<code>--version</code> [254]		Display version information and exit	
<code>--vertical</code> [254]	<code>vertical</code> [254]	Print query output rows vertically (one line per column value)	
<code>--wait</code> [254]	<code>wait</code> [254]	If the connection cannot be established, wait and retry instead of aborting	
<code>--xml</code> [254]	<code>xml</code> [254]	Produce XML output	

- `--help` [248], `-?`

Display a help message and exit.

- `--auto-rehash` [248]

Enable automatic rehashing. This option is on by default, which enables database, table, and column name completion. Use `--disable-auto-rehash` [248] to disable rehashing. That causes `mysql` to start faster, but you must issue the `rehash` command if you want to use name completion.

To complete a name, enter the first part and press Tab. If the name is unambiguous, `mysql` completes it. Otherwise, you can press Tab again to see the possible names that begin with what you have typed so far. Completion does not occur if there is no default database.

- `--auto-vertical-output` [248]

Cause result sets to be displayed vertically if they are too wide for the current window, and using normal tabular format otherwise. (This applies to statements terminated by `;` or `\G`.)

- `--batch` [248], `-B`

Print results using tab as the column separator, with each row on a new line. With this option, `mysql` does not use the history file.

Batch mode results in nontabular output format and escaping of special characters. Escaping may be disabled by using raw mode; see the description for the `--raw` [252] option.

- `--binary-mode` [248]

This option helps when processing `mysqlbinlog` output that may contain `BLOB` values. By default, `mysql` translates `\r\n` in statement strings to `\n` and interprets `\0` as the statement terminator. `--binary-mode` [248] disables both features. It also disables all `mysql` commands except `charset` and `delimiter` in non-interactive mode (for input piped to `mysql` or loaded using the `source` command).

This option was added in MySQL 5.6.3.

- `--bind-address=ip_address` [248]

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--character-sets-dir=path` [248]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--column-names` [249]

Write column names in results.

- `--column-type-info` [249], `-m`

Display result set metadata.

- `--comments` [249], `-c`

Whether to preserve comments in statements sent to the server. The default is `--skip-comments` (discard comments), enable with `--comments` (preserve comments).

- `--compress` [249], `-C`

Compress all information sent between the client and the server if both support compression.

- `--connect-expired-password` [249]

Indicate to the server that the client can handle sandbox mode if the account used to connect has an expired password. This can be useful for noninteractive invocations of `mysql` because normally the server disconnects noninteractive clients that attempt to connect using an account with an expired password. (See [Section 6.3.6, “Password Expiration and Sandbox Mode”](#).) This option was added in MySQL 5.6.12.

- `--database=db_name` [249], `-D db_name`

The database to use. This is useful primarily in an option file.

- `--debug[=debug_options]` [249], `-# [debug_options]`

Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysql.trace'`.

- `--debug-check` [249]

Print some debugging information when the program exits.

- `--debug-info` [249], `-T`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-auth=plugin` [249]

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--default-character-set=charset_name` [249]

Use `charset_name` as the default character set for the client and connection.

A common issue that can occur when the operating system uses `utf8` or another multi-byte character set is that output from the `mysql` client is formatted incorrectly, due to the fact that the MySQL client uses the `latin1` character set by default. You can usually fix such issues by using this option to force the client to use the system character set instead.

See [Section 10.5, “Character Set Configuration”](#), for more information.

- `--delimiter=str` [249]

Set the statement delimiter. The default is the semicolon character (`;`).

- `--disable-named-commands` [250]

Disable named commands. Use the `\*` form only, or use named commands only at the beginning of a line ending with a semicolon (“;”). `mysql` starts with this option *enabled* by default. However, even with this option, long-format commands still work from the first line. See [Section 4.5.1.2, “mysql Commands”](#).

- `--enable-cleartext-plugin` [250]

Enable the `mysql_clear_password` cleartext authentication plugin. (See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).) This option was added in MySQL 5.6.7.

- `--execute=statement` [250], `-e statement`

Execute the statement and quit. The default output format is like that produced with `--batch` [248]. See [Section 4.2.3.1, “Using Options on the Command Line”](#), for some examples. With this option, `mysql` does not use the history file.

- `--force` [250], `-f`

Continue even if an SQL error occurs.

- `--histignore` [250]

A colon-separated list of one or more patterns specifying statements to ignore for logging purposes. These patterns are added to the default pattern list (“`*IDENTIFIED*:*PASSWORD*`”). The value specified for this option affects logging of statements written to the history file. For more information, see [Section 4.5.1.3, “mysql Logging”](#). This option was added in MySQL 5.6.8.

- `--host=host_name` [250], `-h host_name`

Connect to the MySQL server on the given host.

- `--html` [250], `-H`

Produce HTML output.

- `--ignore-spaces` [250], `-i`

Ignore spaces after function names. The effect of this is described in the discussion for the `IGNORE_SPACE` [604] SQL mode (see [Section 5.1.7, “Server SQL Modes”](#)).

- `--init-command=str` [250]

SQL statement to execute after connecting to the server. If auto-reconnect is enabled, the statement is executed again after reconnection occurs.

- `--line-numbers` [250]

Write line numbers for errors. Disable this with `--skip-line-numbers` [253].

- `--local-infile[={0|1}]` [250]

Enable or disable `LOCAL` capability for `LOAD DATA INFILE`. With no value, the option enables `LOCAL`. The option may be given as `--local-infile=0` [250] or `--local-infile=1` [250] to explicitly disable or enable `LOCAL`. Enabling `LOCAL` has no effect if the server does not also support it.

- `--named-commands` [250], `-G`

Enable named `mysql` commands. Long-format commands are permitted, not just short-format commands. For example, `quit` and `\q` both are recognized. Use `--skip-named-commands` [250] to disable named commands. See [Section 4.5.1.2, “mysql Commands”](#).

- `--no-auto-rehash` [248], `-A`

This has the same effect as `--skip-auto-rehash`. See the description for `--auto-rehash` [248].

- `--no-beep` [251], `-b`

Do not beep when errors occur.

- `--one-database` [251], `-o`

Ignore statements except those that occur while the default database is the one named on the command line. This option is rudimentary and should be used with care. Statement filtering is based only on `USE` statements.

Initially, `mysql` executes statements in the input because specifying a database `db_name` on the command line is equivalent to inserting `USE db_name` at the beginning of the input. Then, for each `USE` statement encountered, `mysql` accepts or rejects following statements depending on whether the database named is the one on the command line. The content of the statements is immaterial.

Suppose that `mysql` is invoked to process this set of statements:

```
DELETE FROM db2.t2;
USE db2;
DROP TABLE db1.t1;
CREATE TABLE db1.t1 (i INT);
USE db1;
INSERT INTO t1 (i) VALUES(1);
CREATE TABLE db2.t1 (j INT);
```

If the command line is `mysql --force --one-database db1`, `mysql` handles the input as follows:

- The `DELETE` statement is executed because the default database is `db1`, even though the statement names a table in a different database.
- The `DROP TABLE` and `CREATE TABLE` statements are not executed because the default database is not `db1`, even though the statements name a table in `db1`.
- The `INSERT` and `CREATE TABLE` statements are executed because the default database is `db1`, even though the `CREATE TABLE` statement names a table in a different database.
- `--pager[=command]` [251]

Use the given command for paging query output. If the command is omitted, the default pager is the value of your `PAGER` environment variable. Valid pagers are `less`, `more`, `cat` [`> filename`], and so forth. This option works only on Unix and only in interactive mode. To disable paging, use `--skip-pager` [251]. Section 4.5.1.2, “`mysql` Commands”, discusses output paging further.

- `--password[=password]` [251], `-p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` [251] or `-p` option on the command line, `mysql` prompts for one.

Specifying a password on the command line should be considered insecure. See Section 6.1.2.1, “`End-User Guidelines for Password Security`”. You can use an option file to avoid giving the password on the command line.

- `--pipe` [251], `-W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path` [252]

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [249] option is used to specify an authentication plugin but `mysql` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--port=port_num` [252], `-P port_num`

The TCP/IP port number to use for the connection.

- `--prompt=format_str` [252]

Set the prompt to the specified format. The default is `mysql>`. The special sequences that the prompt can contain are described in [Section 4.5.1.2, “mysql Commands”](#).

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [252]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--quick` [252], `-q`

Do not cache each query result, print each row as it is received. This may slow down the server if the output is suspended. With this option, `mysql` does not use the history file.

- `--raw` [252], `-r`

For tabular output, the “boxing” around columns enables one column value to be distinguished from another. For nontabular output (such as is produced in batch mode or when the `--batch` [248] or `--silent` [253] option is given), special characters are escaped in the output so they can be identified easily. Newline, tab, `NUL`, and backslash are written as `\n`, `\t`, `\0`, and `\\`. The `--raw` [252] option disables this character escaping.

The following example demonstrates tabular versus nontabular output and the use of raw mode to disable escaping:

```
% mysql
mysql> SELECT CHAR(92);
+-----+
| CHAR(92) |
+-----+
| \       |
+-----+

% mysql -s
mysql> SELECT CHAR(92);
CHAR(92)
\\

% mysql -s -r
mysql> SELECT CHAR(92);
CHAR(92)
\
```

- `--reconnect` [252]

If the connection to the server is lost, automatically try to reconnect. A single reconnect attempt is made each time the connection is lost. To suppress reconnection behavior, use `--skip-reconnect` [252].

- `--safe-updates` [252], `--i-am-a-dummy` [252], `-U`

Permit only those `UPDATE` and `DELETE` statements that specify which rows to modify by using key values. If you have set this option in an option file, you can override it by using `--safe-`

`updates` [252] on the command line. See [Section 4.5.1.6, “mysql Tips”](#), for more information about this option.

- `--secure-auth` [253]

Do not send passwords to the server in old (pre-4.1.1) format. This prevents connections except for servers that use the newer password format. As of MySQL 5.6.7, this option is enabled by default; use `--skip-secure-auth` [253] to disable it.

- `--server-public-key-path=file_name` [253]

The path name to a file containing the server RSA public key. The file must be in PEM format. The public key is used for RSA encryption of the client password for connections to the server made using accounts that authenticate with the `sha256_password` plugin. This option is ignored for client accounts that do not authenticate with that plugin. It is also ignored if password encryption is not needed, as is the case when the client connects to the server using an SSL connection.

The server sends the public key to the client as needed, so it is not necessary to use this option for RSA password encryption to occur. It is more efficient to do so because then the server need not send the key.

For additional discussion regarding use of the `sha256_password` plugin, including how to get the RSA public key, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

This option is available only if MySQL was built using OpenSSL. It was added in MySQL 5.6.6 under the name `--server-public-key` and renamed in 5.6.7 to `--server-public-key-path`.

- `--show-warnings` [253]

Cause warnings to be shown after each statement if there are any. This option applies to interactive and batch mode.

- `--sigint-ignore` [253]

Ignore `SIGINT` signals (typically the result of typing **Control+C**).

- `--silent` [253], `-s`

Silent mode. Produce less output. This option can be given multiple times to produce less and less output.

This option results in nontabular output format and escaping of special characters. Escaping may be disabled by using raw mode; see the description for the `--raw` [252] option.

- `--skip-column-names` [253], `-N`

Do not write column names in results.

- `--skip-line-numbers` [253], `-L`

Do not write line numbers for errors. Useful when you want to compare result files that include error messages.

- `--socket=path` [253], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).



- `--table [254]`, `-t`

Display output in table format. This is the default for interactive use, but can be used to produce table output in batch mode.

- `--tee=file_name [254]`

Append a copy of output to the given file. This option works only in interactive mode. [Section 4.5.1.2, “mysql Commands”](#), discusses tee files further.

- `--unbuffered [254]`, `-n`

Flush the buffer after each query.

- `--user=user_name [254]`, `-u user_name`

The MySQL user name to use when connecting to the server.

- `--verbose [254]`, `-v`

Verbose mode. Produce more output about what the program does. This option can be given multiple times to produce more and more output. (For example, `-v -v -v` produces table output format even in batch mode.)

- `--version [254]`, `-V`

Display version information and exit.

- `--vertical [254]`, `-E`

Print query output rows vertically (one line per column value). Without this option, you can specify vertical output for individual statements by terminating them with `\G`.

- `--wait [254]`, `-w`

If the connection cannot be established, wait and retry instead of aborting.

- `--xml [254]`, `-X`

Produce XML output.

```
<field name="column_name">NULL</field>
```

The output when `--xml [254]` is used with `mysql` matches that of `mysqldump --xml [291]`. See [Section 4.5.4, “mysqldump — A Database Backup Program”](#) for details.

The XML output also uses an XML namespace, as shown here:

```
shell> mysql --xml -uroot -e "SHOW VARIABLES LIKE 'version%'"
<?xml version="1.0"?>

<resultset statement="SHOW VARIABLES LIKE 'version%'" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
<row>
<field name="Variable_name">version</field>
<field name="Value">5.0.40-debug</field>
</row>

<row>
<field name="Variable_name">version_comment</field>
<field name="Value">Source distribution</field>
</row>

<row>
<field name="Variable_name">version_compile_machine</field>
<field name="Value">i686</field>
</row>
```

```
<row>
<field name="Variable_name">version_compile_os</field>
<field name="Value">suse-linux-gnu</field>
</row>
</resultset>
```

(See Bug #25946.)

You can also set the following variables by using `--var_name=value`.

- `connect_timeout`

The number of seconds before connection timeout. (Default value is 0.)

- `max_allowed_packet`

The maximum packet length to send to or receive from the server. (Default value is 16MB.)

- `max_join_size` [505]

The automatic limit for rows in a join when using `--safe-updates` [252]. (Default value is 1,000,000.)

- `net_buffer_length` [515]

The buffer size for TCP/IP and socket communication. (Default value is 16KB.)

- `select_limit`

The automatic limit for `SELECT` statements when using `--safe-updates` [252]. (Default value is 1,000.)

#### 4.5.1.2. `mysql` Commands

`mysql` sends each SQL statement that you issue to the server to be executed. There is also a set of commands that `mysql` itself interprets. For a list of these commands, type `help` or `\h` at the `mysql>` prompt:

```
mysql> help

List of all MySQL commands:
Note that all text commands must be first on line and end with ';'
?          (? ) Synonym for `help`.
clear      (\c) Clear command.
connect    (\r) Reconnect to the server. Optional arguments are db and host.
delimiter (\d) Set statement delimiter.
edit       (\e) Edit command with $EDITOR.
ego        (\G) Send command to mysql server, display result vertically.
exit       (\q) Exit mysql. Same as quit.
go         (\g) Send command to mysql server.
help       (\h) Display this help.
nopager    (\n) Disable pager, print to stdout.
notee     (\t) Don't write into outfile.
pager      (\P) Set PAGER [to_pager]. Print the query results via PAGER.
print      (\p) Print current command.
prompt     (\R) Change your mysql prompt.
quit       (\q) Quit mysql.
rehash     (\#) Rebuild completion hash.
source     (\.) Execute an SQL script file. Takes a file name as an argument.
status     (\s) Get status information from the server.
system     (\!) Execute a system shell command.
tee        (\T) Set outfile [to_outfile]. Append everything into given
         outfile.
use        (\u) Use another database. Takes database name as argument.
charset    (\C) Switch to another charset. Might be needed for processing
         binlog with multi-byte charsets.
warnings   (\W) Show warnings after every statement.
```

```
nowarning (\w) Don't show warnings after every statement.
```

```
For server side help, type 'help contents'
```

If `mysql` is invoked with the `--binary-mode` [248] option, all `mysql` commands are disabled except `charset` and `delimiter` in non-interactive mode (for input piped to `mysql` or loaded using the `source` command).

Each command has both a long and short form. The long form is not case sensitive; the short form is. The long form can be followed by an optional semicolon terminator, but the short form should not.

The use of short-form commands within multi-line `/* ... */` comments is not supported.

- `help [arg], \h [arg], \? [arg], ? [arg]`

Display a help message listing the available `mysql` commands.

If you provide an argument to the `help` command, `mysql` uses it as a search string to access server-side help from the contents of the MySQL Reference Manual. For more information, see [Section 4.5.1.4, “mysql Server-Side Help”](#).

- `charset charset_name, \C charset_name`

Change the default character set and issue a `SET NAMES` statement. This enables the character set to remain synchronized on the client and server if `mysql` is run with auto-reconnect enabled (which is not recommended), because the specified character set is used for reconnects.

- `clear, \c`

Clear the current input. Use this if you change your mind about executing the statement that you are entering.

- `connect [db_name host_name]], \r [db_name host_name]]`

Reconnect to the server. The optional database name and host name arguments may be given to specify the default database or the host where the server is running. If omitted, the current values are used.

- `delimiter str, \d str`

Change the string that `mysql` interprets as the separator between SQL statements. The default is the semicolon character (“;”).

The delimiter string can be specified as an unquoted or quoted argument on the `delimiter` command line. Quoting can be done with either single quote (‘), double quote (“), or backtick (`) characters. To include a quote within a quoted string, either quote the string with a different quote character or escape the quote with a backslash (“\”) character. Backslash should be avoided outside of quoted strings because it is the escape character for MySQL. For an unquoted argument, the delimiter is read up to the first space or end of line. For a quoted argument, the delimiter is read up to the matching quote on the line.

`mysql` interprets instances of the delimiter string as a statement delimiter anywhere it occurs, except within quoted strings. Be careful about defining a delimiter that might occur within other words. For example, if you define the delimiter as `X`, you will be unable to use the word `INDEX` in statements. `mysql` interprets this as `INDE` followed by the delimiter `X`.

When the delimiter recognized by `mysql` is set to something other than the default of “;”, instances of that character are sent to the server without interpretation. However, the server itself still interprets “;” as a statement delimiter and processes statements accordingly. This behavior on the server side comes into play for multiple-statement execution (see [Section 21.8.16, “C API Support for Multiple Statement Execution”](#)), and for parsing the body of stored procedures and functions, triggers, and events (see [Section 18.1, “Defining Stored Programs”](#)).

- `edit, \e`

Edit the current input statement. `mysql` checks the values of the `EDITOR` and `VISUAL` environment variables to determine which editor to use. The default editor is `vi` if neither variable is set.

The `edit` command works only in Unix.

- `ego, \G`

Send the current statement to the server to be executed and display the result using vertical format.

- `exit, \q`

Exit `mysql`.

- `go, \g`

Send the current statement to the server to be executed.

- `nopager, \n`

Disable output paging. See the description for `pager`.

The `nopager` command works only in Unix.

- `notee, \t`

Disable output copying to the tee file. See the description for `tee`.

- `nowarning, \w`

Enable display of warnings after each statement.

- `pager [command], \P [command]`

Enable output paging. By using the `--pager [251]` option when you invoke `mysql`, it is possible to browse or search query results in interactive mode with Unix programs such as `less`, `more`, or any other similar program. If you specify no value for the option, `mysql` checks the value of the `PAGER` environment variable and sets the pager to that. Pager functionality works only in interactive mode.

Output paging can be enabled interactively with the `pager` command and disabled with `nopager`. The command takes an optional argument; if given, the paging program is set to that. With no argument, the pager is set to the pager that was set on the command line, or `stdout` if no pager was specified.

Output paging works only in Unix because it uses the `popen()` function, which does not exist on Windows. For Windows, the `tee` option can be used instead to save query output, although it is not as convenient as `pager` for browsing output in some situations.

- `print, \p`

Print the current input statement without executing it.

- `prompt [str], \R [str]`

Reconfigure the `mysql` prompt to the given string. The special character sequences that can be used in the prompt are described later in this section.

If you specify the `prompt` command with no argument, `mysql` resets the prompt to the default of `mysql>`.

- `quit, \q`

Exit `mysql`.

- `rehash, \#`

Rebuild the completion hash that enables database, table, and column name completion while you are entering statements. (See the description for the `--auto-rehash` [248] option.)

- `source file_name, \. file_name`

Read the named file and executes the statements contained therein. On Windows, you can specify path name separators as `/` or `\\`.

- `status, \s`

Provide status information about the connection and the server you are using. If you are running in `--safe-updates` [252] mode, `status` also prints the values for the `mysql` variables that affect your queries.

- `system command, \! command`

Execute the given command using your default command interpreter.

The `system` command works only in Unix.

- `tee [file_name], \T [file_name]`

By using the `--tee` [254] option when you invoke `mysql`, you can log statements and their output. All the data displayed on the screen is appended into a given file. This can be very useful for debugging purposes also. `mysql` flushes results to the file after each statement, just before it prints its next prompt. Tee functionality works only in interactive mode.

You can enable this feature interactively with the `tee` command. Without a parameter, the previous file is used. The `tee` file can be disabled with the `notee` command. Executing `tee` again re-enables logging.

- `use db_name, \u db_name`

Use `db_name` as the default database.

- `warnings, \W`

Enable display of warnings after each statement (if there are any).

Here are a few tips about the `pager` command:

- You can use it to write to a file and the results go only to the file:

```
mysql> pager cat > /tmp/log.txt
```

You can also pass any options for the program that you want to use as your pager:

```
mysql> pager less -n -i -S
```

- In the preceding example, note the `-S` option. You may find it very useful for browsing wide query results. Sometimes a very wide result set is difficult to read on the screen. The `-S` option to `less` can make the result set much more readable because you can scroll it horizontally using the left-arrow and right-arrow keys. You can also use `-S` interactively within `less` to switch the horizontal-browse mode on and off. For more information, read the `less` manual page:

```
shell> man less
```

- The `-F` and `-X` options may be used with `less` to cause it to exit if output fits on one screen, which is convenient when no scrolling is necessary:

```
mysql> pager less -n -i -S -F -X
```

- You can specify very complex pager commands for handling query output:

```
mysql> pager cat | tee /dr1/tmp/res.txt \  
      | tee /dr2/tmp/res2.txt | less -n -i -S
```

In this example, the command would send query results to two files in two different directories on two different file systems mounted on `/dr1` and `/dr2`, yet still display the results onscreen using `less`.

You can also combine the `tee` and `pager` functions. Have a `tee` file enabled and `pager` set to `less`, and you are able to browse the results using the `less` program and still have everything appended into a file the same time. The difference between the Unix `tee` used with the `pager` command and the `mysql` built-in `tee` command is that the built-in `tee` works even if you do not have the Unix `tee` available. The built-in `tee` also logs everything that is printed on the screen, whereas the Unix `tee` used with `pager` does not log quite that much. Additionally, `tee` file logging can be turned on and off interactively from within `mysql`. This is useful when you want to log some queries to a file, but not others.

The `prompt` command reconfigures the default `mysql>` prompt. The string for defining the prompt can contain the following special sequences.

Option	Description
<code>\c</code>	A counter that increments for each statement you issue
<code>\D</code>	The full current date
<code>\d</code>	The default database
<code>\h</code>	The server host
<code>\l</code>	The current delimiter
<code>\m</code>	Minutes of the current time
<code>\n</code>	A newline character
<code>\O</code>	The current month in three-letter format (Jan, Feb, ...)
<code>\o</code>	The current month in numeric format
<code>\P</code>	am/pm
<code>\p</code>	The current TCP/IP port or socket file
<code>\R</code>	The current time, in 24-hour military time (0–23)
<code>\r</code>	The current time, standard 12-hour time (1–12)
<code>\S</code>	Semicolon
<code>\s</code>	Seconds of the current time
<code>\t</code>	A tab character
<code>\U</code>	Your full <code>user_name@host_name</code> account name
<code>\u</code>	Your user name
<code>\v</code>	The server version
<code>\w</code>	The current day of the week in three-letter format (Mon, Tue, ...)
<code>\Y</code>	The current year, four digits
<code>\y</code>	The current year, two digits
<code>\_</code>	A space
<code>\</code>	A space (a space follows the backslash)
<code>\'</code>	Single quote
<code>\"</code>	Double quote
<code>\\</code>	A literal “\” backslash character
<code>\x</code>	<code>x</code> , for any “ <code>x</code> ” not listed above

You can set the prompt in several ways:

- *Use an environment variable.* You can set the `MYSQL_PS1` environment variable to a prompt string. For example:

```
shell> export MYSQL_PS1="(\\u@\\h) [\\d]> "
```

- *Use a command-line option.* You can set the `--prompt` [252] option on the command line to `mysql`. For example:

```
shell> mysql --prompt="(\\u@\\h) [\\d]> "  
(user@host) [database]>
```

- *Use an option file.* You can set the `prompt` option in the `[mysql]` group of any MySQL option file, such as `/etc/my.cnf` or the `.my.cnf` file in your home directory. For example:

```
[mysql]  
prompt=(\\u@\\h) [\\d]>\\_
```

In this example, note that the backslashes are doubled. If you set the prompt using the `prompt` option in an option file, it is advisable to double the backslashes when using the special prompt options. There is some overlap in the set of permissible prompt options and the set of special escape sequences that are recognized in option files. (The rules for escape sequences in option files are listed in [Section 4.2.3.3, “Using Option Files”](#).) The overlap may cause you problems if you use single backslashes. For example, `\\s` is interpreted as a space rather than as the current seconds value. The following example shows how to define a prompt within an option file to include the current time in `HH:MM:SS>` format:

```
[mysql]  
prompt="\\r:\\m:\\s> "
```

- *Set the prompt interactively.* You can change your prompt interactively by using the `prompt` (or `\\R`) command. For example:

```
mysql> prompt (\\u@\\h) [\\d]>\\_  
PROMPT set to '(\\u@\\h) [\\d]>\\_'  
(user@host) [database]>  
(user@host) [database]> prompt  
Returning to default PROMPT of mysql>  
mysql>
```

### 4.5.1.3. mysql Logging

On Unix, the `mysql` client logs statements executed interactively to a history file. By default, this file is named `.mysql_history` in your home directory. To specify a different file, set the value of the `MYSQL_HISTFILE` environment variable.

#### How Logging Occurs

Statement logging occurs as follows:

- Statements are logged only when executed interactively. Statements are noninteractive, for example, when read from a file or a pipe. It is also possible to suppress statement logging by using the `--batch` [248] or `--execute` [250] option.
- Statements are ignored and not logged if they match any pattern in the “ignore” list. This list is described later.
- `mysql` logs each nonignored, nonempty statement line individually.
- If a nonignored statement spans multiple lines (not including the terminating delimiter), `mysql` concatenates the lines to form the complete statement, maps newlines to spaces, and logs the result, plus a delimiter.

Consequently, an input statement that spans multiple lines can be logged twice. Consider this input:



```
mysql> SELECT
-> 'Today is'
-> ,
-> CONCAT()
-> ;
```

In this case, `mysql` logs the “SELECT”, “Today is”, “,”, “CONCAT()”, and “;” lines as it reads them. It also logs the complete statement, after mapping `SELECT\n'Today is'\n,\nCURDATE()` to `SELECT 'Today is' , CURDATE()`, plus a delimiter. Thus, these lines appear in logged output:

```
SELECT
'Today is'
,
CURDATE()
;
SELECT 'Today is' , CURDATE();
```

As of MySQL 5.6.8, `mysql` ignores for logging purposes statements that match any pattern in the “ignore” list. By default, the pattern list is “`*IDENTIFIED*:*PASSWORD*`”, to ignore statements that refer to passwords. Pattern matching is not case sensitive. Within patterns, two characters are special:

- `?` matches any single character.
- `*` matches any sequence of zero or more characters.

To specify additional patterns, use the `--histignore [250]` option or set the `MYSQL_HISTIGNORE` environment variable. (If both are specified, the option value takes precedence.) The value should be a colon-separated list of one or more patterns, which are appended to the default pattern list.

Patterns specified on the command line might need to be quoted or escaped to prevent your command interpreter from treating them specially. For example, to suppress logging for `UPDATE` and `DELETE` statements in addition to statements that refer to passwords, invoke `mysql` like this:

```
shell> mysql --histignore="*UPDATE*:*DELETE*"
```

## Controlling the History File

The `.mysql_history` file should be protected with a restrictive access mode because sensitive information might be written to it, such as the text of SQL statements that contain passwords. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#).

If you do not want to maintain a history file, first remove `.mysql_history` if it exists. Then use either of the following techniques to prevent it from being created again:

- Set the `MYSQL_HISTFILE` environment variable to `/dev/null`. To cause this setting to take effect each time you log in, put it in one of your shell's startup files.
- Create `.mysql_history` as a symbolic link to `/dev/null`; this need be done only once:

```
shell> ln -s /dev/null $HOME/.mysql_history
```

### 4.5.1.4. mysql Server-Side Help

```
mysql> help search_string
```

If you provide an argument to the `help` command, `mysql` uses it as a search string to access server-side help from the contents of the MySQL Reference Manual. The proper operation of this command requires that the help tables in the `mysql` database be initialized with help topic information (see [Section 5.1.10, “Server-Side Help”](#)).

If there is no match for the search string, the search fails:

```
mysql> help me

Nothing found
Please try to run 'help contents' for a list of all accessible topics
```

Use `help contents` to see a list of the help categories:

```
mysql> help contents
You asked for help about help category: "Contents"
For more information, type 'help <item>', where <item> is one of the
following categories:
  Account Management
  Administration
  Data Definition
  Data Manipulation
  Data Types
  Functions
  Functions and Modifiers for Use with GROUP BY
  Geographic Features
  Language Structure
  Plugins
  Storage Engines
  Stored Routines
  Table Maintenance
  Transactions
  Triggers
```

If the search string matches multiple items, `mysql` shows a list of matching topics:

```
mysql> help logs
Many help items for your request exist.
To make a more specific request, please type 'help <item>',
where <item> is one of the following topics:
  SHOW
  SHOW BINARY LOGS
  SHOW ENGINE
  SHOW LOGS
```

Use a topic as the search string to see the help entry for that topic:

```
mysql> help show binary logs
Name: 'SHOW BINARY LOGS'
Description:
Syntax:
SHOW BINARY LOGS
SHOW MASTER LOGS

Lists the binary log files on the server. This statement is used as
part of the procedure described in [purge-binary-logs], that shows how
to determine which logs can be purged.

mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name          | File_size |
+-----+-----+
| binlog.000015    | 724935   |
| binlog.000016    | 733481   |
+-----+-----+
```

The search string can contain the the wildcard characters “%” and “\_”. These have the same meaning as for pattern-matching operations performed with the [LIKE \[1190\]](#) operator. For example, `HELP rep%` returns a list of topics that begin with `rep`:

```
mysql> HELP rep%
Many help items for your request exist.
To make a more specific request, please type 'help <item>',
where <item> is one of the following
topics:
  REPAIR TABLE
  REPEAT FUNCTION
  REPEAT LOOP
  REPLACE
  REPLACE FUNCTION
```

#### 4.5.1.5. Executing SQL Statements from a Text File

The `mysql` client typically is used interactively, like this:

```
shell> mysql db_name
```

However, it is also possible to put your SQL statements in a file and then tell `mysql` to read its input from that file. To do so, create a text file `text_file` that contains the statements you wish to execute. Then invoke `mysql` as shown here:

```
shell> mysql db_name < text_file
```

If you place a `USE db_name` statement as the first statement in the file, it is unnecessary to specify the database name on the command line:

```
shell> mysql < text_file
```

If you are already running `mysql`, you can execute an SQL script file using the `source` command or `\.` command:

```
mysql> source file_name
mysql> \. file_name
```

Sometimes you may want your script to display progress information to the user. For this you can insert statements like this:

```
SELECT '<info_to_display>' AS ' ';
```

The statement shown outputs `<info_to_display>`.

You can also invoke `mysql` with the `--verbose` [254] option, which causes each statement to be displayed before the result that it produces.

`mysql` ignores Unicode byte order mark (BOM) characters at the beginning of input files. Previously, it read them and sent them to the server, resulting in a syntax error. Presence of a BOM does not cause `mysql` to change its default character set. To do that, invoke `mysql` with an option such as `--default-character-set=utf8` [249].

For more information about batch mode, see [Section 3.5, “Using `mysql` in Batch Mode”](#).

#### 4.5.1.6. `mysql` Tips

This section describes some techniques that can help you use `mysql` more effectively.

##### 4.5.1.6.1. Input-Line Editing

`mysql` supports input-line editing, which enables you to modify the current input line in place or recall previous input lines. For example, the **left-arrow** and **right-arrow** keys move horizontally within the current input line, and the **up-arrow** and **down-arrow** keys move up and down through the set of previously entered lines. **Backspace** deletes the character before the cursor and typing new characters enters them at the cursor position. To enter the line, press **Enter**.

On Windows, the editing key sequences are the same as supported for command editing in console windows. On Unix, the key sequences depend on the input library used to build `mysql` (for example, the `libedit` or `readline` library).

Documentation for the `libedit` and `readline` libraries is available online. To change the set of key sequences permitted by a given input library, define key bindings in the library startup file. This is a file in your home directory: `.editrc` for `libedit` and `.inputrc` for `readline`.

For example, in `libedit`, **Control+W** deletes everything before the current cursor position and **Control+U** deletes the entire line. In `readline`, **Control+W** deletes the word before the cursor and **Control+U** deletes everything before the current cursor position. If `mysql` was built using `libedit`, a user who prefers the `readline` behavior for these two keys can put the following lines in the `.editrc` file (creating the file if necessary):

```
bind "^W" ed-delete-prev-word
```

```
bind "^U" vi-kill-line-prev
```

To see the current set of key bindings, temporarily put a line that says only `bind` at the end of `.editrc`. `mysql` will show the bindings when it starts.

#### 4.5.1.6.2. Unicode Support on Windows

Windows provides APIs based on UTF-16LE for reading from and writing to the console. As of MySQL 5.6.2, the `mysql` client for Windows is able to use these APIs. As of 5.6.3, the Windows installer creates an item in the MySQL menu named `MySQL command line client - Unicode`. This item invokes the `mysql` client with properties set to communicate through the console to the MySQL server using Unicode.

To take advantage of this support manually, run `mysql` within a console that uses a compatible Unicode font and set the default character set to a Unicode character set that is supported for communication with the server:

1. Open a console window.
2. Go to the console window properties, select the font tab, and choose Lucida Console or some other compatible Unicode font. This is necessary because console windows start by default using a DOS raster font that is inadequate for Unicode.
3. Execute `mysql.exe` with the `--default-character-set=utf8` [249] (or `utf8mb4`) option. This option is necessary because `utf16le` is not supported as a connection character set.

With those changes, `mysql` will use the Windows APIs to communicate with the console using UTF-16LE, and communicate with the server using UTF-8. (The menu item mentioned previously sets the font and character set as just described.)

To avoid those steps each time you run `mysql`, you can create a shortcut that invokes `mysql.exe`. The shortcut should set the console font to Lucida Console or some other compatible Unicode font, and pass the `--default-character-set=utf8` [249] (or `utf8mb4`) option to `mysql.exe`.

Alternatively, create a shortcut that only sets the console font, and set the character set in the `[mysql]` group of your `my.ini` file:

```
[mysql]
default-character-set=utf8
```

#### 4.5.1.6.3. Displaying Query Results Vertically

Some query results are much more readable when displayed vertically, instead of in the usual horizontal table format. Queries can be displayed vertically by terminating the query with `\G` instead of a semicolon. For example, longer text values that include newlines often are much easier to read with vertical output:

```
mysql> SELECT * FROM mails WHERE LENGTH(txt) < 300 LIMIT 300,1\G
***** 1. row *****
  msg_nro: 3068
    date: 2000-03-01 23:29:50
time_zone: +0200
mail_from: Monty
   reply: monty@no.spam.com
 mail_to: "Thimble Smith" <tim@no.spam.com>
    sbj: UTF-8
    txt: >>>> "Thimble" == Thimble Smith writes:

Thimble> Hi. I think this is a good idea. Is anyone familiar
Thimble> with UTF-8 or Unicode? Otherwise, I'll put this on my
Thimble> TODO list and see what happens.

Yes, please do that.

Regards,
Monty
```

```
file: inbox-jani-1
hash: 190402944
1 row in set (0.09 sec)
```

#### 4.5.1.6.4. Using the `--safe-updates` [252] Option

For beginners, a useful startup option is `--safe-updates` [252] (or `--i-am-a-dummy` [252], which has the same effect). It is helpful for cases when you might have issued a `DELETE FROM tbl_name` statement but forgotten the `WHERE` clause. Normally, such a statement deletes all rows from the table. With `--safe-updates` [252], you can delete rows only by specifying the key values that identify them. This helps prevent accidents.

When you use the `--safe-updates` [252] option, `mysql` issues the following statement when it connects to the MySQL server:

```
SET sql_safe_updates=1, sql_select_limit=1000, max_join_size=1000000;
```

See [Section 5.1.4, “Server System Variables”](#).

The `SET` statement has the following effects:

- You are not permitted to execute an `UPDATE` or `DELETE` statement unless you specify a key constraint in the `WHERE` clause or provide a `LIMIT` clause (or both). For example:

```
UPDATE tbl_name SET not_key_column=val WHERE key_column=val;

UPDATE tbl_name SET not_key_column=val LIMIT 1;
```

- The server limits all large `SELECT` results to 1,000 rows unless the statement includes a `LIMIT` clause.
- The server aborts multiple-table `SELECT` statements that probably need to examine more than 1,000,000 row combinations.

To specify limits different from 1,000 and 1,000,000, you can override the defaults by using the `--select_limit` and `--max_join_size` options:

```
shell> mysql --safe-updates --select_limit=500 --max_join_size=10000
```

#### 4.5.1.6.5. Disabling `mysql` Auto-Reconnect

If the `mysql` client loses its connection to the server while sending a statement, it immediately and automatically tries to reconnect once to the server and send the statement again. However, even if `mysql` succeeds in reconnecting, your first connection has ended and all your previous session objects and settings are lost: temporary tables, the autocommit mode, and user-defined and session variables. Also, any current transaction rolls back. This behavior may be dangerous for you, as in the following example where the server was shut down and restarted between the first and second statements without you knowing it:

```
mysql> SET @a=1;
Query OK, 0 rows affected (0.05 sec)

mysql> INSERT INTO t VALUES(@a);
ERROR 2006: MySQL server has gone away
No connection. Trying to reconnect...
Connection id: 1
Current database: test

Query OK, 1 row affected (1.30 sec)

mysql> SELECT * FROM t;
+-----+
| a     |
+-----+
| NULL |
+-----+
1 row in set (0.05 sec)
```

The `@a` user variable has been lost with the connection, and after the reconnection it is undefined. If it is important to have `mysql` terminate with an error if the connection has been lost, you can start the `mysql` client with the `--skip-reconnect` [252] option.

For more information about auto-reconnect and its effect on state information when a reconnection occurs, see [Section 21.8.15, “Controlling Automatic Reconnection Behavior”](#).

## 4.5.2. `mysqladmin` — Client for Administering a MySQL Server

`mysqladmin` is a client for performing administrative operations. You can use it to check the server's configuration and current status, to create and drop databases, and more.

Invoke `mysqladmin` like this:

```
shell> mysqladmin [options] command [command-arg] [command [command-arg]] ...
```

`mysqladmin` supports the following commands. Some of the commands take an argument following the command name.

- `create db_name`  
Create a new database named `db_name`.
- `debug`  
Tell the server to write debug information to the error log. Format and content of this information is subject to change.  
  
This includes information about the Event Scheduler. See [Section 18.4.5, “Event Scheduler Status”](#).
- `drop db_name`  
Delete the database named `db_name` and all its tables.
- `extended-status`  
Display the server status variables and their values.
- `flush-hosts`  
Flush all information in the host cache.
- `flush-logs`  
Flush all logs.
- `flush-privileges`  
Reload the grant tables (same as `reload`).
- `flush-status`  
Clear status variables.
- `flush-tables`  
Flush all tables.
- `flush-threads`  
Flush the thread cache.
- `kill id,id,...`  
Kill server threads. If multiple thread ID values are given, there must be no spaces in the list.

- `old-password new-password`

This is like the `password` command but stores the password using the old (pre-4.1) password-hashing format. (See [Section 6.1.2.4, “Password Hashing in MySQL”](#).)

- `password new-password`

Set a new password. This changes the password to `new-password` for the account that you use with `mysqladmin` for connecting to the server. Thus, the next time you invoke `mysqladmin` (or any other client program) using the same account, you will need to specify the new password.

If the `new-password` value contains spaces or other characters that are special to your command interpreter, you need to enclose it within quotation marks. On Windows, be sure to use double quotation marks rather than single quotation marks; single quotation marks are not stripped from the password, but rather are interpreted as part of the password. For example:

```
shell> mysqladmin password "my new password"
```

In MySQL 5.6, the new password can be omitted following the `password` command. In this case, `mysqladmin` prompts for the password value, which enables you to avoid specifying the password on the command line. Omitting the password value should be done only if `password` is the final command on the `mysqladmin` command line. Otherwise, the next argument is taken as the password.



#### Caution

Do not use this command used if the server was started with the `--skip-grant-tables` [439] option. No password change will be applied. This is true even if you precede the `password` command with `flush-privileges` on the same command line to re-enable the grant tables because the flush operation occurs after you connect. However, you can use `mysqladmin flush-privileges` to re-enable the grant table and then use a separate `mysqladmin password` command to change the password.

- `ping`

Check whether the server is available. The return status from `mysqladmin` is 0 if the server is running, 1 if it is not. This is 0 even in case of an error such as `Access denied`, because this means that the server is running but refused the connection, which is different from the server not running.

- `processlist`

Show a list of active server threads. This is like the output of the `SHOW PROCESSLIST` statement. If the `--verbose` [272] option is given, the output is like that of `SHOW FULL PROCESSLIST`. (See [Section 13.7.5.30, “SHOW PROCESSLIST Syntax”](#).)

- `reload`

Reload the grant tables.

- `refresh`

Flush all tables and close and open log files.

- `shutdown`

Stop the server.

- `start-slave`

Start replication on a slave server.



- `status`  
Display a short server status message.
- `stop-slave`  
Stop replication on a slave server.
- `variables`  
Display the server system variables and their values.
- `version`  
Display version information from the server.

All commands can be shortened to any unique prefix. For example:

```
shell> mysqladmin proc stat
+-----+-----+-----+-----+-----+-----+-----+
| Id | User | Host | db | Command | Time | State | Info |
+-----+-----+-----+-----+-----+-----+-----+
| 51 | monty | localhost | | Query | 0 | | show processlist |
+-----+-----+-----+-----+-----+-----+-----+
Uptime: 1473624  Threads: 1  Questions: 39487
Slow queries: 0  Opens: 541  Flush tables: 1
Open tables: 19  Queries per second avg: 0.0268
```

The `mysqladmin status` command result displays the following values:

- `Uptime [602]`  
The number of seconds the MySQL server has been running.
- `Threads`  
The number of active threads (clients).
- `Questions [596]`  
The number of questions (queries) from clients since the server was started.
- `Slow queries`  
The number of queries that have taken more than `long_query_time [500]` seconds. See [Section 5.2.5, “The Slow Query Log”](#).
- `Opens`  
The number of tables the server has opened.
- `Flush tables`  
The number of `flush-*`, `refresh`, and `reload` commands the server has executed.
- `Open tables`  
The number of tables that currently are open.

If you execute `mysqladmin shutdown` when connecting to a local server using a Unix socket file, `mysqladmin` waits until the server's process ID file has been removed, to ensure that the server has stopped properly.

`mysqladmin` supports the following options, which can be specified on the command line or in the `[mysqladmin]` and `[client]` groups of an option file. `mysqladmin` also supports the options for

processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.4. mysqladmin Options**

Format	Option File	Description	Introduced
<code>--bind-address=ip_address</code> [270]	<code>bind-address</code> [270]	Use the specified network interface to connect to the MySQL Server	
<code>--compress</code> [270]	<code>compress</code> [270]	Compress all information sent between the client and the server	
<code>--connect_timeout=seconds</code> [272]	<code>connect_timeout</code> [272]	The number of seconds before connection timeout	
<code>--count=#</code> [270]	<code>count</code> [270]	The number of iterations to make for repeated command execution	
<code>--debug[=debug_options]</code> [271]	<code>debug</code> [271]	Write a debugging log	
<code>--debug-check</code> [271]	<code>debug-check</code> [271]	Print debugging information when the program exits	
<code>--debug-info</code> [271]	<code>debug-info</code> [271]	Print debugging information, memory and CPU statistics when the program exits	
<code>--default-auth=plugin</code> [271]	<code>default-auth=plugin</code> [271]	The authentication plugin to use	
<code>--default-character-set=charset_name</code>	<code>default-character-set</code> [271]	Use <code>charset_name</code> as the default character set	
<code>--enable-cleartext-plugin</code> [271]	<code>enable-cleartext-plugin</code> [271]	Enable cleartext authentication plugin	5.6.7
<code>--force</code> [271]	<code>force</code> [271]	Continue even if an SQL error occurs	
<code>--help</code> [270]		Display help message and exit	
<code>--host=host_name</code> [271]	<code>host</code> [271]	Connect to the MySQL server on the given host	
<code>--login-path=name</code> [219]		Read login path options from <code>.mylogin.cnf</code>	5.6.6
<code>--no-beep</code> [271]	<code>no-beep</code> [271]	Do not beep when errors occur	
<code>--password[=password]</code> [271]	<code>password</code> [271]	The password to use when connecting to the server	
<code>--pipe</code> [271]		On Windows, connect to server using a named pipe	
<code>--plugin-dir=path</code> [271]	<code>plugin-dir=path</code> [271]	The directory where plugins are located	
<code>--port=port_num</code> [271]	<code>port</code> [271]	The TCP/IP port number to use for the connection	
<code>--protocol=type</code> [272]	<code>protocol</code> [272]	The connection protocol to use	
<code>--relative</code> [272]	<code>relative</code> [272]	Show the difference between the current and previous values when used with the <code>--sleep</code> option	
<code>--shutdown_timeout=seconds</code> [272]	<code>shutdown_timeout</code> [272]	The maximum number of seconds to wait for server shutdown	
<code>--silent</code> [272]	<code>silent</code> [272]	Silent mode	

Format	Option File	Description	Introduced
-- sleep=delay [272]	sleep [272]	Execute commands repeatedly, sleeping for delay seconds in between	
-- socket=path [272]	socket [272]	For connections to localhost	
--ssl- ca=file_name [272]	ssl-ca [272]	The path to a file that contains a list of trusted SSL CAs	
--ssl- capath=dir_name [272]	ssl-capath [272]	The path to a directory that contains trusted SSL CA certificates in PEM format	
--ssl- cert=file_name [272]	ssl-cert [272]	The name of the SSL certificate file to use for establishing a secure connection	
--ssl- cipher=cipher_list [272]	ssl-cipher [272]	A list of allowable ciphers to use for SSL encryption	
--ssl- crl=file_name [272]	ssl-crl [272]	The path to a file that contains certificate revocation lists	5.6.3
--ssl- crlpath=dir_name [272]	ssl-crlpath [272]	The path to a directory that contains certificate revocation list files	5.6.3
--ssl- key=file_name [272]	ssl-key [272]	The name of the SSL key file to use for establishing a secure connection	
--ssl-verify-server- cert [272]	ssl-verify-server- cert [272]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
-- user=user_name, [272]	user [272]	MySQL user name to use when connecting to server	
--verbose [272]		Verbose mode	
--version [272]		Display version information and exit	
--vertical [272]	vertical [272]	Print query output rows vertically (one line per column value)	
--wait [272]	wait [272]	If the connection cannot be established, wait and retry instead of aborting	

- `--help [270]`, `-?`

Display a help message and exit.

- `--bind-address=ip_address [248]`

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--character-sets-dir=path [270]`

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--compress [270]`, `-C`

Compress all information sent between the client and the server if both support compression.

- `--count=N [270]`, `-c N`

The number of iterations to make for repeated command execution if the `--sleep [272]` option is given.

- `--debug[=debug_options] [271], -# [debug_options]`

Write a debugging log. A typical *debug\_options* string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysqladmin.trace'`.

- `--debug-check [271]`

Print some debugging information when the program exits.

- `--debug-info [271]`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-auth=plugin [271]`

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--default-character-set=charset_name [271]`

Use *charset\_name* as the default character set. See [Section 10.5, “Character Set Configuration”](#).

- `--enable-cleartext-plugin [271]`

Enable the `mysql_clear_password` cleartext authentication plugin. (See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).) This option was added in MySQL 5.6.7.

- `--force [271], -f`

Do not ask for confirmation for the `drop db_name` command. With multiple commands, continue even if an error occurs.

- `--host=host_name [271], -h host_name`

Connect to the MySQL server on the given host.

- `--no-beep [271], -b`

Suppress the warning beep that is emitted by default for errors such as a failure to connect to the server.

- `--password[=password] [271], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the *password* value following the `--password [271]` or `-p` option on the command line, `mysqladmin` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe [271], -W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path [271]`

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth [271]` option is used to specify an authentication plugin but `mysqladmin` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--port=port_num [271], -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [272]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--relative` [272], `-r`

Show the difference between the current and previous values when used with the `--sleep` [272] option. This option works only with the `extended-status` command.

- `--silent` [272], `-s`

Exit silently if a connection to the server cannot be established.

- `--sleep=delay` [272], `-i delay`

Execute commands repeatedly, sleeping for *delay* seconds in between. The `--count` [270] option determines the number of iterations. If `--count` [270] is not given, `mysqladmin` executes commands indefinitely until interrupted.

- `--socket=path` [272], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--user=user_name` [272], `-u user_name`

The MySQL user name to use when connecting to the server.

- `--verbose` [272], `-v`

Verbose mode. Print more information about what the program does.

- `--version` [272], `-V`

Display version information and exit.

- `--vertical` [272], `-E`

Print output vertically. This is similar to `--relative` [272], but prints output vertically.

- `--wait[=count]` [272], `-w[count]`

If the connection cannot be established, wait and retry instead of aborting. If a *count* value is given, it indicates the number of times to retry. The default is one time.

You can also set the following variables by using `--var_name=value`.

- `connect_timeout`

The maximum number of seconds before connection timeout. The default value is 43200 (12 hours).

- `shutdown_timeout`

The maximum number of seconds to wait for server shutdown. The default value is 3600 (1 hour).

### 4.5.3. mysqlcheck — A Table Maintenance Program

The `mysqlcheck` client performs table maintenance: It checks, repairs, optimizes, or analyzes tables.

Each table is locked and therefore unavailable to other sessions while it is being processed, although for check operations, the table is locked with a `READ` lock only (see [Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”](#), for more information about `READ` and `WRITE` locks). Table maintenance operations can be time-consuming, particularly for large tables. If you use the `--databases` [276] or `--all-databases` [276] option to process all tables in one or more databases, an invocation of `mysqlcheck` might take a long time. (This is also true for `mysql_upgrade` because that program invokes `mysqlcheck` to check all tables and repair them if necessary.)

`mysqlcheck` is similar in function to `myisamchk`, but works differently. The main operational difference is that `mysqlcheck` must be used when the `mysqld` server is running, whereas `myisamchk` should be used when it is not. The benefit of using `mysqlcheck` is that you do not have to stop the server to perform table maintenance.

`mysqlcheck` uses the SQL statements `CHECK TABLE`, `REPAIR TABLE`, `ANALYZE TABLE`, and `OPTIMIZE TABLE` in a convenient way for the user. It determines which statements to use for the operation you want to perform, and then sends the statements to the server to be executed. For details about which storage engines each statement works with, see the descriptions for those statements in [Section 13.7.2, “Table Maintenance Statements”](#).

The `MyISAM` storage engine supports all four maintenance operations, so `mysqlcheck` can be used to perform any of them on `MyISAM` tables. Other storage engines do not necessarily support all operations. In such cases, an error message is displayed. For example, if `test.t` is a `MEMORY` table, an attempt to check it produces this result:

```
shell> mysqlcheck test t
test.t
note      : The storage engine for the table doesn't support check
```

If `mysqlcheck` is unable to repair a table, see [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#) for manual table repair strategies. This will be the case, for example, for `InnoDB` tables, which can be checked with `CHECK TABLE`, but not repaired with `REPAIR TABLE`.



### Caution

It is best to make a backup of a table before performing a table repair operation; under some circumstances the operation might cause data loss. Possible causes include but are not limited to file system errors.

There are three general ways to invoke `mysqlcheck`:

```
shell> mysqlcheck [options] db_name [tbl_name ...]
shell> mysqlcheck [options] --databases db_name ...
shell> mysqlcheck [options] --all-databases
```

If you do not name any tables following `db_name` or if you use the `--databases` [276] or `--all-databases` [276] option, entire databases are checked.

`mysqlcheck` has a special feature compared to other client programs. The default behavior of checking tables (`--check` [276]) can be changed by renaming the binary. If you want to have a tool that repairs tables by default, you should just make a copy of `mysqlcheck` named `mysqlrepair`, or make a symbolic link to `mysqlcheck` named `mysqlrepair`. If you invoke `mysqlrepair`, it repairs tables.

The names shown in the following table can be used to change `mysqlcheck` default behavior.

Command	Meaning
<code>mysqlrepair</code>	The default option is <code>--repair</code> [278]
<code>mysqlanalyze</code>	The default option is <code>--analyze</code> [276]
<code>mysqloptimize</code>	The default option is <code>--optimize</code> [277]

mysqlcheck supports the following options, which can be specified on the command line or in the [mysqlcheck] and [client] groups of an option file. mysqlcheck also supports the options for processing option files described at Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”.

**Table 4.5. mysqlcheck Options**

Format	Option File	Description	Introduced
--all-databases [276]	all-databases [276]	Check all tables in all databases	
--all-in-1 [276]	all-in-1 [276]	Execute a single statement for each database that names all the tables from that database	
--analyze [276]	analyze [276]	Analyze the tables	
--auto-repair [276]	auto-repair [276]	If a checked table is corrupted, automatically fix it	
--bind-address=ip_address [276]	bind-address [276]	Use the specified network interface to connect to the MySQL Server	
--character-sets-dir=path [276]	character-sets-dir [276]	The directory where character sets are installed	
--check [276]	check [276]	Check the tables for errors	
--check-only-changed [276]	check-only-changed [276]	Check only tables that have changed since the last check	
--check-upgrade [276]	check-upgrade [276]	Invoke CHECK TABLE with the FOR UPGRADE option	
--compress [276]	compress [276]	Compress all information sent between the client and the server	
--databases [276]	databases [276]	Process all tables in the named databases	
--debug [277]	debug [277]	Write a debugging log	
--debug-check [277]	debug-check [277]	Print debugging information when the program exits	
--debug-info [277]	debug-info [277]	Print debugging information, memory and CPU statistics when the program exits	
--default-auth=plugin [277]	default-auth=plugin [277]	The authentication plugin to use	5.6.2
--default-character-set=charset_name [277]	default-character-set [277]	Use charset_name as the default character set	
--extended [277]	extended [277]	Check and repair tables	
--fast [277]	fast [277]	Check only tables that have not been closed properly	
--fix-db-names [277]	fix-db-names [277]	Convert database names to 5.1 format	
--fix-table-names [277]	fix-table-names [277]	Convert table names to 5.1 format	
--force [277]	force [277]	Continue even if an SQL error occurs	
--help [276]		Display help message and exit	
--host=host_name [277]	host [277]	Connect to the MySQL server on the given host	



Format	Option File	Description	Introduced
<a href="#">--login-path=name</a> [219]		Read login path options from .mylogin.cnf	5.6.6
<a href="#">--medium-check</a> [277]	<a href="#">medium-check</a> [277]	Do a check that is faster than an --extended operation	
<a href="#">--optimize</a> [277]	<a href="#">optimize</a> [277]	Optimize the tables	
<a href="#">--password[=password]</a> [277]	<a href="#">password</a> [277]	The password to use when connecting to the server	
<a href="#">--pipe</a> [278]		On Windows, connect to server using a named pipe	
<a href="#">--plugin-dir=path</a> [278]	<a href="#">plugin-dir=path</a> [278]	The directory where plugins are located	5.6.2
<a href="#">--port=port_num</a> [278]	<a href="#">port</a> [278]	The TCP/IP port number to use for the connection	
<a href="#">--protocol=type</a> [278]	<a href="#">protocol</a> [278]	The connection protocol to use	
<a href="#">--quick</a> [278]	<a href="#">quick</a> [278]	The fastest method of checking	
<a href="#">--repair</a> [278]	<a href="#">repair</a> [278]	Perform a repair that can fix almost anything except unique keys that are not unique	
<a href="#">--silent</a> [278]	<a href="#">silent</a> [278]	Silent mode	
<a href="#">--skip-database=db_name</a> [278]	<a href="#">skip-database</a> [278]	Omit this database from performed operations	5.6.11
<a href="#">--socket=path</a> [278]	<a href="#">socket</a> [278]	For connections to localhost	
<a href="#">--ssl-ca=file_name</a> [278]	<a href="#">ssl-ca</a> [278]	The path to a file that contains a list of trusted SSL CAs	
<a href="#">--ssl-capath=dir_name</a> [278]	<a href="#">ssl-capath</a> [278]	The path to a directory that contains trusted SSL CA certificates in PEM format	
<a href="#">--ssl-cert=file_name</a> [278]	<a href="#">ssl-cert</a> [278]	The name of the SSL certificate file to use for establishing a secure connection	
<a href="#">--ssl-cipher=cipher_list</a> [278]	<a href="#">ssl-cipher</a> [278]	A list of allowable ciphers to use for SSL encryption	
<a href="#">--ssl-crl=file_name</a> [278]	<a href="#">ssl-crl</a> [278]	The path to a file that contains certificate revocation lists	5.6.3
<a href="#">--ssl-crlpath=dir_name</a> [278]	<a href="#">ssl-crlpath</a> [278]	The path to a directory that contains certificate revocation list files	5.6.3
<a href="#">--ssl-key=file_name</a> [278]	<a href="#">ssl-key</a> [278]	The name of the SSL key file to use for establishing a secure connection	
<a href="#">--ssl-verify-server-cert</a> [278]	<a href="#">ssl-verify-server-cert</a> [278]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<a href="#">--tables</a> [278]	<a href="#">tables</a> [278]	Overrides the --databases or -B option	
<a href="#">--use-frm</a> [279]	<a href="#">use-frm</a> [279]	For repair operations on MyISAM tables	
<a href="#">--user=user_name,</a> [279]	<a href="#">user</a> [279]	MySQL user name to use when connecting to server	
<a href="#">--verbose</a> [279]		Verbose mode	
<a href="#">--version</a> [279]		Display version information and exit	

Format	Option File	Description	Introduced
<code>--write-binlog</code> [279]	<code>write-binlog</code> [279]	Log ANALYZE, OPTIMIZE, REPAIR statements to binary log. <code>--skip-write-binlog</code> adds NO_WRITE_TO_BINLOG to these statements.	

- `--help` [276], `-?`

Display a help message and exit.

- `--all-databases` [276], `-A`

Check all tables in all databases. This is the same as using the `--databases` [276] option and naming all the databases on the command line.

- `--all-in-1` [276], `-1`

Instead of issuing a statement for each table, execute a single statement for each database that names all the tables from that database to be processed.

- `--analyze` [276], `-a`

Analyze the tables.

- `--auto-repair` [276]

If a checked table is corrupted, automatically fix it. Any necessary repairs are done after all tables have been checked.

- `--bind-address=ip_address`

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--character-sets-dir=path` [276]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--check` [276], `-c`

Check the tables for errors. This is the default operation.

- `--check-only-changed` [276], `-C`

Check only tables that have changed since the last check or that have not been closed properly.

- `--check-upgrade` [276], `-g`

Invoke `CHECK TABLE` with the `FOR UPGRADE` option to check tables for incompatibilities with the current version of the server. This option automatically enables the `--fix-db-names` [277] and `--fix-table-names` [277] options.

- `--compress` [276]

Compress all information sent between the client and the server if both support compression.

- `--databases` [276], `-B`

Process all tables in the named databases. Normally, `mysqlcheck` treats the first name argument on the command line as a database name and following names as table names. With this option, it treats all name arguments as database names.

- `--debug[=debug_options] [277], -# [debug_options]`

Write a debugging log. A typical *debug\_options* string is '*d:t:o,file\_name*'. The default is '*d:t:o*'.

- `--debug-check [277]`

Print some debugging information when the program exits.

- `--debug-info [277]`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-character-set=charset_name [277]`

Use *charset\_name* as the default character set. See [Section 10.5, “Character Set Configuration”](#).

- `--extended [277], -e`

If you are using this option to check tables, it ensures that they are 100% consistent but takes a long time.

If you are using this option to repair tables, it runs an extended repair that may not only take a long time to execute, but may produce a lot of garbage rows also!

- `--default-auth=plugin [277]`

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--fast [277], -F`

Check only tables that have not been closed properly.

- `--fix-db-names [277]`

Convert database names to 5.1 format. Only database names that contain special characters are affected.

- `--fix-table-names [277]`

Convert table names to 5.1 format. Only table names that contain special characters are affected. This option also applies to views.

- `--force [277], -f`

Continue even if an SQL error occurs.

- `--host=host_name [277], -h host_name`

Connect to the MySQL server on the given host.

- `--medium-check [277], -m`

Do a check that is faster than an `--extended [277]` operation. This finds only 99.99% of all errors, which should be good enough in most cases.

- `--optimize [277], -o`

Optimize the tables.

- `--password[=password] [277], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password [277]` or `-p` option on the command line, `mysqlcheck` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe [278], -W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path [278]`

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth [277]` option is used to specify an authentication plugin but `mysqlcheck` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--port=port_num [278], -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY} [278]`

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--quick [278], -q`

If you are using this option to check tables, it prevents the check from scanning the rows to check for incorrect links. This is the fastest check method.

If you are using this option to repair tables, it tries to repair only the index tree. This is the fastest repair method.

- `--repair [278], -r`

Perform a repair that can fix almost anything except unique keys that are not unique.

- `--silent [278], -s`

Silent mode. Print only error messages.

- `--skip-database=db_name [278]`

Do not include the named database (case sensitive) in the operations performed by `mysqlcheck`. This option was added in MySQL 5.6.11.

- `--socket=path [278], -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--tables [278]`

Override the `--databases` [276] or `-B` option. All name arguments following the option are regarded as table names.

- `--use-frn` [279]

For repair operations on `MyISAM` tables, get the table structure from the `.frm` file so that the table can be repaired even if the `.MYI` header is corrupted.

- `--user=user_name` [279], `-u user_name`

The MySQL user name to use when connecting to the server.

- `--verbose` [279], `-v`

Verbose mode. Print information about the various stages of program operation.

- `--version` [279], `-V`

Display version information and exit.

- `--write-binlog` [279]

This option is enabled by default, so that `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements generated by `mysqlcheck` are written to the binary log. Use `--skip-write-binlog` [279] to cause `NO_WRITE_TO_BINLOG` to be added to the statements so that they are not logged. Use the `--skip-write-binlog` [279] when these statements should not be sent to replication slaves or run when using the binary logs for recovery from backup.

#### 4.5.4. `mysqldump` — A Database Backup Program

The `mysqldump` client is a utility that performs **logical backups**, producing a set of SQL statements that can be run to reproduce the original schema objects, table data, or both. It dumps one or more MySQL database for backup or transfer to another SQL server. The `mysqldump` command can also generate output in CSV, other delimited text, or XML format.

`mysqldump` requires at least the `SELECT` [755] privilege for dumped tables, `SHOW VIEW` [756] for dumped views, `TRIGGER` [756] for dumped triggers, and `LOCK TABLES` [755] if the `--single-transaction` [295] option is not used. Certain options might require other privileges as noted in the option descriptions.

To reload a dump file, you must have the same privileges needed to create each of the dumped objects by issuing `CREATE` statements manually.

### Performance and Scalability Considerations

`mysqldump` advantages include the convenience and flexibility of viewing or even editing the output before restoring. You can clone databases for development and DBA work, or produce slight variations of an existing database for testing. It is not intended as a fast or scalable solution for backing up substantial amounts of data. With large data sizes, even if the backup step takes a reasonable time, restoring the data can be very slow because replaying the SQL statements involves disk I/O for insertion, index creation, and so on.

For large-scale backup and restore, a **physical** backup is more appropriate, to copy the data files in their original format that can be restored quickly:

- If your tables are primarily `InnoDB` tables, or if you have a mix of `InnoDB` and `MyISAM` tables, consider using the `mysqlbackup` command of the MySQL Enterprise Backup product. (Available as part of the Enterprise subscription.) It provides the best performance for `InnoDB` backups with minimal disruption; it can also back up tables from `MyISAM` and other storage engines; and it provides a number of convenient options to accommodate different backup scenarios. See [MySQL Enterprise Backup User's Guide \(Version 3.8.1\)](#).

- If your tables are primarily MyISAM tables, consider using the `mysqlhotcopy` instead, for better performance than `mysqldump` of backup and restore operations. See [Section 4.6.10](#), “`mysqlhotcopy` — A Database Backup Program”.

`mysqldump` can retrieve and dump table contents row by row, or it can retrieve the entire content from a table and buffer it in memory before dumping it. Buffering in memory can be a problem if you are dumping large tables. To dump tables row by row, use the `--quick` [294] option (or `--opt` [294], which enables `--quick` [294]). The `--opt` [294] option (and hence `--quick` [294]) is enabled by default, so to enable memory buffering, use `--skip-quick` [294].

If you are using a recent version of `mysqldump` to generate a dump to be reloaded into a very old MySQL server, use the `--skip-opt` [294] option instead of the `--opt` [294] or `--extended-insert` [294] option.

For additional information about `mysqldump`, see [Section 7.4](#), “Using `mysqldump` for Backups”.

## Syntax

There are three general ways to invoke `mysqldump`, to dump progressively more data:

```
shell> mysqldump [options] db_name [tbl_name ...]
shell> mysqldump [options] --databases db_name ...
shell> mysqldump [options] --all-databases
```

To dump entire databases, do not name any tables following `db_name`, or use the `--databases` [292] or `--all-databases` [292] option.

To see a list of the options your version of `mysqldump` supports, issue the command `mysqldump --help`.

## Option Syntax - Alphabetical Summary

`mysqldump` supports the following options, which can be specified on the command line or in the `[mysqldump]` and `[client]` groups of an option file. `mysqldump` also supports the options for processing option files described at [Section 4.2.3.4](#), “Command-Line Options that Affect Option-File Handling”.

**Table 4.6.** `mysqldump` Options

Format	Option File	Description	Introduced
<code>--add-drop-database</code> [286]	<code>add-drop-database</code> [286]	Add a DROP DATABASE statement before each CREATE DATABASE statement	
<code>--add-drop-table</code> [286]	<code>add-drop-table</code> [286]	Add a DROP TABLE statement before each CREATE TABLE statement	
<code>--add-drop-trigger</code> [286]	<code>add-drop-trigger</code> [286]	Add a DROP TRIGGER statement before each CREATE TRIGGER statement	
<code>--add-locks</code> [294]	<code>add-locks</code> [294]	Surround each table dump with LOCK TABLES and UNLOCK TABLES statements	
<code>--all-databases</code> [292]	<code>all-databases</code> [292]	Dump all tables in all databases	
<code>--allow-keywords</code> [286]	<code>allow-keywords</code> [286]	Allow creation of column names that are keywords	
<code>--apply-slave-statements</code> [288]	<code>apply-slave-statements</code> [288]	Include STOP SLAVE prior to CHANGE MASTER statement and START SLAVE at end of output	
<code>--bind-address=ip_address</code> [284]	<code>bind-address</code> [284]	Use the specified network interface to connect to the MySQL Server	
<code>--comments</code> [286]	<code>comments</code> [286]	Add comments to the dump file	

Format	Option File	Description	Introduced
<code>--compact</code> [290]	<code>compact</code> [290]	Produce more compact output	
<code>--compatible=name[,name,...]</code> [290]	<code>compatible</code> [290] <code>name,...</code> [290]	Produce output that is more compatible with other database systems or with older MySQL servers	
<code>--complete-insert</code> [290]	<code>complete-insert</code> [290]	Use complete INSERT statements that include column names	
<code>--create-options</code> [290]	<code>create-options</code> [290]	Include all MySQL-specific table options in CREATE TABLE statements	
<code>--databases</code> [292]	<code>databases</code> [292]	Dump several databases	
<code>--debug[=debug_options]</code> [286]	<code>debug</code> [286]	Write a debugging log	
<code>--debug-check</code> [287]	<code>debug-check</code> [287]	Print debugging information when the program exits	
<code>--debug-info</code> [287]	<code>debug-info</code> [287]	Print debugging information, memory and CPU statistics when the program exits	
<code>--default-auth=plugin</code> [284]	<code>default-auth=plugin</code> [284]	The authentication plugin to use	
<code>--default-character-set=charset_name</code> [288]	<code>default-character-set</code> [288] <code>charset_name</code> [288]	Use <code>charset_name</code> as the default character set	
<code>--delayed-insert</code> [294]	<code>delayed-insert</code> [294]	Write INSERT DELAYED statements rather than INSERT statements	
<code>--delete-master-logs</code> [288]	<code>delete-master-logs</code> [288]	On a master replication server, delete the binary logs after performing the dump operation	
<code>--disable-keys</code> [294]	<code>disable-keys</code> [294]	For each table, surround the INSERT statements with statements to disable and enable keys	
<code>--dump-date</code> [287]	<code>dump-date</code> [287]	Include dump date as "Dump completed on" comment if <code>--comments</code> is given	
<code>--dump-slave[=value]</code> [288]	<code>dump-slave</code> [288]	Include CHANGE MASTER statement that lists binary log coordinates of slave's master	
<code>--events</code> [293]	<code>events</code> [293]	Dump events from the dumped databases	
<code>--extended-insert</code> [294]	<code>extended-insert</code> [294]	Use multiple-row INSERT syntax that include several VALUES lists	
<code>--fields-enclosed-by=string</code> [290]	<code>fields-enclosed-by</code> [290]	This option is used with the <code>--tab</code> option and has the same meaning as the corresponding clause for LOAD DATA INFILE	
<code>--fields-escaped-by</code> [290]	<code>fields-escaped-by</code> [290]	This option is used with the <code>--tab</code> option and has the same meaning as the corresponding clause for LOAD DATA INFILE	
<code>--fields-optionally-enclosed-by=string</code> [290]	<code>fields-optionally-enclosed-by</code> [290]	This option is used with the <code>--tab</code> option and has the same meaning as the corresponding clause for LOAD DATA INFILE	
<code>--fields-terminated-by=string</code> [290]	<code>fields-terminated-by</code> [290]	This option is used with the <code>--tab</code> option and has the same meaning as the corresponding clause for LOAD DATA INFILE	
<code>--flush-logs</code> [294]	<code>flush-logs</code> [294]	Flush the MySQL server log files before starting the dump	



Format	Option File	Description	Introduced
<a href="#">--flush-privileges [295]</a>	<a href="#">flush-privileges [295]</a>	Emit a FLUSH PRIVILEGES statement after dumping the mysql database	
<a href="#">--help [287]</a>		Display help message and exit	
<a href="#">--hex-blob [290]</a>	<a href="#">hex-blob [290]</a>	Dump binary columns using hexadecimal notation (for example, 'abc' becomes 0x616263)	
<a href="#">--host [284]</a>	<a href="#">host [284]</a>	Host to connect to (IP address or hostname)	
<a href="#">--ignore-table=db_name.tbl [293]</a>	<a href="#">ignore-table [293]</a>	Do not dump the given table	
<a href="#">--include-master-host-port [288]</a>	<a href="#">include-master-host-port [288]</a>	Include MASTER_HOST/MASTER_PORT options in CHANGE MASTER statement produced with --dump-slave	
<a href="#">--insert-ignore [294]</a>	<a href="#">insert-ignore [294]</a>	Write INSERT IGNORE statements rather than INSERT statements	
<a href="#">--lines-terminated-by=string [290]</a>	<a href="#">lines-terminated-by [290]</a>	This option is used with the --tab option and has the same meaning as the corresponding clause for LOAD DATA INFILE	
<a href="#">--lock-all-tables [295]</a>	<a href="#">lock-all-tables [295]</a>	Lock all tables across all databases	
<a href="#">--lock-tables [295]</a>	<a href="#">lock-tables [295]</a>	Lock all tables before dumping them	
<a href="#">--log-error=file_name [287]</a>	<a href="#">log-error [287]</a>	Append warnings and errors to the named file	
<a href="#">--login-path=name [219]</a>		Read login path options from .mylogin.cnf	5.6.6
<a href="#">--master-data[=value] [288]</a>	<a href="#">master-data [288]</a>	Write the binary log file name and position to the output	
<a href="#">--max_allowed_packet=value</a>	<a href="#">max_allowed_packet [285]</a>	The maximum packet length to send to or receive from the server	
<a href="#">--net_buffer_length=value [285]</a>	<a href="#">net_buffer_length [285]</a>	The buffer size for TCP/IP and socket communication	
<a href="#">--no-autocommit [295]</a>	<a href="#">no-autocommit [295]</a>	Enclose the INSERT statements for each dumped table within SET autocommit = 0 and COMMIT statements	
<a href="#">--no-create-db [286]</a>	<a href="#">no-create-db [286]</a>	This option suppresses the CREATE DATABASE statements	
<a href="#">--no-create-info [286]</a>	<a href="#">no-create-info [286]</a>	Do not write CREATE TABLE statements that re-create each dumped table	
<a href="#">--no-data [293]</a>	<a href="#">no-data [293]</a>	Do not dump table contents	
<a href="#">--no-set-names [288]</a>	<a href="#">no-set-names [288]</a>	Same as --skip-set-charset	
<a href="#">--no-tablespaces [286]</a>	<a href="#">no-tablespaces [286]</a>	Do not write any CREATE LOGFILE GROUP or CREATE TABLESPACE statements in output	
<a href="#">--opt [294]</a>	<a href="#">opt [294]</a>	Shorthand for --add-drop-table --add-locks --create-options --disable-keys --extended-insert --lock-tables --quick --set-charset.	
<a href="#">--order-by-primary [295]</a>	<a href="#">order-by-primary [295]</a>	Dump each table's rows sorted by its primary key, or by its first unique index	

Format	Option File	Description	Introduced
--password[=password] [285]	password [285]	The password to use when connecting to the server	
--pipe [285]		On Windows, connect to server using a named pipe	
--plugin-dir=path [285]	plugin-dir=path [285]	The directory where plugins are located	
--port=port_num [285]	port [285]	The TCP/IP port number to use for the connection	
--quick [294]	quick [294]	Retrieve rows for a table from the server a row at a time	
--quote-names [291]	quote-names [291]	Quote identifiers within backtick characters	
--replace [286]	replace [286]	Write REPLACE statements rather than INSERT statements	
--result-file=file [291]	result-file [291]	Direct output to a given file	
--routines [293]	routines [293]	Dump stored routines (procedures and functions) from the dumped databases	
--set-charset [288]	set-charset [288]	Add SET NAMES default_character_set to output	
--set-gtid-purged=value [289]	set-gtid-purged [289]	Whether to add SET @@GLOBAL.GTID_PURGED to output	5.6.9
--single-transaction [295]	single-transaction [295]	This option issues a BEGIN SQL statement before dumping data from the server	
--skip-add-drop-table [286]	skip-add-drop-table [286]	Do not add a DROP TABLE statement before each CREATE TABLE statement	
--skip-add-locks [294]	skip-add-locks [294]	Do not add locks	
--skip-comments [287]	skip-comments [287]	Do not add comments to the dump file	
--skip-compact [290]	skip-compact [290]	Do not produce more compact output	
--skip-disable-keys [294]	skip-disable-keys [294]	Do not disable keys	
--skip-extended-insert [294]	skip-extended-insert [294]	Turn off extended-insert	
--skip-opt [294]	skip-opt [294]	Turn off the options set by --opt	
--skip-quick [294]	skip-quick [294]	Do not retrieve rows for a table from the server a row at a time	
--skip-quote-names [291]	skip-quote-names [291]	Do not quote identifiers	
--skip-set-charset [288]	skip-set-charset [288]	Suppress the SET NAMES statement	
--skip-triggers [293]	skip-triggers [293]	Do not dump triggers	
--skip-tz-utc [291]	skip-tz-utc [291]	Turn off tz-utc	

Format	Option File	Description	Introduced
<code>--ssl-ca=file_name</code> [285]	<code>ssl-ca</code> [285]	The path to a file that contains a list of trusted SSL CAs	
<code>--ssl-capath=dir_name</code> [285]	<code>ssl-capath</code> [285]	The path to a directory that contains trusted SSL CA certificates in PEM format	
<code>--ssl-cert=file_name</code> [285]	<code>ssl-cert</code> [285]	The name of the SSL certificate file to use for establishing a secure connection	
<code>--ssl-cipher=cipher_list</code> [285]	<code>ssl-cipher</code> [285]	A list of allowable ciphers to use for SSL encryption	
<code>--ssl-crl=file_name</code> [285]	<code>ssl-crl</code> [285]	The path to a file that contains certificate revocation lists	5.6.3
<code>--ssl-crlpath=dir_name</code> [285]	<code>ssl-crlpath</code> [285]	The path to a directory that contains certificate revocation list files	5.6.3
<code>--ssl-key=file_name</code> [285]	<code>ssl-key</code> [285]	The name of the SSL key file to use for establishing a secure connection	
<code>--ssl-verify-server-cert</code> [285]	<code>ssl-verify-server-cert</code> [285]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<code>--tab=path</code> [291]	<code>tab</code> [291]	Produce tab-separated data files	
<code>--tables</code> [293]	<code>tables</code> [293]	Override the <code>--databases</code> or <code>-B</code> option	
<code>--triggers</code> [293]	<code>triggers</code> [293]	Dump triggers for each dumped table	
<code>--tz-utc</code> [291]	<code>tz-utc</code> [291]	Add <code>SET TIME_ZONE='+00:00'</code> to the dump file	
<code>--user=user_name</code> [285]	<code>user</code> [285]	MySQL user name to use when connecting to server	
<code>--verbose</code> [287]		Verbose mode	
<code>--version</code> [287]		Display version information and exit	
<code>--where='where_condition'</code> [293]	<code>where</code> [293]	Dump only rows selected by the given WHERE condition	
<code>--xml</code> [291]	<code>xml</code> [291]	Produce XML output	

## Connection Options

The `mysqldump` command logs into a MySQL server to extract information. The following options specify how to connect to the MySQL server, either on the same machine or a remote system.

- `--bind-address=ip_address` [248]

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--compress, -C`

Compress all information sent between the client and the server if both support compression.

- `--default-auth=plugin`

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--host=host_name, -h host_name`

Dump data from the MySQL server on the given host. The default host is `localhost`.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the *password* value following the `--password` or `-p` option on the command line, `mysqldump` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe, -W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path`

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [284] option is used to specify an authentication plugin but `mysqldump` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP | SOCKET | PIPE | MEMORY}`

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--user=user_name, -u user_name`

The MySQL user name to use when connecting to the server.

You can also set the following variables by using `--var_name=value` syntax:

- `max_allowed_packet`

The maximum size of the buffer for client/server communication. The maximum is 1GB.

- `net_buffer_length`

The initial size of the buffer for client/server communication. When creating multiple-row `INSERT` statements (as with the `--extended-insert` [294] or `--opt` [294] option), `mysqldump` creates rows up to `net_buffer_length` [515] length. If you increase this variable, ensure that the `net_buffer_length` [515] variable in the MySQL server is at least this large.

## DDL Options

Usage scenarios for `mysqldump` include setting up an entire new MySQL instance (including database tables), and replacing data inside an existing instance with existing databases and tables. The following

options let you specify which things to tear down and set up when restoring a dump, by encoding various DDL statements within the dump file.

- `--add-drop-database`

Add a `DROP DATABASE` statement before each `CREATE DATABASE` statement. This option is typically used in conjunction with the `--all-databases` [292] or `--databases` [292] option because no `CREATE DATABASE` statements are written unless one of those options is specified.

- `--add-drop-table`

Add a `DROP TABLE` statement before each `CREATE TABLE` statement.

- `--add-drop-trigger`

Add a `DROP TRIGGER` statement before each `CREATE TRIGGER` statement.

- `--all-tablespaces, -Y`

Adds to a table dump all SQL statements needed to create any tablespaces used by an `NDB` table. This information is not otherwise included in the output from `mysqldump`. This option is currently relevant only to MySQL Cluster tables.

- `--no-create-db, -n`

This option suppresses the `CREATE DATABASE` statements that are otherwise included in the output if the `--databases` [292] or `--all-databases` [292] option is given.

- `--no-create-info, -t`

Do not write `CREATE TABLE` statements that re-create each dumped table.



#### Note

This option does *not* not exclude statements creating log file groups or tablespaces from `mysqldump` output; however, you can use the `--no-tablespaces` [286] option for this purpose.

- `--no-tablespaces, -y`

This option suppresses all `CREATE LOGFILE GROUP` and `CREATE TABLESPACE` statements in the output of `mysqldump`.

- `--replace`

Write `REPLACE` statements rather than `INSERT` statements.

## Debug Options

The following options print debugging information, encode debugging information in the dump file, or let the dump operation proceed regardless of potential problems.

- `--allow-keywords`

Permit creation of column names that are keywords. This works by prefixing each column name with the table name.

- `--comments, -i`

Write additional information in the dump file such as program version, server version, and host. This option is enabled by default. To suppress this additional information, use `--skip-comments` [287].

- `--debug[=debug_options], -# [debug_options]`

Write a debugging log. A typical *debug\_options* string is `'d:t:o,file_name'`. The default value is `'d:t:o,/tmp/mysqldump.trace'`.

- `--debug-check`

Print some debugging information when the program exits.

- `--debug-info`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--dump-date`

If the `--comments` [286] option is given, `mysqldump` produces a comment at the end of the dump of the following form:

```
-- Dump completed on DATE
```

However, the date causes dump files taken at different times to appear to be different, even if the data are otherwise identical. `--dump-date` [287] and `--skip-dump-date` [287] control whether the date is added to the comment. The default is `--dump-date` [287] (include the date in the comment). `--skip-dump-date` [287] suppresses date printing.

- `--force, -f`

Continue even if an SQL error occurs during a table dump.

One use for this option is to cause `mysqldump` to continue executing even when it encounters a view that has become invalid because the definition refers to a table that has been dropped. Without `--force`, `mysqldump` exits with an error message. With `--force`, `mysqldump` prints the error message, but it also writes an SQL comment containing the view definition to the dump output and continues executing.

- `--log-error=file_name`

Log warnings and errors by appending them to the named file. The default is to do no logging.

- `--skip-comments`

See the description for the `--comments` [286] option.

- `--verbose, -v`

Verbose mode. Print more information about what the program does.

## Help Options

The following options display information about the `mysqldump` command itself.

- `--help, -?`

Display a help message and exit.

- `--version, -V`

Display version information and exit.

## Internationalization Options

The following options change how the `mysqldump` command represents character data with national language settings.

- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--default-character-set=charset_name`

Use `charset_name` as the default character set. See [Section 10.5, “Character Set Configuration”](#). If no character set is specified, `mysqldump` uses `utf8`, and earlier versions use `latin1`.

- `--no-set-names, -N`

Turns off the `--set-charset` [\[288\]](#) setting, the same as specifying `--skip-set-charset`.

- `--set-charset`

Add `SET NAMES default_character_set` to the output. This option is enabled by default. To suppress the `SET NAMES` statement, use `--skip-set-charset` [\[288\]](#).

## Replication Options

The `mysqldump` command is frequently used to create an empty instance, or an instance including data, on a slave server in a replication configuration. The following options apply to dumping and restoring data on replication master and slave servers.

- `--apply-slave-statements`

For a slave dump produced with the `--dump-slave` [\[288\]](#) option, add a `STOP SLAVE` statement before the `CHANGE MASTER TO` statement and a `START SLAVE` statement at the end of the output.

- `--delete-master-logs`

On a master replication server, delete the binary logs by sending a `PURGE BINARY LOGS` statement to the server after performing the dump operation. This option automatically enables `--master-data` [\[288\]](#).

- `--dump-slave[=value]`

This option is similar to `--master-data` [\[288\]](#) except that it is used to dump a replication slave server to produce a dump file that can be used to set up another server as a slave that has the same master as the dumped server. It causes the dump output to include a `CHANGE MASTER TO` statement that indicates the binary log coordinates (file name and position) of the dumped slave's master (rather than the coordinates of the dumped server, as is done by the `--master-data` [\[288\]](#) option). These are the master server coordinates from which the slave should start replicating.

The option value is handled the same way as for `--master-data` [\[288\]](#) and has the same effect as `--master-data` in terms of enabling or disabling other options and in how locking is handled.

This option causes `mysqldump` to stop the slave SQL thread before the dump and restart it again after.

In conjunction with `--dump-slave`, the `--apply-slave-statements` [\[288\]](#) and `--include-master-host-port` [\[288\]](#) options can also be used.

- `--include-master-host-port`

For the `CHANGE MASTER TO` statement in a slave dump produced with the `--dump-slave` [\[288\]](#) option, add `MASTER_PORT` and `MASTER_HOST` options for the host name and TCP/IP port number of the slave's master.

- `--master-data[=value]`

Use this option to dump a master replication server to produce a dump file that can be used to set up another server as a slave of the master. It causes the dump output to include a `CHANGE MASTER`



`TO` statement that indicates the binary log coordinates (file name and position) of the dumped server. These are the master server coordinates from which the slave should start replicating after you load the dump file into the slave.

If the option value is 2, the `CHANGE MASTER TO` statement is written as an SQL comment, and thus is informative only; it has no effect when the dump file is reloaded. If the option value is 1, the statement is not written as a comment and takes effect when the dump file is reloaded. If no option value is specified, the default value is 1.

This option requires the `RELOAD [755]` privilege and the binary log must be enabled.

The `--master-data` option automatically turns off `--lock-tables [295]`. It also turns on `--lock-all-tables [295]`, unless `--single-transaction [295]` also is specified, in which case, a global read lock is acquired only for a short time at the beginning of the dump (see the description for `--single-transaction [295]`). In all cases, any action on logs happens at the exact moment of the dump.

It is also possible to set up a slave by dumping an existing slave of the master. To do this, use the following procedure on the existing slave:

1. Stop the slave's SQL thread and get its current status:

```
mysql> STOP SLAVE SQL_THREAD;
mysql> SHOW SLAVE STATUS;
```

2. From the output of the `SHOW SLAVE STATUS` statement, the binary log coordinates of the master server from which the new slave should start replicating are the values of the `Relay_Master_Log_File` and `Exec_Master_Log_Pos` fields. Denote those values as `file_name` and `file_pos`.

3. Dump the slave server:

```
shell> mysqldump --master-data=2 --all-databases > dumpfile
```

Using `--master-data=2` works only if binary logging has been enabled on the slave. Otherwise, `mysqldump` fails with the error `Binlogging on server not active`. In this case you must handle any locking issues in another manner, using one or more of `--add-locks [294]`, `--lock-tables [295]`, `--lock-all-tables [295]`, or `--single-transaction [295]`, as required by your application and environment.

4. Restart the slave:

```
mysql> START SLAVE;
```

5. On the new slave, load the dump file:

```
shell> mysql < dumpfile
```

6. On the new slave, set the replication coordinates to those of the master server obtained earlier:

```
mysql> CHANGE MASTER TO
-> MASTER_LOG_FILE = 'file_name', MASTER_LOG_POS = file_pos;
```

The `CHANGE MASTER TO` statement might also need other parameters, such as `MASTER_HOST` to point the slave to the correct master server host. Add any such parameters as necessary.

Prior to MySQL 5.6.4, this option was required for dumping the replication log tables (see [Section 16.2.2, “Replication Relay and Status Logs”](#)).

- `--set-gtid-purged=value`

This option enables control over global transaction ID (GTID) information written to the dump file, by indicating whether to add a `SET @@global.gtid_purged` statement to the output.

The following table shows the permitted option values. The default value is `AUTO`.

Value	Meaning
<code>OFF</code>	Add no <code>SET</code> statement to the output.
<code>ON</code>	Add a <code>SET</code> statement to the output. An error occurs if GTIDs are not enabled on the server.
<code>AUTO</code>	Add a <code>SET</code> statement to the output if GTIDs are enabled on the server.

This option was added in MySQL 5.6.9.

## Format Options

The following options specify how to represent the entire dump file or certain kinds of data in the dump file. They also control whether certain optional information is written to the dump file.

- `--compact`

Produce more compact output. This option enables the `--skip-add-drop-table` [286], `--skip-add-locks` [294], `--skip-comments` [287], `--skip-disable-keys` [294], and `--skip-set-charset` [288] options.

- `--compatible=name`

Produce output that is more compatible with other database systems or with older MySQL servers. The value of *name* can be `ansi`, `mysql323`, `mysql40`, `postgresql`, `oracle`, `mssql`, `db2`, `maxdb`, `no_key_options`, `no_table_options`, or `no_field_options`. To use several values, separate them by commas. These values have the same meaning as the corresponding options for setting the server SQL mode. See [Section 5.1.7, “Server SQL Modes”](#).

This option does not guarantee compatibility with other servers. It only enables those SQL mode values that are currently available for making dump output more compatible. For example, `--compatible=oracle` does not map data types to Oracle types or use Oracle comment syntax.

*This option requires a server version of 4.1.0 or higher. With older servers, it does nothing.*

- `--complete-insert, -c`

Use complete `INSERT` statements that include column names.

- `--create-options`

Include all MySQL-specific table options in the `CREATE TABLE` statements.

- `--fields-terminated-by=...`, `--fields-enclosed-by=...`, `--fields-optionally-enclosed-by=...`, `--fields-escaped-by=...`

These options are used with the `--tab` [291] option and have the same meaning as the corresponding `FIELDS` clauses for `LOAD DATA INFILE`. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

- `--hex-blob`

Dump binary columns using hexadecimal notation (for example, `'abc'` becomes `0x616263`). The affected data types are `BINARY`, `VARBINARY`, the `BLOB` types, and `BIT`.

- `--lines-terminated-by=...`

This option is used with the `--tab` [291] option and has the same meaning as the corresponding `LINES` clause for `LOAD DATA INFILE`. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

- `--quote-names, -Q`

Quote identifiers (such as database, table, and column names) within “`” characters. If the [ANSI\\_QUOTES \[603\]](#) SQL mode is enabled, identifiers are quoted within “” characters. This option is enabled by default. It can be disabled with `--skip-quote-names`, but this option should be given after any option such as `--compatible [290]` that may enable `--quote-names [291]`.

- `--result-file=file_name, -r file_name`

Direct output to a given file. This option should be used on Windows to prevent newline “\n” characters from being converted to “\r\n” carriage return/newline sequences. The result file is created and its previous contents overwritten, even if an error occurs while generating the dump.

- `--tab=path, -T path`

Produce tab-separated text-format data files. For each dumped table, `mysqldump` creates a `tbl_name.sql` file that contains the `CREATE TABLE` statement that creates the table, and the server writes a `tbl_name.txt` file that contains its data. The option value is the directory in which to write the files.



#### Note

This option should be used only when `mysqldump` is run on the same machine as the `mysqld` server. You must have the [FILE \[754\]](#) privilege, and the server must have permission to write files in the directory that you specify.

By default, the `.txt` data files are formatted using tab characters between column values and a newline at the end of each line. The format can be specified explicitly using the `--fields-xxx` and `--lines-terminated-by [290]` options.

Column values are converted to the character set specified by the `--default-character-set [288]` option.

- `--tz-utc`

This option enables `TIMESTAMP` columns to be dumped and reloaded between servers in different time zones. `mysqldump` sets its connection time zone to UTC and adds `SET TIME_ZONE= '+00:00'` to the dump file. Without this option, `TIMESTAMP` columns are dumped and reloaded in the time zones local to the source and destination servers, which can cause the values to change if the servers are in different time zones. `--tz-utc` also protects against changes due to daylight saving time. `--tz-utc` is enabled by default. To disable it, use `--skip-tz-utc`.

- `--xml, -X`

Write dump output as well-formed XML.

**NULL, 'NULL', and Empty Values:** For a column named `column_name`, the `NULL` value, an empty string, and the string value `'NULL'` are distinguished from one another in the output generated by this option as follows.

Value:	XML Representation:
<code>NULL</code> ( <i>unknown value</i> )	<code>&lt;field name="column_name" xsi:nil="true" /&gt;</code>
<code>' '</code> ( <i>empty string</i> )	<code>&lt;field name="column_name"&gt;&lt;/field&gt;</code>
<code>'NULL'</code> ( <i>string value</i> )	<code>&lt;field name="column_name"&gt;NULL&lt;/field&gt;</code>

The output from the `mysql` client when run using the `--xml` [254] option also follows the preceding rules. (See [Section 4.5.1.1](#), “`mysql` Options”.)

XML output from `mysqldump` includes the XML namespace, as shown here:

```
shell> mysqldump --xml -u root world City
<?xml version="1.0"?>
<mysqldump xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<database name="world">
<table_structure name="City">
<field Field="ID" Type="int(11)" Null="NO" Key="PRI" Extra="auto_increment" />
<field Field="Name" Type="char(35)" Null="NO" Key="" Default="" Extra="" />
<field Field="CountryCode" Type="char(3)" Null="NO" Key="" Default="" Extra="" />
<field Field="District" Type="char(20)" Null="NO" Key="" Default="" Extra="" />
<field Field="Population" Type="int(11)" Null="NO" Key="" Default="0" Extra="" />
<key Table="City" Non_unique="0" Key_name="PRIMARY" Seq_in_index="1" Column_name="ID"
Collation="A" Cardinality="4079" Null="" Index_type="BTREE" Comment="" />
<options Name="City" Engine="MyISAM" Version="10" Row_format="Fixed" Rows="4079"
Avg_row_length="67" Data_length="273293" Max_data_length="18858823439613951"
Index_length="43008" Data_free="0" Auto_increment="4080"
Create_time="2007-03-31 01:47:01" Update_time="2007-03-31 01:47:02"
Collation="latin1_swedish_ci" Create_options="" Comment="" />
</table_structure>
<table_data name="City">
<row>
<field name="ID">1</field>
<field name="Name">Kabul</field>
<field name="CountryCode">AFG</field>
<field name="District">Kabul</field>
<field name="Population">1780000</field>
</row>
...
<row>
<field name="ID">4079</field>
<field name="Name">Rafah</field>
<field name="CountryCode">PSE</field>
<field name="District">Rafah</field>
<field name="Population">92020</field>
</row>
</table_data>
</database>
</mysqldump>
```

Prior to MySQL 5.6.5, this option prevented the `--routines` [293] option from working correctly—that is, no stored routines, triggers, or events could be dumped in XML format. (Bug #11760384, Bug #52792)

## Filtering Options

The following options control which kinds of schema objects are written to the dump file: by category, such as triggers or events; by name, for example, choosing which databases and tables to dump; or even filtering rows from the table data using a `WHERE` clause.

- `--all-databases, -A`

Dump all tables in all databases. This is the same as using the `--databases` [292] option and naming all the databases on the command line.

Prior to MySQL 5.6.4, the `slave_master_info` and `slave_relay_log_info` tables (see [Section 16.2.2](#), “`Replication Relay and Status Logs`”) were not included by this option.

- `--databases, -B`

Dump several databases. Normally, `mysqldump` treats the first name argument on the command line as a database name and following names as table names. With this option, it treats all name

arguments as database names. `CREATE DATABASE` and `USE` statements are included in the output before each new database.

- `--events, -E`

Include Event Scheduler events for the dumped databases in the output.

- `--ignore-table=db_name.tbl_name`

Do not dump the given table, which must be specified using both the database and table names. To ignore multiple tables, use this option multiple times. This option also can be used to ignore views.

- `--no-data, -d`

Do not write any table row information (that is, do not dump table contents). This is useful if you want to dump only the `CREATE TABLE` statement for the table (for example, to create an empty copy of the table by loading the dump file).

- `--routines, -R`

Include stored routines (procedures and functions) for the dumped databases in the output. Use of this option requires the `SELECT [755]` privilege for the `mysql.proc` table. The output generated by using `--routines` contains `CREATE PROCEDURE` and `CREATE FUNCTION` statements to re-create the routines. However, these statements do not include attributes such as the routine creation and modification timestamps. This means that when the routines are reloaded, they will be created with the timestamps equal to the reload time.

If you require routines to be re-created with their original timestamp attributes, do not use `--routines`. Instead, dump and reload the contents of the `mysql.proc` table directly, using a MySQL account that has appropriate privileges for the `mysql` database.

Prior to MySQL 5.6.5, this option had no effect when used together with the `--xml [291]` option. (Bug #11760384, Bug #52792)

- `--tables`

Override the `--databases [292]` or `-B` option. `mysqldump` regards all name arguments following the option as table names.

- `--triggers`

Include triggers for each dumped table in the output. This option is enabled by default; disable it with `--skip-triggers`.

- `--where='where_condition', -w 'where_condition'`

Dump only rows selected by the given `WHERE` condition. Quotes around the condition are mandatory if it contains spaces or other characters that are special to your command interpreter.

Examples:

```
--where="user='jimf'"
-w"userid>1"
-w"userid<1"
```

## Performance Options

The following options are the most relevant for the performance particularly of the restore operations. For large data sets, restore operation (processing the `INSERT` statements in the dump file) is the most time-consuming part. When it is urgent to restore data quickly, plan and test the performance of this stage in advance. For restore times measured in hours, you might prefer an alternative backup and restore solution, such as [MySQL Enterprise Backup](#) for `InnoDB`-only and mixed-use databases, or [mysqlhotcopy](#) for `MyISAM`-only databases.

Performance is also affected by the [transactional options](#), primarily for the dump operation.

- `--delayed-insert`

For those nontransactional tables that support the `INSERT DELAYED` syntax, use that statement rather than regular `INSERT` statements.

As of MySQL 5.6.6, `DELAYED` inserts are deprecated, so this option will be removed in a future release.

- `--disable-keys, -K`

For each table, surround the `INSERT` statements with `/*!40000 ALTER TABLE tbl_name DISABLE KEYS */;` and `/*!40000 ALTER TABLE tbl_name ENABLE KEYS */;` statements. This makes loading the dump file faster because the indexes are created after all rows are inserted. This option is effective only for nonunique indexes of `MyISAM` tables.

- `--extended-insert, -e`

Use multiple-row `INSERT` syntax that include several `VALUES` lists. This results in a smaller dump file and speeds up inserts when the file is reloaded.

- `--insert-ignore`

Write `INSERT IGNORE` statements rather than `INSERT` statements.

- `--opt`

This option, enabled by default, is shorthand for the combination of `--add-drop-table` [286] `--add-locks` [294] `--create-options` [290] `--disable-keys` [294] `--extended-insert` [294] `--lock-tables` [295] `--quick` [294] `--set-charset` [288]. It gives a fast dump operation and produces a dump file that can be reloaded into a MySQL server quickly.

Because the `--opt` option is enabled by default, you only specify its converse, the `--skip-opt` [294] to turn off several default settings. See the discussion of [mysqldump option groups](#) for information about selectively enabling or disabling a subset of the options affected by `--opt`.

- `--quick, -q`

This option is useful for dumping large tables. It forces `mysqldump` to retrieve rows for a table from the server a row at a time rather than retrieving the entire row set and buffering it in memory before writing it out.

- `--skip-opt`

See the description for the `--opt` [294] option.

## Transactional Options

The following options trade off the performance of the dump operation, against the reliability and consistency of the exported data.

- `--add-locks`

Surround each table dump with `LOCK TABLES` and `UNLOCK TABLES` statements. This results in faster inserts when the dump file is reloaded. See [Section 8.2.2.1, “Speed of INSERT Statements”](#).

- `--flush-logs, -F`

Flush the MySQL server log files before starting the dump. This option requires the `RELOAD` [755] privilege. If you use this option in combination with the `--all-databases` [292] option, the logs are flushed *for each database dumped*. The exception is when using `--lock-all-tables` [295],

`--master-data` [288], or `--single-transaction` [295]: In this case, the logs are flushed only once, corresponding to the moment that all tables are locked. If you want your dump and the log flush to happen at exactly the same moment, you should use `--flush-logs` together with `--lock-all-tables` [295], `--master-data` [288], or `--single-transaction` [295].

- `--flush-privileges`

Send a `FLUSH PRIVILEGES` statement to the server after dumping the `mysql` database. This option should be used any time the dump contains the `mysql` database and any other database that depends on the data in the `mysql` database for proper restoration.

- `--lock-all-tables, -x`

Lock all tables across all databases. This is achieved by acquiring a global read lock for the duration of the whole dump. This option automatically turns off `--single-transaction` [295] and `--lock-tables` [295].

- `--lock-tables, -l`

For each dumped database, lock all tables to be dumped before dumping them. The tables are locked with `READ LOCAL` to permit concurrent inserts in the case of `MyISAM` tables. For transactional tables such as `InnoDB`, `--single-transaction` [295] is a much better option than `--lock-tables` because it does not need to lock the tables at all.

Because `--lock-tables` locks tables for each database separately, this option does not guarantee that the tables in the dump file are logically consistent between databases. Tables in different databases may be dumped in completely different states.

Some options, such as `--opt` [294], automatically enable `--lock-tables`. If you want to override this, use `--skip-lock-tables` at the end of the option list.

- `--no-autocommit`

Enclose the `INSERT` statements for each dumped table within `SET autocommit = 0` and `COMMIT` statements.

- `--order-by-primary`

Dump each table's rows sorted by its primary key, or by its first unique index, if such an index exists. This is useful when dumping a `MyISAM` table to be loaded into an `InnoDB` table, but makes the dump operation take considerably longer.

- `--single-transaction`

This option sets the transaction isolation mode to `REPEATABLE READ` and sends a `START TRANSACTION` SQL statement to the server before dumping data. It is useful only with transactional tables such as `InnoDB`, because then it dumps the consistent state of the database at the time when `START TRANSACTION` was issued without blocking any applications.

When using this option, you should keep in mind that only `InnoDB` tables are dumped in a consistent state. For example, any `MyISAM` or `MEMORY` tables dumped while using this option may still change state.

The `--single-transaction` option and the `--lock-tables` [295] option are mutually exclusive because `LOCK TABLES` causes any pending transactions to be committed implicitly.

To dump large tables, combine the `--single-transaction` option with the `--quick` [294].

## Option Groups

- The `--opt` [294] option turns on several settings that work together to perform a fast dump operation. All of these settings are on by default, because `--opt` is on by default. Thus you rarely if



ever specify `--opt`. Instead, you can turn these settings off as a group by specifying `--skip-opt`, the optionally re-enable certain settings by specifying the associated options later on the command line.

- The `--compact` [290] option turns off several settings that control whether optional statements and comments appear in the output. Again, you can follow this option with other options that re-enable certain settings, or turn all the settings on by using the `--skip-compact` form.

When you selectively enable or disable the effect of a group option, order is important because options are processed first to last. For example, `--disable-keys` [294] `--lock-tables` [295] `--skip-opt` [294] would not have the intended effect; it is the same as `--skip-opt` [294] by itself.

## Examples

To make a backup of an entire database:

```
shell> mysqldump db_name > backup-file.sql
```

To load the dump file back into the server:

```
shell> mysql db_name < backup-file.sql
```

Another way to reload the dump file:

```
shell> mysql -e "source /path-to-backup/backup-file.sql" db_name
```

`mysqldump` is also very useful for populating databases by copying data from one MySQL server to another:

```
shell> mysqldump --opt db_name | mysql --host=remote_host -C db_name
```

You can dump several databases with one command:

```
shell> mysqldump --databases db_name1 [db_name2 ...] > my_databases.sql
```

To dump all databases, use the `--all-databases` [292] option:

```
shell> mysqldump --all-databases > all_databases.sql
```

For InnoDB tables, `mysqldump` provides a way of making an online backup:

```
shell> mysqldump --all-databases --single-transaction > all_databases.sql
```

This backup acquires a global read lock on all tables (using `FLUSH TABLES WITH READ LOCK`) at the beginning of the dump. As soon as this lock has been acquired, the binary log coordinates are read and the lock is released. If long updating statements are running when the `FLUSH` statement is issued, the MySQL server may get stalled until those statements finish. After that, the dump becomes lock free and does not disturb reads and writes on the tables. If the update statements that the MySQL server receives are short (in terms of execution time), the initial lock period should not be noticeable, even with many updates.

For point-in-time recovery (also known as “roll-forward,” when you need to restore an old backup and replay the changes that happened since that backup), it is often useful to rotate the binary log (see [Section 5.2.4, “The Binary Log”](#)) or at least know the binary log coordinates to which the dump corresponds:

```
shell> mysqldump --all-databases --master-data=2 > all_databases.sql
```

Or:

```
shell> mysqldump --all-databases --flush-logs --master-data=2
> all_databases.sql
```

The `--master-data` [288] and `--single-transaction` [295] options can be used simultaneously, which provides a convenient way to make an online backup suitable for use prior to point-in-time recovery if tables are stored using the InnoDB storage engine.

For more information on making backups, see [Section 7.2, “Database Backup Methods”](#), and [Section 7.3, “Example Backup and Recovery Strategy”](#).

- To select the effect of `--opt [294]` except for some features, use the `--skip` option for each feature. To disable extended inserts and memory buffering, use `--opt [294] --skip-extended-insert [294] --skip-quick [294]`. (Actually, `--skip-extended-insert [294] --skip-quick [294]` is sufficient because `--opt [294]` is on by default.)
- To reverse `--opt [294]` for all features except index disabling and table locking, use `--skip-opt [294] --disable-keys [294] --lock-tables [295]`.

## Restrictions

`mysqldump` does not dump the `INFORMATION_SCHEMA` database by default. To dump `INFORMATION_SCHEMA`, name it explicitly on the command line and also use the `--skip-lock-tables [295]` option.

`mysqldump` never dumps the `performance_schema` database.

Before MySQL 5.6.6, `mysqldump` does not dump the `general_log` or `slow_query_log` tables for dumps of the `mysql` database. As of 5.6.6, the dump includes statements to recreate those tables so that they are not missing after reloading the dump file. Log table contents are not dumped.

It is not recommended to restore from a dump made using `mysqldump` to a MySQL Server that has GTIDs enabled. See [Section 16.1.3.4, “Restrictions on Replication with GTIDs”](#).

If you encounter problems backing up views due to insufficient privileges, see [Section E.5, “Restrictions on Views”](#) for a workaround.

## 4.5.5. mysqlimport — A Data Import Program

The `mysqlimport` client provides a command-line interface to the `LOAD DATA INFILE SQL` statement. Most options to `mysqlimport` correspond directly to clauses of `LOAD DATA INFILE` syntax. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

Invoke `mysqlimport` like this:

```
shell> mysqlimport [options] db_name textfile1 [textfile2 ...]
```

For each text file named on the command line, `mysqlimport` strips any extension from the file name and uses the result to determine the name of the table into which to import the file's contents. For example, files named `patient.txt`, `patient.text`, and `patient` all would be imported into a table named `patient`.

`mysqlimport` supports the following options, which can be specified on the command line or in the `[mysqlimport]` and `[client]` groups of an option file. `mysqlimport` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.7. `mysqlimport` Options**

Format	Option File	Description	Introduced
<code>--bind-address=ip_address [299]</code>	<code>bind-address [299]</code>	Use the specified network interface to connect to the MySQL Server	
<code>--columns=column_list [299]</code>	<code>columns [299]</code>	This option takes a comma-separated list of column names as its value	
<code>--compress [299]</code>	<code>compress [299]</code>	Compress all information sent between the client and the server	
<code>--debug[=debug_options] [300]</code>	<code>debug [300]</code>	Write a debugging log	

Format	Option File	Description	Introduced
--debug-check [300]	debug-check [300]	Print debugging information when the program exits	
--debug-info [300]	debug-info [300]	Print debugging information, memory and CPU statistics when the program exits	
--default-auth=plugin [300]	default-auth=plugin [300]	The authentication plugin to use	5.6.2
--default-character-set=charset_name [300]	default-character-set [300]	Use charset_name as the default character set	
--delete [300]	delete [300]	Empty the table before importing the text file	
--fields-enclosed-by=string [300]	fields-enclosed-by [300]	This option has the same meaning as the corresponding clause for LOAD DATA INFILE	
--fields-escaped-by [300]	fields-escaped-by [300]	This option has the same meaning as the corresponding clause for LOAD DATA INFILE	
--fields-optionally-enclosed-by=string [300]	fields-optionally-enclosed-by [300]	This option has the same meaning as the corresponding clause for LOAD DATA INFILE	
--fields-terminated-by=string [300]	fields-terminated-by [300]	-- This option has the same meaning as the corresponding clause for LOAD DATA INFILE	
--force [300]	force [300]	Continue even if an SQL error occurs	
--help [299]		Display help message and exit	
--host=host_name [300]	host [300]	Connect to the MySQL server on the given host	
--ignore [300]	ignore [300]	See the description for the --replace option	
--ignore-lines=# [300]	ignore-lines [300]	Ignore the first N lines of the data file	
--lines-terminated-by=string [300]	lines-terminated-by [300]	This option has the same meaning as the corresponding clause for LOAD DATA INFILE	
--local [300]	local [300]	Read input files locally from the client host	
--lock-tables [300]	lock-tables [300]	Lock all tables for writing before processing any text files	
--login-path=name [219]		Read login path options from .mylogin.cnf	5.6.6
--low-priority [301]	low-priority [301]	Use LOW_PRIORITY when loading the table.	
--password[=password] [301]	password [301]	The password to use when connecting to the server	
--pipe [301]		On Windows, connect to server using a named pipe	
--plugin-dir=path [301]	plugin-dir=path [301]	The directory where plugins are located	5.6.2
--port=port_num [301]	port [301]	The TCP/IP port number to use for the connection	
--protocol=type [301]	protocol [301]	The connection protocol to use	

Format	Option File	Description	Introduced
<code>--replace</code> [301]	<code>replace</code> [301]	The <code>--replace</code> and <code>--ignore</code> options control handling of input rows that duplicate existing rows on unique key values	
<code>--silent</code> [301]	<code>silent</code> [301]	Produce output only when errors occur	
<code>--socket=</code> <i>path</i> [301]	<code>socket</code> [301]	For connections to localhost	
<code>--ssl-ca=</code> <i>file_name</i> [301]	<code>ssl-ca</code> [301]	The path to a file that contains a list of trusted SSL CAs	
<code>--ssl-capath=</code> <i>dir_name</i> [301]	<code>ssl-capath</code> [301]	The path to a directory that contains trusted SSL CA certificates in PEM format	
<code>--ssl-cert=</code> <i>file_name</i> [301]	<code>ssl-cert</code> [301]	The name of the SSL certificate file to use for establishing a secure connection	
<code>--ssl-cipher=</code> <i>cipher_list</i> [301]	<code>ssl-cipher</code> [301]	A list of allowable ciphers to use for SSL encryption	
<code>--ssl-crl=</code> <i>file_name</i> [301]	<code>ssl-crl</code> [301]	The path to a file that contains certificate revocation lists	5.6.3
<code>--ssl-crlpath=</code> <i>dir_name</i> [301]	<code>ssl-crlpath</code> [301]	The path to a directory that contains certificate revocation list files	5.6.3
<code>--ssl-key=</code> <i>file_name</i> [301]	<code>ssl-key</code> [301]	The name of the SSL key file to use for establishing a secure connection	
<code>--ssl-verify-server-cert</code> [301]	<code>ssl-verify-server-cert</code> [301]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<code>--use-threads=</code> <i>#</i> [302]	<code>use-threads</code> [302]	The number of threads for parallel file-loading	
<code>--user=</code> <i>user_name</i> , [302]	<code>user</code> [302]	MySQL user name to use when connecting to server	
<code>--verbose</code> [302]		Verbose mode	
<code>--version</code> [302]		Display version information and exit	

- `--help` [299], `-?`

Display a help message and exit.

- `--bind-address=`*ip\_address* [248]

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--character-sets-dir=`*path* [299]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--columns=`*column\_list* [299], `-c` *column\_list*

This option takes a comma-separated list of column names as its value. The order of the column names indicates how to match data file columns with table columns.

- `--compress` [299], `-C`

Compress all information sent between the client and the server if both support compression.

- `--debug[=debug_options] [300], -# [debug_options]`

Write a debugging log. A typical *debug\_options* string is '*d:t:o,file\_name*'. The default is '*d:t:o*'.

- `--debug-check [300]`

Print some debugging information when the program exits.

- `--debug-info [300]`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-character-set=charset_name [300]`

Use *charset\_name* as the default character set. See [Section 10.5, “Character Set Configuration”](#).

- `--default-auth=plugin [300]`

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--delete [300], -D`

Empty the table before importing the text file.

- `--fields-terminated-by=... [300], --fields-enclosed-by=... [300], --fields-optionally-enclosed-by=... [300], --fields-escaped-by=... [300]`

These options have the same meaning as the corresponding clauses for `LOAD DATA INFILE`. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

- `--force [300], -f`

Ignore errors. For example, if a table for a text file does not exist, continue processing any remaining files. Without `--force [300]`, `mysqlimport` exits if a table does not exist.

- `--host=host_name [300], -h host_name`

Import data to the MySQL server on the given host. The default host is `localhost`.

- `--ignore [300], -i`

See the description for the `--replace [301]` option.

- `--ignore-lines=N [300]`

Ignore the first *N* lines of the data file.

- `--lines-terminated-by=... [300]`

This option has the same meaning as the corresponding clause for `LOAD DATA INFILE`. For example, to import Windows files that have lines terminated with carriage return/linefeed pairs, use `--lines-terminated-by="\r\n" [300]`. (You might have to double the backslashes, depending on the escaping conventions of your command interpreter.) See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

- `--local [300], -L`

Read input files locally from the client host.

- `--lock-tables [295], -l`

Lock *all* tables for writing before processing any text files. This ensures that all tables are synchronized on the server.

- `--low-priority` [301]

Use `LOW_PRIORITY` when loading the table. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).

- `--password[=password]` [301], `-p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` [301] or `-p` option on the command line, `mysqlimport` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe` [301], `-W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path` [301]

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [300] option is used to specify an authentication plugin but `mysqlimport` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--port=port_num` [301], `-P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [301]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--replace` [301], `-r`

The `--replace` [301] and `--ignore` [300] options control handling of input rows that duplicate existing rows on unique key values. If you specify `--replace` [301], new rows replace existing rows that have the same unique key value. If you specify `--ignore` [300], input rows that duplicate an existing row on a unique key value are skipped. If you do not specify either option, an error occurs when a duplicate key value is found, and the rest of the text file is ignored.

- `--silent` [301], `-s`

Silent mode. Produce output only when errors occur.

- `--socket=path` [301], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--user=user_name` [302], `-u user_name`

The MySQL user name to use when connecting to the server.

- `--use-threads=N` [302]

Load files in parallel using *N* threads.

- `--verbose` [302], `-v`

Verbose mode. Print more information about what the program does.

- `--version` [302], `-V`

Display version information and exit.

Here is a sample session that demonstrates use of `mysqlimport`:

```
shell> mysql -e 'CREATE TABLE impptest(id INT, n VARCHAR(30))' test
shell> ed
a
100      Max Sydow
101      Count Dracula
.
w impptest.txt
32
q
shell> od -c impptest.txt
0000000  1  0  0  \t  M  a  x           S  y  d  o  w  \n  1  0
0000020  1  \t  C  o  u  n  t           D  r  a  c  u  l  a  \n
0000040
shell> mysqlimport --local test impptest.txt
test.impptest: Records: 2 Deleted: 0 Skipped: 0 Warnings: 0
shell> mysql -e 'SELECT * FROM impptest' test
+-----+-----+
| id  | n                |
+-----+-----+
| 100 | Max Sydow        |
| 101 | Count Dracula    |
+-----+-----+
```

## 4.5.6. mysqlshow — Display Database, Table, and Column Information

The `mysqlshow` client can be used to quickly see which databases exist, their tables, or a table's columns or indexes.

`mysqlshow` provides a command-line interface to several SQL `SHOW` statements. See [Section 13.7.5, “SHOW Syntax”](#). The same information can be obtained by using those statements directly. For example, you can issue them from the `mysql` client program.

Invoke `mysqlshow` like this:

```
shell> mysqlshow [options] [db_name [tbl_name [col_name]]]
```

- If no database is given, a list of database names is shown.
- If no table is given, all matching tables in the database are shown.
- If no column is given, all matching columns and column types in the table are shown.

The output displays only the names of those databases, tables, or columns for which you have some privileges.



If the last argument contains shell or SQL wildcard characters (“\*”, “?”, “%”, or “\_”), only those names that are matched by the wildcard are shown. If a database name contains any underscores, those should be escaped with a backslash (some Unix shells require two) to get a list of the proper tables or columns. “\*” and “?” characters are converted into SQL “%” and “\_” wildcard characters. This might cause some confusion when you try to display the columns for a table with a “\_” in the name, because in this case, `mysqlshow` shows you only the table names that match the pattern. This is easily fixed by adding an extra “%” last on the command line as a separate argument.

`mysqlshow` supports the following options, which can be specified on the command line or in the `[mysqlshow]` and `[client]` groups of an option file. `mysqlshow` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.8. `mysqlshow` Options**

Format	Option File	Description	Introduced
<code>--bind-address=ip_address</code> [304]	<code>bind-address</code> [304]	Use the specified network interface to connect to the MySQL Server	
<code>--compress</code> [304]	<code>compress</code> [304]	Compress all information sent between the client and the server	
<code>--count</code> [304]	<code>count</code> [304]	Show the number of rows per table	
<code>--debug[=debug_options]</code> [304]	<code>debug</code> [304]	Write a debugging log	
<code>--debug-check</code> [304]	<code>debug-check</code> [304]	Print debugging information when the program exits	
<code>--debug-info</code> [304]	<code>debug-info</code> [304]	Print debugging information, memory and CPU statistics when the program exits	
<code>--default-auth=plugin</code> [305]	<code>default-auth=plugin</code> [305]	The authentication plugin to use	5.6.2
<code>--default-character-set=charset_name</code> [305]	<code>default-character-set</code> [305]	Use <code>charset_name</code> as the default character set	
<code>--help</code> [304]		Display help message and exit	
<code>--host=host_name</code> [305]	<code>host</code> [305]	Connect to the MySQL server on the given host	
<code>--keys</code> [305]	<code>keys</code> [305]	Show table indexes	
<code>--login-path=name</code> [219]		Read login path options from <code>.mylogin.cnf</code>	5.6.6
<code>--password[=password]</code> [305]	<code>password</code> [305]	The password to use when connecting to the server	
<code>--pipe</code> [305]		On Windows, connect to server using a named pipe	
<code>--plugin-dir=path</code> [305]	<code>plugin-dir=path</code> [305]	The directory where plugins are located	5.6.2
<code>--port=port_num</code> [305]	<code>port</code> [305]	The TCP/IP port number to use for the connection	
<code>--protocol=type</code> [305]	<code>protocol</code> [305]	The connection protocol to use	
<code>--show-table-type</code> [305]		Show a column indicating the table type	
<code>--socket=path</code> [305]	<code>socket</code> [305]	For connections to localhost	

Format	Option File	Description	Introduced
<code>--ssl-ca=file_name [306]</code>	<code>ssl-ca [306]</code>	The path to a file that contains a list of trusted SSL CAs	
<code>--ssl-capath=dir_name [306]</code>	<code>ssl-capath [306]</code>	The path to a directory that contains trusted SSL CA certificates in PEM format	
<code>--ssl-cert=file_name [306]</code>	<code>ssl-cert [306]</code>	The name of the SSL certificate file to use for establishing a secure connection	
<code>--ssl-cipher=cipher_list [306]</code>	<code>ssl-cipher [306]</code>	A list of allowable ciphers to use for SSL encryption	
<code>--ssl-crl=file_name [306]</code>	<code>ssl-crl [306]</code>	The path to a file that contains certificate revocation lists	5.6.3
<code>--ssl-crlpath=dir_name [306]</code>	<code>ssl-crlpath [306]</code>	The path to a directory that contains certificate revocation list files	5.6.3
<code>--ssl-key=file_name [306]</code>	<code>ssl-key [306]</code>	The name of the SSL key file to use for establishing a secure connection	
<code>--ssl-verify-server-cert [306]</code>	<code>ssl-verify-server-cert [306]</code>	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<code>--status [306]</code>	<code>status [306]</code>	Display extra information about each table	
<code>--user=user_name, [306]</code>	<code>user [306]</code>	MySQL user name to use when connecting to server	
<code>--verbose [306]</code>		Verbose mode	
<code>--version [306]</code>		Display version information and exit	

- `--help [304], -?`

Display a help message and exit.

- `--bind-address=ip_address [248]`

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--character-sets-dir=path [304]`

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--compress [304], -C`

Compress all information sent between the client and the server if both support compression.

- `--count [304]`

Show the number of rows per table. This can be slow for non-MyISAM tables.

- `--debug[=debug_options] [304], -# [debug_options]`

Write a debugging log. A typical *debug\_options* string is `'d:t:o,file_name'`. The default is `'d:t:o'`.

- `--debug-check [304]`

Print some debugging information when the program exits.

- `--debug-info [304]`

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-character-set=charset_name` [305]

Use *charset\_name* as the default character set. See [Section 10.5, “Character Set Configuration”](#).

- `--default-auth=plugin` [305]

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--host=host_name` [305], `-h host_name`

Connect to the MySQL server on the given host.

- `--keys` [305], `-k`

Show table indexes.

- `--password[=password]` [305], `-p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the *password* value following the `--password` [305] or `-p` option on the command line, `mysqlshow` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe` [305], `-W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path` [305]

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [305] option is used to specify an authentication plugin but `mysqlshow` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--port=port_num` [305], `-P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [305]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--show-table-type` [305], `-t`

Show a column indicating the table type, as in `SHOW FULL TABLES`. The type is `BASE TABLE` or `VIEW`.

- `--socket=path` [305], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--status` [306], `-i`

Display extra information about each table.

- `--user=user_name` [306], `-u user_name`

The MySQL user name to use when connecting to the server.

- `--verbose` [306], `-v`

Verbose mode. Print more information about what the program does. This option can be used multiple times to increase the amount of information.

- `--version` [306], `-V`

Display version information and exit.

## 4.5.7. mysqlslap — Load Emulation Client

`mysqlslap` is a diagnostic program designed to emulate client load for a MySQL server and to report the timing of each stage. It works as if multiple clients are accessing the server.

Invoke `mysqlslap` like this:

```
shell> mysqlslap [options]
```

Some options such as `--create` [310] or `--query` [312] enable you to specify a string containing an SQL statement or a file containing statements. If you specify a file, by default it must contain one statement per line. (That is, the implicit statement delimiter is the newline character.) Use the `--delimiter` [311] option to specify a different delimiter, which enables you to specify statements that span multiple lines or place multiple statements on a single line. You cannot include comments in a file; `mysqlslap` does not understand them.

`mysqlslap` runs in three stages:

1. Create schema, table, and optionally any stored programs or data to use for the test. This stage uses a single client connection.
2. Run the load test. This stage can use many client connections.
3. Clean up (disconnect, drop table if specified). This stage uses a single client connection.

Examples:

Supply your own create and query SQL statements, with 50 clients querying and 200 selects for each (enter the command on a single line):

```
mysqlslap --delimiter=";"
--create="CREATE TABLE a (b int);INSERT INTO a VALUES (23)"
--query="SELECT * FROM a" --concurrency=50 --iterations=200
```

Let `mysqlslap` build the query SQL statement with a table of two `INT` columns and three `VARCHAR` columns. Use five clients querying 20 times each. Do not create the table or insert the data (that is, use the previous test's schema and data):

```
mysqlslap --concurrency=5 --iterations=20
--number-int-cols=2 --number-char-cols=3
--auto-generate-sql
```

Tell the program to load the create, insert, and query SQL statements from the specified files, where the `create.sql` file has multiple table creation statements delimited by `;` and multiple insert

statements delimited by ';'. The `--query` file will have multiple queries delimited by ';'. Run all the load statements, then run all the queries in the query file with five clients (five times each):

```
mysqlslap --concurrency=5
--iterations=5 --query=query.sql --create=create.sql
--delimiter=";"
```

`mysqlslap` supports the following options, which can be specified on the command line or in the `[mysqlslap]` and `[client]` groups of an option file. `mysqlslap` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.9. `mysqlslap` Options**

Format	Option File	Description	Introduced
<code>--auto-generate-sql</code> [309]	<code>auto-generate-sql</code> [309]	Generate SQL statements automatically when they are not supplied in files or using command options	
<code>--auto-generate-sql-add-autoincrement</code> [309]	<code>auto-generate-sql-add-autoincrement</code> [309]	Add AUTO_INCREMENT column to automatically generated tables	
<code>--auto-generate-sql-execute-number=#</code> [309]	<code>auto-generate-sql-execute-number</code> [309]	Specify how many queries to generate automatically	
<code>--auto-generate-sql-guid-primary</code> [309]	<code>auto-generate-sql-guid-primary</code> [309]	Add a GUID-based primary key to automatically generated tables	
<code>--auto-generate-sql-load-type=type</code> [309]	<code>auto-generate-sql-load-type</code> [309]	Specify how many queries to generate automatically	
<code>--auto-generate-sql-secondary-indexes=#</code> [310]	<code>auto-generate-sql-secondary-indexes</code> [310]	Specify how many secondary indexes to add to automatically generated tables	
<code>--auto-generate-sql-unique-query-number=#</code> [310]	<code>auto-generate-sql-unique-query-number</code> [310]	How many different queries to generate for automatic tests.	
<code>--auto-generate-sql-unique-write-number=#</code> [310]	<code>auto-generate-sql-unique-write-number</code> [310]	How many different queries to generate for <code>--auto-generate-sql-write-number</code>	
<code>--auto-generate-sql-write-number=#</code> [310]	<code>auto-generate-sql-write-number</code> [310]	How many row inserts to perform on each thread	
<code>--commit=#</code> [310]	<code>commit</code> [310]	How many statements to execute before committing.	
<code>--compress</code> [310]	<code>compress</code> [310]	Compress all information sent between the client and the server	
<code>--concurrency=#</code> [310]	<code>concurrency</code> [310]	The number of clients to simulate when issuing the SELECT statement	
<code>--create=value</code> [310]	<code>create</code> [310]	The file or string containing the statement to use for creating the table	
<code>--create-schema=value</code> [310]	<code>create-schema</code> [310]	The schema in which to run the tests	
<code>--csv=[file]</code> [310]	<code>csv</code> [310]	Generate output in comma-separated values format	

Format	Option File	Description	Introduced
-- debug[=debug_options] [310]	debug [310]	Write a debugging log	
--debug-check [310]	debug-check [310]	Print debugging information when the program exits	
--debug-info [311]	debug-info [311]	Print debugging information, memory and CPU statistics when the program exits	
--default-auth=plugin [311]	default-auth=plugin [311]	The authentication plugin to use	5.6.2
-- delimiter=str [311]	delimiter [311]	The delimiter to use in SQL statements	
-- detach=# [311]	detach [311]	Detach (close and reopen) each connection after each N statements	
--enable-cleartext-plugin [311]	enable-cleartext-plugin [311]	Enable cleartext authentication plugin	5.6.7
-- engine=engine_name [311]	engine [311]	The storage engine to use for creating the table	
--help [309]		Display help message and exit	
-- host=host_name [311]	host [311]	Connect to the MySQL server on the given host	
-- iterations=# [311]	iterations [311]	The number of times to run the tests	
--login-path=name [219]		Read login path options from .mylogin.cnf	5.6.6
--no-drop [311]	no-drop [311]	Do not drop any schema created during the test run	5.6.3
--number-char-cols=# [311]	number-char-cols [311]	The number of VARCHAR columns to use if --auto-generate-sql is specified	
--number-int-cols=# [311]	number-int-cols [311]	The number of INT columns to use if --auto-generate-sql is specified	
--number-of-queries=# [311]	number-of-queries [311]	Limit each client to approximately this number of queries	
--only-print [311]	only-print [311]	Do not connect to databases. mysqlslap only prints what it would have done	
-- password[=password] [311]	password [311]	The password to use when connecting to the server	
--pipe [312]		On Windows, connect to server using a named pipe	
--plugin-dir=path [312]	plugin-dir=path [312]	The directory where plugins are located	5.6.2
-- port=port_num [312]	port [312]	The TCP/IP port number to use for the connection	
--post-query=value [312]	post-query [312]	The file or string containing the statement to execute after the tests have completed	
--post-system=str [312]	post-system [312]	The string to execute using system() after the tests have completed	
--pre-query=value [312]	pre-query [312]	The file or string containing the statement to execute before running the tests	

Format	Option File	Description	Introduced
<code>--pre-system=str</code> [312]	<code>pre-system</code> [312]	The string to execute using <code>system()</code> before running the tests	
<code>--protocol=type</code> [312]	<code>protocol</code> [312]	The connection protocol to use	
<code>--query=value</code> [312]	<code>query</code> [312]	The file or string containing the SELECT statement to use for retrieving data	
<code>--silent</code> [312]	<code>silent</code> [312]	Silent mode	
<code>--socket=path</code> [313]	<code>socket</code> [313]	For connections to localhost	
<code>--ssl-ca=file_name</code> [313]	<code>ssl-ca</code> [313]	The path to a file that contains a list of trusted SSL CAs	
<code>--ssl-capath=dir_name</code> [313]	<code>ssl-capath</code> [313]	The path to a directory that contains trusted SSL CA certificates in PEM format	
<code>--ssl-cert=file_name</code> [313]	<code>ssl-cert</code> [313]	The name of the SSL certificate file to use for establishing a secure connection	
<code>--ssl-cipher=cipher_list</code> [313]	<code>ssl-cipher</code> [313]	A list of allowable ciphers to use for SSL encryption	
<code>--ssl-crl=file_name</code> [313]	<code>ssl-crl</code> [313]	The path to a file that contains certificate revocation lists	5.6.3
<code>--ssl-crlpath=dir_name</code> [313]	<code>ssl-crlpath</code> [313]	The path to a directory that contains certificate revocation list files	5.6.3
<code>--ssl-key=file_name</code> [313]	<code>ssl-key</code> [313]	The name of the SSL key file to use for establishing a secure connection	
<code>--ssl-verify-server-cert</code> [313]	<code>ssl-verify-server-cert</code> [313]	The server's Common Name value in its certificate is verified against the host name used when connecting to the server	
<code>--user=user_name,</code> [313]	<code>user</code> [313]	MySQL user name to use when connecting to server	
<code>--verbose</code> [313]		Verbose mode	
<code>--version</code> [313]		Display version information and exit	

- `--help` [309], `-?`

Display a help message and exit.

- `--auto-generate-sql` [309], `-a`

Generate SQL statements automatically when they are not supplied in files or using command options.

- `--auto-generate-sql-add-autoincrement` [309]

Add an `AUTO_INCREMENT` column to automatically generated tables.

- `--auto-generate-sql-execute-number=N` [309]

Specify how many queries to generate automatically.

- `--auto-generate-sql-guid-primary` [309]

Add a GUID-based primary key to automatically generated tables.

- `--auto-generate-sql-load-type=type` [309]



Specify the test load type. The permissible values are `read` (scan tables), `write` (insert into tables), `key` (read primary keys), `update` (update primary keys), or `mixed` (half inserts, half scanning selects). The default is `mixed`.

- `--auto-generate-sql-secondary-indexes=N` [310]

Specify how many secondary indexes to add to automatically generated tables. By default, none are added.

- `--auto-generate-sql-unique-query-number=N` [310]

How many different queries to generate for automatic tests. For example, if you run a `key` test that performs 1000 selects, you can use this option with a value of 1000 to run 1000 unique queries, or with a value of 50 to perform 50 different selects. The default is 10.

- `--auto-generate-sql-unique-write-number=N` [310]

How many different queries to generate for `--auto-generate-sql-write-number` [310]. The default is 10.

- `--auto-generate-sql-write-number=N` [310]

How many row inserts to perform on each thread. The default is 100.

- `--commit=N` [310]

How many statements to execute before committing. The default is 0 (no commits are done).

- `--compress` [310], `-C`

Compress all information sent between the client and the server if both support compression.

- `--concurrency=N` [310], `-c N`

The number of clients to simulate when issuing the `SELECT` statement.

- `--create=value` [310]

The file or string containing the statement to use for creating the table.

- `--create-schema=value` [310]

The schema in which to run the tests.



#### Note

If the `--auto-generate-sql` [309] option is also given, `mysqlslap` drops the schema at the end of the test run. To avoid this, use the `--no-drop` [311] option as well.

- `--csv[=file_name]` [310]

Generate output in comma-separated values format. The output goes to the named file, or to the standard output if no file is given.

- `--debug[=debug_options]` [310], `-# [debug_options]`

Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysqlslap.trace'`.

- `--debug-check` [310]

Print some debugging information when the program exits.

- `--debug-info [311], -T`  
Print debugging information and memory and CPU usage statistics when the program exits.
- `--default-auth=plugin [311]`  
The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).  
This option was added in MySQL 5.6.2.
- `--delimiter=str [311], -F str`  
The delimiter to use in SQL statements supplied in files or using command options.
- `--detach=N [311]`  
Detach (close and reopen) each connection after each *N* statements. The default is 0 (connections are not detached).
- `--enable-cleartext-plugin [311]`  
Enable the `mysql_clear_password` cleartext authentication plugin. (See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).) This option was added in MySQL 5.6.7.
- `--engine=engine_name [311], -e engine_name`  
The storage engine to use for creating tables.
- `--host=host_name [311], -h host_name`  
Connect to the MySQL server on the given host.
- `--iterations=N [311], -i N`  
The number of times to run the tests.
- `--no-drop [311]`  
Prevent `mysqlslap` from dropping any schema it creates during the test run. This option was added in MySQL 5.6.3.
- `--number-char-cols=N [311], -x N`  
The number of `VARCHAR` columns to use if `--auto-generate-sql [309]` is specified.
- `--number-int-cols=N [311], -y N`  
The number of `INT` columns to use if `--auto-generate-sql [309]` is specified.
- `--number-of-queries=N [311]`  
Limit each client to approximately this many queries. Query counting takes into account the statement delimiter. For example, if you invoke `mysqlslap` as follows, the `;` delimiter is recognized so that each instance of the query string counts as two queries. As a result, 5 rows (not 10) are inserted.  

```
shell> mysqlslap --delimiter=";" --number-of-queries=10
--query="use test;insert into t values(null)"
```
- `--only-print [311]`  
Do not connect to databases. `mysqlslap` only prints what it would have done.
- `--password[=password] [311], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` [311] or `-p` option on the command line, `mysqlslap` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--pipe` [312], `-W`

On Windows, connect to the server using a named pipe. This option applies only if the server supports named-pipe connections.

- `--plugin-dir=path` [312]

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [311] option is used to specify an authentication plugin but `mysqlslap` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--port=port_num` [312], `-P port_num`

The TCP/IP port number to use for the connection.

- `--post-query=value` [312]

The file or string containing the statement to execute after the tests have completed. This execution is not counted for timing purposes.

- `--shared-memory-base-name=name` [312]

On Windows, the shared-memory name to use, for connections made using shared memory to a local server. This option applies only if the server supports shared-memory connections.

- `--post-system=str` [312]

The string to execute using `system()` after the tests have completed. This execution is not counted for timing purposes.

- `--pre-query=value` [312]

The file or string containing the statement to execute before running the tests. This execution is not counted for timing purposes.

- `--pre-system=str` [312]

The string to execute using `system()` before running the tests. This execution is not counted for timing purposes.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [312]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--query=value` [312], `-q value`

The file or string containing the `SELECT` statement to use for retrieving data.

- `--silent` [312], `-s`

Silent mode. No output.

- `--socket=path` [313], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to connect to the server using SSL and indicate where to find SSL keys and certificates. See [Section 6.3.9.4, “SSL Command Options”](#).

- `--user=user_name` [313], `-u user_name`

The MySQL user name to use when connecting to the server.

- `--verbose` [313], `-v`

Verbose mode. Print more information about what the program does. This option can be used multiple times to increase the amount of information.

- `--version` [313], `-V`

Display version information and exit.

## 4.6. MySQL Administrative and Utility Programs

This section describes administrative programs and programs that perform miscellaneous utility operations.

### 4.6.1. `innochecksum` — Offline InnoDB File Checksum Utility

`innochecksum` prints checksums for `InnoDB` files. This tool reads an `InnoDB` tablespace file, calculates the checksum for each page, compares the calculated checksum to the stored checksum, and reports mismatches, which indicate damaged pages. It was originally developed to speed up verifying the integrity of tablespace files after power outages but can also be used after file copies. Because checksum mismatches will cause `InnoDB` to deliberately shut down a running server, it can be preferable to use this tool rather than waiting for a server in production usage to encounter the damaged pages.

`innochecksum` cannot be used on tablespace files that the server already has open. For such files, you should use `CHECK TABLE` to check tables within the tablespace.

If checksum mismatches are found, you would normally restore the tablespace from backup or start the server and attempt to use `mysqldump` to make a backup of the tables within the tablespace.

Invoke `innochecksum` like this:

```
shell> innochecksum [options] file_name
```

`innochecksum` supports the following options. For options that refer to page numbers, the numbers are zero-based.

- `-c`

Print a count of the number of pages in the file.

- `-d`

Debug mode; prints checksums for each page.

- `-e num`

End at this page number.

- `-p num`  
Check only this page number.
- `-s num`  
Start at this page number.
- `-v`  
Verbose mode; print a progress indicator every five seconds.

## 4.6.2. myisam\_ftdump — Display Full-Text Index information

`myisam_ftdump` displays information about `FULLTEXT` indexes in `MyISAM` tables. It reads the `MyISAM` index file directly, so it must be run on the server host where the table is located. Before using `myisam_ftdump`, be sure to issue a `FLUSH TABLES` statement first if the server is running.

`myisam_ftdump` scans and dumps the entire index, which is not particularly fast. On the other hand, the distribution of words changes infrequently, so it need not be run often.

Invoke `myisam_ftdump` like this:

```
shell> myisam_ftdump [options] tbl_name index_num
```

The `tbl_name` argument should be the name of a `MyISAM` table. You can also specify a table by naming its index file (the file with the `.MYI` suffix). If you do not invoke `myisam_ftdump` in the directory where the table files are located, the table or index file name must be preceded by the path name to the table's database directory. Index numbers begin with 0.

Example: Suppose that the `test` database contains a table named `mytexttable` that has the following definition:

```
CREATE TABLE mytexttable
(
  id   INT NOT NULL,
  txt  TEXT NOT NULL,
  PRIMARY KEY (id),
  FULLTEXT (txt)
) ENGINE=MyISAM;
```

The index on `id` is index 0 and the `FULLTEXT` index on `txt` is index 1. If your working directory is the `test` database directory, invoke `myisam_ftdump` as follows:

```
shell> myisam_ftdump mytexttable 1
```

If the path name to the `test` database directory is `/usr/local/mysql/data/test`, you can also specify the table name argument using that path name. This is useful if you do not invoke `myisam_ftdump` in the database directory:

```
shell> myisam_ftdump /usr/local/mysql/data/test/mytexttable 1
```

You can use `myisam_ftdump` to generate a list of index entries in order of frequency of occurrence like this:

```
shell> myisam_ftdump -c mytexttable 1 | sort -r
```

`myisam_ftdump` supports the following options:

- `--help [314]`, `-h -?`  
Display a help message and exit.
- `--count [314]`, `-c`  
Calculate per-word statistics (counts and global weights).

- `--dump` [315], `-d`  
Dump the index, including data offsets and word weights.
- `--length` [315], `-l`  
Report the length distribution.
- `--stats` [315], `-s`  
Report global index statistics. This is the default operation if no other operation is specified.
- `--verbose` [315], `-v`  
Verbose mode. Print more output about what the program does.

### 4.6.3. myisamchk — MyISAM Table-Maintenance Utility

The `myisamchk` utility gets information about your database tables or checks, repairs, or optimizes them. `myisamchk` works with MyISAM tables (tables that have `.MYD` and `.MYI` files for storing data and indexes).

You can also use the `CHECK TABLE` and `REPAIR TABLE` statements to check and repair MyISAM tables. See [Section 13.7.2.2, “CHECK TABLE Syntax”](#), and [Section 13.7.2.5, “REPAIR TABLE Syntax”](#).

The use of `myisamchk` with partitioned tables is not supported.



#### Caution

It is best to make a backup of a table before performing a table repair operation; under some circumstances the operation might cause data loss. Possible causes include but are not limited to file system errors.

Invoke `myisamchk` like this:

```
shell> myisamchk [options] tbl_name ...
```

The `options` specify what you want `myisamchk` to do. They are described in the following sections. You can also get a list of options by invoking `myisamchk --help`.

With no options, `myisamchk` simply checks your table as the default operation. To get more information or to tell `myisamchk` to take corrective action, specify options as described in the following discussion.

`tbl_name` is the database table you want to check or repair. If you run `myisamchk` somewhere other than in the database directory, you must specify the path to the database directory, because `myisamchk` has no idea where the database is located. In fact, `myisamchk` does not actually care whether the files you are working on are located in a database directory. You can copy the files that correspond to a database table into some other location and perform recovery operations on them there.

You can name several tables on the `myisamchk` command line if you wish. You can also specify a table by naming its index file (the file with the `.MYI` suffix). This enables you to specify all tables in a directory by using the pattern `*.MYI`. For example, if you are in a database directory, you can check all the MyISAM tables in that directory like this:

```
shell> myisamchk *.MYI
```

If you are not in the database directory, you can check all the tables there by specifying the path to the directory:

```
shell> myisamchk /path/to/database_dir/*.MYI
```

You can even check all tables in all databases by specifying a wildcard with the path to the MySQL data directory:

```
shell> myisamchk /path/to/datadir/*/*.MYI
```

The recommended way to quickly check all MyISAM tables is:

```
shell> myisamchk --silent --fast /path/to/datadir/*/*.MYI
```

If you want to check all MyISAM tables and repair any that are corrupted, you can use the following command:

```
shell> myisamchk --silent --force --fast --update-state \
  --key_buffer_size=64M --myisam_sort_buffer_size=64M \
  --read_buffer_size=1M --write_buffer_size=1M \
  /path/to/datadir/*/*.MYI
```

This command assumes that you have more than 64MB free. For more information about memory allocation with `myisamchk`, see [Section 4.6.3.6, “myisamchk Memory Usage”](#).

For additional information about using `myisamchk`, see [Section 7.6, “MyISAM Table Maintenance and Crash Recovery”](#).



### Important

*You must ensure that no other program is using the tables while you are running `myisamchk`.* The most effective means of doing so is to shut down the MySQL server while running `myisamchk`, or to lock all tables that `myisamchk` is being used on.

Otherwise, when you run `myisamchk`, it may display the following error message:

```
warning: clients are using or haven't closed the table properly
```

This means that you are trying to check a table that has been updated by another program (such as the `mysqld` server) that hasn't yet closed the file or that has died without closing the file properly, which can sometimes lead to the corruption of one or more MyISAM tables.

If `mysqld` is running, you must force it to flush any table modifications that are still buffered in memory by using `FLUSH TABLES`. You should then ensure that no one is using the tables while you are running `myisamchk`.

However, the easiest way to avoid this problem is to use `CHECK TABLE` instead of `myisamchk` to check tables. See [Section 13.7.2.2, “CHECK TABLE Syntax”](#).

`myisamchk` supports the following options, which can be specified on the command line or in the `[myisamchk]` group of an option file. `myisamchk` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.10. `myisamchk` Options**

Format	Option File	Description	Introduced	Deprecated
<code>--analyze [323]</code>	<code>analyze [323]</code>	Analyze the distribution of key values		
<code>--backup [321]</code>	<code>backup [321]</code>	Make a backup of the .MYD file as file_name-time.BAK		
<code>--block-search=offset [323]</code>	<code>block-search [323]</code>	Find the record that a block at the given offset belongs to		
<code>--check [320]</code>	<code>check [320]</code>	Check the table for errors		
<code>--check-only-changed [321]</code>	<code>check-only-changed [321]</code>	Check only tables that have changed since the last check		



Format	Option File	Description	Introduced	Deprecated
<a href="#">--correct-checksum [321]</a>	<a href="#">correct-checksum [321]</a>	Correct the checksum information for the table		
<a href="#">--data-file-length=len [322]</a>	<a href="#">data-file-length [322]</a>	Maximum length of the data file (when re-creating data file when it is full)		
<a href="#">--debug[=debug_options] [319]</a>	<a href="#">debug [319]</a>	Write a debugging log		
<a href="#">decode_bits=#</a>	<a href="#">decode_bits</a>	<a href="#">Decode_bits</a>		
<a href="#">--description [323]</a>	<a href="#">description [323]</a>	Print some descriptive information about the table		
<a href="#">--extend-check [321]</a>	<a href="#">extend-check [321]</a>	Do very thorough table check or repair that tries to recover every possible row from the data file		
<a href="#">--fast [321]</a>	<a href="#">fast [321]</a>	Check only tables that haven't been closed properly		
<a href="#">--force [321]</a>	<a href="#">force [321]</a>	Do a repair operation automatically if myisamchk finds any errors in the table		
<a href="#">--force</a>	<a href="#">force-recover</a>	Overwrite old temporary files. For use with the -r or -o option		
<a href="#">ft_max_word_len=#</a>	<a href="#">ft_max_word_len=#</a>	Maximum word length for FULLTEXT indexes		
<a href="#">ft_min_word_len=#</a>	<a href="#">ft_min_word_len=#</a>	Minimum word length for FULLTEXT indexes		
<a href="#">ft_stopword_file=filename</a>	<a href="#">ft_stopword_file=filename</a>	Use stopwords from this file instead of built-in list		
<a href="#">--HELP [319]</a>		Display help message and exit		
<a href="#">--help [318]</a>		Display help message and exit		
<a href="#">--information [321]</a>	<a href="#">information [321]</a>	Print informational statistics about the table that is checked		
<a href="#">key_buffer_size=#</a>	<a href="#">key_buffer_size=#</a>	The size of the buffer used for index blocks for MyISAM tables		
<a href="#">--keys-used=val [322]</a>	<a href="#">keys-used [322]</a>	A bit-value that indicates which indexes to update		
<a href="#">--max-record-length=len [322]</a>	<a href="#">max-record-length [322]</a>	Skip rows larger than the given length if myisamchk cannot allocate memory to hold them		
<a href="#">--medium-check [321]</a>	<a href="#">medium-check [321]</a>	Do a check that is faster than an --extend-check operation		
<a href="#">myisam_block_size=#</a>	<a href="#">myisam_block_size=#</a>	Block size to be used for MyISAM index pages		
<a href="#">myisam_sort_buffer_size=#</a>	<a href="#">myisam_sort_buffer_size=#</a>	The size of the buffer that is allocated when sorting the index when doing a REPAIR or when creating indexes with CREATE INDEX or ALTER TABLE	5.6.9	
<a href="#">--parallel-recover [322]</a>	<a href="#">parallel-recover [322]</a>	Uses the same technique as -r and -n, but creates all the keys in parallel, using different threads (beta)		
<a href="#">--quick [322]</a>	<a href="#">quick [322]</a>	Achieve a faster repair by not modifying the data file.		

Format	Option File	Description	Introduced	Deprecated
read_buffer_size=#	read_buffer_size	Each thread that does a sequential scan allocates a buffer of this size for each table it scans		
--read-only [321]	read-only [321]	Don't mark the table as checked		
--recover [322]	recover [322]	Do a repair that can fix almost any problem except unique keys that aren't unique		
--safe-recover [322]	safe-recover [322]	Do a repair using an old recovery method that reads through all rows in order and updates all index trees based on the rows found		
--set-auto-increment[=value] [323]	set-auto-increment [323]	Force AUTO_INCREMENT numbering for new records to start at the given value		
--set-collation=name [323]	set-collation [323]	Specify the collation to use for sorting table indexes		
--silent [319]	silent [319]	Silent mode		
sort_buffer_size=#	sort_buffer_size	The buffer that is allocated when sorting the index when doing a REPAIR or when creating indexes with CREATE INDEX or ALTER TABLE		5.6.9
--sort-index [323]	sort-index [323]	Sort the index tree blocks in high-low order		
sort_key_blocks=#	sort_key_blocks	sort_key_blocks		
--sort-records=# [324]	sort-records [324]	Sort records according to a particular index		
--sort-recover [323]	sort-recover [323]	Force myisamchk to use sorting to resolve the keys even if the temporary files would be very large		
stats_method=value	stats_method	Specifies how MyISAM index statistics collection code should treat NULLs		
--tmpdir=path [323]	tmpdir [323]	Path of the directory to be used for storing temporary files		
--unpack [323]	unpack [323]	Unpack a table that was packed with myisampack		
--update-state [321]	update-state [321]	Store information in the .MYI file to indicate when the table was checked and whether the table crashed		
--verbose [319]		Verbose mode		
--version [319]		Display version information and exit		
write_buffer_size=#	write_buffer_size	Write buffer size		

#### 4.6.3.1. myisamchk General Options

The options described in this section can be used for any type of table maintenance operation performed by `myisamchk`. The sections following this one describe options that pertain only to specific operations, such as table checking or repairing.

- `--help [318]`, `-?`

Display a help message and exit. Options are grouped by type of operation.

- `--HELP` [319], `-H`

Display a help message and exit. Options are presented in a single list.

- `--debug=debug_options` [319], `-# debug_options`

Write a debugging log. A typical *debug\_options* string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/myisamchk.trace'`.

- `--silent` [319], `-s`

Silent mode. Write output only when errors occur. You can use `-s` twice (`-ss`) to make `myisamchk` very silent.

- `--verbose` [319], `-v`

Verbose mode. Print more information about what the program does. This can be used with `-d` and `-e`. Use `-v` multiple times (`-vv`, `-vvv`) for even more output.

- `--version` [319], `-V`

Display version information and exit.

- `--wait` [319], `-w`

Instead of terminating with an error if the table is locked, wait until the table is unlocked before continuing. If you are running `mysqld` with external locking disabled, the table can be locked only by another `myisamchk` command.

You can also set the following variables by using `--var_name=value` syntax:

Variable	Default Value
<code>decode_bits</code>	9
<code>ft_max_word_len</code>	version-dependent
<code>ft_min_word_len</code>	4
<code>ft_stopword_file</code>	built-in list
<code>key_buffer_size</code>	523264
<code>myisam_block_size</code>	1024
<code>myisam_sort_key_blocks</code>	16
<code>read_buffer_size</code>	262136
<code>sort_buffer_size</code>	2097144
<code>sort_key_blocks</code>	16
<code>stats_method</code>	nulls_unequal
<code>write_buffer_size</code>	262136

The possible `myisamchk` variables and their default values can be examined with `myisamchk --help`:

`sort_buffer_size` is used when the keys are repaired by sorting keys, which is the normal case when you use `--recover` [322]. As of MySQL 5.6.9, `myisam_sort_buffer_size` is available as an alternative name to `sort_buffer_size`. `myisam_sort_buffer_size` is preferable to `sort_buffer_size` because its name corresponds to the `myisam_sort_buffer_size` [513] server system variable that has a similar meaning. `sort_buffer_size` should be considered deprecated.

`key_buffer_size` is used when you are checking the table with `--extend-check` [321] or when the keys are repaired by inserting keys row by row into the table (like when doing normal inserts). Repairing through the key buffer is used in the following cases:

- You use `--safe-recover` [322].
- The temporary files needed to sort the keys would be more than twice as big as when creating the key file directly. This is often the case when you have large key values for `CHAR`, `VARCHAR`, or `TEXT` columns, because the sort operation needs to store the complete key values as it proceeds. If you have lots of temporary space and you can force `myisamchk` to repair by sorting, you can use the `--sort-recover` [323] option.

Repairing through the key buffer takes much less disk space than using sorting, but is also much slower.

If you want a faster repair, set the `key_buffer_size` and `myisam_sort_buffer_size` variables to about 25% of your available memory. You can set both variables to large values, because only one of them is used at a time.

`myisam_block_size` is the size used for index blocks.

`stats_method` influences how `NULL` values are treated for index statistics collection when the `--analyze` [323] option is given. It acts like the `myisam_stats_method` system variable. For more information, see the description of `myisam_stats_method` in Section 5.1.4, “Server System Variables”, and Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”.

`ft_min_word_len` and `ft_max_word_len` indicate the minimum and maximum word length for `FULLTEXT` indexes on `MyISAM` tables. `ft_stopword_file` names the stopwords file. These need to be set under the following circumstances.

If you use `myisamchk` to perform an operation that modifies table indexes (such as repair or analyze), the `FULLTEXT` indexes are rebuilt using the default full-text parameter values for minimum and maximum word length and the stopwords file unless you specify otherwise. This can result in queries failing.

The problem occurs because these parameters are known only by the server. They are not stored in `MyISAM` index files. To avoid the problem if you have modified the minimum or maximum word length or the stopwords file in the server, specify the same `ft_min_word_len`, `ft_max_word_len`, and `ft_stopword_file` values to `myisamchk` that you use for `mysqld`. For example, if you have set the minimum word length to 3, you can repair a table with `myisamchk` like this:

```
shell> myisamchk --recover --ft_min_word_len=3 tbl_name.MYI
```

To ensure that `myisamchk` and the server use the same values for full-text parameters, you can place each one in both the `[mysqld]` and `[myisamchk]` sections of an option file:

```
[mysqld]
ft_min_word_len=3

[myisamchk]
ft_min_word_len=3
```

An alternative to using `myisamchk` is to use the `REPAIR TABLE`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `ALTER TABLE`. These statements are performed by the server, which knows the proper full-text parameter values to use.

### 4.6.3.2. myisamchk Check Options

`myisamchk` supports the following options for table checking operations:

- `--check` [320], `-c`

Check the table for errors. This is the default operation if you specify no option that selects an operation type explicitly.

- `--check-only-changed` [321], `-C`

Check only tables that have changed since the last check.

- `--extend-check` [321], `-e`

Check the table very thoroughly. This is quite slow if the table has many indexes. This option should only be used in extreme cases. Normally, `myisamchk` or `myisamchk --medium-check` should be able to determine whether there are any errors in the table.

If you are using `--extend-check` [321] and have plenty of memory, setting the `key_buffer_size` variable to a large value helps the repair operation run faster.

See also the description of this option under table repair options.

For a description of the output format, see [Section 4.6.3.5, “Obtaining Table Information with myisamchk”](#).

- `--fast` [321], `-F`

Check only tables that haven't been closed properly.

- `--force` [321], `-f`

Do a repair operation automatically if `myisamchk` finds any errors in the table. The repair type is the same as that specified with the `--recover` [322] or `-r` option.

- `--information` [321], `-i`

Print informational statistics about the table that is checked.

- `--medium-check` [321], `-m`

Do a check that is faster than an `--extend-check` [321] operation. This finds only 99.99% of all errors, which should be good enough in most cases.

- `--read-only` [321], `-T`

Do not mark the table as checked. This is useful if you use `myisamchk` to check a table that is in use by some other application that does not use locking, such as `mysqld` when run with external locking disabled.

- `--update-state` [321], `-U`

Store information in the `.MYI` file to indicate when the table was checked and whether the table crashed. This should be used to get full benefit of the `--check-only-changed` [321] option, but you shouldn't use this option if the `mysqld` server is using the table and you are running it with external locking disabled.

### 4.6.3.3. myisamchk Repair Options

`myisamchk` supports the following options for table repair operations (operations performed when an option such as `--recover` [322] or `--safe-recover` [322] is given):

- `--backup` [321], `-B`

Make a backup of the `.MYD` file as `file_name-time.BAK`

- `--character-sets-dir=path` [321]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--correct-checksum` [321]

Correct the checksum information for the table.

- `--data-file-length=len` [322], `-D len`

The maximum length of the data file (when re-creating data file when it is “full”).

- `--extend-check` [321], `-e`

Do a repair that tries to recover every possible row from the data file. Normally, this also finds a lot of garbage rows. Do not use this option unless you are desperate.

See also the description of this option under table checking options.

For a description of the output format, see [Section 4.6.3.5, “Obtaining Table Information with myisamchk”](#).

- `--force` [321], `-f`

Overwrite old intermediate files (files with names like `tbl_name.TMD`) instead of aborting.

- `--keys-used=val` [322], `-k val`

For `myisamchk`, the option value is a bit-value that indicates which indexes to update. Each binary bit of the option value corresponds to a table index, where the first index is bit 0. An option value of 0 disables updates to all indexes, which can be used to get faster inserts. Deactivated indexes can be reactivated by using `myisamchk -r`.

- `--no-symlinks` [322], `-l`

Do not follow symbolic links. Normally `myisamchk` repairs the table that a symlink points to. This option does not exist as of MySQL 4.0 because versions from 4.0 on do not remove symlinks during repair operations.

- `--max-record-length=len` [322]

Skip rows larger than the given length if `myisamchk` cannot allocate memory to hold them.

- `--parallel-recover` [322], `-p`

Use the same technique as `-r` and `-n`, but create all the keys in parallel, using different threads. *This is beta-quality code. Use at your own risk!*

- `--quick` [322], `-q`

Achieve a faster repair by modifying only the index file, not the data file. You can specify this option twice to force `myisamchk` to modify the original data file in case of duplicate keys.

- `--recover` [322], `-r`

Do a repair that can fix almost any problem except unique keys that are not unique (which is an extremely unlikely error with MyISAM tables). If you want to recover a table, this is the option to try first. You should try `--safe-recover` [322] only if `myisamchk` reports that the table cannot be recovered using `--recover` [322]. (In the unlikely case that `--recover` [322] fails, the data file remains intact.)

If you have lots of memory, you should increase the value of `myisam_sort_buffer_size`.

- `--safe-recover` [322], `-o`

Do a repair using an old recovery method that reads through all rows in order and updates all index trees based on the rows found. This is an order of magnitude slower than `--recover` [322], but can handle a couple of very unlikely cases that `--recover` [322] cannot. This recovery method

also uses much less disk space than `--recover` [322]. Normally, you should repair first using `--recover` [322], and then with `--safe-recover` [322] only if `--recover` [322] fails.

If you have lots of memory, you should increase the value of `key_buffer_size`.

- `--set-character-set=name` [323]

Change the character set used by the table indexes. This option was replaced by `--set-collation` [323] in MySQL 5.0.3.

- `--set-collation=name` [323]

Specify the collation to use for sorting table indexes. The character set name is implied by the first part of the collation name.

- `--sort-recover` [323], `-n`

Force `myisamchk` to use sorting to resolve the keys even if the temporary files would be very large.

- `--tmpdir=path` [323], `-t path`

The path of the directory to be used for storing temporary files. If this is not set, `myisamchk` uses the value of the `TMPDIR` environment variable. `--tmpdir` [323] can be set to a list of directory paths that are used successively in round-robin fashion for creating temporary files. The separator character between directory names is the colon (“:”) on Unix and the semicolon (“;”) on Windows.

- `--unpack` [323], `-u`

Unpack a table that was packed with `myisampack`.

#### 4.6.3.4. Other `myisamchk` Options

`myisamchk` supports the following options for actions other than table checks and repairs:

- `--analyze` [323], `-a`

Analyze the distribution of key values. This improves join performance by enabling the join optimizer to better choose the order in which to join the tables and which indexes it should use. To obtain information about the key distribution, use a `myisamchk --description --verbose tbl_name` command or the `SHOW INDEX FROM tbl_name` statement.

- `--block-search=offset` [323], `-b offset`

Find the record that a block at the given offset belongs to.

- `--description` [323], `-d`

Print some descriptive information about the table. Specifying the `--verbose` [319] option once or twice produces additional information. See [Section 4.6.3.5, “Obtaining Table Information with `myisamchk`”](#).

- `--set-auto-increment[=value]` [323], `-A[value]`

Force `AUTO_INCREMENT` numbering for new records to start at the given value (or higher, if there are existing records with `AUTO_INCREMENT` values this large). If `value` is not specified, `AUTO_INCREMENT` numbers for new records begin with the largest value currently in the table, plus one.

- `--sort-index` [323], `-S`

Sort the index tree blocks in high-low order. This optimizes seeks and makes table scans that use indexes faster.



- `--sort-records=N [324], -R N`

Sort records according to a particular index. This makes your data much more localized and may speed up range-based `SELECT` and `ORDER BY` operations that use this index. (The first time you use this option to sort a table, it may be very slow.) To determine a table's index numbers, use `SHOW INDEX`, which displays a table's indexes in the same order that `myisamchk` sees them. Indexes are numbered beginning with 1.

If keys are not packed (`PACK_KEYS=0`), they have the same length, so when `myisamchk` sorts and moves records, it just overwrites record offsets in the index. If keys are packed (`PACK_KEYS=1`), `myisamchk` must unpack key blocks first, then re-create indexes and pack the key blocks again. (In this case, re-creating indexes is faster than updating offsets for each index.)

#### 4.6.3.5. Obtaining Table Information with `myisamchk`

To obtain a description of a `MyISAM` table or statistics about it, use the commands shown here. The output from these commands is explained later in this section.

- `myisamchk -d tbl_name`

Runs `myisamchk` in “describe mode” to produce a description of your table. If you start the MySQL server with external locking disabled, `myisamchk` may report an error for a table that is updated while it runs. However, because `myisamchk` does not change the table in describe mode, there is no risk of destroying data.

- `myisamchk -dv tbl_name`

Adding `-v` runs `myisamchk` in verbose mode so that it produces more information about the table. Adding `-v` a second time produces even more information.

- `myisamchk -eis tbl_name`

Shows only the most important information from a table. This operation is slow because it must read the entire table.

- `myisamchk -eiv tbl_name`

This is like `-eis`, but tells you what is being done.

The `tbl_name` argument can be either the name of a `MyISAM` table or the name of its index file, as described in [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#). Multiple `tbl_name` arguments can be given.

Suppose that a table named `person` has the following structure. (The `MAX_ROWS` table option is included so that in the example output from `myisamchk` shown later, some values are smaller and fit the output format more easily.)

```
CREATE TABLE person
(
  id          INT NOT NULL AUTO_INCREMENT,
  last_name   VARCHAR(20) NOT NULL,
  first_name  VARCHAR(20) NOT NULL,
  birth       DATE,
  death       DATE,
  PRIMARY KEY (id),
  INDEX (last_name, first_name),
  INDEX (birth)
) MAX_ROWS = 1000000;
```

Suppose also that the table has these data and index file sizes:

```
-rw-rw---- 1 mysql mysql 9347072 Aug 19 11:47 person.MYD
-rw-rw---- 1 mysql mysql 6066176 Aug 19 11:47 person.MYI
```

Example of `myisamchk -dvv` output:

```

MyISAM file:      person
Record format:    Packed
Character set:    latin1_swedish_ci (8)
File-version:     1
Creation time:    2009-08-19 16:47:41
Recover time:    2009-08-19 16:47:56
Status:           checked,analyzed,optimized keys
Auto increment key: 1 Last value: 306688
Data records:     306688 Deleted blocks: 0
Datafile parts:   306688 Deleted data: 0
Datafile pointer (bytes): 4 Keyfile pointer (bytes): 3
Datafile length: 9347072 Keyfile length: 6066176
Max datafile length: 4294967294 Max keyfile length: 17179868159
Recordlength:    54

```

## table description:

Key	Start	Len	Index	Type	Rec/key	Root	Blocksize
1	2	4	unique	long	1	99328	1024
2	6	20	multip.	varchar prefix	512	3563520	1024
	27	20		varchar	512		
3	48	3	multip.	uint24 NULL	306688	6065152	1024

Field	Start	Length	Nullpos	Nullbit	Type
1	1	1			
2	2	4			no zeros
3	6	21			varchar
4	27	21			varchar
5	48	3	1	1	no zeros
6	51	3	1	2	no zeros

Explanations for the types of information `myisamchk` produces are given here. “Keyfile” refers to the index file. “Record” and “row” are synonymous, as are “field” and “column.”

The initial part of the table description contains these values:

- `MyISAM file`

Name of the `MyISAM` (index) file.

- `Record format`

The format used to store table rows. The preceding examples use `Fixed length`. Other possible values are `Compressed` and `Packed`. (`Packed` corresponds to what `SHOW TABLE STATUS` reports as `Dynamic`.)

- `Character set`

The table default character set.

- `File-version`

Version of `MyISAM` format. Currently always 1.

- `Creation time`

When the data file was created.

- `Recover time`

When the index/data file was last reconstructed.

- `Status`

Table status flags. Possible values are `crashed`, `open`, `changed`, `analyzed`, `optimized keys`, and `sorted index pages`.

- `Auto increment key, Last value`

The key number associated the table's `AUTO_INCREMENT` column, and the most recently generated value for this column. These fields do not appear if there is no such column.

- `Data records`

The number of rows in the table.

- `Deleted blocks`

How many deleted blocks still have reserved space. You can optimize your table to minimize this space. See [Section 7.6.4, “MyISAM Table Optimization”](#).

- `Datafile parts`

For dynamic-row format, this indicates how many data blocks there are. For an optimized table without fragmented rows, this is the same as `Data records`.

- `Deleted data`

How many bytes of unreclaimed deleted data there are. You can optimize your table to minimize this space. See [Section 7.6.4, “MyISAM Table Optimization”](#).

- `Datafile pointer`

The size of the data file pointer, in bytes. It is usually 2, 3, 4, or 5 bytes. Most tables manage with 2 bytes, but this cannot be controlled from MySQL yet. For fixed tables, this is a row address. For dynamic tables, this is a byte address.

- `Keyfile pointer`

The size of the index file pointer, in bytes. It is usually 1, 2, or 3 bytes. Most tables manage with 2 bytes, but this is calculated automatically by MySQL. It is always a block address.

- `Max datafile length`

How long the table data file can become, in bytes.

- `Max keyfile length`

How long the table index file can become, in bytes.

- `Recordlength`

How much space each row takes, in bytes.

The `table description` part of the output includes a list of all keys in the table. For each key, `myisamchk` displays some low-level information:

- `Key`

This key's number. This value is shown only for the first column of the key. If this value is missing, the line corresponds to the second or later column of a multiple-column key. For the table shown in the example, there are two `table description` lines for the second index. This indicates that it is a multiple-part index with two parts.

- `Start`

Where in the row this portion of the index starts.

- `Len`

How long this portion of the index is. For packed numbers, this should always be the full length of the column. For strings, it may be shorter than the full length of the indexed column, because you can

index a prefix of a string column. The total length of a multiple-part key is the sum of the `Len` values for all key parts.

- `Index`

Whether a key value can exist multiple times in the index. Possible values are `unique` or `multip.` (multiple).

- `Type`

What data type this portion of the index has. This is a `MyISAM` data type with the possible values `packed`, `stripped`, or `empty`.

- `Root`

Address of the root index block.

- `Blocksize`

The size of each index block. By default this is 1024, but the value may be changed at compile time when MySQL is built from source.

- `Rec/key`

This is a statistical value used by the optimizer. It tells how many rows there are per value for this index. A unique index always has a value of 1. This may be updated after a table is loaded (or greatly changed) with `myisamchk -a`. If this is not updated at all, a default value of 30 is given.

The last part of the output provides information about each column:

- `Field`

The column number.

- `Start`

The byte position of the column within table rows.

- `Length`

The length of the column in bytes.

- `Nullpos, Nullbit`

For columns that can be `NULL`, `MyISAM` stores `NULL` values as a flag in a byte. Depending on how many nullable columns there are, there can be one or more bytes used for this purpose. The `Nullpos` and `Nullbit` values, if nonempty, indicate which byte and bit contains that flag indicating whether the column is `NULL`.

The position and number of bytes used to store `NULL` flags is shown in the line for field 1. This is why there are six `Field` lines for the `person` table even though it has only five columns.

- `Type`

The data type. The value may contain any of the following descriptors:

- `constant`

All rows have the same value.

- `no endspace`

Do not store endspace.

- `no endspace, not_always`

Do not store endspace and do not do endspace compression for all values.

- `no endspace, no empty`

Do not store endspace. Do not store empty values.

- `table-lookup`

The column was converted to an `ENUM`.

- `zerofill(N)`

The most significant `N` bytes in the value are always 0 and are not stored.

- `no zeros`

Do not store zeros.

- `always zero`

Zero values are stored using one bit.

- `Huff tree`

The number of the Huffman tree associated with the column.

- `Bits`

The number of bits used in the Huffman tree.

The `Huff tree` and `Bits` fields are displayed if the table has been compressed with `myisampack`. See [Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#), for an example of this information.

Example of `myisamchk -eiv` output:

```
Checking MyISAM file: person
Data records: 306688 Deleted blocks: 0
- check file-size
- check record delete-chain
No recordlinks
- check key delete-chain
block_size 1024:
- check index reference
- check data record references index: 1
Key: 1: Keyblocks used: 98% Packed: 0% Max levels: 3
- check data record references index: 2
Key: 2: Keyblocks used: 99% Packed: 97% Max levels: 3
- check data record references index: 3
Key: 3: Keyblocks used: 98% Packed: -14% Max levels: 3
Total: Keyblocks used: 98% Packed: 89%

- check records and index references*** LOTS OF ROW NUMBERS DELETED ***

Records:          306688 M.recordlength:      25 Packed:          83%
Recordspace used: 97% Empty space:           2% Blocks/Record: 1.00
Record blocks:    306688 Delete blocks:       0
Record data:      7934464 Deleted data:       0
Lost space:       256512 Linkdata:           1156096

User time 43.08, System time 1.68
Maximum resident set size 0, Integral resident set size 0
Non-physical pagefaults 0, Physical pagefaults 0, Swaps 0
Blocks in 0 out 7, Messages in 0 out 0, Signals 0
```

```
Voluntary context switches 0, Involuntary context switches 0
Maximum memory usage: 1046926 bytes (1023k)
```

`myisamchk -eiv` output includes the following information:

- `Data records`

The number of rows in the table.

- `Deleted blocks`

How many deleted blocks still have reserved space. You can optimize your table to minimize this space. See [Section 7.6.4, “MyISAM Table Optimization”](#).

- `Key`

The key number.

- `Keyblocks used`

What percentage of the keyblocks are used. When a table has just been reorganized with `myisamchk`, the values are very high (very near theoretical maximum).

- `Packed`

MySQL tries to pack key values that have a common suffix. This can only be used for indexes on `CHAR` and `VARCHAR` columns. For long indexed strings that have similar leftmost parts, this can significantly reduce the space used. In the preceding example, the second key is 40 bytes long and a 97% reduction in space is achieved.

- `Max levels`

How deep the B-tree for this key is. Large tables with long key values get high values.

- `Records`

How many rows are in the table.

- `M.recordlength`

The average row length. This is the exact row length for tables with fixed-length rows, because all rows have the same length.

- `Packed`

MySQL strips spaces from the end of strings. The `Packed` value indicates the percentage of savings achieved by doing this.

- `Recordspace used`

What percentage of the data file is used.

- `Empty space`

What percentage of the data file is unused.

- `Blocks/Record`

Average number of blocks per row (that is, how many links a fragmented row is composed of). This is always 1.0 for fixed-format tables. This value should stay as close to 1.0 as possible. If it gets too large, you can reorganize the table. See [Section 7.6.4, “MyISAM Table Optimization”](#).

- `Recordblocks`

How many blocks (links) are used. For fixed-format tables, this is the same as the number of rows.

- `Deleteblocks`  
How many blocks (links) are deleted.
- `Recorddata`  
How many bytes in the data file are used.
- `Deleted data`  
How many bytes in the data file are deleted (unused).
- `Lost space`  
If a row is updated to a shorter length, some space is lost. This is the sum of all such losses, in bytes.
- `Linkdata`  
When the dynamic table format is used, row fragments are linked with pointers (4 to 7 bytes each). `Linkdata` is the sum of the amount of storage used by all such pointers.

#### 4.6.3.6. myisamchk Memory Usage

Memory allocation is important when you run `myisamchk`. `myisamchk` uses no more memory than its memory-related variables are set to. If you are going to use `myisamchk` on very large tables, you should first decide how much memory you want it to use. The default is to use only about 3MB to perform repairs. By using larger values, you can get `myisamchk` to operate faster. For example, if you have more than 512MB RAM available, you could use options such as these (in addition to any other options you might specify):

```
shell> myisamchk --myisam_sort_buffer_size=256M \
  --key_buffer_size=512M \
  --read_buffer_size=64M \
  --write_buffer_size=64M ...
```

Using `--myisam_sort_buffer_size=16M` is probably enough for most cases.

Be aware that `myisamchk` uses temporary files in `TMPDIR`. If `TMPDIR` points to a memory file system, out of memory errors can easily occur. If this happens, run `myisamchk` with the `--tmpdir=path` [323] option to specify a directory located on a file system that has more space.

When performing repair operations, `myisamchk` also needs a lot of disk space:

- Twice the size of the data file (the original file and a copy). This space is not needed if you do a repair with `--quick` [322]; in this case, only the index file is re-created. *This space must be available on the same file system as the original data file*, as the copy is created in the same directory as the original.
- Space for the new index file that replaces the old one. The old index file is truncated at the start of the repair operation, so you usually ignore this space. This space must be available on the same file system as the original data file.
- When using `--recover` [322] or `--sort-recover` [323] (but not when using `--safe-recover` [322]), you need space on disk for sorting. This space is allocated in the temporary directory (specified by `TMPDIR` or `--tmpdir=path` [323]). The following formula yields the amount of space required:

```
(largest_key + row_pointer_length) * number_of_rows * 2
```

You can check the length of the keys and the `row_pointer_length` with `myisamchk -dv tbl_name` (see Section 4.6.3.5, “Obtaining Table Information with `myisamchk`”). The `row_pointer_length` and `number_of_rows` values are the `Datafile pointer` and `Data records` values in the table description. To determine the `largest_key` value, check the `Key`



lines in the table description. The `Len` column indicates the number of bytes for each key part. For a multiple-column index, the key size is the sum of the `Len` values for all key parts.

If you have a problem with disk space during repair, you can try `--safe-recover` [322] instead of `--recover` [322].

#### 4.6.4. myisamlog — Display MyISAM Log File Contents

`myisamlog` processes the contents of a MyISAM log file. To create such a file, start the server with a `--log-isam=log_file` [427] option.

Invoke `myisamlog` like this:

```
shell> myisamlog [options] [file_name [tbl_name] ...]
```

The default operation is update (`-u`). If a recovery is done (`-r`), all writes and possibly updates and deletes are done and errors are only counted. The default log file name is `myisam.log` if no `log_file` argument is given. If tables are named on the command line, only those tables are updated.

`myisamlog` supports the following options:

- `-?, -I`  
Display a help message and exit.
- `-c N`  
Execute only `N` commands.
- `-f N`  
Specify the maximum number of open files.
- `-i`  
Display extra information before exiting.
- `-o offset`  
Specify the starting offset.
- `-p N`  
Remove `N` components from path.
- `-r`  
Perform a recovery operation.
- `-R record_pos_file record_pos`  
Specify record position file and record position.
- `-u`  
Perform an update operation.
- `-v`  
Verbose mode. Print more output about what the program does. This option can be given multiple times to produce more and more output.
- `-w write_file`  
Specify the write file.

- `-V`

Display version information.

#### 4.6.5. `myisampack` — Generate Compressed, Read-Only MyISAM Tables

The `myisampack` utility compresses MyISAM tables. `myisampack` works by compressing each column in the table separately. Usually, `myisampack` packs the data file 40% to 70%.

When the table is used later, the server reads into memory the information needed to decompress columns. This results in much better performance when accessing individual rows, because you only have to uncompress exactly one row.

MySQL uses `mmap()` when possible to perform memory mapping on compressed tables. If `mmap()` does not work, MySQL falls back to normal read/write file operations.

Please note the following:

- If the `mysqld` server was invoked with external locking disabled, it is not a good idea to invoke `myisampack` if the table might be updated by the server during the packing process. It is safest to compress tables with the server stopped.
- After packing a table, it becomes read only. This is generally intended (such as when accessing packed tables on a CD).

Invoke `myisampack` like this:

```
shell> myisampack [options] file_name ...
```

Each file name argument should be the name of an index (`.MYI`) file. If you are not in the database directory, you should specify the path name to the file. It is permissible to omit the `.MYI` extension.

After you compress a table with `myisampack`, you should use `myisamchk -rq` to rebuild its indexes. [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

`myisampack` supports the following options. It also reads option files and supports the options for processing them described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

- `--help` [332], `-?`

Display a help message and exit.

- `--backup` [332], `-b`

Make a backup of each table's data file using the name `tbl_name.OLD`.

- `--character-sets-dir=path` [332]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--debug[=debug_options]` [332], `-#` [debug\_options]

Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o'`.

- `--force` [332], `-f`

Produce a packed table even if it becomes larger than the original or if the intermediate file from an earlier invocation of `myisampack` exists. (`myisampack` creates an intermediate file named `tbl_name.TMD` in the database directory while it compresses the table. If you kill `myisampack`, the `.TMD` file might not be deleted.) Normally, `myisampack` exits with an error if it finds that `tbl_name.TMD` exists. With `--force` [332], `myisampack` packs the table anyway.

- `--join=big_tbl_name [333], -j big_tbl_name`

Join all tables named on the command line into a single packed table `big_tbl_name`. All tables that are to be combined *must* have identical structure (same column names and types, same indexes, and so forth).

`big_tbl_name` must not exist prior to the join operation. All source tables named on the command line to be merged into `big_tbl_name` must exist. The source tables are read for the join operation but not modified. The join operation does not create a `.frm` file for `big_tbl_name`, so after the join operation finishes, copy the `.frm` file from one of the source tables and name it `big_tbl_name.frm`.

- `--silent [333], -s`

Silent mode. Write output only when errors occur.

- `--test [333], -t`

Do not actually pack the table, just test packing it.

- `--tmpdir=path [333], -T path`

Use the named directory as the location where `myisampack` creates temporary files.

- `--verbose [333], -v`

Verbose mode. Write information about the progress of the packing operation and its result.

- `--version [333], -V`

Display version information and exit.

- `--wait [333], -w`

Wait and retry if the table is in use. If the `mysqld` server was invoked with external locking disabled, it is not a good idea to invoke `myisampack` if the table might be updated by the server during the packing process.

The following sequence of commands illustrates a typical table compression session:

```
shell> ls -l station.*
-rw-rw-r-- 1 monty my          994128 Apr 17 19:00 station.MYD
-rw-rw-r-- 1 monty my          53248 Apr 17 19:00 station.MYI
-rw-rw-r-- 1 monty my           5767 Apr 17 19:00 station.frm

shell> myisamchk -dvv station

MyISAM file:      station
Isam-version:    2
Creation time:   1996-03-13 10:08:58
Recover time:   1997-02-02 3:06:43
Data records:   1192 Deleted blocks:      0
Datafile parts: 1192 Deleted data:        0
Datafile pointer (bytes): 2 Keyfile pointer (bytes): 2
Max datafile length: 54657023 Max keyfile length: 33554431
Recordlength:   834
Record format:  Fixed length

table description:
Key Start Len Index Type Root Blocksize Rec/key
1 2 4 unique unsigned long 1024 1024 1
2 32 30 multip. text 10240 1024 1

Field Start Length Type
1 1 1
2 2 4
```

```
3      6      4
4     10      1
5     11     20
6     31      1
7     32     30
8     62     35
9     97     35
10    132     35
11    167      4
12    171     16
13    187     35
14    222      4
15    226     16
16    242     20
17    262     20
18    282     20
19    302     30
20    332      4
21    336      4
22    340      1
23    341      8
24    349      8
25    357      8
26    365      2
27    367      2
28    369      4
29    373      4
30    377      1
31    378      2
32    380      8
33    388      4
34    392      4
35    396      4
36    400      4
37    404      1
38    405      4
39    409      4
40    413      4
41    417      4
42    421      4
43    425      4
44    429     20
45    449     30
46    479      1
47    480      1
48    481     79
49    560     79
50    639     79
51    718     79
52    797      8
53    805      1
54    806      1
55    807     20
56    827      4
57    831      4

shell> myisampack station.MYI
Compressing station.MYI: (1192 records)
- Calculating statistics

normal:      20  empty-space:   16  empty-zero:    12  empty-fill:   11
pre-space:   0  end-space:    12  table-lookups: 5  zero:         7
Original trees: 57 After join: 17
- Compressing file
87.14%
Remember to run myisamchk -rq on compressed tables

shell> ls -l station.*
-rw-rw-r-- 1 monty my      127874 Apr 17 19:00 station.MYD
-rw-rw-r-- 1 monty my      55296 Apr 17 19:04 station.MYI
-rw-rw-r-- 1 monty my       5767 Apr 17 19:00 station.frm
```

```
shell> myisamchk -dvv station
```

```
MyISAM file:      station
Isam-version:    2
Creation time:   1996-03-13 10:08:58
Recover time:   1997-04-17 19:04:26
Data records:   1192 Deleted blocks:      0
Datafile parts: 1192 Deleted data:      0
Datafile pointer (bytes): 3 Keyfile pointer (bytes): 1
Max datafile length: 16777215 Max keyfile length: 131071
Recordlength:   834
Record format:  Compressed
```

table description:

Key	Start	Len	Index	Type	Root	Blocksize	Rec/key
1	2	4	unique	unsigned long	10240	1024	1
2	32	30	multip.	text	54272	1024	1

Field	Start	Length	Type	Huff tree	Bits
1	1	1	constant	1	0
2	2	4	zerofill(1)	2	9
3	6	4	no zeros, zerofill(1)	2	9
4	10	1		3	9
5	11	20	table-lookup	4	0
6	31	1		3	9
7	32	30	no endspace, not_always	5	9
8	62	35	no endspace, not_always, no empty	6	9
9	97	35	no empty	7	9
10	132	35	no endspace, not_always, no empty	6	9
11	167	4	zerofill(1)	2	9
12	171	16	no endspace, not_always, no empty	5	9
13	187	35	no endspace, not_always, no empty	6	9
14	222	4	zerofill(1)	2	9
15	226	16	no endspace, not_always, no empty	5	9
16	242	20	no endspace, not_always	8	9
17	262	20	no endspace, no empty	8	9
18	282	20	no endspace, no empty	5	9
19	302	30	no endspace, no empty	6	9
20	332	4	always zero	2	9
21	336	4	always zero	2	9
22	340	1		3	9
23	341	8	table-lookup	9	0
24	349	8	table-lookup	10	0
25	357	8	always zero	2	9
26	365	2		2	9
27	367	2	no zeros, zerofill(1)	2	9
28	369	4	no zeros, zerofill(1)	2	9
29	373	4	table-lookup	11	0
30	377	1		3	9
31	378	2	no zeros, zerofill(1)	2	9
32	380	8	no zeros	2	9
33	388	4	always zero	2	9
34	392	4	table-lookup	12	0
35	396	4	no zeros, zerofill(1)	13	9
36	400	4	no zeros, zerofill(1)	2	9
37	404	1		2	9
38	405	4	no zeros	2	9
39	409	4	always zero	2	9
40	413	4	no zeros	2	9
41	417	4	always zero	2	9
42	421	4	no zeros	2	9
43	425	4	always zero	2	9
44	429	20	no empty	3	9
45	449	30	no empty	3	9
46	479	1		14	4
47	480	1		14	4
48	481	79	no endspace, no empty	15	9
49	560	79	no empty	2	9
50	639	79	no empty	2	9
51	718	79	no endspace	16	9
52	797	8	no empty	2	9
53	805	1		17	1

54	806	1		3	9
55	807	20	no empty	3	9
56	827	4	no zeros, zerofill(2)	2	9
57	831	4	no zeros, zerofill(1)	2	9

myisampack displays the following kinds of information:

- `normal`

The number of columns for which no extra packing is used.

- `empty-space`

The number of columns containing values that are only spaces. These occupy one bit.

- `empty-zero`

The number of columns containing values that are only binary zeros. These occupy one bit.

- `empty-fill`

The number of integer columns that do not occupy the full byte range of their type. These are changed to a smaller type. For example, a `BIGINT` column (eight bytes) can be stored as a `TINYINT` column (one byte) if all its values are in the range from `-128` to `127`.

- `pre-space`

The number of decimal columns that are stored with leading spaces. In this case, each value contains a count for the number of leading spaces.

- `end-space`

The number of columns that have a lot of trailing spaces. In this case, each value contains a count for the number of trailing spaces.

- `table-lookup`

The column had only a small number of different values, which were converted to an `ENUM` before Huffman compression.

- `zero`

The number of columns for which all values are zero.

- `Original trees`

The initial number of Huffman trees.

- `After join`

The number of distinct Huffman trees left after joining trees to save some header space.

After a table has been compressed, the `Field` lines displayed by `myisamchk -dvv` include additional information about each column:

- `Type`

The data type. The value may contain any of the following descriptors:

- `constant`

All rows have the same value.

- `no endspace`

Do not store endspace.

- `no endspace, not_always`

Do not store endspace and do not do endspace compression for all values.

- `no endspace, no empty`

Do not store endspace. Do not store empty values.

- `table-lookup`

The column was converted to an `ENUM`.

- `zerofill(N)`

The most significant *N* bytes in the value are always 0 and are not stored.

- `no zeros`

Do not store zeros.

- `always zero`

Zero values are stored using one bit.

- `Huff tree`

The number of the Huffman tree associated with the column.

- `Bits`

The number of bits used in the Huffman tree.

After you run `myisampack`, you must run `myisamchk` to re-create any indexes. At this time, you can also sort the index blocks and create statistics needed for the MySQL optimizer to work more efficiently:

```
shell> myisamchk -rq --sort-index --analyze tbl_name.MYI
```

After you have installed the packed table into the MySQL database directory, you should execute `mysqladmin flush-tables` to force `mysqld` to start using the new table.

To unpack a packed table, use the `--unpack [323]` option to `myisamchk`.

## 4.6.6. mysql\_config\_editor — MySQL Configuration Utility

The `mysql_config_editor` utility (available as of MySQL 5.6.6) enables you to store authentication credentials in an encrypted login file named `.mylogin.cnf`. The file location is the `%APPDATA%\MySQL` directory on Windows and the current user's home directory on non-Windows systems. The file can be read later by MySQL client programs to obtain authentication credentials for connecting to MySQL Server.

To specify an alternate file name, set the `MYSQL_TEST_LOGIN_FILE` environment variable. This variable is used by the `mysql-test-run.pl` testing utility, but also is recognized by `mysql_config_editor` and by MySQL clients such as `mysql`, `mysqladmin`, and so forth.

`mysql_config_editor` encrypts the `.mylogin.cnf` file so it cannot be read as clear text, and its contents when decrypted by client programs are used only in memory. In this way, passwords can be stored in a file in non-clear-text format and used later without ever needing to be exposed on the command line or in an environment variable. `mysql_config_editor` provides a `print` command



that enables the user to display the file contents, but even in this case, password values are masked so as never to appear in a way that other users can see them.

The encryption used by `mysql_config_editor` prevents passwords from appearing in `.mylogin.cnf` as clear text and provides a measure of security by preventing inadvertent password exposure. For example, if you display a regular unencrypted `my.cnf` option file on the screen, any passwords it contains are visible for anyone to see. With `.mylogin.cnf`, that is not true. But the encryption used will not deter a determined attacker and you should not consider it unbreakable. A user who can gain system administration privileges on your machine to access your files could decrypt the `.mylogin.cnf` file with some effort.

The login file must be readable and writable to the current user, and inaccessible to other users. Otherwise, `mysql_config_editor` ignores it, and the file is not used by client programs, either. On Windows, this constraint does not apply; instead, the user must have access to the `%APPDATA%\MySQL` directory.

The unencrypted format of the `.mylogin.cnf` login file consists of option groups, similar to other option files. Each option group in `.mylogin.cnf` is called a “login path,” which is a group that permits only a limited set of options: `host`, `user`, and `password`. Think of a login path as a set of values that indicate the server host and the credentials for authenticating with the server. Here is an example:

```
[myloginpath]
user = myname
password = mypass
host = 127.0.0.1
```

When you invoke a client program to connect to the server, `.mylogin.cnf` is used in conjunction with other option files. Its precedence is higher than other option files, but less than options specified explicitly on the client command line. For information about the order in which option files are used, see [Section 4.2.3.3, “Using Option Files”](#).

Invoke `mysql_config_editor` like this:

```
shell> mysql_config_editor [program_options] command [command_options]
```

`program_options` consists of general `mysql_config_editor` options. `command` indicates what command to perform, and `command_options` indicates any additional options needed by the command.

The command indicates what action to perform on the `.mylogin.cnf` login file. For example, `set` writes a login path to the file, `remove` removes a login path, and `print` displays login path contents. Any options given provide information to the command, such as the login path name and the values to use in the login path.

The position of the command name within the set of program arguments is significant. For example, these command lines have the same arguments, but produce different results:

```
mysql_config_editor --help set
mysql_config_editor set --help
```

The first command line displays general `mysql_config_editor` help, and ignores the `set` command. The second command line displays help for the `set` command.

Suppose that you want to establish two login paths named `local` and `remote` for connecting to the local MySQL server and a server on the host `remote.example.com`. You want to authenticate to the local server with a user name and password of `localuser` and `localpass`, and to the remote server with a user name and password of `remoteuser` and `remotepass`. To set up the login paths in the `.mylogin.cnf` file, use the following `set` commands. Enter each command on a single line, then enter the appropriate password when prompted.

```
shell> mysql_config_editor set --login-path=local
--host=localhost --user=localuser --password
Enter password: enter password "localpass" here
shell> mysql_config_editor set --login-path=remote
```

```
--host=remote.example.com --user=remoteuser --password
Enter password: enter password "remotepass" here
```

To see what `mysql_config_editor` wrote to the `.mylogin.cnf` file, use the `print` command:

```
shell> mysql_config_editor print --all
[local]
user = localuser
password = *****
host = localhost
[remote]
user = remoteuser
password = *****
host = remote.example.com
```

The `print` command displays each login path as a set of lines beginning with a group header indicating the login path name in square brackets, followed by the option values for the login path. Password values are masked and do not appear as clear text.

As shown by the preceding examples, the `.mylogin.cnf` file can contain multiple login paths. In this way, `mysql_config_editor` makes it easy to set up multiple “personalities” for connecting to different MySQL servers. Any of these can be selected by name later using the `--login-path` option when you invoke a client program. For example, to connect to the local server, use this command:

```
shell> mysql --login-path=local
```

To connect to the remote server, use this command:

```
shell> mysql --login-path=remote
```

When you use the `set` command with `mysql_config_editor` to create a login path, you need not specify all three possible option values (host name, user name, and password). Only those values given are written to the path. Any missing values required later can be specified when you invoke a client path to connect to the MySQL server, either in other option files or on the command line. Also, any options specified on the command line override those in option files, including the `.mylogin.cnf` file. For example, if the credentials in the `remote` login path also apply for the host `remote2.example.com`, you can connect to the server on that host like this:

```
shell> mysql --login-path=remote --host=remote2.example.com
```

The `.mylogin.cnf` file, if it exists, is read in all cases, even when the `--no-defaults` [220] option is used. This permits passwords to be specified in a safer way than on the command line even if `--no-defaults` [220] is present.

## mysql\_config\_editor Commands

This section describes the permitted `mysql_config_editor` commands, and the interpretation of options that have a command-specific meaning. In addition, `mysql_config_editor` takes other options that can be used with any command, such as `--verbose` to produce more information as `mysql_config_editor` executes. This option may be helpful in diagnosing problems if an operation does not have the effect you expect. For a list of supported options, see [mysql\\_config\\_editor Options](#).

`mysql_config_editor` supports these commands:

- `help`

Display a help message and exit.

- `print [options]`

Print the contents of `.mylogin.cnf` in unencrypted form. Passwords are displayed as `*****`.

The `print` command takes these options:

- `--all` [342]

Print all login paths.

- `--login-path=name` [342]

Print the named login path.

If no login path is specified, the default path name is `client`. If both `--all` [342] and `--login-path` [342] are given, `--all` [342] takes precedence.

- `remove [options]`

Remove a login path from the `.mylogin.cnf` file.

The `remove` command takes these options:

- `--host` [342]

Remove the host name from the login path.

- `--login-path=name` [342]

The login path to remove. If this option is not given, the default path name is `client`.

- `--password` [342]

Remove the password from the login path.

- `--port` [342]

Remove the TCP/IP port number from the login path.

- `--socket` [343]

Remove the Unix socket file name from the login path.

- `--user` [343]

Remove the user name from the login path.

The `--host`, `--user`, and `--password` options are supported for the `remove` command as of MySQL 5.6.9. The `--port` [342] and `--socket` [343] options are supported for the `remove` command as of MySQL 5.6.11

The `remove` command removes from the login path only such values as are specified with the `--host` [342], `--password` [342], `--port` [342], `--socket` [343], and `--user` [343] options. If none of them is given, `remove` removes the entire login path. For example, this command removes only the `user` value from the `client` login path rather than the entire `client` login path:

```
mysql_config_editor remove --login-path=client --user
```

- `reset`

Empty the contents of the `.mylogin.cnf` file. The file is created if it does not exist.

- `set [options]`

Write a login path to the `.mylogin.cnf` file.

The `set` command takes these options:

- `--host=host_name` [342]

The host name to write to the login path.

- `--login-path=name` [342]

The login path to create. If this option is not given, the default path name is `client`.

- `--password` [342]

Prompt for a password to write to the login path.

- `--port=port_num` [342]

The TCP/IP port number to write to the login path.

- `--socket=file_name` [343]

The Unix socket file to write to the login path.

- `--user=user_name` [343]

The user name to write to the login path.

The `--port` [342] and `--socket` [343] options are supported for the `set` command as of MySQL 5.6.11

The `set` command writes to the login path only such values as are specified with the `--host` [342], `--password` [342], `--port` [342], `--socket` [343], and `--user` [343] options. If none of those options are given, `mysql_config_editor` writes the login path as an empty group.

To specify an empty password, use the `set` command with the `--password` [342] option, then press Enter at the password prompt. The resulting login path written to `.mylogin.cnf` will include a line like this:

```
password =
```

If the login path already exists in `.mylogin.cnf`, the `set` command replaces it. To ensure that this is what the user wants, `mysql_config_editor` prints a warning and prompts for confirmation. To suppress the warning and prompt, use the `--skip-warn` [343] option.

## mysql\_config\_editor Options

`mysql_config_editor` supports the following options.

**Table 4.11. mysql\_config\_editor Options**

Format	Option File	Description	Introduced
<code>--all</code> [342]		Print all login paths	
<code>--debug[=debug_options]</code> [342]		Write a debugging log	
<code>--help</code> [342]		Display help message and exit	
<code>--host=host_name</code> [342]		Host to write to login file	
<code>--login-path=name</code> [342]		Login path name	
<code>--password</code> [342]		Solicit password to write to login file	

Format	Option File	Description	Introduced
<code>--port=port_num [342]</code>	<code>port [342]</code>	The TCP/IP port number to write to login file	5.6.11
<code>--socket=path [343]</code>	<code>socket [343]</code>	The Unix socket file name to write to login file	5.6.11
<code>--user=user_name [343]</code>		User name to write to login file	
<code>--verbose [343]</code>		Verbose mode	
<code>--version [343]</code>		Display version information and exit	
<code>--warn [343]</code>		Warn and solicit confirmation for overwriting login path	

- `--help [342], -?`  
Display a help message and exit. If preceded by a command name such as `set` or `remove`, displays information about that command.
- `--all [342]`  
For the `print` command, print all login paths in the login file.
- `--debug[=debug_options] [342], -# debug_options`  
Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysql_config_editor.trace'`.
- `--host=host_name [342], -h host_name`  
For the `set` command, the host name to write to to the login path. For the `remove` command, removes the host name from the login path.
- `--login-path=name [342], -G name`  
For the `print`, `remove`, and `set` commands, the login path to use in the `.mylogin.cnf` login file.  
  
Client programs also support the `--login-path` option, to enable users to specify which login path to use for connecting to a MySQL server. For client programs, `--login-path` must be the first option given, which is not true for `mysql_config_editor`. See [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).
- `--password [342], -p`  
For the `set` command, cause `mysql_config_editor` to prompt for a password and write the value entered by the user to the login path. After `mysql_config_editor` starts and displays the prompt, the user should type the password and press Enter. To prevent other users from seeing the password, `mysql_config_editor` does not echo it.  
  
This option does not permit a password value following the option name. That is, with `mysql_config_editor`, you never enter a password on the command line where it might be seen by other users. This differs from most other MySQL programs, which permit the password to be given on the command line as `--password=pass_val` or `-ppass_val`. (That practice is insecure and should be avoided, however.)  
  
For the `remove` command, removes the password from the login path.
- `--port=port_num [342], -P port_num`  
For the `set` command, the TCP/IP port number to write to the login path. For the `remove` command, removes the port number from the login path.

- `--socket=file_name` [343], `-S file_name`

For the `set` command, the Unix socket file name to write to the login path. For the `remove` command, removes the socket file from the login path.

- `--user=user_name` [343], `-u user_name`

For the `set` command, the user name to write to the login path. For the `remove` command, removes the user name from the login path.

- `--verbose` [343], `-v`

Verbose mode. Print more information about what the program does.

- `--version` [343], `-V`

Display version information and exit.

- `--warn` [343], `-w`

For the `set` command, warn and prompt the user for confirmation if the command attempts to overwrite an existing login path. This option is enabled by default; use `--skip-warn` [343] to disable it.

## 4.6.7. mysqlaccess — Client for Checking Access Privileges

`mysqlaccess` is a diagnostic tool that Yves Carlier has provided for the MySQL distribution. It checks the access privileges for a host name, user name, and database combination. Note that `mysqlaccess` checks access using only the `user` and `db` tables. It does not check table, column, or routine privileges specified in the `tables_priv`, `columns_priv`, or `procs_priv` tables.

Invoke `mysqlaccess` like this:

```
shell> mysqlaccess [host_name [user_name [db_name]]] [options]
```

`mysqlaccess` supports the following options.

**Table 4.12. mysqlaccess Options**

Format	Option File	Description
<code>--brief</code> [344]	<code>brief</code> [344]	Generate reports in single-line tabular format
<code>--commit</code> [344]	<code>commit</code> [344]	Copy the new access privileges from the temporary tables to the original grant tables
<code>--copy</code> [344]	<code>copy</code> [344]	Reload the temporary grant tables from original ones
<code>--db=db_name</code> [344]	<code>db</code> [344]	Specify the database name
<code>--debug=#</code> [344]	<code>debug</code> [344]	Specify the debug level
<code>--help</code> [344]		Display help message and exit
<code>--host=host_name</code> [344]	<code>host</code> [344]	Connect to the MySQL server on the given host
<code>--howto</code> [344]	<code>howto</code> [344]	Display some examples that show how to use <code>mysqlaccess</code>
<code>--old_server</code> [344]	<code>old_server</code> [344]	Assume that the server is an old MySQL server (prior to MySQL 3.21)
<code>--password[=password]</code> [344]	<code>password</code> [344]	The password to use when connecting to the server

Format	Option File	Description
<code>--plan [345]</code>	<code>plan [345]</code>	Display suggestions and ideas for future releases
<code>--preview [345]</code>	<code>preview [345]</code>	Show the privilege differences after making changes to the temporary grant tables
<code>--relnotes [345]</code>	<code>relnotes [345]</code>	Display the release notes
<code>--rhost=host_name [345]</code>	<code>rhost [345]</code>	Connect to the MySQL server on the given host
<code>--rollback [345]</code>	<code>rollback [345]</code>	Undo the most recent changes to the temporary grant tables.
<code>--spassword[=password] [345]</code>	<code>spassword [345]</code>	The password to use when connecting to the server as the superuser
<code>--superuser=user_name [345]</code>	<code>superuser [345]</code>	Specify the user name for connecting as the superuser
<code>--table [345]</code>	<code>table [345]</code>	Generate reports in table format
<code>--user=user_name, [345]</code>	<code>user [345]</code>	MySQL user name to use when connecting to server
<code>--version [345]</code>		Display version information and exit

- `--help [344], -?`  
Display a help message and exit.
- `--brief [344], -b`  
Generate reports in single-line tabular format.
- `--commit [344]`  
Copy the new access privileges from the temporary tables to the original grant tables. The grant tables must be flushed for the new privileges to take effect. (For example, execute a `mysqladmin reload` command.)
- `--copy [344]`  
Reload the temporary grant tables from original ones.
- `--db=db_name [344], -d db_name`  
Specify the database name.
- `--debug=N [344]`  
Specify the debug level. *N* can be an integer from 0 to 3.
- `--host=host_name [344], -h host_name`  
The host name to use in the access privileges.
- `--howto [344]`  
Display some examples that show how to use `mysqlaccess`.
- `--old_server [344]`  
Assume that the server is an old MySQL server (before MySQL 3.21) that does not yet know how to handle full `WHERE` clauses.
- `--password[=password] [344], -p[password]`



The password to use when connecting to the server. If you omit the *password* value following the `--password [344]` or `-p` option on the command line, `mysqlaccess` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#).

- `--plan [345]`

Display suggestions and ideas for future releases.

- `--preview [345]`

Show the privilege differences after making changes to the temporary grant tables.

- `--relnotes [345]`

Display the release notes.

- `--rhost=host_name [345], -H host_name`

Connect to the MySQL server on the given host.

- `--rollback [345]`

Undo the most recent changes to the temporary grant tables.

- `--spassword[=password] [345], -P[password]`

The password to use when connecting to the server as the superuser. If you omit the *password* value following the `--spassword [345]` or `-p` option on the command line, `mysqlaccess` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#).

- `--superuser=user_name [345], -U user_name`

Specify the user name for connecting as the superuser.

- `--table [345], -t`

Generate reports in table format.

- `--user=user_name [345], -u user_name`

The user name to use in the access privileges.

- `--version [345], -v`

Display version information and exit.

If your MySQL distribution is installed in some nonstandard location, you must change the location where `mysqlaccess` expects to find the `mysql` client. Edit the `mysqlaccess` script at approximately line 18. Search for a line that looks like this:

```
$MYSQL = '/usr/local/bin/mysql'; # path to mysql executable
```

Change the path to reflect the location where `mysql` actually is stored on your system. If you do not do this, a [Broken pipe](#) error will occur when you run `mysqlaccess`.

## 4.6.8. mysqlbinlog — Utility for Processing Binary Log Files

The server's binary log consists of files containing “events” that describe modifications to database contents. The server writes these files in binary format. To display their contents in text format, use the `mysqlbinlog` utility. You can also use `mysqlbinlog` to display the contents of relay log files written by a slave server in a replication setup because relay logs have the same format as binary logs. The binary log and relay log are discussed further in [Section 5.2.4, “The Binary Log”](#), and [Section 16.2.2, “Replication Relay and Status Logs”](#).

Invoke `mysqlbinlog` like this:

```
shell> mysqlbinlog [options] log_file ...
```

For example, to display the contents of the binary log file named `binlog.000003`, use this command:

```
shell> mysqlbinlog binlog.000003
```

The output includes events contained in `binlog.000003`. For statement-based logging, event information includes the SQL statement, the ID of the server on which it was executed, the timestamp when the statement was executed, how much time it took, and so forth. For row-based logging, the event indicates a row change rather than an SQL statement. See [Section 16.1.2, “Replication Formats”](#), for information about logging modes.

Events are preceded by header comments that provide additional information. For example:

```
# at 141
#100309 9:28:36 server id 123  end_log_pos 245
  Query thread_id=3350  exec_time=11  error_code=0
```

In the first line, the number following `at` indicates the starting position of the event in the binary log file.

The second line starts with a date and time indicating when the statement started on the server where the event originated. For replication, this timestamp is propagated to slave servers. `server id` is the `server_id` [539] value of the server where the event originated. `end_log_pos` indicates where the next event starts (that is, it is the end position of the current event + 1). `thread_id` indicates which thread executed the event. `exec_time` is the time spent executing the event, on a master server. On a slave, it is the difference of the end execution time on the slave minus the beginning execution time on the master. The difference serves as an indicator of how much replication lags behind the master. `error_code` indicates the result from executing the event. Zero means that no error occurred.

The output from `mysqlbinlog` can be re-executed (for example, by using it as input to `mysql`) to redo the statements in the log. This is useful for recovery operations after a server crash. For other usage examples, see the discussion later in this section and in [Section 7.5, “Point-in-Time \(Incremental\) Recovery Using the Binary Log”](#).

Normally, you use `mysqlbinlog` to read binary log files directly and apply them to the local MySQL server. It is also possible to read binary logs from a remote server by using the `--read-from-remote-server` [352] option. To read remote binary logs, the connection parameter options can be given to indicate how to connect to the server. These options are `--host` [351], `--password` [351], `--port` [352], `--protocol` [352], `--socket` [353], and `--user` [354]; they are ignored except when you also use the `--read-from-remote-server` [352] option.

`mysqlbinlog` supports the following options, which can be specified on the command line or in the `[mysqlbinlog]` and `[client]` groups of an option file. `mysqlbinlog` also supports the options for processing option files described at [Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”](#).

**Table 4.13. `mysqlbinlog` Options**

Format	Option File	Description	Introduced
<code>--base64-output=value</code> [348]	<code>base64-output</code> [348]	Print binary log entries using base-64 encoding	
<code>--bind-address=ip_address</code> [349]	<code>bind-address</code> [349]	Use the specified network interface to connect to the MySQL Server	

Format	Option File	Description	Introduced
<code>--binlog-row-event-max-size=#</code> [349]	<code>binlog-row-event-max-size</code> [349]	Binary log max event size	
<code>--character-sets-dir=path</code> [349]	<code>character-sets-dir</code> [349]	The directory where character sets are installed	
<code>--database=db_name</code> [349]	<code>database</code> [349]	List entries for just this database	
<code>--debug[=debug_options]</code> [350]	<code>debug</code> [350]	Write a debugging log	
<code>--debug-check</code> [350]	<code>debug-check</code> [350]	Print debugging information when the program exits	
<code>--debug-info</code> [350]	<code>debug-info</code> [350]	Print debugging information, memory and CPU statistics when the program exits	
<code>--default-auth=plugin</code> [351]	<code>default-auth=plugin</code> [351]	The authentication plugin to use	5.6.2
<code>--disable-log-bin</code> [351]	<code>disable-log-bin</code> [351]	Disable binary logging	
<code>--exclude-gtids=gtid_set</code> [351]	<code>exclude-gtids</code> [351]	Do not show any of the groups in the GTID set provided	5.6.5
<code>--force-if-open</code> [351]	<code>force-if-open</code> [351]	Read binary log files even if open or not closed properly	
<code>--force-read</code> [351]	<code>force-read</code> [351]	If mysqlbinlog reads a binary log event that it does not recognize, it prints a warning	
<code>--help</code> [348]		Display help message and exit	
<code>--hexdump</code> [351]	<code>hexdump</code> [351]	Display a hex dump of the log in comments	
<code>--host=host_name</code> [351]	<code>host</code> [351]	Connect to the MySQL server on the given host	
<code>--idempotent</code>	<code>idempotent</code>	Cause the server to use idempotent mode while processing binary log updates from this session only	
<code>--include-gtids=gtid_set</code> [351]	<code>include-gtids</code> [351]	Show only the groups in the GTID set provided	5.6.5
<code>--local-load=path</code> [351]	<code>local-load</code> [351]	Prepare local temporary files for LOAD DATA INFILE in the specified directory	
<code>--login-path=name</code> [219]		Read login path options from <code>.mylogin.cnf</code>	5.6.6
<code>--offset=#</code> [351]	<code>offset</code> [351]	Skip the first N entries in the log	
<code>--password[=password]</code> [351]	<code>password</code> [351]	The password to use when connecting to the server	
<code>--plugin-dir=path</code> [352]	<code>plugin-dir=path</code> [352]	The directory where plugins are located	5.6.2
<code>--port=port_num</code> [352]	<code>port</code> [352]	The TCP/IP port number to use for the connection	
<code>--protocol=type</code> [352]	<code>protocol</code> [352]	The connection protocol to use	
<code>--raw</code> [352]	<code>raw</code> [352]	Write events in raw (binary) format to output files	

Format	Option File	Description	Introduced
<code>--read-from-remote-master=type</code> [352]	<code>read-from-remote-master</code> [352]	Read the binary log from a MySQL master rather than reading a local log file	5.6.5
<code>--read-from-remote-server</code> [352]	<code>read-from-remote-server</code> [352]	Read binary log from MySQL server rather than local log file	
<code>--result-file=name</code> [353]	<code>result-file</code> [353]	Direct output to the given file	
<code>--server-id=id</code> [353]	<code>server-id</code> [353]	Extract only those events created by the server having the given server ID	
<code>--set-charset=charset_name</code> [353]	<code>set-charset</code> [353]	Add a SET NAMES charset_name statement to the output	
<code>--short-form</code> [353]	<code>short-form</code> [353]	Display only the statements contained in the log	
<code>--skip-gtids[=true false]</code> [353]	<code>skip-gtids</code> [353]	Do not show any GTIDs; not recommended in production	5.6.5
<code>--socket=path</code> [353]	<code>socket</code> [353]	For connections to localhost	
<code>--start-datetime=datetime</code> [353]	<code>start-datetime</code> [353]	Read binary log from first event with timestamp equal to or later than datetime argument	
<code>--start-position=#</code> [353]	<code>start-position</code> [353]	Read binary log from first event with position equal to or greater than argument	
<code>--stop-datetime=datetime</code> [353]	<code>stop-datetime</code> [353]	Stop reading binary log at first event with timestamp equal to or greater than datetime argument	
<code>--stop-never</code> [354]	<code>stop-never</code> [354]	Stay connected to server after reading last binary log file	
<code>--stop-never-slave-server-id=#</code> [354]	<code>stop-never-slave-server-id</code> [354]	Slave server ID to report when connecting to server	
<code>--stop-position=#</code> [354]	<code>stop-position</code> [354]	Stop reading binary log at first event with position equal to or greater than argument	
<code>--to-last-log</code> [354]	<code>to-last-log</code> [354]	Do not stop at the end of requested binary log from a MySQL server, but rather continue printing to end of last binary log	
<code>--user=user_name,</code> [354]	<code>user</code> [354]	MySQL user name to use when connecting to server	
<code>--verbose</code> [354]		Reconstruct row events as SQL statements	
<code>--verify-binlog-checksum</code> [354]		Verify checksums in binary log	5.6.1
<code>--version</code> [354]		Display version information and exit	

- `--help` [348], `-?`

Display a help message and exit.

- `--base64-output=value` [348]

This option determines when events should be displayed encoded as base-64 strings using `BINLOG` statements. The option has these permissible values (not case sensitive):

- `AUTO` ("automatic") or `UNSPEC` ("unspecified") displays `BINLOG` statements automatically when necessary (that is, for format description events and row events). If no `--base64-output` [348] option is given, the effect is the same as `--base64-output=AUTO` [348].



#### Note

Automatic `BINLOG` display is the only safe behavior if you intend to use the output of `mysqlbinlog` to re-execute binary log file contents. The other option values are intended only for debugging or testing purposes because they may produce output that does not include all events in executable form.

- `NEVER` causes `BINLOG` statements not to be displayed. `mysqlbinlog` exits with an error if a row event is found that must be displayed using `BINLOG`.
- `DECODE-ROWS` specifies to `mysqlbinlog` that you intend for row events to be decoded and displayed as commented SQL statements by also specifying the `--verbose` [354] option. Like `NEVER`, `DECODE-ROWS` suppresses display of `BINLOG` statements, but unlike `NEVER`, it does not exit with an error if a row event is found.

For examples that show the effect of `--base64-output` [348] and `--verbose` [354] on row event output, see [Section 4.6.8.2, “mysqlbinlog Row Event Display”](#).

- `--bind-address=ip_address` [248]

On a computer having multiple network interfaces, this option can be used to select which interface is employed when connecting to the MySQL server.

This option is supported beginning with MySQL 5.6.1.

- `--binlog-row-event-max-size=N` [349]

<b>Command-Line Format</b>	<code>--binlog-row-event-max-size=#</code>
<b>Option-File Format</b>	<code>binlog-row-event-max-size</code>
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	4294967040
<b>Range</b>	256 .. 18446744073709547520

Specify the maximum size of a row-based binary log event, in bytes. Rows are grouped into events smaller than this size if possible. The value should be a multiple of 256. The default is 4GB.

- `--character-sets-dir=path` [349]

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--database=db_name` [349], `-d db_name`

This option causes `mysqlbinlog` to output entries from the binary log (local log only) that occur while `db_name` is been selected as the default database by `USE`.

The `--database` [349] option for `mysqlbinlog` is similar to the `--binlog-do-db` [2004] option for `mysqld`, but can be used to specify only one database. If `--database` [349] is given multiple times, only the last instance is used.

The effects of this option depend on whether the statement-based or row-based logging format is in use, in the same way that the effects of `--binlog-do-db` [2004] depend on whether statement-based or row-based logging is in use.

**Statement-based logging.** The `--database` [349] option works as follows:

- While `db_name` is the default database, statements are output whether they modify tables in `db_name` or a different database.
- Unless `db_name` is selected as the default database, statements are not output, even if they modify tables in `db_name`.
- There is an exception for `CREATE DATABASE`, `ALTER DATABASE`, and `DROP DATABASE`. The database being *created*, *altered*, or *dropped* is considered to be the default database when determining whether to output the statement.

Suppose that the binary log was created by executing these statements using statement-based logging:

```
INSERT INTO test.t1 (i) VALUES(100);
INSERT INTO db2.t2 (j) VALUES(200);
USE test;
INSERT INTO test.t1 (i) VALUES(101);
INSERT INTO t1 (i) VALUES(102);
INSERT INTO db2.t2 (j) VALUES(201);
USE db2;
INSERT INTO test.t1 (i) VALUES(103);
INSERT INTO db2.t2 (j) VALUES(202);
INSERT INTO t2 (j) VALUES(203);
```

`mysqlbinlog --database=test` does not output the first two `INSERT` statements because there is no default database. It outputs the three `INSERT` statements following `USE test`, but not the three `INSERT` statements following `USE db2`.

`mysqlbinlog --database=db2` does not output the first two `INSERT` statements because there is no default database. It does not output the three `INSERT` statements following `USE test`, but does output the three `INSERT` statements following `USE db2`.

**Row-based logging.** `mysqlbinlog` outputs only entries that change tables belonging to `db_name`. The default database has no effect on this. Suppose that the binary log just described was created using row-based logging rather than statement-based logging. `mysqlbinlog --database=test` outputs only those entries that modify `t1` in the test database, regardless of whether `USE` was issued or what the default database is.

If a server is running with `binlog_format` [2010] set to `MIXED` and you want it to be possible to use `mysqlbinlog` with the `--database` [349] option, you must ensure that tables that are modified are in the database selected by `USE`. (In particular, no cross-database updates should be used.)

Prior to MySQL 5.6.10, the `--database` option did not work correctly with a log written by a GTID-enabled MySQL server. (Bug #15912728)

- `--debug[=debug_options]` [350],  `-# [debug_options]`

Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysqlbinlog.trace'`.

- `--debug-check` [350]

Print some debugging information when the program exits.

- `--debug-info` [350]

Print debugging information and memory and CPU usage statistics when the program exits.

- `--default-auth=plugin` [351]

The client-side authentication plugin to use. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--disable-log-bin` [351], `-D`

Disable binary logging. This is useful for avoiding an endless loop if you use the `--to-last-log` [354] option and are sending the output to the same MySQL server. This option also is useful when restoring after a crash to avoid duplication of the statements you have logged.

This option requires that you have the `SUPER` [756] privilege. It causes `mysqlbinlog` to include a `SET sql_log_bin = 0` statement in its output to disable binary logging of the remaining output. The `SET` statement is ineffective unless you have the `SUPER` [756] privilege.

- `--exclude-gtids=gtid_set` [351]

Do not display any of the groups listed in the `gtid_set`. Added in MySQL 5.6.5.

- `--force-if-open` [351], `-F`

Read binary log files even if they are open or were not closed properly.

- `--force-read` [351], `-f`

With this option, if `mysqlbinlog` reads a binary log event that it does not recognize, it prints a warning, ignores the event, and continues. Without this option, `mysqlbinlog` stops if it reads such an event.

- `--hexdump` [351], `-H`

Display a hex dump of the log in comments, as described in [Section 4.6.8.1, “mysqlbinlog Hex Dump Format”](#). The hex output can be helpful for replication debugging.

- `--host=host_name` [351], `-h host_name`

Get the binary log from the MySQL server on the given host.

- `--include-gtids=gtid_set` [351]

Display only the groups listed in the `gtid_set`. Added in MySQL 5.6.5.

- `--local-load=path` [351], `-l path`

Prepare local temporary files for `LOAD DATA INFILE` in the specified directory.



### Important

These temporary files are not automatically removed by `mysqlbinlog` or any other MySQL program.

- `--offset=N` [351], `-o N`

Skip the first `N` entries in the log.

- `--password[=password]` [351], `-p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value



following the `--password` [351] or `-p` option on the command line, `mysqlbinlog` prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--plugin-dir=path` [352]

The directory in which to look for plugins. It may be necessary to specify this option if the `--default-auth` [351] option is used to specify an authentication plugin but `mysqlbinlog` does not find it. See [Section 6.3.7, “Pluggable Authentication”](#).

This option was added in MySQL 5.6.2.

- `--port=port_num` [352], `-P port_num`

The TCP/IP port number to use for connecting to a remote server.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}` [352]

The connection protocol to use for connecting to the server. It is useful when the other connection parameters normally would cause a protocol to be used other than the one you want. For details on the permissible values, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

- `--raw` [352]

By default, `mysqlbinlog` reads binary log files and writes events in text format. The `--raw` [352] option tells `mysqlbinlog` to write them in their original binary format. Its use requires that `--read-from-remote-server` [352] also be used because the files are requested from a server. `mysqlbinlog` writes one output file for each file read from the server. The `--raw` [352] option can be used to make a backup of a server's binary log. With the `--stop-never` [354] option, the backup is “live” because `mysqlbinlog` stays connected to the server. By default, output files are written in the current directory with the same names as the original log files. Output file names can be modified using the `--result-file` [353] option. For more information, see [Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”](#).

This option was added in MySQL 5.6.0.

- `--read-from-remote-master=type` [352]

Read binary logs from a MySQL server with the `COM_BINLOG_DUMP` or `COM_BINLOG_DUMP_GTID` commands by setting the option value to either `BINLOG-DUMP-NON-GTIDS` or `BINLOG-DUMP-GTIDS`, respectively. If `--read-from-remote-master=BINLOG-DUMP-GTIDS` [352] is combined with `--exclude-gtids` [351], transactions can be filtered out on the master, avoiding unnecessary network traffic.

See also the description for `--read-from-remote-server` [352].

This option was added in MySQL 5.6.5.

- `--read-from-remote-server` [352], `-R`

Read the binary log from a MySQL server rather than reading a local log file. Any connection parameter options are ignored unless this option is given as well. These options are `--host` [351], `--password` [351], `--port` [352], `--protocol` [352], `--socket` [353], and `--user` [354].

This option requires that the remote server be running. It works only for binary log files on the remote server, not relay log files.

As of MySQL 5.6.5, this option is like `--read-from-remote-master=BINLOG-DUMP-NON-GTIDS` [352].

- `--result-file=name` [353], `-r name`

Without the `--raw` [352] option, this option indicates the file to which `mysqlbinlog` writes text output. With `--raw` [352], `mysqlbinlog` writes one binary output file for each log file transferred from the server, writing them by default in the current directory using the same names as the original log file. In this case, the `--result-file` [353] option value is treated as a prefix that modifies output file names.

- `--server-id=id` [353]

Display only those events created by the server having the given server ID.

- `--set-charset=charset_name` [353]

Add a `SET NAMES charset_name` statement to the output to specify the character set to be used for processing log files.

- `--short-form` [353], `-s`

Display only the statements contained in the log, without any extra information or row-based events. This is for testing only, and should not be used in production systems.

- `--skip-gtids[=(true|false)]` [353]

Do not display any GTIDs. Not recommended in production. Added in MySQL 5.6.5.

- `--socket=path` [353], `-S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--start-datetime=datetime` [353]

Start reading the binary log at the first event having a timestamp equal to or later than the `datetime` argument. The `datetime` value is relative to the local time zone on the machine where you run `mysqlbinlog`. The value should be in a format accepted for the `DATETIME` or `TIMESTAMP` data types. For example:

```
shell> mysqlbinlog --start-datetime="2005-12-25 11:25:56" binlog.000003
```

This option is useful for point-in-time recovery. See [Section 7.3, “Example Backup and Recovery Strategy”](#).

- `--start-position=N` [353], `-j N`

Start reading the binary log at the first event having a position equal to or greater than `N`. This option applies to the first log file named on the command line.

This option is useful for point-in-time recovery. See [Section 7.3, “Example Backup and Recovery Strategy”](#).

- `--stop-datetime=datetime` [353]

Stop reading the binary log at the first event having a timestamp equal to or later than the `datetime` argument. This option is useful for point-in-time recovery. See the description of the `--start-datetime` [353] option for information about the `datetime` value.

This option is useful for point-in-time recovery. See [Section 7.3, “Example Backup and Recovery Strategy”](#).

- `--stop-never` [354]

This option is used with `--read-from-remote-server` [352]. It tells `mysqlbinlog` to remain connected to the server. Otherwise `mysqlbinlog` exits when the last log file has been transferred from the server. `--stop-never` [354] implies `--to-last-log` [354], so only the first log file to transfer need be named on the command line.

`--stop-never` [354] is commonly used with `--raw` [352] to make a live binary log backup, but also can be used without `--raw` [352] to maintain a continuous text display of log events as the server generates them.

This option was added in MySQL 5.6.0.

- `--stop-never-slave-server-id=id` [354]

With `--stop-never` [354], `mysqlbinlog` reports a server ID of 65535 when it connects to the server. `--stop-never-slave-server-id` [354] explicitly specifies the server ID to report. It can be used to avoid a conflict with the ID of a slave server or another `mysqlbinlog` process. See Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”.

This option was added in MySQL 5.6.0.

- `--stop-position=N` [354]

Stop reading the binary log at the first event having a position equal to or greater than `N`. This option applies to the last log file named on the command line.

This option is useful for point-in-time recovery. See Section 7.3, “Example Backup and Recovery Strategy”.

- `--to-last-log` [354], `-t`

Do not stop at the end of the requested binary log from a MySQL server, but rather continue printing until the end of the last binary log. If you send the output to the same MySQL server, this may lead to an endless loop. This option requires `--read-from-remote-server` [352].

- `--user=user_name` [354], `-u user_name`

The MySQL user name to use when connecting to a remote server.

- `--verbose` [354], `-v`

Reconstruct row events and display them as commented SQL statements. If this option is given twice, the output includes comments to indicate column data types and some metadata.

For examples that show the effect of `--base64-output` [348] and `--verbose` [354] on row event output, see Section 4.6.8.2, “`mysqlbinlog` Row Event Display”.

- `--verify-binlog-checksum` [354], `-c`

Verify checksums in binary log files. This option was added in MySQL 5.6.1.

- `--version` [354], `-V`

Display version information and exit.

Prior to MySQL 5.6.11, the `mysqlbinlog` version number shown was 3.3. In MySQL 5.6.11 and later, this is 3.4. (Bug #15894381, Bug #67643)

You can also set the following variable by using `--var_name=value` syntax:

- `open_files_limit`

Specify the number of open file descriptors to reserve.

You can pipe the output of `mysqlbinlog` into the `mysql` client to execute the events contained in the binary log. This technique is used to recover from a crash when you have an old backup (see [Section 7.5, “Point-in-Time \(Incremental\) Recovery Using the Binary Log”](#)). For example:

```
shell> mysqlbinlog binlog.000001 | mysql -u root -p
```

Or:

```
shell> mysqlbinlog binlog.[0-9]* | mysql -u root -p
```

If the statements produced by `mysqlbinlog` may contain `BLOB` values, these may cause problems when `mysql` processes them. In this case, invoke `mysql` with the `--binary-mode` [\[248\]](#) option.

You can also redirect the output of `mysqlbinlog` to a text file instead, if you need to modify the statement log first (for example, to remove statements that you do not want to execute for some reason). After editing the file, execute the statements that it contains by using it as input to the `mysql` program:

```
shell> mysqlbinlog binlog.000001 > tmpfile
shell> ... edit tmpfile ...
shell> mysql -u root -p < tmpfile
```

When `mysqlbinlog` is invoked with the `--start-position` [\[353\]](#) option, it displays only those events with an offset in the binary log greater than or equal to a given position (the given position must match the start of one event). It also has options to stop and start when it sees an event with a given date and time. This enables you to perform point-in-time recovery using the `--stop-datetime` [\[353\]](#) option (to be able to say, for example, “roll forward my databases to how they were today at 10:30 a.m.”).

If you have more than one binary log to execute on the MySQL server, the safe method is to process them all using a single connection to the server. Here is an example that demonstrates what may be *unsafe*:

```
shell> mysqlbinlog binlog.000001 | mysql -u root -p # DANGER!!
shell> mysqlbinlog binlog.000002 | mysql -u root -p # DANGER!!
```

Processing binary logs this way using multiple connections to the server causes problems if the first log file contains a `CREATE TEMPORARY TABLE` statement and the second log contains a statement that uses the temporary table. When the first `mysql` process terminates, the server drops the temporary table. When the second `mysql` process attempts to use the table, the server reports “unknown table.”

To avoid problems like this, use a *single* `mysql` process to execute the contents of all binary logs that you want to process. Here is one way to do so:

```
shell> mysqlbinlog binlog.000001 binlog.000002 | mysql -u root -p
```

Another approach is to write all the logs to a single file and then process the file:

```
shell> mysqlbinlog binlog.000001 > /tmp/statements.sql
shell> mysqlbinlog binlog.000002 >> /tmp/statements.sql
shell> mysql -u root -p -e "source /tmp/statements.sql"
```

`mysqlbinlog` can produce output that reproduces a `LOAD DATA INFILE` operation without the original data file. `mysqlbinlog` copies the data to a temporary file and writes a `LOAD DATA LOCAL INFILE` statement that refers to the file. The default location of the directory where these files are written is system-specific. To specify a directory explicitly, use the `--local-load` [\[351\]](#) option.

Because `mysqlbinlog` converts `LOAD DATA INFILE` statements to `LOAD DATA LOCAL INFILE` statements (that is, it adds `LOCAL`), both the client and the server that you use to process the

statements must be configured with the `LOCAL` capability enabled. See [Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”](#).



### Warning

The temporary files created for `LOAD DATA LOCAL` statements are *not* automatically deleted because they are needed until you actually execute those statements. You should delete the temporary files yourself after you no longer need the statement log. The files can be found in the temporary file directory and have names like `original_file_name-#-#`.

#### 4.6.8.1. mysqlbinlog Hex Dump Format

The `--hexdump` [\[351\]](#) option causes `mysqlbinlog` to produce a hex dump of the binary log contents:

```
shell> mysqlbinlog --hexdump master-bin.000001
```

The hex output consists of comment lines beginning with `#`, so the output might look like this for the preceding command:

```
/*!40019 SET @@session.max_insert_delayed_threads=0*/;
/*!50003 SET @OLD_COMPLETION_TYPE=@@COMPLETION_TYPE,COMPLETION_TYPE=0*/;
# at 4
#051024 17:24:13 server id 1  end_log_pos 98
# Position Timestamp Type Master ID Size Master Pos Flags
# 00000004 9d fc 5c 43 0f 01 00 00 00 5e 00 00 00 62 00 00 00 00 00
# 00000017 04 00 35 2e 30 2e 31 35 2d 64 65 62 75 67 2d 6c |..5.0.15.debug.1|
# 00000027 6f 67 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |og.....|
# 00000037 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
# 00000047 00 00 00 00 9d fc 5c 43 13 38 0d 00 08 00 12 00 |.....C.8.....|
# 00000057 04 04 04 04 12 00 00 4b 00 04 1a |.....K...|
# Start: binlog v 4, server v 5.0.15-debug-log created 051024 17:24:13
# at startup
ROLLBACK;
```

Hex dump output currently contains the elements in the following list. This format is subject to change. (For more information about binary log format, see [MySQL Internals: The Binary Log](#).)

- **Position:** The byte position within the log file.
- **Timestamp:** The event timestamp. In the example shown, '9d fc 5c 43' is the representation of '051024 17:24:13' in hexadecimal.
- **Type:** The event type code. In the example shown, '0f' indicates a `FORMAT_DESCRIPTION_EVENT`. The following table lists the possible type codes.

TypeName	Meaning
00 UNKNOWN_EVENT	This event should never be present in the log.
01 START_EVENT_V3	This indicates the start of a log file written by MySQL 4 or earlier.
02 QUERY_EVENT	The most common type of events. These contain statements executed on the master.
03 STOP_EVENT	Indicates that master has stopped.
04 ROTATE_EVENT	Written when the master switches to a new log file.
05 INTVAR_EVENT	Used for <code>AUTO_INCREMENT</code> values or when the <code>LAST_INSERT_ID()</code> <a href="#">[1269]</a> function is used in the statement.
06 LOAD_EVENT	Used for <code>LOAD DATA INFILE</code> in MySQL 3.23.

Type	Name	Meaning
07	SLAVE_EVENT	Reserved for future use.
08	CREATE_FILE_EVENT	Used for <code>LOAD DATA INFILE</code> statements. This indicates the start of execution of such a statement. A temporary file is created on the slave. Used in MySQL 4 only.
09	APPEND_BLOCK_EVENT	Contains data for use in a <code>LOAD DATA INFILE</code> statement. The data is stored in the temporary file on the slave.
0a	EXEC_LOAD_EVENT	Used for <code>LOAD DATA INFILE</code> statements. The contents of the temporary file is stored in the table on the slave. Used in MySQL 4 only.
0b	DELETE_FILE_EVENT	Rollback of a <code>LOAD DATA INFILE</code> statement. The temporary file should be deleted on the slave.
0c	NEW_LOAD_EVENT	Used for <code>LOAD DATA INFILE</code> in MySQL 4 and earlier.
0d	RAND_EVENT	Used to send information about random values if the <code>RAND() [1206]</code> function is used in the statement.
0e	USER_VAR_EVENT	Used to replicate user variables.
0f	FORMAT_DESCRIPTION_EVENT	This indicates the start of a log file written by MySQL 5 or later.
10	XID_EVENT	Event indicating commit of an XA transaction.
11	BEGIN_LOAD_QUERY_EVENT	Used for <code>LOAD DATA INFILE</code> statements in MySQL 5 and later.
12	EXECUTE_LOAD_QUERY_EVENT	Used for <code>LOAD DATA INFILE</code> statements in MySQL 5 and later.
13	TABLE_MAP_EVENT	Information about a table definition. Used in MySQL 5.1.5 and later.
14	PRE_GA_WRITE_ROWS_EVENT	Row data for a single table that should be created. Used in MySQL 5.1.5 to 5.1.17.
15	PRE_GA_UPDATE_ROWS_EVENT	Row data for a single table that needs to be updated. Used in MySQL 5.1.5 to 5.1.17.
16	PRE_GA_DELETE_ROWS_EVENT	Row data for a single table that should be deleted. Used in MySQL 5.1.5 to 5.1.17.
17	WRITE_ROWS_EVENT	Row data for a single table that should be created. Used in MySQL 5.1.18 and later.
18	UPDATE_ROWS_EVENT	Row data for a single table that needs to be updated. Used in MySQL 5.1.18 and later.
19	DELETE_ROWS_EVENT	Row data for a single table that should be deleted. Used in MySQL 5.1.18 and later.
1a	INCIDENT_EVENT	Something out of the ordinary happened. Added in MySQL 5.1.18.

- **Master ID:** The server ID of the master that created the event.
- **Size:** The size in bytes of the event.
- **Master Pos:** The position of the next event in the original master log file.
- **Flags:** 16 flags. Currently, the following flags are used. The others are reserved for future use.

Flag	Name	Meaning
01	LOG_EVENT_BINLOG_IN_USE_F	Log file correctly closed. (Used only in <code>FORMAT_DESCRIPTION_EVENT</code> .) If this flag is

Flag	Name	Meaning
		set (if the flags are, for example, '01 00') in a <code>FORMAT_DESCRIPTION_EVENT</code> , the log file has not been properly closed. Most probably this is because of a master crash (for example, due to power failure).
02		Reserved for future use.
04	<code>LOG_EVENT_THREAD_SPECIFIC</code>	Set if the event is dependent on the connection it was executed in (for example, '04 00'), for example, if the event uses temporary tables.
08	<code>LOG_EVENT_SUPPRESS_USE_F</code>	Set in some circumstances when the event is not dependent on the default database.

#### 4.6.8.2. mysqlbinlog Row Event Display

The following examples illustrate how `mysqlbinlog` displays row events that specify data modifications. These correspond to events with the `WRITE_ROWS_EVENT`, `UPDATE_ROWS_EVENT`, and `DELETE_ROWS_EVENT` type codes. The `--base64-output=DECODE-ROWS` [348] and `--verbose` [354] options may be used to affect row event output.

Suppose that the server is using row-based binary logging and that you execute the following sequence of statements:

```
CREATE TABLE t
(
  id INT NOT NULL,
  name VARCHAR(20) NOT NULL,
  date DATE NULL
) ENGINE = InnoDB;

START TRANSACTION;
INSERT INTO t VALUES(1, 'apple', NULL);
UPDATE t SET name = 'pear', date = '2009-01-01' WHERE id = 1;
DELETE FROM t WHERE id = 1;
COMMIT;
```

By default, `mysqlbinlog` displays row events encoded as base-64 strings using `BINLOG` statements. Omitting extraneous lines, the output for the row events produced by the preceding statement sequence looks like this:

```
shell> mysqlbinlog log_file
...
# at 218
#080828 15:03:08 server id 1  end_log_pos 258  Write_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAANoAAAAABEAAAAAAAAABHRlc3QAAXQAAwMPCgIUAAQ=
fAS3SBcBAAAAKAAAAIBAAAQABEAAAAAAAAEAA//8AQAAAAvhcHBsZQ==
'/*!*/;
...
# at 302
#080828 15:03:08 server id 1  end_log_pos 356  Update_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAAC4BAAAAABEAAAAAAAAABHRlc3QAAXQAAwMPCgIUAAQ=
fAS3SBgBAAAANGAAAGQABAAAQABEAAAAAAAAEAA///AEAAAAAFYXBwbGx4AQAAARwZWFyIbIP
'/*!*/;
...
# at 400
#080828 15:03:08 server id 1  end_log_pos 442  Delete_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAABAAAAABEAAAAAAAAABHRlc3QAAXQAAwMPCgIUAAQ=
fAS3SBkBAAAAAGAAALoBAAAQABEAAAAAAAAEAA//4AQAAARwZWFyIbIP
'/*!*/;
```



To see the row events as comments in the form of “pseudo-SQL” statements, run `mysqlbinlog` with the `--verbose` [354] or `-v` option. The output will contain lines beginning with `###`:

```
shell> mysqlbinlog -v log_file
...
# at 218
#080828 15:03:08 server id 1  end_log_pos 258  Write_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAANoAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBcBAAAAKAAAAAIBAAAQABEAAAAAAAAEAA//8AQAAAAVhcHBsZQ==
'/*!*/;
### INSERT INTO test.t
### SET
###   @1=1
###   @2='apple'
###   @3=NULL
...
# at 302
#080828 15:03:08 server id 1  end_log_pos 356  Update_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAAC4BAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBgBAAAAANGAAAGQBAABEAAAAAAAAEAA///AEAAAAFYXBwbGX4AQAAAArWZWFyIbIP
'/*!*/;
### UPDATE test.t
### WHERE
###   @1=1
###   @2='apple'
###   @3=NULL
### SET
###   @1=1
###   @2='pear'
###   @3='2009:01:01'
...
# at 400
#080828 15:03:08 server id 1  end_log_pos 442  Delete_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAAJABAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBkBAAAAKgAAALoBAAAQABEAAAAAAAAEAA//4AQAAAArWZWFyIbIP
'/*!*/;
### DELETE FROM test.t
### WHERE
###   @1=1
###   @2='pear'
###   @3='2009:01:01'
```

Specify `--verbose` [354] or `-v` twice to also display data types and some metadata for each column. The output will contain an additional comment following each column change:

```
shell> mysqlbinlog -vv log_file
...
# at 218
#080828 15:03:08 server id 1  end_log_pos 258  Write_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAANoAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBcBAAAAKAAAAAIBAAAQABEAAAAAAAAEAA//8AQAAAAVhcHBsZQ==
'/*!*/;
### INSERT INTO test.t
### SET
###   @1=1 /* INT meta=0 nullable=0 is_null=0 */
###   @2='apple' /* VARSTRING(20) meta=20 nullable=0 is_null=0 */
###   @3=NULL /* VARSTRING(20) meta=0 nullable=1 is_null=1 */
...
# at 302
#080828 15:03:08 server id 1  end_log_pos 356  Update_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAAC4BAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBgBAAAAANGAAAGQBAABEAAAAAAAAEAA///AEAAAAFYXBwbGX4AQAAAArWZWFyIbIP
```

```

'/*!*/;
### UPDATE test.t
### WHERE
### @1=1 /* INT meta=0 nullable=0 is_null=0 */
### @2='apple' /* VARSTRING(20) meta=20 nullable=0 is_null=0 */
### @3=NULL /* VARSTRING(20) meta=0 nullable=1 is_null=1 */
### SET
### @1=1 /* INT meta=0 nullable=0 is_null=0 */
### @2='pear' /* VARSTRING(20) meta=20 nullable=0 is_null=0 */
### @3='2009:01:01' /* DATE meta=0 nullable=1 is_null=0 */
...
# at 400
#080828 15:03:08 server id 1 end_log_pos 442 Delete_rows: table id 17 flags: STMT_END_F

BINLOG '
fAS3SBMBAAAAALAAAAJABAAAAABEAAAAAAAAABHRLc3QAAXQAawMPCgIUAAQ=
fAS3SBkBAAAAKgAAALoBAAQAABEAAAAAAAAEAA//4AQAAAAARwZWFyIbIP
'/*!*/;
### DELETE FROM test.t
### WHERE
### @1=1 /* INT meta=0 nullable=0 is_null=0 */
### @2='pear' /* VARSTRING(20) meta=20 nullable=0 is_null=0 */
### @3='2009:01:01' /* DATE meta=0 nullable=1 is_null=0 */

```

You can tell `mysqlbinlog` to suppress the `BINLOG` statements for row events by using the `--base64-output=DECODE-ROWS` [348] option. This is similar to `--base64-output=NEVER` [348] but does not exit with an error if a row event is found. The combination of `--base64-output=DECODE-ROWS` [348] and `--verbose` [354] provides a convenient way to see row events only as SQL statements:

```

shell> mysqlbinlog -v --base64-output=DECODE-ROWS log_file
...
# at 218
#080828 15:03:08 server id 1 end_log_pos 258 Write_rows: table id 17 flags: STMT_END_F
### INSERT INTO test.t
### SET
### @1=1
### @2='apple'
### @3=NULL
...
# at 302
#080828 15:03:08 server id 1 end_log_pos 356 Update_rows: table id 17 flags: STMT_END_F
### UPDATE test.t
### WHERE
### @1=1
### @2='apple'
### @3=NULL
### SET
### @1=1
### @2='pear'
### @3='2009:01:01'
...
# at 400
#080828 15:03:08 server id 1 end_log_pos 442 Delete_rows: table id 17 flags: STMT_END_F
### DELETE FROM test.t
### WHERE
### @1=1
### @2='pear'
### @3='2009:01:01'

```



### Note

You should not suppress `BINLOG` statements if you intend to re-execute `mysqlbinlog` output.

The SQL statements produced by `--verbose` [354] for row events are much more readable than the corresponding `BINLOG` statements. However, they do not correspond exactly to the original SQL statements that generated the events. The following limitations apply:

- The original column names are lost and replaced by `@N`, where `N` is a column number.

- Character set information is not available in the binary log, which affects string column display:
  - There is no distinction made between corresponding binary and nonbinary string types (`BINARY` and `CHAR`, `VARBINARY` and `VARCHAR`, `BLOB` and `TEXT`). The output uses a data type of `STRING` for fixed-length strings and `VARSTRING` for variable-length strings.
  - For multi-byte character sets, the maximum number of bytes per character is not present in the binary log, so the length for string types is displayed in bytes rather than in characters. For example, `STRING(4)` will be used as the data type for values from either of these column types:

```
CHAR(4) CHARACTER SET latin1
CHAR(2) CHARACTER SET ucs2
```

- Due to the storage format for events of type `UPDATE_ROWS_EVENT`, `UPDATE` statements are displayed with the `WHERE` clause preceding the `SET` clause.

Proper interpretation of row events requires the information from the format description event at the beginning of the binary log. Because `mysqlbinlog` does not know in advance whether the rest of the log contains row events, by default it displays the format description event using a `BINLOG` statement in the initial part of the output.

If the binary log is known not to contain any events requiring a `BINLOG` statement (that is, no row events), the `--base64-output=NEVER` [348] option can be used to prevent this header from being written.

#### 4.6.8.3. Using `mysqlbinlog` to Back Up Binary Log Files

By default, `mysqlbinlog` reads binary log files and displays their contents in text format. This enables you to examine events within the files more easily or re-execute them (for example, by using the output as input to `mysql`). `mysqlbinlog` can read log files directly from the local file system, or, with the `--read-from-remote-server` [352] option, it can connect to a server and request binary log contents from that server. `mysqlbinlog` writes text output to its standard output, or to the file named as the value of the `--result-file=file_name` [353] option if that option is given.

As of MySQL 5.6, `mysqlbinlog` can read binary log files and write new files containing the same content—that is, in binary format rather than text format. This capability enables you to easily back up a binary log in its original format. `mysqlbinlog` can make a static backup, backing up a set of log files and stopping when the end of the last file is reached. It can also make a continuous (“live”) backup, staying connected to the server when it reaches the end of the last log file and continuing to copy new events as they are generated. In continuous-backup operation, `mysqlbinlog` runs until the connection ends (for example, when the server exits) or `mysqlbinlog` is forcibly terminated. When the connection ends, `mysqlbinlog` does not wait and retry the connection, unlike a slave replication server. To continue a live backup after the server has been restarted, you must also restart `mysqlbinlog`.

Binary log backup requires that you invoke `mysqlbinlog` with two options at minimum:

- The `--read-from-remote-server` [352] (or `-R`) option tells `mysqlbinlog` to connect to a server and request its binary log. (This is similar to a slave replication server connecting to its master server.)
- The `--raw` [352] option tells `mysqlbinlog` to write raw (binary) output, not text output.

Along with `--read-from-remote-server` [352], it is common to specify other options: `--host` [351] indicates where the server is running, and you may also need to specify connection options such as `--user` [354] and `--password` [351].

Several other options are useful in conjunction with `--raw` [352]:

- `--stop-never` [354]: Stay connected to the server after reaching the end of the last log file and continue to read new events.

- `--stop-never-slave-server-id=id` [354]: The server ID that `mysqlbinlog` reports to the server when `--stop-never` [354] is used. The default is 65535. This can be used to avoid a conflict with the ID of a slave server or another `mysqlbinlog` process. See [Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”](#).
- `--result-file` [353]: A prefix for output file names, as described later.

To back up a server's binary log files with `mysqlbinlog`, you must specify file names that actually exist on the server. If you do not know the names, connect to the server and use the `SHOW BINARY LOGS` statement to see the current names. Suppose that the statement produces this output:

```
mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name      | File_size |
+-----+-----+
| binlog.000130 |      27459 |
| binlog.000131 |      13719 |
| binlog.000132 |      43268 |
+-----+-----+
```

With that information, you can use `mysqlbinlog` to back up the binary log to the current directory as follows (enter each command on a single line):

- To make a static backup of `binlog.000130` through `binlog.000132`, use either of these commands:

```
mysqlbinlog --read-from-remote-server --host=host_name --raw
  binlog.000130 binlog.000131 binlog.000132

mysqlbinlog --read-from-remote-server --host=host_name --raw
  --to-last-log binlog.000130
```

The first command specifies every file name explicitly. The second names only the first file and uses `--to-last-log` [354] to read through the last. A difference between these commands is that if the server happens to open `binlog.000133` before `mysqlbinlog` reaches the end of `binlog.000132`, the first command will not read it, but the second command will.

- To make a live backup in which `mysqlbinlog` starts with `binlog.000130` to copy existing log files, then stays connected to copy new events as the server generates them:

```
mysqlbinlog --read-from-remote-server --host=host_name --raw
  --stop-never binlog.000130
```

With `--stop-never` [354], it is not necessary to specify `--to-last-log` [354] to read to the last log file because that option is implied.

## Output File Naming

Without `--raw` [352], `mysqlbinlog` produces text output and the `--result-file` [353] option, if given, specifies the name of the single file to which all output is written. With `--raw` [352], `mysqlbinlog` writes one binary output file for each log file transferred from the server. By default, `mysqlbinlog` writes the files in the current directory with the same names as the original log files. To modify the output file names, use the `--result-file` [353] option. In conjunction with `--raw` [352], the `--result-file` [353] option value is treated as a prefix that modifies the output file names.

Suppose that a server currently has binary log files named `binlog.000999` and up. If you use `mysqlbinlog --raw` to back up the files, the `--result-file` [353] option produces output file names as shown in the following table. You can write the files to a specific directory by beginning the `--result-file` [353] value with the directory path. If the `--result-file` [353] value consists only of a directory name, the value must end with the pathname separator character. Output files are overwritten if they exist.

<code>--result-file</code> [353] Option	Output File Names
<code>--result-file=x</code> [353]	<code>xbinlog.000999</code> and up
<code>--result-file=/tmp/</code> [353]	<code>/tmp/binlog.000999</code> and up
<code>--result-file=/tmp/x</code> [353]	<code>/tmp/xbinlog.000999</code> and up

### Example: `mysqldump` + `mysqlbinlog` for Backup and Restore

The following example describes a simple scenario that shows how to use `mysqldump` and `mysqlbinlog` together to back up a server's data and binary log, and how to use the backup to restore the server if data loss occurs. The example assumes that the server is running on host `host_name` and its first binary log file is named `binlog.000999`. Enter each command on a single line.

Use `mysqlbinlog` to make a continuous backup of the binary log:

```
mysqlbinlog --read-from-remote-server --host=host_name --raw
--stop-never binlog.000999
```

Use `mysqldump` to create a dump file as a snapshot of the server's data. Use `--all-databases` [292], `--events` [293], and `--routines` [293] to back up all data, and `--master-data=2` [288] to include the current binary log coordinates in the dump file.

```
mysqldump --host=host_name --all-databases --events --routines --master-data=2 > dump_file
```

Execute the `mysqldump` command periodically to create newer snapshots as desired.

If data loss occurs (for example, if the server crashes), use the most recent dump file to restore the data:

```
mysql --host=host_name -u root -p < dump_file
```

Then use the binary log backup to re-execute events that were written after the coordinates listed in the dump file. Suppose that the coordinates in the file look like this:

```
-- CHANGE MASTER TO MASTER_LOG_FILE='binlog.001002', MASTER_LOG_POS=27284;
```

If the most recent backed-up log file is named `binlog.001004`, re-execute the log events like this:

```
mysqlbinlog --start-position=27284 binlog.001002 binlog.001003 binlog.001004
| mysql --host=host_name -u root -p
```

You might find it easier to copy the backup files (dump file and binary log files) to the server host to make it easier to perform the restore operation, or if MySQL does not allow remote `root` access.

#### 4.6.8.4. Specifying the `mysqlbinlog` Server ID

When invoked with the `--read-from-remote-server` option, `mysqlbinlog` connects to a MySQL server, specifies a server ID to identify itself, and requests binary log files from the server. You can use `mysqlbinlog` to request log files from a server in several ways:

- Specify an explicitly named set of files: For each file, `mysqlbinlog` connects and issues a `Binlog dump` command. The server sends the file and disconnects. There is one connection per file.
- Specify the beginning file and `--to-last-log` [354]: `mysqlbinlog` connects and issues a `Binlog dump` command for all files. The server sends all files and disconnects.
- Specify the beginning file and `--stop-never` [354] (which implies `--to-last-log` [354]): `mysqlbinlog` connects and issues a `Binlog dump` command for all files. The server sends all files, but does not disconnect after sending the last one.

With `--read-from-remote-server` [352] only, `mysqlbinlog` connects using a server ID of 0, which tells the server to disconnect after sending the last requested log file.

With `--read-from-remote-server` [352] and `--stop-never` [354], `mysqlbinlog` connects using a nonzero server ID, so the server does not disconnect after sending the last log file. The server ID is 65535 by default, but this can be changed with `--stop-never-slave-server-id` [354].

Thus, for the first two ways of requesting files, the server disconnects because `mysqlbinlog` specifies a server ID of 0. It does not disconnect if `--stop-never` [354] is given because `mysqlbinlog` specifies a nonzero server ID.

#### 4.6.9. `mysqldumpslow` — Summarize Slow Query Log Files

The MySQL slow query log contains information about queries that take a long time to execute (see Section 5.2.5, “The Slow Query Log”). `mysqldumpslow` parses MySQL slow query log files and prints a summary of their contents.

Normally, `mysqldumpslow` groups queries that are similar except for the particular values of number and string data values. It “abstracts” these values to `N` and `'S'` when displaying summary output. The `-a` and `-n` options can be used to modify value abstracting behavior.

Invoke `mysqldumpslow` like this:

```
shell> mysqldumpslow [options] [log_file ...]
```

`mysqldumpslow` supports the following options.

**Table 4.14. `mysqldumpslow` Options**

Format	Option File	Description
<code>-a</code> [364]		Do not abstract all numbers to <code>N</code> and strings to <code>S</code>
<code>-n num</code> [365]		Abstract numbers with at least the specified digits
<code>--debug</code> [364]	<code>debug</code> [364]	Write debugging information
<code>-g pattern</code> [364]		Only consider statements that match the pattern
<code>--help</code> [364]		Display help message and exit
<code>-h name</code> [365]		Host name of the server in the log file name
<code>-i name</code> [365]		Name of the server instance
<code>-l</code> [365]		Do not subtract lock time from total time
<code>-r</code> [365]		Reverse the sort order
<code>-s value</code> [365]		How to sort output
<code>-t num</code> [365]		Display only first <code>num</code> queries
<code>--verbose</code> [365]	<code>verbose</code> [365]	Verbose mode

- `--help` [364]

Display a help message and exit.

- `-a`

Do not abstract all numbers to `N` and strings to `'S'`.

- `--debug` [364], `-d`

Run in debug mode.

- `-g pattern`

Consider only queries that match the (`grep`-style) pattern.

- `-h host_name`

Host name of MySQL server for `*-slow.log` file name. The value can contain a wildcard. The default is `*` (match all).

- `-i name`

Name of server instance (if using `mysql.server` startup script).

- `-l`

Do not subtract lock time from total time.

- `-n N`

Abstract numbers with at least `N` digits within names.

- `-r`

Reverse the sort order.

- `-s sort_type`

How to sort the output. The value of `sort_type` should be chosen from the following list:

- `t, at`: Sort by query time or average query time
- `l, al`: Sort by lock time or average lock time
- `r, ar`: Sort by rows sent or average rows sent
- `c`: Sort by count

By default, `mysqldumpslow` sorts by average query time (equivalent to `-s at`).

- `-t N`

Display only the first `N` queries in the output.

- `--verbose [365], -v`

Verbose mode. Print more information about what the program does.

Example of usage:

```
shell> mysqldumpslow

Reading mysql slow query log from /usr/local/mysql/data/mysql51-apple-slow.log
Count: 1  Time=4.32s (4s)  Lock=0.00s (0s)  Rows=0.0 (0), root[root]@localhost
insert into t2 select * from t1

Count: 3  Time=2.53s (7s)  Lock=0.00s (0s)  Rows=0.0 (0), root[root]@localhost
insert into t2 select * from t1 limit N

Count: 3  Time=2.13s (6s)  Lock=0.00s (0s)  Rows=0.0 (0), root[root]@localhost
insert into t1 select * from t1
```

#### 4.6.10. mysqlhotcopy — A Database Backup Program

`mysqlhotcopy` is a Perl script that was originally written and contributed by Tim Bunce. It uses `FLUSH TABLES`, `LOCK TABLES`, and `cp` or `scp` to make a database backup. It is a fast way to make a backup of the database or single tables, but it can be run only on the same machine where the database directories are located. `mysqlhotcopy` works only for backing up `MyISAM` and `ARCHIVE` tables. It runs on Unix.



To use `mysqlhotcopy`, you must have read access to the files for the tables that you are backing up, the `SELECT` [755] privilege for those tables, the `RELOAD` [755] privilege (to be able to execute `FLUSH TABLES`), and the `LOCK TABLES` [755] privilege (to be able to lock the tables).

```
shell> mysqlhotcopy db_name [/path/to/new_directory]
```

```
shell> mysqlhotcopy db_name_1 ... db_name_n /path/to/new_directory
```

Back up tables in the given database that match a regular expression:

```
shell> mysqlhotcopy db_name./regex/
```

The regular expression for the table name can be negated by prefixing it with a tilde (“~”):

```
shell> mysqlhotcopy db_name./~regex/
```

`mysqlhotcopy` supports the following options, which can be specified on the command line or in the `[mysqlhotcopy]` and `[client]` groups of an option file. For information about option files, see Section 4.2.3.3, “Using Option Files”.

**Table 4.15. `mysqlhotcopy` Options**

Format	Option File	Description
<code>--addtodest</code> [367]	<code>addtodest</code> [367]	Do not rename target directory (if it exists); merely add files to it
<code>--allowold</code> [367]	<code>allowold</code> [367]	Do not abort if a target exists; rename it by adding an <code>_old</code> suffix
<code>--checkpoint=db_name.tbl_name</code> [367]	<code>checkpoint</code> [367] <code>tbl_name</code> [367]	Insert checkpoint entries
<code>--chroot=path</code> [367]	<code>chroot</code> [367]	Base directory of the chroot jail in which <code>mysqld</code> operates
<code>--debug</code> [367]	<code>debug</code> [367]	Write a debugging log
<code>--dryrun</code> [367]	<code>dryrun</code> [367]	Report actions without performing them
<code>--flushlog</code> [367]	<code>flushlog</code> [367]	Flush logs after all tables are locked
<code>--help</code> [367]		Display help message and exit
<code>--host=host_name</code> [367]	<code>host</code> [367]	Connect to the MySQL server on the given host
<code>--keepold</code> [367]	<code>keepold</code> [367]	Do not delete previous (renamed) target when done
<code>--method</code> [367]	<code>method</code> [367]	The method for copying files
<code>--noindices</code> [367]	<code>noindices</code> [367]	Do not include full index files in the backup
<code>--old_server</code> [368]	<code>old_server</code> [368]	Connect to server that does not support <code>FLUSH TABLES tbl_list WITH READ LOCK</code>
<code>--password[=password]</code> [367]	<code>password</code> [367] [367]	The password to use when connecting to the server
<code>--port=port_num</code> [368]	<code>port</code> [368]	The TCP/IP port number to use for the connection
<code>--quiet</code> [368]	<code>quiet</code> [368]	Be silent except for errors
<code>--regex</code> [368]	<code>regex</code> [368]	Copy all databases with names that match the given regular expression
<code>--resetmaster</code> [368]	<code>resetmaster</code> [368]	Reset the binary log after locking all the tables
<code>--resetslave</code> [368]	<code>resetslave</code> [368]	Reset the master.info file after locking all the tables
<code>--socket=path</code> [368]	<code>socket</code> [368]	For connections to localhost
<code>--tmpdir=path</code> [368]	<code>tmpdir</code> [368]	The temporary directory

Format	Option File	Description
<code>--user=user_name, [368]</code>	<code>user [368]</code>	MySQL user name to use when connecting to server

- `--help [367], -?`  
Display a help message and exit.
- `--addtodest [367]`  
Do not rename target directory (if it exists); merely add files to it.
- `--allowold [367]`  
Do not abort if a target exists; rename it by adding an `_old` suffix.
- `--checkpoint=db_name.tbl_name [367]`  
Insert checkpoint entries into the specified database `db_name` and table `tbl_name`.
- `--chroot=path [367]`  
Base directory of the `chroot` jail in which `mysqld` operates. The `path` value should match that of the `--chroot [417]` option given to `mysqld`.
- `--debug [367]`  
Enable debug output.
- `--dryrun [367], -n`  
Report actions without performing them.
- `--flushlog [367]`  
Flush logs after all tables are locked.
- `--host=host_name [367], -h host_name`  
The host name of the local host to use for making a TCP/IP connection to the local server. By default, the connection is made to `localhost` using a Unix socket file.
- `--keepold [367]`  
Do not delete previous (renamed) target when done.
- `--method=command [367]`  
The method for copying files (`cp` or `scp`). The default is `cp`.
- `--noindices [367]`  
Do not include full index files for `MyISAM` tables in the backup. This makes the backup smaller and faster. The indexes for reloaded tables can be reconstructed later with `myisamchk -rq`.
- `--password=password [367], -ppassword`  
The password to use when connecting to the server. The password value is not optional for this option, unlike for other MySQL programs.  
  
Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--port=port_num [368]`, `-P port_num`

The TCP/IP port number to use when connecting to the local server.

- `--old_server [368]`

In MySQL 5.6, `mysqlhotcopy` uses `FLUSH TABLES tbl_list WITH READ LOCK` to flush and lock tables. Use the `--old_server [368]` option if the server is older than 5.5.3, which is when that statement was introduced.

- `--quiet [368]`, `-q`

Be silent except for errors.

- `--record_log_pos=db_name.tbl_name [368]`

Record master and slave status in the specified database `db_name` and table `tbl_name`.

- `--regexp=expr [368]`

Copy all databases with names that match the given regular expression.

- `--resetmaster [368]`

Reset the binary log after locking all the tables.

- `--resetslave [368]`

Reset the `master.info` file after locking all the tables.

- `--socket=path [368]`, `-S path`

The Unix socket file to use for connections to `localhost`.

- `--suffix=str [368]`

The suffix to use for names of copied databases.

- `--tmpdir=path [368]`

The temporary directory. The default is `/tmp`.

- `--user=user_name [368]`, `-u user_name`

The MySQL user name to use when connecting to the server.

Use `perldoc` for additional `mysqlhotcopy` documentation, including information about the structure of the tables needed for the `--checkpoint [367]` and `--record_log_pos [368]` options:

```
shell> perldoc mysqlhotcopy
```

#### 4.6.11. `mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine

`mysql_convert_table_format` converts the tables in a database to use a particular storage engine (MyISAM by default). `mysql_convert_table_format` is written in Perl and requires that the `DBI` and `DBD::mysql` Perl modules be installed (see [Section 2.13, “Perl Installation Notes”](#)).

Invoke `mysql_convert_table_format` like this:

```
shell> mysql_convert_table_format [options]db_name
```

The `db_name` argument indicates the database containing the tables to be converted.

`mysql_convert_table_format` supports the options described in the following list.

- `--help` [369]  
Display a help message and exit.
- `--force` [369]  
Continue even if errors occur.
- `--host=host_name` [369]  
Connect to the MySQL server on the given host.
- `--password=password` [369]  
The password to use when connecting to the server. Note that the password value is not optional for this option, unlike for other MySQL programs.  
  
Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.
- `--port=port_num` [369]  
The TCP/IP port number to use for the connection.
- `--socket=path` [369]  
For connections to `localhost`, the Unix socket file to use.
- `--type=engine_name` [369]  
Specify the storage engine that the tables should be converted to use. The default is `MyISAM` if this option is not given.
- `--user=user_name` [369]  
The MySQL user name to use when connecting to the server.
- `--verbose` [369]  
Verbose mode. Print more information about what the program does.
- `--version` [369]  
Display version information and exit.

#### 4.6.12. `mysql_find_rows` — Extract SQL Statements from Files

`mysql_find_rows` reads files containing SQL statements and extracts statements that match a given regular expression or that contain `USE db_name` or `SET` statements. The utility expects statements to be terminated with semicolon (`;`) characters.

Invoke `mysql_find_rows` like this:

```
shell> mysql_find_rows [options] [file_name ...]
```

Each `file_name` argument should be the name of file containing SQL statements. If no file names are given, `mysql_find_rows` reads the standard input.

Examples:

```
mysql_find_rows --regexp=problem_table --rows=20 < update.log
```

```
mysql_find_rows --regexp=problem_table update-log.1 update-log.2
```

`mysql_find_rows` supports the following options:

- `--help` [370], `--Information` [370]  
Display a help message and exit.
- `--regexp=pattern` [370]  
Display queries that match the pattern.
- `--rows=N` [370]  
Quit after displaying *N* queries.
- `--skip-use-db` [370]  
Do not include `USE db_name` statements in the output.
- `--start_row=N` [370]  
Start output from this row.

### 4.6.13. `mysql_fix_extensions` — Normalize Table File Name Extensions

`mysql_fix_extensions` converts the extensions for MyISAM (or ISAM) table files to their canonical forms. It looks for files with extensions matching any lettercase variant of `.frm`, `.myd`, `.myi`, `.isd`, and `.ism` and renames them to have extensions of `.frm`, `.MYD`, `.MYI`, `.ISD`, and `.ISM`, respectively. This can be useful after transferring the files from a system with case-insensitive file names (such as Windows) to a system with case-sensitive file names.

Invoke `mysql_fix_extensions` like this, where `data_dir` is the path name to the MySQL data directory.

```
shell> mysql_fix_extensions data_dir
```

### 4.6.14. `mysql_setpermission` — Interactively Set Permissions in Grant Tables

`mysql_setpermission` is a Perl script that was originally written and contributed by Luuk de Boer. It interactively sets permissions in the MySQL grant tables. `mysql_setpermission` is written in Perl and requires that the `DBI` and `DBD:mysql` Perl modules be installed (see [Section 2.13, “Perl Installation Notes”](#)).

Invoke `mysql_setpermission` like this:

```
shell> mysql_setpermission [options]
```

`options` should be either `--help` [370] to display the help message, or options that indicate how to connect to the MySQL server. The account used when you connect determines which permissions you have when attempting to modify existing permissions in the grant tables.

`mysql_setpermissions` also reads options from the `[client]` and `[perl]` groups in the `.my.cnf` file in your home directory, if the file exists.

`mysql_setpermission` supports the following options:

- `--help` [370]  
Display a help message and exit.

- `--host=host_name` [371]

Connect to the MySQL server on the given host.

- `--password=password` [371]

The password to use when connecting to the server. Note that the password value is not optional for this option, unlike for other MySQL programs.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

- `--port=port_num` [371]

The TCP/IP port number to use for the connection.

- `--socket=path` [371]

For connections to `localhost`, the Unix socket file to use.

- `--user=user_name` [371]

The MySQL user name to use when connecting to the server.

#### 4.6.15. `mysql_waitpid` — Kill Process and Wait for Its Termination

`mysql_waitpid` signals a process to terminate and waits for the process to exit. It uses the `kill()` system call and Unix signals, so it runs on Unix and Unix-like systems.

Invoke `mysql_waitpid` like this:

```
shell> mysql_waitpid [options] pid wait_time
```

`mysql_waitpid` sends signal 0 to the process identified by `pid` and waits up to `wait_time` seconds for the process to terminate. `pid` and `wait_time` must be positive integers.

If process termination occurs within the wait time or the process does not exist, `mysql_waitpid` returns 0. Otherwise, it returns 1.

If the `kill()` system call cannot handle signal 0, `mysql_waitpid()` uses signal 1 instead.

`mysql_waitpid` supports the following options:

- `--help` [371], `-?`, `-I`

Display a help message and exit.

- `--verbose` [371], `-v`

Verbose mode. Display a warning if signal 0 could not be used and signal 1 is used instead.

- `--version` [371], `-V`

Display version information and exit.

#### 4.6.16. `mysql_zap` — Kill Processes That Match a Pattern

`mysql_zap` kills processes that match a pattern. It uses the `ps` command and Unix signals, so it runs on Unix and Unix-like systems.

Invoke `mysql_zap` like this:

```
shell> mysql_zap [-signal] [-?Ift] pattern
```

A process matches if its output line from the `ps` command contains the pattern. By default, `mysql_zap` asks for confirmation for each process. Respond `y` to kill the process, or `q` to exit `mysql_zap`. For any other response, `mysql_zap` does not attempt to kill the process.

If the `-signal` option is given, it specifies the name or number of the signal to send to each process. Otherwise, `mysql_zap` tries first with `TERM` (signal 15) and then with `KILL` (signal 9).

`mysql_zap` supports the following additional options:

- `--help [372]`, `-?`, `-I`

Display a help message and exit.

- `-f`

Force mode. `mysql_zap` attempts to kill each process without confirmation.

- `-t`

Test mode. Display information about each process but do not kill it.

## 4.7. MySQL Program Development Utilities

This section describes some utilities that you may find useful when developing MySQL programs.

In shell scripts, you can use the `my_print_defaults` program to parse option files and see what options would be used by a given program. The following example shows the output that `my_print_defaults` might produce when asked to show the options found in the `[client]` and `[mysql]` groups:

```
shell> my_print_defaults client mysql
--port=3306
--socket=/tmp/mysql.sock
--no-auto-rehash
```

Note for developers: Option file handling is implemented in the C client library simply by processing all options in the appropriate group or groups before any command-line arguments. This works well for programs that use the last instance of an option that is specified multiple times. If you have a C or C++ program that handles multiply specified options this way but that doesn't read option files, you need add only two lines to give it that capability. Check the source code of any of the standard MySQL clients to see how to do this.

Several other language interfaces to MySQL are based on the C client library, and some of them provide a way to access option file contents. These include Perl and Python. For details, see the documentation for your preferred interface.

### 4.7.1. `msql2mysql` — Convert mSQL Programs for Use with MySQL

Initially, the MySQL C API was developed to be very similar to that for the mSQL database system. Because of this, mSQL programs often can be converted relatively easily for use with MySQL by changing the names of the C API functions.

The `msql2mysql` utility performs the conversion of mSQL C API function calls to their MySQL equivalents. `msql2mysql` converts the input file in place, so make a copy of the original before converting it. For example, use `msql2mysql` like this:

```
shell> cp client-prog.c client-prog.c.orig
shell> msql2mysql client-prog.c
client-prog.c converted
```



Then examine `client-prog.c` and make any post-conversion revisions that may be necessary.

`mysql2mysql` uses the `replace` utility to make the function name substitutions. See [Section 4.8.2, “replace — A String-Replacement Utility”](#).

## 4.7.2. `mysql_config` — Display Options for Compiling Clients

`mysql_config` provides you with useful information for compiling your MySQL client and connecting it to MySQL. It is a shell script, so it is available only on Unix and Unix-like systems.

`mysql_config` supports the following options.

- `--cflags` [373]

C Compiler flags to find include files and critical compiler flags and defines used when compiling the `libmysqlclient` library. The options returned are tied to the specific compiler that was used when the library was created and might clash with the settings for your own compiler. Use `--include` [373] for more portable options that contain only include paths.

- `--cxxflags` [373]

Like `--cflags` [373], but for C++ compiler flags. This option was added in MySQL 5.6.4.

- `--include` [373]

Compiler options to find MySQL include files.

- `--libmysqld-libs` [373], `--embedded` [373]

Libraries and options required to link with the MySQL embedded server.

- `--libs` [373]

Libraries and options required to link with the MySQL client library.

- `--libs_r` [373]

Libraries and options required to link with the thread-safe MySQL client library. In MySQL 5.6, all client libraries are thread-safe, so this option need not be used. The `--libs` [373] option can be used in all cases.

- `--plugindir` [373]

The default plugin directory path name, defined when configuring MySQL.

- `--port` [373]

The default TCP/IP port number, defined when configuring MySQL.

- `--socket` [373]

The default Unix socket file, defined when configuring MySQL.

- `--version` [373]

Version number for the MySQL distribution.

If you invoke `mysql_config` with no options, it displays a list of all options that it supports, and their values:

```
shell> mysql_config
Usage: /usr/local/mysql/bin/mysql_config [options]
Options:
```

```

--cflags      [-I/usr/local/mysql/include/mysql -mcpu=pentiumpro]
--include     [-I/usr/local/mysql/include/mysql]
--libs       [-L/usr/local/mysql/lib/mysql -lmysqlclient
             -lpthread -lm -lrt -lssl -lcrypto -ldl]
--libs_r     [-L/usr/local/mysql/lib/mysql -lmysqlclient_r
             -lpthread -lm -lrt -lssl -lcrypto -ldl]
--socket      [/tmp/mysql.sock]
--port       [3306]
--version    [5.6.11]
--libmysqld-libs [-L/usr/local/mysql/lib/mysql -lmysqld
                -lpthread -lm -lrt -lssl -lcrypto -ldl -lcrypt]

```

You can use `mysql_config` within a command line using backticks to include the output that it produces for a particular option. For example, to compile and link a MySQL client program, use `mysql_config` as follows:

```

shell> gcc -c `mysql_config --cflags` progname.c
shell> gcc -o progname progname.o `mysql_config --libs`

```

### 4.7.3. my\_print\_defaults — Display Options from Option Files

`my_print_defaults` displays the options that are present in option groups of option files. The output indicates what options will be used by programs that read the specified option groups. For example, the `mysqlcheck` program reads the `[mysqlcheck]` and `[client]` option groups. To see what options are present in those groups in the standard option files, invoke `my_print_defaults` like this:

```

shell> my_print_defaults mysqlcheck client
--user=myusername
--password=secret
--host=localhost

```

The output consists of options, one per line, in the form that they would be specified on the command line.

`my_print_defaults` supports the following options.

- `--help [374], -?`

Display a help message and exit.

- `--config-file=file_name [374], --defaults-file=file_name [374], -c file_name`

Read only the given option file.

- `--debug=debug_options [374], -# debug_options`

Write a debugging log. A typical `debug_options` string is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/my_print_defaults.trace'`.

- `--defaults-extra-file=file_name [374], --extra-file=file_name [374], -e file_name`

Read this option file after the global option file but (on Unix) before the user option file.

- `--defaults-group-suffix=suffix [374], -g suffix`

In addition to the groups named on the command line, read groups that have the given suffix.

- `--login-path=name [374], -l name`

Read options from the named login path in the `.mylogin.cnf` login file. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#). This option was added in MySQL 5.6.6.

- `--no-defaults [220], -n`

Return an empty string.

- `--verbose` [375], `-v`

Verbose mode. Print more information about what the program does.

- `--version` [375], `-V`

Display version information and exit.

#### 4.7.4. `resolve_stack_dump` — Resolve Numeric Stack Trace Dump to Symbols

`resolve_stack_dump` resolves a numeric stack dump to symbols.

Invoke `resolve_stack_dump` like this:

```
shell> resolve_stack_dump [options] symbols_file [numeric_dump_file]
```

The symbols file should include the output from the `nm --numeric-sort mysqld` command. The numeric dump file should contain a numeric stack track from `mysqld`. If no numeric dump file is named on the command line, the stack trace is read from the standard input.

`resolve_stack_dump` supports the following options.

- `--help` [375], `-h`

Display a help message and exit.

- `--numeric-dump-file=file_name` [375], `-n file_name`

Read the stack trace from the given file.

- `--symbols-file=file_name` [375], `-s file_name`

Use the given symbols file.

- `--version` [375], `-V`

Display version information and exit.

## 4.8. Miscellaneous Programs

### 4.8.1. `pererror` — Explain Error Codes

For most system errors, MySQL displays, in addition to an internal text message, the system error code in one of the following styles:

```
message ... (errno: #)
message ... (Errcode: #)
```

You can find out what the error code means by examining the documentation for your system or by using the `pererror` utility.

`pererror` prints a description for a system error code or for a storage engine (table handler) error code.

Invoke `pererror` like this:

```
shell> pererror [options] errorcode ...
```

Example:

```
shell> perror 13 64
OS error code 13: Permission denied
OS error code 64: Machine is not on the network
```

To obtain the error message for a MySQL Cluster error code, invoke `perror` with the `--ndb` [376] option:

```
shell> perror --ndb errorcode
```

Note that the meaning of system error messages may be dependent on your operating system. A given error code may mean different things on different operating systems.

`perror` supports the following options.

- `--help` [376], `--info` [376], `-I`, `-?`

Display a help message and exit.

- `--ndb` [376]

Print the error message for a MySQL Cluster error code.

- `--silent` [376], `-s`

Silent mode. Print only the error message.

- `--verbose` [376], `-v`

Verbose mode. Print error code and message. This is the default behavior.

- `--version` [376], `-V`

Display version information and exit.

## 4.8.2. `replace` — A String-Replacement Utility

The `replace` utility program changes strings in place in files or on the standard input.

Invoke `replace` in one of the following ways:

```
shell> replace from to [from to] ... -- file_name [file_name] ...
shell> replace from to [from to] ... < file_name
```

`from` represents a string to look for and `to` represents its replacement. There can be one or more pairs of strings.

Use the `--` option to indicate where the string-replacement list ends and the file names begin. In this case, any file named on the command line is modified in place, so you may want to make a copy of the original before converting it. `replace` prints a message indicating which of the input files it actually modifies.

If the `--` option is not given, `replace` reads the standard input and writes to the standard output.

`replace` uses a finite state machine to match longer strings first. It can be used to swap strings. For example, the following command swaps `a` and `b` in the given files, `file1` and `file2`:

```
shell> replace a b b a -- file1 file2 ...
```

The `replace` program is used by `mysql2mysql`. See [Section 4.7.1, “mysql2mysql — Convert mSQL Programs for Use with MySQL”](#).

`replace` supports the following options.

- `-?`, `-I`

Display a help message and exit.

- `-#debug_options`

Enable debugging.

- `-s`

Silent mode. Print less information what the program does.

- `-v`

Verbose mode. Print more information about what the program does.

- `-V`

Display version information and exit.

### 4.8.3. resolveip — Resolve Host name to IP Address or Vice Versa

The `resolveip` utility resolves host names to IP addresses and vice versa.

Invoke `resolveip` like this:

```
shell> resolveip [options] {host_name|ip-addr} ...
```

`resolveip` supports the following options.

- `--help [377]`, `--info [377]`, `-?`, `-I`

Display a help message and exit.

- `--silent [377]`, `-s`

Silent mode. Produce less output.

- `--version [377]`, `-V`

Display version information and exit.



---

# Chapter 5. MySQL Server Administration

## Table of Contents

5.1. The MySQL Server .....	380
5.1.1. Server Option and Variable Reference .....	380
5.1.2. Server Configuration Defaults .....	410
5.1.3. Server Command Options .....	412
5.1.4. Server System Variables .....	447
5.1.5. Using System Variables .....	564
5.1.6. Server Status Variables .....	576
5.1.7. Server SQL Modes .....	602
5.1.8. Server Plugins .....	609
5.1.9. IPv6 Support .....	613
5.1.10. Server-Side Help .....	617
5.1.11. Server Response to Signals .....	617
5.1.12. The Shutdown Process .....	618
5.2. MySQL Server Logs .....	619
5.2.1. Selecting General Query and Slow Query Log Output Destinations .....	620
5.2.2. The Error Log .....	622
5.2.3. The General Query Log .....	623
5.2.4. The Binary Log .....	624
5.2.5. The Slow Query Log .....	635
5.2.6. Server Log Maintenance .....	637
5.3. Managing Disk I/O and File Space for <a href="#">InnoDB</a> Tables .....	638
5.3.1. <a href="#">InnoDB</a> Disk I/O .....	638
5.3.2. File Space Management .....	639
5.3.3. <a href="#">InnoDB</a> Checkpoints .....	640
5.3.4. Defragmenting a Table .....	640
5.4. Creating and Using <a href="#">InnoDB</a> Tables and Indexes .....	640
5.4.1. Managing <a href="#">InnoDB</a> Tablespaces .....	641
5.4.2. Grouping DML Operations with Transactions .....	645
5.4.3. Converting Tables from <a href="#">MyISAM</a> to <a href="#">InnoDB</a> .....	645
5.4.4. <a href="#">AUTO_INCREMENT</a> Handling in <a href="#">InnoDB</a> .....	650
5.4.5. <a href="#">InnoDB</a> and <a href="#">FOREIGN KEY</a> Constraints .....	655
5.4.6. Working with <a href="#">InnoDB</a> Compressed Tables .....	656
5.4.7. <a href="#">InnoDB</a> File-Format Management .....	667
5.4.8. How <a href="#">InnoDB</a> Stores Variable-Length Columns .....	672
5.5. Online DDL for <a href="#">InnoDB</a> Tables .....	674
5.5.1. Overview of Online DDL .....	674
5.5.2. Performance and Concurrency Considerations for Online DDL .....	679
5.5.3. SQL Syntax for Online DDL .....	681
5.5.4. Combining or Separating DDL Statements .....	682
5.5.5. Examples of Online DDL .....	682
5.5.6. Implementation Details of Online DDL .....	702
5.5.7. How Crash Recovery Works with Online DDL .....	703
5.5.8. DDL for Partitioned Tables .....	704
5.5.9. Limitations of Online DDL .....	704
5.6. Running Multiple MySQL Instances on One Machine .....	705
5.6.1. Setting Up Multiple Data Directories .....	706
5.6.2. Running Multiple MySQL Instances on Windows .....	707
5.6.3. Running Multiple MySQL Instances on Unix .....	710
5.6.4. Using Client Programs in a Multiple-Server Environment .....	711
5.7. Tracing <a href="#">mysqld</a> Using DTrace .....	711
5.7.1. <a href="#">mysqld</a> DTrace Probe Reference .....	713



This chapter provides an overview of MySQL Server, the `mysqld` command that does most of the work in a MySQL installation. It covers administration topics for `mysqld` and for the MySQL tables and indexes containing your data:

- Configuring the server through command options and system variables. Checking the runtime status through status variables.
- Understanding server activity and troubleshooting problems with the server log files.
- Managing the physical aspects of MySQL server operations: disk space, memory, I/O, and CPU utilization related to [queries](#), [DML](#), and [DDL](#).
- Managing the logical structures within a MySQL server: databases, tables, tablespaces, and indexes.
- Managing multiple servers on a single machine.

For additional information on administrative topics, see also:

- [Chapter 6, Security](#)
- [Section 14.2.1.2, “Configuring InnoDB”](#) and [Section 14.2.2, “Administering InnoDB”](#)
- [Chapter 7, Backup and Recovery](#)
- [Chapter 16, Replication](#)
- [Chapter 17, Partitioning](#)

## 5.1. The MySQL Server

`mysqld` is the MySQL server. The following discussion covers these MySQL server configuration topics:

- Startup options that the server supports. You can specify these options on the command line, through configuration files, or both.
- Server system variables. These variables reflect the current state and values of the startup options, some of which can be modified while the server is running.
- Server status variables. These variables contain counters and statistics about runtime operation.
- How to set the server SQL mode. This setting modifies certain aspects of SQL syntax and semantics, for example for compatibility with code from other database systems, or to control the error handling for particular situations.
- The server shutdown process. There are performance and reliability considerations depending on the type of table (transactional or nontransactional) and whether you use replication.



### Note

Not all storage engines are supported by all MySQL server binaries and configurations. To find out how to determine which storage engines your MySQL server installation supports, see [Section 13.7.5.17, “SHOW ENGINES Syntax”](#).

### 5.1.1. Server Option and Variable Reference

The following table provides a list of all the command line options, server and status variables applicable within `mysqld`.

The table lists command-line options (Cmd-line), options valid in configuration files (Option file), server system variables (System Var), and status variables (Status var) in one unified list, with notification of where each option/variable is valid. If a server option set on the command line or in an option file differs from the name of the corresponding server system or status variable, the variable name is noted

immediately below the corresponding option. For status variables, the scope of the variable is shown (Scope) as either global, session, or both. Please see the corresponding sections for details on setting and using the options and variables. Where appropriate, a direct link to further information on the item as available.

**Table 5.1. Option/Variable Summary**

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">abort-slave-event-count [1971]</a>	Yes	Yes				
<a href="#">Aborted_clients [585]</a>				Yes	Global	No
<a href="#">Aborted_connects [585]</a>				Yes	Global	No
<a href="#">allow-suspicious-udfs [413]</a>	Yes	Yes				
<a href="#">ansi [413]</a>	Yes	Yes				
<a href="#">audit-log [824]</a>	Yes	Yes				
<a href="#">audit_log_buffer_size [824]</a>			Yes		Global	No
<a href="#">audit_log_file [825]</a>			Yes		Global	No
<a href="#">audit_log_flush [825]</a>			Yes		Global	Yes
<a href="#">audit_log_policy [826]</a>			Yes		Global	Yes
<a href="#">audit_log_rotate_on_size [826]</a>			Yes		Global	Yes
<a href="#">audit_log_strategy [827]</a>			Yes		Global	No
<a href="#">authentication_windows_log_level [463]</a>	Yes	Yes				
<a href="#">authentication_windows_use_principal_name [464]</a>	Yes	Yes				
<a href="#">auto_increment_increment [1965]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">auto_increment_offset [1968]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">autocommit [464]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">automatic_sp_privileges [464]</a>			Yes		Global	Yes
<a href="#">back_log [465]</a>	Yes	Yes	Yes		Global	No
<a href="#">basedir [414]</a>	Yes	Yes	Yes		Global	No
<a href="#">big-tables [414]</a>	Yes	Yes			Both	Yes
- Variable: <a href="#">big_tables</a>			Yes		Both	Yes
<a href="#">bind-address [414]</a>	Yes	Yes	Yes		Global	No
<a href="#">Binlog_cache_disk_use [586]</a>				Yes	Global	No
<a href="#">binlog_cache_size [588]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Binlog_cache_use [586]</a>				Yes	Global	No
<a href="#">binlog-checksum [2006]</a>	Yes	Yes				
<a href="#">binlog_checksum [2009]</a>			Yes		Global	Yes
<a href="#">binlog_direct_non_transactional_updates [2010]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">binlog-do-db [2004]</a>	Yes	Yes				
<a href="#">binlog-format [415]</a>	Yes	Yes			Both	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
- Variable: binlog_format			Yes		Both	Yes
binlog-ignore-db [2005]	Yes	Yes				
binlog_max_flush_queue_time [2011]			Yes		Global	Yes
binlog_order_commits [2012]			Yes		Global	Yes
binlog-row-event-max-size [2001]	Yes	Yes				
binlog_row_image [2012]	Yes		Yes		Both	Yes
binlog_rows_query_log_events [2014]			Yes		Both	Yes
binlog-rows-query-log-events [2007]	Yes	Yes				
- Variable: binlog_rows_query_log_events						
Binlog_stmt_cache_disk_use [586]				Yes	Global	No
binlog_stmt_cache_size [2014]	Yes	Yes	Yes		Global	Yes
Binlog_stmt_cache_use [586]				Yes	Global	No
bootstrap [416]	Yes	Yes				
bulk_insert_buffer_size [466]	Yes	Yes	Yes		Both	Yes
Bytes_received [586]				Yes	Both	No
Bytes_sent [586]				Yes	Both	No
character_set_client [467]			Yes		Both	Yes
character-set-client-handshake [416]	Yes	Yes				
character_set_connection [467]			Yes		Both	Yes
character_set_database [468] <sup>a</sup>			Yes		Both	Yes
character-set-filesystem [416]	Yes	Yes			Both	Yes
- Variable: character_set_filesystem			Yes		Both	Yes
character_set_results [468]			Yes		Both	Yes
character-set-server [417]	Yes	Yes			Both	Yes
- Variable: character_set_server			Yes		Both	Yes
character_set_system [469]			Yes		Global	No
character-sets-dir [416]	Yes	Yes			Global	No
- Variable: character_sets_dir			Yes		Global	No
chroot [417]	Yes	Yes				
collation_connection [469]			Yes		Both	Yes
collation_database [469] <sup>b</sup>			Yes		Both	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">collation-server [417]</a>	Yes	Yes			Both	Yes
- Variable: <a href="#">collation_server</a>			Yes		Both	Yes
<a href="#">Com_admin_commands [586]</a>				Yes	Both	No
<a href="#">Com_alter_db [586]</a>				Yes	Both	No
<a href="#">Com_alter_db_upgrade [586]</a>				Yes	Both	No
<a href="#">Com_alter_event [586]</a>				Yes	Both	No
<a href="#">Com_alter_function [586]</a>				Yes	Both	No
<a href="#">Com_alter_procedure [586]</a>				Yes	Both	No
<a href="#">Com_alter_server [586]</a>				Yes	Both	No
<a href="#">Com_alter_table [586]</a>				Yes	Both	No
<a href="#">Com_alter_tablespace [586]</a>				Yes	Both	No
<a href="#">Com_alter_user [586]</a>				Yes	Both	No
<a href="#">Com_analyze [586]</a>				Yes	Both	No
<a href="#">Com_assign_to_keycache [586]</a>				Yes	Both	No
<a href="#">Com_begin [586]</a>				Yes	Both	No
<a href="#">Com_binlog [586]</a>				Yes	Both	No
<a href="#">Com_call_procedure [586]</a>				Yes	Both	No
<a href="#">Com_change_db [586]</a>				Yes	Both	No
<a href="#">Com_change_master [586]</a>				Yes	Both	No
<a href="#">Com_check [586]</a>				Yes	Both	No
<a href="#">Com_checksum [586]</a>				Yes	Both	No
<a href="#">Com_commit [586]</a>				Yes	Both	No
<a href="#">Com_create_db [586]</a>				Yes	Both	No
<a href="#">Com_create_event [586]</a>				Yes	Both	No
<a href="#">Com_create_function [586]</a>				Yes	Both	No
<a href="#">Com_create_index [586]</a>				Yes	Both	No
<a href="#">Com_create_procedure [586]</a>				Yes	Both	No
<a href="#">Com_create_server [586]</a>				Yes	Both	No
<a href="#">Com_create_table [586]</a>				Yes	Both	No
<a href="#">Com_create_trigger [586]</a>				Yes	Both	No
<a href="#">Com_create_udf [586]</a>				Yes	Both	No
<a href="#">Com_create_user [586]</a>				Yes	Both	No
<a href="#">Com_create_view [586]</a>				Yes	Both	No
<a href="#">Com_dealloc_sql [586]</a>				Yes	Both	No
<a href="#">Com_delete [586]</a>				Yes	Both	No
<a href="#">Com_delete_multi [586]</a>				Yes	Both	No
<a href="#">Com_do [586]</a>				Yes	Both	No
<a href="#">Com_drop_db [586]</a>				Yes	Both	No
<a href="#">Com_drop_event [586]</a>				Yes	Both	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
Com_drop_function [586]				Yes	Both	No
Com_drop_index [586]				Yes	Both	No
Com_drop_procedure [586]				Yes	Both	No
Com_drop_server [586]				Yes	Both	No
Com_drop_table [586]				Yes	Both	No
Com_drop_trigger [586]				Yes	Both	No
Com_drop_user [586]				Yes	Both	No
Com_drop_view [586]				Yes	Both	No
Com_empty_query [586]				Yes	Both	No
Com_execute_sql [586]				Yes	Both	No
Com_flush [586]				Yes	Both	No
Com_get_diagnostics [586]				Yes	Both	No
Com_grant [586]				Yes	Both	No
Com_ha_close [586]				Yes	Both	No
Com_ha_open [586]				Yes	Both	No
Com_ha_read [586]				Yes	Both	No
Com_help [586]				Yes	Both	No
Com_insert [586]				Yes	Both	No
Com_insert_select [586]				Yes	Both	No
Com_install_plugin [586]				Yes	Both	No
Com_kill [586]				Yes	Both	No
Com_load [586]				Yes	Both	No
Com_lock_tables [586]				Yes	Both	No
Com_optimize [586]				Yes	Both	No
Com_preload_keys [586]				Yes	Both	No
Com_prepare_sql [586]				Yes	Both	No
Com_purge [586]				Yes	Both	No
Com_purge_before_date [586]				Yes	Both	No
Com_release_savepoint [586]				Yes	Both	No
Com_rename_table [586]				Yes	Both	No
Com_rename_user [586]				Yes	Both	No
Com_repair [586]				Yes	Both	No
Com_replace [586]				Yes	Both	No
Com_replace_select [586]				Yes	Both	No
Com_reset [586]				Yes	Both	No
Com_resignal [586]				Yes	Both	No
Com_revoke [586]				Yes	Both	No
Com_revoke_all [586]				Yes	Both	No
Com_rollback [586]				Yes	Both	No
Com_rollback_to_savepoint [586]				Yes	Both	No
Com_savepoint [586]				Yes	Both	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">Com_select [586]</a>				Yes	Both	No
<a href="#">Com_set_option [586]</a>				Yes	Both	No
<a href="#">Com_show_authors [586]</a>				Yes	Both	No
<a href="#">Com_show_binlog_events [586]</a>				Yes	Both	No
<a href="#">Com_show_binlogs [586]</a>				Yes	Both	No
<a href="#">Com_show_charsets [586]</a>				Yes	Both	No
<a href="#">Com_show_collations [586]</a>				Yes	Both	No
<a href="#">Com_show_contributors [586]</a>				Yes	Both	No
<a href="#">Com_show_create_db [586]</a>				Yes	Both	No
<a href="#">Com_show_create_event [586]</a>				Yes	Both	No
<a href="#">Com_show_create_func [586]</a>				Yes	Both	No
<a href="#">Com_show_create_proc [586]</a>				Yes	Both	No
<a href="#">Com_show_create_table [586]</a>				Yes	Both	No
<a href="#">Com_show_create_trigger [586]</a>				Yes	Both	No
<a href="#">Com_show_databases [586]</a>				Yes	Both	No
<a href="#">Com_show_engine_logs [586]</a>				Yes	Both	No
<a href="#">Com_show_engine_mutex [586]</a>				Yes	Both	No
<a href="#">Com_show_engine_status [586]</a>				Yes	Both	No
<a href="#">Com_show_errors [586]</a>				Yes	Both	No
<a href="#">Com_show_events [586]</a>				Yes	Both	No
<a href="#">Com_show_fields [586]</a>				Yes	Both	No
<a href="#">Com_show_function_code [586]</a>				Yes	Both	No
<a href="#">Com_show_function_status [586]</a>				Yes	Both	No
<a href="#">Com_show_grants [586]</a>				Yes	Both	No
<a href="#">Com_show_keys [586]</a>				Yes	Both	No
<a href="#">Com_show_logs [586]</a>				Yes	Both	No
<a href="#">Com_show_master_status [586]</a>				Yes	Both	No
<a href="#">Com_show_new_master [586]</a>				Yes	Both	No
<a href="#">Com_show_open_tables [586]</a>				Yes	Both	No
<a href="#">Com_show_plugins [586]</a>				Yes	Both	No
<a href="#">Com_show_privileges [586]</a>				Yes	Both	No
<a href="#">Com_show_procedure_code [586]</a>				Yes	Both	No
<a href="#">Com_show_procedure_status [586]</a>				Yes	Both	No
<a href="#">Com_show_processlist [586]</a>				Yes	Both	No
<a href="#">Com_show_profile [586]</a>				Yes	Both	No
<a href="#">Com_show_profiles [586]</a>				Yes	Both	No
<a href="#">Com_show_relaylog_events [586]</a>				Yes	Both	No
<a href="#">Com_show_slave_hosts [586]</a>				Yes	Both	No
<a href="#">Com_show_slave_status [586]</a>				Yes	Both	No
<a href="#">Com_show_status [586]</a>				Yes	Both	No
<a href="#">Com_show_storage_engines [586]</a>				Yes	Both	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">Com_show_table_status [586]</a>				Yes	Both	No
<a href="#">Com_show_tables [586]</a>				Yes	Both	No
<a href="#">Com_show_triggers [586]</a>				Yes	Both	No
<a href="#">Com_show_variables [586]</a>				Yes	Both	No
<a href="#">Com_show_warnings [586]</a>				Yes	Both	No
<a href="#">Com_signal [586]</a>				Yes	Both	No
<a href="#">Com_slave_start [586]</a>				Yes	Both	No
<a href="#">Com_slave_stop [586]</a>				Yes	Both	No
<a href="#">Com_stmt_close [586]</a>				Yes	Both	No
<a href="#">Com_stmt_execute [586]</a>				Yes	Both	No
<a href="#">Com_stmt_fetch [586]</a>				Yes	Both	No
<a href="#">Com_stmt_prepare [586]</a>				Yes	Both	No
<a href="#">Com_stmt_reprepare [586]</a>				Yes	Both	No
<a href="#">Com_stmt_reset [586]</a>				Yes	Both	No
<a href="#">Com_stmt_send_long_data [586]</a>				Yes	Both	No
<a href="#">Com_truncate [586]</a>				Yes	Both	No
<a href="#">Com_uninstall_plugin [586]</a>				Yes	Both	No
<a href="#">Com_unlock_tables [586]</a>				Yes	Both	No
<a href="#">Com_update [586]</a>				Yes	Both	No
<a href="#">Com_update_multi [586]</a>				Yes	Both	No
<a href="#">Com_xa_commit [586]</a>				Yes	Both	No
<a href="#">Com_xa_end [586]</a>				Yes	Both	No
<a href="#">Com_xa_prepare [586]</a>				Yes	Both	No
<a href="#">Com_xa_recover [586]</a>				Yes	Both	No
<a href="#">Com_xa_rollback [586]</a>				Yes	Both	No
<a href="#">Com_xa_start [586]</a>				Yes	Both	No
<a href="#">completion_type [476]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">Compression [587]</a>				Yes	Session	No
<a href="#">concurrent_insert [476]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">connect_timeout [476]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Connection_errors_accept [587]</a>				Yes	Global	No
<a href="#">Connection_errors_internal [587]</a>				Yes	Global	No
<a href="#">Connection_errors_max_connections [587]</a>				Yes	Global	No
<a href="#">Connection_errors_peer_addr [587]</a>				Yes	Global	No
<a href="#">Connection_errors_select [587]</a>				Yes	Global	No
<a href="#">Connection_errors_tcpwrap [587]</a>				Yes	Global	No
<a href="#">Connections [587]</a>				Yes	Global	No
<a href="#">console [417]</a>	Yes	Yes				
<a href="#">core_file [472]</a>			Yes		Global	No
<a href="#">core-file [418]</a>	Yes	Yes				
<a href="#">Created_tmp_disk_tables [587]</a>				Yes	Both	No



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">Created_tmp_files [588]</a>				Yes	Global	No
<a href="#">Created_tmp_tables [588]</a>				Yes	Both	No
<a href="#">daemon_memcached_enable_binlog [1698]</a>	Yes	Yes	Yes		Global	No
<a href="#">daemon_memcached_enable_engine_library [1698]</a>	Yes	Yes	Yes		Global	No
<a href="#">daemon_memcached_enable_engine_library_path [1699]</a>	Yes	Yes	Yes		Global	No
<a href="#">daemon_memcached_option [1699]</a>	Yes	Yes	Yes		Global	No
<a href="#">daemon_memcached_r_batch_size [1699]</a>	Yes	Yes	Yes		Global	No
<a href="#">daemon_memcached_w_batch_size [1700]</a>	Yes	Yes	Yes		Global	No
<a href="#">datadir [418]</a>	Yes	Yes	Yes		Global	No
<a href="#">date_format [472]</a>			Yes		Global	No
<a href="#">datetime_format [472]</a>			Yes		Global	No
<a href="#">debug [418]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">debug_sync [473]</a>			Yes		Session	Yes
<a href="#">debug-sync-timeout [419]</a>	Yes	Yes				
<a href="#">default-authentication-plugin [419]</a>	Yes	Yes				
<a href="#">default-storage-engine [420]</a>	Yes	Yes			Both	Yes
- Variable: <a href="#">default_storage_engine</a>			Yes		Both	Yes
<a href="#">default-time-zone [420]</a>	Yes	Yes				
<a href="#">default_tmp_storage_engine [474]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">default_week_format [475]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">defaults-extra-file [219]</a>	Yes					
<a href="#">defaults-file [219]</a>	Yes					
<a href="#">defaults-group-suffix [219]</a>	Yes					
<a href="#">delay-key-write [421]</a>	Yes	Yes			Global	Yes
- Variable: <a href="#">delay_key_write</a>			Yes		Global	Yes
<a href="#">Delayed_errors [588]</a>				Yes	Global	No
<a href="#">delayed_insert_limit [476]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Delayed_insert_threads [588]</a>				Yes	Global	No
<a href="#">delayed_insert_timeout [476]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">delayed_queue_size [477]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Delayed_writes [588]</a>				Yes	Global	No
<a href="#">des-key-file [421]</a>	Yes	Yes				
<a href="#">disable_gtid_unsafe_statements [2020]</a>	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">disable-gtid-unsafe-statements [2018]</a>	Yes	Yes	Yes		Global	No
<a href="#">disconnect_on_expire_passwords [477]</a>	Yes	Yes	Yes		Session	No
<a href="#">disconnect-slave-event-count [1971]</a>	Yes	Yes				
<a href="#">div_precision_increment [478]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">enable-named-pipe [421]</a>	Yes	Yes				
- Variable: named_pipe						
<a href="#">end_markers_in_json [479]</a>			Yes		Both	Yes
<a href="#">enforce_gtid_consistency [2021]</a>	Yes	Yes	Yes		Global	No
<a href="#">enforce-gtid-consistency [2018]</a>	Yes	Yes	Yes		Global	No
<a href="#">engine-condition-pushdown [421]</a>	Yes	Yes			Both	Yes
- Variable: engine_condition_pushdown			Yes		Both	Yes
<a href="#">eq_range_index_dive_limit [479]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">error_count [480]</a>			Yes		Session	No
<a href="#">event-scheduler [422]</a>	Yes	Yes			Global	Yes
- Variable: event_scheduler			Yes		Global	Yes
<a href="#">exit-info [422]</a>	Yes	Yes				
<a href="#">expire_logs_days [483]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">explicit_defaults_for_timestamp [481]</a>	Yes	Yes	Yes		Session	No
<a href="#">external-locking [422]</a>	Yes	Yes				
- Variable: skip_external_locking						
<a href="#">external_user [482]</a>			Yes		Session	No
<a href="#">federated</a>	Yes	Yes				
<a href="#">flush [482]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Flush_commands [588]</a>				Yes	Global	No
<a href="#">flush_time [482]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">foreign_key_checks [483]</a>			Yes		Both	Yes
<a href="#">ft_boolean_syntax [483]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">ft_max_word_len [483]</a>	Yes	Yes	Yes		Global	No
<a href="#">ft_min_word_len [483]</a>	Yes	Yes	Yes		Global	No
<a href="#">ft_query_expansion_limit [484]</a>	Yes	Yes	Yes		Global	No
<a href="#">ft_stopword_file [483]</a>	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">gdb [423]</a>	Yes	Yes				
<a href="#">general-log [423]</a>	Yes	Yes			Global	Yes
- Variable: <a href="#">general_log</a>			Yes		Global	Yes
<a href="#">general_log_file [486]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">group_concat_max_len [486]</a>	Yes	Yes	Yes	Yes	Both	Yes
<a href="#">gtid_done [2020]</a>			Yes		Both	No
<a href="#">gtid_executed [2021]</a>			Yes		Both	No
<a href="#">gtid_lost [2022]</a>			Yes		Global	No
<a href="#">gtid_mode [2022]</a>			Yes		Global	No
<a href="#">gtid-mode [2019]</a>	Yes	Yes			Global	No
- Variable: <a href="#">gtid_mode</a>			Yes		Global	No
<a href="#">gtid_next [2022]</a>			Yes		Session	Yes
<a href="#">gtid_owned [2023]</a>			Yes		Both	No
<a href="#">gtid_purged [2023]</a>			Yes		Global	Yes
<a href="#">Handler_commit [588]</a>				Yes	Both	No
<a href="#">Handler_delete [589]</a>				Yes	Both	No
<a href="#">Handler_discover</a>				Yes	Both	No
<a href="#">Handler_external_lock [589]</a>				Yes	Both	No
<a href="#">Handler_mrr_init [589]</a>				Yes	Both	No
<a href="#">Handler_prepare [589]</a>				Yes	Both	No
<a href="#">Handler_read_first [589]</a>				Yes	Both	No
<a href="#">Handler_read_key [589]</a>				Yes	Both	No
<a href="#">Handler_read_last [589]</a>				Yes	Both	No
<a href="#">Handler_read_next [589]</a>				Yes	Both	No
<a href="#">Handler_read_prev [589]</a>				Yes	Both	No
<a href="#">Handler_read_rnd [589]</a>				Yes	Both	No
<a href="#">Handler_read_rnd_next [589]</a>				Yes	Both	No
<a href="#">Handler_rollback [590]</a>				Yes	Both	No
<a href="#">Handler_savepoint [590]</a>				Yes	Both	No
<a href="#">Handler_savepoint_rollback [590]</a>				Yes	Both	No
<a href="#">Handler_update [590]</a>				Yes	Both	No
<a href="#">Handler_write [590]</a>				Yes	Both	No
<a href="#">have_compress [486]</a>			Yes		Global	No
<a href="#">have_crypt [486]</a>			Yes		Global	No
<a href="#">have_csv [486]</a>			Yes		Global	No
<a href="#">have_dynamic_loading [487]</a>			Yes		Global	No
<a href="#">have_geometry [487]</a>			Yes		Global	No
<a href="#">have_innodb [487]</a>			Yes		Global	No
<a href="#">have_ndbcluster</a>			Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">have_openssl [487]</a>			Yes		Global	No
<a href="#">have_partitioning [487]</a>			Yes		Global	No
<a href="#">have_profiling [487]</a>			Yes		Global	No
<a href="#">have_query_cache [487]</a>			Yes		Global	No
<a href="#">have_rtree_keys [487]</a>			Yes		Global	No
<a href="#">have_ssl [487]</a>			Yes		Global	No
<a href="#">have_symlink [487]</a>			Yes		Global	No
<a href="#">help [413]</a>	Yes	Yes				
<a href="#">host_cache_size [487]</a>		Yes	Yes		Global	Yes
<a href="#">hostname [488]</a>			Yes		Global	No
<a href="#">identity [488]</a>			Yes		Session	Yes
<a href="#">ignore-builtin-innodb [1697]</a>	Yes	Yes			Global	No
- Variable: <a href="#">ignore_builtin_innodb</a>			Yes		Global	No
<a href="#">ignore-db-dir [424]</a>	Yes	Yes				
<a href="#">ignore_db_dirs [488]</a>			Yes		Global	No
<a href="#">init_connect [489]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">init-file [424]</a>	Yes	Yes			Global	No
- Variable: <a href="#">init_file</a>			Yes		Global	No
<a href="#">init_slave [1987]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb [1698]</a>	Yes	Yes				
<a href="#">innodb_adaptive_flushing [1700]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_lwm [1701]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_index [1701]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_sleep_delay [1702]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_additional_pool_size [1702]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_api_bk_commit_interval [1702]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_api_disable_rowlock [1703]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_api_enable_binlog [1703]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_api_enable_ddl [1703]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_api_trx_level [1704]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_autoextend_increment [1704]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_autoinc_lock_mode [1705]</a>	Yes	Yes	Yes		Global	No
<a href="#">Innodb_available_undo_logs [590]</a>				Yes	Global	No
<a href="#">Innodb_buffer_pool_bytes_data [590]</a>				Yes	Global	No
<a href="#">Innodb_buffer_pool_bytes_dirty [590]</a>				Yes	Global	No
<a href="#">innodb_buffer_pool_dump_at_shutdown [1705]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_buffer_pool_dump_now [1706]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Innodb_buffer_pool_dump_status [590]</a>				Yes	Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
innodb_buffer_pool_filename [1706]	Yes	Yes	Yes		Global	Yes
innodb_buffer_pool_instances [1707]	Yes	Yes	Yes		Global	No
innodb_buffer_pool_read_abort [1707]	Yes	Yes	Yes		Global	Yes
innodb_buffer_pool_read_at_startup [1707]	Yes	Yes	Yes		Global	No
innodb_buffer_pool_read_now [1707]	Yes	Yes	Yes		Global	Yes
Innodb_buffer_pool_load_status [590]				Yes	Global	No
Innodb_buffer_pool_pages_data [590]				Yes	Global	No
Innodb_buffer_pool_pages_dirty [590]				Yes	Global	No
Innodb_buffer_pool_pages_flushed [591]				Yes	Global	No
Innodb_buffer_pool_pages_free [591]				Yes	Global	No
Innodb_buffer_pool_pages_latched [591]				Yes	Global	No
Innodb_buffer_pool_pages_misc [591]				Yes	Global	No
Innodb_buffer_pool_pages_total [591]				Yes	Global	No
Innodb_buffer_pool_read_ahead [591]				Yes	Global	No
Innodb_buffer_pool_read_ahead_evicted [591]				Yes	Global	No
Innodb_buffer_pool_read_requests [591]				Yes	Global	No
Innodb_buffer_pool_reads [591]				Yes	Global	No
innodb_buffer_pool_size [1708]	Yes	Yes	Yes		Global	No
Innodb_buffer_pool_wait_free [591]				Yes	Global	No
Innodb_buffer_pool_write_requests [591]				Yes	Global	No
innodb_change_buffer_max_size [1709]	Yes	Yes	Yes		Global	Yes
innodb_change_buffering [1709]	Yes	Yes	Yes		Global	Yes
innodb_checksum_algorithm [1710]	Yes	Yes	Yes		Global	Yes
innodb_checksums [1711]	Yes	Yes	Yes		Global	No
innodb_cmp_per_index_enabled [1712]	Yes	Yes	Yes		Global	Yes
innodb_commit_concurrency [1712]	Yes	Yes	Yes		Global	Yes
innodb_compression_failure_threshold_pct [1713]	Yes	Yes	Yes		Global	Yes
innodb_compression_level [1713]	Yes	Yes	Yes		Global	Yes
innodb_compression_pad_pct_max [1713]	Yes	Yes	Yes		Global	Yes
innodb_concurrency_tickets [1714]	Yes	Yes	Yes		Global	Yes
innodb_data_file_path [1714]	Yes	Yes	Yes		Global	No
Innodb_data_fsyncs [591]				Yes	Global	No
innodb_data_home_dir [1715]	Yes	Yes	Yes		Global	No
Innodb_data_pending_fsyncs [591]				Yes	Global	No
Innodb_data_pending_reads [592]				Yes	Global	No
Innodb_data_pending_writes [592]				Yes	Global	No
Innodb_data_read [592]				Yes	Global	No
Innodb_data_reads [592]				Yes	Global	No
Innodb_data_writes [592]				Yes	Global	No
Innodb_data_written [592]				Yes	Global	No
Innodb_dblwr_pages_written [592]				Yes	Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
InnoDB_dblwr_writes [592]				Yes	Global	No
innodb_disable_sort_file_cache [Yes 5]	Yes	Yes	Yes		Global	Yes
innodb_doublewrite [Yes 15]	Yes	Yes	Yes		Global	No
innodb_fast_shutdown [Yes 1716]	Yes	Yes	Yes		Global	Yes
innodb_file_format [Yes 16]	Yes	Yes	Yes		Global	Yes
innodb_file_format_check [Yes 1717]	Yes	Yes	Yes		Global	No
innodb_file_format_max [Yes x 1717]	Yes	Yes	Yes		Global	Yes
innodb_file_per_table [Yes 1718]	Yes	Yes	Yes		Global	Yes
innodb_flush_log_at_timeout [1718]			Yes		Global	Yes
innodb_flush_log_at_trx_commit [Yes 18]	Yes	Yes	Yes		Global	Yes
innodb_flush_method [Yes 1719]	Yes	Yes	Yes		Global	No
innodb_flush_neighbors [Yes 1721]	Yes	Yes	Yes		Global	Yes
innodb_flushing_avg_loops [1722]	Yes	Yes	Yes		Global	Yes
innodb_force_load_corrupted [Yes 1723]	Yes	Yes	Yes		Global	No
innodb_force_recovery [Yes 1722]	Yes	Yes	Yes		Global	No
innodb_ft_aux_table [Yes 1723]	Yes	Yes	Yes		Global	Yes
innodb_ft_cache_size [Yes 1724]	Yes	Yes	Yes		Global	No
innodb_ft_enable_stopword [Yes 1724]	Yes	Yes	Yes		Global	Yes
innodb_ft_max_token_size [Yes 1724]	Yes	Yes	Yes		Global	No
innodb_ft_min_token_size [Yes 1725]	Yes	Yes	Yes		Global	No
innodb_ft_num_word_optimize [Yes 1725]	Yes	Yes	Yes		Global	Yes
innodb_ft_server_stopword_table [Yes 1725]	Yes	Yes	Yes		Global	Yes
innodb_ft_sort_pll_degree [Yes 1726]	Yes	Yes	Yes		Global	No
innodb_ft_user_stopword_table [Yes 1726]	Yes	Yes	Yes		Both	Yes
InnoDB_have_atomic_builtins [592]				Yes	Global	No
innodb_io_capacity [Yes 1727]	Yes	Yes	Yes		Global	Yes
innodb_io_capacity_max [Yes 1728]	Yes	Yes	Yes		Global	Yes
innodb_large_prefix [Yes 1728]	Yes	Yes	Yes		Global	Yes
innodb_lock_wait_timeout [Yes 1729]	Yes	Yes	Yes		Both	Yes
innodb_locks_unsafe_for_binlog [Yes 1729]	Yes	Yes	Yes		Global	No
innodb_log_buffer_size [Yes 1732]	Yes	Yes	Yes		Global	No
innodb_log_compressed_pages [Yes 1732]	Yes	Yes	Yes		Global	Yes
innodb_log_file_size [Yes 1732]	Yes	Yes	Yes		Global	No
innodb_log_files_in_group [Yes 1733]	Yes	Yes	Yes		Global	No
innodb_log_group_home_dir [Yes 1733]	Yes	Yes	Yes		Global	No
InnoDB_log_waits [592]				Yes	Global	No
InnoDB_log_write_requests [592]				Yes	Global	No
InnoDB_log_writes [592]				Yes	Global	No
innodb_lru_scan_depth [Yes 1734]	Yes	Yes	Yes		Global	Yes
innodb_max_dirty_pages_pct [Yes 1735]	Yes	Yes	Yes		Global	Yes
innodb_max_dirty_pages_pct_lwm [Yes 1735]	Yes	Yes	Yes		Global	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
innodb_max_purge_age	Yes [1735]	Yes	Yes		Global	Yes
innodb_max_purge_delay	Yes [1735]	Yes	Yes		Global	Yes
innodb_mirrored_log_groups	Yes [1736]	Yes	Yes		Global	No
innodb_monitor_disable	Yes [1736]	Yes	Yes		Global	Yes
innodb_monitor_enable	Yes [1737]	Yes	Yes		Global	Yes
innodb_monitor_reset	Yes [1737]	Yes	Yes		Global	Yes
innodb_monitor_reset_all	Yes [1737]	Yes	Yes		Global	Yes
Innodb_num_open_files	[592]			Yes	Global	No
innodb_old_blocks_test	Yes [1738]	Yes	Yes		Global	Yes
innodb_old_blocks_time	Yes [1738]	Yes	Yes		Global	Yes
innodb_online_alter_log_max_size	Yes [1739]	Yes	Yes		Global	Yes
innodb_open_files	Yes [1739]	Yes	Yes		Global	No
innodb_optimize_fulltext_only	Yes [1740]	Yes	Yes		Global	Yes
Innodb_os_log_fsyncs	[592]			Yes	Global	No
Innodb_os_log_pending_fsyncs	[592]			Yes	Global	No
Innodb_os_log_pending_writes	[593]			Yes	Global	No
Innodb_os_log_written	[593]			Yes	Global	No
innodb_page_size	Yes [1740]	Yes	Yes		Global	No
Innodb_page_size	[593]			Yes	Global	No
Innodb_pages_created	[593]			Yes	Global	No
Innodb_pages_read	[593]			Yes	Global	No
Innodb_pages_written	[593]			Yes	Global	No
innodb_print_all_deadlocks	Yes [1741]	Yes	Yes		Global	Yes
innodb_purge_batch_size	Yes [1741]	Yes	Yes		Global	Yes
innodb_purge_threads	Yes [1742]	Yes	Yes		Global	No
innodb_random_read_ahead	Yes [1742]	Yes	Yes		Global	Yes
innodb_read_ahead_threshold	Yes [1742]	Yes	Yes		Global	Yes
innodb_read_io_threads	Yes [1743]	Yes	Yes		Global	No
innodb_read_only	Yes [1743]	Yes	Yes		Global	No
innodb_replication_delay	Yes [1744]	Yes	Yes		Global	Yes
innodb_rollback_on_timeout	Yes [1744]	Yes	Yes		Global	No
innodb_rollback_segments	Yes [1744]	Yes	Yes		Global	Yes
Innodb_row_lock_current_waits	[593]			Yes	Global	No
Innodb_row_lock_time	[593]			Yes	Global	No
Innodb_row_lock_time_avg	[593]			Yes	Global	No
Innodb_row_lock_time_max	[593]			Yes	Global	No
Innodb_row_lock_waits	[593]			Yes	Global	No
Innodb_rows_deleted	[593]			Yes	Global	No
Innodb_rows_inserted	[593]			Yes	Global	No
Innodb_rows_read	[593]			Yes	Global	No
Innodb_rows_updated	[593]			Yes	Global	No



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">innodb_sort_buffer_size [1745]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_spin_wait_delay [1745]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_auto_refresh [1746]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_method [1746]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_on_metadata [1747]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_persistent [1747]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_persistent_sample_pages [1748]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_sample_pages [1748]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_stats_transient_sample_pages [1748]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb-status-file [1698]</a>	Yes	Yes				
<a href="#">innodb_strict_mode [1749]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">innodb_support_xa [1749]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">innodb_sync_array_size [1750]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_sync_spin_ops [1750]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_table_locks [1751]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">innodb_thread_concurrency [1751]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_thread_sleep_delay [1752]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">InnoDB_truncated_status_writes [593]</a>				Yes	Global	No
<a href="#">innodb_undo_directory [1752]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_undo_logs [1752]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_undo_tablespaces [1753]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">innodb_use_native_aio [1753]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_use_sys_malloc [1754]</a>	Yes	Yes	Yes		Global	No
<a href="#">innodb_version [1754]</a>			Yes		Global	No
<a href="#">innodb_write_io_threads [1754]</a>	Yes	Yes	Yes		Global	No
<a href="#">insert_id [490]</a>			Yes		Session	Yes
<a href="#">install [424]</a>	Yes					
<a href="#">install-manual [424]</a>	Yes					
<a href="#">interactive_timeout [490]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">join_buffer_size [490]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">keep_files_on_create [491]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">Key_blocks_not_flushed [594]</a>				Yes	Global	No
<a href="#">Key_blocks_unused [594]</a>				Yes	Global	No
<a href="#">Key_blocks_used [594]</a>				Yes	Global	No
<a href="#">key_buffer_size [492]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">key_cache_age_threshold [493]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">key_cache_block_size [493]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">key_cache_division_limit [494]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Key_read_requests [594]</a>				Yes	Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">Key_reads [594]</a>				Yes	Global	No
<a href="#">Key_write_requests [594]</a>				Yes	Global	No
<a href="#">Key_writes [594]</a>				Yes	Global	No
<a href="#">language [425]</a>	Yes	Yes	Yes		Global	No
<a href="#">large_files_support [494]</a>			Yes		Global	No
<a href="#">large_page_size [494]</a>			Yes		Global	No
<a href="#">large-pages [425]</a>	Yes	Yes			Global	No
- Variable: <a href="#">large_pages</a>			Yes		Global	No
<a href="#">last_insert_id [495]</a>			Yes		Session	Yes
<a href="#">Last_query_cost [594]</a>				Yes	Session	No
<a href="#">Last_query_partial_plans [594]</a>				Yes	Session	No
<a href="#">lc-messages [426]</a>	Yes	Yes			Both	Yes
- Variable: <a href="#">lc_messages</a>			Yes		Both	Yes
<a href="#">lc-messages-dir [426]</a>	Yes	Yes			Global	No
- Variable: <a href="#">lc_messages_dir</a>			Yes		Global	No
<a href="#">lc_time_names [495]</a>			Yes		Both	Yes
<a href="#">license [496]</a>			Yes		Global	No
<a href="#">local_infile [496]</a>			Yes		Global	Yes
<a href="#">lock_wait_timeout [496]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">locked_in_memory [497]</a>			Yes		Global	No
<a href="#">log [426]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">log_bin [2008]</a>			Yes		Global	No
<a href="#">log-bin [2002]</a>	Yes	Yes	Yes		Global	No
<a href="#">log_bin_basename [2015]</a>			Yes		Global	No
<a href="#">log-bin-index [2002]</a>	Yes	Yes				
<a href="#">log_bin_use_v1_row_events [2015]</a>	Yes	Yes	Yes		Global	No
<a href="#">log-bin-use-v1-row-events [2003]</a>	Yes	Yes			Global	No
- Variable: <a href="#">log_bin_use_v1_row_events</a>			Yes		Global	No
<a href="#">log-error [427]</a>	Yes	Yes			Global	No
- Variable: <a href="#">log_error</a>			Yes		Global	No
<a href="#">log-isam [427]</a>	Yes	Yes				
<a href="#">log-output [427]</a>	Yes	Yes			Global	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
- Variable: log_output			Yes		Global	Yes
log-queries-not-using-indexes [428]	Yes	Yes			Global	Yes
- Variable: log_queries_not_using_indexes			Yes		Global	Yes
log-raw [428]	Yes	Yes				
log-short-format [428]	Yes	Yes				
log-slave-updates [1971]	Yes	Yes			Global	No
- Variable: log_slave_updates			Yes		Global	No
log_slave_updates [2008]	Yes	Yes	Yes		Global	No
log_slow_admin_statements [498]			Yes		Global	Yes
log-slow-admin-statements [429]	Yes	Yes				
log-slow-queries [429]	Yes	Yes			Global	Yes
- Variable: log_slow_queries			Yes		Global	Yes
log_slow_slave_statements [1988]			Yes		Global	Yes
log-slow-slave-statements [1972]	Yes	Yes				
log-tc [429]	Yes	Yes				
log-tc-size [429]	Yes	Yes				
log_throttle_queries_not_using_indexes [499]			Yes		Global	Yes
log-warnings [430]	Yes	Yes			Global	Yes
- Variable: log_warnings			Yes		Global	Yes
long_query_time [500]	Yes	Yes	Yes		Both	Yes
low-priority-updates [430]	Yes	Yes			Both	Yes
- Variable: low_priority_updates			Yes		Both	Yes
lower_case_file_system [501]	Yes	Yes	Yes		Global	No
lower_case_table_names [501]	Yes	Yes	Yes		Global	No
master-info-file [1973]	Yes	Yes				
master_info_repository [1988]			Yes		Global	Yes
master-info-repository [2000]	Yes	Yes				

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
- Variable: master_info_repository						
master-retry-count [1973]	Yes	Yes				
master_verify_checksum [2015]			Yes		Global	Yes
master-verify-checksum [2006]	Yes	Yes				
- Variable: master_verify_checksum						
max_allowed_packet [502]	Yes	Yes	Yes		Global	Yes
max_binlog_cache_size [2016]	Yes	Yes	Yes		Global	Yes
max-binlog-dump-events [2007]	Yes	Yes				
max_binlog_size [2016]	Yes	Yes	Yes		Global	Yes
max_binlog_stmt_cache_size [2017]	Yes	Yes	Yes		Global	Yes
max_connect_errors [502]	Yes	Yes	Yes		Global	Yes
max_connections [503]	Yes	Yes	Yes		Global	Yes
max_delayed_threads [504]	Yes	Yes	Yes		Both	Yes
max_error_count [505]	Yes	Yes	Yes		Both	Yes
max_heap_table_size [504]	Yes	Yes	Yes		Both	Yes
max_insert_delayed_threads [505]			Yes		Both	Yes
max_join_size [505]	Yes	Yes	Yes		Both	Yes
max_length_for_sort_data [506]	Yes	Yes	Yes		Both	Yes
max_prepared_stmt_count [506]	Yes	Yes	Yes		Global	Yes
max_relay_log_size [507]	Yes	Yes	Yes		Global	Yes
max_seeks_for_key [507]	Yes	Yes	Yes		Both	Yes
max_sort_length [508]	Yes	Yes	Yes		Both	Yes
max_sp_recursion_depth [508]	Yes	Yes	Yes		Both	Yes
Max_used_connections [594]				Yes	Global	No
max_user_connections [509]	Yes	Yes	Yes		Both	Yes
max_write_lock_count [509]	Yes	Yes	Yes		Global	Yes
memlock [431]	Yes	Yes	Yes		Global	No
metadata_locks_cache_size [510]			Yes		Global	No
metadata_locks_hash_instances [510]			Yes		Global	No
min-examined-row-limit [431]	Yes	Yes	Yes		Both	Yes
myisam-block-size [432]	Yes	Yes				
myisam_data_pointer_size [511]	Yes	Yes	Yes		Global	Yes
myisam_max_sort_size [511]	Yes	Yes	Yes		Global	Yes
myisam_mmap_size [512]	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">mysam-recover-options [432]</a>	Yes	Yes				
- Variable: mysam_recover_options						
<a href="#">mysam_recover_options [512]</a>			Yes		Global	No
<a href="#">mysam_repair_thread_size [512]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">mysam_sort_buffer_size [513]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">mysam_stats_method [513]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">mysam_use_mmapped [514]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">named_pipe [514]</a>			Yes		Global	No
<a href="#">ndb_cache_check_time [515]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Ndb_conflict_fn_max</a>				Yes	Global	No
<a href="#">Ndb_conflict_fn_old</a>				Yes	Global	No
<a href="#">Ndb_number_of_data_nodes</a>				Yes	Global	No
<a href="#">net_buffer_length [515]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">net_read_timeout [515]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">net_retry_count [515]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">net_write_timeout [516]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">new [516]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">no-defaults [220]</a>	Yes					
<a href="#">Not_flushed_delayed_rows [594]</a>				Yes	Global	No
<a href="#">old [517]</a>	Yes	Yes	Yes		Global	No
<a href="#">old-alter-table [433]</a>	Yes	Yes			Both	Yes
- Variable: old_alter_table			Yes		Both	Yes
<a href="#">old_passwords [517]</a>			Yes		Both	Yes
<a href="#">old-style-user-limits [433]</a>	Yes	Yes				
<a href="#">one-thread [434]</a>	Yes	Yes				
<a href="#">Open_files [594]</a>				Yes	Global	No
<a href="#">open-files-limit [434]</a>	Yes	Yes			Global	No
- Variable: open_files_limit			Yes		Global	No
<a href="#">Open_streams [595]</a>				Yes	Global	No
<a href="#">Open_table_definitions [595]</a>				Yes	Global	No
<a href="#">Open_tables [595]</a>				Yes	Both	No
<a href="#">Opened_files [595]</a>				Yes	Global	No
<a href="#">Opened_table_definitions [595]</a>				Yes	Both	No
<a href="#">Opened_tables [595]</a>				Yes	Both	No
<a href="#">optimizer_join_cache_level [519]</a>	Yes	Yes	Yes		Both	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">optimizer_prune_level [519]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">optimizer_search_depth [520]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">optimizer_switch [521]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">optimizer_trace [523]</a>			Yes		Both	Yes
<a href="#">optimizer_trace_features [523]</a>			Yes		Both	Yes
<a href="#">optimizer_trace_limit [524]</a>			Yes		Both	Yes
<a href="#">optimizer_trace_max_mem_size [524]</a>			Yes		Both	Yes
<a href="#">optimizer_trace_offset [524]</a>			Yes		Both	Yes
<a href="#">partition [434]</a>	Yes	Yes			Global	No
- Variable: <a href="#">have_partitioning</a>			Yes		Global	No
<a href="#">performance_schema [2307]</a>	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_accounts_lost [2318]</a>				Yes	Global	No
<a href="#">performance_schema_accounts_size [2308]</a>	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_cond_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_cond_instances_lost [2318]</a>				Yes	Global	No
<a href="#">performance_schema_events_statements_history_size [2308]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_size [2309]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_long_size [2309]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_size [2309]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_waits_history_long_size [2310]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_waits_history_size [2310]</a>	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_file_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_file_handles_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_file_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_hosts_lost [2318]</a>				Yes	Global	No
<a href="#">performance_schema_hosts_size [2310]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance-schema-instrument [2306]</a>	Yes	Yes				
<a href="#">Performance_schema_locker_lost [2319]</a>				Yes	Global	No
<a href="#">performance_schema_max_cond_classes [2311]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_cond_instances [2311]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_classes [2311]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_handles [2312]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_instances [2312]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_mutex_classes [2313]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_mutex_instances [2313]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_rwlock_classes [2313]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_rwlock_instances [2313]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_socket_classes [2314]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_socket_instances [2314]</a>	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">performance_schema_max_stage_classes [2314]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_statement_classes [2314]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_table_handles [2314]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_table_instances [2315]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_thread_classes [2316]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_thread_instances [2316]</a>	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_mutex_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_mutex_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_rwlock_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_rwlock_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_session_connect_attrs_lost [2319]</a>				Yes	Global	No
<a href="#">performance_schema_session_connect_attrs [2316]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_setup_actor_size [2317]</a>	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_setup_object_size [2317]</a>	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_socket_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_socket_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_stage_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_statement_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_table_handles_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_table_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_thread_classes_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_thread_instances_lost [2318]</a>				Yes	Global	No
<a href="#">Performance_schema_users_lost [2319]</a>				Yes	Global	No
<a href="#">performance_schema_users_size [2317]</a>	Yes	Yes	Yes		Global	No
<a href="#">pid-file [435]</a>	Yes	Yes			Global	No
- Variable: <a href="#">pid_file</a>			Yes		Global	No
<a href="#">plugin [435]</a>	Yes	Yes				
<a href="#">plugin_dir [525]</a>	Yes	Yes	Yes		Global	No
<a href="#">plugin-load [435]</a>	Yes	Yes				
<a href="#">plugin-load-add [436]</a>	Yes	Yes				
<a href="#">port [437]</a>	Yes	Yes	Yes		Global	No
<a href="#">port-open-timeout [437]</a>	Yes	Yes				
<a href="#">preload_buffer_size [525]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">Prepared_stmt_count [595]</a>				Yes	Global	No
<a href="#">print-defaults [220]</a>	Yes					
<a href="#">profiling [526]</a>			Yes		Both	Yes
<a href="#">profiling_history_size [526]</a>	Yes	Yes	Yes		Both	Yes



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">protocol_version [526]</a>			Yes		Global	No
<a href="#">proxy_user [526]</a>			Yes		Session	No
<a href="#">pseudo_slave_mode [526]</a>			Yes		Session	Yes
<a href="#">pseudo_thread_id [527]</a>			Yes		Session	Yes
<a href="#">Qcache_free_blocks [595]</a>				Yes	Global	No
<a href="#">Qcache_free_memory [595]</a>				Yes	Global	No
<a href="#">Qcache_hits [595]</a>				Yes	Global	No
<a href="#">Qcache_inserts [595]</a>				Yes	Global	No
<a href="#">Qcache_lowmem_prunes [595]</a>				Yes	Global	No
<a href="#">Qcache_not_cached [595]</a>				Yes	Global	No
<a href="#">Qcache_queries_in_cache [596]</a>				Yes	Global	No
<a href="#">Qcache_total_blocks [596]</a>				Yes	Global	No
<a href="#">Queries [596]</a>				Yes	Both	No
<a href="#">query_alloc_block_size [527]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">query_cache_limit [527]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">query_cache_min_size [528]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">query_cache_size [528]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">query_cache_type [529]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">query_cache_wlock_invalidate [530]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">query_prealloc_size [531]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">Questions [596]</a>				Yes	Both	No
<a href="#">rand_seed1 [531]</a>			Yes		Session	Yes
<a href="#">rand_seed2 [531]</a>			Yes		Session	Yes
<a href="#">range_alloc_block_size [531]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">read_buffer_size [532]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">read_only [532]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">read_rnd_buffer_size [533]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">relay-log [1974]</a>	Yes	Yes			Global	No
- Variable: relay_log			Yes		Global	No
<a href="#">relay_log_basename [1989]</a>			Yes		Global	No
<a href="#">relay-log-index [1974]</a>	Yes	Yes			Global	No
- Variable: relay_log_index			Yes		Global	No
<a href="#">relay_log_index [1989]</a>	Yes	Yes	Yes		Global	No
<a href="#">relay-log-info-file [1975]</a>	Yes	Yes				
- Variable: relay_log_info_file						
<a href="#">relay_log_info_file [1989]</a>	Yes	Yes	Yes		Global	No
<a href="#">relay-log-info-repository [2000]</a>	Yes	Yes				

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
- Variable: relay_log_info_repository						
relay_log_info_repository [1990]			Yes		Global	Yes
relay_log_purge [535]	Yes	Yes	Yes		Global	Yes
relay_log_recovery [1990]	Yes	Yes	Yes		Global	Yes
relay-log-recovery [1975]	Yes	Yes				
- Variable: relay_log_recovery						
relay_log_space_limit [534]	Yes	Yes	Yes		Global	No
remove [437]	Yes					
replicate-do-db [1976]	Yes	Yes				
replicate-do-table [1978]	Yes	Yes				
replicate-ignore-db [1977]	Yes	Yes				
replicate-ignore-table [1978]	Yes	Yes				
replicate-rewrite-db [1979]	Yes	Yes				
replicate-same-server-id [1979]	Yes	Yes				
replicate-wild-do-table [1979]	Yes	Yes				
replicate-wild-ignore-table [1980]	Yes	Yes				
report-host [1980]	Yes	Yes			Global	No
- Variable: report_host			Yes		Global	No
report-password [1981]	Yes	Yes			Global	No
- Variable: report_password			Yes		Global	No
report-port [1981]	Yes	Yes			Global	No
- Variable: report_port			Yes		Global	No
report-user [1982]	Yes	Yes			Global	No
- Variable: report_user			Yes		Global	No
Rpl_semi_sync_master_clients [596]				Yes	Global	No
rpl_semi_sync_master_enabled [536]			Yes		Global	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
Rpl_semi_sync_master_net_avg_wait_time [596]				Yes	Global	No
Rpl_semi_sync_master_net_wait_time [596]				Yes	Global	No
Rpl_semi_sync_master_net_waits [596]				Yes	Global	No
Rpl_semi_sync_master_no_times [596]				Yes	Global	No
Rpl_semi_sync_master_no_tx [596]				Yes	Global	No
Rpl_semi_sync_master_status [596]				Yes	Global	No
Rpl_semi_sync_master_timefunc_failures [597]				Yes	Global	No
rpl_semi_sync_master_timeout [536]			Yes		Global	Yes
rpl_semi_sync_master_trace_level [536]			Yes		Global	Yes
Rpl_semi_sync_master_tx_avg_wait_time [597]				Yes	Global	No
Rpl_semi_sync_master_tx_wait_time [597]				Yes	Global	No
Rpl_semi_sync_master_tx_waits [597]				Yes	Global	No
rpl_semi_sync_master_wait_no_slave [537]			Yes		Global	Yes
Rpl_semi_sync_master_wait_pos_backtraverse [597]				Yes	Global	No
Rpl_semi_sync_master_wait_sessions [597]				Yes	Global	No
Rpl_semi_sync_master_yes_tx [597]				Yes	Global	No
rpl_semi_sync_slave_enabled [537]			Yes		Global	Yes
Rpl_semi_sync_slave_status [597]				Yes	Global	No
rpl_semi_sync_slave_trace_level [538]			Yes		Global	Yes
Rsa_public_key [597]				Yes	Global	No
safe-mode [437]	Yes	Yes				
safe-user-create [438]	Yes	Yes				
secure-auth [438]	Yes	Yes			Global	Yes
- Variable: secure_auth			Yes		Global	Yes
secure-file-priv [438]	Yes	Yes			Global	No
- Variable: secure_file_priv			Yes		Global	No
Select_full_join [598]				Yes	Both	No
Select_full_range_join [598]				Yes	Both	No
Select_range [598]				Yes	Both	No
Select_range_check [598]				Yes	Both	No
Select_scan [598]				Yes	Both	No
server-id [1955]	Yes	Yes			Global	Yes
- Variable: server_id			Yes		Global	Yes
server_uuid [1955]			Yes		Global	No
sha256_password_private_key_path [539]			Yes		Global	No
sha256_password_public_key_path [539]			Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">shared_memory</a> [540]			Yes		Global	No
<a href="#">shared_memory_base_name</a> [540]			Yes		Global	No
<a href="#">show-slave-auth-info</a> [1982]	Yes	Yes				
<a href="#">skip-character-set-client-handshake</a> [416]	Yes	Yes				
<a href="#">skip-concurrent-insert</a> [439]	Yes	Yes				
- Variable: concurrent_insert						
<a href="#">skip-event-scheduler</a> [439]	Yes	Yes				
<a href="#">skip-external-locking</a> [439]	Yes	Yes			Global	No
- Variable: skip_external_locking			Yes		Global	No
<a href="#">skip-grant-tables</a> [439]	Yes	Yes				
<a href="#">skip-host-cache</a> [439]	Yes	Yes				
<a href="#">skip-log-warnings</a> [430]	Yes					
<a href="#">skip-name-resolve</a> [440]	Yes	Yes			Global	No
- Variable: skip_name_resolve			Yes		Global	No
<a href="#">skip-networking</a> [440]	Yes	Yes			Global	No
- Variable: skip_networking			Yes		Global	No
<a href="#">skip-new</a>	Yes	Yes				
<a href="#">skip-partition</a> [440]	Yes	Yes				
<a href="#">skip-show-database</a> [441]	Yes	Yes			Global	No
- Variable: skip_show_database			Yes		Global	No
<a href="#">skip-slave-start</a> [1984]	Yes	Yes				
<a href="#">skip-ssl</a> [810]	Yes	Yes				
<a href="#">skip-stack-trace</a> [441]	Yes	Yes				
<a href="#">skip-symbolic-links</a> [441]	Yes					
<a href="#">skip-thread-priority</a> [441]	Yes	Yes				

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">slave_allow_batching [1968]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">slave_checkpoint_group [1991]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">slave-checkpoint-group [1982]</a>	Yes	Yes				
- Variable: slave_checkpoint_group						
<a href="#">slave_checkpoint_period [1991]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">slave-checkpoint-period [1983]</a>	Yes	Yes				
- Variable: slave_checkpoint_period						
<a href="#">slave_compressed_protocol [1992]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">slave_exec_mode [1992]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">Slave_heartbeat_period [598]</a>				Yes	Global	No
<a href="#">Slave_last_heartbeat [598]</a>				Yes	Global	No
<a href="#">slave-load-tmpdir [1985]</a>	Yes	Yes			Global	No
- Variable: slave_load_tmpdir			Yes		Global	No
<a href="#">slave_max_allowed_packet [1969]</a>			Yes		Global	Yes
<a href="#">slave-max-allowed-packet [1964]</a>	Yes	Yes				
- Variable: slave_max_allowed_packet						
<a href="#">slave-net-timeout [1985]</a>	Yes	Yes			Global	Yes
- Variable: slave_net_timeout			Yes		Global	Yes
<a href="#">Slave_open_temp_tables [598]</a>				Yes	Global	No
<a href="#">slave_parallel_workers [1993]</a>			Yes		Global	Yes
<a href="#">slave-parallel-workers [1983]</a>	Yes	Yes				
- Variable: slave_parallel_workers						
<a href="#">slave_pending_jobs_size_max [1994]</a>			Yes		Global	Yes
<a href="#">slave-pending-jobs-size-max [1984]</a>	Yes					
- Variable: slave_pending_jobs_size_max						
<a href="#">Slave_received_heartbeats [598]</a>				Yes	Global	No
<a href="#">Slave_retried_transactions [598]</a>				Yes	Global	No
<a href="#">slave-rows-search-algorithms [1964]</a>	Yes	Yes				

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
- Variable: slave_rows_search_algorithms						
slave_rows_search_algorithms [1969]			Yes		Global	Yes
Slave_running [598]				Yes	Global	No
slave-skip-errors [1986]	Yes	Yes			Global	No
- Variable: slave_skip_errors			Yes		Global	No
slave_sql_verify_checksum [1995]			Yes		Global	Yes
slave-sql-verify-checksum [1987]	Yes	Yes				
slave_transaction_retries [1995]	Yes	Yes	Yes		Global	Yes
slave_type_convert [1996]	Yes	Yes	Yes		Global	No
Slow_launch_threads [598]				Yes	Both	No
slow_launch_time [544]	Yes	Yes	Yes		Global	Yes
Slow_queries [599]				Yes	Both	No
slow-query-log [441]	Yes	Yes			Global	Yes
- Variable: slow_query_log			Yes		Global	Yes
slow_query_log_file [541]	Yes	Yes	Yes		Global	Yes
slow-start-timeout [442]	Yes	Yes				
socket [442]	Yes	Yes	Yes		Global	No
sort_buffer_size [542]	Yes	Yes	Yes		Both	Yes
Sort_merge_passes [599]				Yes	Both	No
Sort_range [599]				Yes	Both	No
Sort_rows [599]				Yes	Both	No
Sort_scan [599]				Yes	Both	No
sporadic-binlog-dump-fail [2007]	Yes	Yes				
sql_auto_is_null [543]			Yes		Both	Yes
sql_big_selects [543]			Yes		Both	Yes
sql_big_tables [414]			Yes		Both	Yes
sql_buffer_result [544]			Yes		Both	Yes
sql_log_bin [544]			Yes		Both	Yes
sql_log_off [544]			Yes		Both	Yes
sql_low_priority_updates [500]			Yes		Both	Yes
sql_max_join_size [505]			Yes		Both	Yes
sql-mode [442]	Yes	Yes			Both	Yes
- Variable: sql_mode			Yes		Both	Yes
sql_notes [546]			Yes		Both	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">sql_quote_show_create [546]</a>			Yes		Both	Yes
<a href="#">sql_safe_updates [547]</a>			Yes		Both	Yes
<a href="#">sql_select_limit [547]</a>			Yes		Both	Yes
<a href="#">sql_slave_skip_counter [1996]</a>			Yes		Global	Yes
<a href="#">sql_warnings [547]</a>			Yes		Both	Yes
<a href="#">ssl [810]</a>	Yes	Yes				
<a href="#">Ssl_accept_renegotiates [599]</a>				Yes	Global	No
<a href="#">Ssl_accepts [599]</a>				Yes	Global	No
<a href="#">ssl-ca [810]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_ca</code>			Yes		Global	No
<a href="#">Ssl_callback_cache_hits [599]</a>				Yes	Global	No
<a href="#">ssl-capath [810]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_capath</code>			Yes		Global	No
<a href="#">ssl-cert [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_cert</code>			Yes		Global	No
<a href="#">ssl-cipher [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_cipher</code>			Yes		Global	No
<a href="#">Ssl_cipher [599]</a>				Yes	Both	No
<a href="#">Ssl_cipher_list [599]</a>				Yes	Both	No
<a href="#">Ssl_client_connects [599]</a>				Yes	Global	No
<a href="#">Ssl_connect_renegotiates [599]</a>				Yes	Global	No
<a href="#">ssl-crl [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_crl</code>			Yes		Global	No
<a href="#">ssl-crlpath [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_crlpath</code>			Yes		Global	No
<a href="#">Ssl_ctx_verify_depth [599]</a>				Yes	Global	No
<a href="#">Ssl_ctx_verify_mode [599]</a>				Yes	Global	No
<a href="#">Ssl_default_timeout [599]</a>				Yes	Both	No
<a href="#">Ssl_finished_accepts [599]</a>				Yes	Global	No
<a href="#">Ssl_finished_connects [600]</a>				Yes	Global	No
<a href="#">ssl-key [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_key</code>			Yes		Global	No
<a href="#">Ssl_server_not_after [600]</a>				Yes	Both	No
<a href="#">Ssl_server_not_before [600]</a>				Yes	Both	No
<a href="#">Ssl_session_cache_hits [600]</a>				Yes	Global	No
<a href="#">Ssl_session_cache_misses [600]</a>				Yes	Global	No



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
Ssl_session_cache_mode [600]				Yes	Global	No
Ssl_session_cache_overflows [600]				Yes	Global	No
Ssl_session_cache_size [600]				Yes	Global	No
Ssl_session_cache_timeouts [600]				Yes	Global	No
Ssl_sessions_reused [600]				Yes	Both	No
Ssl_used_session_cache_entries [600]				Yes	Global	No
Ssl_verify_depth [600]				Yes	Both	No
Ssl_verify_mode [600]				Yes	Both	No
ssl-verify-server-cert [811]	Yes	Yes				
Ssl_version [600]				Yes	Both	No
standalone [440]	Yes	Yes				
storage_engine [549]			Yes		Both	Yes
stored_program_cache_size [549]	Yes	Yes	Yes		Global	Yes
symbolic-links [441]	Yes	Yes				
sync_binlog [2017]	Yes	Yes	Yes		Global	Yes
sync_frm [550]	Yes	Yes	Yes		Global	Yes
sync_master_info [597]	Yes	Yes	Yes		Global	Yes
sync_relay_log [1999]	Yes	Yes	Yes		Global	Yes
sync_relay_log_info [1999]	Yes	Yes	Yes		Global	Yes
sysdate-is-now [444]	Yes	Yes				
system_time_zone [550]			Yes		Global	No
table_definition_cache [550]	Yes	Yes	Yes		Global	Yes
Table_locks_immediate [600]				Yes	Global	No
Table_locks_waited [600]				Yes	Global	No
table_open_cache [551]	Yes	Yes	Yes		Global	Yes
Table_open_cache_hits [601]				Yes	Both	No
table_open_cache_instances [552]	Yes	Yes	Yes		Global	No
Table_open_cache_misses [601]				Yes	Both	No
Table_open_cache_overflows [601]				Yes	Both	No
tc-heuristic-recover [444]	Yes	Yes				
Tc_log_max_pages_used [601]				Yes	Global	No
Tc_log_page_size [601]				Yes	Global	No
Tc_log_page_waits [601]				Yes	Global	No
temp-pool [445]	Yes	Yes				
thread_cache_size [552]	Yes	Yes	Yes		Global	Yes
thread_concurrency [553]	Yes	Yes	Yes		Global	No
thread_handling [553]	Yes	Yes	Yes		Global	No
thread_pool_algorithm [554]	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">thread_pool_high_priority_connections [554]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">thread_pool_max_disconnected_threads [555]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">thread_pool_priority_timer [555]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">thread_pool_size [556]</a>	Yes	Yes	Yes		Global	No
<a href="#">thread_pool_stall_time [556]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">thread_stack [556]</a>	Yes	Yes	Yes		Global	No
<a href="#">Threads_cached [601]</a>				Yes	Global	No
<a href="#">Threads_connected [601]</a>				Yes	Global	No
<a href="#">Threads_created [601]</a>				Yes	Global	No
<a href="#">Threads_running [601]</a>				Yes	Global	No
<a href="#">time_format [557]</a>			Yes		Global	No
<a href="#">time_zone [557]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">timed_mutexes [557]</a>	Yes	Yes	Yes		Global	Yes
<a href="#">timestamp [558]</a>			Yes		Session	Yes
<a href="#">tmp_table_size [558]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">tmpdir [446]</a>	Yes	Yes	Yes		Global	No
<a href="#">transaction_alloc_block_size [559]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">transaction-isolation [445]</a>	Yes	Yes				
- Variable: tx_isolation						
<a href="#">transaction_prealloc_size [560]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">transaction-read-only [445]</a>	Yes	Yes				
- Variable: tx_read_only						
<a href="#">tx_isolation [560]</a>			Yes		Both	Yes
<a href="#">tx_read_only [561]</a>			Yes		Both	Yes
<a href="#">unique_checks [561]</a>			Yes		Both	Yes
<a href="#">updatable_views_with_limit [562]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">Uptime [602]</a>				Yes	Global	No
<a href="#">Uptime_since_flush_status [602]</a>				Yes	Global	No
<a href="#">user [446]</a>	Yes	Yes				
<a href="#">validate-password [742]</a>	Yes	Yes				
<a href="#">validate_password_dictionary_file [743]</a>			Yes		Global	No
<a href="#">validate_password_length [744]</a>			Yes		Global	Yes
<a href="#">validate_password_mixed_case_count [744]</a>			Yes		Global	Yes
<a href="#">validate_password_number_count [745]</a>			Yes		Global	Yes
<a href="#">validate_password_policy [745]</a>			Yes		Global	Yes
<a href="#">validate_password_special_char_count [745]</a>			Yes		Global	Yes
<a href="#">validate_user_plugins [562]</a>			Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">verbose [447]</a>	Yes	Yes				
<a href="#">version [563]</a>			Yes		Global	No
<a href="#">version_comment [563]</a>			Yes		Global	No
<a href="#">version_compile_machine [563]</a>			Yes		Global	No
<a href="#">version_compile_os [563]</a>			Yes		Global	No
<a href="#">wait_timeout [563]</a>	Yes	Yes	Yes		Both	Yes
<a href="#">warning_count [564]</a>			Yes		Session	No

<sup>a</sup>This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.

<sup>b</sup>This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.

## 5.1.2. Server Configuration Defaults

The MySQL server has many operating parameters, which you can change at server startup using command-line options or configuration files (option files). It is also possible to change many parameters at runtime. For general instructions on setting parameters at startup or runtime, see [Section 5.1.3, “Server Command Options”](#), and [Section 5.1.4, “Server System Variables”](#).

### 5.1.2.1. Changes to Server Defaults

Beginning with MySQL 5.6.6, several MySQL Server parameter defaults differ from the defaults in previous releases. The motivation for these changes is to provide better out-of-box performance and to reduce the need for database administrators to change settings manually.

In some cases, a parameter has a different fixed default value. In other cases, the server autosizes a parameter at startup using a formula based on other related parameters or server host configuration, rather than using a fixed value. For example, the setting for [back\\_log \[465\]](#) is its previous default of 50, adjusted up by an amount proportional to the value of [max\\_connections \[503\]](#). The idea behind autosizing is that when the server has information available to make a decision about a parameter setting likely to be better than a fixed default, it will.

The following table summarizes changes to defaults. The Version column indicates when each default changed. For variables that are autosized, the main variable description provides additional detail about the sizing algorithm. See [Section 5.1.4, “Server System Variables”](#), and [Section 14.2.6, “InnoDB Startup Options and System Variables”](#). Any of these default settings can be overridden by specifying an explicit value at server startup.

**Table 5.2. Changes to Server Defaults in MySQL 5.6**

Parameter	Old Default	New Default	Version
<a href="#">back_log [465]</a>	50	Autosized using <a href="#">max_connections [503]</a>	5.6.6
<a href="#">binlog_checksum [2009]</a>	NONE	CRC32	5.6.6
<a href="#">--binlog-row-event-max-size [2001]</a>	1024	8192	5.6.6
<a href="#">flush_time [482]</a>	1800 (on Windows)	0	5.6.6
<a href="#">host_cache_size [487]</a>	128	Autosized using <a href="#">max_connections [503]</a>	5.6.8
<a href="#">innodb_autoextend_increment [1708]</a>	8	64	5.6.6
<a href="#">innodb_buffer_pool_instances [1706]</a>	1	8 (platform dependent)	5.6.6
<a href="#">innodb_concurrency_tickets [1714]</a>	500	5000	5.6.6
<a href="#">innodb_data_file_path [1714]</a>	ibdata1:10M:autoextend	ibdata1:12M:autoextend	5.6.7

Parameter	Old Default	New Default	Version
<a href="#">innodb_file_per_table [1718]</a>	0	1	5.6.6
<a href="#">innodb_log_file_size [1732]</a>	5MB	48MB	5.6.8
<a href="#">innodb_old_blocks_time [1738]</a>	0	1000	5.6.6
<a href="#">innodb_open_files [1739]</a>	300	Autosized using <a href="#">innodb_file_per_table [1718]</a> , <a href="#">table_open_cache [551]</a>	5.6.6
<a href="#">innodb_stats_on_metadata [1747]</a>	ON	OFF	5.6.6
<a href="#">join_buffer_size [490]</a>	128KB	256KB	5.6.6
<a href="#">max_allowed_packet [502]</a>	1MB	4MB	5.6.6
<a href="#">max_connect_errors [502]</a>	10	100	5.6.6
<a href="#">open_files_limit [518]</a>	0	Autosized using <a href="#">max_connections [503]</a>	5.6.8
<a href="#">performance_schema [2307]</a>	OFF	ON	5.6.6
<a href="#">performance_schema_events_waits_history_long_size [2310]</a>	10000	Autosize	5.6.6
<a href="#">performance_schema_events_waits_history_size [2310]</a>	10	Autosize	5.6.6
<a href="#">performance_schema_max_cond_instances [2311]</a>	10000	Autosize	5.6.6
<a href="#">performance_schema_max_file_instances [2312]</a>	100000	Autosize	5.6.6
<a href="#">performance_schema_max_mutex_instances [2313]</a>	1000000	Autosize	5.6.6
<a href="#">performance_schema_max_rwlock_instances [2313]</a>	1000000	Autosize	5.6.6
<a href="#">performance_schema_max_table_handles [2315]</a>	100000	Autosize	5.6.6
<a href="#">performance_schema_max_table_instances [2315]</a>	50000	Autosize	5.6.6
<a href="#">performance_schema_max_thread_instances [2316]</a>	1000	Autosize	5.6.6
<a href="#">query_cache_size [528]</a>	0	1M	5.6.8
<a href="#">query_cache_type [529]</a>	ON	OFF	5.6.8
<a href="#">secure_auth [538]</a>	OFF	ON	5.6.7
<a href="#">sql_mode [545]</a>	" (empty string)	NO_ENGINE_SUBSTITUTION	5.6.6
<a href="#">sync_master_info [1997]</a>	0	10000	5.6.6
<a href="#">sync_relay_log [1998]</a>	0	10000	5.6.6
<a href="#">sync_relay_log_info [1999]</a>	0	10000	5.6.6
<a href="#">table_definition_cache [550]</a>	400	Autosized using <a href="#">table_open_cache [551]</a>	5.6.8
<a href="#">table_open_cache [551]</a>	400	2000	5.6.8
<a href="#">thread_cache_size [552]</a>	0	Autosized using <a href="#">max_connections [503]</a>	5.6.8

In MySQL 5.6.6, the default for [innodb\\_checksum\\_algorithm \[1710\]](#) was changed from `INNODB` to `CRC32`. For compatibility reasons, the default was returned in 5.6.7 to `INNODB`.

### 5.1.2.2. Using a Sample Default Server Configuration File

As of MySQL 5.6.8, on Unix platforms, `mysql_install_db` creates a default option file named `my.cnf` in the base installation directory. This file is created from a template included in the distribution package named `my-default.cnf`. You can find the template in or under the base installation directory. When started using `mysqld_safe`, the server uses `my.cnf` file by default. If `my.cnf` already exists, `mysql_install_db` assumes it to be in use and writes a new file named `my-new.cnf` instead.

With one exception, the settings in the default option file are commented and have no effect. The exception is that the file changes the `sql_mode` [545] system variable from its default of `NO_ENGINE_SUBSTITUTION` [605] to also include `STRICT_TRANS_TABLES` [607]. This setting produces a server configuration that results in errors rather than warnings for bad data in operations that modify transactional tables. See [Section 5.1.7, “Server SQL Modes”](#).

The `my-default.cnf` template replaces the older sample option files formerly supplied with MySQL distributions (`my-small.cnf`, `my-medium.cnf`, `my-large.cnf`, and `my-huge.cnf`). As of MySQL 5.6.8, these older files are no longer distributed.

On Windows, MySQL Installer interacts with the user and creates a file named `my.ini` in the base installation directory as the default option file. If you install on Windows from a Zip archive, you can copy the `my-default.ini` template file in the base installation directory to `my.ini` and use the latter as the default option file.



#### Note

On Windows, the `.ini` or `.cnf` option file extension might not be displayed.

On any platform, after completing the installation process, you can edit the default option file at any time to modify the parameters used by the server. For example, to use a parameter setting in the file that is commented with a `#` character at the beginning of the line, remove the `#`, and modify the parameter value if necessary. To disable a setting, either add a `#` to the beginning of the line or remove it.

For additional information about option file format and syntax, see [Section 4.2.3.3, “Using Option Files”](#).

Before MySQL 5.6.8, MySQL distributions include several sample option files that can be used as a basis for tuning the MySQL server. Look for files named `my-small.cnf`, `my-medium.cnf`, `my-large.cnf`, and `my-huge.cnf`, which are sample files for small, medium, large, and very large systems. On Windows, the extension is `.ini` rather than `.cnf`.

For a binary distribution, look for the sample files in or under your installation directory. If you have a source distribution, look in the `support-files` directory. To use a sample file as a base configuration file, rename a copy of it and place the copy in the appropriate location. Regarding names and appropriate location, see the general information provided in [Section 4.2.3.3, “Using Option Files”](#).

### 5.1.3. Server Command Options

When you start the `mysqld` server, you can specify program options using any of the methods described in [Section 4.2.3, “Specifying Program Options”](#). The most common methods are to provide options in an option file or on the command line. However, in most cases it is desirable to make sure that the server uses the same options each time it runs. The best way to ensure this is to list them in an option file. See [Section 4.2.3.3, “Using Option Files”](#). That section also describes option file format and syntax.

`mysqld` reads options from the `[mysqld]` and `[server]` groups. `mysqld_safe` reads options from the `[mysqld]`, `[server]`, `[mysqld_safe]`, and `[safe_mysqld]` groups. `mysql.server` reads options from the `[mysqld]` and `[mysql.server]` groups.

An embedded MySQL server usually reads options from the `[server]`, `[embedded]`, and `[xxxxx_SERVER]` groups, where `xxxxx` is the name of the application into which the server is embedded.

`mysqld` accepts many command options. For a brief summary, execute `mysqld --help`. To see the full list, use `mysqld --verbose --help`.

The following list shows some of the most common server options. Additional options are described in other sections:

- Options that affect security: See [Section 6.1.4, “Security-Related `mysqld` Options and Variables”](#).

- SSL-related options: See [Section 6.3.9.4, “SSL Command Options”](#).
- Binary log control options: See [Section 5.2.4, “The Binary Log”](#).
- Replication-related options: See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).
- Options for loading plugins such as pluggable storage engines: See [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).
- Options specific to particular storage engines: See [Section 14.2.6, “InnoDB Startup Options and System Variables”](#) and [Section 14.3.1, “MyISAM Startup Options”](#).

You can also set the values of server system variables by using variable names as options, as described at the end of this section.

Some options control the size of buffers or caches. For a given buffer, the server might need to allocate internal data structures. These structures typically are allocated from the total memory allocated to the buffer, and the amount of space required might be platform dependent. This means that when you assign a value to an option that controls a buffer size, the amount of space actually available might differ from the value assigned. In some cases, the amount might be less than the value assigned. It is also possible that the server will adjust a value upward. For example, if you assign a value of 0 to an option for which the minimal value is 1024, the server will set the value to 1024.

Values for buffer sizes, lengths, and stack sizes are given in bytes unless otherwise specified.

Some options take file name values. Unless otherwise specified, the default file location is the data directory if the value is a relative path name. To specify the location explicitly, use an absolute path name. Suppose that the data directory is `/var/mysql/data`. If a file-valued option is given as a relative path name, it will be located under `/var/mysql/data`. If the value is an absolute path name, its location is as given by the path name.

- `--help [413], -?`

<b>Command-Line Format</b>	<code>-?</code>
	<code>--help</code>
<b>Option-File Format</b>	<code>help</code>

Display a short help message and exit. Use both the `--verbose [447]` and `--help [413]` options to see the full message.

- `--allow-suspicious-udfs [413]`

<b>Command-Line Format</b>	<code>--allow-suspicious-udfs</code>	
<b>Option-File Format</b>	<code>allow-suspicious-udfs</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

This option controls whether user-defined functions that have only an `xxx` symbol for the main function can be loaded. By default, the option is off and only UDFs that have at least one auxiliary symbol can be loaded; this prevents attempts at loading functions from shared object files other than those containing legitimate UDFs. See [Section 22.3.2.6, “User-Defined Function Security Precautions”](#).

- `--ansi [413]`

<b>Command-Line Format</b>	<code>--ansi</code>
----------------------------	---------------------

	-a
<b>Option-File Format</b>	ansi

Use standard (ANSI) SQL syntax instead of MySQL syntax. For more precise control over the server SQL mode, use the `--sql-mode` [442] option instead. See [Section 1.8.3, “Running MySQL in ANSI Mode”](#), and [Section 5.1.7, “Server SQL Modes”](#).

- `--basedir=path` [414], `-b path` [414]

<b>Command-Line Format</b>	<code>--basedir=path</code>	
	<code>-b</code>	
<b>Option-File Format</b>	basedir	
<b>Option Sets Variable</b>	Yes, <code>basedir</code> [466]	
<b>Variable Name</b>	basedir	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path to the MySQL installation directory. All paths are usually resolved relative to this directory.

- `--big-tables` [414]

<b>Version Removed</b>	5.6.1	
<b>Command-Line Format</b>	<code>--big-tables</code>	
<b>Option-File Format</b>	big-tables	
<b>Option Sets Variable</b>	Yes, <code>big_tables</code> [466]	
<b>Variable Name</b>	big-tables	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean

Enable large result sets by saving all temporary sets in files. This option prevents most “table full” errors, but also slows down queries for which in-memory tables would suffice. Since MySQL 3.23.2, the server is able to handle large result sets automatically by using memory for small temporary tables and switching to disk tables where necessary.

- `--bind-address=addr` [414]

<b>Command-Line Format</b>	<code>--bind-address=addr</code>	
<b>Option-File Format</b>	bind-address=addr	
<b>Variable Name</b>	bind-address	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	string
	<b>Default</b>	0.0.0.0
	<b>Permitted Values (&gt;= 5.6.6)</b>	



<b>Type</b>	string
<b>Default</b>	*

The MySQL server listens on a single network socket for TCP/IP connections. This socket is bound to a single address, but it is possible for an address to map onto multiple network interfaces. To specify an address, use the `--bind-address=addr` [414] option at server startup, where `addr` is an IPv4 or IPv6 address or a host name. If `addr` is a host name, the server resolves the name to an IP address and binds to that address.

The server treats different types of addresses as follows:

- If the address is `*`, the server accepts TCP/IP connections on all server host IPv6 and IPv4 interfaces if the server host supports IPv6, or accepts TCP/IP connections on all IPv4 addresses otherwise. Use this address to permit both IPv4 and IPv6 connections on all server interfaces. This value is permitted (and is the default) as of MySQL 5.6.6.
- If the address is `0.0.0.0`, the server accepts TCP/IP connections on all server host IPv4 interfaces. This is the default before MySQL 5.6.6.
- If the address is `::`, the server accepts TCP/IP connections on all server host IPv4 and IPv6 interfaces.
- If the address is an IPv4-mapped address, the server accepts TCP/IP connections for that address, in either IPv4 or IPv6 format. For example, if the server is bound to `::ffff:127.0.0.1`, clients can connect using `--host=127.0.0.1` or `--host>::ffff:127.0.0.1`.
- If the address is a “regular” IPv4 or IPv6 address (such as `127.0.0.1` or `::1`), the server accepts TCP/IP connections only for that IPv4 or IPv6 address.

If you intend to bind the server to a specific address, be sure that the `mysql.user` grant table contains an account with administrative privileges that you can use to connect to that address. Otherwise, you will not be able to shut down the server. For example, if you bind the server to `*`, you can connect to it using all existing accounts. But if you bind the server to `::1`, it accepts connections only on that address. In that case, first make sure that the `'root'@'::1'` account is present in the `mysql.user` table so you can still connect to the server to shut it down.

- `--binlog-format={ROW | STATEMENT | MIXED}` [415]

<b>Command-Line Format</b>	<code>--binlog-format=format</code>	
<b>Option-File Format</b>	<code>binlog-format=format</code>	
<b>Option Sets Variable</b>	Yes, <code>binlog_format</code> [2010]	
<b>Variable Name</b>	<code>binlog_format</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	STATEMENT
	<b>Valid Values</b>	ROW STATEMENT MIXED

Specify whether to use row-based, statement-based, or mixed replication. Statement-based is the default in MySQL 5.6. See [Section 16.1.2, “Replication Formats”](#).

Under some conditions, changing this variable at runtime is not possible, or causes replication to fail. See [Section 5.2.4.2, “Setting The Binary Log Format”](#), for more information.

Setting the binary logging format without enabling binary logging sets the `binlog_format` [2010] global system variable and logs a warning.

- `--bootstrap` [416]

<b>Command-Line Format</b>	<code>--bootstrap</code>
<b>Option-File Format</b>	<code>bootstrap</code>

This option is used by the `mysql_install_db` script to create the MySQL privilege tables without having to start a full MySQL server.

In MySQL 5.6.6 and later, replication and global transaction identifiers are automatically disabled whenever this option is used (Bug #1332602). See [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

- `--character-sets-dir=path` [416]

<b>Command-Line Format</b>	<code>--character-sets-dir=path</code>	
<b>Option-File Format</b>	<code>character-sets-dir=path</code>	
<b>Option Sets Variable</b>	Yes, <code>character_sets_dir</code> [469]	
<b>Variable Name</b>	<code>character-sets-dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The directory where character sets are installed. See [Section 10.5, “Character Set Configuration”](#).

- `--character-set-client-handshake` [416]

<b>Command-Line Format</b>	<code>--character-set-client-handshake</code>	
<b>Option-File Format</b>	<code>character-set-client-handshake</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>TRUE</code>

Do not ignore character set information sent by the client. To ignore client information and use the default server character set, use `--skip-character-set-client-handshake` [416]; this makes MySQL behave like MySQL 4.0.

- `--character-set-filesystem=charset_name` [416]

<b>Command-Line Format</b>	<code>--character-set-filesystem=name</code>	
<b>Option-File Format</b>	<code>character-set-filesystem</code>	
<b>Option Sets Variable</b>	Yes, <code>character_set_filesystem</code> [468]	
<b>Variable Name</b>	<code>character_set_filesystem</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	416

	<b>Permitted Values</b>	
	<b>Type</b>	string

The file system character set. This option sets the `character_set_filesystem` [468] system variable.

- `--character-set-server=charset_name` [417], `-C charset_name`

<b>Command-Line Format</b>	<code>--character-set-server</code>	
<b>Option-File Format</b>	<code>character-set-server</code>	
<b>Option Sets Variable</b>	Yes, <code>character_set_server</code> [468]	
<b>Variable Name</b>	<code>character_set_server</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Use `charset_name` as the default server character set. See [Section 10.5, “Character Set Configuration”](#). If you use this option to specify a nondefault character set, you should also use `--collation-server` [417] to specify the collation.

- `--chroot=path` [417], `-r path`

<b>Command-Line Format</b>	<code>--chroot=name</code>	
	<code>-r name</code>	
<b>Option-File Format</b>	<code>chroot</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

Put the `mysqld` server in a closed environment during startup by using the `chroot()` system call. This is a recommended security measure. Note that use of this option somewhat limits `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE`.

- `--collation-server=collation_name` [417]

<b>Command-Line Format</b>	<code>--collation-server</code>	
<b>Option-File Format</b>	<code>collation-server</code>	
<b>Option Sets Variable</b>	Yes, <code>collation_server</code> [470]	
<b>Variable Name</b>	<code>collation_server</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Use `collation_name` as the default server collation. See [Section 10.5, “Character Set Configuration”](#).

- `--console` [417]

<b>Command-Line Format</b>	<code>--console</code>	
<b>Option-File Format</b>	<code>console</code>	

<b>Platform Specific</b>	windows
--------------------------	---------

(Windows only.) Write error log messages to `stderr` and `stdout`. `mysqld` does not close the console window if this option is used.

Passing in both `--log-error` [427] and `--console` [417] will write to the log file, but not to the console.

- `--core-file` [418]

<b>Command-Line Format</b>	<code>--core-file</code>	
<b>Option-File Format</b>	<code>core-file</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Write a core file if `mysqld` dies. The name and location of the core file is system dependent. On Linux, a core file named `core.pid` is written to the current working directory of the process, which for `mysqld` is the data directory. `pid` represents the process ID of the server process. On Mac OS X, a core file named `core.pid` is written to the `/cores` directory. On Solaris, use the `coreadm` command to specify where to write the core file and how to name it.

For some systems, to get a core file you must also specify the `--core-file-size` [227] option to `mysqld_safe`. See [Section 4.3.2, “mysqld\\_safe — MySQL Server Startup Script”](#). On some systems, such as Solaris, you do not get a core file if you are also using the `--user` [446] option. There might be additional restrictions or limitations. For example, it might be necessary to execute `ulimit -c unlimited` before starting the server. Consult your system documentation.

- `--datadir=path` [418], `-h path`

<b>Command-Line Format</b>	<code>--datadir=path</code>	
	<code>-h</code>	
<b>Option-File Format</b>	<code>datadir</code>	
<b>Option Sets Variable</b>	Yes, <code>datadir</code> [472]	
<b>Variable Name</b>	<code>datadir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path to the data directory.

- `--debug[=debug_options]` [418], `-# [debug_options]`

<b>Command-Line Format</b>	<code>--debug[=debug_options]</code>	
<b>Option-File Format</b>	<code>debug</code>	
<b>Variable Name</b>	<code>debug</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

	<b>Default</b>	'd:t:o,/tmp/mysqld.trace'
--	----------------	---------------------------

If MySQL is configured with `-DWITH_DEBUG=1` [136], you can use this option to get a trace file of what `mysqld` is doing. A typical *debug\_options* string is `'d:t:o,file_name'`. The default is `'d:t:i:o,mysqld.trace'`. See [MySQL Internals: Porting to Other Systems](#).

Using `-DWITH_DEBUG=1` [136] to configure MySQL with debugging support enables you to use the `--debug="d,parser_debug"` [418] option when you start the server. This causes the Bison parser that is used to process SQL statements to dump a parser trace to the server's standard error output. Typically, this output is written to the error log.

This option may be given multiple times. Values that begin with `+` or `-` are added to or subtracted from the previous value. For example, `--debug=T` [418] `--debug=+P` [418] sets the value to `P:T`.

- `--debug-sync-timeout[=N]` [419]

<b>Command-Line Format</b>	<code>--debug-sync-timeout[=#]</code>
<b>Option-File Format</b>	<code>debug-sync-timeout</code>
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>

Controls whether the Debug Sync facility for testing and debugging is enabled. Use of Debug Sync requires that MySQL be configured with the `-DENABLE_DEBUG_SYNC=1` [134] option (see [Section 2.9.4, "MySQL Source-Configuration Options"](#)). If Debug Sync is not compiled in, this option is not available. The option value is a timeout in seconds. The default value is 0, which disables Debug Sync. To enable it, specify a value greater than 0; this value also becomes the default timeout for individual synchronization points. If the option is given without a value, the timeout is set to 300 seconds.

For a description of the Debug Sync facility and how to use synchronization points, see [MySQL Internals: Test Synchronization](#).

- `--default-authentication-plugin=plugin_name` [419]

<b>Version Introduced</b>	5.6.6
<b>Command-Line Format</b>	<code>--default-authentication-plugin=plugin_name</code>
<b>Option-File Format</b>	<code>default-authentication-plugin=plugin_name</code>
	<b>Permitted Values</b>
	<b>Type</b> <code>enumeration</code>
	<b>Default</b> <code>mysql_native_password</code>
	<b>Valid Values</b> <code>mysql_native_password</code> <code>sha256_password</code>

The default authentication plugin. Acceptable values are `mysql_native_password` (use MySQL native passwords) and `sha256_password` (use SHA-256 passwords). For more information about these plugins, see [Section 6.3.7.1, "The Native Authentication Plugins"](#), and [Section 6.3.7.2, "The SHA-256 Authentication Plugin"](#). This option was added in MySQL 5.6.6.

The value of this option affects these aspects of server operation:

- It determines which plugin the server assigns to new accounts for which no plugin is explicitly specified.

- It sets the `old_passwords` [517] system variable at startup to the value that is consistent with the password hashing format required by the default plugin. This in turn affects the password hashing method used by the `PASSWORD()` [1263] function.
- For an account created with either of the following statements, the server associates the account with the default authentication plugin and assigns the account the given password, hashed according to the value of `old_passwords` [517].

```
CREATE USER ... IDENTIFIED BY 'cleartext password';
GRANT ... IDENTIFIED BY 'cleartext password';
```

- For an account created with either of the following statements, the statement fails if the password hash is not encrypted using the hash format required by the default authentication plugin. Otherwise, the server associates the account with the default authentication plugin and assigns the account the given password hash.

```
CREATE USER ... IDENTIFIED BY 'encrypted password';
GRANT ... IDENTIFIED BY 'encrypted password';
```



**Note**

If you use this option to change the default authentication method to a value other than `mysql_native_password`, clients older than MySQL 5.5.6 will no longer be able to connect because they will not understand the change to the authentication protocol.

- `--default-storage-engine=type` [420]

<b>Command-Line Format</b>	<code>--default-storage-engine=name</code>	
<b>Option-File Format</b>	<code>default-storage-engine</code>	
<b>Option Sets Variable</b>	Yes, <code>default_storage_engine</code> [474]	
<b>Variable Name</b>	<code>default_storage_engine</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>InnoDB</code>

Set the default storage engine for tables. See [Chapter 14, Storage Engines](#). As of MySQL 5.6.3, this option sets the storage engine for permanent tables only. To set the storage engine for `TEMPORARY` tables, set the `default_tmp_storage_engine` [474] system variable.

If you disable the default storage engine at server startup, you must set the default engine for both permanent and `TEMPORARY` tables to a different engine or the server will not start.

- `--default-time-zone=timezone` [420]

<b>Command-Line Format</b>	<code>--default-time-zone=name</code>	
<b>Option-File Format</b>	<code>default-time-zone</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

Set the default server time zone. This option sets the global `time_zone` [557] system variable. If this option is not given, the default time zone is the same as the system time zone (given by the value of the `system_time_zone` [550] system variable).

- `--delay-key-write[={OFF|ON|ALL}]` [421]

<b>Command-Line Format</b>	<code>--delay-key-write[=name]</code>
<b>Option-File Format</b>	<code>delay-key-write</code>
<b>Option Sets Variable</b>	Yes, <code>delay_key_write</code> [475]
<b>Variable Name</b>	<code>delay-key-write</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> enumeration
	<b>Default</b> ON
	<b>Valid Values</b> ON OFF ALL

Specify how to use delayed key writes. Delayed key writing causes key buffers not to be flushed between writes for `MyISAM` tables. `OFF` disables delayed key writes. `ON` enables delayed key writes for those tables that were created with the `DELAY_KEY_WRITE` option. `ALL` delays key writes for all `MyISAM` tables. See [Section 8.11.2, “Tuning Server Parameters”](#), and [Section 14.3.1, “MyISAM Startup Options”](#).



**Note**

If you set this variable to `ALL`, you should not use `MyISAM` tables from within another program (such as another MySQL server or `myisamchk`) when the tables are in use. Doing so leads to index corruption.

- `--des-key-file=file_name` [421]

<b>Command-Line Format</b>	<code>--des-key-file=file_name</code>
<b>Option-File Format</b>	<code>des-key-file=file_name</code>

Read the default DES keys from this file. These keys are used by the `DES_ENCRYPT()` [1262] and `DES_DECRYPT()` [1261] functions.

- `--enable-named-pipe` [421]

<b>Command-Line Format</b>	<code>--enable-named-pipe</code>
<b>Option-File Format</b>	<code>enable-named-pipe</code>
<b>Option Sets Variable</b>	Yes, <code>named_pipe</code> [514]
<b>Platform Specific</b>	windows

Enable support for named pipes. This option applies only on Windows.

- `--engine-condition-pushdown={ON|OFF}`

<b>Version Removed</b>	5.6.1
<b>Command-Line Format</b>	<code>--engine-condition-pushdown</code>
<b>Option-File Format</b>	<code>engine-condition-pushdown</code>
<b>Option Sets Variable</b>	Yes, <code>engine_condition_pushdown</code> [478]
<b>Variable Name</b>	<code>engine_condition_pushdown</code>



<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
<b>Deprecated</b>	5.5.3, by <a href="#">optimizer_switch</a>
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>ON</code>

Formerly, this option set the [engine\\_condition\\_pushdown \[478\]](#) system variable; the option and the variable were both removed in MySQL 5.6.1. For more information, see [Section 8.13.3, “Engine Condition Pushdown Optimization”](#).

- [--event-scheduler\[=value\] \[422\]](#)

<b>Command-Line Format</b>	<code>--event-scheduler[=value]</code>
<b>Option-File Format</b>	<code>event-scheduler</code>
<b>Option Sets Variable</b>	Yes, <a href="#">event_scheduler [480]</a>
<b>Variable Name</b>	<code>event_scheduler</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>enumeration</code>
	<b>Default</b> <code>OFF</code>
	<b>Valid Values</b> <code>ON</code> <code>OFF</code> <code>DISABLED</code>

Enable or disable, and start or stop, the event scheduler.

For detailed information, see [The --event-scheduler Option \[2159\]](#).

- [--exit-info\[=flags\] \[422\]](#), `-T [flags]`

<b>Command-Line Format</b>	<code>--exit-info[=flags]</code>
	<code>-T [flags]</code>
<b>Option-File Format</b>	<code>exit-info</code>
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>

This is a bit mask of different flags that you can use for debugging the `mysqld` server. Do not use this option unless you know *exactly* what it does!

- [--external-locking \[422\]](#)

<b>Command-Line Format</b>	<code>--external-locking</code>
<b>Option-File Format</b>	<code>external-locking</code>
<b>Option Sets Variable</b>	Yes, <a href="#">skip_external_locking [540]</a>
<b>Disabled by</b>	<code>skip-external-locking</code>
	<b>Permitted Values</b>

	<b>Type</b>	boolean
	<b>Default</b>	FALSE

Enable external locking (system locking), which is disabled by default as of MySQL 4.0. Note that if you use this option on a system on which `lockd` does not fully work (such as Linux), it is easy for `mysqld` to deadlock.

External locking affects only `MyISAM` table access. For more information, including conditions under which it can and cannot be used, see [Section 8.10.5, “External Locking”](#).

- `--flush` [\[423\]](#)

<b>Command-Line Format</b>	<code>--flush</code>	
<b>Option-File Format</b>	<code>flush</code>	
<b>Variable Name</b>	<code>flush</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Flush (synchronize) all changes to disk after each SQL statement. Normally, MySQL does a write of all changes to disk only after each SQL statement and lets the operating system handle the synchronizing to disk. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#).

- `--gdb` [\[423\]](#)

<b>Command-Line Format</b>	<code>--gdb</code>	
<b>Option-File Format</b>	<code>gdb</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	FALSE

Install an interrupt handler for `SIGINT` (needed to stop `mysqld` with `^C` to set breakpoints) and disable stack tracing and core file handling. See [MySQL Internals: Porting to Other Systems](#).

- `--general-log[={0|1}]` [\[423\]](#)

<b>Command-Line Format</b>	<code>--general-log</code>	
<b>Option-File Format</b>	<code>general-log</code>	
<b>Option Sets Variable</b>	Yes, <code>general_log</code> <a href="#">[485]</a>	
<b>Variable Name</b>	<code>general_log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Specify the initial general query log state. With no argument or an argument of 1, the `--general-log` [\[423\]](#) option enables the log. If omitted or given with an argument of 0, the option disables the log.

- `--ignore-db-dir=dir_name` [424]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--ignore-db-dir</code>	
<b>Option-File Format</b>	<code>ignore-db-dir</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

This option tells the server to ignore the given directory name for purposes of the `SHOW DATABASES` statement or `INFORMATION_SCHEMA` tables. For example, if a MySQL configuration locates the data directory at the root of a file system on Unix, the system might create a `lost+found` directory there that the server should ignore. Starting the server with `--ignore-db-dir=lost+found` [424] causes that name not to be listed as a database.

To specify more than one name, use this option multiple times, once for each name. Specifying the option with an empty value (that is, as `--ignore-db-dir=` [424]) resets the directory list to the empty list.

Instances of this option given at server startup are used to set the `ignore_db_dirs` [488] system variable.

This option was added in MySQL 5.6.3.

- `--init-file=file_name` [424]

<b>Command-Line Format</b>	<code>--init-file=file_name</code>	
<b>Option-File Format</b>	<code>init-file=file_name</code>	
<b>Option Sets Variable</b>	Yes, <code>init_file</code> [489]	
<b>Variable Name</b>	<code>init_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

Read SQL statements from this file at startup. Each statement must be on a single line and should not include comments.

- `--innodb-xxx`

Set an option for the InnoDB storage engine. The InnoDB options are listed in Section 14.2.6, “InnoDB Startup Options and System Variables”.

- `--install [service_name]` [424]

<b>Command-Line Format</b>	<code>--install [service_name]</code>
----------------------------	---------------------------------------

(Windows only) Install the server as a Windows service that starts automatically during Windows startup. The default service name is `MySQL` if no `service_name` value is given. For more information, see Section 2.3.4.8, “Starting MySQL as a Windows Service”.

- `--install-manual [service_name]` [424]

<b>Command-Line Format</b>	<code>--install-manual [service_name]</code>
----------------------------	--

(Windows only) Install the server as a Windows service that must be started manually. It does not start automatically during Windows startup. The default service name is `MySQL` if no `service_name` value is given. For more information, see [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

- `--language=lang_name`, `-L lang_name` [425]

<b>Version Deprecated</b>	5.6.1	
<b>Command-Line Format</b>	<code>--language=name</code>	
	<code>-L</code>	
<b>Option-File Format</b>	<code>language</code>	
<b>Option Sets Variable</b>	Yes, <code>language</code>	
<b>Variable Name</b>	<code>language</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Deprecated</b>	5.6.1	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>
	<b>Default</b>	<code>/usr/local/mysql/share/mysql/english/</code>

The language to use for error messages. `lang_name` can be given as the language name or as the full path name to the directory where the language files are installed. See [Section 10.2, “Setting the Error Message Language”](#).

In MySQL 5.6, `--lc-messages-dir` [426] and `--lc-messages` [426] should be used rather than `--language` [425], which is deprecated as of MySQL 5.6.1 and handled as an alias for `--lc-messages-dir` [426]. The `--language` [425] option will be removed in a future MySQL release.

- `--large-pages` [425]

<b>Command-Line Format</b>	<code>--large-pages</code>	
<b>Option-File Format</b>	<code>large-pages</code>	
<b>Option Sets Variable</b>	Yes, <code>large_pages</code> [494]	
<b>Variable Name</b>	<code>large_pages</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Platform Specific</b>	linux	
	<b>Permitted Values</b>	
	<b>Type</b> (linux)	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Some hardware/operating system architectures support memory pages greater than the default (usually 4KB). The actual implementation of this support depends on the underlying hardware and operating system. Applications that perform a lot of memory accesses may obtain performance improvements by using large pages due to reduced Translation Lookaside Buffer (TLB) misses.

MySQL 5.6 supports the Linux implementation of large page support (which is called HugeTLB in Linux). See [Section 8.11.4.2, “Enabling Large Page Support”](#). For Solaris support of large pages, see the description of the `--super-large-pages` [440] option.

`--large-pages` [425] is disabled by default.

- `--lc-messages=locale_name` [426]

<b>Command-Line Format</b>	<code>--lc-messages=name</code>	
<b>Option-File Format</b>	<code>lc-messages</code>	
<b>Option Sets Variable</b>	Yes, <code>lc_messages</code> [495]	
<b>Variable Name</b>	<code>lc-messages</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The locale to use for error messages. The server converts the argument to a language name and combines it with the value of the `--lc-messages-dir` to produce the location for the error message file. See [Section 10.2, “Setting the Error Message Language”](#).

- `--lc-messages-dir=path` [426]

<b>Command-Line Format</b>	<code>--lc-messages-dir=path</code>	
<b>Option-File Format</b>	<code>lc-messages-dir</code>	
<b>Option Sets Variable</b>	Yes, <code>lc_messages_dir</code> [495]	
<b>Variable Name</b>	<code>lc-messages-dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The directory where error messages are located. The value is used together with the value of `--lc-messages` to produce the location for the error message file. See [Section 10.2, “Setting the Error Message Language”](#).

- `--log[=file_name]` [426], `-l [file_name]`

<b>Version Removed</b>	5.6.1	
<b>Command-Line Format</b>	<code>--log[=name]</code>	
	<code>-l</code>	
<b>Option-File Format</b>	<code>log</code>	
<b>Option Sets Variable</b>	Yes, <code>log</code> [497]	
<b>Variable Name</b>	<code>log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.1.29, by <code>general-log</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>OFF</code>

The `--log` [426] option was removed in MySQL 5.6.1 (along with the `log` [497] system variable). Instead, use the `--general_log` [485] option to enable the general query log and the `--general_log_file=file_name` [486] option to set the general query log file name.

- `--log-error[=file_name]` [427]

<b>Command-Line Format</b>	<code>--log-error[=name]</code>	
<b>Option-File Format</b>	<code>log-error</code>	
<b>Option Sets Variable</b>	Yes, <code>log_error</code> [497]	
<b>Variable Name</b>	<code>log_error</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

Log errors and startup messages to this file. See Section 5.2.2, “The Error Log”. If you omit the file name, MySQL uses `host_name.err`. If the file name has no extension, the server adds an extension of `.err`.

- `--log-isam[=file_name]` [427]

<b>Command-Line Format</b>	<code>--log-isam[=name]</code>	
<b>Option-File Format</b>	<code>log-isam</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

Log all MyISAM changes to this file (used only when debugging MyISAM).

- `--log-output=value,...` [427]

<b>Command-Line Format</b>	<code>--log-output=name</code>	
<b>Option-File Format</b>	<code>log-output</code>	
<b>Option Sets Variable</b>	Yes, <code>log_output</code> [498]	
<b>Variable Name</b>	<code>log_output</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>set</code>
	<b>Default</b>	<code>FILE</code>
	<b>Valid Values</b>	<code>TABLE</code> <code>FILE</code> <code>NONE</code>

This option determines the destination for general query log and slow query log output. The option value can be given as one or more of the words `TABLE`, `FILE`, or `NONE`. `TABLE` selects logging to the `general_log` [485] and `slow_log` tables in the `mysql` database as a destination. `FILE` selects logging to log files as a destination. `NONE` disables logging. If `NONE` is present in the option value, it takes precedence over any other words that are present. `TABLE` and `FILE` can both be given to select to both log output destinations.

This option selects log output destinations, but does not enable log output. To do that, use the `--general_log` [485] and `--slow_query_log` [541] options. For FILE logging, the `--general_log_file` [486] and `--slow_query_log_file` options determine the log file location. For more information, see Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”.

- `--log-queries-not-using-indexes` [428]

<b>Command-Line Format</b>	<code>--log-queries-not-using-indexes</code>	
<b>Option-File Format</b>	<code>log-queries-not-using-indexes</code>	
<b>Option Sets Variable</b>	Yes, <code>log_queries_not_using_indexes</code> [498]	
<b>Variable Name</b>	<code>log_queries_not_using_indexes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

If you are using this option with the slow query log enabled, queries that are expected to retrieve all rows are logged. See Section 5.2.5, “The Slow Query Log”. This option does not necessarily mean that no index is used. For example, a query that uses a full index scan uses an index but would be logged because the index would not limit the number of rows.

- `--log-raw` [428]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--log-raw[=value]</code>	
<b>Option-File Format</b>	<code>log-raw[=value]</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

As of MySQL 5.6.3, passwords in certain statements written to the general query log, slow query log, and binary log are rewritten by the server not to occur literally in plain text. Password rewriting can be suppressed for the general query log by starting the server with the `--log-raw` [428] option. This option may be useful for diagnostic purposes, to see the exact text of statements as received by the server, but for security reasons is not recommended for production use.

Before MySQL 5.6.3, passwords in statements are not rewritten and the general query log should be protected. See Section 6.1.2.3, “Passwords and Logging”.

This option was added in MySQL 5.6.3.

- `--log-short-format` [428]

<b>Command-Line Format</b>	<code>--log-short-format</code>	
<b>Option-File Format</b>	<code>log-short-format</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Log less information to the binary log and slow query log, if they have been activated.



- `--log-slow-admin-statements` [429]

<b>Command-Line Format</b>	<code>--log-slow-admin-statements</code>	
<b>Option-File Format</b>	<code>log-slow-admin-statements</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Include slow administrative statements in the statements written to the slow query log. Administrative statements include `ALTER TABLE`, `ANALYZE TABLE`, `CHECK TABLE`, `CREATE INDEX`, `DROP INDEX`, `OPTIMIZE TABLE`, and `REPAIR TABLE`.

- `--log-slow-queries[=file_name]` [429]

<b>Version Removed</b>	5.6.1	
<b>Command-Line Format</b>	<code>--log-slow-queries[=name]</code>	
<b>Option-File Format</b>	<code>log-slow-queries</code>	
<b>Option Sets Variable</b>	Yes, <code>log_slow_queries</code> [499]	
<b>Variable Name</b>	<code>log_slow_queries</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.1.29, by <code>slow-query-log</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>

The `--log-slow-queries` [429] option was removed in MySQL 5.6.1 (along with the `log_slow_queries` [499] system variable). Instead, use the `--slow_query_log` [541] option to enable the slow query log and the `--slow_query_log_file=file_name` [541] option to set the slow query log file name.

- `--log-tc=file_name` [429]

<b>Command-Line Format</b>	<code>--log-tc=name</code>	
<b>Option-File Format</b>	<code>log-tc</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>tc.log</code>

The name of the memory-mapped transaction coordinator log file (for XA transactions that affect multiple storage engines when the binary log is disabled). The default name is `tc.log`. The file is created under the data directory if not given as a full path name. Currently, this option is unused.

- `--log-tc-size=size` [429]

<b>Command-Line Format</b>	<code>--log-tc-size=#</code>	
<b>Option-File Format</b>	<code>log-tc-size</code>	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>32</code>
	<b>Type</b>	<code>numeric</code> <sup>429</sup>

	<b>Default</b>	24576
	<b>Max Value</b>	4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	24576
	<b>Max Value</b>	18446744073709547520

The size in bytes of the memory-mapped transaction coordinator log. The default size is 24KB.

- `--log-warnings[=level] [430], -W [level]`

<b>Command-Line Format</b>	<code>--log-warnings[=#]</code>	
	<code>-W [#]</code>	
<b>Option-File Format</b>	<code>log-warnings</code>	
<b>Option Sets Variable</b>	Yes, <code>log_warnings</code> [499]	
<b>Variable Name</b>	<code>log_warnings</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Disabled by</b>	<code>skip-log-warnings</code>	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Range</b>	0 .. 18446744073709547520

Print out warnings such as `Aborted connection...` to the error log. Enabling this option is recommended, for example, if you use replication (you get more information about what is happening, such as messages about network failures and reconnections). This option is enabled (1) by default, and the default `level` value if omitted is 1. To disable this option, use `--log-warnings=0` [430]. If the value is greater than 1, aborted connections are written to the error log, and access-denied errors for new connection attempts are written. See [Section C.5.2.11, “Communication Errors and Aborted Connections”](#).

If a slave server was started with `--log-warnings` [430] enabled, the slave prints messages to the error log to provide information about its status, such as the binary log and relay log coordinates where it starts its job, when it is switching to another relay log, when it reconnects after a disconnect, and so forth. The server logs messages about statements that are unsafe for statement-based logging only if `--log-warnings` [430] is enabled.

- `--low-priority-updates` [430]

<b>Command-Line Format</b>	<code>--low-priority-updates</code>
<b>Option-File Format</b>	<code>low-priority-updates</code>
<b>Option Sets Variable</b>	Yes, <code>low_priority_updates</code> [500]
<b>Variable Name</b>	<code>low_priority_updates</code>

<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	FALSE

Give table-modifying operations (`INSERT`, `REPLACE`, `DELETE`, `UPDATE`) lower priority than selects. This can also be done using `{INSERT | REPLACE | DELETE | UPDATE} LOW_PRIORITY ...` to lower the priority of only one query, or by `SET LOW_PRIORITY_UPDATES=1` to change the priority in one thread. This affects only storage engines that use only table-level locking (`MyISAM`, `MEMORY`, `MERGE`). See [Section 8.10.2, “Table Locking Issues”](#).

- `--min-examined-row-limit=number` [431]

<b>Command-Line Format</b>	<code>--min-examined-row-limit=#</code>
<b>Option-File Format</b>	<code>min-examined-row-limit</code>
<b>Variable Name</b>	<code>min_examined_row_limit</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> numeric
	<b>Default</b> 0
	<b>Range</b> 0 .. 4294967295
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> numeric
	<b>Default</b> 0
	<b>Range</b> 0 .. 18446744073709547520

When this option is set, queries which examine fewer than `number` rows are not written to the slow query log. The default is 0.

- `--memlock` [431]

<b>Command-Line Format</b>	<code>--memlock</code>
<b>Option-File Format</b>	<code>memlock</code>
<b>Variable Name</b>	<code>locked_in_memory</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
	<b>Type</b> boolean
	<b>Default</b> FALSE

Lock the `mysqld` process in memory. This option might help if you have a problem where the operating system is causing `mysqld` to swap to disk.

`--memlock` [431] works on systems that support the `mlockall()` system call; this includes Solaris, most Linux distributions that use a 2.4 or newer kernel, and perhaps other Unix systems. On Linux systems, you can tell whether or not `mlockall()` (and thus this option) is supported by checking to see whether or not it is defined in the system `mman.h` file, like this:

```
shell> grep mlockall /usr/include/sys/mman.h
```

If `mlockall()` is supported, you should see in the output of the previous command something like the following:

```
extern int mlockall (int __flags) __THROW;
```



**Important**

Use of this option may require you to run the server as `root`, which, for reasons of security, is normally not a good idea. See [Section 6.1.5, “How to Run MySQL as a Normal User”](#).

On Linux and perhaps other systems, you can avoid the need to run the server as `root` by changing the `limits.conf` file. See the notes regarding the `memlock` limit in [Section 8.11.4.2, “Enabling Large Page Support”](#).

You must not try to use this option on a system that does not support the `mlockall()` system call; if you do so, `mysqld` will very likely crash as soon as you try to start it.

- `--myisam-block-size=N` [432]

<b>Command-Line Format</b>	<code>--myisam-block-size=#</code>
<b>Option-File Format</b>	<code>myisam-block-size</code>
	<b>Permitted Values</b>
	<b>Type</b> numeric
	<b>Default</b> 1024
	<b>Range</b> 1024 .. 16384

The block size to be used for `MyISAM` index pages.

- `--myisam-recover-options[=option[,option]...]` [432]

<b>Command-Line Format</b>	<code>--myisam-recover-options[=name]</code>
<b>Option-File Format</b>	<code>myisam-recover-options</code>
<b>Option Sets Variable</b>	Yes, <code>myisam_recover_options</code> [512]
	<b>Permitted Values</b>
	<b>Type</b> enumeration
	<b>Default</b> OFF
	<b>Valid Values</b> OFF DEFAULT BACKUP FORCE QUICK

Set the `MyISAM` storage engine recovery mode. The option value is any combination of the values of `OFF`, `DEFAULT`, `BACKUP`, `FORCE`, or `QUICK`. If you specify multiple values, separate them by commas. Specifying the option with no argument is the same as specifying `DEFAULT`, and specifying with an explicit value of `" "` disables recovery (same as a value of `OFF`). If recovery is enabled, each time `mysqld` opens a `MyISAM` table, it checks whether the table is marked as crashed or was not closed properly. (The last option works only if you are running with external locking disabled.) If this is the case, `mysqld` runs a check on the table. If the table was corrupted, `mysqld` attempts to repair it.

The following options affect how the repair works.

Option	Description
<code>OFF</code>	No recovery.
<code>DEFAULT</code>	Recovery without backup, forcing, or quick checking.
<code>BACKUP</code>	If the data file was changed during recovery, save a backup of the <code>tbl_name.MYD</code> file as <code>tbl_name-datetime.BAK</code> .
<code>FORCE</code>	Run recovery even if we would lose more than one row from the <code>.MYD</code> file.
<code>QUICK</code>	Do not check the rows in the table if there are not any delete blocks.

Before the server automatically repairs a table, it writes a note about the repair to the error log. If you want to be able to recover from most problems without user intervention, you should use the options `BACKUP`, `FORCE`. This forces a repair of a table even if some rows would be deleted, but it keeps the old data file as a backup so that you can later examine what happened.

See [Section 14.3.1, “MyISAM Startup Options”](#).

- `--old-alter-table` [\[433\]](#)

<b>Command-Line Format</b>	<code>--old-alter-table</code>	
<b>Option-File Format</b>	<code>old-alter-table</code>	
<b>Option Sets Variable</b>	Yes, <code>old_alter_table</code> <a href="#">[517]</a>	
<b>Variable Name</b>	<code>old_alter_table</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

When this option is given, the server does not use the optimized method of processing an `ALTER TABLE` operation. It reverts to using a temporary table, copying over the data, and then renaming the temporary table to the original, as used by MySQL 5.0 and earlier. For more information on the operation of `ALTER TABLE`, see [Section 13.1.6, “ALTER TABLE Syntax”](#).

- `--old-style-user-limits` [\[433\]](#)

<b>Command-Line Format</b>	<code>--old-style-user-limits</code>	
<b>Option-File Format</b>	<code>old-style-user-limits</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Enable old-style user limits. (Before MySQL 5.0.3, account resource limits were counted separately for each host from which a user connected rather than per account row in the `user` table.) See [Section 6.3.4, “Setting Account Resource Limits”](#).

- `--one-thread` [\[434\]](#)

<b>Version Removed</b>	5.6.1
<b>Command-Line Format</b>	<code>--one-thread</code>
<b>Option-File Format</b>	<code>one-thread</code>

This option was removed in MySQL 5.6.1. Use `--thread_handling=no-threads` [\[553\]](#) instead.

- `--open-files-limit=count` [\[434\]](#)

<b>Command-Line Format</b>	<code>--open-files-limit=#</code>
<b>Option-File Format</b>	<code>open-files-limit</code>
<b>Option Sets Variable</b>	Yes, <code>open_files_limit</code> <a href="#">[518]</a>
<b>Variable Name</b>	<code>open_files_limit</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&lt;= 5.6.7)</b>
<b>Type</b>	<code>numeric</code>
<b>Default</b>	0
<b>Range</b>	0 .. 65535
	<b>Permitted Values (&gt;= 5.6.8)</b>
<b>Type</b>	<code>numeric</code>
<b>Default</b>	-1 (autosized)
<b>Range</b>	0 .. 65535

Changes the number of file descriptors available to `mysqld`. You should try increasing the value of this option if `mysqld` gives you the error `Too many open files`. `mysqld` uses the option value to reserve descriptors with `setrlimit()`. If the requested number of file descriptors cannot be allocated, `mysqld` writes a warning to the error log.

`mysqld` may attempt to allocate more than the requested number of descriptors (if they are available), using the values of `max_connections` [\[503\]](#) and `table_open_cache` [\[551\]](#) to estimate whether more descriptors will be needed.

On Unix, the value cannot be set less than `ulimit -n`.

- `--partition[=value]` [\[434\]](#)

<b>Command-Line Format</b>	<code>--partition</code>
<b>Option-File Format</b>	<code>partition</code>
<b>Option Sets Variable</b>	Yes, <code>have_partitioning</code> <a href="#">[487]</a>
<b>Variable Name</b>	<code>partition</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Disabled by</b>	<code>skip-partition</code>

	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	ON	

Enables or disables user-defined partitioning support in the MySQL Server.

- `--performance-schema-xxx`

Configure a Performance Schema option. For details, see [Section 20.11, “Performance Schema Command Options”](#).

- `--pid-file=path` [435]

<b>Command-Line Format</b>	<code>--pid-file=file_name</code>	
<b>Option-File Format</b>	<code>pid-file=file_name</code>	
<b>Option Sets Variable</b>	Yes, <code>pid_file</code> [524]	
<b>Variable Name</b>	<code>pid_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path name of the process ID file. The server creates the file in the data directory unless an absolute path name is given to specify a different directory. This file is used by other programs such as `mysqld_safe` to determine the server's process ID.

- `--plugin-xxx` [435]

Specifies an option that pertains to a server plugin. For example, many storage engines can be built as plugins, and for such engines, options for them can be specified with a `--plugin` [435] prefix. Thus, the `--innodb_file_per_table` [1718] option for InnoDB can be specified as `--plugin-innodb_file_per_table` [1718].

For boolean options that can be enabled or disabled, the `--skip` prefix and other alternative formats are supported as well (see [Section 4.2.3.2, “Program Option Modifiers”](#)). For example, `--skip-plugin-innodb_file_per_table` [1718] disables `innodb_file_per_table` [1718].

The rationale for the `--plugin` prefix is that it enables plugin options to be specified unambiguously if there is a name conflict with a built-in server option. For example, were a plugin writer to name a plugin “sql” and implement a “mode” option, the option name might be `--sql-mode` [442], which would conflict with the built-in option of the same name. In such cases, references to the conflicting name are resolved in favor of the built-in option. To avoid the ambiguity, users can specify the plugin option as `--plugin-sql-mode`. Use of the `--plugin` prefix for plugin options is recommended to avoid any question of ambiguity.

- `--plugin-load=plugin_list` [435]

<b>Command-Line Format</b>	<code>--plugin-load=plugin_list</code>	
<b>Option-File Format</b>	<code>plugin-load=plugin_list</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	string

This option tells the server to load the named plugins at startup. The option value is a semicolon-separated list of `name=plugin_library` pairs. Each `name` is the name of the plugin, and `plugin_library` is the name of the shared library that contains the plugin code. Each library file



must be located in the directory named by the `plugin_dir` [525] system variable. For example, if plugins named `myplug1` and `myplug2` have library files `myplug1.so` and `myplug2.so`, use this option to load them at startup:

```
shell> mysqld --plugin-load="myplug1=myplug1.so;myplug2=myplug2.so"
```

Quotes are used around the argument value here because semicolon (`;`) is interpreted as a special character by some command interpreters. (Unix shells treat it as a command terminator, for example.)

If multiple `--plugin-load` [435] options are given, only the last one is used. Additional plugins to load may be specified using `--plugin-load-add` [436] options.

If a plugin library is named without any preceding plugin name, the server loads all plugins in the library.

Each plugin is loaded for a single invocation of `mysqld` only. After a restart, the plugin is not loaded unless `--plugin-load` [435] is used again. This is in contrast to `INSTALL PLUGIN`, which adds an entry to the `mysql.plugins` table to cause the plugin to be loaded for every normal server startup.

Under normal startup, the server determines which plugins to load by reading the `mysql.plugins` system table. If the server is started with the `--skip-grant-tables` [439] option, it does not consult the `mysql.plugins` table and does not load plugins listed there. `--plugin-load` [435] enables plugins to be loaded even when `--skip-grant-tables` [439] is given. `--plugin-load` [435] also enables plugins to be loaded at startup under configurations when plugins cannot be loaded at runtime.

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

- `--plugin-load-add=plugin_list` [436]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--plugin-load-add=plugin_list</code>	
<b>Option-File Format</b>	<code>plugin-load-add=plugin_list</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This option complements the `--plugin-load` [435] option. `--plugin-load-add` [436] adds a plugin or plugins to the set of plugins to be loaded at startup. The argument format is the same as for `--plugin-load` [435]. `--plugin-load-add` [436] can be used to avoid specifying a large set of plugins as a single long unwieldy `--plugin-load` [435] argument. This option was added in MySQL 5.6.3.

`--plugin-load-add` [436] can be given in the absence of `--plugin-load` [435], but any instance of `--plugin-load-add` [436] that appears before `--plugin-load` [435] has no effect because `--plugin-load` [435] resets the set of plugins to load. In other words, these options:

```
--plugin-load=x --plugin-load-add=y
```

are equivalent to this option:

```
--plugin-load="x;y"
```

But these options:

```
--plugin-load-add=y --plugin-load=x
```

are equivalent to this option:

```
--plugin-load=x
```

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

- `--port=port_num [437]`, `-P port_num`

<b>Command-Line Format</b>	<code>--port=#</code>	
	<code>-P</code>	
<b>Option-File Format</b>	<code>port</code>	
<b>Option Sets Variable</b>	Yes, <code>port</code> [525]	
<b>Variable Name</b>	<code>port</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>3306</code>
	<b>Range</b>	<code>0 .. 65535</code>

The port number to use when listening for TCP/IP connections. The port number must be 1024 or higher unless the server is started by the `root` system user.

- `--port-open-timeout=num [437]`

<b>Command-Line Format</b>	<code>--port-open-timeout=#</code>	
<b>Option-File Format</b>	<code>port-open-timeout</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

On some systems, when the server is stopped, the TCP/IP port might not become available immediately. If the server is restarted quickly afterward, its attempt to reopen the port can fail. This option indicates how many seconds the server should wait for the TCP/IP port to become free if it cannot be opened. The default is not to wait.

- `--remove [service_name] [437]`

<b>Command-Line Format</b>	<code>--remove [service_name]</code>
----------------------------	--------------------------------------

(Windows only) Remove a MySQL Windows service. The default service name is `MySQL` if no `service_name` value is given. For more information, see [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

- `--safe-mode [437]`

<b>Version Removed</b>	5.6.6
<b>Command-Line Format</b>	<code>--safe-mode</code>
<b>Option-File Format</b>	<code>safe-mode</code>
<b>Deprecated</b>	5.5.26

Skip some optimization stages. This option was removed in MySQL 5.6.6.

- `--safe-user-create` [438]

<b>Command-Line Format</b>	<code>--safe-user-create</code>	
<b>Option-File Format</b>	<code>safe-user-create</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	FALSE

If this option is enabled, a user cannot create new MySQL users by using the `GRANT` statement unless the user has the `INSERT` [755] privilege for the `mysql.user` table or any column in the table. If you want a user to have the ability to create new users that have those privileges that the user has the right to grant, you should grant the user the following privilege:

```
GRANT INSERT(user) ON mysql.user TO 'user_name'@'host_name';
```

This ensures that the user cannot change any privilege columns directly, but has to use the `GRANT` statement to give privileges to other users.

- `--secure-auth` [438]

<b>Command-Line Format</b>	<code>--secure-auth</code>	
<b>Option-File Format</b>	<code>secure-auth</code>	
<b>Option Sets Variable</b>	Yes, <code>secure_auth</code> [538]	
<b>Variable Name</b>	<code>secure_auth</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.4)</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF
	<b>Permitted Values (&gt;= 5.6.5)</b>	
	<b>Type</b>	boolean
	<b>Default</b>	ON

This option causes the server to block connections by clients that attempt to use accounts that have passwords stored in the old (pre-4.1) format. Use it to prevent all use of passwords employing the old format (and hence insecure communication over the network). Before MySQL 5.6.5, this option is disabled by default. As of MySQL 5.6.5, it is enabled by default; to disable it, use `--skip-secure-auth` [438].

Server startup fails with an error if this option is enabled and the privilege tables are in pre-4.1 format. See [Section C.5.2.4](#), “Client does not support authentication protocol”.

The `mysql` client also has a `--secure-auth` [253] option, which prevents connections to a server if the server requires a password in old format for the client account.

- `--secure-file-priv=path` [438]

<b>Command-Line Format</b>	<code>--secure-file-priv=path</code>	
<b>Option-File Format</b>	<code>secure-file-priv=path</code>	
<b>Option Sets Variable</b>	Yes, <code>secure_file_priv</code> [538]	

<b>Variable Name</b>	<code>secure-file-priv</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This option limits the effect of the `LOAD_FILE()` [1181] function and the `LOAD DATA` and `SELECT ... INTO OUTFILE` statements to work only with files in the specified directory.

- `--shared-memory` [439]

Enable shared-memory connections by local clients. This option is available only on Windows.

- `--shared-memory-base-name=name` [439]

The name of shared memory to use for shared-memory connections. This option is available only on Windows. The default name is `MYSQL`. The name is case sensitive.

- `--skip-concurrent-insert` [439]

Turn off the ability to select and insert at the same time on `MyISAM` tables. (This is to be used only if you think you have found a bug in this feature.) See [Section 8.10.3, “Concurrent Inserts”](#).

- `--skip-external-locking` [439]

Do not use external locking (system locking). This affects only `MyISAM` table access. For more information, including conditions under which it can and cannot be used, see [Section 8.10.5, “External Locking”](#).

External locking has been disabled by default since MySQL 4.0.

- `--skip-event-scheduler` [439]

<b>Command-Line Format</b>	<code>--skip-event-scheduler</code>
	<code>--disable-event-scheduler</code>
<b>Option-File Format</b>	<code>skip-event-scheduler</code>

Turns the Event Scheduler `OFF`. This is not the same as disabling the Event Scheduler, which requires setting `--event-scheduler=DISABLED` [422]; see [The `--event-scheduler` Option \[2159\]](#), for more information.

- `--skip-grant-tables` [439]

This option causes the server to start without using the privilege system at all, which gives anyone with access to the server *unrestricted access to all databases*. You can cause a running server to start using the grant tables again by executing `mysqladmin flush-privileges` or `mysqladmin reload` command from a system shell, or by issuing a MySQL `FLUSH PRIVILEGES` statement after connecting to the server. This option also suppresses loading of plugins that were installed with the `INSTALL PLUGIN` statement, user-defined functions (UDFs), and scheduled events. To cause plugins to be loaded anyway, use the `--plugin-load` [435] option.

Note that `FLUSH PRIVILEGES` might be executed implicitly by other actions performed after startup. For example, `mysql_upgrade` flushes the privileges during the upgrade procedure.

- `--skip-host-cache` [439]

Disable use of the internal host cache for faster name-to-IP resolution. In this case, the server performs a DNS lookup every time a client connects. See [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#).

Use of `--skip-host-cache` [439] is similar to setting the `host_cache_size` [487] system variable to 0, but `host_cache_size` [487] is more flexible because it can also be used to resize, enable, or disable the host cache at runtime, not just at server startup.

If you start the server with `--skip-host-cache` [439], that does not prevent changes to the value of `host_cache_size` [487], but such changes have no effect and the cache is not re-enabled even if `host_cache_size` [487] is set larger than 0.

- `--skip-innodb` [1698]

Disable the InnoDB storage engine. In this case, because the default storage engine is InnoDB, the server will not start unless you also use `--default-storage-engine` [474] and `--default-tmp-storage-engine` [474] to set the default to some other engine for both permanent and TEMPORARY tables.

- `--skip-name-resolve` [440]

Do not resolve host names when checking client connections. Use only IP addresses. If you use this option, all `Host` column values in the grant tables must be IP addresses or `localhost`. See Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”.

- `--skip-networking` [440]

Do not listen for TCP/IP connections at all. All interaction with `mysqld` must be made using named pipes or shared memory (on Windows) or Unix socket files (on Unix). This option is highly recommended for systems where only local clients are permitted. See Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”.

- `--skip-partition` [440]

<b>Command-Line Format</b>	<code>--skip-partition</code>
	<code>--disable-partition</code>
<b>Option-File Format</b>	<code>skip-partition</code>

Disables user-defined partitioning. Partitioned tables can be seen using `SHOW TABLES` or by querying the `INFORMATION_SCHEMA.TABLES` table, but cannot be created or modified, nor can data in such tables be accessed. All partition-specific columns in the `INFORMATION_SCHEMA.PARTITIONS` table display `NULL`.

Since `DROP TABLE` removes table definition (`.frm`) files, this statement works on partitioned tables even when partitioning is disabled using the option. The statement, however, does not remove `.par` files associated with partitioned tables in such cases. For this reason, you should avoid dropping partitioned tables with partitioning disabled, or take action to remove the orphaned `.par` files manually.

- `--ssl*`

Options that begin with `--ssl` [810] specify whether to permit clients to connect using SSL and indicate where to find SSL keys and certificates. See Section 6.3.9.4, “SSL Command Options”.

- `--standalone` [440]

<b>Command-Line Format</b>	<code>--standalone</code>
<b>Option-File Format</b>	<code>standalone</code>
<b>Platform Specific</b>	windows

Available on Windows only; instructs the MySQL server not to run as a service.

- `--super-large-pages` [440]

Standard use of large pages in MySQL attempts to use the largest size supported, up to 4MB. Under Solaris, a “super large pages” feature enables uses of pages up to 256MB. This feature is available for recent SPARC platforms. It can be enabled or disabled by using the `--super-large-pages` [440] or `--skip-super-large-pages` [440] option.

- `--symbolic-links` [441], `--skip-symbolic-links` [441]

<b>Command-Line Format</b>	<code>--symbolic-links</code>
<b>Option-File Format</b>	<code>symbolic-links</code>

Enable or disable symbolic link support. This option has different effects on Windows and Unix:

- On Windows, enabling symbolic links enables you to establish a symbolic link to a database directory by creating a `db_name.sym` file that contains the path to the real directory. See [Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”](#).
- On Unix, enabling symbolic links means that you can link a `MyISAM` index file or data file to another directory with the `INDEX DIRECTORY` or `DATA DIRECTORY` options of the `CREATE TABLE` statement. If you delete or rename the table, the files that its symbolic links point to also are deleted or renamed. See [Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”](#).
- `--skip-show-database` [441]

<b>Command-Line Format</b>	<code>--skip-show-database</code>
<b>Option-File Format</b>	<code>skip-show-database</code>
<b>Option Sets Variable</b>	Yes, <code>skip_show_database</code> [540]
<b>Variable Name</b>	<code>skip_show_database</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

This option sets the `skip_show_database` [540] system variable that controls who is permitted to use the `SHOW DATABASES` statement. See [Section 5.1.4, “Server System Variables”](#).

- `--skip-stack-trace` [441]

<b>Command-Line Format</b>	<code>--skip-stack-trace</code>
<b>Option-File Format</b>	<code>skip-stack-trace</code>

Do not write stack traces. This option is useful when you are running `mysqld` under a debugger. On some systems, you also must use this option to get a core file. See [MySQL Internals: Porting to Other Systems](#).

- `--skip-thread-priority` [441]

<b>Version Removed</b>	5.6.1
<b>Command-Line Format</b>	<code>--skip-thread-priority</code>
<b>Option-File Format</b>	<code>skip-thread-priority</code>
<b>Deprecated</b>	5.1.29

Disable using thread priorities for faster response time. This option was unused and was removed in MySQL 5.6.1.

- `--slow-query-log[={0|1}]` [441]

<b>Command-Line Format</b>	<code>--slow-query-log</code>
----------------------------	-------------------------------

<b>Option-File Format</b>	<code>slow-query-log</code>
<b>Option Sets Variable</b>	Yes, <code>slow_query_log</code> [541]
<b>Variable Name</b>	<code>slow_query_log</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	<code>boolean</code>
<b>Default</b>	<code>OFF</code>

Specify the initial slow query log state. With no argument or an argument of 1, the `--slow-query-log` [441] option enables the log. If omitted or given with an argument of 0, the option disables the log.

- `--slow-start-timeout=timeout` [442]

<b>Version Introduced</b>	5.6.5
<b>Command-Line Format</b>	<code>--slow-start-timeout=#</code>
<b>Option-File Format</b>	<code>slow-start-timeout=#</code>
	<b>Permitted Values</b>
<b>Type</b> (windows)	<code>numeric</code>
<b>Default</b>	<code>15000</code>

This option controls the Windows service control manager's service start timeout. The value is the maximum number of milliseconds that the service control manager waits before trying to kill the windows service during startup. The default value is 15000 (15 seconds). If the MySQL service takes too long to start, you may need to increase this value. A value of 0 means there is no timeout.

- `--socket=path` [442]

<b>Command-Line Format</b>	<code>--socket=name</code>
<b>Option-File Format</b>	<code>socket</code>
<b>Option Sets Variable</b>	Yes, <code>socket</code> [542]
<b>Variable Name</b>	<code>socket</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	<code>file name</code>
<b>Default</b>	<code>/tmp/mysql.sock</code>

On Unix, this option specifies the Unix socket file to use when listening for local connections. The default value is `/tmp/mysql.sock`. If this option is given, the server creates the file in the data directory unless an absolute path name is given to specify a different directory. On Windows, the option specifies the pipe name to use when listening for local connections that use a named pipe. The default value is `MySQL` (not case sensitive).

- `--sql-mode=value[,value[,value...]]` [442]

<b>Command-Line Format</b>	<code>--sql-mode=name</code>
<b>Option-File Format</b>	<code>sql-mode</code> 442



<b>Option Sets Variable</b>	Yes, <a href="#">sql_mode</a> [545]
<b>Variable Name</b>	<a href="#">sql_mode</a>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Type</b>	set
<b>Default</b>	''
<b>Valid Values</b>	ALLOW_INVALID_DATES ANSI_QUOTES ERROR_FOR_DIVISION_BY_ZERO HIGH_NOT_PRECEDENCE IGNORE_SPACE NO_AUTO_CREATE_USER NO_AUTO_VALUE_ON_ZERO NO_BACKSLASH_ESCAPES NO_DIR_IN_CREATE NO_ENGINE_SUBSTITUTION NO_FIELD_OPTIONS NO_KEY_OPTIONS NO_TABLE_OPTIONS NO_UNSIGNED_SUBTRACTION NO_ZERO_DATE NO_ZERO_IN_DATE ONLY_FULL_GROUP_BY PAD_CHAR_TO_FULL_LENGTH PIPES_AS_CONCAT REAL_AS_FLOAT STRICT_ALL_TABLES STRICT_TRANS_TABLES
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Type</b>	set
<b>Default</b>	NO_ENGINE_SUBSTITUTION
<b>Valid Values</b>	ALLOW_INVALID_DATES ANSI_QUOTES

	ERROR_FOR_DIVISION_BY_ZERO
	HIGH_NOT_PRECEDENCE
	IGNORE_SPACE
	NO_AUTO_CREATE_USER
	NO_AUTO_VALUE_ON_ZERO
	NO_BACKSLASH_ESCAPES
	NO_DIR_IN_CREATE
	NO_ENGINE_SUBSTITUTION
	NO_FIELD_OPTIONS
	NO_KEY_OPTIONS
	NO_TABLE_OPTIONS
	NO_UNSIGNED_SUBTRACTION
	NO_ZERO_DATE
	NO_ZERO_IN_DATE
	ONLY_FULL_GROUP_BY
	PAD_CHAR_TO_FULL_LENGTH
	PIPES_AS_CONCAT
	REAL_AS_FLOAT
	STRICT_ALL_TABLES
	STRICT_TRANS_TABLES

Set the SQL mode. Prior to MySQL 5.6.6, the default was '' (an empty string); beginning with MySQL 5.6.6, the default is `NO_ENGINE_SUBSTITUTION`. See [Section 5.1.7, “Server SQL Modes”](#).

- `--sysdate-is-now` [\[444\]](#)

<b>Command-Line Format</b>	<code>--sysdate-is-now</code>
<b>Option-File Format</b>	<code>sysdate-is-now</code>
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	FALSE

`SYSDATE()` [\[1223\]](#) by default returns the time at which it executes, not the time at which the statement in which it occurs begins executing. This differs from the behavior of `NOW()` [\[1220\]](#). This option causes `SYSDATE()` [\[1223\]](#) to be an alias for `NOW()` [\[1220\]](#). For information about the implications for binary logging and replication, see the description for `SYSDATE()` [\[1223\]](#) in [Section 12.7, “Date and Time Functions”](#) and for `SET TIMESTAMP` in [Section 5.1.4, “Server System Variables”](#).

- `--tc-heuristic-recover={COMMIT|ROLLBACK}` [\[444\]](#)

<b>Command-Line Format</b>	<code>--tc-heuristic-recover=name</code>
<b>Option-File Format</b>	<code>tc-heuristic-recover</code>
	<b>Permitted Values</b>
	<b>Type</b> enumeration
	<b>Valid Values</b> COMMIT RECOVER

The type of decision to use in the heuristic recovery process. Currently, this option is unused.

- `--temp-pool` [445]

<b>Command-Line Format</b>	<code>--temp-pool</code>
<b>Option-File Format</b>	<code>temp-pool</code>
	<b>Permitted Values</b>
	<b>Type</b> boolean
	<b>Default</b> TRUE

This option causes most temporary files created by the server to use a small set of names, rather than a unique name for each new file. This works around a problem in the Linux kernel dealing with creating many new files with different names. With the old behavior, Linux seems to “leak” memory, because it is being allocated to the directory entry cache rather than to the disk cache. This option is ignored except on Linux.

- `--transaction-isolation=level` [445]

<b>Command-Line Format</b>	<code>--transaction-isolation=name</code>
<b>Option-File Format</b>	<code>transaction-isolation</code>
<b>Option Sets Variable</b>	Yes, <code>tx_isolation</code> [560]
	<b>Permitted Values</b>
	<b>Type</b> enumeration
	<b>Valid Values</b> READ-UNCOMMITTED READ-COMMITTED REPEATABLE-READ SERIALIZABLE

Sets the default transaction isolation level. The `level` value can be `READ-UNCOMMITTED` [1470], `READ-COMMITTED` [1469], `REPEATABLE-READ` [1469], or `SERIALIZABLE` [1470]. See Section 13.3.6, “`SET TRANSACTION Syntax`”.

The default transaction isolation level can also be set at runtime using the `SET TRANSACTION` statement or by setting the `tx_isolation` [560] system variable.

- `--transaction-read-only` [445]

<b>Version Introduced</b>	5.6.5
<b>Command-Line Format</b>	<code>--transaction-read-only</code>
<b>Option-File Format</b>	<code>transaction-read-only</code>
<b>Option Sets Variable</b>	Yes, <code>tx_read_only</code> [561]

	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	OFF	

Sets the default transaction access mode. By default, read-only mode is disabled, so the mode is read/write.

To set the default transaction access mode at runtime, use the `SET TRANSACTION` statement or set the `tx_read_only` [561] system variable. See Section 13.3.6, “`SET TRANSACTION Syntax`”.

This option was added in MySQL 5.6.5.

- `--tmpdir=path` [446], `-t path`

<b>Command-Line Format</b>	<code>--tmpdir=path</code>	
	<code>-t</code>	
<b>Option-File Format</b>	<code>tmpdir</code>	
<b>Option Sets Variable</b>	Yes, <code>tmpdir</code> [559]	
<b>Variable Name</b>	<code>tmpdir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path of the directory to use for creating temporary files. It might be useful if your default `/tmp` directory resides on a partition that is too small to hold temporary tables. This option accepts several paths that are used in round-robin fashion. Paths should be separated by colon characters (“:”) on Unix and semicolon characters (“;”) on Windows. If the MySQL server is acting as a replication slave, you should not set `--tmpdir` [446] to point to a directory on a memory-based file system or to a directory that is cleared when the server host restarts. For more information about the storage location of temporary files, see Section C.5.4.4, “Where MySQL Stores Temporary Files”. A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails.

- `--user={user_name|user_id}` [446], `-u {user_name|user_id}`

<b>Command-Line Format</b>	<code>--user=name</code>	
	<code>-u name</code>	
<b>Option-File Format</b>	<code>user</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Run the `mysqld` server as the user having the name `user_name` or the numeric user ID `user_id`. (“User” in this context refers to a system login account, not a MySQL user listed in the grant tables.)

This option is *mandatory* when starting `mysqld` as `root`. The server changes its user ID during its startup sequence, causing it to run as that particular user rather than as `root`. See Section 6.1.1, “Security Guidelines”.

To avoid a possible security hole where a user adds a `--user=root` [446] option to a `my.cnf` file (thus causing the server to run as `root`), `mysqld` uses only the first `--user` [446] option specified and produces a warning if there are multiple `--user` [446] options. Options in `/etc/my.cnf` and

`$MYSQL_HOME/my.cnf` are processed before command-line options, so it is recommended that you put a `--user [446]` option in `/etc/my.cnf` and specify a value other than `root`. The option in `/etc/my.cnf` is found before any other `--user [446]` options, which ensures that the server runs as a user other than `root`, and that a warning results if any other `--user [446]` option is found.

- `--verbose [447]`, `-v [447]`

Use this option with the `--help [413]` option for detailed help.

- `--version [447]`, `-V`

Display version information and exit.

You can assign a value to a server system variable by using an option of the form `--var_name=value`. For example, `--key_buffer_size=32M [492]` sets the `key_buffer_size [492]` variable to a value of 32MB.

Note that when you assign a value to a variable, MySQL might automatically correct the value to stay within a given range, or adjust the value to the closest permissible value if only certain values are permitted.

If you want to restrict the maximum value to which a variable can be set at runtime with `SET`, you can define this by using the `--maximum-var_name=value` command-line option.

You can change the values of most system variables for a running server with the `SET` statement. See [Section 13.7.4, “SET Syntax”](#).

[Section 5.1.4, “Server System Variables”](#), provides a full description for all variables, and additional information for setting them at server startup and runtime. [Section 8.11.2, “Tuning Server Parameters”](#), includes information on optimizing the server by tuning system variables.

## 5.1.4. Server System Variables

The MySQL server maintains many system variables that indicate how it is configured. Each system variable has a default value. System variables can be set at server startup using options on the command line or in an option file. Most of them can be changed dynamically while the server is running by means of the `SET` statement, which enables you to modify operation of the server without having to stop and restart it. You can refer to system variable values in expressions.

There are several ways to see the names and values of system variables:

- To see the values that a server will use based on its compiled-in defaults and any option files that it reads, use this command:

```
mysqld --verbose --help
```

- To see the values that a server will use based on its compiled-in defaults, ignoring the settings in any option files, use this command:

```
mysqld --no-defaults --verbose --help
```

- To see the current values used by a running server, use the `SHOW VARIABLES` statement.

This section provides a description of each system variable. Variables with no version indicated are present in all MySQL 5.6 releases. For historical information concerning their implementation, please see <http://dev.mysql.com/doc/refman/5.0/en/>, and <http://dev.mysql.com/doc/refman/4.1/en/>.

The following table lists all available system variables.

**Table 5.3. System Variable Summary**

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<code>audit_log_buffer_size [824]</code>			Yes	Global	No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">audit_log_file [825]</a>			Yes	Global	No
<a href="#">audit_log_flush [825]</a>			Yes	Global	Yes
<a href="#">audit_log_policy [826]</a>			Yes	Global	Yes
<a href="#">audit_log_rotate_on_size [826]</a>			Yes	Global	Yes
<a href="#">audit_log_strategy [827]</a>			Yes	Global	No
<a href="#">auto_increment_increment [1965]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">auto_increment_offset [1968]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">autocommit [464]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">automatic_sp_privileges [464]</a>			Yes	Global	Yes
<a href="#">back_log [465]</a>	Yes	Yes	Yes	Global	No
<a href="#">basedir [414]</a>	Yes	Yes	Yes	Global	No
<a href="#">big-tables [414]</a>	Yes	Yes			Yes
- Variable: <a href="#">big_tables</a>			Yes	Both	Yes
<a href="#">bind-address [414]</a>	Yes	Yes	Yes	Global	No
<a href="#">binlog_cache_size [2008]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">binlog_checksum [2009]</a>			Yes	Global	Yes
<a href="#">binlog_direct_non_transactional_updates [2010]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">binlog-format [415]</a>	Yes	Yes			Yes
- Variable: <a href="#">binlog_format</a>			Yes	Both	Yes
<a href="#">binlog_max_flush_queue_time [2011]</a>			Yes	Global	Yes
<a href="#">binlog_order_commits [2012]</a>			Yes	Global	Yes
<a href="#">binlog_row_image [2013]</a>	Yes		Yes	Both	Yes
<a href="#">binlog_rows_query_log_events [2014]</a>			Yes	Both	Yes
<a href="#">binlog_stmt_cache_size [2014]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">bulk_insert_buffer_size [466]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">character_set_client [467]</a>			Yes	Both	Yes
<a href="#">character_set_connection [467]</a>			Yes	Both	Yes
<a href="#">character_set_database [468]<sup>a</sup></a>			Yes	Both	Yes
<a href="#">character-set-filesystem [416]</a>	Yes	Yes			Yes
- Variable: <a href="#">character_set_filesystem</a>			Yes	Both	Yes
<a href="#">character_set_results [468]</a>			Yes	Both	Yes
<a href="#">character-set-server [417]</a>	Yes	Yes			Yes
- Variable: <a href="#">character_set_server</a>			Yes	Both	Yes
<a href="#">character_set_system [469]</a>			Yes	Global	No
<a href="#">character-sets-dir [416]</a>	Yes	Yes			No

Server System Variables

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
- Variable: character_sets_dir			Yes	Global	No
collation_connection [469]			Yes	Both	Yes
collation_database [469] <sup>b</sup>			Yes	Both	Yes
collation-server [417]	Yes	Yes			Yes
- Variable: collation_server			Yes	Both	Yes
completion_type [470]	Yes	Yes	Yes	Both	Yes
concurrent_insert [471]	Yes	Yes	Yes	Global	Yes
connect_timeout [472]	Yes	Yes	Yes	Global	Yes
core_file [472]			Yes	Global	No
daemon_memcached_enable_binlog [1698]	Yes	Yes	Yes	Global	No
daemon_memcached_engine_lib_name [1698]	Yes	Yes	Yes	Global	No
daemon_memcached_engine_lib_path [1699]	Yes	Yes	Yes	Global	No
daemon_memcached_extern [1699]	Yes	Yes	Yes	Global	No
daemon_memcached_flush_batch_size [1700]	Yes	Yes	Yes	Global	No
daemon_memcached_flush_batch_size [1700]	Yes	Yes	Yes	Global	No
datadir [418]	Yes	Yes	Yes	Global	No
date_format [472]			Yes	Global	No
datetime_format [472]			Yes	Global	No
debug [418]	Yes	Yes	Yes	Both	Yes
debug_sync [473]			Yes	Session	Yes
default-storage-engine [420]	Yes	Yes			Yes
- Variable: default_storage_engine			Yes	Both	Yes
default_tmp_storage_engine [474]	Yes	Yes	Yes	Both	Yes
default_week_format [475]	Yes	Yes	Yes	Both	Yes
delay-key-write [421]	Yes	Yes			Yes
- Variable: delay_key_write			Yes	Global	Yes
delayed_insert_limit [476]	Yes	Yes	Yes	Global	Yes
delayed_insert_timeout [476]	Yes	Yes	Yes	Global	Yes
delayed_queue_size [477]	Yes	Yes	Yes	Global	Yes
disable_gtid_unsafe_statements [2020]	Yes	Yes	Yes	Global	No
disable-gtid-unsafe-statements [2018]	Yes	Yes	Yes	Global	No
disconnect_on_expired_passwords [478]	Yes	Yes	Yes	Session	No
div_precision_increment [478]	Yes	Yes	Yes	Both	Yes
end_markers_in_json [479]			Yes	Both	Yes
enforce_gtid_consistency [2021]	Yes	Yes	Yes	Global	No



Server System Variables

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">enforce-gtid-consistency [2018]</a>	Yes	Yes	Yes	Global	No
<a href="#">engine-condition-pushdown [421]</a>	Yes	Yes			Yes
- Variable: <a href="#">engine_condition_pushdown</a>			Yes	Both	Yes
<a href="#">eq_range_index_divisibility [479]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">error_count [480]</a>			Yes	Session	No
<a href="#">event-scheduler [422]</a>	Yes	Yes			Yes
- Variable: <a href="#">event_scheduler</a>			Yes	Global	Yes
<a href="#">expire_logs_days [480]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">explicit_defaults_for_timestamp [481]</a>	Yes	Yes	Yes	Session	No
<a href="#">external_user [482]</a>			Yes	Session	No
<a href="#">flush [482]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">flush_time [482]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">foreign_key_checks [483]</a>			Yes	Both	Yes
<a href="#">ft_boolean_syntax [483]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">ft_max_word_len [484]</a>	Yes	Yes	Yes	Global	No
<a href="#">ft_min_word_len [484]</a>	Yes	Yes	Yes	Global	No
<a href="#">ft_query_expansion_limit [484]</a>	Yes	Yes	Yes	Global	No
<a href="#">ft_stopword_file [485]</a>	Yes	Yes	Yes	Global	No
<a href="#">general-log [423]</a>	Yes	Yes			Yes
- Variable: <a href="#">general_log</a>			Yes	Global	Yes
<a href="#">general_log_file [486]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">group_concat_max_len [486]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">gtid_done [2020]</a>			Yes	Both	No
<a href="#">gtid_executed [2021]</a>			Yes	Both	No
<a href="#">gtid_lost [2022]</a>			Yes	Global	No
<a href="#">gtid_mode [2022]</a>			Yes	Global	No
<a href="#">gtid-mode [2019]</a>	Yes	Yes			No
- Variable: <a href="#">gtid_mode</a>			Yes	Global	No
<a href="#">gtid_next [2022]</a>			Yes	Session	Yes
<a href="#">gtid_owned [2023]</a>			Yes	Both	No
<a href="#">gtid_purged [2023]</a>			Yes	Global	Yes
<a href="#">have_compress [486]</a>			Yes	Global	No
<a href="#">have_crypt [486]</a>			Yes	Global	No
<a href="#">have_csv [486]</a>			Yes	Global	No
<a href="#">have_dynamic_loading [487]</a>			Yes	Global	No
<a href="#">have_geometry [487]</a>			Yes	Global	No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">have_innodb [487]</a>			Yes	Global	No
have_ndbcluster			Yes	Global	No
<a href="#">have_openssl [487]</a>			Yes	Global	No
<a href="#">have_partitioning [487]</a>			Yes	Global	No
<a href="#">have_profiling [487]</a>			Yes	Global	No
<a href="#">have_query_cache [487]</a>			Yes	Global	No
<a href="#">have_rtree_keys [487]</a>			Yes	Global	No
<a href="#">have_ssl [487]</a>			Yes	Global	No
<a href="#">have_symlink [487]</a>			Yes	Global	No
<a href="#">host_cache_size [487]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">hostname [488]</a>			Yes	Global	No
<a href="#">identity [488]</a>			Yes	Session	Yes
<a href="#">ignore-builtin-innodb [1697]</a>	Yes	Yes			No
- Variable: <a href="#">ignore_builtin_innodb</a>			Yes	Global	No
<a href="#">ignore_db_dirs [488]</a>			Yes	Global	No
<a href="#">init_connect [489]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">init-file [424]</a>	Yes	Yes			No
- Variable: <a href="#">init_file</a>			Yes	Global	No
<a href="#">init_slave [1987]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_adaptive_flushing [1700]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_adaptive_flushing_lwm [1701]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_adaptive_hash_index [1701]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_adaptive_max_concurrency [1702]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_additional_mem_pool_size [1703]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_api_bk_commit_interval [1702]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_api_disable_rowlock [1703]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_api_enable_binlog [1703]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_api_enable_mutex [1703]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_api_trx_level [1704]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_autoextend_increment [1704]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_autoinc_lock_mode [1705]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_buffer_pool_dump_at_shutdown [1705]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_buffer_pool_dump_now [1706]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_buffer_pool_file_per_group [1706]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_buffer_pool_instances [1706]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_buffer_pool_load_abort [1707]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_buffer_pool_load_at_startup [1707]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_buffer_pool_load_now [1708]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_buffer_pool_size [1708]</a>	Yes	Yes	Yes	Global	No

Server System Variables

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
innodb_change_buffer_max_size [1709]	Yes	Yes	Yes	Global	Yes
innodb_change_buffer_size [1709]	Yes	Yes	Yes	Global	Yes
innodb_checksum_algorithm [1710]	Yes	Yes	Yes	Global	Yes
innodb_checksums [1710]	Yes	Yes	Yes	Global	No
innodb_cmp_per_indexenabled [1712]	Yes	Yes	Yes	Global	Yes
innodb_commit_concurrency [1712]	Yes	Yes	Yes	Global	Yes
innodb_compression_failure_threshold_pct [1713]	Yes	Yes	Yes	Global	Yes
innodb_compression_level [1713]	Yes	Yes	Yes	Global	Yes
innodb_compression_pct_max [1713]	Yes	Yes	Yes	Global	Yes
innodb_concurrency_tickets [1714]	Yes	Yes	Yes	Global	Yes
innodb_data_file_path [1714]	Yes	Yes	Yes	Global	No
innodb_data_home_dir [1715]	Yes	Yes	Yes	Global	No
innodb_disable_sort_filecache [1715]	Yes	Yes	Yes	Global	Yes
innodb_doublewrite [1715]	Yes	Yes	Yes	Global	No
innodb_fast_shutdown [1716]	Yes	Yes	Yes	Global	Yes
innodb_file_format [1716]	Yes	Yes	Yes	Global	Yes
innodb_file_format_check [1717]	Yes	Yes	Yes	Global	No
innodb_file_format_max [1717]	Yes	Yes	Yes	Global	Yes
innodb_file_per_table [1718]	Yes	Yes	Yes	Global	Yes
innodb_flush_log_at_timeout [1718]	Yes		Yes	Global	Yes
innodb_flush_log_at_commit [1718]	Yes	Yes	Yes	Global	Yes
innodb_flush_method [1719]	Yes	Yes	Yes	Global	No
innodb_flush_neighbors [1721]	Yes	Yes	Yes	Global	Yes
innodb_flushing_avg_pages [1722]	Yes	Yes	Yes	Global	Yes
innodb_force_load_consistent [1722]	Yes	Yes	Yes	Global	No
innodb_force_recovery [1722]	Yes	Yes	Yes	Global	No
innodb_ft_aux_table [1723]	Yes	Yes	Yes	Global	Yes
innodb_ft_cache_size [1724]	Yes	Yes	Yes	Global	No
innodb_ft_enable_stopword [1724]	Yes	Yes	Yes	Global	Yes
innodb_ft_max_token_size [1724]	Yes	Yes	Yes	Global	No
innodb_ft_min_token_size [1725]	Yes	Yes	Yes	Global	No
innodb_ft_num_word_optimize [1725]	Yes	Yes	Yes	Global	Yes
innodb_ft_server_stopword_table [1725]	Yes	Yes	Yes	Global	Yes
innodb_ft_sort_pll_degree [1726]	Yes	Yes	Yes	Global	No
innodb_ft_user_stopword_table [1726]	Yes	Yes	Yes	Both	Yes
innodb_io_capacity [1728]	Yes	Yes	Yes	Global	Yes
innodb_io_capacity_max [1728]	Yes	Yes	Yes	Global	Yes
innodb_large_prefix [1728]	Yes	Yes	Yes	Global	Yes
innodb_lock_wait_timeout [1729]	Yes	Yes	Yes	Both	Yes
innodb_locks_unsafe_for_binlog [1729]	Yes	Yes	Yes	Global	No
innodb_log_buffer_size [1732]	Yes	Yes	Yes	Global	No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
innodb_log_compressed_pages [1732]	Yes	Yes	Yes	Global	Yes
innodb_log_file_size [1732]	Yes	Yes	Yes	Global	No
innodb_log_files_in_group [1733]	Yes	Yes	Yes	Global	No
innodb_log_group_home_dir [1733]	Yes	Yes	Yes	Global	No
innodb_lru_scan_depth [1734]	Yes	Yes	Yes	Global	Yes
innodb_max_dirty_pages_pct [1734]	Yes	Yes	Yes	Global	Yes
innodb_max_dirty_pages_pct_lwm [1735]	Yes	Yes	Yes	Global	Yes
innodb_max_purge_lag [1735]	Yes	Yes	Yes	Global	Yes
innodb_max_purge_lag_delay [1736]	Yes	Yes	Yes	Global	Yes
innodb_mirrored_log_groups [1736]	Yes	Yes	Yes	Global	No
innodb_monitor_disable [1736]	Yes	Yes	Yes	Global	Yes
innodb_monitor_enable [1737]	Yes	Yes	Yes	Global	Yes
innodb_monitor_reset [1737]	Yes	Yes	Yes	Global	Yes
innodb_monitor_reset_reset [1737]	Yes	Yes	Yes	Global	Yes
innodb_old_blocks_pct [1738]	Yes	Yes	Yes	Global	Yes
innodb_old_blocks_time [1738]	Yes	Yes	Yes	Global	Yes
innodb_online_alter_max_size [1739]	Yes	Yes	Yes	Global	Yes
innodb_open_files [1739]	Yes	Yes	Yes	Global	No
innodb_optimize_fulltext_only [1740]	Yes	Yes	Yes	Global	Yes
innodb_page_size [1740]	Yes	Yes	Yes	Global	No
innodb_print_all_deadlocks [1741]	Yes	Yes	Yes	Global	Yes
innodb_purge_batch_size [1741]	Yes	Yes	Yes	Global	Yes
innodb_purge_threads [1742]	Yes	Yes	Yes	Global	No
innodb_random_read_ahead [1742]	Yes	Yes	Yes	Global	Yes
innodb_read_ahead_threshold [1742]	Yes	Yes	Yes	Global	Yes
innodb_read_io_threads [1743]	Yes	Yes	Yes	Global	No
innodb_read_only [1743]	Yes	Yes	Yes	Global	No
innodb_replication_delay [1744]	Yes	Yes	Yes	Global	Yes
innodb_rollback_on_timeout [1744]	Yes	Yes	Yes	Global	No
innodb_rollback_segments [1744]	Yes	Yes	Yes	Global	Yes
innodb_sort_buffer_size [1745]	Yes	Yes	Yes	Global	Yes
innodb_spin_wait_delay [1745]	Yes	Yes	Yes	Global	Yes
innodb_stats_auto_recalc [1746]	Yes	Yes	Yes	Global	Yes
innodb_stats_method [1746]	Yes	Yes	Yes	Global	Yes
innodb_stats_on_metadata [1747]	Yes	Yes	Yes	Global	Yes
innodb_stats_persistent [1747]	Yes	Yes	Yes	Global	Yes
innodb_stats_persistent_sample_pages [1748]	Yes	Yes	Yes	Global	Yes
innodb_stats_sample_pages [1748]	Yes	Yes	Yes	Global	Yes
innodb_stats_transient_sample_pages [1748]	Yes	Yes	Yes	Global	Yes
innodb_strict_mode [1749]	Yes	Yes	Yes	Both	Yes
innodb_support_xa [1749]	Yes	Yes	Yes	Both	Yes

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">innodb_sync_array_size [1750]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_sync_spin_locks [1750]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_table_locks [1751]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">innodb_thread_concurrency [1751]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_thread_sleep_delay [1752]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_undo_directory [1752]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_undo_logs [1753]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_undo_tablespace [1753]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">innodb_use_native_aio [1753]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_use_sys_malloc [1754]</a>	Yes	Yes	Yes	Global	No
<a href="#">innodb_version [1754]</a>			Yes	Global	No
<a href="#">innodb_write_io_threads [1754]</a>	Yes	Yes	Yes	Global	No
<a href="#">insert_id [490]</a>			Yes	Session	Yes
<a href="#">interactive_timeout [490]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">join_buffer_size [490]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">keep_files_on_create [491]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">key_buffer_size [492]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">key_cache_age_threshold [493]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">key_cache_block_size [493]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">key_cache_division_limit [494]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">language [425]</a>	Yes	Yes	Yes	Global	No
<a href="#">large_files_support [494]</a>			Yes	Global	No
<a href="#">large_page_size [494]</a>			Yes	Global	No
<a href="#">large-pages [425]</a>	Yes	Yes			No
- Variable: <a href="#">large_pages</a>			Yes	Global	No
<a href="#">last_insert_id [495]</a>			Yes	Session	Yes
<a href="#">lc-messages [426]</a>	Yes	Yes			Yes
- Variable: <a href="#">lc_messages</a>			Yes	Both	Yes
<a href="#">lc-messages-dir [426]</a>	Yes	Yes			No
- Variable: <a href="#">lc_messages_dir</a>			Yes	Global	No
<a href="#">lc_time_names [495]</a>			Yes	Both	Yes
<a href="#">license [496]</a>			Yes	Global	No
<a href="#">local_infile [496]</a>			Yes	Global	Yes
<a href="#">lock_wait_timeout [497]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">locked_in_memory [497]</a>			Yes	Global	No
<a href="#">log [426]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">log_bin [2008]</a>			Yes	Global	No
<a href="#">log-bin [2002]</a>	Yes	Yes	Yes	Global	No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">log_bin_basename [2015]</a>			Yes	Global	No
<a href="#">log_bin_use_v1_row_events [2015]</a>	Yes	Yes	Yes	Global	No
<a href="#">log-bin-use-v1-row-events [2003]</a>	Yes	Yes			No
- Variable: <a href="#">log_bin_use_v1_row_events</a>			Yes	Global	No
<a href="#">log-error [427]</a>	Yes	Yes			No
- Variable: <a href="#">log_error</a>			Yes	Global	No
<a href="#">log-output [427]</a>	Yes	Yes			Yes
- Variable: <a href="#">log_output</a>			Yes	Global	Yes
<a href="#">log-queries-not-using-indexes [428]</a>	Yes	Yes			Yes
- Variable: <a href="#">log_queries_not_using_indexes</a>			Yes	Global	Yes
<a href="#">log-slave-updates [1971]</a>	Yes	Yes			No
- Variable: <a href="#">log_slave_updates</a>			Yes	Global	No
<a href="#">log_slave_updates [2003]</a>	Yes	Yes	Yes	Global	No
<a href="#">log_slow_admin_statements [498]</a>			Yes	Global	Yes
<a href="#">log-slow-queries [429]</a>	Yes	Yes			Yes
- Variable: <a href="#">log_slow_queries</a>			Yes	Global	Yes
<a href="#">log_slow_slave_statements [1988]</a>			Yes	Global	Yes
<a href="#">log_throttle_queries_not_using_indexes [499]</a>			Yes	Global	Yes
<a href="#">log-warnings [430]</a>	Yes	Yes			Yes
- Variable: <a href="#">log_warnings</a>			Yes	Global	Yes
<a href="#">long_query_time [500]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">low-priority-updates [430]</a>	Yes	Yes			Yes
- Variable: <a href="#">low_priority_updates</a>			Yes	Both	Yes
<a href="#">lower_case_file_system [501]</a>	Yes	Yes	Yes	Global	No
<a href="#">lower_case_table_names [501]</a>	Yes	Yes	Yes	Global	No
<a href="#">master_info_repository [1988]</a>			Yes	Global	Yes
<a href="#">master_verify_checksum [2015]</a>			Yes	Global	Yes
<a href="#">max_allowed_packet [502]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_binlog_cache_size [2016]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_binlog_size [2016]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_binlog_stmt_cache_size [2017]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_connect_errors [503]</a>	Yes	Yes	Yes	Global	Yes

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">max_connections [503]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_delayed_threads [504]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_error_count [504]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_heap_table_size [504]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_insert_delayed_threads [505]</a>			Yes	Both	Yes
<a href="#">max_join_size [505]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_length_for_sort_files [506]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_prepared_stmt_count [506]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_relay_log_size [506]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">max_seeks_for_key [506]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_sort_length [508]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_sp_recursion_depth [508]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_user_connections [509]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">max_write_lock_count [509]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">memlock [431]</a>	Yes	Yes	Yes	Global	No
<a href="#">metadata_locks_cache_size [510]</a>			Yes	Global	No
<a href="#">metadata_locks_hash_instances [510]</a>			Yes	Global	No
<a href="#">min-examined-row-limit [431]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">myisam_data_pointer_size [511]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">myisam_max_sort_file_size [511]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">myisam_mmap_size [511]</a>	Yes	Yes	Yes	Global	No
<a href="#">myisam_recover_options [512]</a>			Yes	Global	No
<a href="#">myisam_repair_threads [512]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">myisam_sort_buffer_size [513]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">myisam_stats_method [513]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">myisam_use_mmap [513]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">named_pipe [514]</a>			Yes	Global	No
<a href="#">ndb_cache_check_time [514]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">net_buffer_length [515]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">net_read_timeout [515]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">net_retry_count [515]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">net_write_timeout [516]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">new [516]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">old [517]</a>	Yes	Yes	Yes	Global	No
<a href="#">old-alter-table [433]</a>	Yes	Yes			Yes
- Variable: <a href="#">old_alter_table</a>			Yes	Both	Yes
<a href="#">old_passwords [517]</a>			Yes	Both	Yes
<a href="#">open-files-limit [434]</a>	Yes	Yes			No



Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
- Variable: open_files_limit			Yes	Global	No
optimizer_join_cache	Yes [519]	Yes	Yes	Both	Yes
optimizer_prune_level	Yes [519]	Yes	Yes	Both	Yes
optimizer_search_depth	Yes [520]	Yes	Yes	Both	Yes
optimizer_switch [520]	Yes	Yes	Yes	Both	Yes
optimizer_trace [523]			Yes	Both	Yes
optimizer_trace_features [523]			Yes	Both	Yes
optimizer_trace_limit [524]			Yes	Both	Yes
optimizer_trace_max_mem_size [524]			Yes	Both	Yes
optimizer_trace_offset [524]			Yes	Both	Yes
partition [434]	Yes	Yes			No
- Variable: have_partitioning			Yes	Global	No
performance_schema	Yes [2307]	Yes	Yes	Global	No
performance_schema_accounts_size [2308]	Yes	Yes	Yes	Global	No
performance_schema_events_stages_history_long_size [2308]	Yes	Yes	Yes	Global	No
performance_schema_events_stages_history_size [2308]	Yes	Yes	Yes	Global	No
performance_schema_events_statements_history_long_size [2309]	Yes	Yes	Yes	Global	No
performance_schema_events_statements_history_size [2309]	Yes	Yes	Yes	Global	No
performance_schema_events_waits_history_long_size [2310]	Yes	Yes	Yes	Global	No
performance_schema_events_waits_history_size [2310]	Yes	Yes	Yes	Global	No
performance_schema_events_size [2310]	Yes	Yes	Yes	Global	No
performance_schema_locks_size [2311]	Yes	Yes	Yes	Global	No
performance_schema_mutex_cond_classes [2311]	Yes	Yes	Yes	Global	No
performance_schema_mutex_cond_instances [2311]	Yes	Yes	Yes	Global	No
performance_schema_mutex_file_classes [2311]	Yes	Yes	Yes	Global	No
performance_schema_mutex_file_handles [2312]	Yes	Yes	Yes	Global	No
performance_schema_mutex_file_instances [2312]	Yes	Yes	Yes	Global	No
performance_schema_mutex_mutex_classes [2312]	Yes	Yes	Yes	Global	No
performance_schema_mutex_mutex_instances [2313]	Yes	Yes	Yes	Global	No
performance_schema_mutex_rwlock_classes [2313]	Yes	Yes	Yes	Global	No
performance_schema_mutex_rwlock_instances [2313]	Yes	Yes	Yes	Global	No
performance_schema_mutex_socket_classes [2314]	Yes	Yes	Yes	Global	No
performance_schema_mutex_socket_instances [2314]	Yes	Yes	Yes	Global	No
performance_schema_mutex_stage_classes [2314]	Yes	Yes	Yes	Global	No
performance_schema_mutex_statement_classes [2314]	Yes	Yes	Yes	Global	No
performance_schema_mutex_table_handles [2315]	Yes	Yes	Yes	Global	No
performance_schema_mutex_table_instances [2315]	Yes	Yes	Yes	Global	No
performance_schema_mutex_thread_classes [2316]	Yes	Yes	Yes	Global	No
performance_schema_mutex_thread_instances [2316]	Yes	Yes	Yes	Global	No
performance_schema_session_connect_attrs_size [2316]	Yes	Yes	Yes	Global	No



Server System Variables

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">performance_schema_actors_size [2317]</a>	Yes	Yes	Yes	Global	No
<a href="#">performance_schema_objects_size [2317]</a>	Yes	Yes	Yes	Global	No
<a href="#">performance_schema_users_size [2317]</a>	Yes	Yes	Yes	Global	No
<a href="#">pid-file [435]</a>	Yes	Yes			No
- Variable: pid_file			Yes	Global	No
<a href="#">plugin_dir [525]</a>	Yes	Yes	Yes	Global	No
<a href="#">port [437]</a>	Yes	Yes	Yes	Global	No
<a href="#">preload_buffer_size [525]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">profiling [526]</a>			Yes	Both	Yes
<a href="#">profiling_history_size [526]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">protocol_version [526]</a>			Yes	Global	No
<a href="#">proxy_user [526]</a>			Yes	Session	No
<a href="#">pseudo_slave_mode [526]</a>			Yes	Session	Yes
<a href="#">pseudo_thread_id [527]</a>			Yes	Session	Yes
<a href="#">query_alloc_block_size [527]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">query_cache_limit [527]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">query_cache_min_res_limit [528]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">query_cache_size [528]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">query_cache_type [528]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">query_cache_wlock_invalidate [530]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">query_prealloc_size [530]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">rand_seed1 [531]</a>			Yes	Session	Yes
<a href="#">rand_seed2 [531]</a>			Yes	Session	Yes
<a href="#">range_alloc_block_size [531]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">read_buffer_size [532]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">read_only [532]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">read_rnd_buffer_size [533]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">relay-log [1974]</a>	Yes	Yes			No
- Variable: relay_log			Yes	Global	No
<a href="#">relay_log_basename [1989]</a>			Yes	Global	No
<a href="#">relay-log-index [1974]</a>	Yes	Yes			No
- Variable: relay_log_index			Yes	Global	No
<a href="#">relay_log_index [1989]</a>	Yes	Yes	Yes	Global	No
<a href="#">relay_log_info_file [1990]</a>	Yes	Yes	Yes	Global	No
<a href="#">relay_log_info_repository [1990]</a>			Yes	Global	Yes
<a href="#">relay_log_purge [534]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">relay_log_recovery [534]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">relay_log_space_limit [534]</a>	Yes	Yes	Yes	Global	No
<a href="#">report-host [1980]</a>	Yes	Yes			No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
- Variable: report_host			Yes	Global	No
<a href="#">report-password [1981]</a>	Yes	Yes			No
- Variable: report_password			Yes	Global	No
<a href="#">report-port [1981]</a>	Yes	Yes			No
- Variable: report_port			Yes	Global	No
<a href="#">report-user [1982]</a>	Yes	Yes			No
- Variable: report_user			Yes	Global	No
<a href="#">rpl_semi_sync_master_enabled [536]</a>			Yes	Global	Yes
<a href="#">rpl_semi_sync_master_timeout [536]</a>			Yes	Global	Yes
<a href="#">rpl_semi_sync_master_trace_level [536]</a>			Yes	Global	Yes
<a href="#">rpl_semi_sync_master_wait_no_slave [537]</a>			Yes	Global	Yes
<a href="#">rpl_semi_sync_slave_enabled [537]</a>			Yes	Global	Yes
<a href="#">rpl_semi_sync_slave_trace_level [538]</a>			Yes	Global	Yes
<a href="#">secure-auth [438]</a>	Yes	Yes			Yes
- Variable: secure_auth			Yes	Global	Yes
<a href="#">secure-file-priv [438]</a>	Yes	Yes			No
- Variable: secure_file_priv			Yes	Global	No
<a href="#">server-id [1955]</a>	Yes	Yes			Yes
- Variable: server_id			Yes	Global	Yes
<a href="#">server_uuid [1955]</a>			Yes	Global	No
<a href="#">sha256_password_private_key_path [539]</a>			Yes	Global	No
<a href="#">sha256_password_public_key_path [539]</a>			Yes	Global	No
<a href="#">shared_memory [540]</a>			Yes	Global	No
<a href="#">shared_memory_base_name [540]</a>			Yes	Global	No
<a href="#">skip-external-locking [439]</a>	Yes	Yes			No
- Variable: skip_external_locking			Yes	Global	No
<a href="#">skip-name-resolve [440]</a>	Yes	Yes			No
- Variable: skip_name_resolve			Yes	Global	No
<a href="#">skip-networking [440]</a>	Yes	Yes			No
- Variable: skip_networking			Yes	Global	No

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">skip-show-database [441]</a>	Yes	Yes			No
- Variable: skip_show_database			Yes	Global	No
<a href="#">slave_allow_batching [1968]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave_checkpoint_group [1991]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave_checkpoint_period [1991]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave_compressed_protocol [1992]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave_exec_mode [1992]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave-load-tmpdir [1985]</a>	Yes	Yes			No
- Variable: slave_load_tmpdir			Yes	Global	No
<a href="#">slave_max_allowed_packet [1969]</a>			Yes	Global	Yes
<a href="#">slave-net-timeout [1985]</a>	Yes	Yes			Yes
- Variable: slave_net_timeout			Yes	Global	Yes
<a href="#">slave_parallel_workers [1993]</a>			Yes	Global	Yes
<a href="#">slave_pending_jobs_size_max [1994]</a>			Yes	Global	Yes
<a href="#">slave_rows_search_algorithms [1969]</a>			Yes	Global	Yes
<a href="#">slave-skip-errors [1986]</a>	Yes	Yes			No
- Variable: slave_skip_errors			Yes	Global	No
<a href="#">slave_sql_verify_checksum [1995]</a>			Yes	Global	Yes
<a href="#">slave_transaction_retries [1995]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slave_type_conversion [1996]</a>	Yes	Yes	Yes	Global	No
<a href="#">slow_launch_time [544]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">slow-query-log [441]</a>	Yes	Yes			Yes
- Variable: slow_query_log			Yes	Global	Yes
<a href="#">slow_query_log_file [544]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">socket [442]</a>	Yes	Yes	Yes	Global	No
<a href="#">sort_buffer_size [542]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">sql_auto_is_null [543]</a>			Yes	Both	Yes
<a href="#">sql_big_selects [543]</a>			Yes	Both	Yes
<a href="#">sql_big_tables [414]</a>			Yes	Both	Yes
<a href="#">sql_buffer_result [544]</a>			Yes	Both	Yes
<a href="#">sql_log_bin [544]</a>			Yes	Both	Yes
<a href="#">sql_log_off [544]</a>			Yes	Both	Yes
<a href="#">sql_low_priority_updates [500]</a>			Yes	Both	Yes
<a href="#">sql_max_join_size [505]</a>			Yes	Both	Yes

Server System Variables

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">sql-mode [442]</a>	Yes	Yes			Yes
- Variable: <a href="#">sql_mode</a>			Yes	Both	Yes
<a href="#">sql_notes [546]</a>			Yes	Both	Yes
<a href="#">sql_quote_show_create [546]</a>			Yes	Both	Yes
<a href="#">sql_safe_updates [547]</a>			Yes	Both	Yes
<a href="#">sql_select_limit [547]</a>			Yes	Both	Yes
<a href="#">sql_slave_skip_counter [1996]</a>			Yes	Global	Yes
<a href="#">sql_warnings [547]</a>			Yes	Both	Yes
<a href="#">ssl-ca [810]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_ca</a>			Yes	Global	No
<a href="#">ssl-capath [810]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_capath</a>			Yes	Global	No
<a href="#">ssl-cert [811]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_cert</a>			Yes	Global	No
<a href="#">ssl-cipher [811]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_cipher</a>			Yes	Global	No
<a href="#">ssl-crl [811]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_crl</a>			Yes	Global	No
<a href="#">ssl-crlpath [811]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_crlpath</a>			Yes	Global	No
<a href="#">ssl-key [811]</a>	Yes	Yes			No
- Variable: <a href="#">ssl_key</a>			Yes	Global	No
<a href="#">storage_engine [549]</a>			Yes	Both	Yes
<a href="#">stored_program_cache [549]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">sync_binlog [2017]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">sync_frm [550]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">sync_master_info [1997]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">sync_relay_log [1998]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">sync_relay_log_info [1999]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">system_time_zone [550]</a>			Yes	Global	No
<a href="#">table_definition_cache [550]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">table_open_cache [551]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">table_open_cache_instances [552]</a>	Yes	Yes	Yes	Global	No
<a href="#">thread_cache_size [552]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">thread_concurrency [553]</a>	Yes	Yes	Yes	Global	No
<a href="#">thread_handling [553]</a>	Yes	Yes	Yes	Global	No
<a href="#">thread_pool_algorithm [554]</a>	Yes	Yes	Yes	Global	No
<a href="#">thread_pool_high_priority_connection [554]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">thread_pool_max_unused_threads [555]</a>	Yes	Yes	Yes	Global	Yes

Name	Cmd-Line	Option file	System Var	Var Scope	Dynamic
<a href="#">thread_pool_prio_kickoff_timer [555]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">thread_pool_size [556]</a>	Yes	Yes	Yes	Global	No
<a href="#">thread_pool_stall_limit [556]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">thread_stack [556]</a>	Yes	Yes	Yes	Global	No
<a href="#">time_format [557]</a>			Yes	Global	No
<a href="#">time_zone [557]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">timed_mutexes [557]</a>	Yes	Yes	Yes	Global	Yes
<a href="#">timestamp [558]</a>			Yes	Session	Yes
<a href="#">tmp_table_size [558]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">tmpdir [446]</a>	Yes	Yes	Yes	Global	No
<a href="#">transaction_alloc_block_size [559]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">transaction_prealloc_size [560]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">tx_isolation [560]</a>			Yes	Both	Yes
<a href="#">tx_read_only [561]</a>			Yes	Both	Yes
<a href="#">unique_checks [561]</a>			Yes	Both	Yes
<a href="#">updatable_views_with_limit [562]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">validate_password_dictionary_file [743]</a>			Yes	Global	No
<a href="#">validate_password_length [744]</a>			Yes	Global	Yes
<a href="#">validate_password_mixed_case_count [744]</a>			Yes	Global	Yes
<a href="#">validate_password_number_count [745]</a>			Yes	Global	Yes
<a href="#">validate_password_policy [745]</a>			Yes	Global	Yes
<a href="#">validate_password_special_char_count [745]</a>			Yes	Global	Yes
<a href="#">validate_user_plugins [562]</a>			Yes	Global	No
<a href="#">version [563]</a>			Yes	Global	No
<a href="#">version_comment [563]</a>			Yes	Global	No
<a href="#">version_compile_machine [563]</a>			Yes	Global	No
<a href="#">version_compile_os [563]</a>			Yes	Global	No
<a href="#">wait_timeout [563]</a>	Yes	Yes	Yes	Both	Yes
<a href="#">warning_count [564]</a>			Yes	Session	No

<sup>a</sup>This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.

<sup>b</sup>This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.

For additional system variable information, see these sections:

- [Section 5.1.5, “Using System Variables”](#), discusses the syntax for setting and displaying system variable values.
- [Section 5.1.5.2, “Dynamic System Variables”](#), lists the variables that can be set at runtime.
- Information on tuning system variables can be found in [Section 8.11.2, “Tuning Server Parameters”](#).
- [Section 14.2.6, “InnoDB Startup Options and System Variables”](#), lists InnoDB system variables.
- [MySQL Cluster System Variables](#), lists system variables which are specific to MySQL Cluster.
- For information on server system variables specific to replication, see [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).

**Note**

Some of the following variable descriptions refer to “enabling” or “disabling” a variable. These variables can be enabled with the `SET` statement by setting them to `ON` or `1`, or disabled by setting them to `OFF` or `0`. However, before MySQL 5.6.2, to set such a variable on the command line or in an option file, you must set it to `1` or `0`; setting it to `ON` or `OFF` will not work. For example, on the command line, `--delay_key_write=1` [475] works but `--delay_key_write=ON` [475] does not. As of MySQL 5.6.2, boolean variables can be set at startup to the values `ON`, `TRUE`, `OFF`, and `FALSE` (not case sensitive). See [Section 4.2.3.2, “Program Option Modifiers”](#).

Some system variables control the size of buffers or caches. For a given buffer, the server might need to allocate internal data structures. These structures typically are allocated from the total memory allocated to the buffer, and the amount of space required might be platform dependent. This means that when you assign a value to a system variable that controls a buffer size, the amount of space actually available might differ from the value assigned. In some cases, the amount might be less than the value assigned. It is also possible that the server will adjust a value upward. For example, if you assign a value of 0 to a variable for which the minimal value is 1024, the server will set the value to 1024.

Values for buffer sizes, lengths, and stack sizes are given in bytes unless otherwise specified.

Some system variables take file name values. Unless otherwise specified, the default file location is the data directory if the value is a relative path name. To specify the location explicitly, use an absolute path name. Suppose that the data directory is `/var/mysql/data`. If a file-valued variable is given as a relative path name, it will be located under `/var/mysql/data`. If the value is an absolute path name, its location is as given by the path name.

- `authentication_windows_log_level` [463]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--authentication_windows_log_level</code>	
<b>Option-File Format</b>	<code>authentication_windows_log_level</code>	
<b>Option Sets Variable</b>	Yes, <code>authentication_windows_log_level</code> [463]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 4</code>

This variable is available only if the `authentication_windows` Windows authentication plugin is enabled and debugging code is enabled. See [Section 6.3.7.4, “The Windows Native Authentication Plugin”](#).

This variable sets the logging level for the Windows authentication plugin. The following table shows the permitted values.

Value	Description
0	No logging
1	Log only error messages
2	Log level 1 messages and warning messages
3	Log level 2 messages and information notes
4	Log level 3 messages and debug messages

This variable was added in MySQL 5.6.10.

- `authentication_windows_use_principal_name` [464]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--authentication_windows_use_principal_name</code>	
<b>Option-File Format</b>	<code>authentication_windows_use_principal_name</code>	
<b>Option Sets Variable</b>	Yes, <code>authentication_windows_use_principal_name</code> [464]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

This variable is available only if the `authentication_windows` Windows authentication plugin is enabled. See [Section 6.3.7.4, “The Windows Native Authentication Plugin”](#).

A client that authenticates using the `InitSecurityContext()` function should provide a string identifying the service to which it connects (`targetName`). MySQL uses the principal name (UPN) of the account under which the server is running. The UPN has the form `user_id@computer_name` and need not be registered anywhere to be used. This UPN is sent by the server at the beginning of authentication handshake.

This variable controls whether the server sends the UPN in the initial challenge. By default, the variable is enabled. For security reasons, it can be disabled to avoid sending the server's account name to a client in clear text. If the variable is disabled, the server always sends a `0x00` byte in the first challenge, the client does not specify `targetName`, and as a result, NTLM authentication is used.

If the server fails to obtain its UPN (which will happen primarily in environments that do not support Kerberos authentication), the UPN is not sent by the server and NTLM authentication is used.

This variable was added in MySQL 5.6.10.

- `autocommit` [464]

<b>Command-Line Format</b>	<code>--autocommit[=#]</code>	
<b>Option-File Format</b>	<code>autocommit</code>	
<b>Option Sets Variable</b>	Yes, <code>autocommit</code> [464]	
<b>Variable Name</b>	<code>autocommit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>

The autocommit mode. If set to 1, all changes to a table take effect immediately. If set to 0, you must use `COMMIT` to accept a transaction or `ROLLBACK` to cancel it. If `autocommit` [464] is 0 and you change it to 1, MySQL performs an automatic `COMMIT` of any open transaction. Another way to begin a transaction is to use a `START TRANSACTION` or `BEGIN` statement. See [Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#).

By default, client connections begin with `autocommit` [464] set to 1. To cause clients to begin with a default of 0, set the global `autocommit` [464] value by starting the server with the `--autocommit=0` [464] option. To set the variable using an option file, include these lines:

```
[mysqld]
autocommit=0
```

- `automatic_sp_privileges` [464]

<b>Variable Name</b>	<code>automatic_sp_privileges</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>TRUE</code>

When this variable has a value of 1 (the default), the server automatically grants the `EXECUTE` [754] and `ALTER ROUTINE` [754] privileges to the creator of a stored routine, if the user cannot already execute and alter or drop the routine. (The `ALTER ROUTINE` [754] privilege is required to drop the routine.) The server also automatically drops those privileges from the creator when the routine is dropped. If `automatic_sp_privileges` [464] is 0, the server does not automatically add or drop these privileges.

The creator of a routine is the account used to execute the `CREATE` statement for it. This might not be the same as the account named as the `DEFINER` in the routine definition.

See also [Section 18.2.2, “Stored Routines and MySQL Privileges”](#).

- `back_log` [465]

<b>Command-Line Format</b>	<code>--back_log=#</code>	
<b>Option-File Format</b>	<code>back_log</code>	
<b>Option Sets Variable</b>	Yes, <code>back_log</code> [465]	
<b>Variable Name</b>	<code>back_log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>50</code>
	<b>Range</b>	<code>1 .. 65535</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>
	<b>Range</b>	<code>1 .. 65535</code>

The number of outstanding connection requests MySQL can have. This comes into play when the main MySQL thread gets very many connection requests in a very short time. It then takes some time (although very little) for the main thread to check the connection and start a new thread. The `back_log` [465] value indicates how many requests can be stacked during this short time before MySQL momentarily stops answering new requests. You need to increase this only if you expect a large number of connections in a short period of time.

In other words, this value is the size of the listen queue for incoming TCP/IP connections. Your operating system has its own limit on the size of this queue. The manual page for the Unix `listen()` system call should have more details. Check your OS documentation for the maximum value for this variable. `back_log` [465] cannot be set higher than your operating system limit.

As of MySQL 5.6.6, the default value is based on the following formula, capped to a limit of 900:

```
50 + (max_connections / 5)
```



Before 5.6.6, the default is 50.

- [basedir \[466\]](#)

<b>Command-Line Format</b>	<code>--basedir=path</code>	
	<code>-b</code>	
<b>Option-File Format</b>	<code>basedir</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">basedir [466]</a>	
<b>Variable Name</b>	<code>basedir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The MySQL installation base directory. This variable can be set with the `--basedir [414]` option. Relative path names for other variables usually are resolved relative to the base directory.

- [big\\_tables \[466\]](#)

If set to 1, all temporary tables are stored on disk rather than in memory. This is a little slower, but the error `The table tbl_name is full` does not occur for `SELECT` operations that require a large temporary table. The default value for a new connection is 0 (use in-memory temporary tables). Normally, you should never need to set this variable, because in-memory tables are automatically converted to disk-based tables as required.

- [bind\\_address \[466\]](#)

<b>Command-Line Format</b>	<code>--bind-address=addr</code>	
<b>Option-File Format</b>	<code>bind-address=addr</code>	
<b>Variable Name</b>	<code>bind-address</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>0.0.0.0</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>*</code>

The value of the `--bind-address [414]` option. This variable was added in MySQL 5.6.1.

- [bulk\\_insert\\_buffer\\_size \[466\]](#)

<b>Command-Line Format</b>	<code>--bulk_insert_buffer_size=#</code>
<b>Option-File Format</b>	<code>bulk_insert_buffer_size</code>
<b>Option Sets Variable</b>	Yes, <a href="#">bulk_insert_buffer_size [466]</a>
<b>Variable Name</b>	<code>bulk_insert_buffer_size</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes

	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	8388608
<b>Range</b>	0 .. 4294967295
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	8388608
<b>Range</b>	0 .. 18446744073709547520

MyISAM uses a special tree-like cache to make bulk inserts faster for `INSERT ... SELECT`, `INSERT ... VALUES (...), (...), ...`, and `LOAD DATA INFILE` when adding data to nonempty tables. This variable limits the size of the cache tree in bytes per thread. Setting it to 0 disables this optimization. The default value is 8MB.

- `character_set_client` [467]

<b>Variable Name</b>	<code>character_set_client</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	string

The character set for statements that arrive from the client. The session value of this variable is set using the character set requested by the client when the client connects to the server. (Many clients support a `--default-character-set` option to enable this character set to be specified explicitly. See also [Section 10.1.4, “Connection Character Sets and Collations”](#).) The global value of the variable is used to set the session value in cases when the client-requested value is unknown or not available, or the server is configured to ignore client requests:

- The client is from a version of MySQL older than MySQL 4.1, and thus does not request a character set.
- The client requests a character set not known to the server. For example, a Japanese-enabled client requests `sjis` when connecting to a server not configured with `sjis` support.
- `mysqld` was started with the `--skip-character-set-client-handshake` [416] option, which causes it to ignore client character set configuration. This reproduces MySQL 4.0 behavior and is useful should you wish to upgrade the server without upgrading all the clients.

`ucs2`, `utf16`, `utf16le`, and `utf32` cannot be used as a client character set, which means that they also do not work for `SET NAMES` or `SET CHARACTER SET`.

- `character_set_connection` [467]

<b>Variable Name</b>	<code>character_set_connection</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>

	<b>Type</b>	string
--	-------------	--------

The character set used for literals that do not have a character set introducer and for number-to-string conversion.

- [character\\_set\\_database \[468\]](#)

<b>Variable Name</b>	character_set_database	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
<b>Footnote</b>	This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The character set used by the default database. The server sets this variable whenever the default database changes. If there is no default database, the variable has the same value as [character\\_set\\_server \[468\]](#).

- [character\\_set\\_filesystem \[468\]](#)

<b>Command-Line Format</b>	--character-set-filesystem=name	
<b>Option-File Format</b>	character-set-filesystem	
<b>Option Sets Variable</b>	Yes, <a href="#">character_set_filesystem [468]</a>	
<b>Variable Name</b>	character_set_filesystem	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The file system character set. This variable is used to interpret string literals that refer to file names, such as in the `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` statements and the `LOAD_FILE()` [1181] function. Such file names are converted from [character\\_set\\_client \[467\]](#) to [character\\_set\\_filesystem \[468\]](#) before the file opening attempt occurs. The default value is `binary`, which means that no conversion occurs. For systems on which multi-byte file names are permitted, a different value may be more appropriate. For example, if the system represents file names using UTF-8, set [character\\_set\\_filesystem \[468\]](#) to `'utf8'`.

- [character\\_set\\_results \[468\]](#)

<b>Variable Name</b>	character_set_results	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The character set used for returning query results such as result sets or error messages to the client.

- [character\\_set\\_server \[468\]](#)

<b>Command-Line Format</b>	468	--character-set-server
----------------------------	-----	------------------------

<b>Option-File Format</b>	<code>character-set-server</code>	
<b>Option Sets Variable</b>	Yes, <code>character_set_server</code> [468]	
<b>Variable Name</b>	<code>character_set_server</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The server's default character set.

- `character_set_system` [469]

<b>Variable Name</b>	<code>character_set_system</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The character set used by the server for storing identifiers. The value is always `utf8`.

- `character_sets_dir` [469]

<b>Command-Line Format</b>	<code>--character-sets-dir=path</code>	
<b>Option-File Format</b>	<code>character-sets-dir=path</code>	
<b>Option Sets Variable</b>	Yes, <code>character_sets_dir</code> [469]	
<b>Variable Name</b>	<code>character-sets-dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	directory name

The directory where character sets are installed.

- `collation_connection` [469]

<b>Variable Name</b>	<code>collation_connection</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The collation of the connection character set.

- `collation_database` [469]

<b>Variable Name</b>	<code>collation_database</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	

<b>Footnote</b>	This option is dynamic, but only the server should set this information. You should not set the value of this variable manually.	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The collation used by the default database. The server sets this variable whenever the default database changes. If there is no default database, the variable has the same value as `collation_server` [470].

- `collation_server` [470]

<b>Command-Line Format</b>	<code>--collation-server</code>	
<b>Option-File Format</b>	<code>collation-server</code>	
<b>Option Sets Variable</b>	Yes, <code>collation_server</code> [470]	
<b>Variable Name</b>	<code>collation_server</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The server's default collation.

- `completion_type` [470]

<b>Command-Line Format</b>	<code>--completion_type=#</code>	
<b>Option-File Format</b>	<code>completion_type</code>	
<b>Option Sets Variable</b>	Yes, <code>completion_type</code> [470]	
<b>Variable Name</b>	<code>completion_type</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	NO_CHAIN
	<b>Valid Values</b>	NO_CHAIN CHAIN RELEASE 0 1 2

The transaction completion type. This variable can take the values shown in the following table. The variable can be assigned using either the name values or corresponding integer values.

Value	Description
NO_CHAIN (or 0)	COMMIT and ROLLBACK are unaffected. This is the default value.

Value	Description
CHAIN (or 1)	COMMIT and ROLLBACK are equivalent to COMMIT AND CHAIN and ROLLBACK AND CHAIN, respectively. (A new transaction starts immediately with the same isolation level as the just-terminated transaction.)
RELEASE (or 2)	COMMIT and ROLLBACK are equivalent to COMMIT RELEASE and ROLLBACK RELEASE, respectively. (The server disconnects after terminating the transaction.)

completion\_type [470] affects transactions that begin with START TRANSACTION or BEGIN and end with COMMIT or ROLLBACK. It does not apply to implicit commits resulting from execution of the statements listed in Section 13.3.3, “Statements That Cause an Implicit Commit”. It also does not apply for XA COMMIT, XA ROLLBACK, or when autocommit=1 [464].

- concurrent\_insert [471]

<b>Command-Line Format</b>	--concurrent_insert[#]	
<b>Option-File Format</b>	concurrent_insert	
<b>Option Sets Variable</b>	Yes, concurrent_insert [471]	
<b>Variable Name</b>	concurrent_insert	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	AUTO
	<b>Valid Values</b>	NEVER AUTO ALWAYS 0 1 2

If AUTO (the default), MySQL permits INSERT and SELECT statements to run concurrently for MyISAM tables that have no free blocks in the middle of the data file. If you start mysqld with --skip-new [516], this variable is set to NEVER.

This variable can take the values shown in the following table. The variable can be assigned using either the name values or corresponding integer values.

Value	Description
NEVER (or 0)	Disables concurrent inserts
AUTO (or 1)	(Default) Enables concurrent insert for MyISAM tables that do not have holes
ALWAYS (or 2)	Enables concurrent inserts for all MyISAM tables, even those that have holes. For a table with a hole, new rows are inserted at the end of the table if it is in use by another thread. Otherwise, MySQL acquires a normal write lock and inserts the row into the hole.

See also Section 8.10.3, “Concurrent Inserts”.

- [connect\\_timeout \[472\]](#)

<b>Command-Line Format</b>	<code>--connect_timeout=#</code>	
<b>Option-File Format</b>	<code>connect_timeout</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">connect_timeout [472]</a>	
<b>Variable Name</b>	<code>connect_timeout</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10</code>

The number of seconds that the `mysqld` server waits for a connect packet before responding with `Bad handshake`. The default value is 10 seconds.

Increasing the [connect\\_timeout \[472\]](#) value might help if clients frequently encounter errors of the form `Lost connection to MySQL server at 'XXX', system error: errno`.

- [core\\_file \[472\]](#)

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	<code>core_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether to write a core file if the server crashes. This variable is set by the `--core-file [418]` option. It was added in MySQL 5.6.2.

- [datadir \[472\]](#)

<b>Command-Line Format</b>	<code>--datadir=path</code>	
	<code>-h</code>	
<b>Option-File Format</b>	<code>datadir</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">datadir [472]</a>	
<b>Variable Name</b>	<code>datadir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The MySQL data directory. This variable can be set with the `--datadir [418]` option.

- [date\\_format \[472\]](#)

This variable is unused. It is deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release.

- [datetime\\_format \[472\]](#)

This variable is unused. It is deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release.

- [debug \[473\]](#)

<b>Command-Line Format</b>	<code>--debug[=debug_options]</code>	
<b>Option-File Format</b>	<code>debug</code>	
<b>Variable Name</b>	<code>debug</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>'d:t:o,/tmp/mysqld.trace'</code>

This variable indicates the current debugging settings. It is available only for servers built with debugging support. The initial value comes from the value of instances of the `--debug [418]` option given at server startup. The global and session values may be set at runtime; the [SUPER \[756\]](#) privilege is required, even for the session value.

Assigning a value that begins with `+` or `-` cause the value to added to or subtracted from the current value:

```
mysql> SET debug = 'T';
mysql> SELECT @@debug;
+-----+
| @@debug |
+-----+
| T       |
+-----+

mysql> SET debug = '+P';
mysql> SELECT @@debug;
+-----+
| @@debug |
+-----+
| P:T     |
+-----+

mysql> SET debug = '-P';
mysql> SELECT @@debug;
+-----+
| @@debug |
+-----+
| T       |
+-----+
```

- [debug\\_sync \[473\]](#)

<b>Variable Name</b>	<code>debug_sync</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This variable is the user interface to the Debug Sync facility. Use of Debug Sync requires that MySQL be configured with the `-DENABLE_DEBUG_SYNC=1 [134]` option (see [Section 2.9.4, "MySQL Source-Configuration Options"](#)). If Debug Sync is not compiled in, this system variable is not available.



The global variable value is read only and indicates whether the facility is enabled. By default, Debug Sync is disabled and the value of `debug_sync` [473] is `OFF`. If the server is started with `--debug-sync-timeout=N` [419], where *N* is a timeout value greater than 0, Debug Sync is enabled and the value of `debug_sync` [473] is `ON - current signal` followed by the signal name. Also, *N* becomes the default timeout for individual synchronization points.

The session value can be read by any user and will have the same value as the global variable. The session value can be set by users that have the `SUPER` [756] privilege to control synchronization points.

For a description of the Debug Sync facility and how to use synchronization points, see [MySQL Internals: Test Synchronization](#).

- `default_storage_engine` [474]

<b>Command-Line Format</b>	<code>--default-storage-engine=name</code>	
<b>Option-File Format</b>	<code>default-storage-engine</code>	
<b>Option Sets Variable</b>	Yes, <code>default_storage_engine</code> [474]	
<b>Variable Name</b>	<code>default_storage_engine</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>InnoDB</code>

The default storage engine. As of MySQL 5.6.3, this variable sets the storage engine for permanent tables only. To set the storage engine for `TEMPORARY` tables, set the `default_tmp_storage_engine` [474] system variable.

`default_storage_engine` [474] should be used in preference to `storage_engine` [549], which is deprecated.

If you disable the default storage engine at server startup, you must set the default engine for both permanent and `TEMPORARY` tables to a different engine or the server will not start.

- `default_tmp_storage_engine` [474]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--default_tmp_storage_engine=name</code>	
<b>Option-File Format</b>	<code>default_tmp_storage_engine</code>	
<b>Option Sets Variable</b>	Yes, <code>default_tmp_storage_engine</code> [474]	
<b>Variable Name</b>	<code>default_tmp_storage_engine</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>InnoDB</code>

The default storage engine for `TEMPORARY` tables (created with `CREATE TEMPORARY TABLE`). To set the storage engine for permanent tables, set the `default_storage_engine` [474] system variable.

If you disable the default storage engine at server startup, you must set the default engine for both permanent and `TEMPORARY` tables to a different engine or the server will not start.

`default_tmp_storage_engine` [474] was added in MySQL 5.6.3.

- `default_week_format` [475]

<b>Command-Line Format</b>	<code>--default-week-format=#</code>	
<b>Option-File Format</b>	<code>default-week-format</code>	
<b>Option Sets Variable</b>	Yes, <code>default-week-format</code> [475]	
<b>Variable Name</b>	<code>default-week-format</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 7</code>

The default mode value to use for the `WEEK()` [1228] function. See Section 12.7, “Date and Time Functions”.

- `delay_key_write` [475]

<b>Command-Line Format</b>	<code>--delay-key-write[=name]</code>	
<b>Option-File Format</b>	<code>delay-key-write</code>	
<b>Option Sets Variable</b>	Yes, <code>delay-key-write</code> [475]	
<b>Variable Name</b>	<code>delay-key-write</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>ON</code>
	<b>Valid Values</b>	<code>ON</code> <code>OFF</code> <code>ALL</code>

This option applies only to `MyISAM` tables. It can have one of the following values to affect handling of the `DELAY_KEY_WRITE` table option that can be used in `CREATE TABLE` statements.

Option	Description
<code>OFF</code>	<code>DELAY_KEY_WRITE</code> is ignored.
<code>ON</code>	MySQL honors any <code>DELAY_KEY_WRITE</code> option specified in <code>CREATE TABLE</code> statements. This is the default value.
<code>ALL</code>	All new opened tables are treated as if they were created with the <code>DELAY_KEY_WRITE</code> option enabled.

If `DELAY_KEY_WRITE` is enabled for a table, the key buffer is not flushed for the table on every index update, but only when the table is closed. This speeds up writes on keys a lot, but if you

use this feature, you should add automatic checking of all `MyISAM` tables by starting the server with the `--myisam-recover-options` [432] option (for example, `--myisam-recover-options=BACKUP,FORCE` [432]). See Section 5.1.3, “Server Command Options”, and Section 14.3.1, “MyISAM Startup Options”.



### Warning

If you enable external locking with `--external-locking` [422], there is no protection against index corruption for tables that use delayed key writes.

- `delayed_insert_limit` [476]

<b>Version Deprecated</b>	5.6.7
<b>Command-Line Format</b>	<code>--delayed_insert_limit=#</code>
<b>Option-File Format</b>	<code>delayed_insert_limit</code>
<b>Option Sets Variable</b>	Yes, <code>delayed_insert_limit</code> [476]
<b>Variable Name</b>	<code>delayed_insert_limit</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
<b>Deprecated</b>	5.6.7
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 100
	<b>Range</b> 1 .. 4294967295
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 100
	<b>Range</b> 1 .. 18446744073709547520

After inserting `delayed_insert_limit` [476] delayed rows into a nontransactional table, the `INSERT DELAYED` handler thread checks whether there are any `SELECT` statements pending. If so, it permits them to execute before continuing to insert delayed rows.

As of MySQL 5.6.7, this system variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `delayed_insert_timeout` [476]

<b>Version Deprecated</b>	5.6.7
<b>Command-Line Format</b>	<code>--delayed_insert_timeout=#</code>
<b>Option-File Format</b>	<code>delayed_insert_timeout</code>
<b>Option Sets Variable</b>	Yes, <code>delayed_insert_timeout</code> [476]
<b>Variable Name</b>	<code>delayed_insert_timeout</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
<b>Deprecated</b>	5.6.7

	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	300	

How many seconds an `INSERT DELAYED` handler thread should wait for `INSERT` statements before terminating.

As of MySQL 5.6.7, this system variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `delayed_queue_size` [477]

<b>Version Deprecated</b>	5.6.7	
<b>Command-Line Format</b>	<code>--delayed_queue_size=#</code>	
<b>Option-File Format</b>	<code>delayed_queue_size</code>	
<b>Option Sets Variable</b>	Yes, <code>delayed_queue_size</code> [477]	
<b>Variable Name</b>	<code>delayed_queue_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.6.7	
	<b>Permitted Values</b>	
<b>Platform Bit Size</b>	32	
<b>Type</b>	numeric	
<b>Default</b>	1000	
<b>Range</b>	1 .. 4294967295	
	<b>Permitted Values</b>	
<b>Platform Bit Size</b>	64	
<b>Type</b>	numeric	
<b>Default</b>	1000	
<b>Range</b>	1 .. 18446744073709547520	

This is a per-table limit on the number of rows to queue when handling `INSERT DELAYED` statements for nontransactional tables. If the queue becomes full, any client that issues an `INSERT DELAYED` statement waits until there is room in the queue again.

As of MySQL 5.6.7, this system variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `disconnect_on_expired_passwords` [477]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--disconnect_on_expired_passwords=#</code>	
<b>Option-File Format</b>	<code>disconnect_on_expired_passwords</code>	
<b>Option Sets Variable</b>	Yes, <code>disconnect_on_expired_passwords</code> [477]	
<b>Variable Name</b>	<code>disconnect_on_expired_passwords</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	No	

	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	ON	

This variable controls how the server handles clients with expired passwords:

- If the client indicates that it can handle expires passwords, the value of `disconnect_on_expired_passwords` [477] is irrelevant. The server permits the client to connect but puts it in sandbox mode.
- If the client does not indicate that it can handle expires passwords, the server handles the client according to the value of `disconnect_on_expired_passwords` [477]:
  - If `disconnect_on_expired_passwords` [477]: is enabled, the server disconnects the client.
  - If `disconnect_on_expired_passwords` [477]: is disabled, the server permits the client to connect but puts it in sandbox mode.

For more information about the interaction of client and server settings relating to expired-password handling, see [Section 6.3.6, “Password Expiration and Sandbox Mode”](#).

This variable was added in MySQL 5.6.10.

- `div_precision_increment` [478]

<b>Command-Line Format</b>	<code>--div_precision_increment=#</code>
<b>Option-File Format</b>	<code>div_precision_increment</code>
<b>Option Sets Variable</b>	Yes, <code>div_precision_increment</code> [478]
<b>Variable Name</b>	<code>div_precision_increment</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	4
<b>Range</b>	0 .. 30

This variable indicates the number of digits by which to increase the scale of the result of division operations performed with the `/` [1200] operator. The default value is 4. The minimum and maximum values are 0 and 30, respectively. The following example illustrates the effect of increasing the default value.

```
mysql> SELECT 1/7;
+-----+
| 1/7   |
+-----+
| 0.1429 |
+-----+
mysql> SET div_precision_increment = 12;
mysql> SELECT 1/7;
+-----+
| 1/7   |
+-----+
| 0.142857142857 |
+-----+
```

- `engine_condition_pushdown` [478]

<b>Version Removed</b>	5.6.1	
<b>Command-Line Format</b>	<code>--engine-condition-pushdown</code>	
<b>Option-File Format</b>	<code>engine-condition-pushdown</code>	
<b>Option Sets Variable</b>	Yes, <code>engine_condition_pushdown</code> [478]	
<b>Variable Name</b>	<code>engine_condition_pushdown</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.5.3, by <code>optimizer_switch</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

This variable was removed in MySQL 5.6.1. Use the `engine_condition_pushdown` flag of the `optimizer_switch` [520] variable instead. See Section 8.8.5.2, “Controlling Switchable Optimizations”.

- `end_markers_in_json` [479]

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<code>end_markers_in_json</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether optimizer JSON output should add end markers.

- `eq_range_index_dive_limit` [479]

<b>Version Introduced</b>	5.6.5	
<b>Command-Line Format</b>	<code>--eq_range_index_dive_limit=#</code>	
<b>Option-File Format</b>	<code>eq_range_index_dive_limit</code>	
<b>Option Sets Variable</b>	Yes, <code>eq_range_index_dive_limit</code> [479]	
<b>Variable Name</b>	<code>eq_range_index_dive_limit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10</code>
	<b>Range</b>	<code>0 .. 4294967295</code>

This variable indicates the number of equality ranges in an equality comparison condition when the optimizer should switch from using index dives to index statistics in estimating the number of qualifying rows. It applies to evaluation of expressions that have either of these equivalent forms, where the optimizer uses a nonunique index to look up `col_name` values:

```
col_name IN(val1, ..., valN)
```

```
col_name = val1 OR ... OR col_name = valN
```

In both cases, the expression contains  $N$  equality ranges. The optimizer can make row estimates using index dives or index statistics. If `eq_range_index_dive_limit` [479] is greater than 0, the optimizer uses existing index statistics instead of index dives if there are `eq_range_index_dive_limit` [479] or more equality ranges. Thus, to permit use of index dives for up to  $N$  equality ranges, set `eq_range_index_dive_limit` [479] to  $N + 1$ . Set `eq_range_index_dive_limit` [479] to 0 to disable use of index statistics and always use index dives regardless of  $N$ .

For more information, see [Section 8.13.1.3, “Equality Range Optimization of Many-Valued Comparisons”](#).

This variable was added in MySQL 5.6.5. Before 5.6.5, the optimizer makes row estimates using index dives in all cases.

To update table index statistics for best estimates, use `ANALYZE TABLE`.

- `error_count` [480]

The number of errors that resulted from the last statement that generated messages. This variable is read only. See [Section 13.7.5.18, “SHOW ERRORS Syntax”](#).

- `event_scheduler` [480]

<b>Command-Line Format</b>	<code>--event-scheduler[=value]</code>	
<b>Option-File Format</b>	<code>event-scheduler</code>	
<b>Option Sets Variable</b>	Yes, <code>event_scheduler</code> [480]	
<b>Variable Name</b>	<code>event_scheduler</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	OFF
	<b>Valid Values</b>	ON OFF DISABLED

This variable indicates the status of the Event Scheduler; possible values are `ON`, `OFF`, and `DISABLED`, with the default being `OFF`. This variable and its effects on the Event Scheduler's operation are discussed in greater detail in the [Overview section of the Events chapter \[2159\]](#).

- `expire_logs_days` [480]

<b>Command-Line Format</b>	<code>--expire_logs_days=#</code>	
<b>Option-File Format</b>	<code>expire_logs_days</code>	
<b>Option Sets Variable</b>	Yes, <code>expire_logs_days</code> [480]	
<b>Variable Name</b>	<code>expire_logs_days</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric

<b>Default</b>	0
<b>Range</b>	0 .. 99

The number of days for automatic binary log file removal. The default is 0, which means “no automatic removal.” Possible removals happen at startup and when the binary log is flushed. Log flushing occurs as indicated in [Section 5.2, “MySQL Server Logs”](#).

To remove binary log files manually, use the `PURGE BINARY LOGS` statement. See [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#).

- `explicit_defaults_for_timestamp` [481]

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--explicit_defaults_for_timestamp=#</code>	
<b>Option-File Format</b>	<code>explicit_defaults_for_timestamp</code>	
<b>Option Sets Variable</b>	Yes, <code>explicit_defaults_for_timestamp</code> [481]	
<b>Variable Name</b>	<code>explicit_defaults_for_timestamp</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

In MySQL, the `TIMESTAMP` data type differs in nonstandard ways from other data types:

- `TIMESTAMP` columns not explicitly declared with the `NULL` attribute are assigned the `NOT NULL` attribute. (Columns of other data types, if not explicitly declared as `NOT NULL`, permit `NULL` values.) Setting such a column to `NULL` sets it to the current timestamp.
- The first `TIMESTAMP` column in a table, if not declared with the `NULL` attribute or an explicit `DEFAULT` or `ON UPDATE` clause, is automatically assigned the `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` attributes.
- `TIMESTAMP` columns following the first one, if not declared with the `NULL` attribute or an explicit `DEFAULT` clause, are automatically assigned `DEFAULT '0000-00-00 00:00:00'` (the “zero” timestamp). For inserted rows that specify no explicit value for such a column, the column is assigned `'0000-00-00 00:00:00'` and no warning occurs.

Those nonstandard behaviors remain the default for `TIMESTAMP` but as of MySQL 5.6.6 are deprecated and this warning appears at startup:

```
[Warning] TIMESTAMP with implicit DEFAULT value is deprecated.
Please use --explicit_defaults_for_timestamp server option (see
documentation for more details).
```

As indicated by the warning, to turn off the nonstandard behaviors, enable the new `explicit_defaults_for_timestamp` [481] system variable at server startup. With this variable enabled, the server handles `TIMESTAMP` as follows instead:

- `TIMESTAMP` columns not explicitly declared as `NOT NULL` permit `NULL` values. Setting such a column to `NULL` sets it to `NULL`, not the current timestamp.
- No `TIMESTAMP` column is assigned the `DEFAULT CURRENT_TIMESTAMP` or `ON UPDATE CURRENT_TIMESTAMP` attributes automatically. Those attributes must be explicitly specified.



- `TIMESTAMP` columns declared as `NOT NULL` and without an explicit `DEFAULT` clause are treated as having no default value. For inserted rows that specify no explicit value for such a column, the result depends on the SQL mode. If strict SQL mode is enabled, an error occurs. If strict SQL mode is not enabled, the column is assigned the implicit default of `'0000-00-00 00:00:00'` and a warning occurs. This is similar to how MySQL treats other temporal types such as `DATETIME`.

This variable was added in MySQL 5.6.6.

- `external_user` [482]

<b>Variable Name</b>	<code>external_user</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The external user name used during the authentication process, as set by the plugin used to authenticate the client. With native (built-in) MySQL authentication, or if the plugin does not set the value, this variable is `NULL`. See [Section 6.3.8, “Proxy Users”](#).

- `flush` [482]

<b>Command-Line Format</b>	<code>--flush</code>	
<b>Option-File Format</b>	<code>flush</code>	
<b>Variable Name</b>	<code>flush</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

If `ON`, the server flushes (synchronizes) all changes to disk after each SQL statement. Normally, MySQL does a write of all changes to disk only after each SQL statement and lets the operating system handle the synchronizing to disk. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#). This variable is set to `ON` if you start `mysqld` with the `--flush` [423] option.

- `flush_time` [482]

<b>Command-Line Format</b>	<code>--flush_time=#</code>	
<b>Option-File Format</b>	<code>flush_time</code>	
<b>Option Sets Variable</b>	Yes, <code>flush_time</code> [482]	
<b>Variable Name</b>	<code>flush_time</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b> (windows)	<code>numeric</code>
	<b>Default</b>	<code>1800</code>

	<b>Min Value</b>	0
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b> (windows)	numeric
	<b>Default</b>	0
	<b>Min Value</b>	0

If this is set to a nonzero value, all tables are closed every `flush_time` [482] seconds to free up resources and synchronize unflushed data to disk. This option is best used only on systems with minimal resources. The default is 0 except that before MySQL 5.6.6, the default is 1800 on Windows.

- `foreign_key_checks` [483]

If set to 1 (the default), `foreign key` constraints for `InnoDB` tables are checked. If set to 0, they are ignored. Typically you leave this setting enabled during normal operation, to enforce `referential integrity`. Disabling foreign key checking can be useful for reloading `InnoDB` tables in an order different from that required by their parent/child relationships. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

Setting `foreign_key_checks` to 0 also affects data definition statements: `DROP SCHEMA` drops a schema even if it contains tables that have foreign keys that are referred to by tables outside the schema, and `DROP TABLE` drops tables that have foreign keys that are referred to by other tables.



#### Note

Setting `foreign_key_checks` to 1 does not trigger a scan of the existing table data. Therefore, rows added to the table while `foreign_key_checks = 0` [483] will not be verified for consistency.

- `ft_boolean_syntax` [483]

<b>Command-Line Format</b>	<code>--ft_boolean_syntax=name</code>	
<b>Option-File Format</b>	<code>ft_boolean_syntax</code>	
<b>Variable Name</b>	<code>ft_boolean_syntax</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string
	<b>Default</b>	<code>+-&gt;&lt;()~*:""&amp;</code>

The list of operators supported by boolean full-text searches performed using `IN BOOLEAN MODE`. See [Section 12.9.2, “Boolean Full-Text Searches”](#).

The default variable value is `'+-><()~*:""&|'`. The rules for changing the value are as follows:

- Operator function is determined by position within the string.
- The replacement value must be 14 characters.
- Each character must be an ASCII nonalphanumeric character.
- Either the first or second character must be a space.

- No duplicates are permitted except the phrase quoting operators in positions 11 and 12. These two characters are not required to be the same, but they are the only two that may be.
- Positions 10, 13, and 14 (which by default are set to “:”, “&”, and “|”) are reserved for future extensions.
- `ft_max_word_len` [484]

<b>Command-Line Format</b>	<code>--ft_max_word_len=#</code>	
<b>Option-File Format</b>	<code>ft_max_word_len</code>	
<b>Option Sets Variable</b>	Yes, <code>ft_max_word_len</code> [484]	
<b>Variable Name</b>	<code>ft_max_word_len</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Min Value</b>	10

The maximum length of the word to be included in a `MyISAM FULLTEXT` index.



**Note**

`FULLTEXT` indexes on `MyISAM` tables must be rebuilt after changing this variable. Use `REPAIR TABLE tbl_name QUICK`.

- `ft_min_word_len` [484]

<b>Command-Line Format</b>	<code>--ft_min_word_len=#</code>	
<b>Option-File Format</b>	<code>ft_min_word_len</code>	
<b>Option Sets Variable</b>	Yes, <code>ft_min_word_len</code> [484]	
<b>Variable Name</b>	<code>ft_min_word_len</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	4
	<b>Min Value</b>	1

The minimum length of the word to be included in a `MyISAM FULLTEXT` index.



**Note**

`FULLTEXT` indexes on `MyISAM` tables must be rebuilt after changing this variable. Use `REPAIR TABLE tbl_name QUICK`.

- `ft_query_expansion_limit` [484]

<b>Command-Line Format</b>	<code>--ft_query_expansion_limit=#</code>
<b>Option-File Format</b>	<code>ft_query_expansion_limit</code>

<b>Option Sets Variable</b>	Yes, <code>ft_query_expansion_limit</code> [484]	
<b>Variable Name</b>	<code>ft_query_expansion_limit</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>20</code>
	<b>Range</b>	<code>0 .. 1000</code>

The number of top matches to use for full-text searches performed using `WITH QUERY EXPANSION`.

- `ft_stopword_file` [485]

<b>Command-Line Format</b>	<code>--ft_stopword_file=file_name</code>	
<b>Option-File Format</b>	<code>ft_stopword_file=file_name</code>	
<b>Option Sets Variable</b>	Yes, <code>ft_stopword_file</code> [485]	
<b>Variable Name</b>	<code>ft_stopword_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The file from which to read the list of stopwords for full-text searches on `MyISAM` tables. The server looks for the file in the data directory unless an absolute path name is given to specify a different directory. All the words from the file are used; comments are *not* honored. By default, a built-in list of stopwords is used (as defined in the `storage/myisam/ft_static.c` file). Setting this variable to the empty string ( `'` ) disables stopwords filtering. See also [Section 12.9.4, “Full-Text Stopwords”](#).



#### Note

`FULLTEXT` indexes on `MyISAM` tables must be rebuilt after changing this variable or the contents of the stopwords file. Use `REPAIR TABLE tbl_name QUICK`.

- `general_log` [485]

<b>Command-Line Format</b>	<code>--general-log</code>	
<b>Option-File Format</b>	<code>general-log</code>	
<b>Option Sets Variable</b>	Yes, <code>general_log</code> [485]	
<b>Variable Name</b>	<code>general_log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether the general query log is enabled. The value can be 0 (or `OFF`) to disable the log or 1 (or `ON`) to enable the log. The default value depends on whether the `--general_log` [485] option is given. The destination for log output is controlled by the `log_output` [498] system variable; if that value is `NONE`, no log entries are written even if the log is enabled.

- [general\\_log\\_file \[486\]](#)

<b>Command-Line Format</b>	<code>--general-log-file=file_name</code>	
<b>Option-File Format</b>	<code>general_log_file</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">general_log_file [486]</a>	
<b>Variable Name</b>	<code>general_log_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>host_name.log</code>

The name of the general query log file. The default value is `host_name.log`, but the initial value can be changed with the `--general_log_file [486]` option.

- [group\\_concat\\_max\\_len \[486\]](#)

<b>Command-Line Format</b>	<code>--group_concat_max_len=#</code>	
<b>Option-File Format</b>	<code>group_concat_max_len</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">group_concat_max_len [486]</a>	
<b>Variable Name</b>	<code>group_concat_max_len</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	1024
	<b>Range</b>	4 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	1024
	<b>Range</b>	4 .. 18446744073709547520

The maximum permitted result length in bytes for the `GROUP_CONCAT()` [1283] function. The default is 1024.

- [have\\_compress \[486\]](#)

`YES` if the `zlib` compression library is available to the server, `NO` if not. If not, the `COMPRESS()` [1261] and `UNCOMPRESS()` [1265] functions cannot be used.

- [have\\_crypt \[486\]](#)

`YES` if the `crypt()` system call is available to the server, `NO` if not. If not, the `ENCRYPT()` [1262] function cannot be used.

- [have\\_csv \[486\]](#)

YES if `mysqld` supports CSV tables, NO if not.

This variable was removed in MySQL 5.6.1. Use `SHOW ENGINES` instead.

- `have_dynamic_loading` [487]

YES if `mysqld` supports dynamic loading of plugins, NO if not.

- `have_geometry` [487]

YES if the server supports spatial data types, NO if not.

- `have_innodb` [487]

YES if `mysqld` supports InnoDB tables. DISABLED if `--skip-innodb` [1698] is used.

This variable was removed in MySQL 5.6.1. Use `SHOW ENGINES` instead.

- `have_openssl` [487]

This variable is an alias for `have_ssl` [487].

- `have_partitioning` [487]

YES if `mysqld` supports partitioning.

This variable was removed in MySQL 5.6.1. Use `SHOW ENGINES` instead.

- `have_profiling` [487]

YES if statement profiling capability is present, NO if not. If present, the `profiling` system variable controls whether this capability is enabled or disabled. See [Section 13.7.5.32, "SHOW PROFILES Syntax"](#).

This variable is deprecated in MySQL 5.6.8 and will be removed in a future MySQL release.

- `have_query_cache` [487]

YES if `mysqld` supports the query cache, NO if not.

- `have_rtree_keys` [487]

YES if RTREE indexes are available, NO if not. (These are used for spatial indexes in MyISAM tables.)

- `have_ssl` [487]

YES if `mysqld` supports SSL connections, NO if not. DISABLED indicates that the server was compiled with SSL support, but was not started with the appropriate `--ssl-xxx` options. For more information, see [Section 6.3.9.2, "Configuring MySQL for SSL"](#).

- `have_symlink` [487]

YES if symbolic link support is enabled, NO if not. This is required on Unix for support of the `DATA DIRECTORY` and `INDEX DIRECTORY` table options, and on Windows for support of data directory symlinks. If the server is started with the `--skip-symbolic-links` [441] option, the value is DISABLED.

- `host_cache_size` [487]

<b>Version Introduced</b>	5.6.5
<b>Command-Line Format</b>	<code>--host_cache_size=#</code>
<b>Option-File Format</b>	<code>host_cache_size</code>

<b>Option Sets Variable</b>	Yes, <a href="#">host_cache_size</a> [487]	
<b>Variable Name</b>	<a href="#">host_cache_size</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.7)</b>	
	<b>Type</b>	<a href="#">numeric</a>
	<b>Default</b>	128
	<b>Range</b>	0 .. 65536
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	<a href="#">numeric</a>
	<b>Default</b>	-1 (autosized)
	<b>Range</b>	0 .. 65536

The size of the internal host cache (see [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#)). Setting the size to 0 disables the host cache. Changing the cache size at runtime implicitly causes a `FLUSH HOSTS` operation to clear the host cache and truncate the `host_cache` table.

The default value is 128, plus 1 for a value of `max_connections` [503] up to 500, plus 1 for every increment of 20 over 500 in the `max_connections` [503] value, capped to a limit of 2000. Before MySQL 5.6.8, the default is 128.

Use of `--skip-host-cache` [439] is similar to setting the `host_cache_size` [487] system variable to 0, but `host_cache_size` [487] is more flexible because it can also be used to resize, enable, or disable the host cache at runtime, not just at server startup.

If you start the server with `--skip-host-cache` [439], that does not prevent changes to the value of `host_cache_size` [487], but such changes have no effect and the cache is not re-enabled even if `host_cache_size` [487] is set larger than 0.

This variable was added in MySQL 5.6.5.

- [hostname](#) [488]

<b>Variable Name</b>	<a href="#">hostname</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">string</a>

The server sets this variable to the server host name at startup.

- [identity](#) [488]

This variable is a synonym for the `last_insert_id` [495] variable. It exists for compatibility with other database systems. You can read its value with `SELECT @@identity`, and set it using `SET identity`.

- [ignore\\_db\\_dirs](#) [488]

<b>Version Introduced</b>	5.6.3
<b>Variable Name</b>	<a href="#">ignore_db_dirs</a>
<b>Variable Scope</b>	Global

<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

A comma-separated list of names that are not considered as database directories in the data directory. The value is set from any instances of `--ignore-db-dir` [424] given at server startup.

This variable was added in MySQL 5.6.3.

- [init\\_connect](#) [489]

<b>Command-Line Format</b>	<code>--init-connect=name</code>	
<b>Option-File Format</b>	<code>init_connect</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">init_connect</a> [489]	
<b>Variable Name</b>	<code>init_connect</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

A string to be executed by the server for each client that connects. The string consists of one or more SQL statements, separated by semicolon characters. For example, each client session begins by default with autocommit mode enabled. For older servers (before MySQL 5.5.8), there is no global [autocommit](#) [464] system variable to specify that autocommit should be disabled by default, but as a workaround [init\\_connect](#) [489] can be used to achieve the same effect:

```
SET GLOBAL init_connect='SET autocommit=0';
```

The [init\\_connect](#) [489] variable can also be set on the command line or in an option file. To set the variable as just shown using an option file, include these lines:

```
[mysqld]
init_connect='SET autocommit=0'
```

The content of [init\\_connect](#) [489] is not executed for users that have the [SUPER](#) [756] privilege. This is done so that an erroneous value for [init\\_connect](#) [489] does not prevent all clients from connecting. For example, the value might contain a statement that has a syntax error, thus causing client connections to fail. Not executing [init\\_connect](#) [489] for users that have the [SUPER](#) [756] privilege enables them to open a connection and fix the [init\\_connect](#) [489] value.

- [init\\_file](#) [489]

<b>Command-Line Format</b>	<code>--init-file=file_name</code>	
<b>Option-File Format</b>	<code>init-file=file_name</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">init_file</a> [489]	
<b>Variable Name</b>	<code>init_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name



The name of the file specified with the `--init-file` [424] option when you start the server. This should be a file containing SQL statements that you want the server to execute when it starts. Each statement must be on a single line and should not include comments. No statement terminator such as `;`, `\g`, or `\G` should be given at the end of each statement.

- `innodb_xxx`

InnoDB system variables are listed in [Section 14.2.6, “InnoDB Startup Options and System Variables”](#). These variables control many aspects of storage, memory use, and I/O patterns for InnoDB tables, and are especially important now that [InnoDB is the default storage engine](#).

- `insert_id` [490]

The value to be used by the following `INSERT` or `ALTER TABLE` statement when inserting an `AUTO_INCREMENT` value. This is mainly used with the binary log.

- `interactive_timeout` [490]

<b>Command-Line Format</b>	<code>--interactive_timeout=#</code>
<b>Option-File Format</b>	<code>interactive_timeout</code>
<b>Option Sets Variable</b>	Yes, <code>interactive_timeout</code> [490]
<b>Variable Name</b>	<code>interactive_timeout</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> numeric
	<b>Default</b> 28800
	<b>Min Value</b> 1

The number of seconds the server waits for activity on an interactive connection before closing it. An interactive client is defined as a client that uses the `CLIENT_INTERACTIVE` option to `mysql_real_connect()`. See also `wait_timeout` [563].

- `join_buffer_size` [490]

<b>Command-Line Format</b>	<code>--join_buffer_size=#</code>
<b>Option-File Format</b>	<code>join_buffer_size</code>
<b>Option Sets Variable</b>	Yes, <code>join_buffer_size</code> [490]
<b>Variable Name</b>	<code>join_buffer_size</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&gt;= 5.6.0, &lt;= 5.6.5)</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> numeric
	<b>Default</b> 131072
	<b>Range</b> 128 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.0, &lt;= 5.6.5)</b>

	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	131072
	<b>Range</b>	128 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	262144
	<b>Range</b>	128 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	262144
	<b>Range</b>	128 .. 18446744073709547520

The minimum size of the buffer that is used for plain index scans, range index scans, and joins that do not use indexes and thus perform full table scans. Normally, the best way to get fast joins is to add indexes. Increase the value of `join_buffer_size` [490] to get a faster full join when adding indexes is not possible. One join buffer is allocated for each full join between two tables. For a complex join between several tables for which indexes are not used, multiple join buffers might be necessary. There is no gain from setting the buffer larger than required to hold each matching row, and all joins allocate at least the minimum size, so use caution in setting this variable to a large value globally. It is better to keep the global setting small and change to a larger setting only in sessions that are doing large joins. Memory allocation time can cause substantial performance drops if the global size is larger than needed by most queries that use it.

The default is 256KB as of MySQL 5.6.6, 128KB before that. The maximum permissible setting for `join_buffer_size` [490] is 4GB. Values larger than 4GB are permitted for 64-bit platforms (except 64-bit Windows, for which large values are truncated to 4GB with a warning).

For additional information about join buffering, see [Section 8.13.8, “Nested-Loop Join Algorithms”](#).

- `keep_files_on_create` [491]

<b>Command-Line Format</b>	<code>--keep_files_on_create=#</code>	
<b>Option-File Format</b>	<code>keep_files_on_create</code>	
<b>Option Sets Variable</b>	Yes, <code>keep_files_on_create</code> [491]	
<b>Variable Name</b>	<code>keep_files_on_create</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

If a `MyISAM` table is created with no `DATA DIRECTORY` option, the `.MYD` file is created in the database directory. By default, if `MyISAM` finds an existing `.MYD` file in this case, it overwrites it. The same applies to `.MYI` files for tables created with no `INDEX DIRECTORY` option. To suppress this

behavior, set the `keep_files_on_create` [491] variable to `ON` (1), in which case `MyISAM` will not overwrite existing files and returns an error instead. The default value is `OFF` (0).

If a `MyISAM` table is created with a `DATA DIRECTORY` or `INDEX DIRECTORY` option and an existing `.MYD` or `.MYI` file is found, `MyISAM` always returns an error. It will not overwrite a file in the specified directory.

- `key_buffer_size` [492]

<b>Command-Line Format</b>	<code>--key_buffer_size=#</code>
<b>Option-File Format</b>	<code>key_buffer_size</code>
<b>Option Sets Variable</b>	Yes, <code>key_buffer_size</code> [492]
<b>Variable Name</b>	<code>key_buffer_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>8388608</code>
	<b>Range</b> <code>8 .. 4294967295</code>

Index blocks for `MyISAM` tables are buffered and are shared by all threads.

`key_buffer_size` [492] is the size of the buffer used for index blocks. The key buffer is also known as the key cache.

The maximum permissible setting for `key_buffer_size` [492] is 4GB on 32-bit platforms. Values larger than 4GB are permitted for 64-bit platforms. The effective maximum size might be less, depending on your available physical RAM and per-process RAM limits imposed by your operating system or hardware platform. The value of this variable indicates the amount of memory requested. Internally, the server allocates as much memory as possible up to this amount, but the actual allocation might be less.

You can increase the value to get better index handling for all reads and multiple writes; on a system whose primary function is to run MySQL using the `MyISAM` storage engine, 25% of the machine's total memory is an acceptable value for this variable. However, you should be aware that, if you make the value too large (for example, more than 50% of the machine's total memory), your system might start to page and become extremely slow. This is because MySQL relies on the operating system to perform file system caching for data reads, so you must leave some room for the file system cache. You should also consider the memory requirements of any other storage engines that you may be using in addition to `MyISAM`.

For even more speed when writing many rows at the same time, use `LOCK TABLES`. See [Section 8.2.2.1, "Speed of INSERT Statements"](#).

You can check the performance of the key buffer by issuing a `SHOW STATUS` statement and examining the `Key_read_requests` [594], `Key_reads` [594], `Key_write_requests` [594], and `Key_writes` [594] status variables. (See [Section 13.7.5, "SHOW Syntax"](#).) The `Key_reads/Key_read_requests` ratio should normally be less than 0.01. The `Key_writes/Key_write_requests` ratio is usually near 1 if you are using mostly updates and deletes, but might be much smaller if you tend to do updates that affect many rows at the same time or if you are using the `DELAY_KEY_WRITE` table option.

The fraction of the key buffer in use can be determined using `key_buffer_size` [492] in conjunction with the `Key_blocks_unused` [594] status variable and the buffer block size, which is available from the `key_cache_block_size` [493] system variable:

```
1 - ((Key_blocks_unused * key_cache_block_size) / key_buffer_size)
```

This value is an approximation because some space in the key buffer is allocated internally for administrative structures. Factors that influence the amount of overhead for these structures include block size and pointer size. As block size increases, the percentage of the key buffer lost to overhead tends to decrease. Larger blocks results in a smaller number of read operations (because more keys are obtained per read), but conversely an increase in reads of keys that are not examined (if not all keys in a block are relevant to a query).

It is possible to create multiple `MyISAM` key caches. The size limit of 4GB applies to each cache individually, not as a group. See [Section 8.9.2, “The `MyISAM` Key Cache”](#).

- `key_cache_age_threshold` [493]

<b>Command-Line Format</b>	<code>--key_cache_age_threshold=#</code>
<b>Option-File Format</b>	<code>key_cache_age_threshold</code>
<b>Option Sets Variable</b>	Yes, <code>key_cache_age_threshold</code> [493]
<b>Variable Name</b>	<code>key_cache_age_threshold</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> numeric
	<b>Default</b> 300
	<b>Range</b> 100 .. 4294967295
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> numeric
	<b>Default</b> 300
	<b>Range</b> 100 .. 18446744073709547520

This value controls the demotion of buffers from the hot sublist of a key cache to the warm sublist. Lower values cause demotion to happen more quickly. The minimum value is 100. The default value is 300. See [Section 8.9.2, “The `MyISAM` Key Cache”](#).

- `key_cache_block_size` [493]

<b>Command-Line Format</b>	<code>--key_cache_block_size=#</code>
<b>Option-File Format</b>	<code>key_cache_block_size</code>
<b>Option Sets Variable</b>	Yes, <code>key_cache_block_size</code> [493]
<b>Variable Name</b>	<code>key_cache_block_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>

	<b>Type</b>	numeric
	<b>Default</b>	1024
	<b>Range</b>	512 .. 16384

The size in bytes of blocks in the key cache. The default value is 1024. See [Section 8.9.2, “The MyISAM Key Cache”](#).

- [key\\_cache\\_division\\_limit \[494\]](#)

<b>Command-Line Format</b>	<code>--key_cache_division_limit=#</code>	
<b>Option-File Format</b>	<code>key_cache_division_limit</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">key_cache_division_limit [494]</a>	
<b>Variable Name</b>	<code>key_cache_division_limit</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	100
	<b>Range</b>	1 .. 100

The division point between the hot and warm sublists of the key cache buffer list. The value is the percentage of the buffer list to use for the warm sublist. Permissible values range from 1 to 100. The default value is 100. See [Section 8.9.2, “The MyISAM Key Cache”](#).

- [large\\_files\\_support \[494\]](#)

<b>Variable Name</b>	<code>large_files_support</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

Whether `mysqld` was compiled with options for large file support.

- [large\\_pages \[494\]](#)

<b>Command-Line Format</b>	<code>--large-pages</code>	
<b>Option-File Format</b>	<code>large-pages</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">large_pages [494]</a>	
<b>Variable Name</b>	<code>large_pages</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Platform Specific</b>	linux	
	<b>Permitted Values</b>	
	<b>Type</b> (linux)	boolean
	<b>Default</b>	FALSE

Whether large page support is enabled (via the `--large-pages [425]` option). See [Section 8.11.4.2, “Enabling Large Page Support”](#).

- [large\\_page\\_size \[494\]](#)

<b>Variable Name</b>	<code>large_page_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type (linux)</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

If large page support is enabled, this shows the size of memory pages. Currently, large memory pages are supported only on Linux; on other platforms, the value of this variable is always 0. See [Section 8.11.4.2, “Enabling Large Page Support”](#).

- `last_insert_id` [\[495\]](#)

The value to be returned from `LAST_INSERT_ID()` [\[1269\]](#). This is stored in the binary log when you use `LAST_INSERT_ID()` [\[1269\]](#) in a statement that updates a table. Setting this variable does not update the value returned by the `mysql_insert_id()` C API function.

- `lc_messages` [\[495\]](#)

<b>Command-Line Format</b>	<code>--lc-messages=name</code>	
<b>Option-File Format</b>	<code>lc-messages</code>	
<b>Option Sets Variable</b>	Yes, <code>lc_messages</code> <a href="#">[495]</a>	
<b>Variable Name</b>	<code>lc-messages</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The locale to use for error messages. The server converts the value to a language name and combines it with the value of the `lc_messages_dir` [\[495\]](#) to produce the location for the error message file. See [Section 10.2, “Setting the Error Message Language”](#).

- `lc_messages_dir` [\[495\]](#)

<b>Command-Line Format</b>	<code>--lc-messages-dir=path</code>	
<b>Option-File Format</b>	<code>lc-messages-dir</code>	
<b>Option Sets Variable</b>	Yes, <code>lc_messages_dir</code> <a href="#">[495]</a>	
<b>Variable Name</b>	<code>lc-messages-dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The directory where error messages are located. The value is used together with the value of `lc_messages` [\[495\]](#) to produce the location for the error message file. See [Section 10.2, “Setting the Error Message Language”](#).

- `lc_time_names` [\[495\]](#)

<b>Variable Name</b>	<code>lc_time_names</code>
----------------------	----------------------------

<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

This variable specifies the locale that controls the language used to display day and month names and abbreviations. This variable affects the output from the `DATE_FORMAT()` [1215], `DAYNAME()` [1217] and `MONTHNAME()` [1220] functions. Locale names are POSIX-style values such as 'ja\_JP' or 'pt\_BR'. The default value is 'en\_US' regardless of your system's locale setting. For further information, see [Section 10.7, “MySQL Server Locale Support”](#).

- [license](#) [496]

<b>Variable Name</b>	license	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string
	<b>Default</b>	GPL

The type of license the server has.

- [local\\_infile](#) [496]

<b>Variable Name</b>	local_infile	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean

Whether `LOCAL` is supported for `LOAD DATA INFILE` statements. If this variable is disabled, clients cannot use `LOCAL` in `LOAD DATA` statements. See [Section 6.1.6, “Security Issues with LOAD DATA LOCAL”](#).

- [lock\\_wait\\_timeout](#) [496]

<b>Command-Line Format</b>	--lock_wait_timeout=#	
<b>Option-File Format</b>	lock_wait_timeout	
<b>Option Sets Variable</b>	Yes, <a href="#">lock_wait_timeout</a> [496]	
<b>Variable Name</b>	lock_wait_timeout	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	31536000
	<b>Range</b>	1 .. 31536000

This variable specifies the timeout in seconds for attempts to acquire metadata locks. The permissible values range from 1 to 31536000 (1 year). The default is 31536000.

This timeout applies to all statements that use metadata locks. These include DML and DDL operations on tables, views, stored procedures, and stored functions, as well as `LOCK TABLES`, `FLUSH TABLES WITH READ LOCK`, and `HANDLER` statements.

This timeout does not apply to implicit accesses to system tables in the `mysql` database, such as grant tables modified by `GRANT` or `REVOKE` statements or table logging statements. The timeout does apply to system tables accessed directly, such as with `SELECT` or `UPDATE`.

The timeout value applies separately for each metadata lock attempt. A given statement can require more than one lock, so it is possible for the statement to block for longer than the `lock_wait_timeout` [496] value before reporting a timeout error. When lock timeout occurs, `ER_LOCK_WAIT_TIMEOUT` [3674] is reported.

`lock_wait_timeout` [496] does not apply to delayed inserts, which always execute with a timeout of 1 year. This is done to avoid unnecessary timeouts because a session that issues a delayed insert receives no notification of delayed insert timeouts.

- `locked_in_memory` [497]

<b>Variable Name</b>	<code>locked_in_memory</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

Whether `mysqld` was locked in memory with `--memlock` [431].

- `log` [497]

This variable removed in MySQL 5.6.1. Use `general_log` [485] instead.

- `log_bin_trust_function_creators` [497]

<b>Command-Line Format</b>	<code>--log-bin-trust-function-creators</code>	
<b>Option-File Format</b>	<code>log-bin-trust-function-creators</code>	
<b>Option Sets Variable</b>	Yes, <code>log_bin_trust_function_creators</code> [497]	
<b>Variable Name</b>	<code>log_bin_trust_function_creators</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

This variable applies when binary logging is enabled. It controls whether stored function creators can be trusted not to create stored functions that will cause unsafe events to be written to the binary log. If set to 0 (the default), users are not permitted to create or alter stored functions unless they have the `SUPER` [756] privilege in addition to the `CREATE ROUTINE` [754] or `ALTER ROUTINE` [754] privilege. A setting of 0 also enforces the restriction that a function must be declared with the `DETERMINISTIC` characteristic, or with the `READS SQL DATA` or `NO SQL` characteristic. If the variable is set to 1, MySQL does not enforce these restrictions on stored function creation. This variable also applies to trigger creation. See [Section 18.7, “Binary Logging of Stored Programs”](#).

- `log_error` [497]

<b>Command-Line Format</b>	<code>--log-error[=name]</code>
<b>Option-File Format</b>	<code>log-error</code>



<b>Option Sets Variable</b>	Yes, <code>log_error</code> [497]	
<b>Variable Name</b>	<code>log_error</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The location of the error log.

- `log_output` [498]

<b>Command-Line Format</b>	<code>--log-output=name</code>	
<b>Option-File Format</b>	<code>log-output</code>	
<b>Option Sets Variable</b>	Yes, <code>log_output</code> [498]	
<b>Variable Name</b>	<code>log_output</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>set</code>
	<b>Default</b>	<code>FILE</code>
	<b>Valid Values</b>	<code>TABLE</code> <code>FILE</code> <code>NONE</code>

The destination for general query log and slow query log output. The value can be a comma-separated list of one or more of the words `TABLE` (log to tables), `FILE` (log to files), or `NONE` (do not log to tables or files). The default value is `FILE`. `NONE`, if present, takes precedence over any other specifiers. If the value is `NONE` log entries are not written even if the logs are enabled. If the logs are not enabled, no logging occurs even if the value of `log_output` [498] is not `NONE`. For more information, see [Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”](#).

- `log_queries_not_using_indexes` [498]

<b>Command-Line Format</b>	<code>--log-queries-not-using-indexes</code>	
<b>Option-File Format</b>	<code>log-queries-not-using-indexes</code>	
<b>Option Sets Variable</b>	Yes, <code>log_queries_not_using_indexes</code> [498]	
<b>Variable Name</b>	<code>log_queries_not_using_indexes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether queries that do not use indexes are logged to the slow query log. See [Section 5.2.5, “The Slow Query Log”](#).

- `log_slow_admin_statements` [498]

<b>Version Introduced</b>	5.6.11	
<b>Variable Name</b>	<code>log_slow_admin_statements</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Log slow administrative statements such as `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `ALTER TABLE` to the slow query log. This variable was added in MySQL 5.6.11.

- [log\\_slow\\_queries \[499\]](#)

<b>Version Removed</b>	5.6.1	
<b>Command-Line Format</b>	<code>--log-slow-queries[=name]</code>	
<b>Option-File Format</b>	<code>log-slow-queries</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">log_slow_queries [499]</a>	
<b>Variable Name</b>	<code>log_slow_queries</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.1.29, by <code>slow-query-log</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>

This variable was removed in MySQL 5.6.1. Use [slow\\_query\\_log \[541\]](#) instead.

- [log\\_throttle\\_queries\\_not\\_using\\_indexes \[499\]](#)

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<code>log_throttle_queries_not_using_indexes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

If [log\\_queries\\_not\\_using\\_indexes \[498\]](#) is enabled, the [log\\_throttle\\_queries\\_not\\_using\\_indexes \[499\]](#) variable limits the number of such queries per minute that can be written to the slow query log. A value of 0 (the default) means “no limit”. For more information, see [Section 5.2.5, “The Slow Query Log”](#).

This variable was added in MySQL 5.6.5.

- [log\\_warnings \[499\]](#)

<b>Command-Line Format</b>	<code>--log-warnings[=#]</code>
	<code>-W [#]</code>
<b>Option-File Format</b>	<code>log-warnings</code>

<b>Option Sets Variable</b>	Yes, <a href="#">log_warnings</a> [499]
<b>Variable Name</b>	<code>log_warnings</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
<b>Disabled by</b>	<code>skip-log-warnings</code>
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	1
<b>Range</b>	0 .. 18446744073709547520

Whether to produce additional warning messages to the error log. It is enabled (1) by default and can be disabled by setting it to 0. Aborted connections and access-denied errors for new connection attempts are logged if the value is greater than 1. The server logs messages about statements that are unsafe for statement-based logging only if the value is greater than 0.

- [long\\_query\\_time](#) [500]

<b>Command-Line Format</b>	<code>--long_query_time=#</code>
<b>Option-File Format</b>	<code>long_query_time</code>
<b>Option Sets Variable</b>	Yes, <a href="#">long_query_time</a> [500]
<b>Variable Name</b>	<code>long_query_time</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	10
<b>Min Value</b>	0

If a query takes longer than this many seconds, the server increments the [Slow\\_queries](#) [599] status variable. If the slow query log is enabled, the query is logged to the slow query log file. This value is measured in real time, not CPU time, so a query that is under the threshold on a lightly loaded system might be above the threshold on a heavily loaded one. The minimum and default values of [long\\_query\\_time](#) [500] are 0 and 10, respectively. The value can be specified to a resolution of microseconds. For logging to a file, times are written including the microseconds part. For logging to tables, only integer times are written; the microseconds part is ignored. See [Section 5.2.5, “The Slow Query Log”](#).

- [low\\_priority\\_updates](#) [500]

<b>Command-Line Format</b>	<code>--low-priority-updates</code>
<b>Option-File Format</b>	<code>low-priority-updates</code>
<b>Option Sets Variable</b>	Yes, <a href="#">low_priority_updates</a> [500]
<b>Variable Name</b>	<code>low_priority_updates</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>

	<b>Type</b>	boolean
	<b>Default</b>	FALSE

If set to 1, all `INSERT`, `UPDATE`, `DELETE`, and `LOCK TABLE WRITE` statements wait until there is no pending `SELECT` or `LOCK TABLE READ` on the affected table. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).

- `lower_case_file_system` [501]

<b>Command-Line Format</b>	<code>--lower_case_file_system[=#]</code>	
<b>Option-File Format</b>	<code>lower_case_file_system</code>	
<b>Option Sets Variable</b>	Yes, <code>lower_case_file_system</code> [501]	
<b>Variable Name</b>	<code>lower_case_file_system</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean

This variable describes the case sensitivity of file names on the file system where the data directory is located. `OFF` means file names are case sensitive, `ON` means they are not case sensitive. This variable is read only because it reflects a file system attribute and setting it would have no effect on the file system.

- `lower_case_table_names` [501]

<b>Command-Line Format</b>	<code>--lower_case_table_names[=#]</code>	
<b>Option-File Format</b>	<code>lower_case_table_names</code>	
<b>Option Sets Variable</b>	Yes, <code>lower_case_table_names</code> [501]	
<b>Variable Name</b>	<code>lower_case_table_names</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 2

If set to 0, table names are stored as specified and comparisons are case sensitive. If set to 1, table names are stored in lowercase on disk and comparisons are not case sensitive. If set to 2, table names are stored as given but compared in lowercase. This option also applies to database names and table aliases. For additional information, see [Section 9.2.2, “Identifier Case Sensitivity”](#).

You should *not* set this variable to 0 if you are running MySQL on a system that has case-insensitive file names (such as Windows or Mac OS X). If you set this variable to 0 on such a system and access `MyISAM` table names using different lettercases, index corruption may result. On Windows the default value is 1. On Mac OS X, the default value is 2.

If you are using `InnoDB` tables, you should set this variable to 1 on all platforms to force names to be converted to lowercase.

The setting of this variable in MySQL 5.6 affects the behavior of replication filtering options with regard to case sensitivity. This is a change from previous versions of MySQL. (Bug #51639) See [Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#), for more information.

In previous versions of MySQL, using different settings for `lower_case_table_names` on replication masters and slaves could cause replication to fail when the slave used a case-sensitive file system. This issue is resolved in MySQL 5.6.1. For more information, see [Section 16.4.1.32, “Replication and Variables”](#).

- `max_allowed_packet` [\[502\]](#)

<b>Command-Line Format</b>	<code>--max_allowed_packet=#</code>	
<b>Option-File Format</b>	<code>max_allowed_packet</code>	
<b>Option Sets Variable</b>	Yes, <code>max_allowed_packet</code> <a href="#">[502]</a>	
<b>Variable Name</b>	<code>max_allowed_packet</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1048576</code>
	<b>Range</b>	<code>1024 .. 1073741824</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4194304</code>
	<b>Range</b>	<code>1024 .. 1073741824</code>

The maximum size of one packet or any generated/intermediate string, or any parameter sent by the `mysql_stmt_send_long_data()` C API function. The default is 4MB as of MySQL 5.6.6, 1MB before that.

The packet message buffer is initialized to `net_buffer_length` [\[515\]](#) bytes, but can grow up to `max_allowed_packet` [\[502\]](#) bytes when needed. This value by default is small, to catch large (possibly incorrect) packets.

You must increase this value if you are using large `BLOB` columns or long strings. It should be as big as the largest `BLOB` you want to use. The protocol limit for `max_allowed_packet` [\[502\]](#) is 1GB. The value should be a multiple of 1024; nonmultiples are rounded down to the nearest multiple.

When you change the message buffer size by changing the value of the `max_allowed_packet` [\[502\]](#) variable, you should also change the buffer size on the client side if your client program permits it. On the client side, `max_allowed_packet` [\[502\]](#) has a default of 1GB. Some programs such as `mysql` and `mysqldump` enable you to change the client-side value by setting `max_allowed_packet` [\[502\]](#) on the command line or in an option file.

The session value of this variable is read only.

- `max_connect_errors` [\[502\]](#)

<b>Command-Line Format</b>	<code>--max_connect_errors=#</code>	
<b>Option-File Format</b>	<code>max_connect_errors</code>	
<b>Option Sets Variable</b>	Yes, <code>max_connect_errors</code> <a href="#">[502]</a>	
<b>Variable Name</b>	<code>max_connect_errors</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	

	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	10
	<b>Range</b>	1 .. 4294967295
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	10
	<b>Range</b>	1 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	100
	<b>Range</b>	1 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	100
	<b>Range</b>	1 .. 18446744073709547520

If more than this many successive connection requests from a host are interrupted without a successful connection, the server blocks that host from further connections. You can unblock blocked hosts by flushing the host cache. To do so, issue a `FLUSH HOSTS` statement or execute a `mysqladmin flush-hosts` command. If a connection is established successfully within fewer than `max_connect_errors` [502] attempts after a previous connection was interrupted, the error count for the host is cleared to zero. However, once a host is blocked, flushing the host cache is the only way to unblock it. The default is 100 as of MySQL 5.6.6, 10 before that.

- `max_connections` [503]

<b>Command-Line Format</b>	<code>--max_connections=#</code>	
<b>Option-File Format</b>	<code>max_connections</code>	
<b>Option Sets Variable</b>	Yes, <code>max_connections</code> [503]	
<b>Variable Name</b>	<code>max_connections</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	151
	<b>Range</b>	1 .. 100000

The maximum permitted number of simultaneous client connections. By default, this is 151. See [Section C.5.2.7, “Too many connections”](#), for more information.

Increasing this value increases the number of file descriptors that `mysqld` requires. See [Section 8.4.3.1, “How MySQL Opens and Closes Tables”](#), for comments on file descriptor limits.

Connections refused because the `max_connections` [503] limit is reached increment the `Connection_errors_max_connections` [587] status variable.

- `max_delayed_threads` [504]

<b>Version Deprecated</b>	5.6.7	
<b>Command-Line Format</b>	<code>--max_delayed_threads=#</code>	
<b>Option-File Format</b>	<code>max_delayed_threads</code>	
<b>Option Sets Variable</b>	Yes, <code>max_delayed_threads</code> [504]	
<b>Variable Name</b>	<code>max_delayed_threads</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.6.7	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>20</code>
	<b>Range</b>	<code>0 .. 16384</code>

Do not start more than this number of threads to handle `INSERT DELAYED` statements for nontransactional tables. If you try to insert data into a new table after all `INSERT DELAYED` threads are in use, the row is inserted as if the `DELAYED` attribute was not specified. If you set this to 0, MySQL never creates a thread to handle `DELAYED` rows; in effect, this disables `DELAYED` entirely.

For the `SESSION` value of this variable, the only valid values are 0 or the `GLOBAL` value.

As of MySQL 5.6.7, this system variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `max_error_count` [504]

<b>Command-Line Format</b>	<code>--max_error_count=#</code>	
<b>Option-File Format</b>	<code>max_error_count</code>	
<b>Option Sets Variable</b>	Yes, <code>max_error_count</code> [504]	
<b>Variable Name</b>	<code>max_error_count</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>64</code>
	<b>Range</b>	<code>0 .. 65535</code>

The maximum number of error, warning, and note messages to be stored for display by the `SHOW ERRORS` and `SHOW WARNINGS` statements. This is the same as the number of condition areas in the diagnostics area, and thus the number of conditions that can be inspected by `GET DIAGNOSTICS`.

- `max_heap_table_size` [504]

<b>Command-Line Format</b>	<code>--max_heap_table_size=#</code>
----------------------------	--------------------------------------

<b>Option-File Format</b>	<code>max_heap_table_size</code>	
<b>Option Sets Variable</b>	Yes, <code>max_heap_table_size</code> [504]	
<b>Variable Name</b>	<code>max_heap_table_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	16777216
	<b>Range</b>	16384 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	16777216
	<b>Range</b>	16384 .. 1844674407370954752

This variable sets the maximum size to which user-created `MEMORY` tables are permitted to grow. The value of the variable is used to calculate `MEMORY` table `MAX_ROWS` values. Setting this variable has no effect on any existing `MEMORY` table, unless the table is re-created with a statement such as `CREATE TABLE` or altered with `ALTER TABLE` or `TRUNCATE TABLE`. A server restart also sets the maximum size of existing `MEMORY` tables to the global `max_heap_table_size` [504] value.

This variable is also used in conjunction with `tmp_table_size` [558] to limit the size of internal in-memory tables. See Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”.

`max_heap_table_size` is not replicated. See Section 16.4.1.20, “Replication and `MEMORY` Tables”, and Section 16.4.1.32, “Replication and Variables”, for more information.

- `max_insert_delayed_threads` [505]

<b>Version Deprecated</b>	5.6.7	
<b>Variable Name</b>	<code>max_insert_delayed_threads</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.6.7	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric

This variable is a synonym for `max_delayed_threads` [504].

As of MySQL 5.6.7, this system variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `max_join_size` [505]

<b>Command-Line Format</b>	<code>--max_join_size=#</code>
<b>Option-File Format</b>	<code>max_join_size</code>
<b>Option Sets Variable</b>	Yes, <code>max_join_size</code> [505]



<b>Variable Name</b>	max_join_size	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	18446744073709551615
	<b>Range</b>	1 .. 18446744073709551615

Do not permit statements that probably need to examine more than `max_join_size` [505] rows (for single-table statements) or row combinations (for multiple-table statements) or that are likely to do more than `max_join_size` [505] disk seeks. By setting this value, you can catch statements where keys are not used properly and that would probably take a long time. Set it if your users tend to perform joins that lack a `WHERE` clause, that take a long time, or that return millions of rows.

Setting this variable to a value other than `DEFAULT` resets the value of `sql_big_selects` [543] to 0. If you set the `sql_big_selects` [543] value again, the `max_join_size` [505] variable is ignored.

If a query result is in the query cache, no result size check is performed, because the result has previously been computed and it does not burden the server to send it to the client.

- `max_length_for_sort_data` [506]

<b>Command-Line Format</b>	--max_length_for_sort_data=#	
<b>Option-File Format</b>	max_length_for_sort_data	
<b>Option Sets Variable</b>	Yes, <code>max_length_for_sort_data</code> [506]	
<b>Variable Name</b>	max_length_for_sort_data	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1024
	<b>Range</b>	4 .. 8388608

The cutoff on the size of index values that determines which `filesort` algorithm to use. See [Section 8.13.13, “ORDER BY Optimization”](#).

- `max_prepared_stmt_count` [506]

<b>Command-Line Format</b>	--max_prepared_stmt_count=#	
<b>Option-File Format</b>	max_prepared_stmt_count	
<b>Option Sets Variable</b>	Yes, <code>max_prepared_stmt_count</code> [506]	
<b>Variable Name</b>	max_prepared_stmt_count	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	16382
	<b>Range</b>	0 .. 1048576

This variable limits the total number of prepared statements in the server. (The sum of the number of prepared statements across all sessions.) It can be used in environments where there is the potential for denial-of-service attacks based on running the server out of memory by preparing huge numbers of statements. If the value is set lower than the current number of prepared statements, existing statements are not affected and can be used, but no new statements can be prepared until the current number drops below the limit. The default value is 16,382. The permissible range of values is from 0 to 1 million. Setting the value to 0 disables prepared statements.

- [max\\_relay\\_log\\_size \[507\]](#)

<b>Command-Line Format</b>	<code>--max_relay_log_size=#</code>	
<b>Option-File Format</b>	<code>max_relay_log_size</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">max_relay_log_size [507]</a>	
<b>Variable Name</b>	<code>max_relay_log_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1073741824</code>

If a write by a replication slave to its relay log causes the current log file size to exceed the value of this variable, the slave rotates the relay logs (closes the current file and opens the next one). If [max\\_relay\\_log\\_size \[507\]](#) is 0, the server uses [max\\_binlog\\_size \[2016\]](#) for both the binary log and the relay log. If [max\\_relay\\_log\\_size \[507\]](#) is greater than 0, it constrains the size of the relay log, which enables you to have different sizes for the two logs. You must set [max\\_relay\\_log\\_size \[507\]](#) to between 4096 bytes and 1GB (inclusive), or to 0. The default value is 0. See [Section 16.2.1, “Replication Implementation Details”](#).

- [max\\_seeks\\_for\\_key \[507\]](#)

<b>Command-Line Format</b>	<code>--max_seeks_for_key=#</code>	
<b>Option-File Format</b>	<code>max_seeks_for_key</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">max_seeks_for_key [507]</a>	
<b>Variable Name</b>	<code>max_seeks_for_key</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>32</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4294967295</code>
	<b>Range</b>	<code>1 .. 4294967295</code>
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>18446744073709547520</code>

	<b>Range</b>	1 .. 18446744073709547520
--	--------------	---------------------------

Limit the assumed maximum number of seeks when looking up rows based on a key. The MySQL optimizer assumes that no more than this number of key seeks are required when searching for matching rows in a table by scanning an index, regardless of the actual cardinality of the index (see [Section 13.7.5.23, “SHOW INDEX Syntax”](#)). By setting this to a low value (say, 100), you can force MySQL to prefer indexes instead of table scans.

- [max\\_sort\\_length \[508\]](#)

<b>Command-Line Format</b>	<code>--max_sort_length=#</code>	
<b>Option-File Format</b>	<code>max_sort_length</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">max_sort_length [508]</a>	
<b>Variable Name</b>	<code>max_sort_length</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1024</code>
	<b>Range</b>	<code>4 .. 8388608</code>

The number of bytes to use when sorting data values. Only the first [max\\_sort\\_length \[508\]](#) bytes of each value are used; the rest are ignored.

As of MySQL 5.6.9, [max\\_sort\\_length \[508\]](#) is ignored for integer, decimal, floating-point, and temporal data types.

- [max\\_sp\\_recursion\\_depth \[508\]](#)

<b>Command-Line Format</b>	<code>--max_sp_recursion_depth[=#]</code>	
<b>Option-File Format</b>	<code>max_sp_recursion_depth</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">max_sp_recursion_depth [508]</a>	
<b>Variable Name</b>	<code>max_sp_recursion_depth</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Max Value</b>	<code>255</code>

The number of times that any given stored procedure may be called recursively. The default value for this option is 0, which completely disables recursion in stored procedures. The maximum value is 255.

Stored procedure recursion increases the demand on thread stack space. If you increase the value of [max\\_sp\\_recursion\\_depth \[508\]](#), it may be necessary to increase thread stack size by increasing the value of [thread\\_stack \[556\]](#) at server startup.

- [max\\_tmp\\_tables \[508\]](#)

This variable is unused. It is deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release.

- `max_user_connections` [509]

<b>Command-Line Format</b>	<code>--max_user_connections=#</code>	
<b>Option-File Format</b>	<code>max_user_connections</code>	
<b>Option Sets Variable</b>	Yes, <code>max_user_connections</code> [509]	
<b>Variable Name</b>	<code>max_user_connections</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 4294967295</code>

The maximum number of simultaneous connections permitted to any given MySQL user account. A value of 0 (the default) means “no limit.”

This variable has a global value that can be set at server startup or runtime. It also has a read-only session value that indicates the effective simultaneous-connection limit that applies to the account associated with the current session. The session value is initialized as follows:

- If the user account has a nonzero `MAX_USER_CONNECTIONS` resource limit, the session `max_user_connections` [509] value is set to that limit.
- Otherwise, the session `max_user_connections` [509] value is set to the global value.

Account resource limits are specified using the `GRANT` statement. See [Section 6.3.4, “Setting Account Resource Limits”](#), and [Section 13.7.1.4, “GRANT Syntax”](#).

- `max_write_lock_count` [509]

<b>Command-Line Format</b>	<code>--max_write_lock_count=#</code>	
<b>Option-File Format</b>	<code>max_write_lock_count</code>	
<b>Option Sets Variable</b>	Yes, <code>max_write_lock_count</code> [509]	
<b>Variable Name</b>	<code>max_write_lock_count</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>32</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4294967295</code>
	<b>Range</b>	<code>1 .. 4294967295</code>
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>

	<b>Default</b>	18446744073709547520
	<b>Range</b>	1 .. 18446744073709547520

After this many write locks, permit some pending read lock requests to be processed in between.

- [metadata\\_locks\\_cache\\_size \[510\]](#)

<b>Version Introduced</b>	5.6.4	
<b>Variable Name</b>	<code>metadata_locks_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1024
	<b>Range</b>	1 .. 1048576

The size of the metadata locks cache. The server uses this cache to avoid creation and destruction of synchronization objects. This is particularly helpful on systems where such operations are expensive, such as Windows XP. This variable was added in MySQL 5.6.4.

- [metadata\\_locks\\_hash\\_instances \[510\]](#)

<b>Version Introduced</b>	5.6.8	
<b>Variable Name</b>	<code>metadata_locks_hash_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	8
	<b>Range</b>	1 .. 1024

The set of metadata locks can be partitioned into separate hashes to permit connections accessing different objects to use different locking hashes and reduce contention. The [metadata\\_locks\\_hash\\_instances \[510\]](#) system variable specifies the number of hashes (default 8). This variable was added in MySQL 5.6.8.

- [min\\_examined\\_row\\_limit \[510\]](#)

<b>Command-Line Format</b>	<code>--min-examined-row-limit=#</code>	
<b>Option-File Format</b>	<code>min-examined-row-limit</code>	
<b>Variable Name</b>	<code>min_examined_row_limit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 4294967295

	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 18446744073709547520

Queries that examine fewer than this number of rows are not logged to the slow query log.

- `myisam_data_pointer_size` [511]

<b>Command-Line Format</b>	<code>--myisam_data_pointer_size=#</code>
<b>Option-File Format</b>	<code>myisam_data_pointer_size</code>
<b>Option Sets Variable</b>	Yes, <code>myisam_data_pointer_size</code> [511]
<b>Variable Name</b>	<code>myisam_data_pointer_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	6
<b>Range</b>	2 .. 7

The default pointer size in bytes, to be used by `CREATE TABLE` for `MyISAM` tables when no `MAX_ROWS` option is specified. This variable cannot be less than 2 or larger than 7. The default value is 6. See [Section C.5.2.12](#), “The table is full”.

- `myisam_max_sort_file_size` [511]

<b>Command-Line Format</b>	<code>--myisam_max_sort_file_size=#</code>
<b>Option-File Format</b>	<code>myisam_max_sort_file_size</code>
<b>Option Sets Variable</b>	Yes, <code>myisam_max_sort_file_size</code> [511]
<b>Variable Name</b>	<code>myisam_max_sort_file_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	2147483648
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	9223372036854775807

The maximum size of the temporary file that MySQL is permitted to use while re-creating a `MyISAM` index (during `REPAIR TABLE`, `ALTER TABLE`, or `LOAD DATA INFILE`). If the file size would be

larger than this value, the index is created using the key cache instead, which is slower. The value is given in bytes.

The default value is 2GB. If `MyISAM` index files exceed this size and disk space is available, increasing the value may help performance. The space must be available in the file system containing the directory where the original index file is located.

- `myisam_mmap_size` [512]

<b>Command-Line Format</b>	<code>--myisam_mmap_size=#</code>	
<b>Option-File Format</b>	<code>myisam_mmap_size</code>	
<b>Option Sets Variable</b>	Yes, <code>myisam_mmap_size</code> [512]	
<b>Variable Name</b>	<code>myisam_mmap_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	4294967295
	<b>Range</b>	7 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	18446744073709547520
	<b>Range</b>	7 .. 18446744073709547520

The maximum amount of memory to use for memory mapping compressed `MyISAM` files. If many compressed `MyISAM` tables are used, the value can be decreased to reduce the likelihood of memory-swapping problems.

- `myisam_recover_options` [512]

<b>Variable Name</b>	<code>myisam_recover_options</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

The value of the `--myisam-recover-options` [432] option. See [Section 5.1.3, “Server Command Options”](#).

- `myisam_repair_threads` [512]

<b>Command-Line Format</b>	<code>--myisam_repair_threads=#</code>	
<b>Option-File Format</b>	<code>myisam_repair_threads</code>	
<b>Option Sets Variable</b>	Yes, <code>myisam_repair_threads</code> [512]	
<b>Variable Name</b>	<code>myisam_repair_threads</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	

	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Range</b>	1 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Range</b>	1 .. 18446744073709547520

If this value is greater than 1, `MyISAM` table indexes are created in parallel (each index in its own thread) during the `Repair by sorting` process. The default value is 1.

**Note**

Multi-threaded repair is still *beta-quality* code.

- `myisam_sort_buffer_size` [513]

<b>Command-Line Format</b>	<code>--myisam_sort_buffer_size=#</code>	
<b>Option-File Format</b>	<code>myisam_sort_buffer_size</code>	
<b>Option Sets Variable</b>	Yes, <code>myisam_sort_buffer_size</code> [513]	
<b>Variable Name</b>	<code>myisam_sort_buffer_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	8388608
	<b>Range</b>	4 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	8388608
	<b>Range</b>	4 .. 18446744073709547520

The size of the buffer that is allocated when sorting `MyISAM` indexes during a `REPAIR TABLE` or when creating indexes with `CREATE INDEX` or `ALTER TABLE`.

The maximum permissible setting for `myisam_sort_buffer_size` [513] is 4GB. Values larger than 4GB are permitted for 64-bit platforms (except 64-bit Windows, for which large values are truncated to 4GB with a warning).

- `myisam_stats_method` [513]



<b>Command-Line Format</b>	<code>--myisam_stats_method=name</code>
<b>Option-File Format</b>	<code>myisam_stats_method</code>
<b>Option Sets Variable</b>	Yes, <code>myisam_stats_method</code> [513]
<b>Variable Name</b>	<code>myisam_stats_method</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>enumeration</code>
	<b>Valid Values</b> <code>nulls_equal</code> <code>nulls_unequal</code> <code>nulls_ignored</code>

How the server treats `NULL` values when collecting statistics about the distribution of index values for `MyISAM` tables. This variable has three possible values, `nulls_equal`, `nulls_unequal`, and `nulls_ignored`. For `nulls_equal`, all `NULL` index values are considered equal and form a single value group that has a size equal to the number of `NULL` values. For `nulls_unequal`, `NULL` values are considered unequal, and each `NULL` forms a distinct value group of size 1. For `nulls_ignored`, `NULL` values are ignored.

The method that is used for generating table statistics influences how the optimizer chooses indexes for query execution, as described in [Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”](#).

- `myisam_use_mmap` [514]

<b>Command-Line Format</b>	<code>--myisam_use_mmap</code>
<b>Option-File Format</b>	<code>myisam_use_mmap</code>
<b>Option Sets Variable</b>	Yes, <code>myisam_use_mmap</code> [514]
<b>Variable Name</b>	<code>myisam_use_mmap</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>OFF</code>

Use memory mapping for reading and writing `MyISAM` tables.

- `named_pipe` [514]

<b>Variable Name</b>	<code>named_pipe</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Platform Specific</b>	windows
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code> (windows)
	<b>Default</b> <code>OFF</code>

(Windows only.) Indicates whether the server supports connections over named pipes.

- [net\\_buffer\\_length \[515\]](#)

<b>Command-Line Format</b>	<code>--net_buffer_length=#</code>	
<b>Option-File Format</b>	<code>net_buffer_length</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">net_buffer_length [515]</a>	
<b>Variable Name</b>	<code>net_buffer_length</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>16384</code>
	<b>Range</b>	<code>1024 .. 1048576</code>

Each client thread is associated with a connection buffer and result buffer. Both begin with a size given by [net\\_buffer\\_length \[515\]](#) but are dynamically enlarged up to [max\\_allowed\\_packet \[502\]](#) bytes as needed. The result buffer shrinks to [net\\_buffer\\_length \[515\]](#) after each SQL statement.

This variable should not normally be changed, but if you have very little memory, you can set it to the expected length of statements sent by clients. If statements exceed this length, the connection buffer is automatically enlarged. The maximum value to which [net\\_buffer\\_length \[515\]](#) can be set is 1MB.

The session value of this variable is read only.

- [net\\_read\\_timeout \[515\]](#)

<b>Command-Line Format</b>	<code>--net_read_timeout=#</code>	
<b>Option-File Format</b>	<code>net_read_timeout</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">net_read_timeout [515]</a>	
<b>Variable Name</b>	<code>net_read_timeout</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>30</code>
	<b>Min Value</b>	<code>1</code>

The number of seconds to wait for more data from a connection before aborting the read. When the server is reading from the client, [net\\_read\\_timeout \[515\]](#) is the timeout value controlling when to abort. When the server is writing to the client, [net\\_write\\_timeout \[516\]](#) is the timeout value controlling when to abort. See also [slave\\_net\\_timeout \[1993\]](#).

- [net\\_retry\\_count \[515\]](#)

<b>Command-Line Format</b>	<code>--net_retry_count=#</code>
<b>Option-File Format</b>	<code>net_retry_count</code>

<b>Option Sets Variable</b>	Yes, <a href="#">net_retry_count</a> [515]
<b>Variable Name</b>	<a href="#">net_retry_count</a>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <a href="#">numeric</a>
	<b>Default</b> 10
	<b>Range</b> 1 .. 4294967295
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> <a href="#">numeric</a>
	<b>Default</b> 10
	<b>Range</b> 1 .. 18446744073709547520

If a read or write on a communication port is interrupted, retry this many times before giving up. This value should be set quite high on FreeBSD because internal interrupts are sent to all threads.

- [net\\_write\\_timeout](#) [516]

<b>Command-Line Format</b>	<code>--net_write_timeout=#</code>
<b>Option-File Format</b>	<a href="#">net_write_timeout</a>
<b>Option Sets Variable</b>	Yes, <a href="#">net_write_timeout</a> [516]
<b>Variable Name</b>	<a href="#">net_write_timeout</a>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <a href="#">numeric</a>
	<b>Default</b> 60
	<b>Min Value</b> 1

The number of seconds to wait for a block to be written to a connection before aborting the write. See also [net\\_read\\_timeout](#) [515].

- [new](#) [516]

<b>Command-Line Format</b>	<code>--new</code>
	<code>-n</code>
<b>Option-File Format</b>	<a href="#">new</a>
<b>Option Sets Variable</b>	Yes, <a href="#">new</a> [516]
<b>Variable Name</b>	<a href="#">new</a>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
<b>Disabled by</b>	<a href="#">skip-new</a>

	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	FALSE	

This variable was used in MySQL 4.0 to turn on some 4.1 behaviors, and is retained for backward compatibility. In MySQL 5.6, its value is always `OFF`.

- [old \[517\]](#)

<b>Command-Line Format</b>	<code>--old</code>
<b>Option-File Format</b>	<code>old</code>
<b>Variable Name</b>	<code>old</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

`old [517]` is a compatibility variable. It is disabled by default, but can be enabled at startup to revert the server to behaviors present in older versions.

Currently, when `old [517]` is enabled, it changes the default scope of index hints to that used prior to MySQL 5.1.17. That is, index hints with no `FOR` clause apply only to how indexes are used for row retrieval and not to resolution of `ORDER BY` or `GROUP BY` clauses. (See [Section 13.2.9.3, “Index Hint Syntax”](#).) Take care about enabling this in a replication setup. With statement-based binary logging, having different modes for the master and slaves might lead to replication errors.

- [old\\_alter\\_table \[517\]](#)

<b>Command-Line Format</b>	<code>--old-alter-table</code>	
<b>Option-File Format</b>	<code>old-alter-table</code>	
<b>Option Sets Variable</b>	Yes, <code>old_alter_table [517]</code>	
<b>Variable Name</b>	<code>old_alter_table</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	OFF	

When this variable is enabled, the server does not use the optimized method of processing an `ALTER TABLE` operation. It reverts to using a temporary table, copying over the data, and then renaming the temporary table to the original, as used by MySQL 5.0 and earlier. For more information on the operation of `ALTER TABLE`, see [Section 13.1.6, “ALTER TABLE Syntax”](#).

- [old\\_passwords \[517\]](#)

<b>Variable Name</b>	<code>old_passwords</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
<b>Type</b>	boolean	
<b>Default</b>	0	
	<b>Permitted Values (&gt;= 5.6.6)</b>	

<b>Type</b>	enumeration
<b>Default</b>	0
<b>Valid Values</b>	0 1 2

This variable determines the type of password hashing performed by the `PASSWORD()` [1263] function and statements such as `CREATE USER` and `GRANT`.

The following table shows the permitted values of `old_passwords` [517], as well as the type of password hashing each value produces, and the plugins that use each hashing type. These values are permitted as of MySQL 5.6.6. Before 5.6.6, the value can be 0 (or `OFF`), or 1 (or `ON`).

Value	Password Hashing Format	Intended Use
0	MySQL 4.1 native hashing	Accounts that authenticate with the <code>mysql_native_password</code> plugin
1	Pre-4.1 (“old”) native hashing	Accounts that authenticate with the <code>mysql_old_password</code> plugin
2	SHA-256 hashing	Accounts that authenticate with the <code>sha256_password</code> plugin

If `old_passwords=1`, `PASSWORD('str')` [1263] returns the same value as `OLD_PASSWORD('str')` [1263].

For information about authentication plugins and hashing formats, see [Section 6.3.7, “Pluggable Authentication”](#), and [Section 6.1.2.4, “Password Hashing in MySQL”](#). If you set `old_passwords=2`, follow the instructions for using the `sha256_password` plugin at [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

- `open_files_limit` [518]

<b>Command-Line Format</b>	<code>--open-files-limit=#</code>
<b>Option-File Format</b>	<code>open-files-limit</code>
<b>Option Sets Variable</b>	Yes, <code>open_files_limit</code> [518]
<b>Variable Name</b>	<code>open_files_limit</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&lt;= 5.6.7)</b>
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 65535
	<b>Permitted Values (&gt;= 5.6.8)</b>
<b>Type</b>	numeric
<b>Default</b>	-1 (autosized)
<b>Range</b>	0 .. 65535

The number of files that the operating system permits `mysqld` to open. The value of this variable at runtime is the real value permitted by the system and might be different from the value you specify at server startup. The value is 0 on systems where MySQL cannot change the number of open files.

The effective `open_files_limit` [518] value is based on the value specified at system startup (if any) and the values of `max_connections` [503] and `table_open_cache` [551], using these formulas:

```
1) 10 + max_connections + (table_open_cache * 2)
2) max_connections * 5
3) open_files_limit value specified at startup, 5000 if none
```

The server attempts to obtain the number of file descriptors using the maximum of those three values. If that many descriptors cannot be obtained, the server attempts to obtain as many as the system will permit.

- `optimizer_join_cache_level` [519]

<b>Version Introduced</b>	5.6.1	
<b>Version Removed</b>	5.6.3	
<b>Command-Line Format</b>	<code>--optimizer_join_cache_level=#</code>	
<b>Option-File Format</b>	<code>optimizer_join_cache_level</code>	
<b>Option Sets Variable</b>	Yes, <code>optimizer_join_cache_level</code> [519]	
<b>Variable Name</b>	<code>optimizer_join_cache_level</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4</code>
	<b>Range</b>	<code>0 .. 4</code>

Before MySQL 5.6.3, this variable is used for join buffer management. It controls how join buffers are used for join operations. As of MySQL 5.6.3, it is removed and the `optimizer_switch` [520] variable is used instead. See [Section 8.13.12, “Block Nested-Loop and Batched Key Access Joins”](#).

The following table shows the permissible `optimizer_join_cache_level` [519] values.

Option	Description
0	No join buffer is used for any join operation. This setting can be useful for assessing baseline join performance in comparison to performance with nonzero values that enable use of join buffering.
1	This is the default value. Join buffers are employed only for inner joins that are executed by the original Block Nested-Loop (BNL) join algorithm. When this algorithm is applied, rows of the inner table are accessed through a table scan, a plain index scan, or a range index scan.
2	The server employs an incremental join buffer for a join operation if its first operand is produced by a join operation that uses a join buffer itself.
3	The BNL algorithm is used for outer join operations with one inner table and for inner joins.
4	The BNL algorithm uses incremental buffers for inner tables. In this case, the BNL algorithm can be used for nested outer joins (outer joins with several inner tables). Such an operation can be executed only if incremental join buffers are used to join all inner tables but the first one.

- `optimizer_prune_level` [519]

<b>Command-Line Format</b>	<code>--optimizer_prune_level[=#]</code>
<b>Option-File Format</b>	<code>optimizer_prune_level</code>
<b>Option Sets Variable</b>	Yes, <code>optimizer_prune_level</code> [519]
<b>Variable Name</b>	<code>optimizer_prune_level</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>1</code>

Controls the heuristics applied during query optimization to prune less-promising partial plans from the optimizer search space. A value of 0 disables heuristics so that the optimizer performs an exhaustive search. A value of 1 causes the optimizer to prune plans based on the number of rows retrieved by intermediate plans.

- `optimizer_search_depth` [520]

<b>Command-Line Format</b>	<code>--optimizer_search_depth[=#]</code>
<b>Option-File Format</b>	<code>optimizer_search_depth</code>
<b>Option Sets Variable</b>	Yes, <code>optimizer_search_depth</code> [520]
<b>Variable Name</b>	<code>optimizer_search_depth</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&gt;= 5.6)</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>62</code>
	<b>Range</b> <code>0 .. 62</code>

The maximum depth of search performed by the query optimizer. Values larger than the number of relations in a query result in better query plans, but take longer to generate an execution plan for a query. Values smaller than the number of relations in a query return an execution plan quicker, but the resulting plan may be far from being optimal. If set to 0, the system automatically picks a reasonable value.

- `optimizer_switch` [520]

<b>Command-Line Format</b>	<code>--optimizer_switch=value</code>
<b>Option-File Format</b>	<code>optimizer_switch</code>
	<code>optimizer_switch</code>
	<code>optimizer_switch</code>
<b>Option Sets Variable</b>	Yes, <code>optimizer_switch</code> [520]
<b>Variable Name</b>	<code>optimizer_switch</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&gt;= 5.6.1, &lt;= 5.6.2)</b>
	<b>Type</b> <code>set</code>

	<b>Valid Values</b>	<p>engine_condition_pushdown={on off}</p> <p>index_condition_pushdown={on off}</p> <p>index_merge={on off}</p> <p>index_merge_intersection={on off}</p> <p>index_merge_sort_union={on off}</p> <p>index_merge_union={on off}</p> <p>mrr={on off}</p> <p>mrr_cost_based={on off}</p>
<b>Permitted Values (&gt;= 5.6.3, &lt;= 5.6.4)</b>		
	<b>Type</b>	set
	<b>Valid Values</b>	<p>batched_key_access={on off}</p> <p>block_nested_loop={on off}</p> <p>engine_condition_pushdown={on off}</p> <p>index_condition_pushdown={on off}</p> <p>index_merge={on off}</p> <p>index_merge_intersection={on off}</p> <p>index_merge_sort_union={on off}</p> <p>index_merge_union={on off}</p> <p>mrr={on off}</p> <p>mrr_cost_based={on off}</p>
<b>Permitted Values (&gt;= 5.6.5, &lt;= 5.6.6)</b>		
	<b>Type</b>	set
	<b>Valid Values</b>	<p>batched_key_access={on off}</p> <p>block_nested_loop={on off}</p> <p>engine_condition_pushdown={on off}</p> <p>firstmatch={on off}</p> <p>index_condition_pushdown={on off}</p> <p>index_merge={on off}</p> <p>index_merge_intersection={on off}</p> <p>index_merge_sort_union={on off}</p> <p>index_merge_union={on off}</p> <p>loosescan={on off}</p> <p>materialization={on off}</p>



	<pre>mrr={on off} mrr_cost_based={on off} semijoin={on off}</pre>				
	<p><b>Permitted Values (&gt;= 5.6.7, &lt;= 5.6.8)</b></p> <table border="1"> <tr> <td><b>Type</b></td> <td>set</td> </tr> <tr> <td><b>Valid Values</b></td> <td> <pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off} mrr={on off} mrr_cost_based={on off} semijoin={on off} subquery_materialization_cost_based={on off}</pre> </td> </tr> </table>	<b>Type</b>	set	<b>Valid Values</b>	<pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off} mrr={on off} mrr_cost_based={on off} semijoin={on off} subquery_materialization_cost_based={on off}</pre>
<b>Type</b>	set				
<b>Valid Values</b>	<pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off} mrr={on off} mrr_cost_based={on off} semijoin={on off} subquery_materialization_cost_based={on off}</pre>				
	<p><b>Permitted Values (&gt;= 5.6.9)</b></p> <table border="1"> <tr> <td><b>Type</b></td> <td>set</td> </tr> <tr> <td><b>Valid Values</b></td> <td> <pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off}</pre> </td> </tr> </table>	<b>Type</b>	set	<b>Valid Values</b>	<pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off}</pre>
<b>Type</b>	set				
<b>Valid Values</b>	<pre>batched_key_access={on off} block_nested_loop={on off} engine_condition_pushdown={on off} firstmatch={on off} index_condition_pushdown={on off} index_merge={on off} index_merge_intersection={on off} index_merge_sort_union={on off} index_merge_union={on off} loosescan={on off} materialization={on off}</pre>				

	<pre>mrr={on off} mrr_cost_based={on off} semijoin={on off} subquery_materialization_cost_based={on off} use_index_extensions={on off}</pre>
--	--

The `optimizer_switch` [520] system variable enables control over optimizer behavior. The value of this variable is a set of flags, each of which has a value of `on` or `off` to indicate whether the corresponding optimizer behavior is enabled or disabled. This variable has global and session values and can be changed at runtime. The global default can be set at server startup.

To see the current set of optimizer flags, select the variable value:

```
mysql> SELECT @@optimizer_switch\G
***** 1. row *****
@@optimizer_switch: index_merge=on,index_merge_union=on,
                    index_merge_sort_union=on,
                    index_merge_intersection=on,
                    engine_condition_pushdown=on,
                    index_condition_pushdown=on,
                    mrr=on,mrr_cost_based=on,
                    block_nested_loop=on,batched_key_access=off,
                    materialization=on,semijoin=on,loosescan=on,
                    firstmatch=on,
                    subquery_materialization_cost_based=on,
                    use_index_extensions=on
```

For more information about the syntax of this variable and the optimizer behaviors that it controls, see [Section 8.8.5.2, “Controlling Switchable Optimizations”](#).

- `optimizer_trace` [523]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>optimizer_trace</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This variable controls optimizer tracing. For details, see [MySQL Internals: Tracing the Optimizer](#). This variable was added in MySQL 5.6.3.

- `optimizer_trace_features` [523]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>optimizer_trace_features</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This variable enables or disables selected optimizer tracing features. For details, see [MySQL Internals: Tracing the Optimizer](#). This variable was added in MySQL 5.6.3.

- `optimizer_trace_limit` [524]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>optimizer_trace_limit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>

The maximum number of optimizer traces to display. For details, see [MySQL Internals: Tracing the Optimizer](#). This variable was added in MySQL 5.6.3.

- `optimizer_trace_max_mem_size` [524]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>optimizer_trace_max_mem_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>16384</code>

The maximum cumulative size of stored optimizer traces. For details, see [MySQL Internals: Tracing the Optimizer](#). This variable was added in MySQL 5.6.3.

- `optimizer_trace_offset` [524]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>optimizer_trace_offset</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1</code>

The offset of optimizer traces to display. For details, see [MySQL Internals: Tracing the Optimizer](#). This variable was added in MySQL 5.6.3.

- `performance_schema_xxx`

Performance Schema system variables are listed in [Section 20.12, “Performance Schema System Variables”](#). These variables may be used to configure Performance Schema operation.

- `pid_file` [524]

<b>Command-Line Format</b>	<code>--pid-file=file_name</code>
<b>Option-File Format</b>	<code>pid-file=file_name</code>
<b>Option Sets Variable</b>	Yes, <code>pid_file</code> [524]
<b>Variable Name</b>	<code>pid_file</code>
<b>Variable Scope</b>	Global

<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path name of the process ID (PID) file. This variable can be set with the `--pid-file` [435] option.

- `plugin_dir` [525]

<b>Command-Line Format</b>	<code>--plugin_dir=path</code>	
<b>Option-File Format</b>	<code>plugin_dir</code>	
<b>Option Sets Variable</b>	Yes, <code>plugin_dir</code> [525]	
<b>Variable Name</b>	<code>plugin_dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	directory name
	<b>Default</b>	<code>BASEDIR/lib/plugin</code>

The path name of the plugin directory.

If the plugin directory is writable by the server, it may be possible for a user to write executable code to a file in the directory using `SELECT ... INTO DUMPFILE`. This can be prevented by making `plugin_dir` [525] read only to the server or by setting `--secure-file-priv` [438] to a directory where `SELECT` writes can be made safely.

- `port` [525]

<b>Command-Line Format</b>	<code>--port=#</code>	
	<code>-P</code>	
<b>Option-File Format</b>	<code>port</code>	
<b>Option Sets Variable</b>	Yes, <code>port</code> [525]	
<b>Variable Name</b>	<code>port</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	3306
	<b>Range</b>	0 .. 65535

The number of the port on which the server listens for TCP/IP connections. This variable can be set with the `--port` [437] option.

- `preload_buffer_size` [525]

<b>Command-Line Format</b>	<code>--preload_buffer_size=#</code>	
<b>Option-File Format</b>	<code>preload_buffer_size</code>	
<b>Option Sets Variable</b>	Yes, <code>preload_buffer_size</code> [525]	
<b>Variable Name</b>	<code>preload_buffer_size</code>	

<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	32768
	<b>Range</b>	1024 .. 1073741824

The size of the buffer that is allocated when preloading indexes.

- [profiling \[526\]](#)

If set to 0 or `OFF` (the default), statement profiling is disabled. If set to 1 or `ON`, statement profiling is enabled and the `SHOW PROFILE` and `SHOW PROFILES` statements provide access to profiling information. See [Section 13.7.5.32, “SHOW PROFILES Syntax”](#).

This variable is deprecated in MySQL 5.6.8 and will be removed in a future MySQL release.

- [profiling\\_history\\_size \[526\]](#)

The number of statements for which to maintain profiling information if [profiling \[526\]](#) is enabled. The default value is 15. The maximum value is 100. Setting the value to 0 effectively disables profiling. See [Section 13.7.5.32, “SHOW PROFILES Syntax”](#).

This variable is deprecated in MySQL 5.6.8 and will be removed in a future MySQL release.

- [protocol\\_version \[526\]](#)

<b>Variable Name</b>	<code>protocol_version</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric

The version of the client/server protocol used by the MySQL server.

- [proxy\\_user \[526\]](#)

<b>Variable Name</b>	<code>proxy_user</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

If the current client is a proxy for another user, this variable is the proxy user account name. Otherwise, this variable is `NULL`. See [Section 6.3.8, “Proxy Users”](#).

- [pseudo\\_slave\\_mode \[526\]](#)

<b>Version Introduced</b>	5.6.10
<b>Variable Name</b>	<code>pseudo_slave_mode</code>
<b>Variable Scope</b>	Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>

	<b>Type</b>	numeric
--	-------------	---------

This variable is for internal server use. It was added in MySQL 5.6.10.

- [pseudo\\_thread\\_id \[527\]](#)

<b>Variable Name</b>	pseudo_thread_id	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric

This variable is for internal server use.

- [query\\_alloc\\_block\\_size \[527\]](#)

<b>Command-Line Format</b>	--query_alloc_block_size=#	
<b>Option-File Format</b>	query_alloc_block_size	
<b>Option Sets Variable</b>	Yes, <a href="#">query_alloc_block_size [527]</a>	
<b>Variable Name</b>	query_alloc_block_size	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	8192
	<b>Range</b>	1024 .. 4294967295
	<b>Block Size</b>	1024
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	8192
	<b>Range</b>	1024 .. 18446744073709547520
	<b>Block Size</b>	1024

The allocation size of memory blocks that are allocated for objects created during statement parsing and execution. If you have problems with memory fragmentation, it might help to increase this parameter.

- [query\\_cache\\_limit \[527\]](#)

<b>Command-Line Format</b>	--query_cache_limit=#	
<b>Option-File Format</b>	query_cache_limit	
<b>Option Sets Variable</b>	Yes, <a href="#">query_cache_limit [527]</a>	
<b>Variable Name</b>	query_cache_limit	

<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	0 .. 4294967295
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	0 .. 18446744073709547520

Do not cache results that are larger than this number of bytes. The default value is 1MB.

- [query\\_cache\\_min\\_res\\_unit \[528\]](#)

<b>Command-Line Format</b>	--query_cache_min_res_unit=#
<b>Option-File Format</b>	query_cache_min_res_unit
<b>Option Sets Variable</b>	Yes, <a href="#">query_cache_min_res_unit [528]</a>
<b>Variable Name</b>	query_cache_min_res_unit
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	4096
<b>Range</b>	512 .. 4294967295
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	4096
<b>Range</b>	512 .. 18446744073709547520

The minimum size (in bytes) for blocks allocated by the query cache. The default value is 4096 (4KB). Tuning information for this variable is given in [Section 8.9.3.3, "Query Cache Configuration"](#).

- [query\\_cache\\_size \[528\]](#)

<b>Command-Line Format</b>	--query_cache_size=#
<b>Option-File Format</b>	query_cache_size
<b>Option Sets Variable</b>	Yes, <a href="#">query_cache_size [528]</a>
<b>Variable Name</b>	query_cache_size

<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.7)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&lt;= 5.6.7)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.8)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.8)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	0 .. 4294967295

The amount of memory allocated for caching query results. By default, the query cache is disabled. This is achieved using a default value of 1M, with a default for `query_cache_type` of 0. (Before MySQL 5.6.8, the default size is 0, with a default `query_cache_type` of 1. To reduce overhead significantly, you should also start the server with `query_cache_type=0` [529] if you will not be using the query cache.)

The permissible values are multiples of 1024; other values are rounded down to the nearest multiple. Note that `query_cache_size` [528] bytes of memory are allocated even if `query_cache_type` [529] is set to 0. See Section 8.9.3.3, “Query Cache Configuration”, for more information.

The query cache needs a minimum size of about 40KB to allocate its structures. (The exact size depends on system architecture.) If you set the value of `query_cache_size` [528] too small, a warning will occur, as described in Section 8.9.3.3, “Query Cache Configuration”.

- `query_cache_type` [529]

<b>Command-Line Format</b>	<code>--query_cache_type=#</code>
<b>Option-File Format</b>	<code>query_cache_type</code>
<b>Option Sets Variable</b>	Yes, <code>query_cache_type</code> [529]
<b>Variable Name</b>	<code>query_cache_type</code>
<b>Variable Scope</b>	Global, Session



<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.7)</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	1
	<b>Valid Values</b>	0 1 2
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	0
	<b>Valid Values</b>	0 1 2

Set the query cache type. Setting the `GLOBAL` value sets the type for all clients that connect thereafter. Individual clients can set the `SESSION` value to affect their own use of the query cache. Possible values are shown in the following table.

Option	Description
0 or <code>OFF</code>	Do not cache results in or retrieve results from the query cache. Note that this does not deallocate the query cache buffer. To do that, you should set <code>query_cache_size</code> [528] to 0.
1 or <code>ON</code>	Cache all cacheable query results except for those that begin with <code>SELECT SQL_NO_CACHE</code> .
2 or <code>DEMAND</code>	Cache results only for cacheable queries that begin with <code>SELECT SQL_CACHE</code> .

This variable defaults to `OFF` as of MySQL 5.6.8, `ON` before that.

If the server is started with `query_cache_type` set to 0, it does not acquire the query cache mutex at all, which means that the query cache cannot be enabled at runtime and there is reduced overhead in query execution.

- `query_cache_wlock_invalidate` [530]

<b>Command-Line Format</b>	<code>--query_cache_wlock_invalidate</code>	
<b>Option-File Format</b>	<code>query_cache_wlock_invalidate</code>	
<b>Option Sets Variable</b>	Yes, <code>query_cache_wlock_invalidate</code> [530]	
<b>Variable Name</b>	<code>query_cache_wlock_invalidate</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	<code>FALSE</code>

Normally, when one client acquires a `WRITE` lock on a `MyISAM` table, other clients are not blocked from issuing statements that read from the table if the query results are present in the query cache. Setting this variable to 1 causes acquisition of a `WRITE` lock for a table to invalidate any queries in

the query cache that refer to the table. This forces other clients that attempt to access the table to wait while the lock is in effect.

- [query\\_prealloc\\_size \[531\]](#)

<b>Command-Line Format</b>	<code>--query_prealloc_size=#</code>	
<b>Option-File Format</b>	<code>query_prealloc_size</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">query_prealloc_size [531]</a>	
<b>Variable Name</b>	<code>query_prealloc_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	8192
	<b>Range</b>	8192 .. 4294967295
	<b>Block Size</b>	1024
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	8192
	<b>Range</b>	8192 .. 18446744073709547520
	<b>Block Size</b>	1024

The size of the persistent buffer used for statement parsing and execution. This buffer is not freed between statements. If you are running complex queries, a larger [query\\_prealloc\\_size \[531\]](#) value might be helpful in improving performance, because it can reduce the need for the server to perform memory allocation during query execution operations.

- [rand\\_seed1 \[531\]](#)

The [rand\\_seed1 \[531\]](#) and [rand\\_seed2 \[531\]](#) variables exist as session variables only, and can be set but not read. The variables—but not their values—are shown in the output of `SHOW VARIABLES`.

The purpose of these variables is to support replication of the [RAND\( \) \[1206\]](#) function. For statements that invoke [RAND\( \) \[1206\]](#), the master passes two values to the slave, where they are used to seed the random number generator. The slave uses these values to set the session variables [rand\\_seed1 \[531\]](#) and [rand\\_seed2 \[531\]](#) so that [RAND\( \) \[1206\]](#) on the slave generates the same value as on the master.

- [rand\\_seed2 \[531\]](#)

See the description for [rand\\_seed1 \[531\]](#).

- [range\\_alloc\\_block\\_size \[531\]](#)

<b>Command-Line Format</b>	<code>--range_alloc_block_size=#</code>
----------------------------	---

<b>Option-File Format</b>	<code>range_alloc_block_size</code>
<b>Option Sets Variable</b>	Yes, <a href="#">range_alloc_block_size [531]</a>
<b>Variable Name</b>	<code>range_alloc_block_size</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	4096
<b>Range</b>	4096 .. 4294967295
<b>Block Size</b>	1024

The size of blocks that are allocated when doing range optimization.

- [read\\_buffer\\_size \[532\]](#)

<b>Command-Line Format</b>	<code>--read_buffer_size=#</code>
<b>Option-File Format</b>	<code>read_buffer_size</code>
<b>Option Sets Variable</b>	Yes, <a href="#">read_buffer_size [532]</a>
<b>Variable Name</b>	<code>read_buffer_size</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	131072
<b>Range</b>	8200 .. 2147479552

Each thread that does a sequential scan for a `MyISAM` table allocates a buffer of this size (in bytes) for each table it scans. If you do many sequential scans, you might want to increase this value, which defaults to 131072. The value of this variable should be a multiple of 4KB. If it is set to a value that is not a multiple of 4KB, its value will be rounded down to the nearest multiple of 4KB.

This option is also used in the following context for all search engines:

- For caching the indexes in a temporary file (not a temporary table), when sorting rows for `ORDER BY`.
- For bulk insert into partitions.
- For caching results of nested queries.

and in one other storage engine-specific way: to determine the memory block size for `MEMORY` tables.

The maximum permissible setting for [read\\_buffer\\_size \[532\]](#) is 2GB.

For more information about memory use during different operations, see [Section 8.11.4.1, “How MySQL Uses Memory”](#).

- [read\\_only \[532\]](#)

<b>Command-Line Format</b>	<code>--read-only</code>
<b>Option-File Format</b>	<code>read_only</code>
<b>Option Sets Variable</b>	Yes, <code>read_only</code> [532]
<b>Variable Name</b>	<code>read_only</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>false</code>

This variable is off by default. When it is enabled, the server permits no updates except from users that have the `SUPER` [756] privilege or (on a slave server) from updates performed by slave threads. In replication setups, it can be useful to enable `read_only` [532] on slave servers to ensure that slaves accept updates only from the master server and not from clients.

`read_only` [532] does not apply to `TEMPORARY` tables, nor does it prevent the server from inserting rows into the log tables (see Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”). This variable does not prevent the use of `ANALYZE TABLE` or `OPTIMIZE TABLE` statements because its purpose is to prevent changes to table structure or contents. Analysis and optimization do not qualify as such changes. This means, for example, that consistency checks on read-only slaves can be performed with `mysqlcheck --all-databases --analyze`.

`read_only` [532] exists only as a `GLOBAL` variable, so changes to its value require the `SUPER` [756] privilege. Changes to `read_only` [532] on a master server are not replicated to slave servers. The value can be set on a slave server independent of the setting on the master.



### Important

In MySQL 5.6, enabling `read_only` prevents the use of the `SET PASSWORD` statement by any user not having the `SUPER` [756] privilege. This is not necessarily the case for all MySQL release series. When replicating from one MySQL release series to another (for example, from a MySQL 5.0 master to a MySQL 5.1 or later slave), you should check the documentation for the versions running on both master and slave to determine whether the behavior of `read_only` in this regard is or is not the same, and, if it is different, whether this has an impact on your applications.

The following conditions apply:

- If you attempt to enable `read_only` [532] while you have any explicit locks (acquired with `LOCK TABLES`) or have a pending transaction, an error occurs.
- If you attempt to enable `read_only` [532] while other clients hold explicit table locks or have pending transactions, the attempt blocks until the locks are released and the transactions end. While the attempt to enable `read_only` [532] is pending, requests by other clients for table locks or to begin transactions also block until `read_only` [532] has been set.
- `read_only` [532] can be enabled while you hold a global read lock (acquired with `FLUSH TABLES WITH READ LOCK`) because that does not involve table locks.

In MySQL 5.6, attempts to set `read_only` [532] block for active transactions that hold metadata locks until those transactions end.

- `read_rnd_buffer_size` [533]

<b>Command-Line Format</b>	<code>--read_rnd_buffer_size=#</code>
<b>Option-File Format</b>	<code>read_rnd_buffer_size</code>
<b>Option Sets Variable</b>	Yes, <code>read_rnd_buffer_size</code> [533]
<b>Variable Name</b>	<code>read_rnd_buffer_size</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>262144</code>
	<b>Range</b> <code>8200 .. 4294967295</code>

When reading rows from a `MyISAM` table in sorted order following a key-sorting operation, the rows are read through this buffer to avoid disk seeks. See [Section 8.13.13, “ORDER BY Optimization”](#). Setting the variable to a large value can improve `ORDER BY` performance by a lot. However, this is a buffer allocated for each client, so you should not set the global variable to a large value. Instead, change the session variable only from within those clients that need to run large queries.

The maximum permissible setting for `read_rnd_buffer_size` [533] is 2GB.

For more information about memory use during different operations, see [Section 8.11.4.1, “How MySQL Uses Memory”](#).

- `relay_log_purge` [534]

<b>Command-Line Format</b>	<code>--relay_log_purge</code>
<b>Option-File Format</b>	<code>relay_log_purge</code>
<b>Option Sets Variable</b>	Yes, <code>relay_log_purge</code> [534]
<b>Variable Name</b>	<code>relay_log_purge</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>TRUE</code>

Disables or enables automatic purging of relay log files as soon as they are not needed any more. The default value is 1 (`ON`).

- `relay_log_space_limit` [534]

<b>Command-Line Format</b>	<code>--relay_log_space_limit=#</code>
<b>Option-File Format</b>	<code>relay_log_space_limit</code>
<b>Option Sets Variable</b>	Yes, <code>relay_log_space_limit</code> [534]
<b>Variable Name</b>	<code>relay_log_space_limit</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> <code>32</code>

	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 18446744073709547520

The maximum amount of space to use for all relay logs.

- `report_host` [535]

<b>Command-Line Format</b>	<code>--report-host=host_name</code>	
<b>Option-File Format</b>	<code>report-host</code>	
<b>Option Sets Variable</b>	Yes, <code>report_host</code> [535]	
<b>Variable Name</b>	<code>report-host</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The value of the `--report-host` [1980] option.

- `report_password` [535]

<b>Command-Line Format</b>	<code>--report-password=name</code>	
<b>Option-File Format</b>	<code>report-password</code>	
<b>Option Sets Variable</b>	Yes, <code>report_password</code> [535]	
<b>Variable Name</b>	<code>report-password</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The value of the `--report-password` [1981] option. Not the same as the password used for the MySQL replication user account.

- `report_port` [535]

<b>Command-Line Format</b>	<code>--report-port=#</code>	
<b>Option-File Format</b>	<code>report-port</code>	
<b>Option Sets Variable</b>	Yes, <code>report_port</code> [535]	
<b>Variable Name</b>	<code>report-port</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b> ( $\geq 5.6.5$ )	

	<b>Type</b>	numeric
	<b>Default</b>	[slave_port]
	<b>Range</b>	0 .. 65535

The value of the `--report-port` [1981] option.

- `report_user` [536]

<b>Command-Line Format</b>	<code>--report-user=name</code>	
<b>Option-File Format</b>	<code>report-user</code>	
<b>Option Sets Variable</b>	Yes, <code>report_user</code> [536]	
<b>Variable Name</b>	<code>report-user</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The value of the `--report-user` [1982] option. Not the same as the name for the MySQL replication user account.

- `rpl_semi_sync_master_enabled` [536]

<b>Variable Name</b>	<code>rpl_semi_sync_master_enabled</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Controls whether semisynchronous replication is enabled on the master. To enable or disable the plugin, set this variable to `ON` or `OFF` (or 1 or 0), respectively. The default is `OFF`.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `rpl_semi_sync_master_timeout` [536]

<b>Variable Name</b>	<code>rpl_semi_sync_master_timeout</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10000

A value in milliseconds that controls how long the master waits on a commit for acknowledgment from a slave before timing out and reverting to asynchronous replication. The default value is 10000 (10 seconds).

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `rpl_semi_sync_master_trace_level` [536]

<b>Variable Name</b>	<code>rpl_semi_sync_master_trace_level</code>
----------------------	---

<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	32

The semisynchronous replication debug trace level on the master. Currently, four levels are defined:

- 1 = general level (for example, time function failures)
- 16 = detail level (more verbose information)
- 32 = net wait level (more information about network waits)
- 64 = function level (information about function entry and exit)

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `rpl_semi_sync_master_wait_no_slave` [537]

<b>Variable Name</b>	<code>rpl_semi_sync_master_wait_no_slave</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	ON

With semisynchronous replication, for each transaction, the master waits until timeout for acknowledgment of receipt from some semisynchronous slave. If no response occurs during this period, the master reverts to normal replication. This variable controls whether the master waits for the timeout to expire before reverting to normal replication even if the slave count drops to zero during the timeout period.

If the value is `ON` (the default), it is permissible for the slave count to drop to zero during the timeout period (for example, if slaves disconnect). The master still waits for the timeout, so as long as some slave reconnects and acknowledges the transaction within the timeout interval, semisynchronous replication continues.

If the value is `OFF`, the master reverts to normal replication if the slave count drops to zero during the timeout period.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `rpl_semi_sync_slave_enabled` [537]

<b>Variable Name</b>	<code>rpl_semi_sync_slave_enabled</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	OFF

Controls whether semisynchronous replication is enabled on the slave. To enable or disable the plugin, set this variable to `ON` or `OFF` (or 1 or 0), respectively. The default is `OFF`.



This variable is available only if the slave-side semisynchronous replication plugin is installed.

- [rpl\\_semi\\_sync\\_slave\\_trace\\_level \[538\]](#)

<b>Variable Name</b>	<code>rpl_semi_sync_slave_trace_level</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>32</code>

The semisynchronous replication debug trace level on the slave. See [rpl\\_semi\\_sync\\_master\\_trace\\_level \[536\]](#) for the permissible values.

This variable is available only if the slave-side semisynchronous replication plugin is installed.

- [secure\\_auth \[538\]](#)

<b>Command-Line Format</b>	<code>--secure-auth</code>	
<b>Option-File Format</b>	<code>secure-auth</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">secure_auth [538]</a>	
<b>Variable Name</b>	<code>secure_auth</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.4)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>
	<b>Permitted Values (&gt;= 5.6.5)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

If this variable is enabled, the server blocks connections by clients that attempt to use accounts that have passwords stored in the old (pre-4.1) format.

Enable this variable to prevent all use of passwords employing the old format (and hence insecure communication over the network). Before MySQL 5.6.5, this variable is disabled by default. As of MySQL 5.6.5, it is enabled by default.

Server startup fails with an error if this variable is enabled and the privilege tables are in pre-4.1 format. See [Section C.5.2.4, “Client does not support authentication protocol”](#).

- [secure\\_file\\_priv \[538\]](#)

<b>Command-Line Format</b>	<code>--secure-file-priv=path</code>	
<b>Option-File Format</b>	<code>secure-file-priv=path</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">secure_file_priv [538]</a>	
<b>Variable Name</b>	<code>secure-file-priv</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	

	<b>Type</b>	string
--	-------------	--------

By default, this variable is empty. If set to the name of a directory, it limits the effect of the `LOAD_FILE()` [1181] function and the `LOAD DATA` and `SELECT ... INTO OUTFILE` statements to work only with files in that directory.

- `server_id` [539]

<b>Command-Line Format</b>	<code>--server-id=#</code>	
<b>Option-File Format</b>	<code>server-id</code>	
<b>Option Sets Variable</b>	Yes, <code>server_id</code> [539]	
<b>Variable Name</b>	<code>server_id</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 4294967295

The server ID, used in replication to give each master and slave a unique identity. This variable is set by the `--server-id` [1955] option. For each server participating in replication, you should pick a positive integer in the range from 1 to  $2^{32} - 1$  to act as that server's ID.

- `sha256_password_private_key_path` [539]

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>sha256_password_private_key_path</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name
	<b>Default</b>	<code>private_key.pem</code>

The path name of the RSA private key file for the `sha256_password` authentication plugin. If the file is named as a relative path, it is interpreted relative to the server data directory. The file must be in PEM format. Because this file stores a private key, its access mode should be restricted so that only the MySQL server can read it.

For information about `sha256_password`, including instructions for creating the RSA key files, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

This variable is available only if MySQL was built using OpenSSL. It was added in MySQL 5.6.6.

- `sha256_password_public_key_path` [539]

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>sha256_password_public_key_path</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

	<b>Default</b>	<code>public_key.pem</code>
--	----------------	-----------------------------

The path name of the RSA public key file for the `sha256_password` authentication plugin. If the file is named as a relative path, it is interpreted relative to the server data directory. The file must be in PEM format. Because this file stores a public key, copies can be freely distributed to client users. (Clients that explicitly specify a public key when connecting to the server using RSA password encryption must use the same public key as that used by the server.)

For information about `sha256_password`, including instructions for creating the RSA key files and how clients specify the RSA public key, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

This variable is available only if MySQL was built using OpenSSL. It was added in MySQL 5.6.6.

- [shared\\_memory \[540\]](#)

<b>Variable Name</b>	<code>shared_memory</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Platform Specific</b>	windows

(Windows only.) Whether the server permits shared-memory connections.

- [shared\\_memory\\_base\\_name \[540\]](#)

<b>Variable Name</b>	<code>shared_memory_base_name</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Platform Specific</b>	windows

(Windows only.) The name of shared memory to use for shared-memory connections. This is useful when running multiple MySQL instances on a single physical machine. The default name is `MYSQL`. The name is case sensitive.

- [skip\\_external\\_locking \[540\]](#)

This is `OFF` if `mysqld` uses external locking, `ON` if external locking is disabled. This affects only `MyISAM` table access.

- [skip\\_name\\_resolve \[540\]](#)

This variable is set from the value of the `--skip-name-resolve [440]` option. If it is `ON`, `mysqld` resolves host names when checking client connections. If `OFF`, `mysqld` uses only IP numbers and all `Host` column values in the grant tables must be IP addresses or `localhost`. See [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#).

- [skip\\_networking \[540\]](#)

This is `ON` if the server permits only local (non-TCP/IP) connections. On Unix, local connections use a Unix socket file. On Windows, local connections use a named pipe or shared memory. This variable can be set to `ON` with the `--skip-networking [440]` option.

- [skip\\_show\\_database \[540\]](#)

This prevents people from using the `SHOW DATABASES` statement if they do not have the `SHOW DATABASES [756]` privilege. This can improve security if you have concerns about users being able to see databases belonging to other users. Its effect depends on the `SHOW DATABASES [756]` privilege: If the variable value is `ON`, the `SHOW DATABASES` statement is permitted only to users who have the `SHOW DATABASES [756]` privilege, and the statement displays all database names. If

the value is `OFF`, `SHOW DATABASES` is permitted to all users, but displays the names of only those databases for which the user has the `SHOW DATABASES` [756] or other privilege. (Note that *any* global privilege is considered a privilege for the database.)

- `slow_launch_time` [541]

<b>Command-Line Format</b>	<code>--slow_launch_time=#</code>	
<b>Option-File Format</b>	<code>slow_launch_time</code>	
<b>Option Sets Variable</b>	Yes, <code>slow_launch_time</code> [541]	
<b>Variable Name</b>	<code>slow_launch_time</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>2</code>

If creating a thread takes longer than this many seconds, the server increments the `slow_launch_threads` [598] status variable.

- `slow_query_log` [541]

<b>Command-Line Format</b>	<code>--slow-query-log</code>	
<b>Option-File Format</b>	<code>slow-query-log</code>	
<b>Option Sets Variable</b>	Yes, <code>slow_query_log</code> [541]	
<b>Variable Name</b>	<code>slow_query_log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether the slow query log is enabled. The value can be 0 (or `OFF`) to disable the log or 1 (or `ON`) to enable the log. The default value depends on whether the `--slow_query_log` [541] option is given. The destination for log output is controlled by the `log_output` [498] system variable; if that value is `NONE`, no log entries are written even if the log is enabled.

“Slow” is determined by the value of the `long_query_time` [500] variable. See Section 5.2.5, “The Slow Query Log”.

- `slow_query_log_file` [541]

<b>Command-Line Format</b>	<code>--slow-query-log-file=file_name</code>	
<b>Option-File Format</b>	<code>slow_query_log_file</code>	
<b>Option Sets Variable</b>	Yes, <code>slow_query_log_file</code> [541]	
<b>Variable Name</b>	<code>slow_query_log_file</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The name of the slow query log file. The default value is `host_name-slow.log`, but the initial value can be changed with the `--slow_query_log_file` option.

- [socket \[542\]](#)

<b>Command-Line Format</b>	<code>--socket=name</code>	
<b>Option-File Format</b>	<code>socket</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">socket [542]</a>	
<b>Variable Name</b>	<code>socket</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>/tmp/mysql.sock</code>

On Unix platforms, this variable is the name of the socket file that is used for local client connections. The default is `/tmp/mysql.sock`. (For some distribution formats, the directory might be different, such as `/var/lib/mysql` for RPMs.)

On Windows, this variable is the name of the named pipe that is used for local client connections. The default value is `MySQL` (not case sensitive).

- [sort\\_buffer\\_size \[542\]](#)

<b>Command-Line Format</b>	<code>--sort_buffer_size=#</code>	
<b>Option-File Format</b>	<code>sort_buffer_size</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">sort_buffer_size [542]</a>	
<b>Variable Name</b>	<code>sort_buffer_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.3)</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>2097144</code>
	<b>Range</b>	<code>32768 .. 18446744073709547520</code>
	<b>Permitted Values (&lt;= 5.6.3)</b>	
	<b>Platform Bit Size</b>	<code>32</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>2097144</code>
	<b>Range</b>	<code>32768 .. 4294967295</code>
	<b>Permitted Values (&gt;= 5.6.4)</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>262144</code>

	<b>Range</b>	32768 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.4)</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	262144
	<b>Range</b>	32768 .. 4294967295

Each session that needs to do a sort allocates a buffer of this size. `sort_buffer_size` [542] is not specific to any storage engine and applies in a general manner for optimization. See [Section 8.13.13, “ORDER BY Optimization”](#), for example.

If you see many `Sort_merge_passes` [599] per second in `SHOW GLOBAL STATUS` output, you can consider increasing the `sort_buffer_size` [542] value to speed up `ORDER BY` or `GROUP BY` operations that cannot be improved with query optimization or improved indexing.

As of MySQL 5.6.4, the optimizer tries to work out how much space is needed but can allocate more, up to the limit. Before MySQL 5.6.4, the optimizer allocates the entire buffer even if it is not all needed. In either case, setting it larger than required globally will slow down most queries that sort. It is best to increase it as a session setting, and only for the sessions that need a larger size. On Linux, there are thresholds of 256KB and 2MB where larger values may significantly slow down memory allocation, so you should consider staying below one of those values. Experiment to find the best value for your workload. See [Section C.5.4.4, “Where MySQL Stores Temporary Files”](#).

The maximum permissible setting for `sort_buffer_size` [542] is 4GB. Values larger than 4GB are permitted for 64-bit platforms (except 64-bit Windows, for which large values are truncated to 4GB with a warning).

- `sql_auto_is_null` [543]

<b>Variable Name</b>	<code>sql_auto_is_null</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	0

If this variable is set to 1, then after a statement that successfully inserts an automatically generated `AUTO_INCREMENT` value, you can find that value by issuing a statement of the following form:

```
SELECT * FROM tbl_name WHERE auto_col IS NULL
```

If the statement returns a row, the value returned is the same as if you invoked the `LAST_INSERT_ID()` [1269] function. For details, including the return value after a multiple-row insert, see [Section 12.14, “Information Functions”](#). If no `AUTO_INCREMENT` value was successfully inserted, the `SELECT` statement returns no row.

The behavior of retrieving an `AUTO_INCREMENT` value by using an `IS NULL` [1168] comparison is used by some ODBC programs, such as Access. See [Section 21.1.7.1.1, “Obtaining Auto-Increment Values”](#). This behavior can be disabled by setting `sql_auto_is_null` [543] to 0.

The default value of `sql_auto_is_null` [543] is 0 in MySQL 5.6.

- `sql_big_selects` [543]

<b>Variable Name</b>	<code>sql_big_selects</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>1</code>

If set to 0, MySQL aborts `SELECT` statements that are likely to take a very long time to execute (that is, statements for which the optimizer estimates that the number of examined rows exceeds the value of `max_join_size` [505]). This is useful when an inadvisable `WHERE` statement has been issued. The default value for a new connection is 1, which permits all `SELECT` statements.

If you set the `max_join_size` [505] system variable to a value other than `DEFAULT`, `sql_big_selects` [543] is set to 0.

- `sql_buffer_result` [544]

<b>Variable Name</b>	<code>sql_buffer_result</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>0</code>

If set to 1, `sql_buffer_result` [544] forces results from `SELECT` statements to be put into temporary tables. This helps MySQL free the table locks early and can be beneficial in cases where it takes a long time to send results to the client. The default value is 0.

- `sql_log_bin` [544]

<b>Variable Name</b>	<code>sql_log_bin</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>

This variable controls whether logging to the binary log is done. The default value is 1 (do logging). To change logging for the current session, change the session value of this variable. The session user must have the `SUPER` [756] privilege to set this variable.

In MySQL 5.6, it is not possible to set `@@session.sql_log_bin` within a transaction or subquery. (Bug #53437)

- `sql_log_off` [544]

<b>Variable Name</b>	<code>sql_log_off</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>0</code>

This variable controls whether logging to the general query log is done. The default value is 0 (do logging). To change logging for the current session, change the session value of this variable. The session user must have the `SUPER` [756] privilege to set this option. The default value is 0.

- `sql_mode` [545]

<b>Command-Line Format</b>	<code>--sql-mode=name</code>
<b>Option-File Format</b>	<code>sql-mode</code>
<b>Option Sets Variable</b>	Yes, <code>sql_mode</code> [545]
<b>Variable Name</b>	<code>sql_mode</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
	<b>Type</b> <code>set</code>
	<b>Default</b> <code>''</code>
	<b>Valid Values</b> <code>ALLOW_INVALID_DATES</code> <code>ANSI_QUOTES</code> <code>ERROR_FOR_DIVISION_BY_ZERO</code> <code>HIGH_NOT_PRECEDENCE</code> <code>IGNORE_SPACE</code> <code>NO_AUTO_CREATE_USER</code> <code>NO_AUTO_VALUE_ON_ZERO</code> <code>NO_BACKSLASH_ESCAPES</code> <code>NO_DIR_IN_CREATE</code> <code>NO_ENGINE_SUBSTITUTION</code> <code>NO_FIELD_OPTIONS</code> <code>NO_KEY_OPTIONS</code> <code>NO_TABLE_OPTIONS</code> <code>NO_UNSIGNED_SUBTRACTION</code> <code>NO_ZERO_DATE</code> <code>NO_ZERO_IN_DATE</code> <code>ONLY_FULL_GROUP_BY</code> <code>PAD_CHAR_TO_FULL_LENGTH</code> <code>PIPES_AS_CONCAT</code> <code>REAL_AS_FLOAT</code> <code>STRICT_ALL_TABLES</code>



	STRICT_TRANS_TABLES
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Type</b>	set
<b>Default</b>	NO_ENGINE_SUBSTITUTION
<b>Valid Values</b>	ALLOW_INVALID_DATES ANSI_QUOTES ERROR_FOR_DIVISION_BY_ZERO HIGH_NOT_PRECEDENCE IGNORE_SPACE NO_AUTO_CREATE_USER NO_AUTO_VALUE_ON_ZERO NO_BACKSLASH_ESCAPES NO_DIR_IN_CREATE NO_ENGINE_SUBSTITUTION NO_FIELD_OPTIONS NO_KEY_OPTIONS NO_TABLE_OPTIONS NO_UNSIGNED_SUBTRACTION NO_ZERO_DATE NO_ZERO_IN_DATE ONLY_FULL_GROUP_BY PAD_CHAR_TO_FULL_LENGTH PIPES_AS_CONCAT REAL_AS_FLOAT STRICT_ALL_TABLES STRICT_TRANS_TABLES

The current server SQL mode, which can be set dynamically. The default as of MySQL 5.6.6 is `NO_ENGINE_SUBSTITUTION`; previously it was an empty string. See [Section 5.1.7, “Server SQL Modes”](#).

- [sql\\_notes \[546\]](#)

If set to 1 (the default), warnings of `Note` level increment `warning_count` and the server records them. If set to 0, `Note` warnings do not increment `warning_count` [564] and the server does not record them. `mysqldump` includes output to set this variable to 0 so that reloading the dump file does not produce warnings for events that do not affect the integrity of the reload operation.

- [sql\\_quote\\_show\\_create \[546\]](#)

If set to 1 (the default), the server quotes identifiers for `SHOW CREATE TABLE` and `SHOW CREATE DATABASE` statements. If set to 0, quoting is disabled. This option is enabled by default so that replication works for identifiers that require quoting. See [Section 13.7.5.12, “SHOW CREATE TABLE Syntax”](#), and [Section 13.7.5.8, “SHOW CREATE DATABASE Syntax”](#).

- `sql_safe_updates` [\[547\]](#)

If set to 1, MySQL aborts `UPDATE` or `DELETE` statements that do not use a key in the `WHERE` clause or a `LIMIT` clause. This makes it possible to catch `UPDATE` or `DELETE` statements where keys are not used properly and that would probably change or delete a large number of rows. The default value is 0.

- `sql_select_limit` [\[547\]](#)

<b>Variable Name</b>	<code>sql_select_limit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>

The maximum number of rows to return from `SELECT` statements. The default value for a new connection is the maximum number of rows that the server permits per table. Typical default values are  $(2^{32})-1$  or  $(2^{64})-1$ . If you have changed the limit, the default value can be restored by assigning a value of `DEFAULT`.

If a `SELECT` has a `LIMIT` clause, the `LIMIT` takes precedence over the value of `sql_select_limit` [\[547\]](#).

- `sql_warnings` [\[547\]](#)

This variable controls whether single-row `INSERT` statements produce an information string if warnings occur. The default is 0. Set the value to 1 to produce an information string.

- `ssl_ca` [\[547\]](#)

<b>Command-Line Format</b>	<code>--ssl-ca=name</code>	
<b>Option-File Format</b>	<code>ssl-ca</code>	
<b>Option Sets Variable</b>	Yes, <code>ssl_ca</code> <a href="#">[547]</a>	
<b>Variable Name</b>	<code>ssl-ca</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The path to a file with a list of trusted SSL CAs.

- `ssl_capath` [\[547\]](#)

<b>Command-Line Format</b>	<code>--ssl-capath=name</code>	
<b>Option-File Format</b>	<code>ssl-capath</code>	
<b>Option Sets Variable</b>	Yes, <code>ssl_capath</code> <a href="#">[547]</a>	
<b>Variable Name</b>	<code>ssl-capath</code>	
<b>Variable Scope</b>	Global	

<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	directory name

The path to a directory that contains trusted SSL CA certificates in PEM format.

- [ssl\\_cert \[548\]](#)

<b>Command-Line Format</b>	<code>--ssl-cert=name</code>	
<b>Option-File Format</b>	<code>ssl-cert</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">ssl_cert [548]</a>	
<b>Variable Name</b>	<code>ssl-cert</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The name of the SSL certificate file to use for establishing a secure connection.

- [ssl\\_cipher \[548\]](#)

<b>Command-Line Format</b>	<code>--ssl-cipher=name</code>	
<b>Option-File Format</b>	<code>ssl-cipher</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">ssl_cipher [548]</a>	
<b>Variable Name</b>	<code>ssl-cipher</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

A list of permissible ciphers to use for SSL encryption.

- [ssl\\_crl \[548\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--ssl-crl=name</code>	
<b>Option-File Format</b>	<code>ssl-crl</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">ssl_crl [548]</a>	
<b>Variable Name</b>	<code>ssl-crl</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The path to a file containing certificate revocation lists in PEM format. Revocation lists work for MySQL distributions compiled against OpenSSL (but not yaSSL).

This variable was added in MySQL 5.6.3.

- [ssl\\_crlpath \[548\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--ssl-crlpath=name</code>	
<b>Option-File Format</b>	<code>ssl-crlpath</code>	
<b>Option Sets Variable</b>	Yes, <code>ssl_crlpath</code> [548]	
<b>Variable Name</b>	<code>ssl-crlpath</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The path to a directory that contains files containing certificate revocation lists in PEM format. Revocation lists work for MySQL distributions compiled against OpenSSL (but not yaSSL).

This variable was added in MySQL 5.6.3.

- `ssl_key` [549]

<b>Command-Line Format</b>	<code>--ssl-key=name</code>	
<b>Option-File Format</b>	<code>ssl-key</code>	
<b>Option Sets Variable</b>	Yes, <code>ssl_key</code> [549]	
<b>Variable Name</b>	<code>ssl-key</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The name of the SSL key file to use for establishing a secure connection.

- `storage_engine` [549]

<b>Variable Name</b>	<code>storage_engine</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>InnoDB</code>

The default storage engine (table type). To set the storage engine at server startup, use the `--default-storage-engine` [420] option. See Section 5.1.3, “Server Command Options”.

This variable is deprecated. Use `default_storage_engine` [474] instead.

- `stored_program_cache` [549]

<b>Version Introduced</b>	5.6.5	
<b>Command-Line Format</b>	<code>--stored-program-cache=#</code>	
<b>Option-File Format</b>	<code>stored_program_cache</code>	
<b>Variable Name</b>	<code>stored_program_cache</code>	
<b>Variable Scope</b>	Global	

<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	256
	<b>Range</b>	256 .. 524288

Sets a soft upper limit for the number of cached stored routines per connection. The value of this variable is specified in terms of the number of stored routines held in each of the two caches maintained by the MySQL Server for, respectively, stored procedures and stored functions.

Whenever a stored routine is executed this cache size is checked before the first or top-level statement in the routine is parsed; if the number of routines of the same type (stored procedures or stored functions according to which is being executed) exceeds the limit specified by this variable, the corresponding cache is flushed and memory previously allocated for cached objects is freed. This allows the cache to be flushed safely, even when there are dependencies between stored routines.

- [sync\\_frm \[550\]](#)

<b>Command-Line Format</b>	<code>--sync_frm</code>	
<b>Option-File Format</b>	<code>sync_frm</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">sync_frm [550]</a>	
<b>Variable Name</b>	<code>sync_frm</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	TRUE

If this variable is set to 1, when any nontemporary table is created its `.frm` file is synchronized to disk (using `fdatasync()`). This is slower but safer in case of a crash. The default is 1.

- [system\\_time\\_zone \[550\]](#)

<b>Variable Name</b>	<code>system_time_zone</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	string

The server system time zone. When the server begins executing, it inherits a time zone setting from the machine defaults, possibly modified by the environment of the account used for running the server or the startup script. The value is used to set [system\\_time\\_zone \[550\]](#). Typically the time zone is specified by the `TZ` environment variable. It also can be specified using the `--timezone [229]` option of the `mysqld_safe` script.

The [system\\_time\\_zone \[550\]](#) variable differs from [time\\_zone \[557\]](#). Although they might have the same value, the latter variable is used to initialize the time zone for each client that connects. See [Section 10.6, “MySQL Server Time Zone Support”](#).

- [table\\_definition\\_cache \[550\]](#)

<b>Command-Line Format</b>	<code>--table_definition_cache=#</code>	
<b>Option-File Format</b>	<code>table_definition_cache</code>	
<b>Option Sets Variable</b>	Yes, <code>table_definition_cache</code> [550]	
<b>Variable Name</b>	<code>table_definition_cache</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.7)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>400</code>
	<b>Range</b>	<code>400 .. 524288</code>
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>
	<b>Range</b>	<code>400 .. 524288</code>

The number of table definitions (from `.frm` files) that can be stored in the definition cache. If you use a large number of tables, you can create a large table definition cache to speed up opening of tables. The table definition cache takes less space and does not use file descriptors, unlike the normal table cache. The minimum value is 400. The default value is based on the following formula, capped to a limit of 2000:

```
400 + (table_open_cache / 2)
```

Before MySQL 5.6.8, the default is 400.

- `table_open_cache` [551]

<b>Command-Line Format</b>	<code>--table-open-cache=#</code>	
<b>Option-File Format</b>	<code>table_open_cache</code>	
<b>Variable Name</b>	<code>table_open_cache</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.7)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>400</code>
	<b>Range</b>	<code>400 .. 524288</code>
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>2000</code>
	<b>Range</b>	<code>400 .. 524288</code>

The number of open tables for all threads. Increasing this value increases the number of file descriptors that `mysqld` requires. You can check whether you need to increase the table cache by checking the `Opened_tables` [595] status variable. See Section 5.1.6, “Server Status Variables”. If the value of `Opened_tables` [595] is large and you do not use `FLUSH TABLES` often (which just forces all tables to be closed and reopened), then you should increase the value of the `table_open_cache` [551] variable. For more information about the table cache, see Section 8.4.3.1, “How MySQL Opens and Closes Tables”.

- [table\\_open\\_cache\\_instances](#) [552]

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--table-open-cache_instances=#</code>	
<b>Option-File Format</b>	<code>table_open_cache_instances</code>	
<b>Variable Name</b>	<code>table_open_cache_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>

The number of open tables cache instances (default 1). To improve scalability by reducing contention among sessions, the open tables cache can be partitioned into several smaller cache instances of size `table_open_cache` [551] / `table_open_cache_instances` [552]. A session need lock only one instance to access it for DML statements. This segments cache access among instances, permitting higher performance for operations that need to use the cache when many there are many sessions accessing tables. (DDL statements still require a lock on the entire cache, but such statements are much less frequent than DML statements.)

A value of 8 or 16 is recommended on systems that routinely use 16 or more cores.

This variable was added in MySQL 5.6.6.

- [thread\\_cache\\_size](#) [552]

<b>Command-Line Format</b>	<code>--thread_cache_size=#</code>	
<b>Option-File Format</b>	<code>thread_cache_size</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_cache_size</code> [552]	
<b>Variable Name</b>	<code>thread_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.7)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 16384</code>
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>
	<b>Range</b>	<code>0 .. 16384</code>

How many threads the server should cache for reuse. When a client disconnects, the client's threads are put in the cache if there are fewer than `thread_cache_size` [552] threads there. Requests for threads are satisfied by reusing threads taken from the cache if possible, and only when the cache is empty is a new thread created. This variable can be increased to improve performance if you have a lot of new connections. Normally, this does not provide a notable performance improvement if you have a good thread implementation. However, if your server sees hundreds of connections per second you should normally set `thread_cache_size` [552] high enough so that most new connections use cached threads. By examining the difference between the

`Connections` [587] and `Threads_created` [601] status variables, you can see how efficient the thread cache is. For details, see [Section 5.1.6, “Server Status Variables”](#).

The default value is based on the following formula, capped to a limit of 100:

```
8 + (max_connections / 100)
```

Before MySQL 5.6.8, the default is 0.

- `thread_concurrency` [553]

<b>Version Deprecated</b>	5.6.1	
<b>Command-Line Format</b>	<code>--thread_concurrency=#</code>	
<b>Option-File Format</b>	<code>thread_concurrency</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_concurrency</code> [553]	
<b>Variable Name</b>	<code>thread_concurrency</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Deprecated</b>	5.6.1	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10</code>
	<b>Range</b>	<code>1 .. 512</code>

This variable is specific to Solaris systems, for which `mysqld` invokes the `thr_setconcurrency()` with the variable value. This function enables applications to give the threads system a hint about the desired number of threads that should be run at the same time.

This variable is deprecated as of MySQL 5.6.1 and will be removed in a future MySQL release.

- `thread_handling` [553]

<b>Command-Line Format</b>	<code>--thread_handling=name</code>	
<b>Option-File Format</b>	<code>thread_handling=name</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_handling</code> [553]	
<b>Variable Name</b>	<code>thread_handling</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Valid Values</b>	<code>no-threads</code> <code>one-thread-per-connection</code> <code>dynamically-loaded</code>

The thread-handling model used by the server for connection threads. The permissible values are `no-threads` (the server uses a single thread) and `one-thread-per-connection` (the server uses one thread to handle each client connection). `no-threads` is useful for debugging under Linux; see [MySQL Internals: Porting to Other Systems](#).



If the thread pool plugin is enabled, the server sets the `thread_handling` [553] value to `dynamically-loaded`. See Section 8.11.6.1, “Thread Pool Components and Installation”.

- `thread_pool_algorithm` [554]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_algorithm=#</code>	
<b>Option-File Format</b>	<code>thread_pool_algorithm=#</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_pool_algorithm</code> [554]	
<b>Variable Name</b>	<code>thread_pool_algorithm</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1</code>

This variable controls which algorithm the thread pool plugin uses:

- A value of 0 (the default) uses a conservative low-concurrency algorithm which is most well tested and is known to produce very good results.
- A value of 1 increases the concurrency and uses a more aggressive algorithm which at times has been known to perform 5–10% better on optimal thread counts, but has degrading performance as the number of connections increases. Its use should be considered as experimental and not supported.

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See Section 8.11.6, “The Thread Pool Plugin”

- `thread_pool_high_priority_connection` [554]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_high_priority_connection=#</code>	
<b>Option-File Format</b>	<code>thread_pool_high_priority_connection=#</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_pool_high_priority_connection</code> [554]	
<b>Variable Name</b>	<code>thread_pool_high_priority_connection</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1</code>

This variable affects queuing of new statements prior to execution. If the value is 0 (false, the default), statement queuing uses both the low-priority and high-priority queues. If the value is 1 (true), queued statements always go to the high-priority queue.

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See Section 8.11.6, “The Thread Pool Plugin”

- `thread_pool_max_unused_threads` [555]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_max_unused_threads=#</code>	
<b>Option-File Format</b>	<code>thread_pool_max_unused_threads=#</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_pool_max_unused_threads</code> [555]	
<b>Variable Name</b>	<code>thread_pool_max_unused_threads</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 4096

The maximum permitted number of unused threads in the thread pool. This variable makes it possible to limit the amount of memory used by sleeping threads.

A value of 0 (the default) means no limit on the number of sleeping threads. A value of  $N$  where  $N$  is greater than 0 means 1 consumer thread and  $N-1$  reserve threads. In this case, if a thread is ready to sleep but the number of sleeping threads is already at the maximum, the thread exits rather than going to sleep.

A sleeping thread is either sleeping as a consumer thread or a reserve thread. The thread pool permits one thread to be the consumer thread when sleeping. If a thread goes to sleep and there is no existing consumer thread, it will sleep as a consumer thread. When a thread must be woken up, a consumer thread is selected if there is one. A reserve thread is selected only when there is no consumer thread to wake up.

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See [Section 8.11.6, “The Thread Pool Plugin”](#)

- `thread_pool_prio_kickup_timer` [555]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_prio_kickup_timer=#</code>	
<b>Option-File Format</b>	<code>thread_pool_prio_kickup_timer=#</code>	
<b>Option Sets Variable</b>	Yes, <code>thread_pool_prio_kickup_timer</code> [555]	
<b>Variable Name</b>	<code>thread_pool_prio_kickup_timer</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1000
	<b>Range</b>	0 .. 4294967294

This variable affects statements waiting for execution in the low-priority queue. The value is the number of milliseconds before a waiting statement is moved to the high-priority queue. The default is 1000 (1 second). The range of values is 0 to  $2^{32} - 2$ .

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See [Section 8.11.6, “The Thread Pool Plugin”](#)

- [thread\\_pool\\_size \[556\]](#)

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_size=#</code>	
<b>Option-File Format</b>	<code>thread_pool_size=#</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">thread_pool_size [556]</a>	
<b>Variable Name</b>	<code>thread_pool_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>16</code>
	<b>Range</b>	<code>1 .. 64</code>

The number of thread groups in the thread pool. This is the most important parameter controlling thread pool performance. It affects how many statements can execute simultaneously. The default value is 16, with a range from 1 to 64 of permissible values. If a value outside this range is specified, the thread pool plugin does not load and the server writes a message to the error log.

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See [Section 8.11.6, “The Thread Pool Plugin”](#)

- [thread\\_pool\\_stall\\_limit \[556\]](#)

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--thread_pool_stall_limit=#</code>	
<b>Option-File Format</b>	<code>thread_pool_stall_limit=#</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">thread_pool_stall_limit [556]</a>	
<b>Variable Name</b>	<code>thread_pool_stall_limit</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>6</code>
	<b>Range</b>	<code>4 .. 600</code>

This variable affects executing statements. The value is the amount of time a statement has to finish after starting to execute before it becomes defined as stalled, at which point the thread pool permits the thread group to begin executing another statement. The value is measured in 10 millisecond units, so a value of 6 (the default) means 60ms. The range of values is 4 to 600 (40ms to 6s). Short wait values permit threads to start more quickly. Short values are also better for avoiding deadlock situations. Long wait values are useful for workloads that include long-running statements, to avoid starting too many new statements while the current ones execute.

This variable was added in MySQL 5.6.10. It is available only if the thread pool plugin is enabled. See [Section 8.11.6, “The Thread Pool Plugin”](#)

- [thread\\_stack \[556\]](#)

<b>Command-Line Format</b>	<code>--thread_stack=#</code>
----------------------------	-------------------------------

<b>Option-File Format</b>	<code>thread_stack</code>
<b>Option Sets Variable</b>	Yes, <code>thread_stack</code> [556]
<b>Variable Name</b>	<code>thread_stack</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	<code>numeric</code>
<b>Default</b>	196608
<b>Range</b>	131072 .. 4294967295
<b>Block Size</b>	1024
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	<code>numeric</code>
<b>Default</b>	262144
<b>Range</b>	131072 .. 18446744073709547520
<b>Block Size</b>	1024

The stack size for each thread. Many of the limits detected by the `crash-me` test are dependent on this value. See [Section 8.12.2, “The MySQL Benchmark Suite”](#). The default of 192KB (256KB for 64-bit systems) is large enough for normal operation. If the thread stack size is too small, it limits the complexity of the SQL statements that the server can handle, the recursion depth of stored procedures, and other memory-consuming actions.

- `time_format` [557]

This variable is unused. It is deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release.

- `time_zone` [557]

<b>Command-Line Format</b>	<code>--default_time_zone=string</code>
<b>Option-File Format</b>	<code>default_time_zone</code>
<b>Variable Name</b>	<code>time_zone</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>

The current time zone. This variable is used to initialize the time zone for each client that connects. By default, the initial value of this is `'SYSTEM'` (which means, “use the value of `system_time_zone` [550]”). The value can be specified explicitly at server startup with the `--default-time-zone` [420] option. See [Section 10.6, “MySQL Server Time Zone Support”](#).

- `timed_mutexes` [557]

<b>Command-Line Format</b>	<code>--timed_mutexes</code>	
<b>Option-File Format</b>	<code>timed_mutexes</code>	
<b>Option Sets Variable</b>	Yes, <code>timed_mutexes</code> [557]	
<b>Variable Name</b>	<code>timed_mutexes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

This variable controls whether `InnoDB` mutexes are timed. If this variable is set to 0 or `OFF` (the default), mutex timing is disabled. If the variable is set to 1 or `ON`, mutex timing is enabled. With timing enabled, the `os_wait_times` value in the output from `SHOW ENGINE INNODB MUTEX` indicates the amount of time (in ms) spent in operating system waits. Otherwise, the value is 0.

- `timestamp = {timestamp_value | DEFAULT}` [558]

Set the time for this client. This is used to get the original timestamp if you use the binary log to restore rows. `timestamp_value` should be a Unix epoch timestamp, not a MySQL timestamp.

As of MySQL 5.6.4, `timestamp` [558] is a `DOUBLE` rather than `BIGINT` because its value includes a microseconds part.

`SET timestamp` affects the value returned by `NOW()` [1220] but not by `SYSDATE()` [1223]. This means that timestamp settings in the binary log have no effect on invocations of `SYSDATE()` [1223]. The server can be started with the `--sysdate-is-now` [444] option to cause `SYSDATE()` [1223] to be an alias for `NOW()` [1220], in which case `SET timestamp` affects both functions.

- `tmp_table_size` [558]

<b>Command-Line Format</b>	<code>--tmp_table_size=#</code>	
<b>Option-File Format</b>	<code>tmp_table_size</code>	
<b>Option Sets Variable</b>	Yes, <code>tmp_table_size</code> [558]	
<b>Variable Name</b>	<code>tmp_table_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>system dependent</code>
	<b>Range</b>	<code>1024 .. 4294967295</code>

The maximum size of internal in-memory temporary tables. (The actual limit is determined as the minimum of `tmp_table_size` [558] and `max_heap_table_size` [504].) If an in-memory temporary table exceeds the limit, MySQL automatically converts it to an on-disk `MyISAM` table. Increase the value of `tmp_table_size` [558] (and `max_heap_table_size` [504] if necessary) if you do many advanced `GROUP BY` queries and you have lots of memory. This variable does not apply to user-created `MEMORY` tables.

You can compare the number of internal on-disk temporary tables created to the total number of internal temporary tables created by comparing the values of the `Created_tmp_disk_tables` [587] and `Created_tmp_tables` [588] variables.

See also [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).

- `tmpdir` [\[559\]](#)

<b>Command-Line Format</b>	<code>--tmpdir=path</code>	
	<code>-t</code>	
<b>Option-File Format</b>	<code>tmpdir</code>	
<b>Option Sets Variable</b>	Yes, <code>tmpdir</code> <a href="#">[559]</a>	
<b>Variable Name</b>	<code>tmpdir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The directory used for temporary files and temporary tables. This variable can be set to a list of several paths that are used in round-robin fashion. Paths should be separated by colon characters (“:”) on Unix and semicolon characters (“;”) on Windows.

The multiple-directory feature can be used to spread the load between several physical disks. If the MySQL server is acting as a replication slave, you should not set `tmpdir` [\[559\]](#) to point to a directory on a memory-based file system or to a directory that is cleared when the server host restarts. A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails. You can set the slave's temporary directory using the `slave_load_tmpdir` [\[1992\]](#) variable. In that case, the slave will not use the general `tmpdir` [\[559\]](#) value and you can set `tmpdir` [\[559\]](#) to a nonpermanent location.

- `transaction_alloc_block_size` [\[559\]](#)

<b>Command-Line Format</b>	<code>--transaction_alloc_block_size=#</code>	
<b>Option-File Format</b>	<code>transaction_alloc_block_size</code>	
<b>Option Sets Variable</b>	Yes, <code>transaction_alloc_block_size</code> <a href="#">[559]</a>	
<b>Variable Name</b>	<code>transaction_alloc_block_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	8192
	<b>Range</b>	1024 .. 4294967295
	<b>Block Size</b>	1024
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	8192

	<b>Range</b>	1024 .. 18446744073709547520
	<b>Block Size</b>	1024

The amount in bytes by which to increase a per-transaction memory pool which needs memory. See the description of `transaction_prealloc_size` [560].

- `transaction_prealloc_size` [560]

<b>Command-Line Format</b>	<code>--transaction_prealloc_size=#</code>	
<b>Option-File Format</b>	<code>transaction_prealloc_size</code>	
<b>Option Sets Variable</b>	Yes, <code>transaction_prealloc_size</code> [560]	
<b>Variable Name</b>	<code>transaction_prealloc_size</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	4096
	<b>Range</b>	1024 .. 4294967295
	<b>Block Size</b>	1024
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	4096
	<b>Range</b>	1024 .. 18446744073709547520
	<b>Block Size</b>	1024

There is a per-transaction memory pool from which various transaction-related allocations take memory. The initial size of the pool in bytes is `transaction_prealloc_size` [560]. For every allocation that cannot be satisfied from the pool because it has insufficient memory available, the pool is increased by `transaction_alloc_block_size` [559] bytes. When the transaction ends, the pool is truncated to `transaction_prealloc_size` [560] bytes.

By making `transaction_prealloc_size` [560] sufficiently large to contain all statements within a single transaction, you can avoid many `malloc()` calls.

- `tx_isolation` [560]

<b>Variable Name</b>	<code>tx_isolation</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	REPEATABLE-READ

	<b>Valid Values</b>	<a href="#">READ-UNCOMMITTED</a> <a href="#">READ-COMMITTED</a> <a href="#">REPEATABLE-READ</a> <a href="#">SERIALIZABLE</a>
--	---------------------	---

The default transaction isolation level. Defaults to [REPEATABLE-READ](#) [1469].

This variable can be set directly, or indirectly using the `SET TRANSACTION` statement. See [Section 13.3.6, “SET TRANSACTION Syntax”](#). If you set `tx_isolation` [560] directly to an isolation level name that contains a space, the name should be enclosed within quotation marks, with the space replaced by a dash. For example:

```
SET tx_isolation = 'READ-COMMITTED';
```

Any unique prefix of a valid value may be used to set the value of this variable.

The default transaction isolation level can also be set at startup using the `--transaction-isolation` [445] server option.

- [tx\\_read\\_only](#) [561]

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<a href="#">tx_read_only</a>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">boolean</a>
	<b>Default</b>	<a href="#">OFF</a>

The default transaction access mode. The value can be [OFF](#) (read/write, the default) or [ON](#) (read only).

This variable can be set directly, or indirectly using the `SET TRANSACTION` statement. See [Section 13.3.6, “SET TRANSACTION Syntax”](#).

To set the default transaction access mode at startup, use the `--transaction-read-only` [445] server option.

This variable was added in MySQL 5.6.5.

- [unique\\_checks](#) [561]

<b>Variable Name</b>	<a href="#">unique_checks</a>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">boolean</a>
	<b>Default</b>	<a href="#">1</a>

If set to 1 (the default), uniqueness checks for secondary indexes in [InnoDB](#) tables are performed. If set to 0, storage engines are permitted to assume that duplicate keys are not present in input data. If you know for certain that your data does not contain uniqueness violations, you can set this to 0 to speed up large table imports to [InnoDB](#).



Note that setting this variable to 0 does not *require* storage engines to ignore duplicate keys. An engine is still permitted to check for them and issue duplicate-key errors if it detects them.

- [updatable\\_views\\_with\\_limit](#) [562]

<b>Command-Line Format</b>	<code>--updatable_views_with_limit=#</code>	
<b>Option-File Format</b>	<code>updatable_views_with_limit</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">updatable_views_with_limit</a> [562]	
<b>Variable Name</b>	<code>updatable_views_with_limit</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>1</code>

This variable controls whether updates to a view can be made when the view does not contain all columns of the primary key defined in the underlying table, if the update statement contains a `LIMIT` clause. (Such updates often are generated by GUI tools.) An update is an `UPDATE` or `DELETE` statement. Primary key here means a `PRIMARY KEY`, or a `UNIQUE` index in which no column can contain `NULL`.

The variable can have two values:

- `1` or `YES`: Issue a warning only (not an error message). This is the default value.
  - `0` or `NO`: Prohibit the update.
- [validate\\_password\\_xxx](#)

The `validate_password` plugin implements a set of system variables having names of the form `validate_password_xxx`. These variables affect password testing by that plugin; see [Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”](#).

- [validate\\_user\\_plugins](#) [562]

<b>Version Introduced</b>	5.6.11	
<b>Variable Name</b>	<code>validate_user_plugins</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

If this variable is enabled (the default), the server checks each user account and produces a warning if conditions are found that would make the account unusable:

- The account requires an authentication plugin that is not loaded.
- The account requires the `sha256_password` authentication plugin but the server was started with neither SSL nor RSA enabled as required by this plugin.

Enabling `validate_user_plugins` slows down server initialization and `FLUSH PRIVILEGES`. If you do not require the additional checking, you can disable this variable at startup to avoid the performance decrement.

This variable was added in MySQL 5.6.11.

- `version` [563]

The version number for the server. The value might also include a suffix indicating server build or configuration information. `-log` indicates that one or more of the general log, slow query log, or binary log are enabled. `-debug` indicates that the server was built with debugging support enabled.

- `version_comment` [563]

<b>Variable Name</b>	<code>version_comment</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The CMake configuration program has a `COMPILATION_COMMENT` [134] option that permits a comment to be specified when building MySQL. This variable contains the value of that comment. See Section 2.9.4, “MySQL Source-Configuration Options”.

- `version_compile_machine` [563]

<b>Variable Name</b>	<code>version_compile_machine</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The type of the server binary.

- `version_compile_os` [563]

<b>Variable Name</b>	<code>version_compile_os</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The type of operating system on which MySQL was built.

- `wait_timeout` [563]

<b>Command-Line Format</b>	<code>--wait_timeout=#</code>	
<b>Option-File Format</b>	<code>wait_timeout</code>	
<b>Option Sets Variable</b>	Yes, <code>wait_timeout</code> [563]	
<b>Variable Name</b>	<code>wait_timeout</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>28800</code>

	<b>Range</b>	1 .. 31536000
	<b>Permitted Values</b>	
	<b>Type</b> (windows)	numeric
	<b>Default</b>	28800
	<b>Range</b>	1 .. 2147483

The number of seconds the server waits for activity on a noninteractive connection before closing it.

On thread startup, the session `wait_timeout` [563] value is initialized from the global `wait_timeout` [563] value or from the global `interactive_timeout` [490] value, depending on the type of client (as defined by the `CLIENT_INTERACTIVE` connect option to `mysql_real_connect()`). See also `interactive_timeout` [490].

- `warning_count` [564]

The number of errors, warnings, and notes that resulted from the last statement that generated messages. This variable is read only. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

## 5.1.5. Using System Variables

The MySQL server maintains many system variables that indicate how it is configured. [Section 5.1.4, “Server System Variables”](#), describes the meaning of these variables. Each system variable has a default value. System variables can be set at server startup using options on the command line or in an option file. Most of them can be changed dynamically while the server is running by means of the `SET` statement, which enables you to modify operation of the server without having to stop and restart it. You can refer to system variable values in expressions.

The server maintains two kinds of system variables. Global variables affect the overall operation of the server. Session variables affect its operation for individual client connections. A given system variable can have both a global and a session value. Global and session system variables are related as follows:

- When the server starts, it initializes all global variables to their default values. These defaults can be changed by options specified on the command line or in an option file. (See [Section 4.2.3, “Specifying Program Options”](#).)
- The server also maintains a set of session variables for each client that connects. The client's session variables are initialized at connect time using the current values of the corresponding global variables. For example, the client's SQL mode is controlled by the session `sql_mode` [545] value, which is initialized when the client connects to the value of the global `sql_mode` [545] value.

System variable values can be set globally at server startup by using options on the command line or in an option file. When you use a startup option to set a variable that takes a numeric value, the value can be given with a suffix of `K`, `M`, or `G` (either uppercase or lowercase) to indicate a multiplier of 1024, 1024<sup>2</sup> or 1024<sup>3</sup>; that is, units of kilobytes, megabytes, or gigabytes, respectively. Thus, the following command starts the server with a query cache size of 16 megabytes and a maximum packet size of one gigabyte:

```
mysqld --query_cache_size=16M --max_allowed_packet=1G
```

Within an option file, those variables are set like this:

```
[mysqld]
query_cache_size=16M
max_allowed_packet=1G
```

The lettercase of suffix letters does not matter; `16M` and `16m` are equivalent, as are `1G` and `1g`.

If you want to restrict the maximum value to which a system variable can be set at runtime with the `SET` statement, you can specify this maximum by using an option of the form

`--maximum-var_name=value` at server startup. For example, to prevent the value of `query_cache_size` [528] from being increased to more than 32MB at runtime, use the option `--maximum-query_cache_size=32M`.

Many system variables are dynamic and can be changed while the server runs by using the `SET` statement. For a list, see [Section 5.1.5.2, “Dynamic System Variables”](#). To change a system variable with `SET`, refer to it as `var_name`, optionally preceded by a modifier:

- To indicate explicitly that a variable is a global variable, precede its name by `GLOBAL` or `@@global..`. The `SUPER` [756] privilege is required to set global variables.
- To indicate explicitly that a variable is a session variable, precede its name by `SESSION`, `@@session.`, or `@@`. Setting a session variable requires no special privilege, but a client can change only its own session variables, not those of any other client.
- `LOCAL` and `@@local.` are synonyms for `SESSION` and `@@session..`
- If no modifier is present, `SET` changes the session variable.

A `SET` statement can contain multiple variable assignments, separated by commas. If you set several system variables, the most recent `GLOBAL` or `SESSION` modifier in the statement is used for following variables that have no modifier specified.

Examples:

```
SET sort_buffer_size=10000;
SET @@local.sort_buffer_size=10000;
SET GLOBAL sort_buffer_size=1000000, SESSION sort_buffer_size=1000000;
SET @@sort_buffer_size=1000000;
SET @@global.sort_buffer_size=1000000, @@local.sort_buffer_size=1000000;
```

The `@@var_name` syntax for system variables is supported for compatibility with some other database systems.

If you change a session system variable, the value remains in effect until your session ends or until you change the variable to a different value. The change is not visible to other clients.

If you change a global system variable, the value is remembered and used for new connections until the server restarts. (To make a global system variable setting permanent, you should set it in an option file.) The change is visible to any client that accesses that global variable. However, the change affects the corresponding session variable only for clients that connect after the change. The global variable change does not affect the session variable for any client that is currently connected (not even that of the client that issues the `SET GLOBAL` statement).

To prevent incorrect usage, MySQL produces an error if you use `SET GLOBAL` with a variable that can only be used with `SET SESSION` or if you do not specify `GLOBAL` (or `@@global..`) when setting a global variable.

To set a `SESSION` variable to the `GLOBAL` value or a `GLOBAL` value to the compiled-in MySQL default value, use the `DEFAULT` keyword. For example, the following two statements are identical in setting the session value of `max_join_size` [505] to the global value:

```
SET max_join_size=DEFAULT;
SET @@session.max_join_size=@@global.max_join_size;
```

Not all system variables can be set to `DEFAULT`. In such cases, use of `DEFAULT` results in an error.

You can refer to the values of specific global or session system variables in expressions by using one of the `@@`-modifiers. For example, you can retrieve values in a `SELECT` statement like this:

```
SELECT @@global.sql_mode, @@session.sql_mode, @@sql_mode;
```

When you refer to a system variable in an expression as `@@var_name` (that is, when you do not specify `@@global.` or `@@session.`), MySQL returns the session value if it exists and the global value otherwise. (This differs from `SET @@var_name = value`, which always refers to the session value.)

**Note**

Some variables displayed by `SHOW VARIABLES` may not be available using `SELECT @@var_name` syntax; an `Unknown system variable` occurs. As a workaround in such cases, you can use `SHOW VARIABLES LIKE 'var_name'`.

Suffixes for specifying a value multiplier can be used when setting a variable at server startup, but not to set the value with `SET` at runtime. On the other hand, with `SET` you can assign a variable's value using an expression, which is not true when you set a variable at server startup. For example, the first of the following lines is legal at server startup, but the second is not:

```
shell> mysql --max_allowed_packet=16M
shell> mysql --max_allowed_packet=16*1024*1024
```

Conversely, the second of the following lines is legal at runtime, but the first is not:

```
mysql> SET GLOBAL max_allowed_packet=16M;
mysql> SET GLOBAL max_allowed_packet=16*1024*1024;
```

**Note**

Some system variables can be enabled with the `SET` statement by setting them to `ON` or `1`, or disabled by setting them to `OFF` or `0`. However, to set such a variable on the command line or in an option file, you must set it to `1` or `0`; setting it to `ON` or `OFF` will not work. For example, on the command line, `--delay_key_write=1` [475] works but `--delay_key_write=ON` [475] does not.

To display system variable names and values, use the `SHOW VARIABLES` statement:

```
mysql> SHOW VARIABLES;
+-----+-----+
| Variable_name | Value |
+-----+-----+
| auto_increment_increment | 1 |
| auto_increment_offset | 1 |
| automatic_sp_privileges | ON |
| back_log | 50 |
| basedir | /home/mysql/ |
| binlog_cache_size | 32768 |
| bulk_insert_buffer_size | 8388608 |
| character_set_client | latin1 |
| character_set_connection | latin1 |
| character_set_database | latin1 |
| character_set_results | latin1 |
| character_set_server | latin1 |
| character_set_system | utf8 |
| character_sets_dir | /home/mysql/share/mysql/charsets/ |
| collation_connection | latin1_swedish_ci |
| collation_database | latin1_swedish_ci |
| collation_server | latin1_swedish_ci |
| ... | ... |
| innodb_additional_mem_pool_size | 1048576 |
| innodb_autoextend_increment | 8 |
| innodb_buffer_pool_size | 8388608 |
| innodb_checksums | ON |
| innodb_commit_concurrency | 0 |
| innodb_concurrency_tickets | 500 |
| innodb_data_file_path | ibdata1:10M:autoextend |
| innodb_data_home_dir | |
| ... | ... |
| version | 5.1.6-alpha-log |
| version_comment | Source distribution |
| version_compile_machine | i686 |
| version_compile_os | suse-linux |
| wait_timeout | 28800 |
+-----+-----+
```

With a [LIKE \[1190\]](#) clause, the statement displays only those variables that match the pattern. To obtain a specific variable name, use a [LIKE \[1190\]](#) clause as shown:

```
SHOW VARIABLES LIKE 'max_join_size';
SHOW SESSION VARIABLES LIKE 'max_join_size';
```

To get a list of variables whose name match a pattern, use the “%” wildcard character in a [LIKE \[1190\]](#) clause:

```
SHOW VARIABLES LIKE '%size%';
SHOW GLOBAL VARIABLES LIKE '%size%';
```

Wildcard characters can be used in any position within the pattern to be matched. Strictly speaking, because “\_” is a wildcard that matches any single character, you should escape it as “\\_” to match it literally. In practice, this is rarely necessary.

For [SHOW VARIABLES](#), if you specify neither [GLOBAL](#) nor [SESSION](#), MySQL returns [SESSION](#) values.

The reason for requiring the [GLOBAL](#) keyword when setting [GLOBAL](#)-only variables but not when retrieving them is to prevent problems in the future. If we were to remove a [SESSION](#) variable that has the same name as a [GLOBAL](#) variable, a client with the [SUPER \[756\]](#) privilege might accidentally change the [GLOBAL](#) variable rather than just the [SESSION](#) variable for its own connection. If we add a [SESSION](#) variable with the same name as a [GLOBAL](#) variable, a client that intends to change the [GLOBAL](#) variable might find only its own [SESSION](#) variable changed.

### 5.1.5.1. Structured System Variables

A structured variable differs from a regular system variable in two respects:

- Its value is a structure with components that specify server parameters considered to be closely related.
- There might be several instances of a given type of structured variable. Each one has a different name and refers to a different resource maintained by the server.

MySQL 5.6 supports one structured variable type, which specifies parameters governing the operation of key caches. A key cache structured variable has these components:

- [key\\_buffer\\_size \[492\]](#)
- [key\\_cache\\_block\\_size \[493\]](#)
- [key\\_cache\\_division\\_limit \[494\]](#)
- [key\\_cache\\_age\\_threshold \[493\]](#)

This section describes the syntax for referring to structured variables. Key cache variables are used for syntax examples, but specific details about how key caches operate are found elsewhere, in [Section 8.9.2, “The MyISAM Key Cache”](#).

To refer to a component of a structured variable instance, you can use a compound name in [instance\\_name.component\\_name](#) format. Examples:

```
hot_cache.key_buffer_size
hot_cache.key_cache_block_size
cold_cache.key_cache_block_size
```

For each structured system variable, an instance with the name of [default](#) is always predefined. If you refer to a component of a structured variable without any instance name, the [default](#) instance is used. Thus, [default.key\\_buffer\\_size](#) and [key\\_buffer\\_size \[492\]](#) both refer to the same system variable.

Structured variable instances and components follow these naming rules:

- For a given type of structured variable, each instance must have a name that is unique *within* variables of that type. However, instance names need not be unique *across* structured variable

types. For example, each structured variable has an instance named `default`, so `default` is not unique across variable types.

- The names of the components of each structured variable type must be unique across all system variable names. If this were not true (that is, if two different types of structured variables could share component member names), it would not be clear which default structured variable to use for references to member names that are not qualified by an instance name.
- If a structured variable instance name is not legal as an unquoted identifier, refer to it as a quoted identifier using backticks. For example, `hot-cache` is not legal, but ``hot-cache`` is.
- `global`, `session`, and `local` are not legal instance names. This avoids a conflict with notation such as `@@global.var_name` for referring to nonstructured system variables.

Currently, the first two rules have no possibility of being violated because the only structured variable type is the one for key caches. These rules will assume greater significance if some other type of structured variable is created in the future.

With one exception, you can refer to structured variable components using compound names in any context where simple variable names can occur. For example, you can assign a value to a structured variable using a command-line option:

```
shell> mysqld --hot_cache.key_buffer_size=64K
```

In an option file, use this syntax:

```
[mysqld]
hot_cache.key_buffer_size=64K
```

If you start the server with this option, it creates a key cache named `hot_cache` with a size of 64KB in addition to the default key cache that has a default size of 8MB.

Suppose that you start the server as follows:

```
shell> mysqld --key_buffer_size=256K \
  --extra_cache.key_buffer_size=128K \
  --extra_cache.key_cache_block_size=2048
```

In this case, the server sets the size of the default key cache to 256KB. (You could also have written `--default.key_buffer_size=256K`.) In addition, the server creates a second key cache named `extra_cache` that has a size of 128KB, with the size of block buffers for caching table index blocks set to 2048 bytes.

The following example starts the server with three different key caches having sizes in a 3:1:1 ratio:

```
shell> mysqld --key_buffer_size=6M \
  --hot_cache.key_buffer_size=2M \
  --cold_cache.key_buffer_size=2M
```

Structured variable values may be set and retrieved at runtime as well. For example, to set a key cache named `hot_cache` to a size of 10MB, use either of these statements:

```
mysql> SET GLOBAL hot_cache.key_buffer_size = 10*1024*1024;
mysql> SET @@global.hot_cache.key_buffer_size = 10*1024*1024;
```

To retrieve the cache size, do this:

```
mysql> SELECT @@global.hot_cache.key_buffer_size;
```

However, the following statement does not work. The variable is not interpreted as a compound name, but as a simple string for a `LIKE` [1190] pattern-matching operation:

```
mysql> SHOW GLOBAL VARIABLES LIKE 'hot_cache.key_buffer_size';
```

This is the exception to being able to use structured variable names anywhere a simple variable name may occur.



### 5.1.5.2. Dynamic System Variables

Many server system variables are dynamic and can be set at runtime using `SET GLOBAL` or `SET SESSION`. You can also obtain their values using `SELECT`. See [Section 5.1.5, “Using System Variables”](#).

The following table shows the full list of all dynamic system variables. The last column indicates for each variable whether `GLOBAL` or `SESSION` (or both) apply. The table also lists session options that can be set with the `SET` statement. [Section 5.1.4, “Server System Variables”](#), discusses these options.

Variables that have a type of “string” take a string value. Variables that have a type of “numeric” take a numeric value. Variables that have a type of “boolean” can be set to 0, 1, `ON` or `OFF`. (If you set them on the command line or in an option file, use the numeric values.) Variables that are marked as “enumeration” normally should be set to one of the available values for the variable, but can also be set to the number that corresponds to the desired enumeration value. For enumerated system variables, the first enumeration value corresponds to 0. This differs from `ENUM` columns, for which the first enumeration value corresponds to 1.

**Table 5.4. Dynamic Variable Summary**

Variable Name	Variable Type	Variable Scope
<code>audit_log_flush</code> <a href="#">[825]</a>	boolean	GLOBAL
<code>audit_log_policy</code> <a href="#">[826]</a>	enumeration	GLOBAL
<code>audit_log_rotate_on_size</code> <a href="#">[826]</a>	numeric	GLOBAL
<code>auto_increment_increment</code> <a href="#">[1965]</a>	numeric	GLOBAL   SESSION
<code>auto_increment_offset</code> <a href="#">[1968]</a>	numeric	GLOBAL   SESSION
<code>autocommit</code> <a href="#">[464]</a>	boolean	GLOBAL   SESSION
<code>automatic_sp_privileges</code> <a href="#">[464]</a>	boolean	GLOBAL
<code>big_tables</code> <a href="#">[414]</a>	boolean	GLOBAL   SESSION
<code>binlog_cache_size</code> <a href="#">[2008]</a>	numeric	GLOBAL
<code>binlog_checksum</code> <a href="#">[2009]</a>	string	GLOBAL
<code>binlog_direct_non_transactional_updates</code> <a href="#">[415]</a>	boolean	GLOBAL   SESSION
<code>binlog_format</code> <a href="#">[415]</a>	enumeration	GLOBAL   SESSION
<code>binlog_max_flush_queue_time</code> <a href="#">[2011]</a>	numeric	GLOBAL
<code>binlog_order_commits</code> <a href="#">[2012]</a>	boolean	GLOBAL
<code>binlog_row_image=image_type</code> <a href="#">[2012]</a>	enumeration	GLOBAL   SESSION
<code>binlog_rows_query_log_events</code> <a href="#">[2014]</a>	boolean	GLOBAL   SESSION
<code>binlog_stmt_cache_size</code> <a href="#">[2014]</a>	numeric	GLOBAL
<code>bulk_insert_buffer_size</code> <a href="#">[466]</a>	numeric	GLOBAL   SESSION
<code>character_set_client</code> <a href="#">[467]</a>	string	GLOBAL   SESSION
<code>character_set_connection</code> <a href="#">[467]</a>	string	GLOBAL   SESSION
<code>character_set_database</code> <a href="#">[468]</a>	string	GLOBAL   SESSION
<code>character_set_filesystem</code> <a href="#">[416]</a>	string	GLOBAL   SESSION
<code>character_set_results</code> <a href="#">[468]</a>	string	GLOBAL   SESSION
<code>character_set_server</code> <a href="#">[417]</a>	string	GLOBAL   SESSION
<code>collation_connection</code> <a href="#">[469]</a>	string	GLOBAL   SESSION
<code>collation_database</code> <a href="#">[469]</a>	string	GLOBAL   SESSION
<code>collation_server</code> <a href="#">[417]</a>	string	GLOBAL   SESSION
<code>completion_type</code> <a href="#">[470]</a>	numeric	GLOBAL   SESSION



Variable Name	Variable Type	Variable Scope
<a href="#">concurrent_insert [471]</a>	boolean	GLOBAL
<a href="#">connect_timeout [472]</a>	numeric	GLOBAL
<a href="#">debug [418]</a>	string	GLOBAL   SESSION
<a href="#">debug_sync [473]</a>	string	SESSION
<a href="#">default_storage_engine [420]</a>	enumeration	GLOBAL   SESSION
<a href="#">default_tmp_storage_engine [474]</a>	enumeration	GLOBAL   SESSION
<a href="#">default_week_format [475]</a>	numeric	GLOBAL   SESSION
<a href="#">delay_key_write [421]</a>	enumeration	GLOBAL
<a href="#">delayed_insert_limit [476]</a>	numeric	GLOBAL
<a href="#">delayed_insert_timeout [476]</a>	numeric	GLOBAL
<a href="#">delayed_queue_size [477]</a>	numeric	GLOBAL
<a href="#">div_precision_increment [478]</a>	numeric	GLOBAL   SESSION
<a href="#">end_markers_in_json [479]</a>	boolean	GLOBAL   SESSION
<a href="#">engine_condition_pushdown [421]</a>	boolean	GLOBAL   SESSION
<a href="#">eq_range_index_dive_limit [479]</a>	numeric	GLOBAL   SESSION
<a href="#">event_scheduler [422]</a>	enumeration	GLOBAL
<a href="#">expire_logs_days [480]</a>	numeric	GLOBAL
<a href="#">flush [482]</a>	boolean	GLOBAL
<a href="#">flush_time [482]</a>	numeric	GLOBAL
<a href="#">foreign_key_checks [483]</a>	boolean	GLOBAL   SESSION
<a href="#">ft_boolean_syntax [483]</a>	string	GLOBAL
<a href="#">general_log [423]</a>	boolean	GLOBAL
<a href="#">general_log_file [486]</a>	filename	GLOBAL
<a href="#">group_concat_max_len [486]</a>	numeric	GLOBAL   SESSION
<a href="#">gtid_next [2022]</a>	enumeration	SESSION
<a href="#">gtid_purged [2023]</a>	string	GLOBAL
<a href="#">host_cache_size [487]</a>	numeric	GLOBAL
<a href="#">identity [488]</a>	numeric	SESSION
<a href="#">init_connect [489]</a>	string	GLOBAL
<a href="#">init_slave [1987]</a>	string	GLOBAL
<a href="#">innodb_adaptive_flushing [1700]</a>	boolean	GLOBAL
<a href="#">innodb_adaptive_flushing_lwm [1701]</a>	numeric	GLOBAL
<a href="#">innodb_adaptive_hash_index [1701]</a>	boolean	GLOBAL
<a href="#">innodb_adaptive_max_sleep_delay [1702]</a>	numeric	GLOBAL
<a href="#">innodb_api_bk_commit_interval [1702]</a>	numeric	GLOBAL
<a href="#">innodb_api_trx_level [1704]</a>	numeric	GLOBAL
<a href="#">innodb_autoextend_increment [1704]</a>	numeric	GLOBAL
<a href="#">innodb_buffer_pool_dump_at_shutdown [1705]</a>	boolean	GLOBAL
<a href="#">innodb_buffer_pool_dump_now [1706]</a>	boolean	GLOBAL
<a href="#">innodb_buffer_pool_filename [1706]</a>	string	GLOBAL
<a href="#">innodb_buffer_pool_load_abort [1707]</a>	boolean	GLOBAL

Variable Name	Variable Type	Variable Scope
<a href="#">innodb_buffer_pool_load_now [1708]</a>	boolean	GLOBAL
<a href="#">innodb_change_buffer_max_size [1709]</a>	numeric	GLOBAL
<a href="#">innodb_change_buffering [1709]</a>	enumeration	GLOBAL
<a href="#">innodb_checksum_algorithm [1710]</a>	enumeration	GLOBAL
<a href="#">innodb_cmp_per_index_enabled [1712]</a>	boolean	GLOBAL
<a href="#">innodb_commit_concurrency [1712]</a>	numeric	GLOBAL
<a href="#">innodb_compression_failure_threshold_pages [1713]</a>	numeric	GLOBAL
<a href="#">innodb_compression_level [1713]</a>	numeric	GLOBAL
<a href="#">innodb_compression_pad_pct_max [1713]</a>	numeric	GLOBAL
<a href="#">innodb_concurrency_tickets [1714]</a>	numeric	GLOBAL
<a href="#">innodb_disable_sort_file_cache [1715]</a>	boolean	GLOBAL
<a href="#">innodb_fast_shutdown [1716]</a>	numeric	GLOBAL
<a href="#">innodb_file_format [1716]</a>	string	GLOBAL
<a href="#">innodb_file_format_max [1717]</a>	string	GLOBAL
<a href="#">innodb_file_per_table [1718]</a>	boolean	GLOBAL
<a href="#">innodb_flush_log_at_timeout [1718]</a>	numeric	GLOBAL
<a href="#">innodb_flush_log_at_trx_commit [1718]</a>	enumeration	GLOBAL
<a href="#">innodb_flush_neighbors [1721]</a>	boolean	GLOBAL
<a href="#">innodb_flushing_avg_loops [1722]</a>	numeric	GLOBAL
<a href="#">innodb_ft_aux_table [1723]</a>	string	GLOBAL
<a href="#">innodb_ft_enable_stopword [1724]</a>	boolean	GLOBAL
<a href="#">innodb_ft_num_word_optimize [1725]</a>	numeric	GLOBAL
<a href="#">innodb_ft_server_stopword_table [1725]</a>	string	GLOBAL
<a href="#">innodb_ft_user_stopword_table [1726]</a>	string	GLOBAL   SESSION
<a href="#">innodb_io_capacity [1727]</a>	numeric	GLOBAL
<a href="#">innodb_io_capacity_max [1728]</a>	numeric	GLOBAL
<a href="#">innodb_large_prefix [1728]</a>	boolean	GLOBAL
<a href="#">innodb_lock_wait_timeout [1729]</a>	numeric	GLOBAL   SESSION
<a href="#">innodb_log_compressed_pages [1732]</a>	boolean	GLOBAL
<a href="#">innodb_lru_scan_depth [1734]</a>	numeric	GLOBAL
<a href="#">innodb_max_dirty_pages_pct [1734]</a>	numeric	GLOBAL
<a href="#">innodb_max_dirty_pages_pct_lwm [1735]</a>	numeric	GLOBAL
<a href="#">innodb_max_purge_lag [1735]</a>	numeric	GLOBAL
<a href="#">innodb_max_purge_lag_delay [1736]</a>	numeric	GLOBAL
<a href="#">innodb_monitor_disable [1736]</a>	string	GLOBAL
<a href="#">innodb_monitor_enable [1737]</a>	string	GLOBAL
<a href="#">innodb_monitor_reset [1737]</a>	string	GLOBAL
<a href="#">innodb_monitor_reset_all [1737]</a>	string	GLOBAL
<a href="#">innodb_old_blocks_pct [1738]</a>	numeric	GLOBAL
<a href="#">innodb_old_blocks_time [1738]</a>	numeric	GLOBAL
<a href="#">innodb_online_alter_log_max_size [1739]</a>	numeric	GLOBAL

Variable Name	Variable Type	Variable Scope
<a href="#">innodb_optimize_fulltext_only [1740]</a>	boolean	GLOBAL
<a href="#">innodb_print_all_deadlocks [1741]</a>	boolean	GLOBAL
<a href="#">innodb_purge_batch_size [1741]</a>	numeric	GLOBAL
<a href="#">innodb_random_read_ahead [1742]</a>	boolean	GLOBAL
<a href="#">innodb_read_ahead_threshold [1742]</a>	numeric	GLOBAL
<a href="#">innodb_replication_delay [1744]</a>	numeric	GLOBAL
<a href="#">innodb_rollback_segments [1744]</a>	numeric	GLOBAL
<a href="#">innodb_sort_buffer_size [1745]</a>	numeric	GLOBAL
<a href="#">innodb_spin_wait_delay [1745]</a>	numeric	GLOBAL
<a href="#">innodb_stats_auto_recalc [1746]</a>	boolean	GLOBAL
<a href="#">innodb_stats_method [1746]</a>	enumeration	GLOBAL
<a href="#">innodb_stats_on_metadata [1747]</a>	boolean	GLOBAL
<a href="#">innodb_stats_persistent [1747]</a>	boolean	GLOBAL
<a href="#">innodb_stats_persistent_sample_pages [1748]</a>	numeric	GLOBAL
<a href="#">innodb_stats_sample_pages [1748]</a>	numeric	GLOBAL
<a href="#">innodb_stats_transient_sample_pages [1748]</a>	numeric	GLOBAL
<a href="#">innodb_strict_mode [1749]</a>	boolean	GLOBAL   SESSION
<a href="#">innodb_support_xa [1749]</a>	boolean	GLOBAL   SESSION
<a href="#">innodb_sync_spin_loops [1750]</a>	numeric	GLOBAL
<a href="#">innodb_table_locks [1751]</a>	boolean	GLOBAL   SESSION
<a href="#">innodb_thread_concurrency [1751]</a>	numeric	GLOBAL
<a href="#">innodb_thread_sleep_delay [1752]</a>	numeric	GLOBAL
<a href="#">innodb_undo_logs [1752]</a>	numeric	GLOBAL
<a href="#">innodb_undo_tablespace [1753]</a>	numeric	GLOBAL
<a href="#">insert_id [490]</a>	numeric	SESSION
<a href="#">interactive_timeout [490]</a>	numeric	GLOBAL   SESSION
<a href="#">join_buffer_size [490]</a>	numeric	GLOBAL   SESSION
<a href="#">keep_files_on_create [491]</a>	boolean	GLOBAL   SESSION
<a href="#">key_buffer_size [492]</a>	numeric	GLOBAL
<a href="#">key_cache_age_threshold [493]</a>	numeric	GLOBAL
<a href="#">key_cache_block_size [493]</a>	numeric	GLOBAL
<a href="#">key_cache_division_limit [494]</a>	numeric	GLOBAL
<a href="#">last_insert_id [495]</a>	numeric	SESSION
<a href="#">lc_messages [426]</a>	string	GLOBAL   SESSION
<a href="#">lc_time_names [495]</a>	string	GLOBAL   SESSION
<a href="#">local_infile [496]</a>	boolean	GLOBAL
<a href="#">lock_wait_timeout [496]</a>	numeric	GLOBAL   SESSION
<a href="#">log [426]</a>	string	GLOBAL
<a href="#">log_output [427]</a>	set	GLOBAL
<a href="#">log_queries_not_using_indexes [428]</a>	boolean	GLOBAL
<a href="#">log_slow_admin_statements [498]</a>	boolean	GLOBAL

Variable Name	Variable Type	Variable Scope
<code>log_slow_queries</code> [429]	boolean	GLOBAL
<code>log_slow_slave_statements</code> [1988]	boolean	GLOBAL
<code>log_throttle_queries_not_using_indexes</code> [403]	numeric	GLOBAL
<code>log_warnings</code> [430]	numeric	GLOBAL
<code>long_query_time</code> [500]	numeric	GLOBAL   SESSION
<code>low_priority_updates</code> [430]	boolean	GLOBAL   SESSION
<code>master_info_repository</code> [1988]	string	GLOBAL
<code>master_verify_checksum</code> [2015]	boolean	GLOBAL
<code>max_allowed_packet</code> [502]	numeric	GLOBAL
<code>max_binlog_cache_size</code> [2016]	numeric	GLOBAL
<code>max_binlog_size</code> [2016]	numeric	GLOBAL
<code>max_binlog_stmt_cache_size</code> [2017]	numeric	GLOBAL
<code>max_connect_errors</code> [502]	numeric	GLOBAL
<code>max_connections</code> [503]	numeric	GLOBAL
<code>max_delayed_threads</code> [504]	numeric	GLOBAL   SESSION
<code>max_error_count</code> [504]	numeric	GLOBAL   SESSION
<code>max_heap_table_size</code> [504]	numeric	GLOBAL   SESSION
<code>max_insert_delayed_threads</code> [505]	numeric	GLOBAL   SESSION
<code>max_join_size</code> [505]	numeric	GLOBAL   SESSION
<code>max_length_for_sort_data</code> [506]	numeric	GLOBAL   SESSION
<code>max_prepared_stmt_count</code> [506]	numeric	GLOBAL
<code>max_relay_log_size</code> [507]	numeric	GLOBAL
<code>max_seeks_for_key</code> [507]	numeric	GLOBAL   SESSION
<code>max_sort_length</code> [508]	numeric	GLOBAL   SESSION
<code>max_sp_recursion_depth</code> [508]	numeric	GLOBAL   SESSION
<code>max_user_connections</code> [509]	numeric	GLOBAL   SESSION
<code>max_write_lock_count</code> [509]	numeric	GLOBAL
<code>min_examined_row_limit</code> [431]	numeric	GLOBAL   SESSION
<code>myisam_data_pointer_size</code> [511]	numeric	GLOBAL
<code>myisam_max_sort_file_size</code> [511]	numeric	GLOBAL
<code>myisam_repair_threads</code> [512]	numeric	GLOBAL   SESSION
<code>myisam_sort_buffer_size</code> [513]	numeric	GLOBAL   SESSION
<code>myisam_stats_method</code> [513]	enumeration	GLOBAL   SESSION
<code>myisam_use_mmap</code> [514]	boolean	GLOBAL
<code>ndb_cache_check_time</code>	numeric	GLOBAL
<code>net_buffer_length</code> [515]	numeric	GLOBAL   SESSION
<code>net_read_timeout</code> [515]	numeric	GLOBAL   SESSION
<code>net_retry_count</code> [515]	numeric	GLOBAL   SESSION
<code>net_write_timeout</code> [516]	numeric	GLOBAL   SESSION
<code>new</code> [516]	boolean	GLOBAL   SESSION
<code>old_alter_table</code> [433]	boolean	GLOBAL   SESSION

Variable Name	Variable Type	Variable Scope
<a href="#">old_passwords [517]</a>	boolean	GLOBAL   SESSION
<a href="#">optimizer_join_cache_level [519]</a>	numeric	GLOBAL   SESSION
<a href="#">optimizer_prune_level [519]</a>	boolean	GLOBAL   SESSION
<a href="#">optimizer_search_depth [520]</a>	numeric	GLOBAL   SESSION
<a href="#">optimizer_switch [520]</a>	set	GLOBAL   SESSION
<a href="#">optimizer_trace [523]</a>	string	GLOBAL   SESSION
<a href="#">optimizer_trace_features [523]</a>	string	GLOBAL   SESSION
<a href="#">optimizer_trace_limit [524]</a>	numeric	GLOBAL   SESSION
<a href="#">optimizer_trace_max_mem_size [524]</a>	numeric	GLOBAL   SESSION
<a href="#">optimizer_trace_offset [524]</a>	numeric	GLOBAL   SESSION
<a href="#">preload_buffer_size [525]</a>	numeric	GLOBAL   SESSION
<a href="#">profiling [526]</a>	boolean	GLOBAL   SESSION
<a href="#">profiling_history_size [526]</a>	numeric	GLOBAL   SESSION
<a href="#">pseudo_slave_mode [526]</a>	numeric	SESSION
<a href="#">pseudo_thread_id [527]</a>	numeric	SESSION
<a href="#">query_alloc_block_size [527]</a>	numeric	GLOBAL   SESSION
<a href="#">query_cache_limit [527]</a>	numeric	GLOBAL
<a href="#">query_cache_min_res_unit [528]</a>	numeric	GLOBAL
<a href="#">query_cache_size [528]</a>	numeric	GLOBAL
<a href="#">query_cache_type [529]</a>	enumeration	GLOBAL   SESSION
<a href="#">query_cache_wlock_invalidate [530]</a>	boolean	GLOBAL   SESSION
<a href="#">query_prealloc_size [531]</a>	numeric	GLOBAL   SESSION
<a href="#">rand_seed1 [531]</a>	numeric	SESSION
<a href="#">rand_seed2 [531]</a>	numeric	SESSION
<a href="#">range_alloc_block_size [531]</a>	numeric	GLOBAL   SESSION
<a href="#">read_buffer_size [532]</a>	numeric	GLOBAL   SESSION
<a href="#">read_only [532]</a>	boolean	GLOBAL
<a href="#">read_rnd_buffer_size [533]</a>	numeric	GLOBAL   SESSION
<a href="#">relay_log_info_repository [1990]</a>	string	GLOBAL
<a href="#">relay_log_purge [534]</a>	boolean	GLOBAL
<a href="#">relay_log_recovery [1990]</a>	boolean	GLOBAL
<a href="#">rpl_semi_sync_master_enabled [536]</a>	boolean	GLOBAL
<a href="#">rpl_semi_sync_master_timeout [536]</a>	numeric	GLOBAL
<a href="#">rpl_semi_sync_master_trace_level [536]</a>	numeric	GLOBAL
<a href="#">rpl_semi_sync_master_wait_no_slave [537]</a>	boolean	GLOBAL
<a href="#">rpl_semi_sync_slave_enabled [537]</a>	boolean	GLOBAL
<a href="#">rpl_semi_sync_slave_trace_level [538]</a>	numeric	GLOBAL
<a href="#">secure_auth [438]</a>	boolean	GLOBAL
<a href="#">server_id [1955]</a>	numeric	GLOBAL
<a href="#">slave_allow_batching [1968]</a>	boolean	GLOBAL
<a href="#">slave_checkpoint_group=# [1991]</a>	numeric	GLOBAL

Variable Name	Variable Type	Variable Scope
<a href="#">slave_checkpoint_period=# [1991]</a>	numeric	GLOBAL
<a href="#">slave_compressed_protocol [1992]</a>	boolean	GLOBAL
<a href="#">slave_exec_mode [1992]</a>	enumeration	GLOBAL
<a href="#">slave_max_allowed_packet [1969]</a>	numeric	GLOBAL
<a href="#">slave_net_timeout [1985]</a>	numeric	GLOBAL
<a href="#">slave_parallel_workers [1993]</a>	numeric	GLOBAL
<a href="#">slave_pending_jobs_size_max [1994]</a>	numeric	GLOBAL
<a href="#">slave_rows_search_algorithms=list [1969]</a>	set	GLOBAL
<a href="#">slave_sql_verify_checksum [1995]</a>	boolean	GLOBAL
<a href="#">slave_transaction_retries [1995]</a>	numeric	GLOBAL
<a href="#">slow_launch_time [541]</a>	numeric	GLOBAL
<a href="#">slow_query_log [441]</a>	boolean	GLOBAL
<a href="#">slow_query_log_file [541]</a>	filename	GLOBAL
<a href="#">sort_buffer_size [542]</a>	numeric	GLOBAL   SESSION
<a href="#">sql_auto_is_null [543]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_big_selects [543]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_big_tables [414]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_buffer_result [544]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_log_bin [544]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_log_off [544]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_low_priority_updates [500]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_max_join_size [505]</a>	numeric	GLOBAL   SESSION
<a href="#">sql_mode [442]</a>	set	GLOBAL   SESSION
<a href="#">sql_notes [546]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_quote_show_create [546]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_safe_updates [547]</a>	boolean	GLOBAL   SESSION
<a href="#">sql_select_limit [547]</a>	numeric	GLOBAL   SESSION
<a href="#">sql_slave_skip_counter [1996]</a>	numeric	GLOBAL
<a href="#">sql_warnings [547]</a>	boolean	GLOBAL   SESSION
<a href="#">storage_engine [549]</a>	enumeration	GLOBAL   SESSION
<a href="#">stored_program_cache [549]</a>	numeric	GLOBAL
<a href="#">sync_binlog [2017]</a>	numeric	GLOBAL
<a href="#">sync_frm [550]</a>	boolean	GLOBAL
<a href="#">sync_master_info [1997]</a>	numeric	GLOBAL
<a href="#">sync_relay_log [1998]</a>	numeric	GLOBAL
<a href="#">sync_relay_log_info [1999]</a>	numeric	GLOBAL
<a href="#">table_definition_cache [550]</a>	numeric	GLOBAL
<a href="#">table_open_cache [551]</a>	numeric	GLOBAL
<a href="#">thread_cache_size [552]</a>	numeric	GLOBAL
<a href="#">thread_pool_high_priority_connection [554]</a>	numeric	GLOBAL   SESSION
<a href="#">thread_pool_max_unused_threads [555]</a>	numeric	GLOBAL

Variable Name	Variable Type	Variable Scope
<a href="#">thread_pool_prio_kickup_timer [555]</a>	numeric	GLOBAL   SESSION
<a href="#">thread_pool_stall_limit [556]</a>	numeric	GLOBAL
<a href="#">time_zone [557]</a>	string	GLOBAL   SESSION
<a href="#">timed_mutexes [557]</a>	boolean	GLOBAL
<a href="#">timestamp [558]</a>	numeric	SESSION
<a href="#">tmp_table_size [558]</a>	numeric	GLOBAL   SESSION
<a href="#">transaction_alloc_block_size [559]</a>	numeric	GLOBAL   SESSION
<a href="#">transaction_prealloc_size [560]</a>	numeric	GLOBAL   SESSION
<a href="#">tx_isolation [560]</a>	enumeration	GLOBAL   SESSION
<a href="#">tx_read_only [561]</a>	boolean	GLOBAL   SESSION
<a href="#">unique_checks [561]</a>	boolean	GLOBAL   SESSION
<a href="#">updatable_views_with_limit [562]</a>	boolean	GLOBAL   SESSION
<a href="#">validate_password_length [744]</a>	numeric	GLOBAL
<a href="#">validate_password_mixed_case_count [744]</a>	numeric	GLOBAL
<a href="#">validate_password_number_count [745]</a>	numeric	GLOBAL
<a href="#">validate_password_policy [745]</a>	enumeration	GLOBAL
<a href="#">validate_password_special_char_count [745]</a>	numeric	GLOBAL
<a href="#">wait_timeout [563]</a>	numeric	GLOBAL   SESSION

### 5.1.6. Server Status Variables

The server maintains many status variables that provide information about its operation. You can view these variables and their values by using the `SHOW [GLOBAL | SESSION] STATUS` statement (see [Section 13.7.5.36, “SHOW STATUS Syntax”](#)). The optional `GLOBAL` keyword aggregates the values over all connections, and `SESSION` shows the values for the current connection.

```
mysql> SHOW GLOBAL STATUS;
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Aborted_clients | 0 |
| Aborted_connects | 0 |
| Bytes_received | 155372598 |
| Bytes_sent | 1176560426 |
| ... | |
| Connections | 30023 |
| Created_tmp_disk_tables | 0 |
| Created_tmp_files | 3 |
| Created_tmp_tables | 2 |
| ... | |
| Threads_created | 217 |
| Threads_running | 88 |
| Uptime | 1389872 |
+-----+-----+
```

Many status variables are reset to 0 by the `FLUSH STATUS` statement.

The following table lists all available server status variables:

**Table 5.5. Status Variable Summary**

Variable Name	Variable Type	Variable Scope
<a href="#">Aborted_clients [585]</a>	numeric	GLOBAL
<a href="#">Aborted_connects [585]</a>	numeric	GLOBAL

## Server Status Variables

Variable Name	Variable Type	Variable Scope
<a href="#">Binlog_cache_disk_use [586]</a>	numeric	GLOBAL
<a href="#">Binlog_cache_use [586]</a>	numeric	GLOBAL
<a href="#">Binlog_stmt_cache_disk_use [586]</a>	numeric	GLOBAL
<a href="#">Binlog_stmt_cache_use [586]</a>	numeric	GLOBAL
<a href="#">Bytes_received [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Bytes_sent [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_admin_commands [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_db [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_db_upgrade [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_event [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_function [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_procedure [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_server [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_table [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_tablespace [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_alter_user [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_analyze [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_assign_to_keycache [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_begin [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_binlog [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_call_procedure [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_change_db [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_change_master [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_check [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_checksum [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_commit [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_db [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_event [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_function [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_index [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_procedure [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_server [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_table [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_trigger [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_udf [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_user [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_create_view [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_dealloc_sql [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_delete [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_delete_multi [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_do [586]</a>	numeric	GLOBAL   SESSION



Variable Name	Variable Type	Variable Scope
<a href="#">Com_drop_db [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_event [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_function [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_index [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_procedure [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_server [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_table [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_trigger [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_user [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_drop_view [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_empty_query [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_execute_sql [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_flush [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_get_diagnostics [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_grant [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_ha_close [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_ha_open [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_ha_read [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_help [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_insert [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_insert_select [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_install_plugin [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_kill [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_load [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_lock_tables [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_optimize [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_preload_keys [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_prepare_sql [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_purge [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_purge_before_date [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_release_savepoint [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_rename_table [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_rename_user [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_repair [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_replace [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_replace_select [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_reset [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_resignal [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_revoke [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_revoke_all [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_rollback [586]</a>	numeric	GLOBAL   SESSION

Variable Name	Variable Type	Variable Scope
<a href="#">Com_rollback_to_savepoint [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_savepoint [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_select [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_set_option [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_authors [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_binlog_events [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_binlogs [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_charsets [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_collations [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_contributors [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_db [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_event [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_func [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_proc [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_table [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_create_trigger [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_databases [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_engine_logs [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_engine_mutex [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_engine_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_errors [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_events [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_fields [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_function_code [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_function_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_grants [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_keys [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_logs [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_master_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_new_master [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_open_tables [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_plugins [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_privileges [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_procedure_code [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_procedure_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_processlist [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_profile [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_profiles [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_relaylog_events [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_slave_hosts [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_slave_status [586]</a>	numeric	GLOBAL   SESSION

Variable Name	Variable Type	Variable Scope
<a href="#">Com_show_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_storage_engines [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_table_status [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_tables [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_triggers [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_variables [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_show_warnings [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_signal [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_slave_start [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_slave_stop [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_close [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_execute [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_fetch [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_prepare [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_reprepare [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_reset [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_stmt_send_long_data [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_truncate [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_uninstall_plugin [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_unlock_tables [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_update [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_update_multi [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_commit [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_end [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_prepare [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_recover [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_rollback [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Com_xa_start [586]</a>	numeric	GLOBAL   SESSION
<a href="#">Compression [587]</a>	numeric	SESSION
<a href="#">Connection_errors_accept [587]</a>	numeric	GLOBAL
<a href="#">Connection_errors_internal [587]</a>	numeric	GLOBAL
<a href="#">Connection_errors_max_connections [587]</a>	numeric	GLOBAL
<a href="#">Connection_errors_peer_addr [587]</a>	numeric	GLOBAL
<a href="#">Connection_errors_select [587]</a>	numeric	GLOBAL
<a href="#">Connection_errors_tcpwrap [587]</a>	numeric	GLOBAL
<a href="#">Connections [587]</a>	numeric	GLOBAL
<a href="#">Created_tmp_disk_tables [587]</a>	numeric	GLOBAL   SESSION
<a href="#">Created_tmp_files [588]</a>	numeric	GLOBAL
<a href="#">Created_tmp_tables [588]</a>	numeric	GLOBAL   SESSION
<a href="#">Delayed_errors [588]</a>	numeric	GLOBAL
<a href="#">Delayed_insert_threads [588]</a>	numeric	GLOBAL

Server Status Variables

Variable Name	Variable Type	Variable Scope
<a href="#">Delayed_writes [588]</a>	numeric	GLOBAL
<a href="#">Flush_commands [588]</a>	numeric	GLOBAL
<a href="#">Handler_commit [588]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_delete [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_discover</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_external_lock [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_mrr_init [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_prepare [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_first [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_key [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_last [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_next [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_prev [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_rnd [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_read_rnd_next [589]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_rollback [590]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_savepoint [590]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_savepoint_rollback [590]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_update [590]</a>	numeric	GLOBAL   SESSION
<a href="#">Handler_write [590]</a>	numeric	GLOBAL   SESSION
<a href="#">Innodb_available_undo_logs [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_bytes_data [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_bytes_dirty [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_dump_status [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_load_status [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_data [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_dirty [590]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_flushed [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_free [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_latched [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_misc [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_pages_total [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_read_ahead [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_read_ahead_evicted [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_read_requests [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_reads [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_wait_free [591]</a>	numeric	GLOBAL
<a href="#">Innodb_buffer_pool_write_requests [591]</a>	numeric	GLOBAL
<a href="#">Innodb_data_fsyncs [591]</a>	numeric	GLOBAL
<a href="#">Innodb_data_pending_fsyncs [591]</a>	numeric	GLOBAL
<a href="#">Innodb_data_pending_reads [592]</a>	numeric	GLOBAL

Variable Name	Variable Type	Variable Scope
InnoDB_data_pending_writes <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_data_read <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_data_reads <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_data_writes <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_data_written <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_dblwr_pages_written <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_dblwr_writes <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_have_atomic_builtins <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_log_waits <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_log_write_requests <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_log_writes <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_num_open_files <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_os_log_fsyncs <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_os_log_pending_fsyncs <a href="#">[592]</a>	numeric	GLOBAL
InnoDB_os_log_pending_writes <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_os_log_written <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_page_size <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_pages_created <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_pages_read <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_pages_written <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_row_lock_current_waits <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_row_lock_time <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_row_lock_time_avg <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_row_lock_time_max <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_row_lock_waits <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_rows_deleted <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_rows_inserted <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_rows_read <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_rows_updated <a href="#">[593]</a>	numeric	GLOBAL
InnoDB_truncated_status_writes <a href="#">[593]</a>	numeric	GLOBAL
Key_blocks_not_flushed <a href="#">[594]</a>	numeric	GLOBAL
Key_blocks_unused <a href="#">[594]</a>	numeric	GLOBAL
Key_blocks_used <a href="#">[594]</a>	numeric	GLOBAL
Key_read_requests <a href="#">[594]</a>	numeric	GLOBAL
Key_reads <a href="#">[594]</a>	numeric	GLOBAL
Key_write_requests <a href="#">[594]</a>	numeric	GLOBAL
Key_writes <a href="#">[594]</a>	numeric	GLOBAL
Last_query_cost <a href="#">[594]</a>	numeric	SESSION
Last_query_partial_plans <a href="#">[594]</a>	numeric	SESSION
Max_used_connections <a href="#">[594]</a>	numeric	GLOBAL
Ndb_conflict_fn_max	numeric	GLOBAL

Server Status Variables

Variable Name	Variable Type	Variable Scope
Ndb_conflict_fn_old	numeric	GLOBAL
Ndb_number_of_data_nodes	numeric	GLOBAL
Not_flushed_delayed_rows <a href="#">[594]</a>	numeric	GLOBAL
Open_files <a href="#">[594]</a>	numeric	GLOBAL
Open_streams <a href="#">[595]</a>	numeric	GLOBAL
Open_table_definitions <a href="#">[595]</a>	numeric	GLOBAL
Open_tables <a href="#">[595]</a>	numeric	GLOBAL   SESSION
Opened_files <a href="#">[595]</a>	numeric	GLOBAL
Opened_table_definitions <a href="#">[595]</a>	numeric	GLOBAL   SESSION
Opened_tables <a href="#">[595]</a>	numeric	GLOBAL   SESSION
Performance_schema_accounts_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_cond_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_cond_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_file_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_file_handles_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_file_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_hosts_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_locker_lost <a href="#">[2319]</a>	numeric	GLOBAL
Performance_schema_mutex_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_mutex_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_rwlock_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_rwlock_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_session_connect_attempts <a href="#">[2319]</a>	numeric	GLOBAL
Performance_schema_socket_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_socket_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_stage_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_statement_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_table_handles_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_table_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_thread_classes_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_thread_instances_lost <a href="#">[2318]</a>	numeric	GLOBAL
Performance_schema_users_lost <a href="#">[2319]</a>	numeric	GLOBAL
Prepared_stmt_count <a href="#">[595]</a>	numeric	GLOBAL
Qcache_free_blocks <a href="#">[595]</a>	numeric	GLOBAL
Qcache_free_memory <a href="#">[595]</a>	numeric	GLOBAL
Qcache_hits <a href="#">[595]</a>	numeric	GLOBAL
Qcache_inserts <a href="#">[595]</a>	numeric	GLOBAL
Qcache_lowmem_prunes <a href="#">[595]</a>	numeric	GLOBAL
Qcache_not_cached <a href="#">[595]</a>	numeric	GLOBAL
Qcache_queries_in_cache <a href="#">[596]</a>	numeric	GLOBAL
Qcache_total_blocks <a href="#">[596]</a>	numeric	GLOBAL

## Server Status Variables

Variable Name	Variable Type	Variable Scope
<a href="#">Queries [596]</a>	numeric	GLOBAL   SESSION
<a href="#">Questions [596]</a>	numeric	GLOBAL   SESSION
<a href="#">Rpl_semi_sync_master_clients [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_net_avg_wait_time [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_net_wait_time [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_net_waits [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_no_times [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_no_tx [596]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_status [596]</a>	boolean	GLOBAL
<a href="#">Rpl_semi_sync_master_timefunc_failures [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_tx_avg_wait_time [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_tx_wait_time [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_tx_waits [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_wait_pos_backtraverse [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_wait_sessions [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_master_yes_tx [597]</a>	numeric	GLOBAL
<a href="#">Rpl_semi_sync_slave_status [597]</a>	boolean	GLOBAL
<a href="#">Rsa_public_key [597]</a>	string	GLOBAL
<a href="#">Select_full_join [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Select_full_range_join [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Select_range [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Select_range_check [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Select_scan [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Slave_heartbeat_period [598]</a>		GLOBAL
<a href="#">Slave_last_heartbeat [598]</a>		GLOBAL
<a href="#">Slave_open_temp_tables [598]</a>	numeric	GLOBAL
<a href="#">Slave_received_heartbeats [598]</a>		GLOBAL
<a href="#">Slave_retried_transactions [598]</a>	numeric	GLOBAL
<a href="#">Slave_running [598]</a>	boolean	GLOBAL
<a href="#">Slow_launch_threads [598]</a>	numeric	GLOBAL   SESSION
<a href="#">Slow_queries [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Sort_merge_passes [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Sort_range [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Sort_rows [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Sort_scan [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_accept_renegotiates [599]</a>	numeric	GLOBAL
<a href="#">Ssl_accepts [599]</a>	numeric	GLOBAL
<a href="#">Ssl_callback_cache_hits [599]</a>	numeric	GLOBAL
<a href="#">Ssl_cipher [599]</a>	string	GLOBAL   SESSION
<a href="#">Ssl_cipher_list [599]</a>	string	GLOBAL   SESSION
<a href="#">Ssl_client_connects [599]</a>	numeric	GLOBAL

Variable Name	Variable Type	Variable Scope
<a href="#">Ssl_connect_renegotiates [599]</a>	numeric	GLOBAL
<a href="#">Ssl_ctx_verify_depth [599]</a>	numeric	GLOBAL
<a href="#">Ssl_ctx_verify_mode [599]</a>	numeric	GLOBAL
<a href="#">Ssl_default_timeout [599]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_finished_accepts [599]</a>	numeric	GLOBAL
<a href="#">Ssl_finished_connects [600]</a>	numeric	GLOBAL
<a href="#">Ssl_server_not_after [600]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_server_not_before [600]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_session_cache_hits [600]</a>	numeric	GLOBAL
<a href="#">Ssl_session_cache_misses [600]</a>	numeric	GLOBAL
<a href="#">Ssl_session_cache_mode [600]</a>	string	GLOBAL
<a href="#">Ssl_session_cache_overflows [600]</a>	numeric	GLOBAL
<a href="#">Ssl_session_cache_size [600]</a>	numeric	GLOBAL
<a href="#">Ssl_session_cache_timeouts [600]</a>	numeric	GLOBAL
<a href="#">Ssl_sessions_reused [600]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_used_session_cache_entries [600]</a>	numeric	GLOBAL
<a href="#">Ssl_verify_depth [600]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_verify_mode [600]</a>	numeric	GLOBAL   SESSION
<a href="#">Ssl_version [600]</a>	string	GLOBAL   SESSION
<a href="#">Table_locks_immediate [600]</a>	numeric	GLOBAL
<a href="#">Table_locks_waited [600]</a>	numeric	GLOBAL
<a href="#">Table_open_cache_hits [601]</a>	numeric	GLOBAL   SESSION
<a href="#">Table_open_cache_misses [601]</a>	numeric	GLOBAL   SESSION
<a href="#">Table_open_cache_overflows [601]</a>	numeric	GLOBAL   SESSION
<a href="#">Tc_log_max_pages_used [601]</a>	numeric	GLOBAL
<a href="#">Tc_log_page_size [601]</a>	numeric	GLOBAL
<a href="#">Tc_log_page_waits [601]</a>	numeric	GLOBAL
<a href="#">Threads_cached [601]</a>	numeric	GLOBAL
<a href="#">Threads_connected [601]</a>	numeric	GLOBAL
<a href="#">Threads_created [601]</a>	numeric	GLOBAL
<a href="#">Threads_running [601]</a>	numeric	GLOBAL
<a href="#">Uptime [602]</a>	numeric	GLOBAL
<a href="#">Uptime_since_flush_status [602]</a>	numeric	GLOBAL

The status variables have the following meanings.

- [Aborted\\_clients \[585\]](#)

The number of connections that were aborted because the client died without closing the connection properly. See [Section C.5.2.11, “Communication Errors and Aborted Connections”](#).

- [Aborted\\_connects \[585\]](#)

The number of failed attempts to connect to the MySQL server. See [Section C.5.2.11, “Communication Errors and Aborted Connections”](#).



For additional connection-related information, check the `Connection_errors_xxx` [587] status variables and the `host_cache` table.

- `Binlog_cache_disk_use` [586]

The number of transactions that used the temporary binary log cache but that exceeded the value of `binlog_cache_size` [2008] and used a temporary file to store statements from the transaction.

The number of nontransactional statements that caused the binary log transaction cache to be written to disk is tracked separately in the `Binlog_stmt_cache_disk_use` [586] status variable.

- `Binlog_cache_use` [586]

The number of transactions that used the binary log cache.

- `Binlog_stmt_cache_disk_use` [586]

The number of nontransaction statements that used the binary log statement cache but that exceeded the value of `binlog_stmt_cache_size` [2014] and used a temporary file to store those statements.

- `Binlog_stmt_cache_use` [586]

The number of nontransactional statements that used the binary log statement cache.

- `Bytes_received` [586]

The number of bytes received from all clients.

- `Bytes_sent` [586]

The number of bytes sent to all clients.

- `Com_xxx`

The `Com_xxx` statement counter variables indicate the number of times each `xxx` statement has been executed. There is one status variable for each type of statement. For example, `Com_delete` and `Com_update` count `DELETE` and `UPDATE` statements, respectively. `Com_delete_multi` and `Com_update_multi` are similar but apply to `DELETE` and `UPDATE` statements that use multiple-table syntax.

If a query result is returned from query cache, the server increments the `Qcache_hits` [595] status variable, not `Com_select`. See [Section 8.9.3.4, “Query Cache Status and Maintenance”](#).

All of the `Com_stmt_xxx` variables are increased even if a prepared statement argument is unknown or an error occurred during execution. In other words, their values correspond to the number of requests issued, not to the number of requests successfully completed.

The `Com_stmt_xxx` status variables are as follows:

- `Com_stmt_prepare`
- `Com_stmt_execute`
- `Com_stmt_fetch`
- `Com_stmt_send_long_data`
- `Com_stmt_reset`
- `Com_stmt_close`

Those variables stand for prepared statement commands. Their names refer to the `COM_xxx` command set used in the network layer. In other words, their values increase whenever prepared statement API calls such as `mysql_stmt_prepare()`, `mysql_stmt_execute()`, and so forth are executed. However, `Com_stmt_prepare`, `Com_stmt_execute` and `Com_stmt_close` also increase for `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE`, respectively. Additionally, the values of the older statement counter variables `Com_prepare_sql`, `Com_execute_sql`, and `Com_dealloc_sql` increase for the `PREPARE`, `EXECUTE`, and `DEALLOCATE PREPARE` statements. `Com_stmt_fetch` stands for the total number of network round-trips issued when fetching from cursors.

`Com_stmt_reprepare` indicates the number of times statements were automatically reprepared by the server after metadata changes to tables or views referred to by the statement. A reprepare operation increments `Com_stmt_reprepare`, and also `Com_stmt_prepare`.

- [Compression \[587\]](#)

Whether the client connection uses compression in the client/server protocol.

- [Connection\\_errors\\_xxx \[587\]](#)

These variables provide information about errors that occur during the client connection process. They are global only and represent error counts aggregated across connections from all hosts. These variables track errors not accounted for by the host cache (see [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#)), such as errors that are not associated with TCP connections, occur very early in the connection process (even before an IP address is known), or are not specific to any particular IP address (such as out-of-memory conditions). These variables were added in MySQL 5.6.5.

- [Connection\\_errors\\_accept \[587\]](#)

The number of errors that occurred during calls to `accept()` on the listening port.

- [Connection\\_errors\\_internal \[587\]](#)

The number of connections refused due to internal errors in the server, such as failure to start a new thread or an out-of-memory condition.

- [Connection\\_errors\\_max\\_connections \[587\]](#)

The number of connections refused because the server `max_connections` [\[503\]](#) limit was reached.

- [Connection\\_errors\\_peer\\_addr \[587\]](#)

The number of errors that occurred while searching for connecting client IP addresses.

- [Connection\\_errors\\_select \[587\]](#)

The number of errors that occurred during calls to `select()` or `poll()` on the listening port. (Failure of this operation does not necessarily mean a client connection was rejected.)

- [Connection\\_errors\\_tcpwrap \[587\]](#)

The number of connections refused by the `libwrap` library.

- [Connections \[587\]](#)

The number of connection attempts (successful or not) to the MySQL server.

- [Created\\_tmp\\_disk\\_tables \[587\]](#)

The number of internal on-disk temporary tables created by the server while executing statements.

If an internal temporary table is created initially as an in-memory table but becomes too large, MySQL automatically converts it to an on-disk table. The maximum size for in-memory temporary tables is the minimum of the `tmp_table_size` [558] and `max_heap_table_size` [504] values. If `Created_tmp_disk_tables` [587] is large, you may want to increase the `tmp_table_size` [558] or `max_heap_table_size` [504] value to lessen the likelihood that internal temporary tables in memory will be converted to on-disk tables.

You can compare the number of internal on-disk temporary tables created to the total number of internal temporary tables created by comparing the values of the `Created_tmp_disk_tables` [587] and `Created_tmp_tables` [588] variables.

See also [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).

- `Created_tmp_files` [588]

How many temporary files `mysqld` has created.

- `Created_tmp_tables` [588]

The number of internal temporary tables created by the server while executing statements.

You can compare the number of internal on-disk temporary tables created to the total number of internal temporary tables created by comparing the values of the `Created_tmp_disk_tables` [587] and `Created_tmp_tables` [588] variables.

See also [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).

Each invocation of the `SHOW STATUS` statement uses an internal temporary table and increments the global `Created_tmp_tables` [588] value.

- `Delayed_errors` [588]

The number of rows written with `INSERT DELAYED` for which some error occurred (probably `duplicate key`).

As of MySQL 5.6.7, this status variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `Delayed_insert_threads` [588]

The number of `INSERT DELAYED` handler threads in use for nontransactional tables.

As of MySQL 5.6.7, this status variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `Delayed_writes` [588]

The number of `INSERT DELAYED` rows written to nontransactional tables.

As of MySQL 5.6.7, this status variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- `Flush_commands` [588]

The number of times the server flushes tables, whether because a user executed a `FLUSH TABLES` statement or due to internal server operation. It is also incremented by receipt of a `COM_REFRESH` packet. This is in contrast to `Com_flush` [586], which indicates how many `FLUSH` statements have been executed, whether `FLUSH TABLES`, `FLUSH LOGS`, and so forth.

- `Handler_commit` [588]

The number of internal `COMMIT` statements.

- [Handler\\_delete \[589\]](#)

The number of times that rows have been deleted from tables.

- [Handler\\_external\\_lock \[589\]](#)

The server increments this variable for each call to its `external_lock()` function, which generally occurs at the beginning and end of access to a table instance. There might be differences among storage engines. This variable can be used, for example, to discover for a statement that accesses a partitioned table how many partitions were pruned before locking occurred: Check how much the counter increased for the statement, subtract 2 (2 calls for the table itself), then divide by 2 to get the number of partitions locked. This variable was added in MySQL 5.6.2.

- [Handler\\_mrr\\_init \[589\]](#)

The number of times the server uses a storage engine's own Multi-Range Read implementation for table access. This variable was added in MySQL 5.6.1.

- [Handler\\_prepare \[589\]](#)

A counter for the prepare phase of two-phase commit operations.

- [Handler\\_read\\_first \[589\]](#)

The number of times the first entry in an index was read. If this value is high, it suggests that the server is doing a lot of full index scans; for example, `SELECT col1 FROM foo`, assuming that `col1` is indexed.

- [Handler\\_read\\_key \[589\]](#)

The number of requests to read a row based on a key. If this value is high, it is a good indication that your tables are properly indexed for your queries.

- [Handler\\_read\\_last \[589\]](#)

The number of requests to read the last key in an index. With `ORDER BY`, the server will issue a first-key request followed by several next-key requests, whereas with `ORDER BY DESC`, the server will issue a last-key request followed by several previous-key requests. This variable was added in MySQL 5.6.1.

- [Handler\\_read\\_next \[589\]](#)

The number of requests to read the next row in key order. This value is incremented if you are querying an index column with a range constraint or if you are doing an index scan.

- [Handler\\_read\\_prev \[589\]](#)

The number of requests to read the previous row in key order. This read method is mainly used to optimize `ORDER BY ... DESC`.

- [Handler\\_read\\_rnd \[589\]](#)

The number of requests to read a row based on a fixed position. This value is high if you are doing a lot of queries that require sorting of the result. You probably have a lot of queries that require MySQL to scan entire tables or you have joins that do not use keys properly.

- [Handler\\_read\\_rnd\\_next \[589\]](#)

The number of requests to read the next row in the data file. This value is high if you are doing a lot of table scans. Generally this suggests that your tables are not properly indexed or that your queries are not written to take advantage of the indexes you have.

- [Handler\\_rollback \[590\]](#)

The number of requests for a storage engine to perform a rollback operation.

- [Handler\\_savepoint \[590\]](#)

The number of requests for a storage engine to place a savepoint.

- [Handler\\_savepoint\\_rollback \[590\]](#)

The number of requests for a storage engine to roll back to a savepoint.

- [Handler\\_update \[590\]](#)

The number of requests to update a row in a table.

- [Handler\\_write \[590\]](#)

The number of requests to insert a row in a table.

- [Innodb\\_available\\_undo\\_logs \[590\]](#)

The total number of available [InnoDB undo logs](#). Supplements the [innodb\\_undo\\_logs \[1752\]](#) system variable, which reports the number of active undo logs.

- [Innodb\\_buffer\\_pool\\_dump\\_status \[590\]](#)

The progress of an operation to record the [pages](#) held in the [InnoDB buffer pool](#), triggered by the setting of [innodb\\_buffer\\_pool\\_dump\\_at\\_shutdown](#) or [innodb\\_buffer\\_pool\\_dump\\_now](#).

- [Innodb\\_buffer\\_pool\\_load\\_status \[590\]](#)

The progress of an operation to [warm up](#) the [InnoDB buffer pool](#) by reading in a set of [pages](#) corresponding to an earlier point in time, triggered by the setting of [innodb\\_buffer\\_pool\\_load\\_at\\_startup \[1707\]](#) or [innodb\\_buffer\\_pool\\_load\\_now \[1708\]](#). If the operation introduces too much overhead, you can cancel it by setting [innodb\\_buffer\\_pool\\_load\\_abort \[1707\]](#).

- [Innodb\\_buffer\\_pool\\_bytes\\_data \[590\]](#)

The total number of bytes in the [InnoDB buffer pool](#) containing data. The number includes both [dirty](#) and clean pages. For more accurate memory usage calculations than with [Innodb\\_buffer\\_pool\\_pages\\_data \[590\]](#), when [compressed](#) tables cause the buffer pool to hold pages of different sizes.

- [Innodb\\_buffer\\_pool\\_pages\\_data \[590\]](#)

The number of [pages](#) in the [InnoDB buffer pool](#) containing data. The number includes both [dirty](#) and clean pages.

- [Innodb\\_buffer\\_pool\\_bytes\\_dirty \[590\]](#)

The total current number of bytes held in [dirty pages](#) in the [InnoDB buffer pool](#). For more accurate memory usage calculations than with [Innodb\\_buffer\\_pool\\_pages\\_dirty \[590\]](#), when [compressed](#) tables cause the buffer pool to hold pages of different sizes.

- [Innodb\\_buffer\\_pool\\_pages\\_dirty \[590\]](#)

The current number of [dirty pages](#) in the [InnoDB buffer pool](#).

- [Innodb\\_buffer\\_pool\\_pages\\_flushed \[591\]](#)

The number of requests to [flush pages](#) from the [InnoDB buffer pool](#).

- [Innodb\\_buffer\\_pool\\_pages\\_free \[591\]](#)

The number of free [pages](#) in the [InnoDB buffer pool](#).

- [Innodb\\_buffer\\_pool\\_pages\\_latched \[591\]](#)

The number of latched [pages](#) in the [InnoDB buffer pool](#). These are pages currently being read or written, or that cannot be [flushed](#) or removed for some other reason. Calculation of this variable is expensive, so it is available only when the [UNIV\\_DEBUG](#) system is defined at server build time.

- [Innodb\\_buffer\\_pool\\_pages\\_misc \[591\]](#)

The number of [pages](#) in the [InnoDB buffer pool](#) that are busy because they have been allocated for administrative overhead, such as [row locks](#) or the [adaptive hash index](#). This value can also be calculated as [Innodb\\_buffer\\_pool\\_pages\\_total \[591\]](#) – [Innodb\\_buffer\\_pool\\_pages\\_free \[591\]](#) – [Innodb\\_buffer\\_pool\\_pages\\_data \[590\]](#).

- [Innodb\\_buffer\\_pool\\_pages\\_total \[591\]](#)

The total size of the [InnoDB buffer pool](#), in [pages](#).

- [Innodb\\_buffer\\_pool\\_read\\_ahead \[591\]](#)

The number of [pages](#) read into the [InnoDB buffer pool](#) by the [read-ahead](#) background thread.

- [Innodb\\_buffer\\_pool\\_read\\_ahead\\_evicted \[591\]](#)

The number of [pages](#) read into the [InnoDB buffer pool](#) by the [read-ahead](#) background thread that were subsequently [evicted](#) without having been accessed by queries.

- [Innodb\\_buffer\\_pool\\_read\\_requests \[591\]](#)

The number of logical read requests made to the [InnoDB buffer pool](#). These requests could be serviced by returning data that was already in memory, or by reading the data from disk into memory first.

- [Innodb\\_buffer\\_pool\\_reads \[591\]](#)

The number of logical reads that [InnoDB](#) could not satisfy from the [buffer pool](#), and had to read directly from disk.

- [Innodb\\_buffer\\_pool\\_wait\\_free \[591\]](#)

Normally, writes to the [InnoDB buffer pool](#) happen in the background. When [InnoDB](#) needs to read or create a [page](#) and no clean pages are available, [InnoDB](#) flushes some [dirty pages](#) first and waits for that operation to finish. This counter counts instances of these waits. If [innodb\\_buffer\\_pool\\_size \[1708\]](#) has been set properly, this value should be small.

- [Innodb\\_buffer\\_pool\\_write\\_requests \[591\]](#)

The number of writes done to the [InnoDB buffer pool](#).

- [Innodb\\_data\\_fsyncs \[591\]](#)

The number of `fsync()` operations so far. The frequency of `fsync()` calls is influenced by the setting of the [innodb\\_flush\\_method \[1719\]](#) configuration option.

- [Innodb\\_data\\_pending\\_fsyncs \[591\]](#)
-

The current number of pending `fsync()` operations. The frequency of `fsync()` calls is influenced by the setting of the `innodb_flush_method` [1719] configuration option.

- `Innodb_data_pending_reads` [592]

The current number of pending reads.

- `Innodb_data_pending_writes` [592]

The current number of pending writes.

- `Innodb_data_read` [592]

The amount of data read since the server was started.

- `Innodb_data_reads` [592]

The total number of data reads.

- `Innodb_data_writes` [592]

The total number of data writes.

- `Innodb_data_written` [592]

The amount of data written so far, in bytes.

- `Innodb_dblwr_pages_written` [592]

The number of `pages` that have been written to the `doublewrite` buffer. See [Section 5.3.1, “InnoDB Disk I/O”](#).

- `Innodb_dblwr_writes` [592]

The number of `doublewrite` operations that have been performed. See [Section 5.3.1, “InnoDB Disk I/O”](#).

- `Innodb_have_atomic_builtins` [592]

Indicates whether the server was built with `atomic instructions`.

- `Innodb_log_waits` [592]

The number of times that the `log buffer` was too small and a `wait` was required for it to be `flushed` before continuing.

- `Innodb_log_write_requests` [592]

The number of write requests for the `InnoDB redo log`.

- `Innodb_log_writes` [592]

The number of physical writes to the `InnoDB redo log` file.

- `Innodb_num_open_files` [592]

The number of files `InnoDB` currently holds open.

- `Innodb_os_log_fsyncs` [592]

The number of `fsync()` writes done to the `InnoDB redo log` files.

- `Innodb_os_log_pending_fsyncs` [592]

The number of pending `fsync()` operations for the `InnoDB redo log` files.

- `InnoDB_os_log_pending_writes` [593]

The number of pending writes to the `InnoDB redo log` files.

- `InnoDB_os_log_written` [593]

The number of bytes written to the `InnoDB redo log` files.

- `InnoDB_page_size` [593]

The compiled-in `InnoDB` page size (default 16KB). Many values are counted in pages; the page size enables them to be easily converted to bytes.

- `InnoDB_pages_created` [593]

The number of pages created by operations on `InnoDB` tables.

- `InnoDB_pages_read` [593]

The number of pages read by operations on `InnoDB` tables.

- `InnoDB_pages_written` [593]

The number of pages written by operations on `InnoDB` tables.

- `InnoDB_row_lock_current_waits` [593]

The number of `row locks` currently being waited for by operations on `InnoDB` tables.

- `InnoDB_row_lock_time` [593]

The total time spent in acquiring `row locks` for `InnoDB` tables, in milliseconds.

- `InnoDB_row_lock_time_avg` [593]

The average time to acquire a `row lock` for `InnoDB` tables, in milliseconds.

- `InnoDB_row_lock_time_max` [593]

The maximum time to acquire a `row lock` for `InnoDB` tables, in milliseconds.

- `InnoDB_row_lock_waits` [593]

The number of times operations on `InnoDB` tables had to wait for a `row lock`.

- `InnoDB_rows_deleted` [593]

The number of rows deleted from `InnoDB` tables.

- `InnoDB_rows_inserted` [593]

The number of rows inserted into `InnoDB` tables.

- `InnoDB_rows_read` [593]

The number of rows read from `InnoDB` tables.

- `InnoDB_rows_updated` [593]

The number of rows updated in `InnoDB` tables.

- `InnoDB_truncated_status_writes` [593]



The number of times output from the `SHOW ENGINE INNODB STATUS` statement has been truncated.

- [Key\\_blocks\\_not\\_flushed \[594\]](#)

The number of key blocks in the `MyISAM` key cache that have changed but have not yet been flushed to disk.

- [Key\\_blocks\\_unused \[594\]](#)

The number of unused blocks in the `MyISAM` key cache. You can use this value to determine how much of the key cache is in use; see the discussion of [key\\_buffer\\_size \[492\]](#) in [Section 5.1.4, “Server System Variables”](#).

- [Key\\_blocks\\_used \[594\]](#)

The number of used blocks in the `MyISAM` key cache. This value is a high-water mark that indicates the maximum number of blocks that have ever been in use at one time.

- [Key\\_read\\_requests \[594\]](#)

The number of requests to read a key block from the `MyISAM` key cache.

- [Key\\_reads \[594\]](#)

The number of physical reads of a key block from disk into the `MyISAM` key cache. If [Key\\_reads \[594\]](#) is large, then your [key\\_buffer\\_size \[492\]](#) value is probably too small. The cache miss rate can be calculated as  $\text{Key\_reads [594]} / \text{Key\_read\_requests [594]}$ .

- [Key\\_write\\_requests \[594\]](#)

The number of requests to write a key block to the `MyISAM` key cache.

- [Key\\_writes \[594\]](#)

The number of physical writes of a key block from the `MyISAM` key cache to disk.

- [Last\\_query\\_cost \[594\]](#)

The total cost of the last compiled query as computed by the query optimizer. This is useful for comparing the cost of different query plans for the same query. The default value of 0 means that no query has been compiled yet. The default value is 0. [Last\\_query\\_cost \[594\]](#) has session scope.

The [Last\\_query\\_cost \[594\]](#) value can be computed accurately only for simple “flat” queries, not complex queries such as those with subqueries or `UNION`. For the latter, the value is set to 0.

- [Last\\_query\\_partial\\_plans \[594\]](#)

The number of iterations the query optimizer made in execution plan construction for the previous query. [Last\\_query\\_cost \[594\]](#) has session scope. This variable was added in MySQL 5.6.5.

- [Max\\_used\\_connections \[594\]](#)

The maximum number of connections that have been in use simultaneously since the server started.

- [Not\\_flushed\\_delayed\\_rows \[594\]](#)

The number of rows waiting to be written to nontransactional tables in `INSERT DELAYED` queues.

As of MySQL 5.6.7, this status variable is deprecated (because `DELAYED` inserts are deprecated), and will be removed in a future release.

- [Open\\_files \[594\]](#)

The number of files that are open. This count includes regular files opened by the server. It does not include other types of files such as sockets or pipes. Also, the count does not include files that storage engines open using their own internal functions rather than asking the server level to do so.

- [Open\\_streams \[595\]](#)

The number of streams that are open (used mainly for logging).

- [Open\\_table\\_definitions \[595\]](#)

The number of cached `.frm` files.

- [Open\\_tables \[595\]](#)

The number of tables that are open.

- [Opened\\_files \[595\]](#)

The number of files that have been opened with `my_open()` (a `mysys` library function). Parts of the server that open files without using this function do not increment the count.

- [Opened\\_table\\_definitions \[595\]](#)

The number of `.frm` files that have been cached.

- [Opened\\_tables \[595\]](#)

The number of tables that have been opened. If [Opened\\_tables \[595\]](#) is big, your [table\\_open\\_cache \[551\]](#) value is probably too small.

- [Performance\\_schema\\_xxx](#)

Performance Schema status variables are listed in [Section 20.13, "Performance Schema Status Variables"](#). These variables provide information about instrumentation that could not be loaded or created due to memory constraints.

- [Prepared\\_stmt\\_count \[595\]](#)

The current number of prepared statements. (The maximum number of statements is given by the [max\\_prepared\\_stmt\\_count \[506\]](#) system variable.)

- [Qcache\\_free\\_blocks \[595\]](#)

The number of free memory blocks in the query cache.

- [Qcache\\_free\\_memory \[595\]](#)

The amount of free memory for the query cache.

- [Qcache\\_hits \[595\]](#)

The number of query cache hits.

- [Qcache\\_inserts \[595\]](#)

The number of queries added to the query cache.

- [Qcache\\_lowmem\\_prunes \[595\]](#)

The number of queries that were deleted from the query cache because of low memory.

- [Qcache\\_not\\_cached \[595\]](#)

The number of noncached queries (not cacheable, or not cached due to the `query_cache_type` [529] setting).

- `Qcache_queries_in_cache` [596]

The number of queries registered in the query cache.

- `Qcache_total_blocks` [596]

The total number of blocks in the query cache.

- `Queries` [596]

The number of statements executed by the server. This variable includes statements executed within stored programs, unlike the `Questions` variable. It does not count `COM_PING` or `COM_STATISTICS` commands.

- `Questions` [596]

The number of statements executed by the server. This includes only statements sent to the server by clients and not statements executed within stored programs, unlike the `Queries` variable. This variable does not count `COM_PING`, `COM_STATISTICS`, `COM_STMT_PREPARE`, `COM_STMT_CLOSE`, or `COM_STMT_RESET` commands.

- `Rpl_semi_sync_master_clients` [596]

The number of semisynchronous slaves.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_net_avg_wait_time` [596]

The average time in microseconds the master waited for a slave reply.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_net_wait_time` [596]

The total time in microseconds the master waited for slave replies.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_net_waits` [596]

The total number of times the master waited for slave replies.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_no_times` [596]

The number of times the master turned off semisynchronous replication.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_no_tx` [596]

The number of commits that were not acknowledged successfully by a slave.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- `Rpl_semi_sync_master_status` [596]

Whether semisynchronous replication currently is operational on the master. The value is `ON` if the plugin has been enabled and a commit acknowledgment has occurred. It is `OFF` if the plugin is not enabled or the master has fallen back to asynchronous replication due to commit acknowledgment timeout.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_timefunc\\_failures \[597\]](#)

The number of times the master failed when calling time functions such as `gettimeofday()`.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_tx\\_avg\\_wait\\_time \[597\]](#)

The average time in microseconds the master waited for each transaction.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_tx\\_wait\\_time \[597\]](#)

The total time in microseconds the master waited for transactions.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_tx\\_waits \[597\]](#)

The total number of times the master waited for transactions.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_wait\\_pos\\_backtraverse \[597\]](#)

The total number of times the master waited for an event with binary coordinates lower than events waited for previously. This can occur when the order in which transactions start waiting for a reply is different from the order in which their binary log events are written.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_wait\\_sessions \[597\]](#)

The number of sessions currently waiting for slave replies.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_master\\_yes\\_tx \[597\]](#)

The number of commits that were acknowledged successfully by a slave.

This variable is available only if the master-side semisynchronous replication plugin is installed.

- [Rpl\\_semi\\_sync\\_slave\\_status \[597\]](#)

Whether semisynchronous replication currently is operational on the slave. This is `ON` if the plugin has been enabled and the slave I/O thread is running, `OFF` otherwise.

This variable is available only if the slave-side semisynchronous replication plugin is installed.

- [Rsa\\_public\\_key \[597\]](#)

The RSA public key value used by the `sha256_password` authentication plugin.

The value is nonempty only if the server successfully initializes the private and public

keys in the files named by the `sha256_password_private_key_path` [539] and `sha256_password_public_key_path` [539] system variables. The value of `Rsa_public_key` [597] comes from the latter file.

For information about `sha256_password`, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

This variable is available only if MySQL was built using OpenSSL. It was added in MySQL 5.6.6.

- `Select_full_join` [598]

The number of joins that perform table scans because they do not use indexes. If this value is not 0, you should carefully check the indexes of your tables.

- `Select_full_range_join` [598]

The number of joins that used a range search on a reference table.

- `Select_range` [598]

The number of joins that used ranges on the first table. This is normally not a critical issue even if the value is quite large.

- `Select_range_check` [598]

The number of joins without keys that check for key usage after each row. If this is not 0, you should carefully check the indexes of your tables.

- `Select_scan` [598]

The number of joins that did a full scan of the first table.

- `Slave_heartbeat_period` [598]

Shows the replication heartbeat interval (in seconds) on a replication slave.

- `Slave_last_heartbeat` [598]

Shows when the most recent heartbeat signal was received by a replication slave, as a `TIMESTAMP` value.

- `Slave_open_temp_tables` [598]

The number of temporary tables that the slave SQL thread currently has open. If the value is greater than zero, it is not safe to shut down the slave; see [Section 16.4.1.21, “Replication and Temporary Tables”](#).

- `Slave_received_heartbeats` [598]

This counter increments with each replication heartbeat received by a replication slave since the last time that the slave was restarted or reset, or a `CHANGE MASTER TO` statement was issued.

- `Slave_retried_transactions` [598]

The total number of times since startup that the replication slave SQL thread has retried transactions.

- `Slave_running` [598]

This is `ON` if this server is a replication slave that is connected to a replication master, and both the I/O and SQL threads are running; otherwise, it is `OFF`.

- `Slow_launch_threads` [598]

The number of threads that have taken more than `slow_launch_time` [541] seconds to create.

- [Slow\\_queries \[599\]](#)

The number of queries that have taken more than [long\\_query\\_time \[500\]](#) seconds. This counter increments regardless of whether the slow query log is enabled. For information about that log, see [Section 5.2.5, “The Slow Query Log”](#).

- [Sort\\_merge\\_passes \[599\]](#)

The number of merge passes that the sort algorithm has had to do. If this value is large, you should consider increasing the value of the [sort\\_buffer\\_size \[542\]](#) system variable.

- [Sort\\_range \[599\]](#)

The number of sorts that were done using ranges.

- [Sort\\_rows \[599\]](#)

The number of sorted rows.

- [Sort\\_scan \[599\]](#)

The number of sorts that were done by scanning the table.

- [Ssl\\_accept\\_renegotiates \[599\]](#)

The number of negotiates needed to establish the connection.

- [Ssl\\_accepts \[599\]](#)

The number of accepted SSL connections.

- [Ssl\\_callback\\_cache\\_hits \[599\]](#)

The number of callback cache hits.

- [Ssl\\_cipher \[599\]](#)

The current SSL cipher (empty for non-SSL connections).

- [Ssl\\_cipher\\_list \[599\]](#)

The list of possible SSL ciphers.

- [Ssl\\_client\\_connects \[599\]](#)

The number of SSL connection attempts to an SSL-enabled master.

- [Ssl\\_connect\\_renegotiates \[599\]](#)

The number of negotiates needed to establish the connection to an SSL-enabled master.

- [Ssl\\_ctx\\_verify\\_depth \[599\]](#)

The SSL context verification depth (how many certificates in the chain are tested).

- [Ssl\\_ctx\\_verify\\_mode \[599\]](#)

The SSL context verification mode.

- [Ssl\\_default\\_timeout \[599\]](#)

The default SSL timeout.

- [Ssl\\_finished\\_accepts \[599\]](#)

The number of successful SSL connections to the server.

- [Ssl\\_finished\\_connects \[600\]](#)

The number of successful slave connections to an SSL-enabled master.

- [Ssl\\_server\\_not\\_after \[600\]](#)

The last date for which the SSL certificate is valid. This variable was added in MySQL 5.6.3.

- [Ssl\\_server\\_not\\_before \[600\]](#)

The first date for which the SSL certificate is valid. This variable was added in MySQL 5.6.3.

- [Ssl\\_session\\_cache\\_hits \[600\]](#)

The number of SSL session cache hits.

- [Ssl\\_session\\_cache\\_misses \[600\]](#)

The number of SSL session cache misses.

- [Ssl\\_session\\_cache\\_mode \[600\]](#)

The SSL session cache mode.

- [Ssl\\_session\\_cache\\_overflows \[600\]](#)

The number of SSL session cache overflows.

- [Ssl\\_session\\_cache\\_size \[600\]](#)

The SSL session cache size.

- [Ssl\\_session\\_cache\\_timeouts \[600\]](#)

The number of SSL session cache timeouts.

- [Ssl\\_sessions\\_reused \[600\]](#)

How many SSL connections were reused from the cache.

- [Ssl\\_used\\_session\\_cache\\_entries \[600\]](#)

How many SSL session cache entries were used.

- [Ssl\\_verify\\_depth \[600\]](#)

The verification depth for replication SSL connections.

- [Ssl\\_verify\\_mode \[600\]](#)

The verification mode for replication SSL connections.

- [Ssl\\_version \[600\]](#)

The SSL protocol version of the connection.

- [Table\\_locks\\_immediate \[600\]](#)

The number of times that a request for a table lock could be granted immediately.

- [Table\\_locks\\_waited \[600\]](#)

The number of times that a request for a table lock could not be granted immediately and a wait was needed. If this is high and you have performance problems, you should first optimize your queries, and then either split your table or tables or use replication.

- [Table\\_open\\_cache\\_hits](#) [601]

The number of hits for open tables cache lookups. This variable was added in MySQL 5.6.6.

- [Table\\_open\\_cache\\_misses](#) [601]

The number of misses for open tables cache lookups. This variable was added in MySQL 5.6.6.

- [Table\\_open\\_cache\\_overflows](#) [601]

The number of overflows for the open tables cache. This is the number of times, after a table is opened or closed, a cache instance has an unused entry and the size of the instance is larger than [table\\_open\\_cache](#) [551] / [table\\_open\\_cache\\_instances](#) [552]. This variable was added in MySQL 5.6.6.

- [Tc\\_log\\_max\\_pages\\_used](#) [601]

For the memory-mapped implementation of the log that is used by `mysqld` when it acts as the transaction coordinator for recovery of internal XA transactions, this variable indicates the largest number of pages used for the log since the server started. If the product of [Tc\\_log\\_max\\_pages\\_used](#) [601] and [Tc\\_log\\_page\\_size](#) [601] is always significantly less than the log size, the size is larger than necessary and can be reduced. (The size is set by the `--log-tc-size` [429] option. Currently, this variable is unused: It is unneeded for binary log-based recovery, and the memory-mapped recovery log method is not used unless the number of storage engines capable of two-phase commit is greater than one. (InnoDB is the only applicable engine.)

- [Tc\\_log\\_page\\_size](#) [601]

The page size used for the memory-mapped implementation of the XA recovery log. The default value is determined using `getpagesize()`. Currently, this variable is unused for the same reasons as described for [Tc\\_log\\_max\\_pages\\_used](#) [601].

- [Tc\\_log\\_page\\_waits](#) [601]

For the memory-mapped implementation of the recovery log, this variable increments each time the server was not able to commit a transaction and had to wait for a free page in the log. If this value is large, you might want to increase the log size (with the `--log-tc-size` [429] option). For binary log-based recovery, this variable increments each time the binary log cannot be closed because there are two-phase commits in progress. (The close operation waits until all such transactions are finished.)

- [Threads\\_cached](#) [601]

The number of threads in the thread cache.

- [Threads\\_connected](#) [601]

The number of currently open connections.

- [Threads\\_created](#) [601]

The number of threads created to handle connections. If [Threads\\_created](#) [601] is big, you may want to increase the [thread\\_cache\\_size](#) [552] value. The cache miss rate can be calculated as [Threads\\_created](#) [601]/[Connections](#) [587].

- [Threads\\_running](#) [601]

The number of threads that are not sleeping.



- [Uptime](#) [602]

The number of seconds that the server has been up.

- [Uptime\\_since\\_flush\\_status](#) [602]

The number of seconds since the most recent `FLUSH STATUS` statement.

## 5.1.7. Server SQL Modes

The MySQL server can operate in different SQL modes, and can apply these modes differently for different clients. This capability enables each application to tailor the server's operating mode to its own requirements.

For answers to some questions that are often asked about server SQL modes in MySQL, see [Section B.3, “MySQL 5.6 FAQ: Server SQL Mode”](#).

Modes define what SQL syntax MySQL should support and what kind of data validation checks it should perform. This makes it easier to use MySQL in different environments and to use MySQL together with other database servers.

When working with `InnoDB` tables, consider also the `innodb_strict_mode` [1749] configuration option. It enables additional error checks for `InnoDB` tables, as listed in [Section 14.2.5.7, “InnoDB Strict Mode”](#).

### Setting the SQL Mode

You can set the default SQL mode by starting `mysqld` with the `--sql-mode="modes"` [442] option, or by using `sql-mode="modes"` [442] in `my.cnf` (Unix operating systems) or `my.ini` (Windows). `modes` is a list of different modes separated by comma (",") characters. The default SQL mode in MySQL 5.6.6 and later is `NO_ENGINE_SUBSTITUTION`; in MySQL 5.6.5 and earlier, it was empty (no modes set). To clear the mode explicitly, set it to an empty string using `--sql-mode=""` [442] on the command line, or with `sql-mode=""` [442] in `my.cnf` on Unix systems or in `my.ini` on Windows.

You can change the SQL mode at runtime by using a `SET [GLOBAL|SESSION] sql_mode='modes'` statement to set the `sql_mode` [545] system value. Setting the `GLOBAL` variable requires the `SUPER` [756] privilege and affects the operation of all clients that connect from that time on. Setting the `SESSION` variable affects only the current client. Any client can change its own session `sql_mode` [545] value at any time.



#### Important

**SQL mode and user-defined partitioning.** Changing the server SQL mode after creating and inserting data into partitioned tables can cause major changes in the behavior of such tables, and could lead to loss or corruption of data. It is strongly recommended that you never change the SQL mode once you have created tables employing user-defined partitioning.

When replicating partitioned tables, differing SQL modes on master and slave can also lead to problems. For best results, you should always use the same server SQL mode on the master and on the slave.

See [Section 17.6, “Restrictions and Limitations on Partitioning”](#), for more information.

You can retrieve the current global or session `sql_mode` [545] value with the following statements:

```
SELECT @@GLOBAL.sql_mode;
SELECT @@SESSION.sql_mode;
```

### Most Important SQL Modes

The most important `sql_mode` [545] values are probably these:

- [ANSI \[608\]](#)

This mode changes syntax and behavior to conform more closely to standard SQL. It is one of the special [combination modes](#) listed at the end of this section.

- [STRICT\\_TRANS\\_TABLES \[607\]](#)

If a value could not be inserted as given into a transactional table, abort the statement. For a nontransactional table, abort the statement if the value occurs in a single-row statement or the first row of a multiple-row statement. More detail is given later in this section.

- [TRADITIONAL \[609\]](#)

Make MySQL behave like a “traditional” SQL database system. A simple description of this mode is “give an error instead of a warning” when inserting an incorrect value into a column. It is one of the special [combination modes](#) listed at the end of this section.



#### Note

The [INSERT/UPDATE](#) aborts as soon as the error is noticed. This may not be what you want if you are using a nontransactional storage engine, because data changes made prior to the error may not be rolled back, resulting in a “partially done” update.

When this manual refers to “strict mode,” it means a mode where at least one of [STRICT\\_TRANS\\_TABLES \[607\]](#) or [STRICT\\_ALL\\_TABLES \[607\]](#) is enabled.

## Full List of SQL Modes

The following list describes all supported modes:

- [ALLOW\\_INVALID\\_DATES \[603\]](#)

Do not perform full checking of dates. Check only that the month is in the range from 1 to 12 and the day is in the range from 1 to 31. This is very convenient for Web applications where you obtain year, month, and day in three different fields and you want to store exactly what the user inserted (without date validation). This mode applies to [DATE](#) and [DATETIME](#) columns. It does not apply [TIMESTAMP](#) columns, which always require a valid date.

The server requires that month and day values be legal, and not merely in the range 1 to 12 and 1 to 31, respectively. With strict mode disabled, invalid dates such as `'2004-04-31'` are converted to `'0000-00-00'` and a warning is generated. With strict mode enabled, invalid dates generate an error. To permit such dates, enable [ALLOW\\_INVALID\\_DATES \[603\]](#).

- [ANSI\\_QUOTES \[603\]](#)

Treat “`”` as an identifier quote character (like the “```” quote character) and not as a string quote character. You can still use “```” to quote identifiers with this mode enabled. With [ANSI\\_QUOTES \[603\]](#) enabled, you cannot use double quotation marks to quote literal strings, because it is interpreted as an identifier.

- [ERROR\\_FOR\\_DIVISION\\_BY\\_ZERO \[603\]](#)

Produce an error in strict mode (otherwise a warning) when a division by zero (or [MOD\(X, 0\) \[1205\]](#)) occurs during an [INSERT](#) or [UPDATE](#). If this mode is not enabled, MySQL instead returns [NULL](#) for divisions by zero. For [INSERT IGNORE](#) or [UPDATE IGNORE](#), MySQL generates a warning for divisions by zero, but the result of the operation is [NULL](#).

- [HIGH\\_NOT\\_PRECEDENCE \[603\]](#)

The precedence of the `NOT` [1171] operator is such that expressions such as `NOT a BETWEEN b AND c` are parsed as `NOT (a BETWEEN b AND c)`. In some older versions of MySQL, the expression was parsed as `(NOT a) BETWEEN b AND c`. The old higher-precedence behavior can be obtained by enabling the `HIGH_NOT_PRECEDENCE` [603] SQL mode.

```
mysql> SET sql_mode = '';
mysql> SELECT NOT 1 BETWEEN -5 AND 5;
-> 0
mysql> SET sql_mode = 'HIGH_NOT_PRECEDENCE';
mysql> SELECT NOT 1 BETWEEN -5 AND 5;
-> 1
```

- `IGNORE_SPACE` [604]

Permit spaces between a function name and the “(” character. This causes built-in function names to be treated as reserved words. As a result, identifiers that are the same as function names must be quoted as described in [Section 9.2, “Schema Object Names”](#). For example, because there is a `COUNT()` [1283] function, the use of `count` as a table name in the following statement causes an error:

```
mysql> CREATE TABLE count (i INT);
ERROR 1064 (42000): You have an error in your SQL syntax
```

The table name should be quoted:

```
mysql> CREATE TABLE `count` (i INT);
Query OK, 0 rows affected (0.00 sec)
```

The `IGNORE_SPACE` [604] SQL mode applies to built-in functions, not to user-defined functions or stored functions. It is always permissible to have spaces after a UDF or stored function name, regardless of whether `IGNORE_SPACE` [604] is enabled.

For further discussion of `IGNORE_SPACE` [604], see [Section 9.2.4, “Function Name Parsing and Resolution”](#).

- `NO_AUTO_CREATE_USER` [604]

Prevent the `GRANT` statement from automatically creating new users if it would otherwise do so, unless authentication information is specified. The statement must specify a nonempty password using `IDENTIFIED BY` or an authentication plugin using `IDENTIFIED WITH`.

- `NO_AUTO_VALUE_ON_ZERO` [604]

`NO_AUTO_VALUE_ON_ZERO` [604] affects handling of `AUTO_INCREMENT` columns. Normally, you generate the next sequence number for the column by inserting either `NULL` or `0` into it. `NO_AUTO_VALUE_ON_ZERO` [604] suppresses this behavior for `0` so that only `NULL` generates the next sequence number.

This mode can be useful if `0` has been stored in a table's `AUTO_INCREMENT` column. (Storing `0` is not a recommended practice, by the way.) For example, if you dump the table with `mysqldump` and then reload it, MySQL normally generates new sequence numbers when it encounters the `0` values, resulting in a table with contents different from the one that was dumped. Enabling `NO_AUTO_VALUE_ON_ZERO` [604] before reloading the dump file solves this problem. `mysqldump` now automatically includes in its output a statement that enables `NO_AUTO_VALUE_ON_ZERO` [604], to avoid this problem.

- `NO_BACKSLASH_ESCAPES` [604]

Disable the use of the backslash character (“\”) as an escape character within strings. With this mode enabled, backslash becomes an ordinary character like any other.

- `NO_DIR_IN_CREATE` [604]

When creating a table, ignore all `INDEX DIRECTORY` and `DATA DIRECTORY` directives. This option is useful on slave replication servers.

- `NO_ENGINE_SUBSTITUTION` [605]

Control automatic substitution of the default storage engine when a statement such as `CREATE TABLE` or `ALTER TABLE` specifies a storage engine that is disabled or not compiled in.

Because storage engines can be pluggable at runtime, unavailable engines are treated the same way:

With `NO_ENGINE_SUBSTITUTION` [605] disabled, for `CREATE TABLE` the default engine is used and a warning occurs if the desired engine is unavailable. For `ALTER TABLE`, a warning occurs and the table is not altered.

With `NO_ENGINE_SUBSTITUTION` [605] enabled, an error occurs and the table is not created or altered if the desired engine is unavailable.

- `NO_FIELD_OPTIONS` [605]

Do not print MySQL-specific column options in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode.

- `NO_KEY_OPTIONS` [605]

Do not print MySQL-specific index options in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode.

- `NO_TABLE_OPTIONS` [605]

Do not print MySQL-specific table options (such as `ENGINE`) in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode.

- `NO_UNSIGNED_SUBTRACTION` [605]

By default, subtraction between integer operands produces an `UNSIGNED` result if any operand is `UNSIGNED`. When `NO_UNSIGNED_SUBTRACTION` [605] is enabled, the subtraction result is signed, *even if any operand is unsigned*. For example, compare the type of column `c2` in table `t1` with that of column `c2` in table `t2`:

```
mysql> SET sql_mode='';
mysql> CREATE TABLE test (c1 BIGINT UNSIGNED NOT NULL);
mysql> CREATE TABLE t1 SELECT c1 - 1 AS c2 FROM test;
mysql> DESCRIBE t1;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| c2    | bigint(21) unsigned |      |     | 0        |       |
+-----+-----+-----+-----+-----+-----+

mysql> SET sql_mode='NO_UNSIGNED_SUBTRACTION';
mysql> CREATE TABLE t2 SELECT c1 - 1 AS c2 FROM test;
mysql> DESCRIBE t2;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| c2    | bigint(21)          |      |     | 0        |       |
+-----+-----+-----+-----+-----+-----+
```

Note that this means that `BIGINT UNSIGNED` is not 100% usable in all contexts. See [Section 12.10, “Cast Functions and Operators”](#).

```
mysql> SET sql_mode = '';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
```

```

| CAST(0 AS UNSIGNED) - 1 |
+-----+
|          18446744073709551615 |
+-----+

mysql> SET sql_mode = 'NO_UNSIGNED_SUBTRACTION';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
| CAST(0 AS UNSIGNED) - 1 |
+-----+
|                      -1 |
+-----+

```

- [NO\\_ZERO\\_DATE \[606\]](#)

In strict mode, do not permit '0000-00-00' as a valid date. You can still insert zero dates with the [IGNORE](#) option. When not in strict mode, the date is accepted but a warning is generated.

- [NO\\_ZERO\\_IN\\_DATE \[606\]](#)

In strict mode, do not accept dates where the year part is nonzero but the month or day part is 0 (for example, '0000-00-00' is legal but '2010-00-01' and '2010-01-00' are not). If used with the [IGNORE](#) option, MySQL inserts a '0000-00-00' date for any such date. When not in strict mode, the date is accepted but a warning is generated.

- [ONLY\\_FULL\\_GROUP\\_BY \[606\]](#)

Do not permit queries for which the select list or [HAVING](#) list or (as of MySQL 5.6.5) [ORDER](#) list refers to nonaggregated columns that are not named in the [GROUP BY](#) clause.

The following queries are invalid with [ONLY\\_FULL\\_GROUP\\_BY \[606\]](#) enabled. The first is invalid because [address](#) in the select list is not named in the [GROUP BY](#) clause, and the second because [max\\_age](#) in the [HAVING](#) clause is not named in the [GROUP BY](#) clause:

```

mysql> SELECT name, address, MAX(age) FROM t GROUP BY name;
ERROR 1055 (42000): 't.address' isn't in GROUP BY

```

```

mysql> SELECT name, MAX(age) AS max_age FROM t GROUP BY name
-> HAVING max_age < 30;
Empty set (0.00 sec)
ERROR 1463 (42000): Non-grouping field 'max_age' is used in HAVING clause

```

In the second example, the query could be rewritten to use [HAVING MAX\(age\)](#) instead, so that the reference is to a column named in an aggregate function. ([max\\_age](#) fails because it *is* an aggregate function.)

In addition, if a query has aggregate functions and no [GROUP BY](#) clause, it cannot have nonaggregated columns in the select list or [ORDER BY](#) list:

```

mysql> SELECT name, MAX(age) FROM t;
ERROR 1140 (42000): Mixing of GROUP columns (MIN(),MAX(),COUNT(),...)
with no GROUP columns is illegal if there is no GROUP BY clause

```

For more information, see [Section 12.17.3, “MySQL Extensions to GROUP BY”](#).

- [PAD\\_CHAR\\_TO\\_FULL\\_LENGTH \[606\]](#)

By default, trailing spaces are trimmed from [CHAR](#) column values on retrieval. If [PAD\\_CHAR\\_TO\\_FULL\\_LENGTH \[606\]](#) is enabled, trimming does not occur and retrieved [CHAR](#) values are padded to their full length. This mode does not apply to [VARCHAR](#) columns, for which trailing spaces are retained on retrieval.

```
mysql> CREATE TABLE t1 (c1 CHAR(10));
Query OK, 0 rows affected (0.37 sec)

mysql> INSERT INTO t1 (c1) VALUES('xy');
Query OK, 1 row affected (0.01 sec)

mysql> SET sql_mode = '';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT c1, CHAR_LENGTH(c1) FROM t1;
+-----+-----+
| c1   | CHAR_LENGTH(c1) |
+-----+-----+
| xy   |                2 |
+-----+-----+
1 row in set (0.00 sec)

mysql> SET sql_mode = 'PAD_CHAR_TO_FULL_LENGTH';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT c1, CHAR_LENGTH(c1) FROM t1;
+-----+-----+
| c1   | CHAR_LENGTH(c1) |
+-----+-----+
| xy   |                10 |
+-----+-----+
1 row in set (0.00 sec)
```

- [PIPES\\_AS\\_CONCAT \[607\]](#)

Treat `||` [\[1172\]](#) as a string concatenation operator (same as `CONCAT()` [\[1178\]](#)) rather than as a synonym for `OR` [\[1172\]](#).

- [REAL\\_AS\\_FLOAT \[607\]](#)

Treat `REAL` as a synonym for `FLOAT`. By default, MySQL treats `REAL` as a synonym for `DOUBLE`.

- [STRICT\\_ALL\\_TABLES \[607\]](#)

Enable strict mode for all storage engines. Invalid data values are rejected. Additional detail follows.

- [STRICT\\_TRANS\\_TABLES \[607\]](#)

Enable strict mode for transactional storage engines, and when possible for nontransactional storage engines. Additional details follow.

Strict mode controls how MySQL handles input values that are invalid or missing. A value can be invalid for several reasons. For example, it might have the wrong data type for the column, or it might be out of range. A value is missing when a new row to be inserted does not contain a value for a non-`NULL` column that has no explicit `DEFAULT` clause in its definition. (For a `NULL` column, `NULL` is inserted if the value is missing.)

For transactional tables, an error occurs for invalid or missing values in a statement when either of the [STRICT\\_ALL\\_TABLES \[607\]](#) or [STRICT\\_TRANS\\_TABLES \[607\]](#) modes are enabled. The statement is aborted and rolled back.

For nontransactional tables, the behavior is the same for either mode, if the bad value occurs in the first row to be inserted or updated. The statement is aborted and the table remains unchanged. If the statement inserts or modifies multiple rows and the bad value occurs in the second or later row, the result depends on which strict option is enabled:

- For [STRICT\\_ALL\\_TABLES \[607\]](#), MySQL returns an error and ignores the rest of the rows. However, in this case, the earlier rows still have been inserted or updated. This means that you might get a partial update, which might not be what you want. To avoid this, it is best to use single-row statements because these can be aborted without changing the table.

- For `STRICT_TRANS_TABLES` [607], MySQL converts an invalid value to the closest valid value for the column and insert the adjusted value. If a value is missing, MySQL inserts the implicit default value for the column data type. In either case, MySQL generates a warning rather than an error and continues processing the statement. Implicit defaults are described in [Section 11.5, “Data Type Default Values”](#).

Strict mode disallows invalid date values such as `'2004-04-31'`. It does not disallow dates with zero month or day parts such as `'2004-04-00'` or “zero” dates. To disallow these as well, enable the `NO_ZERO_IN_DATE` [606] and `NO_ZERO_DATE` [606] SQL modes in addition to strict mode.

If you are not using strict mode (that is, neither `STRICT_TRANS_TABLES` [607] nor `STRICT_ALL_TABLES` [607] is enabled), MySQL inserts adjusted values for invalid or missing values and produces warnings. In strict mode, you can produce this behavior by using `INSERT IGNORE` or `UPDATE IGNORE`. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

Strict mode does not affect whether foreign key constraints are checked. `foreign_key_checks` [483] can be used for that. (See [Section 5.1.4, “Server System Variables”](#).)

## Combination Modes

The following special modes are provided as shorthand for combinations of mode values from the preceding list.

The descriptions include all mode values that are available in the most recent version of MySQL. For older versions, a combination mode does not include individual mode values that are not available except in newer versions.

- `ANSI` [608]

Equivalent to `REAL_AS_FLOAT` [607], `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604].

`ANSI` [608] mode also causes the server to return an error for queries where a set function `S` with an outer reference `S(outer_ref)` cannot be aggregated in the outer query against which the outer reference has been resolved. This is such a query:

```
SELECT * FROM t1 WHERE t1.a IN (SELECT MAX(t1.b) FROM t2 WHERE ...);
```

Here, `MAX(t1.b)` [1284] cannot be aggregated in the outer query because it appears in the `WHERE` clause of that query. Standard SQL requires an error in this situation. If `ANSI` [608] mode is not enabled, the server treats `S(outer_ref)` in such queries the same way that it would interpret `S(const)`.

See [Section 1.8.3, “Running MySQL in ANSI Mode”](#).

- `DB2` [608]

Equivalent to `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604], `NO_KEY_OPTIONS` [605], `NO_TABLE_OPTIONS` [605], `NO_FIELD_OPTIONS` [605].

- `MAXDB` [608]

Equivalent to `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604], `NO_KEY_OPTIONS` [605], `NO_TABLE_OPTIONS` [605], `NO_FIELD_OPTIONS` [605], `NO_AUTO_CREATE_USER` [604].

- `MSSQL` [608]

Equivalent to `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604], `NO_KEY_OPTIONS` [605], `NO_TABLE_OPTIONS` [605], `NO_FIELD_OPTIONS` [605].

- `MYSQL323` [608]



Equivalent to `NO_FIELD_OPTIONS` [605], `HIGH_NOT_PRECEDENCE` [603].

- `MYSQL40` [609]

Equivalent to `NO_FIELD_OPTIONS` [605], `HIGH_NOT_PRECEDENCE` [603].

- `ORACLE` [609]

Equivalent to `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604], `NO_KEY_OPTIONS` [605], `NO_TABLE_OPTIONS` [605], `NO_FIELD_OPTIONS` [605], `NO_AUTO_CREATE_USER` [604].

- `POSTGRESQL` [609]

Equivalent to `PIPES_AS_CONCAT` [607], `ANSI_QUOTES` [603], `IGNORE_SPACE` [604], `NO_KEY_OPTIONS` [605], `NO_TABLE_OPTIONS` [605], `NO_FIELD_OPTIONS` [605].

- `TRADITIONAL` [609]

Equivalent to `STRICT_TRANS_TABLES` [607], `STRICT_ALL_TABLES` [607], `NO_ZERO_IN_DATE` [606], `NO_ZERO_DATE` [606], `ERROR_FOR_DIVISION_BY_ZERO` [603], `NO_AUTO_CREATE_USER` [604], and `NO_ENGINE_SUBSTITUTION` [605].

## 5.1.8. Server Plugins

MySQL supports a plugin API that enables creation of server components. Plugins can be loaded at server startup, or loaded and unloaded at runtime without restarting the server. The components supported by this interface include, but are not limited to, storage engines, full-text parser plugins, partitioning support, and server extensions.

### 5.1.8.1. Installing and Uninstalling Plugins

Server plugins must be loaded in to the server before they can be used. MySQL enables you to load a plugin at server startup or at runtime. It is also possible to control the activation of loaded plugins at startup, and to unload them at runtime.

- [Installing plugins](#)
- [Controlling plugin activation](#)
- [Uninstalling plugins](#)

#### Installing Plugins

Server plugins must be known to the server before they can be used. A plugin can be made known several ways, as described here. In the following descriptions, *plugin\_name* stands for a plugin name such as `innodb` or `csv`.

##### Built-in plugins:

A plugin that is built in to the server is known by the server automatically. Normally, the server enables the plugin at startup, although this can be changed with the `--plugin_name` option.

##### Plugins registered in the `mysql.plugin` table:

The `mysql.plugin` table serves as a registry of plugins. The server normally enables each plugin listed in the table at startup, although whether a given plugin is enabled can be changed with the `--plugin_name` option. If the server is started with the `--skip-grant-tables` [439] option, it does not consult this table and does not load the plugins listed there.

##### Plugins named with command-line options:



A plugin that is located in a plugin library file can be loaded at server startup with the `--plugin-load [435]` option. Normally, the server enables the plugin at startup, although this can be changed with the `--plugin_name` option.

The option value is a semicolon-separated list of `name=plugin_library` pairs. Each `name` is the name of the plugin, and `plugin_library` is the name of the shared library that contains the plugin code. If a plugin library is named without any preceding plugin name, the server loads all plugins in the library. Each library file must be located in the directory named by the `plugin_dir [525]` system variable.

This option does not register any plugin in the `mysql.plugin` table. For subsequent restarts, the server loads the plugin again only if `--plugin-load [435]` is given again. That is, this option effects a one-time installation that persists only for one server invocation.

`--plugin-load [435]` enables plugins to be loaded even when `--skip-grant-tables [439]` is given (which causes the server to ignore the `mysql.plugin` table). `--plugin-load [435]` also enables plugins to be loaded at startup under configurations when plugins cannot be loaded at runtime.

The `--plugin-load-add [436]` option complements the `--plugin-load [435]` option. `--plugin-load-add [436]` adds a plugin or plugins to the set of plugins to be loaded at startup. The argument format is the same as for `--plugin-load [435]`. `--plugin-load-add [436]` can be used to avoid specifying a large set of plugins as a single long unwieldy `--plugin-load [435]` argument. `--plugin-load-add [436]` can be given in the absence of `--plugin-load [435]`, but any instance of `--plugin-load-add [436]` that appears before `--plugin-load [435]` has no effect because `--plugin-load [435]` resets the set of plugins to load. In other words, these options:

```
--plugin-load=x --plugin-load-add=y
```

are equivalent to this option:

```
--plugin-load="x;y"
```

But these options:

```
--plugin-load-add=y --plugin-load=x
```

are equivalent to this option:

```
--plugin-load=x
```

### Plugins installed with the `INSTALL PLUGIN` statement:

A plugin that is located in a plugin library file can be loaded at runtime with the `INSTALL PLUGIN` statement. The statement also registers the plugin in the `mysql.plugin` table to cause the server to load it on subsequent restarts. For this reason, `INSTALL PLUGIN` requires the `INSERT [755]` privilege for the `mysql.plugin` table.

If a plugin is named both using a `--plugin-load [435]` option and in the `mysql.plugin` table, the server starts but writes these messages to the error log:

```
100310 19:15:44 [ERROR] Function 'plugin_name' already exists
100310 19:15:44 [Warning] Couldn't load plugin named 'plugin_name'
with soname 'plugin_object_file'.
```

Example: The `--plugin-load [435]` option installs a plugin at server startup. To install a plugin named `myplugin` in a plugin library file named `somepluglib.so`, use these lines in a `my.cnf` file:

```
[mysqld]
plugin-load=myplugin=somepluglib.so
```

In this case, the plugin is not registered in `mysql.plugin`. Restarting the server without the `--plugin-load [435]` option causes the plugin not to be loaded at startup.

Alternatively, the `INSTALL PLUGIN` statement causes the server to load the plugin code from the library file at runtime:

```
mysql> INSTALL PLUGIN myplugin SONAME 'somepluglib.so';
```

`INSTALL PLUGIN` also causes “permanent” plugin registration: The server lists the plugin in the `mysql.plugin` table to ensure that it is loaded on subsequent server restarts.

Many plugins can be loaded either at server startup or at runtime. However, if a plugin is designed such that it must be loaded and initialized during server startup, use `--plugin-load` [435] rather than `INSTALL PLUGIN`.

While a plugin is loaded, information about it is available at runtime from several sources, such as the `INFORMATION_SCHEMA.PLUGINS` table and the `SHOW PLUGINS` statement. For more information, see [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).

## Controlling Plugin Activation

If the server knows about a plugin when it starts (for example, because the plugin is named using a `--plugin-load` [435] option or registered in the `mysql.plugin` table), the server loads and enables the plugin by default. It is possible to control activation for such a plugin using a `--plugin_name[=value]` startup option named after the plugin. In the following descriptions, `plugin_name` stands for a plugin name such as `innodb` or `csv`. As with other options, dashes and underscores are interchangeable in option names. For example, `--my_plugin=ON` and `--my-plugin=ON` are equivalent.

- `--plugin_name=OFF`

Tells the server to disable the plugin.

- `--plugin_name[=ON]`

Tells the server to enable the plugin. (Specifying the option as `--plugin_name` without a value has the same effect.) If the plugin fails to initialize, the server runs with the plugin disabled.

- `--plugin_name=FORCE`

Tells the server to enable the plugin, but if plugin initialization fails, the server does not start. In other words, this option forces the server to run with the plugin enabled or not at all.

- `--plugin_name=FORCE_PLUS_PERMANENT`

Like `FORCE`, but in addition prevents the plugin from being unloaded at runtime. If a user attempts to do so with `UNINSTALL PLUGIN`, an error occurs.

The values `OFF`, `ON`, `FORCE`, and `FORCE_PLUS_PERMANENT` are not case sensitive.

The activation state for plugins is visible in the `LOAD_OPTION` column of the `INFORMATION_SCHEMA.PLUGINS` table.

Suppose that `CSV`, `BLACKHOLE`, and `ARCHIVE` are built-in pluggable storage engines and that you want the server to load them at startup, subject to these conditions: The server is permitted to run if `CSV` initialization fails, but must require that `BLACKHOLE` initialization succeeds, and `ARCHIVE` should be disabled. To accomplish that, use these lines in an option file:

```
[mysqld]
csv=ON
blackhole=FORCE
archive=OFF
```

The `--enable-plugin_name` option format is supported as a synonym for `--plugin_name=ON`. The `--disable-plugin_name` and `--skip-plugin_name` option formats are supported as synonyms for `--plugin_name=OFF`.

If a plugin is disabled, either explicitly with `OFF` or implicitly because it was enabled with `ON` but failed to initialize, aspects of server operation that require the plugin will change. For example, if the plugin

implements a storage engine, existing tables for the storage engine become inaccessible, and attempts to create new tables for the storage engine result in tables that use the default storage engine unless the `NO_ENGINE_SUBSTITUTION` [605] SQL mode has been enabled to cause an error to occur instead.

Disabling a plugin may require adjustment to other options. For example, if you start the server using `--skip-innodb` [1698] to disable `InnoDB`, other `innodb_XXX` options likely will need to be omitted from the startup command. In addition, because `InnoDB` is the default storage engine, it will not start unless you specify another available storage engine with `--default_storage_engine` [474]. As of MySQL 5.6.3, you will also need to set `--default_tmp_storage_engine` [474].

## Uninstalling Plugins

A plugin known to the server can be uninstalled to disable it at runtime with the `UNINSTALL PLUGIN` statement. The statement unloads the plugin and removes it from the `mysql.plugin` table if it is registered there. For this reason, `UNINSTALL PLUGIN` statement requires the `DELETE` [754] privilege for the `mysql.plugin` table. With the plugin no longer registered in the table, the server will not load the plugin automatically for subsequent restarts.

`UNINSTALL PLUGIN` can unload plugins regardless of whether they were loaded with `INSTALL PLUGIN` or `--plugin-load` [435].

`UNINSTALL PLUGIN` is subject to these exceptions:

- It cannot unload plugins that are built in to the server. These can be identified as those that have a library name of `NULL` in the output from `INFORMATION_SCHEMA.PLUGINS` or `SHOW PLUGINS`.
- It cannot unload plugins for which the server was started with `--plugin_name=FORCE_PLUS_PERMANENT`, which prevents plugin unloading at runtime. These can be identified from the `LOAD_OPTION` column of the `INFORMATION_SCHEMA.PLUGINS` table.

### 5.1.8.2. Obtaining Server Plugin Information

There are several ways to determine which plugins are installed in the server:

- The `INFORMATION_SCHEMA.PLUGINS` table contains a row for each loaded plugin. Any that have a `PLUGIN_LIBRARY` value of `NULL` are built in and cannot be unloaded.

```
mysql> SELECT * FROM information_schema.PLUGINS\G
***** 1. row *****
      PLUGIN_NAME: binlog
      PLUGIN_VERSION: 1.0
      PLUGIN_STATUS: ACTIVE
      PLUGIN_TYPE: STORAGE ENGINE
      PLUGIN_TYPE_VERSION: 50158.0
      PLUGIN_LIBRARY: NULL
      PLUGIN_LIBRARY_VERSION: NULL
      PLUGIN_AUTHOR: MySQL AB
      PLUGIN_DESCRIPTION: This is a pseudo storage engine to represent the binlog in a transaction
      PLUGIN_LICENSE: GPL
      LOAD_OPTION: FORCE
...
***** 10. row *****
      PLUGIN_NAME: InnoDB
      PLUGIN_VERSION: 1.0
      PLUGIN_STATUS: ACTIVE
      PLUGIN_TYPE: STORAGE ENGINE
      PLUGIN_TYPE_VERSION: 50158.0
      PLUGIN_LIBRARY: ha_innodb_plugin.so
      PLUGIN_LIBRARY_VERSION: 1.0
      PLUGIN_AUTHOR: Innobase Oy
      PLUGIN_DESCRIPTION: Supports transactions, row-level locking,
                        and foreign keys
      PLUGIN_LICENSE: GPL
      LOAD_OPTION: ON
...
```

- The `SHOW PLUGINS` statement displays a row for each loaded plugin. Any that have a `Library` value of `NULL` are built in and cannot be unloaded.

```
mysql> SHOW PLUGINS\G
***** 1. row *****
Name: binlog
Status: ACTIVE
Type: STORAGE ENGINE
Library: NULL
License: GPL
...
***** 10. row *****
Name: InnoDB
Status: ACTIVE
Type: STORAGE ENGINE
Library: ha_innodb_plugin.so
License: GPL
...
```

- The `mysql.plugin` table shows which plugins have been registered with `INSTALL PLUGIN`. The table contains only plugin names and library file names, so it does not provide as much information as the `PLUGINS` table or the `SHOW PLUGINS` statement.

## 5.1.9. IPv6 Support

Support for IPv6 in MySQL includes these capabilities:

- MySQL Server can accept TCP/IP connections from clients connecting over IPv6. For example, this command connects over IPv6 to the MySQL server on the local host:

```
shell> mysql -h ::1
```

To use this capability, two things must be true:

- Your system must be configured to support IPv6. See [Section 5.1.9.1, “Verifying System Support for IPv6”](#).
- The default MySQL server configuration permits only IPv4 connections, so the server must be configured for IPv6 connections. To permit IPv6 connections in addition to or instead of IPv4 connections, start the server with an appropriate `--bind-address [414]` option. See [Section 5.1.4, “Server System Variables”](#).
- MySQL account names permit IPv6 addresses to enable DBAs to specify privileges for clients that connect to the server over IPv6. See [Section 6.2.3, “Specifying Account Names”](#). IPv6 addresses can be specified in account names in statements such as `CREATE USER`, `GRANT`, and `REVOKE`. For example:

```
mysql> CREATE USER 'bill'@':::1' IDENTIFIED BY 'secret';
mysql> GRANT SELECT ON mydb.* TO 'bill'@':::1';
```

- IPv6 functions enable conversion between string and internal format IPv6 address formats, and checking whether values represent valid IPv6 addresses. For example, `INET6_ATON()` [1276] and `INET6_NTOA()` [1277] are similar to `INET_ATON()` [1276] and `INET_NTOA()` [1276], but handle IPv6 addresses in addition to IPv4 addresses. See [Section 12.16, “Miscellaneous Functions”](#).

The following sections describe how to set up MySQL so that clients can connect to the server over IPv6.

### 5.1.9.1. Verifying System Support for IPv6

Before MySQL Server can accept IPv6 connections, the operating system on your server host must support IPv6. As a simple test to determine whether that is true, try this command:

```
shell> ping6 ::1
16 bytes from ::1, icmp_seq=0 hlim=64 time=0.171 ms
```

```
16 bytes from ::1, icmp_seq=1 hlim=64 time=0.077 ms
...
```

To produce a description of your system's network interfaces, invoke `ifconfig -a` and look for IPv6 addresses in the output.

If your host does not support IPv6, consult your system documentation for instructions on enabling it. It might be that you need only reconfigure an existing network interface to add an IPv6 address. Or a more extensive change might be needed, such as rebuilding the kernel with IPv6 options enabled.

These links may be helpful in setting up IPv6 on various platforms:

- [Windows XP](#)
- [Gentoo Linux](#)
- [Ubuntu Linux](#)
- [Linux \(Generic\)](#)
- [Mac OS X](#)

### 5.1.9.2. Configuring the MySQL Server to Permit IPv6 Connections

The MySQL server listens on a single network socket for TCP/IP connections. This socket is bound to a single address, but it is possible for an address to map onto multiple network interfaces. To specify an address, use the `--bind-address=addr` [414] option at server startup, where `addr` is an IPv4 or IPv6 address or a host name. (IPv6 addresses are not supported before MySQL 5.5.3.) If `addr` is a host name, the server resolves the name to an IP address and binds to that address.

The server treats different types of addresses as follows:

- If the address is `*`, the server accepts TCP/IP connections on all server host IPv6 and IPv4 interfaces if the server host supports IPv6, or accepts TCP/IP connections on all IPv4 addresses otherwise. Use this address to permit both IPv4 and IPv6 connections on all server interfaces. This value is permitted (and is the default) as of MySQL 5.6.6.
- If the address is `0.0.0.0`, the server accepts TCP/IP connections on all server host IPv4 interfaces. This is the default before MySQL 5.6.6.
- If the address is `::`, the server accepts TCP/IP connections on all server host IPv4 and IPv6 interfaces. Use this address to permit both IPv4 and IPv6 connections on all server interfaces.
- If the address is an IPv4-mapped address, the server accepts TCP/IP connections for that address, in either IPv4 or IPv6 format. For example, if the server is bound to `::ffff:127.0.0.1`, clients can connect using `--host=127.0.0.1` or `--host>::ffff:127.0.0.1`.
- If the address is a “regular” IPv4 or IPv6 address (such as `127.0.0.1` or `:::1`), the server accepts TCP/IP connections only for that IPv4 or IPv6 address.

If you intend to bind the server to a specific address, be sure that the `mysql.user` grant table contains an account with administrative privileges that you can use to connect to that address. Otherwise, you will not be able to shut down the server. For example, if you bind the server to `:::1`, it accepts connections only on that address. In that case, first make sure that the `'root'@':::1'` account is present in the `mysql.user` table so you can still connect to the server to shut it down.

### 5.1.9.3. Connecting Using the IPv6 Local Host Address

The following procedure shows how to configure MySQL to permit IPv6 connections by clients that connect to the local server using the `:::1` local host address. The instructions given here assume that your system supports IPv6.

1. Start the MySQL server with an appropriate `--bind-address` [414] option to permit it to accept IPv6 connections. For example, put the following lines in your server option file and restart the server:

```
[mysqld]
bind-address = * # before 5.6.6, use :: rather than *
```

Alternatively, you can bind the server to `::1`, but that makes the server more restrictive for TCP/IP connections. It accepts only IPv6 connections for that single address and rejects IPv4 connections. For more information, see [Section 5.1.9.2, “Configuring the MySQL Server to Permit IPv6 Connections”](#).

2. As an administrator, connect to the server and create an account for a local user who will connect from the `::1` local IPv6 host address:

```
mysql> CREATE USER 'ipv6user'@'::1' IDENTIFIED BY 'ipv6pass';
```

For the permitted syntax of IPv6 addresses in account names, see [Section 6.2.3, “Specifying Account Names”](#). In addition to the `CREATE USER` statement, you can issue `GRANT` statements that give specific privileges to the account, although that is not necessary for the remaining steps in this procedure.

3. Invoke the `mysql` client to connect to the server using the new account:

```
shell> mysql -h ::1 -u ipv6user -pipv6pass
```

4. Try some simple statements that show connection information:

```
mysql> STATUS
...
Connection:  ::1 via TCP/IP
...

mysql> SELECT CURRENT_USER(), @@bind_address;
+-----+-----+
| CURRENT_USER() | @@bind_address |
+-----+-----+
| ipv6user@::1   | ::             |
+-----+-----+
```

#### 5.1.9.4. Connecting Using IPv6 Nonlocal Host Addresses

The following procedure shows how to configure MySQL to permit IPv6 connections by remote clients. It is similar to the preceding procedure for local clients, but the server and client hosts are distinct and each has its own nonlocal IPv6 address. The example uses these addresses:

```
Server host: 2001:db8:0:f101::1
Client host: 2001:db8:0:f101::2
```

These addresses are chosen from the nonroutable address range recommended by [IANA](#) for documentation purposes and suffice for testing on your local network. To accept IPv6 connections from clients outside the local network, the server host must have a public address. If your network provider assigns you an IPv6 address, you can use that. Otherwise, another way to obtain an address is to use an IPv6 broker; see [Section 5.1.9.5, “Obtaining an IPv6 Address from a Broker”](#).

1. Start the MySQL server with an appropriate `--bind-address` [414] option to permit it to accept IPv6 connections. For example, put the following lines in your server option file and restart the server:

```
[mysqld]
bind-address = * # before 5.6.6, use :: rather than *
```

Alternatively, you can bind the server to `2001:db8:0:f101::1`, but that makes the server more restrictive for TCP/IP connections. It accepts only IPv6 connections for that single address and

rejects IPv4 connections. For more information, see [Section 5.1.9.2, “Configuring the MySQL Server to Permit IPv6 Connections”](#).

2. On the server host (`2001:db8:0:f101::1`), create an account for a user who will connect from the client host (`2001:db8:0:f101::2`):

```
mysql> CREATE USER 'remoteipv6user'@'2001:db8:0:f101::2' IDENTIFIED BY 'remoteipv6pass';
```

3. On the client host (`2001:db8:0:f101::2`), invoke the `mysql` client to connect to the server using the new account:

```
shell> mysql -h 2001:db8:0:f101::1 -u remoteipv6user -premoteipv6pass
```

4. Try some simple statements that show connection information:

```
mysql> STATUS
...
Connection: 2001:db8:0:f101::1 via TCP/IP
...

mysql> SELECT CURRENT_USER(), @@bind_address;
+-----+-----+
| CURRENT_USER() | @@bind_address |
+-----+-----+
| remoteipv6user@2001:db8:0:f101::2 | :: |
+-----+-----+
```

### 5.1.9.5. Obtaining an IPv6 Address from a Broker

If you do not have a public IPv6 address that enables your system to communicate over IPv6 outside your local network, you can obtain one from an IPv6 broker. The [Wikipedia IPv6 Tunnel Broker page](#) lists several brokers and their features, such as whether they provide static addresses and the supported routing protocols.

After configuring your server host to use a broker-supplied IPv6 address, start the MySQL server with an appropriate `--bind-address [414]` option to permit the server to accept IPv6 connections. For example, put the following lines in the server option file and restart the server:

```
[mysqld]
bind-address = * # before 5.6.6, use :: rather than *
```

Alternatively, you can bind the server to the specific IPv6 address provided by the broker, but that makes the server more restrictive for TCP/IP connections. It accepts only IPv6 connections for that single address and rejects IPv4 connections. For more information, see [Section 5.1.9.2, “Configuring the MySQL Server to Permit IPv6 Connections”](#). In addition, if the broker allocates dynamic addresses, the address provided for your system might change the next time you connect to the broker. If so, any accounts you create that name the original address become invalid. To bind to a specific address but avoid this change-of-address problem, you may be able to arrange with the broker for a static IPv6 address.

The following example shows how to use Freenet6 as the broker and the `gogoc` IPv6 client package on Gentoo Linux.

1. Create a account at Freenet6 by visiting this URL and signing up:

```
http://gogonet.gogo6.com
```

2. After creating the account, go to this URL, sign in, and create a user ID and password for the IPv6 broker:

```
http://gogonet.gogo6.com/page/freenet6-registration
```

3. As `root`, install `gogoc`:

```
shell> emerge gogoc
```



4. Edit `/etc/gogoc/gogoc.conf` to set the `userid` and `password` values. For example:

```
userid=gogouser
passwd=gogopass
```

5. Start `gogoc`:

```
shell> /etc/init.d/gogoc start
```

To start `gogoc` each time your system boots, execute this command:

```
shell> rc-update add gogoc default
```

6. Use `ping6` to try to ping a host:

```
shell> ping6 ipv6.google.com
```

7. To see your IPv6 address:

```
shell> ifconfig tun
```

## 5.1.10. Server-Side Help

MySQL Server supports a `HELP` statement that returns online information from the MySQL Reference manual (see [Section 13.8.3, “HELP Syntax”](#)). The proper operation of this statement requires that the help tables in the `mysql` database be initialized with help topic information, which is done by processing the contents of the `fill_help_tables.sql` script.

If you install MySQL using a binary or source distribution on Unix, help table setup occurs when you run `mysql_install_db`. For an RPM distribution on Linux or binary distribution on Windows, help table setup occurs as part of the MySQL installation process.

If you upgrade MySQL using a binary distribution, the help tables are not upgraded automatically, but you can upgrade them manually. Locate the `fill_help_tables.sql` file in the `share` or `share/mysql` directory. Change location into that directory and process the file with the `mysql` client as follows:

```
shell> mysql -u root mysql < fill_help_tables.sql
```

You can also obtain the latest `fill_help_tables.sql` at any time to upgrade your help tables. Download the proper file for your version of MySQL from <http://dev.mysql.com/doc/index-other.html>. After downloading and uncompressing the file, process it with `mysql` as described previously.

If you are working with Bazaar and a MySQL development source tree, you will need to download the `fill_help_tables.sql` file because the tree contains only a “stub” version.

## 5.1.11. Server Response to Signals

On Unix, signals can be sent to processes. `mysqld` responds to signals sent to it as follows:

- `SIGTERM` causes the server to shut down.
- `SIGHUP` causes the server to reload the grant tables and flush the logs (like `FLUSH PRIVILEGES` and `FLUSH LOGS`). It also writes a status report to the error log that has this format:

```
Status information:

Current dir: /var/mysql/data/
Running threads: 0  Stack size: 196608
Current locks:

Key caches:
default
Buffer_size:      8388600
Block_size:      1024
```



```

Division_limit:      100
Age_limit:          300
blocks used:        0
not flushed:        0
w_requests:         0
writes:             0
r_requests:         0
reads:              0

handler status:
read_key:           0
read_next:          0
read_rnd            0
read_first:         1
write:              0
delete              0
update:             0

Table status:
Opened tables:      5
Open tables:        0
Open files:         7
Open streams:       0

Alarm status:
Active alarms:      1
Max used alarms:    2
Next alarm time:    67

```

On some Mac OS X 10.3 versions, `mysqld` ignores `SIGHUP` and `SIGQUIT`.

## 5.1.12. The Shutdown Process

The server shutdown process takes place as follows:

1. The shutdown process is initiated.

This can occur initiated several ways. For example, a user with the `SHUTDOWN [756]` privilege can execute a `mysqladmin shutdown` command. `mysqladmin` can be used on any platform supported by MySQL. Other operating system-specific shutdown initiation methods are possible as well: The server shuts down on Unix when it receives a `SIGTERM` signal. A server running as a service on Windows shuts down when the services manager tells it to.

2. The server creates a shutdown thread if necessary.

Depending on how shutdown was initiated, the server might create a thread to handle the shutdown process. If shutdown was requested by a client, a shutdown thread is created. If shutdown is the result of receiving a `SIGTERM` signal, the signal thread might handle shutdown itself, or it might create a separate thread to do so. If the server tries to create a shutdown thread and cannot (for example, if memory is exhausted), it issues a diagnostic message that appears in the error log:

```
Error: Can't create thread to kill server
```

3. The server stops accepting new connections.

To prevent new activity from being initiated during shutdown, the server stops accepting new client connections by closing the handlers for the network interfaces to which it normally listens for connections: the TCP/IP port, the Unix socket file, the Windows named pipe, and shared memory on Windows.

4. The server terminates current activity.

For each thread associated with a client connection, the server breaks the connection to the client and marks the thread as killed. Threads die when they notice that they are so marked. Threads for idle connections die quickly. Threads that currently are processing statements check their state periodically and take longer to die. For additional information about thread termination, see

[Section 13.7.6.4, “KILL Syntax”](#), in particular for the instructions about killed `REPAIR TABLE` or `OPTIMIZE TABLE` operations on `MyISAM` tables.

For threads that have an open transaction, the transaction is rolled back. Note that if a thread is updating a nontransactional table, an operation such as a multiple-row `UPDATE` or `INSERT` may leave the table partially updated because the operation can terminate before completion.

If the server is a master replication server, it treats threads associated with currently connected slaves like other client threads. That is, each one is marked as killed and exits when it next checks its state.

If the server is a slave replication server, it stops the I/O and SQL threads, if they are active, before marking client threads as killed. The SQL thread is permitted to finish its current statement (to avoid causing replication problems), and then stops. If the SQL thread is in the middle of a transaction at this point, the server waits until the current replication event group (if any) has finished executing, or until the user issues a `KILL QUERY` or `KILL CONNECTION` statement. See also [Section 13.4.2.6, “STOP SLAVE Syntax”](#). Since nontransactional statements cannot be rolled back, in order to guarantee crash-safe replication, only transactional tables should be used.



#### Note

In order to guarantee crash safety on the slave, you must also run the slave with `--relay-log-recovery [1975]` enabled.

See also [Section 16.2.2, “Replication Relay and Status Logs”](#)).

5. The server shuts down or closes storage engines.

At this stage, the server flushes the table cache and closes all open tables.

Each storage engine performs any actions necessary for tables that it manages. `InnoDB` flushes its buffer pool to disk (unless `innodb_fast_shutdown [1716]` is 2), writes the current LSN to the tablespace, and terminates its own internal threads. `MyISAM` flushes any pending index writes for a table.

6. The server exits.

## 5.2. MySQL Server Logs

MySQL Server has several logs that can help you find out what activity is taking place.

Log Type	Information Written to Log
Error log	Problems encountered starting, running, or stopping <code>mysqld</code>
General query log	Established client connections and statements received from clients
Binary log	Statements that change data (also used for replication)
Relay log	Data changes received from a replication master server
Slow query log	Queries that took more than <code>long_query_time [500]</code> seconds to execute

By default, no logs are enabled (except the error log on Windows). The following log-specific sections provide information about the server options that enable logging.

By default, the server writes files for all enabled logs in the data directory. You can force the server to close and reopen the log files (or in some cases switch to a new log file) by flushing the logs. Log flushing occurs when you issue a `FLUSH LOGS` statement; execute `mysqladmin` with a `flush-logs` or `refresh` argument; or execute `mysqldump` with a `--flush-logs [294]` or `--master-data [288]` option. See [Section 13.7.6.3, “FLUSH Syntax”](#), [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#), and [Section 4.5.4, “mysqldump — A Database Backup Program”](#). In addition,

the binary log is flushed when its size reaches the value of the `max_binlog_size` [2016] system variable.

You can control the general query and slow query logs during runtime. You can enable or disable logging, or change the log file name. You can tell the server to write general query and slow query entries to log tables, log files, or both. For details, see [Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”](#), [Section 5.2.3, “The General Query Log”](#), and [Section 5.2.5, “The Slow Query Log”](#).

The relay log is used only on slave replication servers, to hold data changes from the master server that must also be made on the slave. For discussion of relay log contents and configuration, see [Section 16.2.2.1, “The Slave Relay Log”](#).

For information about log maintenance operations such as expiration of old log files, see [Section 5.2.6, “Server Log Maintenance”](#).

For information about keeping logs secure, see [Section 6.1.2.3, “Passwords and Logging”](#).

## 5.2.1. Selecting General Query and Slow Query Log Output Destinations

MySQL Server provides flexible control over the destination of output to the general query log and the slow query log, if those logs are enabled. Possible destinations for log entries are log files or the `general_log` and `slow_log` tables in the `mysql` database. Either or both destinations can be selected.

Currently, logging to tables incurs significantly more server overhead than logging to files. If you enable the general log or slow query log and require highest performance, you should use file logging, not table logging.

**Log control at server startup.** The `--log-output` [427] option specifies the destination for log output. This option does not in itself enable the logs. Its syntax is `--log-output[=value,...] [427]`:

- If `--log-output` [427] is given with a value, the value should be a comma-separated list of one or more of the words `TABLE` (log to tables), `FILE` (log to files), or `NONE` (do not log to tables or files). `NONE`, if present, takes precedence over any other specifiers.
- If `--log-output` [427] is omitted, the default logging destination is `FILE`.

The `general_log` [485] system variable controls logging to the general query log for the selected log destinations. If specified at server startup, `general_log` [485] takes an optional argument of 1 or 0 to enable or disable the log. To specify a file name other than the default for file logging, set the `general_log_file` [486] variable. Similarly, the `slow_query_log` [541] variable controls logging to the slow query log for the selected destinations and setting `slow_query_log_file` [541] specifies a file name for file logging. If either log is enabled, the server opens the corresponding log file and writes startup messages to it. However, further logging of queries to the file does not occur unless the `FILE` log destination is selected.

Examples:

- To write general query log entries to the log table and the log file, use `--log-output=TABLE,FILE [427]` to select both log destinations and `--general_log [485]` to enable the general query log.
- To write general and slow query log entries only to the log tables, use `--log-output=TABLE [427]` to select tables as the log destination and `--general_log [485]` and `--slow_query_log [541]` to enable both logs.
- To write slow query log entries only to the log file, use `--log-output=FILE [427]` to select files as the log destination and `--slow_query_log [541]` to enable the slow query log. (In this case, because the default log destination is `FILE`, you could omit the `--log-output [427]` option.)

**Log control at runtime.** The system variables associated with log tables and files enable runtime control over logging:

- The global `log_output` [498] system variable indicates the current logging destination. It can be modified at runtime to change the destination.
- The global `general_log` [485] and `slow_query_log` [541] variables indicate whether the general query log and slow query log are enabled (`ON`) or disabled (`OFF`). You can set these variables at runtime to control whether the logs are enabled.
- The global `general_log_file` [486] and `slow_query_log_file` [541] variables indicate the names of the general query log and slow query log files. You can set these variables at server startup or at runtime to change the names of the log files.
- To disable or enable general query logging for the current connection, set the session `sql_log_off` [544] variable to `ON` or `OFF`.

The use of tables for log output offers the following benefits:

- Log entries have a standard format. To display the current structure of the log tables, use these statements:

```
SHOW CREATE TABLE mysql.general_log;
SHOW CREATE TABLE mysql.slow_log;
```

- Log contents are accessible through SQL statements. This enables the use of queries that select only those log entries that satisfy specific criteria. For example, to select log contents associated with a particular client (which can be useful for identifying problematic queries from that client), it is easier to do this using a log table than a log file.
- Logs are accessible remotely through any client that can connect to the server and issue queries (if the client has the appropriate log table privileges). It is not necessary to log in to the server host and directly access the file system.

The log table implementation has the following characteristics:

- In general, the primary purpose of log tables is to provide an interface for users to observe the runtime execution of the server, not to interfere with its runtime execution.
- `CREATE TABLE`, `ALTER TABLE`, and `DROP TABLE` are valid operations on a log table. For `ALTER TABLE` and `DROP TABLE`, the log table cannot be in use and must be disabled, as described later.
- By default, the log tables use the `CSV` storage engine that writes data in comma-separated values format. For users who have access to the `.CSV` files that contain log table data, the files are easy to import into other programs such as spreadsheets that can process CSV input.

The log tables can be altered to use the `MyISAM` storage engine. You cannot use `ALTER TABLE` to alter a log table that is in use. The log must be disabled first. No engines other than `CSV` or `MyISAM` are legal for the log tables.

- To disable logging so that you can alter (or drop) a log table, you can use the following strategy. The example uses the general query log; the procedure for the slow query log is similar but uses the `slow_log` table and `slow_query_log` [541] system variable.

```
SET @old_log_state = @@global.general_log;
SET GLOBAL general_log = 'OFF';
ALTER TABLE mysql.general_log ENGINE = MyISAM;
SET GLOBAL general_log = @old_log_state;
```

- `TRUNCATE TABLE` is a valid operation on a log table. It can be used to expire log entries.
- `RENAME TABLE` is a valid operation on a log table. You can atomically rename a log table (to perform log rotation, for example) using the following strategy:

```
USE mysql;
DROP TABLE IF EXISTS general_log2;
CREATE TABLE general_log2 LIKE general_log;
RENAME TABLE general_log TO general_log_backup, general_log2 TO general_log;
```

- `CHECK TABLE` is a valid operation on a log table.
- `LOCK TABLES` cannot be used on a log table.
- `INSERT`, `DELETE`, and `UPDATE` cannot be used on a log table. These operations are permitted only internally to the server itself.
- `FLUSH TABLES WITH READ LOCK` and the state of the global `read_only` [532] system variable have no effect on log tables. The server can always write to the log tables.
- Entries written to the log tables are not written to the binary log and thus are not replicated to slave servers. (Prior to MySQL 5.6.9, this was not always enforced correctly. See Bug #14741537.)
- To flush the log tables or log files, use `FLUSH TABLES` or `FLUSH LOGS`, respectively.
- Partitioning of log tables is not permitted.
- Before MySQL 5.6.6, `mysqldump` does not dump the `general_log` or `slow_query_log` tables for dumps of the `mysql` database. As of 5.6.6, the dump includes statements to recreate those tables so that they are not missing after reloading the dump file. Log table contents are not dumped.

## 5.2.2. The Error Log

The error log contains information indicating when `mysqld` was started and stopped and also any critical errors that occur while the server is running. If `mysqld` notices a table that needs to be automatically checked or repaired, it writes a message to the error log.

On some operating systems, the error log contains a stack trace if `mysqld` dies. The trace can be used to determine where `mysqld` died. See [MySQL Internals: Porting to Other Systems](#).

You can specify where `mysqld` writes the error log with the `--log-error[=file_name]` [427] option. If the option is given with no `file_name` value, `mysqld` uses the name `host_name.err` by default. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.

If you do not specify `--log-error` [427], or (on Windows) if you use the `--console` [417] option, errors are written to `stderr`, the standard error output. Usually this is your terminal.

On Windows, error output is always written to the error log if `--console` [417] is not given.

In addition, on Windows, events and error messages are written to the Windows Event Log within the Application log. Entries marked as `Warning` and `Note` are written to the Event Log, but informational messages (such as information statements from individual storage engines) are not copied to the Event Log. The log entries have a source of **MySQL**. You cannot disable writing information to the Windows Event Log.

If you flush the logs using `FLUSH LOGS` or `mysqladmin flush-logs` and `mysqld` is writing the error log to a file (for example, if it was started with the `--log-error` [427] option), the server closes and reopens the log file. To rename the file, you can do so manually before flushing. Then flushing the logs reopens a new file with the original file name. For example, you can rename the file and create a new one using the following commands:

```
shell> mv host_name.err host_name.err-old
shell> mysqladmin flush-logs
shell> mv host_name.err-old backup-directory
```

On Windows, use `rename` rather than `mv`.

No error log renaming occurs when the logs are flushed if the server is not writing to a named file.

If you use `mysqld_safe` to start `mysqld`, `mysqld_safe` arranges for `mysqld` to write error messages to a log file or to `syslog`. `mysqld_safe` has three error-logging options, `--syslog` [229], `--skip-syslog` [229], and `--log-error` [227]. The default with no logging options or with `--skip-syslog` [229] is to use the default log file. To explicitly specify use of an error log file, specify `--log-error=file_name` [227] to `mysqld_safe`, and `mysqld_safe` will arrange for `mysqld` to write messages to a log file. To use `syslog` instead, specify the `--syslog` [229] option.

If you specify `--log-error` [427] in an option file in a `[mysqld]`, `[server]`, or `[mysqld_safe]` section, `mysqld_safe` will find and use the option.

If `mysqld_safe` is used to start `mysqld` and `mysqld` dies unexpectedly, `mysqld_safe` notices that it needs to restart `mysqld` and writes a `restarted mysqld` message to the error log.

The `--log-warnings` [430] option or `log_warnings` [499] system variable can be used to control warning logging to the error log. The default value is enabled (1). Warning logging can be disabled using a value of 0. If the value is greater than 1, aborted connections are written to the error log, and access-denied errors for new connection attempts are written. See Section C.5.2.11, “Communication Errors and Aborted Connections”.

### 5.2.3. The General Query Log

The general query log is a general record of what `mysqld` is doing. The server writes information to this log when clients connect or disconnect, and it logs each SQL statement received from clients. The general query log can be very useful when you suspect an error in a client and want to know exactly what the client sent to `mysqld`.

`mysqld` writes statements to the query log in the order that it receives them, which might differ from the order in which they are executed. This logging order contrasts to the binary log, for which statements are written after they are executed but before any locks are released. (Also, the query log contains all statements, whereas the binary log does not contain statements that only select data.)

By default, the general query log is disabled. To specify the initial general query log state explicitly, use `--general_log[={0|1}]` [485]. With no argument or an argument of 1, `--general_log` [485] enables the log. With an argument of 0, this option disables the log. To specify a log file name, use `--general_log_file=file_name` [486]. To specify the log destination, use `--log-output` [427] (as described in Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”).

If you specify no name for the general query log file, the default name is `host_name.log`. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.

To disable or enable the general query log or change the log file name at runtime, use the global `general_log` [485] and `general_log_file` [486] system variables. Set `general_log` [485] to 0 (or `OFF`) to disable the log or to 1 (or `ON`) to enable it. Set `general_log_file` [486] to specify the name of the log file. If a log file already is open, it is closed and the new file is opened.

When the general query log is enabled, the server writes output to any destinations specified by the `--log-output` [427] option or `log_output` [498] system variable. If you enable the log, the server opens the log file and writes startup messages to it. However, further logging of queries to the file does not occur unless the `FILE` log destination is selected. If the destination is `NONE`, the server writes no queries even if the general log is enabled. Setting the log file name has no effect on logging if the log destination value does not contain `FILE`.

Server restarts and log flushing do not cause a new general query log file to be generated (although flushing closes and reopens it). You can rename the file and create a new one by using the following commands:

```
shell> mv host_name.log host_name-old.log
shell> mysqladmin flush-logs
```



```
shell> mv host_name-old.log backup-directory
```

On Windows, use `rename` rather than `mv`.

You can also rename the general query log file at runtime by disabling the log:

```
SET GLOBAL general_log = 'OFF';
```

With the log disabled, rename the log file externally; for example, from the command line. Then enable the log again:

```
SET GLOBAL general_log = 'ON';
```

This method works on any platform and does not require a server restart.

The session `sql_log_off` [544] variable can be set to `ON` or `OFF` to disable or enable general query logging for the current connection.

As of MySQL 5.6.3, passwords in statements written to the general query log are rewritten by the server not to occur literally in plain text. Password rewriting can be suppressed for the general query log by starting the server with the `--log-raw` [428] option. This option may be useful for diagnostic purposes, to see the exact text of statements as received by the server, but for security reasons is not recommended for production use.

Before MySQL 5.6.3, passwords in statements are not rewritten and the general query log should be protected. See [Section 6.1.2.3, “Passwords and Logging”](#).

## 5.2.4. The Binary Log

The binary log contains “events” that describe database changes such as table creation operations or changes to table data. It also contains events for statements that potentially could have made changes (for example, a `DELETE` which matched no rows), unless row-based logging is used. The binary log also contains information about how long each statement took that updated data. The binary log has two important purposes:

- For replication, the binary log on a master replication server provides a record of the data changes to be sent to slave servers. The master server sends the events contained in its binary log to its slaves, which execute those events to make the same data changes that were made on the master. See [Section 16.2, “Replication Implementation”](#).
- Certain data recovery operations require use of the binary log. After a backup has been restored, the events in the binary log that were recorded after the backup was made are re-executed. These events bring databases up to date from the point of the backup. See [Section 7.5, “Point-in-Time \(Incremental\) Recovery Using the Binary Log”](#).

The binary log is not used for statements such as `SELECT` or `SHOW` that do not modify data. To log all statements (for example, to identify a problem query), use the general query log. See [Section 5.2.3, “The General Query Log”](#).

Running a server with binary logging enabled makes performance slightly slower. However, the benefits of the binary log in enabling you to set up replication and for restore operations generally outweigh this minor performance decrement.

Beginning with MySQL 5.6.2, the binary log is crash-safe. Only complete events or transactions are logged or read back.

As of MySQL 5.6.3, passwords in statements written to the binary log are rewritten by the server not to occur literally in plain text. Before MySQL 5.6.3, passwords in statements are not rewritten and the binary log should be protected. See [Section 6.1.2.3, “Passwords and Logging”](#).

The following discussion describes some of the server options and variables that affect the operation of binary logging. For a complete list, see [Section 16.1.4.4, “Binary Log Options and Variables”](#).

To enable the binary log, start the server with the `--log-bin[=base_name]` [2002] option. If no *base\_name* value is given, the default name is the value of the `pid-file` option (which by default is the name of host machine) followed by `-bin`. If the basename is given, the server writes the file in the data directory unless the basename is given with a leading absolute path name to specify a different directory. It is recommended that you specify a basename explicitly rather than using the default of the host name; see [Section C.5.8, “Known Issues in MySQL”](#), for the reason.

If you supply an extension in the log name (for example, `--log-bin=base_name.extension` [2002]), the extension is silently removed and ignored.

`mysqld` appends a numeric extension to the binary log basename to generate binary log file names. The number increases each time the server creates a new log file, thus creating an ordered series of files. The server creates a new file in the series each time it starts or flushes the logs. The server also creates a new binary log file automatically after the current log's size reaches `max_binlog_size` [2016]. A binary log file may become larger than `max_binlog_size` [2016] if you are using large transactions because a transaction is written to the file in one piece, never split between files.

To keep track of which binary log files have been used, `mysqld` also creates a binary log index file that contains the names of all used binary log files. By default, this has the same basename as the binary log file, with the extension `.index`. You can change the name of the binary log index file with the `--log-bin-index[=file_name]` [2002] option. You should not manually edit this file while `mysqld` is running; doing so would confuse `mysqld`.

The term “binary log file” generally denotes an individual numbered file containing database events. The term “binary log” collectively denotes the set of numbered binary log files plus the index file.

A client that has the `SUPER` [756] privilege can disable binary logging of its own statements by using a `SET sql_log_bin=0` statement. See [Section 5.1.4, “Server System Variables”](#).

By default, the server logs the length of the event as well as the event itself and uses this to verify that the event was written correctly. You can also cause the server to write checksums for the events by setting the `binlog_checksum` [2009] system variable. When reading back from the binary log, the master uses the event length by default, but can be made to use checksums if available by enabling the `master_verify_checksum` [2015] system variable. The slave I/O thread also verifies events received from the master. You can cause the slave SQL thread to use checksums if available when reading from the relay log by enabling the `slave_sql_verify_checksum` [1995] system variable.

The format of the events recorded in the binary log is dependent on the binary logging format. Three format types are supported, row-based logging, statement-based logging and mixed-base logging. The binary logging format used depends on the MySQL version. For general descriptions of the logging formats, see [Section 5.2.4.1, “Binary Logging Formats”](#). For detailed information about the format of the binary log, see [MySQL Internals: The Binary Log](#).

The server evaluates the `--binlog-do-db` [2004] and `--binlog-ignore-db` [2005] options in the same way as it does the `--replicate-do-db` [1976] and `--replicate-ignore-db` [1977] options. For information about how this is done, see [Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”](#).

A replication slave server by default does not write to its own binary log any data modifications that are received from the replication master. To log these modifications, start the slave with the `--log-slave-updates` [1971] option in addition to the `--log-bin` [2002] option (see [Section 16.1.4.3, “Replication Slave Options and Variables”](#)). This is done when a slave is also to act as a master to other slaves in chained replication.

You can delete all binary log files with the `RESET MASTER` statement, or a subset of them with `PURGE BINARY LOGS`. See [Section 13.7.6.6, “RESET Syntax”](#), and [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#).

If you are using replication, you should not delete old binary log files on the master until you are sure that no slave still needs to use them. For example, if your slaves never run more than three days



behind, once a day you can execute `mysqladmin flush-logs` on the master and then remove any logs that are more than three days old. You can remove the files manually, but it is preferable to use `PURGE BINARY LOGS`, which also safely updates the binary log index file for you (and which can take a date argument). See [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#).

You can display the contents of binary log files with the `mysqlbinlog` utility. This can be useful when you want to reprocess statements in the log for a recovery operation. For example, you can update a MySQL server from the binary log as follows:

```
shell> mysqlbinlog log_file | mysql -h server_name
```

`mysqlbinlog` also can be used to display replication slave relay log file contents because they are written using the same format as binary log files. For more information on the `mysqlbinlog` utility and how to use it, see [Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#). For more information about the binary log and recovery operations, see [Section 7.5, “Point-in-Time \(Incremental\) Recovery Using the Binary Log”](#).

Binary logging is done immediately after a statement or transaction completes but before any locks are released or any commit is done. This ensures that the log is logged in commit order.

Updates to nontransactional tables are stored in the binary log immediately after execution.

Within an uncommitted transaction, all updates (`UPDATE`, `DELETE`, or `INSERT`) that change transactional tables such as `InnoDB` tables are cached until a `COMMIT` statement is received by the server. At that point, `mysqld` writes the entire transaction to the binary log before the `COMMIT` is executed.

Modifications to nontransactional tables cannot be rolled back. If a transaction that is rolled back includes modifications to nontransactional tables, the entire transaction is logged with a `ROLLBACK` statement at the end to ensure that the modifications to those tables are replicated.

When a thread that handles the transaction starts, it allocates a buffer of `binlog_cache_size` [2008] to buffer statements. If a statement is bigger than this, the thread opens a temporary file to store the transaction. The temporary file is deleted when the thread ends.

The `Binlog_cache_use` [586] status variable shows the number of transactions that used this buffer (and possibly a temporary file) for storing statements. The `Binlog_cache_disk_use` [586] status variable shows how many of those transactions actually had to use a temporary file. These two variables can be used for tuning `binlog_cache_size` [2008] to a large enough value that avoids the use of temporary files.

The `max_binlog_cache_size` [2016] system variable (default 4GB, which is also the maximum) can be used to restrict the total size used to cache a multiple-statement transaction. If a transaction is larger than this many bytes, it fails and rolls back. The minimum value is 4096.

If you are using the binary log and row based logging, concurrent inserts are converted to normal inserts for `CREATE ... SELECT` or `INSERT ... SELECT` statements. This is done to ensure that you can re-create an exact copy of your tables by applying the log during a backup operation. If you are using statement-based logging, the original statement is written to the log.

The binary log format has some known limitations that can affect recovery from backups. See [Section 16.4.1, “Replication Features and Issues”](#).

Binary logging for stored programs is done as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

Note that the binary log format differs in MySQL 5.6 from previous versions of MySQL, due to enhancements in replication. See [Section 16.4.2, “Replication Compatibility Between MySQL Versions”](#).

Writes to the binary log file and binary log index file are handled in the same way as writes to `MyISAM` tables. See [Section C.5.4.3, “How MySQL Handles a Full Disk”](#).

By default, the binary log is not synchronized to disk at each write. So if the operating system or machine (not only the MySQL server) crashes, there is a chance that the last statements of the binary log are lost. To prevent this, you can make the binary log be synchronized to disk after every *N* writes to the binary log, with the `sync_binlog` [2017] system variable. See [Section 5.1.4, “Server System Variables”](#). 1 is the safest value for `sync_binlog` [2017], but also the slowest. Even with `sync_binlog` [2017] set to 1, there is still the chance of an inconsistency between the table content and binary log content in case of a crash. For example, if you are using `InnoDB` tables and the MySQL server processes a `COMMIT` statement, it writes the whole transaction to the binary log and then commits this transaction into `InnoDB`. If the server crashes between those two operations, the transaction is rolled back by `InnoDB` at restart but still exists in the binary log. To resolve this, you should set `--innodb_support_xa` [1749] to 1. Although this option is related to the support of XA transactions in `InnoDB`, it also ensures that the binary log and `InnoDB` data files are synchronized.

For this option to provide a greater degree of safety, the MySQL server should also be configured to synchronize the binary log and the `InnoDB` logs to disk before committing the transaction. The `InnoDB` logs are synchronized by default, and `sync_binlog=1` can be used to synchronize the binary log. The effect of this option is that at restart after a crash, after doing a rollback of transactions, the MySQL server cuts rolled back `InnoDB` transactions from the binary log. This ensures that the binary log reflects the exact data of `InnoDB` tables, and so, that the slave remains in synchrony with the master (not receiving a statement which has been rolled back).

If the MySQL server discovers at crash recovery that the binary log is shorter than it should have been, it lacks at least one successfully committed `InnoDB` transaction. This should not happen if `sync_binlog=1` and the disk/file system do an actual sync when they are requested to (some do not), so the server prints an error message `The binary log file_name is shorter than its expected size`. In this case, this binary log is not correct and replication should be restarted from a fresh snapshot of the master's data.

The session values of the following system variables are written to the binary log and honored by the replication slave when parsing the binary log:

- `sql_mode` [545] (except that the `NO_DIR_IN_CREATE` [604] mode is not replicated; see [Section 16.4.1.32, “Replication and Variables”](#))
- `foreign_key_checks` [483]
- `unique_checks` [561]
- `character_set_client` [467]
- `collation_connection` [469]
- `collation_database` [469]
- `collation_server` [470]
- `sql_auto_is_null` [543]

### 5.2.4.1. Binary Logging Formats

The server uses several logging formats to record information in the binary log. The exact format employed depends on the version of MySQL being used. There are three logging formats:

- Replication capabilities in MySQL originally were based on propagation of SQL statements from master to slave. This is called *statement-based logging*. You can cause this format to be used by starting the server with `--binlog-format=STATEMENT` [415].
- In *row-based logging*, the master writes events to the binary log that indicate how individual table rows are affected. You can cause the server to use row-based logging by starting it with `--binlog-format=ROW` [415].

- A third option is also available: *mixed logging*. With mixed logging, statement-based logging is used by default, but the logging mode switches automatically to row-based in certain cases as described below. You can cause MySQL to use mixed logging explicitly by starting `mysqld` with the option `--binlog-format=MIXED` [415].

In MySQL 5.6, the default binary logging format is `STATEMENT`.

The logging format can also be set or limited by the storage engine being used. This helps to eliminate issues when replicating certain statements between a master and slave which are using different storage engines.

With statement-based replication, there may be issues with replicating nondeterministic statements. In deciding whether or not a given statement is safe for statement-based replication, MySQL determines whether it can guarantee that the statement can be replicated using statement-based logging. If MySQL cannot make this guarantee, it marks the statement as potentially unreliable and issues the warning, `Statement may not be safe to log in statement format`.

You can avoid these issues by using MySQL's row-based replication instead.

#### 5.2.4.2. Setting The Binary Log Format

You can select the binary logging format explicitly by starting the MySQL server with `--binlog-format=type` [415]. The supported values for *type* are:

- `STATEMENT` causes logging to be statement based.
- `ROW` causes logging to be row based.
- `MIXED` causes logging to use mixed format.

In MySQL 5.6, the default binary logging format is `STATEMENT`.

The logging format also can be switched at runtime. To specify the format globally for all clients, set the global value of the `binlog_format` [2010] system variable:

```
mysql> SET GLOBAL binlog_format = 'STATEMENT';
mysql> SET GLOBAL binlog_format = 'ROW';
mysql> SET GLOBAL binlog_format = 'MIXED';
```

An individual client can control the logging format for its own statements by setting the session value of `binlog_format` [2010]:

```
mysql> SET SESSION binlog_format = 'STATEMENT';
mysql> SET SESSION binlog_format = 'ROW';
mysql> SET SESSION binlog_format = 'MIXED';
```



#### Note

Each MySQL Server can set its own and only its own binary logging format (true whether `binlog_format` [2010] is set with global or session scope). This means that changing the logging format on a replication master does not cause a slave to change its logging format to match. (When using `STATEMENT` mode, the `binlog_format` [2010] system variable is not replicated; when using `MIXED` or `ROW` logging mode, it is replicated but is ignored by the slave.) Changing the binary logging format on the master while replication is ongoing, or without also changing it on the slave can thus cause unexpected results, or even cause replication to fail altogether.

To change the global or session `binlog_format` [2010] value, you must have the `SUPER` [756] privilege.

In addition to switching the logging format manually, a slave server may switch the format *automatically*. This happens when the server is running in either `STATEMENT` or `MIXED` format and encounters an event in the binary log that is written in `ROW` logging format. In that case, the slave

switches to row-based replication temporarily for that event, and switches back to the previous format afterward.

There are several reasons why a client might want to set binary logging on a per-session basis:

- A session that makes many small changes to the database might want to use row-based logging.
- A session that performs updates that match many rows in the `WHERE` clause might want to use statement-based logging because it will be more efficient to log a few statements than many rows.
- Some statements require a lot of execution time on the master, but result in just a few rows being modified. It might therefore be beneficial to replicate them using row-based logging.

There are exceptions when you cannot switch the replication format at runtime:

- From within a stored function or a trigger
- If the `NDB` storage engine is enabled
- If the session is currently in row-based replication mode and has open temporary tables

Trying to switch the format in any of these cases results in an error.

If you are using `InnoDB` tables and the transaction isolation level is `READ COMMITTED` [1469] or `READ UNCOMMITTED` [1470], only row-based logging can be used. It is *possible* to change the logging format to `STATEMENT`, but doing so at runtime leads very rapidly to errors because `InnoDB` can no longer perform inserts.

Switching the replication format at runtime is not recommended when any temporary tables exist, because temporary tables are logged only when using statement-based replication, whereas with row-based replication they are not logged. With mixed replication, temporary tables are usually logged; exceptions happen with user-defined functions (UDFs) and with the `UUID()` [1280] function.

With the binary log format set to `ROW`, many changes are written to the binary log using the row-based format. Some changes, however, still use the statement-based format. Examples include all DDL (data definition language) statements such as `CREATE TABLE`, `ALTER TABLE`, or `DROP TABLE`.

The `--binlog-row-event-max-size` [2001] option is available for servers that are capable of row-based replication. Rows are stored into the binary log in chunks having a size in bytes not exceeding the value of this option. The value must be a multiple of 256. The default value is 1024.



#### Warning

When using *statement-based logging* for replication, it is possible for the data on the master and slave to become different if a statement is designed in such a way that the data modification is *nondeterministic*; that is, it is left to the will of the query optimizer. In general, this is not a good practice even outside of replication. For a detailed explanation of this issue, see [Section C.5.8, “Known Issues in MySQL”](#).

For information about logs kept by replication slaves, see [Section 16.2.2, “Replication Relay and Status Logs”](#).

### 5.2.4.3. Mixed Binary Logging Format

When running in `MIXED` logging format, the server automatically switches from statement-based to row-based logging under the following conditions:

- When a function contains `UUID()` [1280].
- When one or more tables with `AUTO_INCREMENT` columns are updated and a trigger or stored function is invoked. Like all other unsafe statements, this generates a warning if `binlog_format = STATEMENT` [2010].

For more information, see [Section 16.4.1.1, “Replication and `AUTO\_INCREMENT`”](#).

- When the body of a view requires row-based replication, the statement creating the view also uses it. For example, this occurs when the statement creating a view uses the `UUID()` [\[1280\]](#) function.
- When a call to a UDF is involved.
- When any `INSERT DELAYED` is executed for a nontransactional table.
- If a statement is logged by row and the session that executed the statement has any temporary tables, logging by row is used for all subsequent statements (except for those accessing temporary tables) until all temporary tables in use by that session are dropped.

This is true whether or not any temporary tables are actually logged.

Temporary tables cannot be logged using row-based format; thus, once row-based logging is used, all subsequent statements using that table are unsafe. The server approximates this condition by treating all statements executed during the session as unsafe until the session no longer holds any temporary tables.

- When `FOUND_ROWS()` [\[1268\]](#) or `ROW_COUNT()` [\[1272\]](#) is used. (Bug #12092, Bug #30244)
- When `USER()` [\[1273\]](#), `CURRENT_USER()` [\[1267\]](#), or `CURRENT_USER` [\[1267\]](#) is used. (Bug #28086)
- When a statement refers to one or more system variables. (Bug #31168)

**Exception.** The following system variables, when used with session scope (only), do not cause the logging format to switch:

- `auto_increment_increment` [\[1965\]](#)
- `auto_increment_offset` [\[1968\]](#)
- `character_set_client` [\[467\]](#)
- `character_set_connection` [\[467\]](#)
- `character_set_database` [\[468\]](#)
- `character_set_server` [\[468\]](#)
- `collation_connection` [\[469\]](#)
- `collation_database` [\[469\]](#)
- `collation_server` [\[470\]](#)
- `foreign_key_checks` [\[483\]](#)
- `identity` [\[488\]](#)
- `last_insert_id` [\[495\]](#)
- `lc_time_names` [\[495\]](#)
- `pseudo_thread_id` [\[527\]](#)
- `sql_auto_is_null` [\[543\]](#)
- `time_zone` [\[557\]](#)
- `timestamp` [\[558\]](#)

- [unique\\_checks](#) [561]

For information about determining system variable scope, see [Section 5.1.5, “Using System Variables”](#).

For information about how replication treats [sql\\_mode](#) [545], see [Section 16.4.1.32, “Replication and Variables”](#).

- When one of the tables involved is a log table in the `mysql` database.
- When the `LOAD_FILE()` [1181] function is used. (Bug #39701)



**Note**

A warning is generated if you try to execute a statement using statement-based logging that should be written using row-based logging. The warning is shown both in the client (in the output of `SHOW WARNINGS`) and through the `mysqld` error log. A warning is added to the `SHOW WARNINGS` table each time such a statement is executed. However, only the first statement that generated the warning for each client session is written to the error log to prevent flooding the log.

In addition to the decisions above, individual engines can also determine the logging format used when information in a table is updated. The logging capabilities of an individual engine can be defined as follows:

- If an engine supports row-based logging, the engine is said to be *row-logging capable*.
- If an engine supports statement-based logging, the engine is said to be *statement-logging capable*.

A given storage engine can support either or both logging formats. The following table lists the formats supported by each engine.

Storage Engine	Row Logging Supported	Statement Logging Supported
<a href="#">ARCHIVE</a>	Yes	Yes
<a href="#">BLACKHOLE</a>	Yes	Yes
<a href="#">CSV</a>	Yes	Yes
<a href="#">EXAMPLE</a>	Yes	No
<a href="#">FEDERATED</a>	Yes	Yes
<a href="#">HEAP</a>	Yes	Yes
<a href="#">InnoDB</a>	Yes	Yes when the transaction isolation level is <a href="#">REPEATABLE READ</a> [1469] or <a href="#">SERIALIZABLE</a> [1470]; No otherwise.
<a href="#">MyISAM</a>	Yes	Yes
<a href="#">MERGE</a>	Yes	Yes
<a href="#">NDB</a>	Yes	No

In MySQL 5.6, whether a statement is to be logged and the logging mode to be used is determined according to the type of statement (safe, unsafe, or binary injected), the binary logging format (`STATEMENT`, `ROW`, or `MIXED`), and the logging capabilities of the storage engine (statement capable, row capable, both, or neither). Statements may be logged with or without a warning; failed statements

are not logged, but generate errors in the log. This is shown in the following decision table, where **SLC** stands for “statement-logging capable” and **RLC** stands for “row-logging capable”.

Condition				Action	
Type	<code>binlog_format</code> [2010]	SLC	RLC	Error / Warning	Logged as
*	*	No	No	<code>Error: Cannot execute statement:</code> Binary logging is impossible since at least one engine is involved that is both row-incapable and statement-incapable.	-
Safe	<code>STATEMENT</code>	Yes	No	-	<code>STATEMENT</code>
Safe	<code>MIXED</code>	Yes	No	-	<code>STATEMENT</code>
Safe	<code>ROW</code>	Yes	No	<code>Error: Cannot execute statement:</code> Binary logging is impossible since <code>BINLOG_FORMAT = ROW</code> and at least one table uses a storage engine that is not capable of row-based logging.	-
Unsafe	<code>STATEMENT</code>	Yes	No	<code>Warning: Unsafe statement binlogged in statement format, since BINLOG_FORMAT = STATEMENT</code>	<code>STATEMENT</code>
Unsafe	<code>MIXED</code>	Yes	No	<code>Error: Cannot execute statement:</code> Binary logging of an unsafe statement is impossible when the storage engine is limited to statement-based logging, even if <code>BINLOG_FORMAT = MIXED</code> .	-
Unsafe	<code>ROW</code>	Yes	No	<code>Error: Cannot execute statement:</code> Binary logging is impossible since <code>BINLOG_FORMAT</code>	-

Condition				Action	
Type	binlog_format [2010]	SLC	RLC	Error / Warning	Logged as
				= ROW and at least one table uses a storage engine that is not capable of row-based logging.	
Row Injection	STATEMENT	Yes	No	Error: Cannot execute row injection: Binary logging is not possible since at least one table uses a storage engine that is not capable of row-based logging.	-
Row Injection	MIXED	Yes	No	Error: Cannot execute row injection: Binary logging is not possible since at least one table uses a storage engine that is not capable of row-based logging.	-
Row Injection	ROW	Yes	No	Error: Cannot execute row injection: Binary logging is not possible since at least one table uses a storage engine that is not capable of row-based logging.	-
Safe	STATEMENT	No	Yes	Error: Cannot execute statement: Binary logging is impossible since BINLOG_FORMAT = STATEMENT and at least one table uses a storage engine that is not capable of statement-based logging.	-
Safe	MIXED	No	Yes	-	ROW
Safe	ROW	No	Yes	-	ROW
Unsafe	STATEMENT	No	Yes	Error: Cannot execute	-



Condition				Action	
Type	<code>binlog_format</code> [2010]	SLC	RLC	Error / Warning	Logged as
				<code>statement:</code> Binary logging is impossible since <code>BINLOG_FORMAT = STATEMENT</code> and at least one table uses a storage engine that is not capable of statement-based logging.	
Unsafe	<code>MIXED</code>	No	Yes	-	<code>ROW</code>
Unsafe	<code>ROW</code>	No	Yes	-	<code>ROW</code>
Row Injection	<code>STATEMENT</code>	No	Yes	<code>Error: Cannot execute row injection:</code> Binary logging is not possible since <code>BINLOG_FORMAT = STATEMENT</code> .	-
Row Injection	<code>MIXED</code>	No	Yes	-	<code>ROW</code>
Row Injection	<code>ROW</code>	No	Yes	-	<code>ROW</code>
Safe	<code>STATEMENT</code>	Yes	Yes	-	<code>STATEMENT</code>
Safe	<code>MIXED</code>	Yes	Yes	-	<code>ROW</code>
Safe	<code>ROW</code>	Yes	Yes	-	<code>ROW</code>
Unsafe	<code>STATEMENT</code>	Yes	Yes	<code>Warning:</code> Unsafe statement binlogged in statement format since <code>BINLOG_FORMAT = STATEMENT</code> .	<code>STATEMENT</code>
Unsafe	<code>MIXED</code>	Yes	Yes	-	<code>ROW</code>
Unsafe	<code>ROW</code>	Yes	Yes	-	<code>ROW</code>
Row Injection	<code>STATEMENT</code>	Yes	Yes	<code>Error: Cannot execute row injection:</code> Binary logging is not possible because <code>BINLOG_FORMAT = STATEMENT</code> .	-
Row Injection	<code>MIXED</code>	Yes	Yes	-	<code>ROW</code>
Row Injection	<code>ROW</code>	Yes	Yes	-	<code>ROW</code>

When a warning is produced by the determination, a standard MySQL warning is produced (and is available using `SHOW WARNINGS`). The information is also written to the `mysqld` error log. Only one error for each error instance per client connection is logged to prevent flooding the log. The log message includes the SQL statement that was attempted.

If a slave server was started with `--log-warnings` [430] enabled, the slave prints messages to the error log to provide information about its status, such as the binary log and relay log coordinates where it starts its job, when it is switching to another relay log, when it reconnects after a disconnect, and so forth.

#### 5.2.4.4. Logging Format for Changes to `mysql` Database Tables

The contents of the grant tables in the `mysql` database can be modified directly (for example, with `INSERT` or `DELETE`) or indirectly (for example, with `GRANT` or `CREATE USER`). Statements that affect `mysql` database tables are written to the binary log using the following rules:

- Data manipulation statements that change data in `mysql` database tables directly are logged according to the setting of the `binlog_format` [2010] system variable. This pertains to statements such as `INSERT`, `UPDATE`, `DELETE`, `REPLACE`, `DO`, `LOAD DATA INFILE`, `SELECT`, and `TRUNCATE TABLE`.
- Statements that change the `mysql` database indirectly are logged as statements regardless of the value of `binlog_format` [2010]. This pertains to statements such as `GRANT`, `REVOKE`, `SET PASSWORD`, `RENAME USER`, `CREATE` (all forms except `CREATE TABLE ... SELECT`), `ALTER` (all forms), and `DROP` (all forms).

`CREATE TABLE ... SELECT` is a combination of data definition and data manipulation. The `CREATE TABLE` part is logged using statement format and the `SELECT` part is logged according to the value of `binlog_format` [2010].

#### 5.2.5. The Slow Query Log

The slow query log consists of SQL statements that took more than `long_query_time` [500] seconds to execute and required at least `min_examined_row_limit` [510] rows to be examined. The minimum and default values of `long_query_time` [500] are 0 and 10, respectively. The value can be specified to a resolution of microseconds. For logging to a file, times are written including the microseconds part. For logging to tables, only integer times are written; the microseconds part is ignored.

By default, administrative statements are not logged, nor are queries that do not use indexes for lookups. This behavior can be changed using `--log-slow-admin-statements` [429] and `log_queries_not_using_indexes` [498], as described later.

The time to acquire the initial locks is not counted as execution time. `mysqld` writes a statement to the slow query log after it has been executed and after all locks have been released, so log order might differ from execution order.

By default, the slow query log is disabled. To specify the initial slow query log state explicitly, use `--slow_query_log[={0|1}]` [541]. With no argument or an argument of 1, `--slow_query_log` [541] enables the log. With an argument of 0, this option disables the log. To specify a log file name, use `--slow_query_log_file=file_name` [541]. To specify the log destination, use `--log-output` [427] (as described in Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”).

If you specify no name for the slow query log file, the default name is `host_name-slow.log`. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.

To disable or enable the slow query log or change the log file name at runtime, use the global `slow_query_log` [541] and `slow_query_log_file` [541] system variables.

Set `slow_query_log` [541] to 0 (or `OFF`) to disable the log or to 1 (or `ON`) to enable it. Set `slow_query_log_file` [541] to specify the name of the log file. If a log file already is open, it is closed and the new file is opened.

When the slow query log is enabled, the server writes output to any destinations specified by the `--log-output` [427] option or `log_output` [498] system variable. If you enable the log, the server opens the log file and writes startup messages to it. However, further logging of queries to the file does not occur unless the `FILE` log destination is selected. If the destination is `NONE`, the server writes no queries even if the slow query log is enabled. Setting the log file name has no effect on logging if the log destination value does not contain `FILE`.

The server writes less information to the slow query log (and binary log) if you use the `--log-short-format` [428] option.

To include slow administrative statements in the statements written to the slow query log, use the `--log-slow-admin-statements` [429] server option. Administrative statements include `ALTER TABLE`, `ANALYZE TABLE`, `CHECK TABLE`, `CREATE INDEX`, `DROP INDEX`, `OPTIMIZE TABLE`, and `REPAIR TABLE`.

To include queries that do not use indexes for row lookups in the statements written to the slow query log, enable the `log_queries_not_using_indexes` [498] system variable. When such queries are logged, the slow query log may grow quickly. It is possible to put a rate limit on these queries by setting the `log_throttle_queries_not_using_indexes` [499] system variable. By default, this variable is 0, which means there is no limit. Positive values impose a per-minute limit on logging of queries that do not use indexes. The first such query opens a 60-second window within which the server logs queries up to the given limit, then suppresses additional queries. If there are suppressed queries when the window ends, the server logs a summary that indicates how many there were and the aggregate time spent in them. The next 60-second window begins when the server logs the next query that does not use indexes.

The server uses the controlling parameters in the following order to determine whether to write a query to the slow query log:

1. The query must either not be an administrative statement, or `--log-slow-admin-statements` [429] must have been specified.
2. The query must have taken at least `long_query_time` [500] seconds, or `log_queries_not_using_indexes` [498] must be enabled and the query used no indexes for row lookups.
3. The query must have examined at least `min_examined_row_limit` [510] rows.
4. The query must not be suppressed according to the `log_throttle_queries_not_using_indexes` [499] setting.

The server does not write queries handled by the query cache to the slow query log, nor queries that would not benefit from the presence of an index because the table has zero rows or one row.

By default, a replication slave does not write replicated queries to the slow query log. To change this, use the `--log-slow-slave-statements` [1972] server option.

As of MySQL 5.6.3, passwords in statements written to the slow query log are rewritten by the server not to occur literally in plain text. Before MySQL 5.6.3, passwords in statements are not rewritten and the slow query log should be protected. See [Section 6.1.2.3, “Passwords and Logging”](#).

The slow query log can be used to find queries that take a long time to execute and are therefore candidates for optimization. However, examining a long slow query log can become a difficult task. To make this easier, you can process a slow query log file using the `mysqldumpslow` command to summarize the queries that appear in the log. See [Section 4.6.9, “mysqldumpslow — Summarize Slow Query Log Files”](#).

## 5.2.6. Server Log Maintenance

As described in [Section 5.2, “MySQL Server Logs”](#), MySQL Server can create several different log files to help you see what activity is taking place. However, you must clean up these files regularly to ensure that the logs do not take up too much disk space.

When using MySQL with logging enabled, you may want to back up and remove old log files from time to time and tell MySQL to start logging to new files. See [Section 7.2, “Database Backup Methods”](#).

On a Linux (Red Hat) installation, you can use the `mysql-log-rotate` script for this. If you installed MySQL from an RPM distribution, this script should have been installed automatically. Be careful with this script if you are using the binary log for replication. You should not remove binary logs until you are certain that their contents have been processed by all slaves.

On other systems, you must install a short script yourself that you start from `cron` (or its equivalent) for handling log files.

For the binary log, you can set the `expire_logs_days` [480] system variable to expire binary log files automatically after a given number of days (see [Section 5.1.4, “Server System Variables”](#)). If you are using replication, you should set the variable no lower than the maximum number of days your slaves might lag behind the master. To remove binary logs on demand, use the `PURGE BINARY LOGS` statement (see [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#)).

You can force MySQL to start using new log files by flushing the logs. Log flushing occurs when you issue a `FLUSH LOGS` statement or execute a `mysqladmin flush-logs`, `mysqladmin refresh`, `mysqldump --flush-logs`, or `mysqldump --master-data` command. See [Section 13.7.6.3, “FLUSH Syntax”](#), [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#), and [Section 4.5.4, “mysqldump — A Database Backup Program”](#). In addition, the binary log is flushed when its size reaches the value of the `max_binlog_size` [2016] system variable.

`FLUSH LOGS` supports optional modifiers to enable selective flushing of individual logs (for example, `FLUSH BINARY LOGS`).

A log-flushing operation does the following:

- If general query logging or slow query logging to a log file is enabled, the server closes and reopens the general query log file or slow query log file.
- If binary logging is enabled, the server closes the current binary log file and opens a new log file with the next sequence number.
- If the server was started with the `--log-error` [427] option to cause the error log to be written to a file, the server closes and reopens the log file.

The server creates a new binary log file when you flush the logs. However, it just closes and reopens the general and slow query log files. To cause new files to be created on Unix, rename the current log files before flushing them. At flush time, the server opens new log files with the original names. For example, if the general and slow query log files are named `mysql.log` and `mysql-slow.log`, you can use a series of commands like this:

```
shell> cd mysql-data-directory
shell> mv mysql.log mysql.old
shell> mv mysql-slow.log mysql-slow.old
shell> mysqladmin flush-logs
```

On Windows, use `rename` rather than `mv`.

At this point, you can make a backup of `mysql.old` and `mysql-slow.old` and then remove them from disk.

A similar strategy can be used to back up the error log file, if there is one.

You can rename the general query log or slow query log at runtime by disabling the log:

```
SET GLOBAL general_log = 'OFF';
SET GLOBAL slow_query_log = 'OFF';
```

With the logs disabled, rename the log files externally; for example, from the command line. Then enable the logs again:

```
SET GLOBAL general_log = 'ON';
SET GLOBAL slow_query_log = 'ON';
```

This method works on any platform and does not require a server restart.

## 5.3. Managing Disk I/O and File Space for InnoDB Tables

As a DBA, you must manage disk I/O to keep the I/O subsystem from becoming saturated, and manage disk space to avoid filling up storage devices. The [ACID](#) design model requires a certain amount of I/O that might seem redundant, but helps to ensure data reliability. Within these constraints, InnoDB tries to optimize the database work and the organization of disk files to minimize the amount of disk I/O. Sometimes, I/O is postponed until the database is not busy, or until everything needs to be brought to a consistent state, such as during a database restart after a [fast shutdown](#).

This section discusses the main considerations for I/O and disk space with the default kind of MySQL tables (also known as InnoDB tables):

- Controlling the amount of background I/O used to improve query performance.
- Enabling or disabling features that provide extra durability at the expense of additional I/O.
- Organizing tables into many small files, a few larger files, or a combination of both.
- Balancing the size of redo log files against the I/O activity that occurs when the log files become full.
- How to reorganize a table for optimal query performance.

### 5.3.1. InnoDB Disk I/O

InnoDB uses asynchronous disk I/O where possible, by creating a number of threads to handle I/O operations, while permitting other database operations to proceed while the I/O is still in progress. On Linux and Windows platforms, InnoDB uses the available OS and library functions to perform “native” asynchronous I/O. On other platforms, InnoDB still uses I/O threads, but the threads may actually wait for I/O requests to complete; this technique is known as “simulated” asynchronous I/O.

#### Read-Ahead

If InnoDB can determine there is a high probability that data might be needed soon, it performs read-ahead operations to bring that data into the buffer pool so that it is available in memory. Making a few large read requests for contiguous data can be more efficient than making several small, spread-out requests. There are two read-ahead heuristics in InnoDB:

- In sequential read-ahead, if InnoDB notices that the access pattern to a segment in the tablespace is sequential, it posts in advance a batch of reads of database pages to the I/O system.
- In random read-ahead, if InnoDB notices that some area in a tablespace seems to be in the process of being fully read into the buffer pool, it posts the remaining reads to the I/O system.

#### Doublewrite Buffer

InnoDB uses a novel file flush technique involving a structure called the [doublewrite buffer](#). It adds safety to recovery following an operating system crash or a power outage, and improves performance on most varieties of Unix by reducing the need for `fsync()` operations.

Before writing pages to a data file, [InnoDB](#) first writes them to a contiguous tablespace area called the doublewrite buffer. Only after the write and the flush to the doublewrite buffer has completed does [InnoDB](#) write the pages to their proper positions in the data file. If the operating system crashes in the middle of a page write (causing a [torn page](#) condition), [InnoDB](#) can later find a good copy of the page from the doublewrite buffer during recovery.

### 5.3.2. File Space Management

The data files that you define in the configuration file form the [InnoDB system tablespace](#). The files are logically concatenated to form the tablespace. There is no striping in use. Currently, you cannot define where within the tablespace your tables are allocated. In a newly created tablespace, [InnoDB](#) allocates space starting from the first data file.

To avoid the issues that come with storing all tables and indexes inside the system tablespace, you can turn on the [innodb\\_file\\_per\\_table](#) configuration option, which stores each newly created table in a separate tablespace file (with extension `.ibd`). For tables stored this way, there is less fragmentation within the disk file, and when the table is truncated, the space is returned to the operating system rather than still being reserved by [InnoDB](#) within the system tablespace.

### Pages, Extents, Segments, and Tablespaces

Each tablespace consists of database [pages](#). Every tablespace in a MySQL instance has the same [page size](#). By default, all tablespaces have a page size of 16KB; you can reduce the page size to 8KB or 4KB by specifying the [innodb\\_page\\_size \[1740\]](#) option when you create the MySQL instance.

The pages are grouped into [extents](#) of size 1MB (64 consecutive 16KB pages, or 128 8KB pages, or 256 4KB pages). The “files” inside a tablespace are called [segments](#) in [InnoDB](#). (These segments are different from the [rollback segment](#), which actually contains many tablespace segments.)

When a segment grows inside the tablespace, [InnoDB](#) allocates the first 32 pages to it one at a time. After that, [InnoDB](#) starts to allocate whole extents to the segment. [InnoDB](#) can add up to 4 extents at a time to a large segment to ensure good sequentiality of data.

Two segments are allocated for each index in [InnoDB](#). One is for nonleaf nodes of the B-tree, the other is for the leaf nodes. Keeping the leaf nodes contiguous on disk enables better sequential I/O operations, because these leaf nodes contain the actual table data.

Some pages in the tablespace contain bitmaps of other pages, and therefore a few extents in an [InnoDB](#) tablespace cannot be allocated to segments as a whole, but only as individual pages.

When you ask for available free space in the tablespace by issuing a [SHOW TABLE STATUS](#) statement, [InnoDB](#) reports the extents that are definitely free in the tablespace. [InnoDB](#) always reserves some extents for cleanup and other internal purposes; these reserved extents are not included in the free space.

When you delete data from a table, [InnoDB](#) contracts the corresponding B-tree indexes. Whether the freed space becomes available for other users depends on whether the pattern of deletes frees individual pages or extents to the tablespace. Dropping a table or deleting all rows from it is guaranteed to release the space to other users, but remember that deleted rows are physically removed only by the [purge](#) operation, which happens automatically some time after they are no longer needed for transaction rollbacks or consistent reads. (See [Section 14.2.3.11](#), “[InnoDB Multi-Versioning](#)”.)

To see information about the tablespace, use the Tablespace Monitor. See [Section 14.2.4.4](#), “[SHOW ENGINE INNODB STATUS and the InnoDB Monitors](#)”.

### How Pages Relate to Table Rows

The maximum row length, except for variable-length columns ([VARBINARY](#), [VARCHAR](#), [BLOB](#) and [TEXT](#)), is slightly less than half of a database page. That is, the maximum row length is about 8000



bytes. `LONGBLOB` and `LONGTEXT` columns must be less than 4GB, and the total row length, including `BLOB` and `TEXT` columns, must be less than 4GB.

If a row is less than half a page long, all of it is stored locally within the page. If it exceeds half a page, variable-length columns are chosen for external off-page storage until the row fits within half a page. For a column chosen for off-page storage, InnoDB stores the first 768 bytes locally in the row, and the rest externally into overflow pages. Each such column has its own list of overflow pages. The 768-byte prefix is accompanied by a 20-byte value that stores the true length of the column and points into the overflow list where the rest of the value is stored.

### 5.3.3. InnoDB Checkpoints

Making your `log files` very large may reduce disk I/O during `checkpointing`. It often makes sense to set the total size of the log files as large as the buffer pool or even larger. Although in the past large log files could make crash recovery take excessive time, starting with MySQL 5.5, performance enhancements to crash recovery make it possible to use large log files with fast startup after a crash. (Strictly speaking, this performance improvement is available for MySQL 5.1 with the InnoDB Plugin 1.0.7 and higher. It is with MySQL 5.5 that this improvement is available in the default InnoDB storage engine.)

### How Checkpoint Processing Works

InnoDB implements a `checkpoint` mechanism known as `fuzzy checkpointing`. InnoDB flushes modified database pages from the buffer pool in small batches. There is no need to flush the buffer pool in one single batch, which would disrupt processing of user SQL statements during the checkpointing process.

During `crash recovery`, InnoDB looks for a checkpoint label written to the log files. It knows that all modifications to the database before the label are present in the disk image of the database. Then InnoDB scans the log files forward from the checkpoint, applying the logged modifications to the database.

### 5.3.4. Defragmenting a Table

Random insertions into or deletions from a secondary index can cause the index to become fragmented. Fragmentation means that the physical ordering of the index pages on the disk is not close to the index ordering of the records on the pages, or that there are many unused pages in the 64-page blocks that were allocated to the index.

One symptom of fragmentation is that a table takes more space than it “should” take. How much that is exactly, is difficult to determine. All InnoDB data and indexes are stored in B-trees, and their `fill factor` may vary from 50% to 100%. Another symptom of fragmentation is that a table scan such as this takes more time than it “should” take:

```
SELECT COUNT(*) FROM t WHERE non_indexed_column <> 12345;
```

The preceding query requires MySQL to perform a full table scan, the slowest type of query for a large table.

To speed up index scans, you can periodically perform a “null” `ALTER TABLE` operation, which causes MySQL to rebuild the table:

```
ALTER TABLE tbl_name ENGINE=INNODB
```

Another way to perform a defragmentation operation is to use `mysqldump` to dump the table to a text file, drop the table, and reload it from the dump file.

If the insertions into an index are always ascending and records are deleted only from the end, the InnoDB filespace management algorithm guarantees that fragmentation in the index does not occur.

## 5.4. Creating and Using InnoDB Tables and Indexes

To create an [InnoDB](#) table, use the `CREATE TABLE` statement without any special clauses. Formerly, you needed the `ENGINE=InnoDB` clause, but not anymore now that [InnoDB](#) is the default storage engine. (You might still use that clause if you plan to use `mysqldump` or replication to replay the `CREATE TABLE` statement on a server running MySQL 5.1 or earlier, where the default storage engine is [MyISAM](#).)

```
-- Default storage engine = InnoDB.
CREATE TABLE t1 (a INT, b CHAR (20), PRIMARY KEY (a));
-- Backwards-compatible with older MySQL.
CREATE TABLE t2 (a INT, b CHAR (20), PRIMARY KEY (a)) ENGINE=InnoDB;
```

Depending on the file-per-table setting, [InnoDB](#) creates each table and associated primary key index either in the [system tablespace](#), or in a separate tablespace (represented by a `.ibd` file) for each table. MySQL creates `t1.frm` and `t2.frm` files in the `test` directory under the MySQL database directory. Internally, [InnoDB](#) adds an entry for the table to its own data dictionary. The entry includes the database name. For example, if `test` is the database in which the `t1` table is created, the entry is for `'test/t1'`. This means you can create a table of the same name `t1` in some other database, and the table names do not collide inside [InnoDB](#).

To see the properties of these tables, issue a `SHOW TABLE STATUS` statement:

```
SHOW TABLE STATUS FROM test LIKE 't%';
```

In the status output, you see the [row format](#) property of these first tables is [Compact](#). Although that setting is fine for basic experimentation, to take advantage of the most powerful [InnoDB](#) performance features, you will quickly graduate to using other row formats such as [Dynamic](#) and [Compressed](#). Using those values requires a little bit of setup first:

```
set global innodb_file_per_table=1;
set global innodb_file_format=barracuda;
CREATE TABLE t3 (a INT, b CHAR (20), PRIMARY KEY (a)) row_format=dynamic;
CREATE TABLE t4 (a INT, b CHAR (20), PRIMARY KEY (a)) row_format=compressed;
```

Always set up a [primary key](#) for each [InnoDB](#) table, specifying the column or columns that:

- Are referenced by the most important queries.
- Are never left blank.
- Never have duplicate values.
- Rarely if ever change value once inserted.

For example, in a table containing information about people, you would not create a primary key on `(firstname, lastname)` because more than one person can have the same name, some people have blank last names, and sometimes people change their names. With so many constraints, often there is not an obvious set of columns to use as a primary key, so you create a new column with a numeric ID to serve as all or part of the primary key. You can declare an [auto-increment](#) column so that ascending values are filled in automatically as rows are inserted:

```
-- The value of ID can act like a pointer between related items in different tables.
CREATE TABLE t5 (id INT AUTO_INCREMENT, b CHAR (20), PRIMARY KEY (id));
-- The primary key can consist of more than one column. Any autoinc column must come first.
CREATE TABLE t6 (id INT AUTO_INCREMENT, a INT, b CHAR (20), PRIMARY KEY (id,a));
```

Although the table *works* correctly without you defining a primary key, the primary key is involved with many aspects of performance and is a crucial design aspect for any large or frequently used table. Make a habit of always specifying one in the `CREATE TABLE` statement. (If you create the table, load data, and then do `ALTER TABLE` to add a primary key later, that operation is much slower than defining the primary key when creating the table.)

### 5.4.1. Managing InnoDB Tablespaces



Historically, all InnoDB tables and indexes were stored in the [system tablespace](#). This monolithic approach was targeted at machines dedicated entirely to database processing, with carefully planned data growth, where any disk storage allocated to MySQL would never be needed for other purposes. InnoDB's [file-per-table](#) mode is a more flexible alternative, where you store each InnoDB table and its indexes in a separate file. Each such `.ibd` file represents a separate [tablespace](#). This mode is controlled by the `innodb_file_per_table` [1718] configuration option, and is the default in MySQL 5.6.6 and higher.

Multiple tablespaces are useful in a number of situations:

- You can reclaim disk space when truncating or dropping a table. For tables created when file-per-table mode is turned off, truncating or dropping them creates free space internally in the [ibdata files](#). That free space can only be used for new InnoDB data.
- You can store specific tables on separate storage devices, for I/O optimization, space management, or backup purposes. In previous releases, you had to move entire database directories to other drives and create symbolic links in the MySQL data directory, as described in [Section 8.11.3.1, “Using Symbolic Links”](#). In MySQL 5.6 and higher, you can specify the location of each table using the syntax `CREATE TABLE ... DATA DIRECTORY = absolute_path_to_directory`, as explained in [Section 5.4.1.2, “Specifying the Location of a Tablespace”](#).
- You can copy individual InnoDB tables from one MySQL instance to another (known as the [transportable tablespace](#) feature).
- You can enable [compression](#) for table and index data, using the [compressed row format](#).
- You can enable more efficient storage for tables with large BLOB or text columns using the [dynamic row format](#).
- You can back up or restore a single table quickly, without interrupting the use of other InnoDB tables, using the MySQL Enterprise Backup product. See [Backing Up and Restoring a Single .ibd File](#) for the procedure and restrictions.

#### 5.4.1.1. Enabling and Disabling File-Per-Table Mode

To make file-per-table mode the default for a MySQL server, start the server with the `--innodb_file_per_table` [1718] command-line option, or add this line to the `[mysqld]` section of `my.cnf`:

```
[mysqld]
innodb_file_per_table
```

You can also issue the command while the server is running:

```
SET GLOBAL innodb_file_per_table=1;
```

With file-per-table mode enabled, InnoDB stores each newly created table in its own `tbl_name.ibd` file in the appropriate database directory. Unlike the [MyISAM](#) storage engine, with its separate `tbl_name.MYD` and `tbl_name.MYI` files for indexes and data, InnoDB stores the data and the indexes together in a single `.ibd` file. The `tbl_name.frm` file is still created as usual.

If you remove `innodb_file_per_table` [1718] from your startup options and restart the server, or turn it off with the `SET GLOBAL` command, InnoDB creates any new tables inside the system tablespace.

You can always read and write any InnoDB tables, regardless of the file-per-table setting.

To move a table from the system tablespace to its own tablespace, or vice versa, change the `innodb_file_per_table` [1718] setting and rebuild the table:

```
-- Move table from system tablespace to its own tablespace.
SET GLOBAL innodb_file_per_table=1;
ALTER TABLE table_name ENGINE=InnoDB;
```

```
-- Move table from its own tablespace to system tablespace.
SET GLOBAL innodb_file_per_table=0;
ALTER TABLE table_name ENGINE=InnoDB;
```



### Note

InnoDB always needs the system tablespace because it puts its internal [data dictionary](#) and [undo logs](#) there. The `.ibd` files are not sufficient for InnoDB to operate.

When a table is moved out of the system tablespace into its own `.ibd` file, the data files that make up the system tablespace remain the same size. The space formerly occupied by the table can be reused for new InnoDB data, but is not reclaimed for use by the operating system. When moving large InnoDB tables out of the system tablespace, where disk space is limited, you might prefer to turn on `innodb_file_per_table` and then recreate the entire instance using the `mysqldump` command.

## 5.4.1.2. Specifying the Location of a Tablespace

To create a new InnoDB table in a specific location outside the MySQL data directory, use the `DATA DIRECTORY = absolute_path_to_directory` clause of the `CREATE TABLE` statement. (Plan the location in advance, because you cannot use this clause with the `ALTER TABLE` statement.) The directory you specify could be on another storage device with particular performance or capacity characteristics, such as a fast [SSD](#) or a high-capacity [HDD](#).

Within the destination directory, MySQL creates a subdirectory corresponding to the database name, and within that a `.ibd` file for the new table. In the database directory underneath the MySQL `DATADIR` [472] directory, MySQL creates a `table_name.isl` file containing the path name for the table. The `.isl` file is treated by MySQL like a symbolic link. ([Using actual symbolic links](#) has never been supported for InnoDB tables.)

The following example shows how you might run a small development or test instance of MySQL on a laptop with a primary hard drive that is 95% full, and place a new table `EXTERNAL` on a different storage device with more free space. The shell commands show the different paths to the `LOCAL` table in its default location under the `DATADIR` directory, and the `EXTERNAL` table in the location you specified:

```
mysql> \! df -k .
Filesystem      1024-blocks      Used Available Capacity  iused   ifree %used  Mounted on
/dev/disk0s2    244277768 231603532  12418236    95% 57964881 3104559   95%  /

mysql> use test;
Database changed
mysql> show variables like 'innodb_file_per_table';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| innodb_file_per_table | ON |
+-----+-----+
1 row in set (0.00 sec)

mysql> \! pwd
/usr/local/mysql
mysql> create table local (x int unsigned not null primary key);
Query OK, 0 rows affected (0.03 sec)

mysql> \! ls -l data/test/local.ibd
-rw-rw----  1 cirrus  staff  98304 Nov 13 15:24 data/test/local.ibd

mysql> create table external (x int unsigned not null primary key) data directory = '/volumes/external1';
Query OK, 0 rows affected (0.03 sec)
```

```
mysql> \! ls -l /volumes/external1/data/test/external.ibd
-rwxrwxrwx 1 cirrus staff 98304 Nov 13 15:34 /volumes/external1/data/test/external.ibd

mysql> select count(*) from local;
+-----+
| count(*) |
+-----+
|          0 |
+-----+
1 row in set (0.01 sec)

mysql> select count(*) from external;
+-----+
| count(*) |
+-----+
|          0 |
+-----+
1 row in set (0.01 sec)
```



#### Notes:

- MySQL initially holds the `.ibd` file open, preventing you from dismounting the device, but might eventually close the table if the server is busy. Be careful not to accidentally dismount the external device while MySQL is running, or to start MySQL while the device is disconnected. Attempting to access a table when the associated `.ibd` file is missing causes a serious error that requires a server restart.

The server restart might fail if the `.ibd` file is still not at the expected path. In this case, manually remove the `table_name.isl` file in the database directory, and after restarting do a `DROP TABLE` to delete the `.frm` file and remove the information about the table from the [data dictionary](#).

- Do not put MySQL tables on an NFS-mounted volume. NFS uses a message-passing protocol to write to files, which could cause data inconsistency if network messages are lost or received out of order.
- If you use an LVM snapshot, file copy, or other file-based mechanism to back up the `.ibd` file, always use the `FLUSH TABLES ... FOR EXPORT` statement first to make sure all changes that were buffered in memory are [flushed](#) to disk before the backup occurs.
- The `DATA DIRECTORY` clause is a supported alternative to [using symbolic links](#), which has always been problematic and was never supported for individual [InnoDB](#) tables.

### 5.4.1.3. Copying Tablespaces to Another Server (Transportable Tablespaces)

There are many reasons why you might copy an [InnoDB](#) table to a different database server:

- To run reports without putting extra load on a production server.
- To set up identical data for a table on a new [slave](#) server.
- To restore a backed-up version of a table after a problem or mistake.
- As a faster way of moving data around than importing the results of a `mysqldump` command. the data is available immediately, rather than having to be re-inserted and the indexes rebuilt.

Use the `FLUSH TABLES` statement with the `FOR EXPORT` clause. This puts the corresponding `.ibd` files into a consistent state so that they can be copied to another server. You copy a `.cfg` file along with the `.ibd` file; this extra file is used by `ALTER TABLE ... IMPORT TABLESPACE` to make any necessary adjustments within the tablespace file during the import process. See [Section 13.7.6.3, “FLUSH Syntax”](#) for the SQL semantics.

#### 5.4.1.4. Moving the Undo Log out of the System Tablespace

Although tablespace management typically involves files holding tables and indexes, you can also divide the [undo log](#) into separate [undo tablespace](#) files. This layout is different from the default configuration where the undo log is part of the [system tablespace](#). See [Section 14.2.4.2.4, “Separate Tablespaces for InnoDB Undo Logs”](#) for details.

#### 5.4.2. Grouping DML Operations with Transactions

By default, connection to the MySQL server begins with [autocommit](#) mode enabled, which automatically commits every SQL statement as you execute it. This mode of operation might be unfamiliar if you have experience with other database systems, where it is standard practice to issue a sequence of [DML](#) statements and commit them or roll them back all together.

To use multiple-statement [transactions](#), switch autocommit off with the SQL statement `SET autocommit = 0` and end each transaction with `COMMIT` or `ROLLBACK` as appropriate. To leave autocommit on, begin each transaction with `START TRANSACTION` and end it with `COMMIT` or `ROLLBACK`. The following example shows two transactions. The first is committed; the second is rolled back.

```
shell> mysql test

mysql> CREATE TABLE customer (a INT, b CHAR (20), INDEX (a));
Query OK, 0 rows affected (0.00 sec)
mysql> -- Do a transaction with autocommit turned on.
mysql> START TRANSACTION;
Query OK, 0 rows affected (0.00 sec)
mysql> INSERT INTO customer VALUES (10, 'Heikki');
Query OK, 1 row affected (0.00 sec)
mysql> COMMIT;
Query OK, 0 rows affected (0.00 sec)
mysql> -- Do another transaction with autocommit turned off.
mysql> SET autocommit=0;
Query OK, 0 rows affected (0.00 sec)
mysql> INSERT INTO customer VALUES (15, 'John');
Query OK, 1 row affected (0.00 sec)
mysql> INSERT INTO customer VALUES (20, 'Paul');
Query OK, 1 row affected (0.00 sec)
mysql> DELETE FROM customer WHERE b = 'Heikki';
Query OK, 1 row affected (0.00 sec)
mysql> -- Now we undo those last 2 inserts and the delete.
mysql> ROLLBACK;
Query OK, 0 rows affected (0.00 sec)
mysql> SELECT * FROM customer;
+-----+-----+
| a    | b      |
+-----+-----+
| 10  | Heikki |
+-----+-----+
1 row in set (0.00 sec)
mysql>
```

#### Transactions in Client-Side Languages

In APIs such as PHP, Perl DBI, JDBC, ODBC, or the standard C call interface of MySQL, you can send transaction control statements such as `COMMIT` to the MySQL server as strings just like any other SQL statements such as `SELECT` or `INSERT`. Some APIs also offer separate special transaction commit and rollback functions or methods.

#### 5.4.3. Converting Tables from MyISAM to InnoDB

If you have existing tables, and applications that use them, that you want to convert to [InnoDB](#) for better reliability and scalability, use the following guidelines and tips. This section assumes most such tables were originally [MyISAM](#), which was formerly the default.

## Reduce Memory Usage for MyISAM, Increase Memory Usage for InnoDB

As you transition away from MyISAM tables, lower the value of the `key_buffer_size` [492] configuration option to free memory no longer needed for caching results. Increase the value of the `innodb_buffer_pool_size` [1708] configuration option, which performs a similar role of allocating cache memory for InnoDB tables. The InnoDB buffer pool caches both table data and index data, so it does double duty in speeding up lookups for queries and keeping query results in memory for reuse.

- Allocate as much memory to this option as you can afford, often up to 80% of physical memory on the server.
- If the operating system runs short of memory for other processes and begins to swap, reduce the `innodb_buffer_pool_size` [1708] value. Swapping is such an expensive operation that it drastically reduces the benefit of the cache memory.
- If the `innodb_buffer_pool_size` [1708] value is several gigabytes or higher, consider increasing the values of `innodb_buffer_pool_instances` [1706]. Doing so helps on busy servers where many connections are reading data into the cache at the same time.
- On a busy server, run benchmarks with the Query Cache turned off. The InnoDB buffer pool provides similar benefits, so the Query Cache might be tying up memory unnecessarily.

## Watch Out for Too-Long Or Too-Short Transactions

Because MyISAM tables do not support [transactions](#), you might not have paid much attention to the `autocommit` [464] configuration option and the `COMMIT` and `ROLLBACK` statements. These keywords are important to allow multiple sessions to read and write InnoDB tables concurrently, providing substantial scalability benefits in write-heavy workloads.

While a transaction is open, the system keeps a snapshot of the data as seen at the beginning of the transaction, which can cause substantial overhead if the system inserts, updates, and deletes millions of rows while a stray transaction keeps running. Thus, take care to avoid transactions that run for too long:

- If you are using a `mysql` session for interactive experiments, always `COMMIT` (to finalize the changes) or `ROLLBACK` (to undo the changes) when finished. Close down interactive sessions rather than leaving them open for long periods, to avoid keeping transactions open for long periods by accident.
- Make sure that any error handlers in your application also `ROLLBACK` incomplete changes or `COMMIT` completed changes.
- `ROLLBACK` is a relatively expensive operation, because `INSERT`, `UPDATE`, and `DELETE` operations are written to InnoDB tables prior to the `COMMIT`, with the expectation that most changes will be committed successfully and rollbacks will be rare. When experimenting with large volumes of data, avoid making changes to large numbers of rows and then rolling back those changes.
- When loading large volumes of data with a sequence of `INSERT` statements, periodically `COMMIT` the results to avoid having transactions that last for hours. In typical load operations for data warehousing, if something goes wrong, you `TRUNCATE TABLE` and start over from the beginning rather than doing a `ROLLBACK`.

The preceding tips save memory and disk space that can be wasted during too-long transactions. When transactions are shorter than they should be, the problem is excessive I/O. With each `COMMIT`, MySQL makes sure each change is safely recorded to disk, which involves some I/O.

- For most operations on InnoDB tables, you should use the setting `autocommit=0` [464]. From an efficiency perspective, this avoids unnecessary I/O when you issue large numbers of consecutive `INSERT`, `UPDATE`, or `DELETE` statements. From a safety perspective, this allows you to issue a `ROLLBACK` statement to recover lost or garbled data if you make a mistake on the `mysql` command line, or in an exception handler in your application.

- The time when `autocommit=1` [464] is suitable for InnoDB tables is when running a sequence of queries for generating reports or analyzing statistics. In this situation, there is no I/O penalty related to `COMMIT` or `ROLLBACK`, and InnoDB can automatically optimize the read-only workload.
- If you make a series of related changes, finalize all those changes at once with a single `COMMIT` at the end. For example, if you insert related pieces of information into several tables, do a single `COMMIT` after making all the changes. Or if you run many consecutive `INSERT` statements, do a single `COMMIT` after all the data is loaded; if you are doing millions of `INSERT` statements, perhaps split up the huge transaction by issuing a `COMMIT` every ten thousand or hundred thousand records, so the transaction does not grow too large.
- Remember that even a `SELECT` statement opens a transaction, so after running some report or debugging queries in an interactive `mysql` session, either issue a `COMMIT` or close the `mysql` session.

## Don't Worry Too Much About Deadlocks

You might see warning messages referring to “deadlocks” in the MySQL error log, or the output of `SHOW ENGINE INNODB STATUS`. Despite the scary-sounding name, a **deadlock** is not a serious issue for InnoDB tables, and often does not require any corrective action. When two transactions start modifying multiple tables, accessing the tables in a different order, they can reach a state where each transaction is waiting for the other and neither can proceed. MySQL immediately detects this condition and cancels (rolls back) the “smaller” transaction, allowing the other to proceed.

Your applications do need error-handling logic to restart a transaction that is forcibly cancelled like this. When you re-issue the same SQL statements as before, the original timing issue no longer applies: either the other transaction has already finished and yours can proceed, or the other transaction is still in progress and your transaction waits until it finishes.

If deadlock warnings occur constantly, you might review the application code to reorder the SQL operations in a consistent way, or to shorten the transactions. You can test with the `innodb_print_all_deadlocks` [1741] option enabled to see all deadlock warnings in the MySQL error log, rather than only the last warning in the `SHOW ENGINE INNODB STATUS` output.

## Plan the Storage Layout

To get the best performance from InnoDB tables, you can adjust a number of parameters related to storage layout.

When you convert MyISAM tables that are large, frequently accessed, and hold vital data, investigate and consider the `innodb_file_per_table` [1718], `innodb_file_format` [1716], and `innodb_page_size` [1740] configuration options, and the `ROW_FORMAT` and `KEY_BLOCK_SIZE` clauses of the `CREATE TABLE` statement.

During your initial experiments, the most important setting is `innodb_file_per_table` [1718]. Enabling this option before creating new InnoDB tables ensures that the InnoDB system **tablespace** files do not allocate disk space permanently for all the InnoDB data. With `innodb_file_per_table` [1718] enabled, `DROP TABLE` and `TRUNCATE TABLE` free disk space as you would expect.

## Converting an Existing Table

To convert a non-InnoDB table to use InnoDB use `ALTER TABLE`:

```
ALTER TABLE table_name ENGINE=InnoDB;
```



### Important

Do not convert MySQL system tables in the `mysql` database (such as `user` or `host`) to the InnoDB type. This is an unsupported operation. The system tables must always be of the MyISAM type.



## Cloning the Structure of a Table

You might make an InnoDB table that is a clone of a MyISAM table, rather than doing the `ALTER TABLE` conversion, to test the old and new table side-by-side before switching.

Create an empty InnoDB table with identical column and index definitions. Use `show create table table_name\G` to see the full `CREATE TABLE` statement to use. Change the `ENGINE` clause to `ENGINE=INNODB`.

## Transferring Existing Data

To transfer a large volume of data into an empty InnoDB table created as shown in the previous section, insert the rows with `INSERT INTO innodb_table SELECT * FROM myisam_table ORDER BY primary_key_columns`.

You can also create the indexes for the InnoDB table after inserting the data. Historically, creating new secondary indexes was a slow operation for InnoDB, but now you can create the indexes after the data is loaded with relatively little overhead from the index creation step.

If you have `UNIQUE` constraints on secondary keys, you can speed up a table import by turning off the uniqueness checks temporarily during the import operation:

```
SET unique_checks=0;... import operation ...
SET unique_checks=1;
```

For big tables, this saves disk I/O because InnoDB can use its `insert buffer` to write secondary index records as a batch. Be certain that the data contains no duplicate keys. `unique_checks` [561] permits but does not require storage engines to ignore duplicate keys.

To get better control over the insertion process, you might insert big tables in pieces:

```
INSERT INTO newtable SELECT * FROM oldtable
WHERE yourkey > something AND yourkey <= somethingelse;
```

After all records have been inserted, you can rename the tables.

During the conversion of big tables, increase the size of the InnoDB buffer pool to reduce disk I/O, to a maximum of 80% of physical memory. You can also increase the sizes of the InnoDB log files.

## Storage Requirements

By this point, as already mentioned, you should already have the `innodb_file_per_table` [1718] option enabled, so that if you temporarily make several copies of your data in InnoDB tables, you can recover all that disk space by dropping unneeded tables afterward.

Whether you convert the MyISAM table directly or create a cloned InnoDB table, make sure that you have sufficient disk space to hold both the old and new tables during the process. InnoDB tables require more disk space than MyISAM tables. If an `ALTER TABLE` operation runs out of space, it starts a rollback, and that can take hours if it is disk-bound. For inserts, InnoDB uses the insert buffer to merge secondary index records to indexes in batches. That saves a lot of disk I/O. For rollback, no such mechanism is used, and the rollback can take 30 times longer than the insertion.

In the case of a runaway rollback, if you do not have valuable data in your database, it may be advisable to kill the database process rather than wait for millions of disk I/O operations to complete. For the complete procedure, see [Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”](#).

## Carefully Choose a PRIMARY KEY for Each Table

The `PRIMARY KEY` clause is a critical factor affecting the performance of MySQL queries and the space usage for tables and indexes. Perhaps you have phoned a financial institution where you are

asked for an account number. If you do not have the number, you are asked for a dozen different pieces of information to “uniquely identify” yourself. The primary key is like that unique account number that lets you get straight down to business when querying or modifying the information in a table. Every row in the table must have a primary key value, and no two rows can have the same primary key value.

Here are some guidelines for the primary key, followed by more detailed explanations.

- Declare a [PRIMARY KEY](#) for each table. Typically, it is the most important column that you refer to in [WHERE](#) clauses when looking up a single row.
- Declare the [PRIMARY KEY](#) clause in the original [CREATE TABLE](#) statement, rather than adding it later through an [ALTER TABLE](#) statement.
- Choose the column and its data type carefully. Prefer numeric columns over character or string ones.
- Consider using an auto-increment column if there is not another stable, unique, non-null, numeric column to use.
- An auto-increment column is also a good choice if there is any doubt whether the value of the primary key column could ever change. Changing the value of a primary key column is an expensive operation, possibly involving rearranging data within the table and within each secondary index.

Consider adding a [primary key](#) to any table that does not already have one. Use the smallest practical numeric type based on the maximum projected size of the table. This can make each row slightly more compact, which can yield substantial space savings for large tables. The space savings are multiplied if the table has any [secondary indexes](#), because the primary key value is repeated in each secondary index entry. In addition to reducing data size on disk, a small primary key also lets more data fit into the [buffer pool](#), speeding up all kinds of operations and improving concurrency.

If the table already has a primary key on some longer column, such as a [VARCHAR](#), consider adding a new unsigned [AUTO\\_INCREMENT](#) column and switching the primary key to that, even if that column is not referenced in queries. This design change can produce substantial space savings in the secondary indexes. You can designate the former primary key columns as [UNIQUE NOT NULL](#) to enforce the same constraints as the [PRIMARY KEY](#) clause, that is, to prevent duplicate or null values across all those columns.

If you spread related information across multiple tables, typically each table uses the same column for its primary key. For example, a personnel database might have several tables, each with a primary key of employee number. A sales database might have some tables with a primary key of customer number, and other tables with a primary key of order number. Because lookups using the primary key are very fast, you can construct efficient join queries for such tables.

If you leave the [PRIMARY KEY](#) clause out entirely, MySQL creates an invisible one for you. It is a 6-byte value that might be longer than you need, thus wasting space. Because it is hidden, you cannot refer to it in queries.

## Application Performance Considerations

The extra reliability and scalability features of [InnoDB](#) do require more disk storage than equivalent [MyISAM](#) tables. You might change the column and index definitions slightly, for better space utilization, reduced I/O and memory consumption when processing result sets, and better query optimization plans making efficient use of index lookups.

If you do set up a numeric ID column for the primary key, use that value to cross-reference with related values in any other tables, particularly for [join](#) queries. For example, rather than accepting a country name as input and doing queries searching for the same name, do one lookup to determine the country ID, then do other queries (or a single join query) to look up relevant information across several tables. Rather than storing a customer or catalog item number as a string of digits, potentially using up several bytes, convert it to a numeric ID for storing and querying. A 4-byte unsigned [INT](#) column can index over 4 billion items (with the US meaning of billion: 1000 million). For the ranges of the different integer



types, see [Section 11.2.1, “Integer Types \(Exact Value\) - INTEGER, INT, SMALLINT, TINYINT, MEDIUMINT, BIGINT”](#).

## Understand Files Associated with InnoDB Tables

InnoDB files require more care and planning than MyISAM files do:

- You must not delete the [ibdata files](#) that represent the [InnoDB system tablespace](#).
- Copying InnoDB tables from one server to another requires issuing the `FLUSH TABLES ... FOR EXPORT` statement first, and copying the `table_name.cfg` file along with the `table_name.ibd` file.

### 5.4.4. AUTO\_INCREMENT Handling in InnoDB

InnoDB provides an optimization that significantly improves scalability and performance of SQL statements that insert rows into tables with `AUTO_INCREMENT` columns. To use the `AUTO_INCREMENT` mechanism with an InnoDB table, an `AUTO_INCREMENT` column `ai_col` must be defined as part of an index such that it is possible to perform the equivalent of an indexed `SELECT MAX(ai_col)` lookup on the table to obtain the maximum column value. Typically, this is achieved by making the column the first column of some table index.

This section provides background information on the original (“traditional”) implementation of auto-increment locking in InnoDB, explains the configurable locking mechanism, documents the parameter for configuring the mechanism, and describes its behavior and interaction with replication.

#### 5.4.4.1. Traditional InnoDB Auto-Increment Locking

The original implementation of auto-increment handling in InnoDB uses the following strategy to prevent problems when using the binary log for statement-based replication or for certain recovery scenarios.

If you specify an `AUTO_INCREMENT` column for an InnoDB table, the table handle in the InnoDB data dictionary contains a special counter called the auto-increment counter that is used in assigning new values for the column. This counter is stored only in main memory, not on disk.

InnoDB uses the following algorithm to initialize the auto-increment counter for a table `t` that contains an `AUTO_INCREMENT` column named `ai_col`: After a server startup, for the first insert into a table `t`, InnoDB executes the equivalent of this statement:

```
SELECT MAX(ai_col) FROM t FOR UPDATE;
```

InnoDB increments the value retrieved by the statement and assigns it to the column and to the auto-increment counter for the table. By default, the value is incremented by one. This default can be overridden by the `auto_increment_increment` [1965] configuration setting.

If the table is empty, InnoDB uses the value 1. This default can be overridden by the `auto_increment_offset` [1968] configuration setting.

If a `SHOW TABLE STATUS` statement examines the table `t` before the auto-increment counter is initialized, InnoDB initializes but does not increment the value and stores it for use by later inserts. This initialization uses a normal exclusive-locking read on the table and the lock lasts to the end of the transaction.

InnoDB follows the same procedure for initializing the auto-increment counter for a freshly created table.

After the auto-increment counter has been initialized, if you do not explicitly specify a value for an `AUTO_INCREMENT` column, InnoDB increments the counter and assigns the new value to the column.

If you insert a row that explicitly specifies the column value, and the value is bigger than the current counter value, the counter is set to the specified column value.

If a user specifies `NULL` or `0` for the `AUTO_INCREMENT` column in an `INSERT`, InnoDB treats the row as if the value was not specified and generates a new value for it.

The behavior of the auto-increment mechanism is not defined if you assign a negative value to the column, or if the value becomes bigger than the maximum integer that can be stored in the specified integer type.

When accessing the auto-increment counter, InnoDB uses a special table-level `AUTO-INC` lock that it keeps to the end of the current SQL statement, not to the end of the transaction. The special lock release strategy was introduced to improve concurrency for inserts into a table containing an `AUTO_INCREMENT` column. Nevertheless, two transactions cannot have the `AUTO-INC` lock on the same table simultaneously, which can have a performance impact if the `AUTO-INC` lock is held for a long time. That might be the case for a statement such as `INSERT INTO t1 ... SELECT ... FROM t2` that inserts all rows from one table into another.

InnoDB uses the in-memory auto-increment counter as long as the server runs. When the server is stopped and restarted, InnoDB reinitializes the counter for each table for the first `INSERT` to the table, as described earlier.

A server restart also cancels the effect of the `AUTO_INCREMENT = N` table option in `CREATE TABLE` and `ALTER TABLE` statements, which you can use with InnoDB tables to set the initial counter value or alter the current counter value.

You may see gaps in the sequence of values assigned to the `AUTO_INCREMENT` column if you roll back transactions that have generated numbers using the counter.

#### 5.4.4.2. Configurable InnoDB Auto-Increment Locking

As described in the previous section, InnoDB uses a special lock called the table-level `AUTO-INC` lock for inserts into tables with `AUTO_INCREMENT` columns. This lock is normally held to the end of the statement (not to the end of the transaction), to ensure that auto-increment numbers are assigned in a predictable and repeatable order for a given sequence of `INSERT` statements.

In the case of statement-based replication, this means that when an SQL statement is replicated on a slave server, the same values are used for the auto-increment column as on the master server. The result of execution of multiple `INSERT` statements is deterministic, and the slave reproduces the same data as on the master. If auto-increment values generated by multiple `INSERT` statements were interleaved, the result of two concurrent `INSERT` statements would be nondeterministic, and could not reliably be propagated to a slave server using statement-based replication.

To make this clear, consider an example that uses this table:

```
CREATE TABLE t1 (
  c1 INT(11) NOT NULL AUTO_INCREMENT,
  c2 VARCHAR(10) DEFAULT NULL,
  PRIMARY KEY (c1)
) ENGINE=InnoDB;
```

Suppose that there are two transactions running, each inserting rows into a table with an `AUTO_INCREMENT` column. One transaction is using an `INSERT ... SELECT` statement that inserts 1000 rows, and another is using a simple `INSERT` statement that inserts one row:

```
Tx1: INSERT INTO t1 (c2) SELECT 1000 rows from another table ...
Tx2: INSERT INTO t1 (c2) VALUES ('xxx');
```

InnoDB cannot tell in advance how many rows will be retrieved from the `SELECT` in the `INSERT` statement in Tx1, and it assigns the auto-increment values one at a time as the statement proceeds. With a table-level lock, held to the end of the statement, only one `INSERT` statement referring to table `t1` can execute at a time, and the generation of auto-increment numbers by different statements is not interleaved. The auto-increment value generated by the Tx1 `INSERT ... SELECT` statement will be

consecutive, and the (single) auto-increment value used by the `INSERT` statement in Tx2 will either be smaller or larger than all those used for Tx1, depending on which statement executes first.

As long as the SQL statements execute in the same order when replayed from the binary log (when using statement-based replication, or in recovery scenarios), the results will be the same as they were when Tx1 and Tx2 first ran. Thus, table-level locks held until the end of a statement make `INSERT` statements using auto-increment safe for use with statement-based replication. However, those locks limit concurrency and scalability when multiple transactions are executing insert statements at the same time.

In the preceding example, if there were no table-level lock, the value of the auto-increment column used for the `INSERT` in Tx2 depends on precisely when the statement executes. If the `INSERT` of Tx2 executes while the `INSERT` of Tx1 is running (rather than before it starts or after it completes), the specific auto-increment values assigned by the two `INSERT` statements are nondeterministic, and may vary from run to run.

InnoDB can avoid using the table-level `AUTO-INC` lock for a class of `INSERT` statements where the number of rows is known in advance, and still preserve deterministic execution and safety for statement-based replication. Further, if you are not using the binary log to replay SQL statements as part of recovery or replication, you can entirely eliminate use of the table-level `AUTO-INC` lock for even greater concurrency and performance, at the cost of permitting gaps in auto-increment numbers assigned by a statement and potentially having the numbers assigned by concurrently executing statements interleaved.

For `INSERT` statements where the number of rows to be inserted is known at the beginning of processing the statement, InnoDB quickly allocates the required number of auto-increment values without taking any lock, but only if there is no concurrent session already holding the table-level `AUTO-INC` lock (because that other statement will be allocating auto-increment values one-by-one as it proceeds). More precisely, such an `INSERT` statement obtains auto-increment values under the control of a mutex (a light-weight lock) that is *not* held until the statement completes, but only for the duration of the allocation process.

This new locking scheme enables much greater scalability, but it does introduce some subtle differences in how auto-increment values are assigned compared to the original mechanism. To describe the way auto-increment works in InnoDB, the following discussion defines some terms, and explains how InnoDB behaves using different settings of the `innodb_autoinc_lock_mode` [1705] configuration parameter, which you can set at server startup. Additional considerations are described following the explanation of auto-increment locking behavior.

First, some definitions:

- “`INSERT-like`” statements

All statements that generate new rows in a table, including `INSERT`, `INSERT ... SELECT`, `REPLACE`, `REPLACE ... SELECT`, and `LOAD DATA`.

- “Simple inserts”

Statements for which the number of rows to be inserted can be determined in advance (when the statement is initially processed). This includes single-row and multiple-row `INSERT` and `REPLACE` statements that do not have a nested subquery, but not `INSERT ... ON DUPLICATE KEY UPDATE`.

- “Bulk inserts”

Statements for which the number of rows to be inserted (and the number of required auto-increment values) is not known in advance. This includes `INSERT ... SELECT`, `REPLACE ... SELECT`, and `LOAD DATA` statements, but not plain `INSERT`. InnoDB will assign new values for the `AUTO_INCREMENT` column one at a time as each row is processed.

- “Mixed-mode inserts”

These are “simple insert” statements that specify the auto-increment value for some (but not all) of the new rows. An example follows, where `c1` is an `AUTO_INCREMENT` column of table `t1`:

```
INSERT INTO t1 (c1,c2) VALUES (1,'a'), (NULL,'b'), (5,'c'), (NULL,'d');
```

Another type of “mixed-mode insert” is `INSERT ... ON DUPLICATE KEY UPDATE`, which in the worst case is in effect an `INSERT` followed by a `UPDATE`, where the allocated value for the `AUTO_INCREMENT` column may or may not be used during the update phase.

There are three possible settings for the `innodb_autoinc_lock_mode` [1705] parameter:

- `innodb_autoinc_lock_mode = 0` (“traditional” lock mode)

This lock mode provides the same behavior as before `innodb_autoinc_lock_mode` [1705] existed. For all “`INSERT`-like” statements, a special table-level `AUTO-INC` lock is obtained and held to the end of the statement. This assures that the auto-increment values assigned by any given statement are consecutive.

This lock mode is provided for:

- Backward compatibility.
- Performance testing.
- Working around issues with “mixed-mode inserts”, due to the possible differences in semantics described later.
- `innodb_autoinc_lock_mode = 1` (“consecutive” lock mode)

This is the default lock mode. In this mode, “bulk inserts” use the special `AUTO-INC` table-level lock and hold it until the end of the statement. This applies to all `INSERT ... SELECT`, `REPLACE ... SELECT`, and `LOAD DATA` statements. Only one statement holding the `AUTO-INC` lock can execute at a time.

With this lock mode, “simple inserts” (only) use a new locking model where a light-weight mutex is used during the allocation of auto-increment values, and no table-level `AUTO-INC` lock is used, unless an `AUTO-INC` lock is held by another transaction. If another transaction does hold an `AUTO-INC` lock, a “simple insert” waits for the `AUTO-INC` lock, as if it too were a “bulk insert”.

This lock mode ensures that, in the presence of `INSERT` statements where the number of rows is not known in advance (and where auto-increment numbers are assigned as the statement progresses), all auto-increment values assigned by any “`INSERT`-like” statement are consecutive, and operations are safe for statement-based replication.

Simply put, the important impact of this lock mode is significantly better scalability. This mode is safe for use with statement-based replication. Further, as with “traditional” lock mode, auto-increment numbers assigned by any given statement are *consecutive*. In this mode, there is *no change* in semantics compared to “traditional” mode for any statement that uses auto-increment, with one important exception.

The exception is for “mixed-mode inserts”, where the user provides explicit values for an `AUTO_INCREMENT` column for some, but not all, rows in a multiple-row “simple insert”. For such inserts, InnoDB will allocate more auto-increment values than the number of rows to be inserted. However, all values automatically assigned are consecutively generated (and thus higher than) the auto-increment value generated by the most recently executed previous statement. “Excess” numbers are lost.

- `innodb_autoinc_lock_mode = 2` (“interleaved” lock mode)

In this lock mode, no “`INSERT`-like” statements use the table-level `AUTO-INC` lock, and multiple statements can execute at the same time. This is the fastest and most scalable lock mode, but it is

*not safe* when using statement-based replication or recovery scenarios when SQL statements are replayed from the binary log.

In this lock mode, auto-increment values are guaranteed to be unique and monotonically increasing across all concurrently executing “INSERT-like” statements. However, because multiple statements can be generating numbers at the same time (that is, allocation of numbers is *interleaved* across statements), the values generated for the rows inserted by any given statement may not be consecutive.

If the only statements executing are “simple inserts” where the number of rows to be inserted is known ahead of time, there will be no gaps in the numbers generated for a single statement, except for “mixed-mode inserts”. However, when “bulk inserts” are executed, there may be gaps in the auto-increment values assigned by any given statement.

The auto-increment locking modes provided by `innodb_autoinc_lock_mode` [1705] have several usage implications:

- Using auto-increment with replication

If you are using statement-based replication, set `innodb_autoinc_lock_mode` [1705] to 0 or 1 and use the same value on the master and its slaves. Auto-increment values are not ensured to be the same on the slaves as on the master if you use `innodb_autoinc_lock_mode` [1705] = 2 (“interleaved”) or configurations where the master and slaves do not use the same lock mode.

If you are using row-based replication, all of the auto-increment lock modes are safe. Row-based replication is not sensitive to the order of execution of the SQL statements.

- “Lost” auto-increment values and sequence gaps

In all lock modes (0, 1, and 2), if a transaction that generated auto-increment values rolls back, those auto-increment values are “lost”. Once a value is generated for an auto-increment column, it cannot be rolled back, whether or not the “INSERT-like” statement is completed, and whether or not the containing transaction is rolled back. Such lost values are not reused. Thus, there may be gaps in the values stored in an `AUTO_INCREMENT` column of a table.

- Gaps in auto-increment values for “bulk inserts”

With `innodb_autoinc_lock_mode` [1705] set to 0 (“traditional”) or 1 (“consecutive”), the auto-increment values generated by any given statement will be consecutive, without gaps, because the table-level `AUTO-INC` lock is held until the end of the statement, and only one such statement can execute at a time.

With `innodb_autoinc_lock_mode` [1705] set to 2 (“interleaved”), there may be gaps in the auto-increment values generated by “bulk inserts,” but only if there are concurrently executing “INSERT-like” statements.

For lock modes 1 or 2, gaps may occur between successive statements because for bulk inserts the exact number of auto-increment values required by each statement may not be known and overestimation is possible.

- Auto-increment values assigned by “mixed-mode inserts”

Consider a “mixed-mode insert,” where a “simple insert” specifies the auto-increment value for some (but not all) resulting rows. Such a statement will behave differently in lock modes 0, 1, and 2. For example, assume `c1` is an `AUTO_INCREMENT` column of table `t1`, and that the most recent automatically generated sequence number is 100. Consider the following “mixed-mode insert” statement:

```
INSERT INTO t1 (c1,c2) VALUES (1,'a'), (NULL,'b'), (5,'c'), (NULL,'d');
```

With `innodb_autoinc_lock_mode` [1705] set to 0 (“traditional”), the four new rows will be:

c1	c2
1	a
101	b
5	c
102	d

The next available auto-increment value will be 103 because the auto-increment values are allocated one at a time, not all at once at the beginning of statement execution. This result is true whether or not there are concurrently executing “INSERT-like” statements (of any type).

With `innodb_autoinc_lock_mode` [1705] set to 1 (“consecutive”), the four new rows will also be:

c1	c2
1	a
101	b
5	c
102	d

However, in this case, the next available auto-increment value will be 105, not 103 because four auto-increment values are allocated at the time the statement is processed, but only two are used. This result is true whether or not there are concurrently executing “INSERT-like” statements (of any type).

With `innodb_autoinc_lock_mode` [1705] set to mode 2 (“interleaved”), the four new rows will be:

c1	c2
1	a
<i>x</i>	b
5	c
<i>y</i>	d

The values of *x* and *y* will be unique and larger than any previously generated rows. However, the specific values of *x* and *y* will depend on the number of auto-increment values generated by concurrently executing statements.

Finally, consider the following statement, issued when the most-recently generated sequence number was the value 4:

```
INSERT INTO t1 (c1,c2) VALUES (1,'a'), (NULL,'b'), (5,'c'), (NULL,'d');
```

With any `innodb_autoinc_lock_mode` [1705] setting, this statement will generate a duplicate-key error 23000 (Can't write; duplicate key in table) because 5 will be allocated for the row (NULL, 'b') and insertion of the row (5, 'c') will fail.

### 5.4.5. InnoDB and FOREIGN KEY Constraints

This section describes differences in the InnoDB storage engine's handling of foreign keys as compared with that of the MySQL Server.

#### Foreign Key Definitions

Foreign key definitions for InnoDB tables are subject to the following conditions:

- InnoDB permits a foreign key to reference any index column or group of columns. However, in the referenced table, there must be an index where the referenced columns are listed as the *first* columns in the same order.



- InnoDB does not currently support foreign keys for tables with user-defined partitioning. This means that no user-partitioned InnoDB table may contain foreign key references or columns referenced by foreign keys.
- InnoDB allows a foreign key constraint to reference a non-unique key. *This is an InnoDB extension to standard SQL.*

## Referential Actions

Referential actions for foreign keys of InnoDB tables are subject to the following conditions:

- While `SET DEFAULT` is allowed by the MySQL Server, it is rejected as invalid by InnoDB. `CREATE TABLE` and `ALTER TABLE` statements using this clause are not allowed for InnoDB tables.
- If there are several rows in the parent table that have the same referenced key value, InnoDB acts in foreign key checks as if the other parent rows with the same key value do not exist. For example, if you have defined a `RESTRICT` type constraint, and there is a child row with several parent rows, InnoDB does not permit the deletion of any of those parent rows.
- InnoDB performs cascading operations through a depth-first algorithm, based on records in the indexes corresponding to the foreign key constraints.
- If `ON UPDATE CASCADE` or `ON UPDATE SET NULL` recurses to update the *same table* it has previously updated during the cascade, it acts like `RESTRICT`. This means that you cannot use self-referential `ON UPDATE CASCADE` or `ON UPDATE SET NULL` operations. This is to prevent infinite loops resulting from cascaded updates. A self-referential `ON DELETE SET NULL`, on the other hand, is possible, as is a self-referential `ON DELETE CASCADE`. Cascading operations may not be nested more than 15 levels deep.
- Like MySQL in general, in an SQL statement that inserts, deletes, or updates many rows, InnoDB checks `UNIQUE` and `FOREIGN KEY` constraints row-by-row. When performing foreign key checks, InnoDB sets shared row-level locks on child or parent records it has to look at. InnoDB checks foreign key constraints immediately; the check is not deferred to transaction commit. According to the SQL standard, the default behavior should be deferred checking. That is, constraints are only checked after the *entire SQL statement* has been processed. Until InnoDB implements deferred constraint checking, some things will be impossible, such as deleting a record that refers to itself using a foreign key.

## Foreign Key Usage and Error Information

You can obtain general information about foreign keys and their usage from querying the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table, and more information more specific to InnoDB tables can be found in the `INNODB_SYS_FOREIGN` and `INNODB_SYS_FOREIGN_COLS` tables, also in the `INFORMATION_SCHEMA` database. See also [Section 13.1.14.2, “Using FOREIGN KEY Constraints”](#).

In addition to `SHOW ERRORS`, in the event of a foreign key error involving InnoDB tables (usually Error 150 in the MySQL Server), you can obtain a detailed explanation of the most recent InnoDB foreign key error by checking the output of `SHOW ENGINE INNODB STATUS`.

### 5.4.6. Working with InnoDB Compressed Tables

By using the SQL syntax and MySQL configuration options for `compression`, you can create tables where the data is stored in compressed form. Compression can help to improve both raw performance and scalability. The compression means less data is transferred between disk and memory, and takes up less space on disk and in memory. The benefits are amplified for tables with `secondary indexes`, because index data is compressed also. Compression can be especially important for `SSD` storage devices, because they tend to have lower capacity than `HDD` devices.

#### 5.4.6.1. Overview of Table Compression

Because processors and cache memories have increased in speed more than disk storage devices, many workloads are `disk-bound`. Data `compression` enables smaller database size, reduced I/O, and

improved throughput, at the small cost of increased CPU utilization. Compression is especially valuable for read-intensive applications, on systems with enough RAM to keep frequently used data in memory.

An InnoDB table created with `ROW_FORMAT=COMPRESSED` can use a smaller [page size](#) on disk than the usual 16KB default. Smaller pages require less I/O to read from and write to disk, which is especially valuable for [SSD](#) devices.

The page size is specified through the `KEY_BLOCK_SIZE` parameter. The different page size means the table must be in its own `.ibd` file rather than in the [system tablespace](#), which requires enabling the `innodb_file_per_table` option. The level of compression is the same regardless of the `KEY_BLOCK_SIZE` value. As you specify smaller values for `KEY_BLOCK_SIZE`, you get the I/O benefits of increasingly smaller pages. But if you specify a value that is too small, there is additional overhead to reorganize the pages when data values cannot be compressed enough to fit multiple rows in each page. There is a hard limit on how small `KEY_BLOCK_SIZE` can be for a table, based on the lengths of the key columns for each of its indexes. Specify a value that is too small, and the `CREATE TABLE` or `ALTER TABLE` statement fails.

In the buffer pool, the compressed data is held in small pages, with a page size based on the `KEY_BLOCK_SIZE` value. For extracting or updating the column values, MySQL also creates a 16KB page in the buffer pool with the uncompressed data. Within the buffer pool, any updates to the uncompressed page are also re-written back to the equivalent compressed page. You might need to size your buffer pool to accommodate the additional data of both compressed and uncompressed pages, although the uncompressed pages are [evicted](#) from the buffer pool when space is needed, and then uncompressed again on the next access.

#### 5.4.6.2. Enabling Compression for a Table

Before creating a compressed table, make sure the `innodb_file_per_table` [\[1718\]](#) configuration option is enabled, and `innodb_file_format` [\[1716\]](#) is set to `Barracuda`. You can set these parameters in the MySQL [configuration file](#) `my.cnf` or `my.ini`, or with the `SET` statement without shutting down the MySQL server.

To enable compression for a table, you use the clauses `ROW_FORMAT=COMPRESSED`, `KEY_BLOCK_SIZE`, or both in a `CREATE TABLE` or `ALTER TABLE` statement.

To create a compressed table, you might use statements like these:

```
SET GLOBAL innodb_file_per_table=1;
SET GLOBAL innodb_file_format=Barracuda;
CREATE TABLE t1
  (c1 INT PRIMARY KEY)
  ROW_FORMAT=COMPRESSED
  KEY_BLOCK_SIZE=8;
```

- If you specify `ROW_FORMAT=COMPRESSED`, you can omit `KEY_BLOCK_SIZE`; the default compressed page size of 8KB is used.
- If you specify `KEY_BLOCK_SIZE`, you can omit `ROW_FORMAT=COMPRESSED`; compression is enabled automatically.
- To determine the best value for `KEY_BLOCK_SIZE`, typically you create several copies of the same table with different values for this clause, then measure the size of the resulting `.ibd` files and see how well each performs with a realistic [workload](#).
- For additional performance-related configuration options, see [Section 5.4.6.3, “Tuning Compression for InnoDB Tables”](#).

The default uncompressed size of InnoDB data [pages](#) is 16KB. Depending on the combination of option values, MySQL uses a page size of 1KB, 2KB, 4KB, 8KB, or 16KB for the `.ibd` file of the table. The actual compression algorithm is not affected by the `KEY_BLOCK_SIZE` value; the value determines how large each compressed chunk is, which in turn affects how many rows can be packed into each compressed page.



Setting `KEY_BLOCK_SIZE=16` typically does not result in much compression, since the normal InnoDB page size is 16KB. This setting may still be useful for tables with many long `BLOB`, `VARCHAR` or `TEXT` columns, because such values often do compress well, and might therefore require fewer overflow pages as described in [Section 5.4.6.5, “How Compression Works for InnoDB Tables”](#).

All indexes of a table (including the [clustered index](#)) are compressed using the same page size, as specified in the `CREATE TABLE` or `ALTER TABLE` statement. Table attributes such as `ROW_FORMAT` and `KEY_BLOCK_SIZE` are not part of the `CREATE INDEX` syntax, and are ignored if they are specified (although you see them in the output of the `SHOW CREATE TABLE` statement).

## Restrictions on Compressed Tables

Because MySQL versions prior to 5.1 cannot process compressed tables, using compression requires specifying the configuration parameter `innodb_file_format=Barracuda` [\[1716\]](#), to avoid accidentally introducing compatibility issues.

Table compression is also not available for the InnoDB [system tablespace](#). The system tablespace (space 0, the `ibdata*` files) can contain user data, but it also contains internal system information, and therefore is never compressed. Thus, compression applies only to tables (and indexes) stored in their own tablespaces, that is, created with the `innodb_file_per_table` [\[1718\]](#) option enabled.

Compression applies to an entire table and all its associated indexes, not to individual rows, despite the clause name `ROW_FORMAT`.

### 5.4.6.3. Tuning Compression for InnoDB Tables

Most often, the internal optimizations described in [Section 5.4.6.5, “InnoDB Data Storage and Compression”](#) ensure that the system runs well with compressed data. However, because the efficiency of compression depends on the nature of your data, you can make decisions that affect the performance of compressed tables:

- Which tables to compress.
- What compressed page size to use.
- Whether to adjust the size of the buffer pool based on run-time performance characteristics, such as the amount of time the system spends compressing and uncompressing data. Whether the workload is more like a [data warehouse](#) (primarily queries) or an [OLTP](#) system (mix of queries and [DML](#)).
- If the system performs DML operations on compressed tables, and the way the data is distributed leads to expensive [compression failures](#) at runtime, you might adjust additional advanced configuration options.

Use the guidelines in this section to help make those architectural and configuration choices. When you are ready to conduct long-term testing and put compressed tables into production, see [Section 5.4.6.4, “Monitoring Compression at Runtime”](#) for ways to verify the effectiveness of those choices under real-world conditions.

## When to Use Compression

In general, compression works best on tables that include a reasonable number of character string columns and where the data is read far more often than it is written. Because there are no guaranteed ways to predict whether or not compression benefits a particular situation, always test with a specific [workload](#) and data set running on a representative configuration. Consider the following factors when deciding which tables to compress.

## Data Characteristics and Compression

A key determinant of the efficiency of compression in reducing the size of data files is the nature of the data itself. Recall that compression works by identifying repeated strings of bytes in a block of data. Completely randomized data is the worst case. Typical data often has repeated values, and so compresses effectively. Character strings often compress well, whether defined in `CHAR`, `VARCHAR`,

**TEXT** or **BLOB** columns. On the other hand, tables containing mostly binary data (integers or floating point numbers) or data that is previously compressed (for example JPEG or PNG images) may not generally compress well, significantly or at all.

You choose whether to turn on compression for each InnoDB table. A table and all of its indexes use the same (compressed) **page size**. It might be that the **primary key** (clustered) index, which contains the data for all columns of a table, compresses more effectively than the secondary indexes. For those cases where there are long rows, the use of compression might result in long column values being stored “off-page”, as discussed in [Section 5.4.8.3, “DYNAMIC and COMPRESSED Row Formats”](#). Those overflow pages may compress well. Given these considerations, for many applications, some tables compress more effectively than others, and you might find that your workload performs best only with a subset of tables compressed.

To determine whether or not to compress a particular table, conduct experiments. You can get a rough estimate of how efficiently your data can be compressed by using a utility that implements LZ77 compression (such as **gzip** or WinZip) on a copy of the **.ibd file** for an uncompressed table. You can expect less compression from a MySQL compressed table than from file-based compression tools, because MySQL compresses data in chunks based on the **page size**, 16KB by default. In addition to user data, the page format includes some internal system data that is not compressed. File-based compression utilities can examine much larger chunks of data, and so might find more repeated strings in a huge file than MySQL can find in an individual page.

Another way to test compression on a specific table is to copy some data from your uncompressed table to a similar, compressed table (having all the same indexes) and look at the size of the resulting **.ibd file**. For example:

```
use test;
set global innodb_file_per_table=1;
set global innodb_file_format=Barracuda;
set global autocommit=0;

-- Create an uncompressed table with a million or two rows.
create table big_table as select * from information_schema.columns;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
commit;
alter table big_table add id int unsigned not null primary key auto_increment;

show create table big_table\G

select count(id) from big_table;

-- Check how much space is needed for the uncompressed table.
\! ls -l data/test/big_table.ibd

create table key_block_size_4 like big_table;
alter table key_block_size_4 key_block_size=4 row_format=compressed;

insert into key_block_size_4 select * from big_table;
commit;

-- Check how much space is needed for a compressed table
-- with particular compression settings.
\! ls -l data/test/key_block_size_4.ibd
```

This experiment produced the following numbers, which of course could vary considerably depending on your table structure and data:

```
-rw-rw---- 1 cirrus staff 310378496 Jan 9 13:44 data/test/big_table.ibd
-rw-rw---- 1 cirrus staff 83886080 Jan 9 15:10 data/test/key_block_size_4.ibd
```

To see whether compression is efficient for your particular [workload](#):

- For simple tests, use a MySQL instance with no other compressed tables and run queries against the `INFORMATION_SCHEMA.INNODB_CMP` table.
- For more elaborate tests involving workloads with multiple compressed tables, run queries against the `INFORMATION_SCHEMA.INNODB_CMP_PER_INDEX` table. Because the statistics in the `INNODB_CMP_PER_INDEX` table are expensive to collect, you must enable the configuration option `innodb_cmp_per_index_enabled` [1712] before querying that table, and you might restrict such testing to a development server or a non-critical [slave server](#).
- Run some typical SQL statements against the compressed table you are testing.
- Examine the ratio of successful compression operations to overall compression operations by querying the `INFORMATION_SCHEMA.INNODB_CMP` or `INFORMATION_SCHEMA.INNODB_CMP_PER_INDEX` table, and comparing `COMPRESS_OPS` to `COMPRESS_OPS_OK`.
- If a high percentage of compression operations complete successfully, the table might be a good candidate for compression.
- If you get a high proportion of [compression failures](#), you can adjust `innodb_compression_level` [1713], `innodb_compression_failure_threshold_pct` [1713], and `innodb_compression_pad_pct_max` [1713] options as described in [Section 14.2.4.2.2](#), “[Compression Enhancements for OLTP Workloads](#)”, and try further tests.

## Database Compression versus Application Compression

Decide whether to compress data in your application or in the table; do not use both types of compression for the same data. When you compress the data in the application and store the results in a compressed table, extra space savings are extremely unlikely, and the double compression just wastes CPU cycles.

## Compressing in the Database

When enabled, MySQL table compression is automatic and applies to all columns and index values. The columns can still be tested with operators such as `LIKE`, and sort operations can still use indexes even when the index values are compressed. Because indexes are often a significant fraction of the total size of a database, compression could result in significant savings in storage, I/O or processor time. The compression and decompression operations happen on the database server, which likely is a powerful system that is sized to handle the expected load.

## Compressing in the Application

If you compress data such as text in your application, before it is inserted into the database, You might save overhead for data that does not compress well by compressing some columns and not others. This approach uses CPU cycles for compression and uncompression on the client machine rather than the database server, which might be appropriate for a distributed application with many clients, or where the client machine has spare CPU cycles.

## Hybrid Approach

Of course, it is possible to combine these approaches. For some applications, it may be appropriate to use some compressed tables and some uncompressed tables. It may be best to externally compress some data (and store it in uncompressed tables) and allow MySQL to compress (some of) the other tables in the application. As always, up-front design and real-life testing are valuable in reaching the right decision.

## Workload Characteristics and Compression

In addition to choosing which tables to compress (and the page size), the workload is another key determinant of performance. If the application is dominated by reads, rather than updates, fewer pages need to be reorganized and recompressed after the index page runs out of room for the per-page “modification log” that MySQL maintains for compressed data. If the updates predominantly change non-indexed columns or those containing `BLOBS` or large strings that happen to be stored “off-page”, the overhead of compression may be acceptable. If the only changes to a table are `INSERTs` that use a monotonically increasing primary key, and there are few secondary indexes, there is little need to reorganize and recompress index pages. Since MySQL can “delete-mark” and delete rows on compressed pages “in place” by modifying uncompressed data, `DELETE` operations on a table are relatively efficient.

For some environments, the time it takes to load data can be as important as run-time retrieval. Especially in data warehouse environments, many tables may be read-only or read-mostly. In those cases, it might or might not be acceptable to pay the price of compression in terms of increased load time, unless the resulting savings in fewer disk reads or in storage cost is significant.

Fundamentally, compression works best when the CPU time is available for compressing and uncompressing data. Thus, if your workload is I/O bound, rather than CPU-bound, you might find that compression can improve overall performance. When you test your application performance with different compression configurations, test on a platform similar to the planned configuration of the production system.

## Configuration Characteristics and Compression

Reading and writing database pages from and to disk is the slowest aspect of system performance. Compression attempts to reduce I/O by using CPU time to compress and uncompress data, and is most effective when I/O is a relatively scarce resource compared to processor cycles.

This is often especially the case when running in a multi-user environment with fast, multi-core CPUs. When a page of a compressed table is in memory, MySQL often uses additional memory, typically 16KB, in the buffer pool for an uncompressed copy of the page. The adaptive LRU algorithm attempts to balance the use of memory between compressed and uncompressed pages to take into account whether the workload is running in an I/O-bound or CPU-bound manner. Still, a configuration with more memory dedicated to the buffer pool tends to run better when using compressed tables than a configuration where memory is highly constrained.

## Choosing the Compressed Page Size

The optimal setting of the compressed page size depends on the type and distribution of data that the table and its indexes contain. The compressed page size should always be bigger than the maximum record size, or operations may fail as noted in [Section 5.4.6.5, “Compression of B-Tree Pages”](#).

Setting the compressed page size too large wastes some space, but the pages do not have to be compressed as often. If the compressed page size is set too small, inserts or updates may require time-consuming recompression, and the B-tree nodes may have to be split more frequently, leading to bigger data files and less efficient indexing.

Typically, you set the compressed page size to 8K or 4K bytes. Given that the maximum row size for an InnoDB table is around 8K, `KEY_BLOCK_SIZE=8` is usually a safe choice.

### 5.4.6.4. Monitoring Compression at Runtime

Overall application performance, CPU and I/O utilization and the size of disk files are good indicators of how effective compression is for your application. This section builds on the performance tuning advice from [Section 5.4.6.3, “Tuning Compression for InnoDB Tables”](#), and shows how to find problems that might not turn up during initial testing.

To dig deeper into performance considerations for compressed tables, you can monitor compression performance at runtime using the [Information Schema](#) tables described in [Example 14.2, “Using the Compression Information Schema Tables”](#). These tables reflect the internal use of memory and the rates of compression used overall.

The `INNODB_CMP` table reports information about compression activity for each compressed page size (`KEY_BLOCK_SIZE`) in use. The information in these tables is system-wide: it summarizes the compression statistics across all compressed tables in your database. You can use this data to help decide whether or not to compress a table by examining these tables when no other compressed tables are being accessed. It involves relatively low overhead on the server, so you might query it periodically on a production server to check the overall efficiency of the compression feature.

The `INNODB_CMP_PER_INDEX` table reports information about compression activity for individual tables and indexes. This information is more targeted and more useful for evaluating compression efficiency and diagnosing performance issues one table or index at a time. (Because that each InnoDB table is represented as a clustered index, MySQL does not make a big distinction between tables and indexes in this context.) The `INNODB_CMP_PER_INDEX` table does involve substantial overhead, so it is more suitable for development servers, where you can compare the effects of different workloads, data, and compression settings in isolation. To guard against imposing this monitoring overhead by accident, you must enable the `innodb_cmp_per_index_enabled` [1712] configuration option before you can query the `INNODB_CMP_PER_INDEX` table.

The key statistics to consider are the number of, and amount of time spent performing, compression and uncompression operations. Since MySQL splits B-tree nodes when they are too full to contain the compressed data following a modification, compare the number of “successful” compression operations with the number of such operations overall. Based on the information in the `INNODB_CMP` and `INNODB_CMP_PER_INDEX` tables and overall application performance and hardware resource utilization, you might make changes in your hardware configuration, adjust the size of the buffer pool, choose a different page size, or select a different set of tables to compress.

If the amount of CPU time required for compressing and uncompressing is high, changing to faster or multi-core CPUs can help improve performance with the same data, application workload and set of compressed tables. Increasing the size of the buffer pool might also help performance, so that more uncompressed pages can stay in memory, reducing the need to uncompress pages that exist in memory only in compressed form.

A large number of compression operations overall (compared to the number of `INSERT`, `UPDATE` and `DELETE` operations in your application and the size of the database) could indicate that some of your compressed tables are being updated too heavily for effective compression. If so, choose a larger page size, or be more selective about which tables you compress.

If the number of “successful” compression operations (`COMPRESS_OPS_OK`) is a high percentage of the total number of compression operations (`COMPRESS_OPS`), then the system is likely performing well. If the ratio is low, then MySQL is reorganizing, recompressing, and splitting B-tree nodes more often than is desirable. In this case, avoid compressing some tables, or increase `KEY_BLOCK_SIZE` for some of the compressed tables. You might turn off compression for tables that cause the number of “compression failures” in your application to be more than 1% or 2% of the total. (Such a failure ratio might be acceptable during a temporary operation such as a data load).

#### 5.4.6.5. How Compression Works for InnoDB Tables

This section describes some internal implementation details about MySQL [compression](#) for InnoDB tables. The information presented here may be helpful in tuning for performance, but is not necessary to know for basic use of compression.

#### Compression Algorithms

Some operating systems implement compression at the file system level. Files are typically divided into fixed-size blocks that are compressed into variable-size blocks, which easily leads into fragmentation. Every time something inside a block is modified, the whole block is recompressed before it is written to disk. These properties make this compression technique unsuitable for use in an update-intensive database system.

MySQL implements compression with the help of the well-known [zlib library](#), which implements the LZ77 compression algorithm. This compression algorithm is mature, robust, and efficient in both CPU



utilization and in reduction of data size. The algorithm is “lossless”, so that the original uncompressed data can always be reconstructed from the compressed form. LZ77 compression works by finding sequences of data that are repeated within the data to be compressed. The patterns of values in your data determine how well it compresses, but typical user data often compresses by 50% or more.

Unlike compression performed by an application, or compression features of some other database management systems, InnoDB compression applies both to user data and to indexes. In many cases, indexes can constitute 40-50% or more of the total database size, so this difference is significant. When compression is working well for a data set, the size of the InnoDB data files (the `.idb` files) is 25% to 50% of the uncompressed size or possibly smaller. Depending on the [workload](#), this smaller database can in turn lead to a reduction in I/O, and an increase in throughput, at a modest cost in terms of increased CPU utilization. You can adjust the balance between compression level and CPU overhead by modifying the `innodb_compression_level` [\[1713\]](#) configuration option.

## InnoDB Data Storage and Compression

All user data in InnoDB tables is stored in pages comprising a [B-tree](#) index (the [clustered index](#)). In some other database systems, this type of index is called an “index-organized table”. Each row in the index node contains the values of the (user-specified or system-generated) [primary key](#) and all the other columns of the table.

[Secondary indexes](#) in InnoDB tables are also B-trees, containing pairs of values: the index key and a pointer to a row in the clustered index. The pointer is in fact the value of the primary key of the table, which is used to access the clustered index if columns other than the index key and primary key are required. Secondary index records must always fit on a single B-tree page.

The compression of B-tree nodes (of both clustered and secondary indexes) is handled differently from compression of [overflow pages](#) used to store long `VARCHAR`, `BLOB`, or `TEXT` columns, as explained in the following sections.

## Compression of B-Tree Pages

Because they are frequently updated, B-tree pages require special treatment. It is important to minimize the number of times B-tree nodes are split, as well as to minimize the need to uncompress and recompress their content.

One technique MySQL uses is to maintain some system information in the B-tree node in uncompressed form, thus facilitating certain in-place updates. For example, this allows rows to be delete-marked and deleted without any compression operation.

In addition, MySQL attempts to avoid unnecessary uncompression and recompression of index pages when they are changed. Within each B-tree page, the system keeps an uncompressed “modification log” to record changes made to the page. Updates and inserts of small records may be written to this modification log without requiring the entire page to be completely reconstructed.

When the space for the modification log runs out, InnoDB uncompresses the page, applies the changes and recompresses the page. If recompression fails (a situation known as a [compression failure](#)), the B-tree nodes are split and the process is repeated until the update or insert succeeds.

To avoid frequent compression failures in write-intensive workloads, such as for [OLTP](#) applications, MySQL sometimes reserves some empty space (padding) in the page, so that the modification log fills up sooner and the page is recompressed while there is still enough room to avoid splitting it. The amount of padding space left in each page varies as the system keeps track of the frequency of page splits. On a busy server doing frequent writes to compressed tables, you can adjust the `innodb_compression_failure_threshold_pct` [\[1713\]](#), and `innodb_compression_pad_pct_max` [\[1713\]](#) configuration options to fine-tune this mechanism.

Generally, MySQL requires that each B-tree page in an InnoDB table can accommodate at least two records. For compressed tables, this requirement has been relaxed. Leaf pages of B-tree nodes (whether of the primary key or secondary indexes) only need to accommodate one record, but that record must fit, in uncompressed form, in the per-page modification log. If

`innodb_strict_mode` [1749] is ON, MySQL checks the maximum row size during `CREATE TABLE` or `CREATE INDEX`. If the row does not fit, the following error message is issued: `ERROR HY000: Too big row`.

If you create a table when `innodb_strict_mode` [1749] is OFF, and a subsequent `INSERT` or `UPDATE` statement attempts to create an index entry that does not fit in the size of the compressed page, the operation fails with `ERROR 42000: Row size too large`. (This error message does not name the index for which the record is too large, or mention the length of the index record or the maximum record size on that particular index page.) To solve this problem, rebuild the table with `ALTER TABLE` and select a larger compressed page size (`KEY_BLOCK_SIZE`), shorten any column prefix indexes, or disable compression entirely with `ROW_FORMAT=DYNAMIC` or `ROW_FORMAT=COMPACT`.

## Compressing BLOB, VARCHAR, and TEXT Columns

In an InnoDB table, `BLOB`, `VARCHAR`, and `TEXT` columns that are not part of the primary key may be stored on separately allocated [overflow pages](#). We refer to these columns as [off-page columns](#). Their values are stored on singly-linked lists of overflow pages.

For tables created in `ROW_FORMAT=DYNAMIC` or `ROW_FORMAT=COMPRESSED`, the values of `BLOB`, `TEXT`, or `VARCHAR` columns may be stored fully off-page, depending on their length and the length of the entire row. For columns that are stored off-page, the clustered index record only contains 20-byte pointers to the overflow pages, one per column. Whether any columns are stored off-page depends on the page size and the total size of the row. When the row is too long to fit entirely within the page of the clustered index, MySQL chooses the longest columns for off-page storage until the row fits on the clustered index page. As noted above, if a row does not fit by itself on a compressed page, an error occurs.

Tables created in older versions of MySQL use the [Antelope](#) file format, which supports only `ROW_FORMAT=REDUNDANT` and `ROW_FORMAT=COMPACT`. In these formats, MySQL stores the first 768 bytes of `BLOB`, `VARCHAR`, and `TEXT` columns in the clustered index record along with the primary key. The 768-byte prefix is followed by a 20-byte pointer to the overflow pages that contain the rest of the column value.

When a table is in `COMPRESSED` format, all data written to overflow pages is compressed “as is”; that is, MySQL applies the `zlib` compression algorithm to the entire data item. Other than the data, compressed overflow pages contain an uncompressed header and trailer comprising a page checksum and a link to the next overflow page, among other things. Therefore, very significant storage savings can be obtained for longer `BLOB`, `TEXT`, or `VARCHAR` columns if the data is highly compressible, as is often the case with text data. Image data, such as `JPEG`, is typically already compressed and so does not benefit much from being stored in a compressed table; the double compression can waste CPU cycles for little or no space savings.

The overflow pages are of the same size as other pages. A row containing ten columns stored off-page occupies ten overflow pages, even if the total length of the columns is only 8K bytes. In an uncompressed table, ten uncompressed overflow pages occupy 160K bytes. In a compressed table with an 8K page size, they occupy only 80K bytes. Thus, it is often more efficient to use compressed table format for tables with long column values.

Using a 16K compressed page size can reduce storage and I/O costs for `BLOB`, `VARCHAR`, or `TEXT` columns, because such data often compress well, and might therefore require fewer overflow pages, even though the B-tree nodes themselves take as many pages as in the uncompressed form.

## Compression and the InnoDB Buffer Pool

In a compressed InnoDB table, every compressed page (whether 1K, 2K, 4K or 8K) corresponds to an uncompressed page of 16K bytes (or a smaller size if `innodb_page_size` [1740] is set). To access the data in a page, MySQL reads the compressed page from disk if it is not already in the [buffer pool](#), then uncompresses the page to its original form. This section describes how InnoDB manages the buffer pool with respect to pages of compressed tables.

To minimize I/O and to reduce the need to uncompress a page, at times the buffer pool contains both the compressed and uncompressed form of a database page. To make room for other required database pages, MySQL can [evict](#) from the buffer pool an uncompressed page, while leaving the compressed page in memory. Or, if a page has not been accessed in a while, the compressed form of the page might be written to disk, to free space for other data. Thus, at any given time, the buffer pool might contain both the compressed and uncompressed forms of the page, or only the compressed form of the page, or neither.

MySQL keeps track of which pages to keep in memory and which to evict using a least-recently-used (LRU) list, so that [hot](#) (frequently accessed) data tends to stay in memory. When compressed tables are accessed, MySQL uses an adaptive LRU algorithm to achieve an appropriate balance of compressed and uncompressed pages in memory. This adaptive algorithm is sensitive to whether the system is running in an [I/O-bound](#) or [CPU-bound](#) manner. The goal is to avoid spending too much processing time uncompressing pages when the CPU is busy, and to avoid doing excess I/O when the CPU has spare cycles that can be used for uncompressing compressed pages (that may already be in memory). When the system is I/O-bound, the algorithm prefers to evict the uncompressed copy of a page rather than both copies, to make more room for other disk pages to become memory resident. When the system is CPU-bound, MySQL prefers to evict both the compressed and uncompressed page, so that more memory can be used for “hot” pages and reducing the need to uncompress data in memory only in compressed form.

## Compression and the InnoDB Redo Log Files

Before a compressed page is written to a [data file](#), MySQL writes a copy of the page to the redo log (if it has been recompressed since the last time it was written to the database). This is done to ensure that redo logs are usable for [crash recovery](#), even in the unlikely case that the [zlib](#) library is upgraded and that change introduces a compatibility problem with the compressed data. Therefore, some increase in the size of [log files](#), or a need for more frequent [checkpoints](#), can be expected when using compression. The amount of increase in the log file size or checkpoint frequency depends on the number of times compressed pages are modified in a way that requires reorganization and recompression.

Note that compressed tables use a different file format for the redo log and the per-table tablespaces than in MySQL 5.1 and earlier. The [MySQL Enterprise Backup](#) product supports this latest [Barracuda](#) file format for compressed InnoDB tables. The older InnoDB Hot Backup product can only back up tables using the file format [Antelope](#), and thus does not support compressed InnoDB tables.

### 5.4.6.6. SQL Compression Syntax Warnings and Errors

Specifying [ROW\\_FORMAT=COMPRESSED](#) or [KEY\\_BLOCK\\_SIZE](#) in [CREATE TABLE](#) or [ALTER TABLE](#) statements produces the following warnings if the Barracuda file format is not enabled. You can view them with the [SHOW WARNINGS](#) statement.

Level	Code	Message
Warning	1478	InnoDB: <a href="#">KEY_BLOCK_SIZE</a> requires <a href="#">innodb_file_per_table</a> .
Warning	1478	InnoDB: <a href="#">KEY_BLOCK_SIZE</a> requires <a href="#">innodb_file_format=1</a>
Warning	1478	InnoDB: ignoring <a href="#">KEY_BLOCK_SIZE=4</a> .
Warning	1478	InnoDB: <a href="#">ROW_FORMAT=COMPRESSED</a> requires <a href="#">innodb_file_per_table</a> .
Warning	1478	InnoDB: assuming <a href="#">ROW_FORMAT=COMPACT</a> .

Notes:

- By default, these messages are only warnings, not errors, and the table is created without compression, as if the options were not specified.
- When [innodb\\_strict\\_mode](#) [1749] is enabled, MySQL generates an error, not a warning, for these cases. The table is not created if the current configuration does not permit using compressed tables.



The “non-strict” behavior lets you import a `mysqldump` file into a database that does not support compressed tables, even if the source database contained compressed tables. In that case, MySQL creates the table in `ROW_FORMAT=COMPACT` instead of preventing the operation.

To import the dump file into a new database, and have the tables re-created as they exist in the original database, ensure the server has the proper settings for the configuration parameters `innodb_file_format` [1716] and `innodb_file_per_table` [1718].

The attribute `KEY_BLOCK_SIZE` is permitted only when `ROW_FORMAT` is specified as `COMPRESSED` or is omitted. Specifying a `KEY_BLOCK_SIZE` with any other `ROW_FORMAT` generates a warning that you can view with `SHOW WARNINGS`. However, the table is non-compressed; the specified `KEY_BLOCK_SIZE` is ignored).

Level	Code	Message
Warning	1478	InnoDB: ignoring KEY_BLOCK_SIZE=n unless ROW_FORMAT=COMPRESSED.

If you are running with `innodb_strict_mode` [1749] enabled, the combination of a `KEY_BLOCK_SIZE` with any `ROW_FORMAT` other than `COMPRESSED` generates an error, not a warning, and the table is not created.

Table 5.6, “Meaning of `CREATE TABLE` and `ALTER TABLE` options” summarizes how the various options on `CREATE TABLE` and `ALTER TABLE` are handled.

**Table 5.6. Meaning of `CREATE TABLE` and `ALTER TABLE` options**

Option	Usage	Description
<code>ROW_FORMAT=REDUNDANT</code>	Storage format used prior to MySQL 5.0.3	Less efficient than <code>ROW_FORMAT=COMPACT</code> ; for backward compatibility
<code>ROW_FORMAT=COMPACT</code>	Default storage format since MySQL 5.0.3	Stores a prefix of 768 bytes of long column values in the clustered index page, with the remaining bytes stored in an overflow page
<code>ROW_FORMAT=DYNAMIC</code>	Available only with <code>innodb_file_format=Barracuda</code>	Store values within the clustered index page if they fit; if not, stores only a 20-byte pointer to an overflow page (no prefix)
<code>ROW_FORMAT=COMPRESSED</code>	Available only with <code>innodb_file_format=Barracuda</code>	Compresses the table and indexes using zlib to default compressed page size of 8K bytes; implies <code>ROW_FORMAT=DYNAMIC</code>
<code>KEY_BLOCK_SIZE=n</code>	Available only with <code>innodb_file_format=Barracuda</code>	Specifies compressed page size of 1, 2, 4, 8 or 16K bytes; implies <code>ROW_FORMAT=DYNAMIC</code> and <code>ROW_FORMAT=COMPRESSED</code>

Table 5.7, “`CREATE/ALTER TABLE Warnings and Errors when InnoDB Strict Mode is OFF`” summarizes error conditions that occur with certain combinations of configuration parameters and options on the `CREATE TABLE` or `ALTER TABLE` statements, and how the options appear in the output of `SHOW TABLE STATUS`.

When `innodb_strict_mode` [1749] is `OFF`, MySQL creates or alters the table, but ignores certain settings as shown below. You can see the warning messages in the MySQL error log. When `innodb_strict_mode` [1749] is `ON`, these specified combinations of options generate errors, and the table is not created or altered. To see the full description of the error condition, issue the `SHOW ERRORS` statement: example:

```
mysql> CREATE TABLE x (id INT PRIMARY KEY, c INT)
-> ENGINE=INNODB KEY_BLOCK_SIZE=33333;
ERROR 1005 (HY000): Can't create table 'test.x' (errno: 1478)
```

```
mysql> SHOW ERRORS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Error | 1478 | InnoDB: invalid KEY_BLOCK_SIZE=33333. |
| Error | 1005 | Can't create table 'test.x' (errno: 1478) |
+-----+-----+-----+

2 rows in set (0.00 sec)
```

**Table 5.7. CREATE/ALTER TABLE Warnings and Errors when InnoDB Strict Mode is OFF**

Syntax	Warning or Error Condition	Resulting ROW_FORMAT, as shown in SHOW TABLE STATUS
ROW_FORMAT=REDUNDANT	None	REDUNDANT
ROW_FORMAT=COMPACT	None	COMPACT
ROW_FORMAT=COMPRESSED or ROW_FORMAT=DYNAMIC or KEY_BLOCK_SIZE is specified	Ignored unless both <code>innodb_file_format [1716]=Barracuda</code> and <code>innodb_file_per_table [1718]</code> are enabled	COMPACT
Invalid KEY_BLOCK_SIZE is specified (not 1, 2, 4, 8 or 16)	KEY_BLOCK_SIZE is ignored	the requested one, or COMPACT by default
ROW_FORMAT=COMPRESSED and valid KEY_BLOCK_SIZE are specified	None; KEY_BLOCK_SIZE specified is used, not the 8K default	COMPRESSED
KEY_BLOCK_SIZE is specified with REDUNDANT, COMPACT or DYNAMIC row format	KEY_BLOCK_SIZE is ignored	REDUNDANT, COMPACT or DYNAMIC
ROW_FORMAT is not one of REDUNDANT, COMPACT, DYNAMIC or COMPRESSED	Ignored if recognized by the MySQL parser. Otherwise, an error is issued.	COMPACT or N/A

When `innodb_strict_mode` is ON, MySQL rejects invalid `ROW_FORMAT` or `KEY_BLOCK_SIZE` parameters. For compatibility with earlier versions of MySQL, strict mode is not enabled by default; instead, MySQL issues warnings (not errors) for ignored invalid parameters.

Note that it is not possible to see the chosen `KEY_BLOCK_SIZE` using `SHOW TABLE STATUS`. The statement `SHOW CREATE TABLE` displays the `KEY_BLOCK_SIZE` (even if it was ignored when creating the table). The real compressed page size of the table cannot be displayed by MySQL.

### 5.4.7. InnoDB File-Format Management

As InnoDB evolves, new on-disk data structures are sometimes required to support new features. Features such as compressed tables (see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#)), and long variable-length columns stored off-page (see [Section 5.4.8, “How InnoDB Stores Variable-Length Columns”](#)) require data file formats that are not compatible with prior versions of InnoDB. These features both require use of the new [Barracuda](#) file format.



#### Note

All other new features are compatible with the original [Antelope](#) file format and do not require the [Barracuda](#) file format.

This section discusses how to specify the file format for new InnoDB tables, compatibility of different file formats between MySQL releases,

**Named File Formats.** InnoDB 1.1 has the idea of a named file format and a configuration parameter to enable the use of features that require use of that format. The new file format is the [Barracuda](#) format, and the original InnoDB file format is called [Antelope](#). Compressed tables and the new row format that stores long columns “off-page” require the use of the Barracuda file format or newer. Future versions of InnoDB may introduce a series of file formats, [identified with the names of animals](#), in ascending alphabetic order.

### 5.4.7.1. Enabling File Formats

The configuration parameter `innodb_file_format` [\[1716\]](#) controls whether such statements as `CREATE TABLE` and `ALTER TABLE` can be used to create tables that depend on support for the [Barracuda](#) file format.

Although Oracle recommends using the Barracuda format for new tables where practical, in MySQL 5.5 the default file format is still [Antelope](#), for maximum compatibility with replication configurations containing different MySQL releases.

The file format is a dynamic, global parameter that can be specified in the MySQL option file (`my.cnf` or `my.ini`) or changed with the `SET GLOBAL` command.

### 5.4.7.2. Verifying File Format Compatibility

InnoDB 1.1 incorporates several checks to guard against the possible crashes and data corruptions that might occur if you run an older release of the MySQL server on InnoDB data files using a newer file format. These checks take place when the server is started, and when you first access a table. This section describes these checks, how you can control them, and error and warning conditions that might arise.

## Backward Compatibility

Considerations of backward compatibility only apply when using a recent version of InnoDB (the InnoDB Plugin, or MySQL 5.5 and higher with InnoDB 1.1) alongside an older one (MySQL 5.1 or earlier, with the built-in InnoDB rather than the InnoDB Plugin). To minimize the chance of compatibility issues, you can standardize on the InnoDB Plugin for all your MySQL 5.1 and earlier database servers.

In general, a newer version of InnoDB may create a table or index that cannot safely be read or written with a prior version of InnoDB without risk of crashes, hangs, wrong results or corruptions. InnoDB 1.1 includes a mechanism to guard against these conditions, and to help preserve compatibility among database files and versions of InnoDB. This mechanism lets you take advantage of some new features of an InnoDB release (such as performance improvements and bug fixes), and still preserve the option of using your database with a prior version of InnoDB, by preventing accidental use of new features that create downward-incompatible disk files.

If a version of InnoDB supports a particular file format (whether or not that format is the default), you can query and update any table that requires that format or an earlier format. Only the creation of new tables using new features is limited based on the particular file format enabled. Conversely, if a tablespace contains a table or index that uses a file format that is not supported by the currently running software, it cannot be accessed at all, even for read access.

The only way to “downgrade” an InnoDB tablespace to an earlier file format is to copy the data to a new table, in a tablespace that uses the earlier format. This can be done with the `ALTER TABLE` statement, as described in [Section 5.4.7.4, “Downgrading the File Format”](#).

The easiest way to determine the file format of an existing InnoDB tablespace is to examine the properties of the table it contains, using the `SHOW TABLE STATUS` command or querying the table `INFORMATION_SCHEMA.TABLES`. If the `Row_format` of the table is reported as `'Compressed'` or `'Dynamic'`, the tablespace containing the table uses the Barracuda format. Otherwise, it uses the prior InnoDB file format, Antelope.

## Internal Details

Every InnoDB per-table tablespace (represented by a `*.ibd` file) file is labeled with a file format identifier. The system tablespace (represented by the `ibdata` files) is tagged with the “highest” file format in use in a group of InnoDB database files, and this tag is checked when the files are opened.

Creating a compressed table, or a table with `ROW_FORMAT=DYNAMIC`, updates the file header for the corresponding `.ibd` file and the table type in the InnoDB data dictionary with the identifier for the Barracuda file format. From that point forward, the table cannot be used with a version of InnoDB that does not support this new file format. To protect against anomalous behavior, InnoDB version 5.0.21 and later performs a compatibility check when the table is opened. (In many cases, the `ALTER TABLE` statement recreates a table and thus changes its properties. The special case of adding or dropping indexes without rebuilding the table is described in [Fast Index Creation in the InnoDB Storage Engine](#).)

### Definition of ib-file set

To avoid confusion, for the purposes of this discussion we define the term “ib-file set” to mean the set of operating system files that InnoDB manages as a unit. The ib-file set includes the following files:

- The system tablespace (one or more `ibdata` files) that contain internal system information (including internal catalogs and undo information) and may include user data and indexes.
- Zero or more single-table tablespaces (also called “file per table” files, named `*.ibd` files).
- InnoDB log files; usually two, `ib_logfile0` and `ib_logfile1`. Used for crash recovery and in backups.

An “ib-file set” does not include the corresponding `.frm` files that contain metadata about InnoDB tables. The `.frm` files are created and managed by MySQL, and can sometimes get out of sync with the internal metadata in InnoDB.

Multiple tables, even from more than one database, can be stored in a single “ib-file set”. (In MySQL, a “database” is a logical collection of tables, what other systems refer to as a “schema” or “catalog”.)

#### 5.4.7.2.1. Compatibility Check When InnoDB Is Started

To prevent possible crashes or data corruptions when InnoDB opens an ib-file set, it checks that it can fully support the file formats in use within the ib-file set. If the system is restarted following a crash, or a “fast shutdown” (i.e., `innodb_fast_shutdown` [1716] is greater than zero), there may be on-disk data structures (such as redo or undo entries, or doublewrite pages) that are in a “too-new” format for the current software. During the recovery process, serious damage can be done to your data files if these data structures are accessed. The startup check of the file format occurs before any recovery process begins, thereby preventing consistency issues with the new tables or startup problems for the MySQL server.

Beginning with version InnoDB 1.0.1, the system tablespace records an identifier or tag for the “highest” file format used by any table in any of the tablespaces that is part of the ib-file set. Checks against this file format tag are controlled by the configuration parameter `innodb_file_format_check` [1717], which is `ON` by default.

If the file format tag in the system tablespace is newer or higher than the highest version supported by the particular currently executing software and if `innodb_file_format_check` [1717] is `ON`, the following error is issued when the server is started:

```
InnoDB: Error: the system tablespace is in a
file format that this version doesn't support
```

You can also set `innodb_file_format` [1716] to a file format name. Doing so prevents InnoDB from starting if the current software does not support the file format specified. It also sets the “high water mark” to the value you specify. The ability to set `innodb_file_format_check` [1717] will be useful (with future releases of InnoDB) if you manually “downgrade” all of the tables in an ib-file set (as described in [Downgrading the InnoDB Storage Engine](#)). You can then rely on the file format check at startup if you subsequently use an older version of InnoDB to access the ib-file set.

In some limited circumstances, you might want to start the server and use an ib-file set that is in a “too new” format (one that is not supported by the software you are using). If you set the configuration parameter `innodb_file_format_check` [1717] to `OFF`, InnoDB opens the database, but issues this warning message in the error log:

```
InnoDB: Warning: the system tablespace is in a
file format that this version doesn't support
```



### Note

This is a very dangerous setting, as it permits the recovery process to run, possibly corrupting your database if the previous shutdown was a crash or “fast shutdown”. You should only set `innodb_file_format_check` [1717] to `OFF` if you are sure that the previous shutdown was done with `innodb_fast_shutdown=0`, so that essentially no recovery process occurs. In a future release, this parameter setting may be renamed from `OFF` to `UNSAFE`. (However, until there are newer releases of InnoDB that support additional file formats, even disabling the startup checking is in fact “safe”.)

The parameter `innodb_file_format_check` [1717] affects only what happens when a database is opened, not subsequently. Conversely, the parameter `innodb_file_format` [1716] (which enables a specific format) only determines whether or not a new table can be created in the enabled format and has no effect on whether or not a database can be opened.

The file format tag is a “high water mark”, and as such it is increased after the server is started, if a table in a “higher” format is created or an existing table is accessed for read or write (assuming its format is supported). If you access an existing table in a format higher than the format the running software supports, the system tablespace tag is not updated, but table-level compatibility checking applies (and an error is issued), as described in [Section 5.4.7.2.2, “Compatibility Check When a Table Is Opened”](#). Any time the high water mark is updated, the value of `innodb_file_format_check` [1717] is updated as well, so the command `SELECT @@innodb_file_format_check;` displays the name of the newest file format known to be used by tables in the currently open ib-file set and supported by the currently executing software.

To best illustrate this behavior, consider the scenario described in [Table 5.8, “InnoDB Data File Compatibility and Related InnoDB Parameters”](#). Imagine that some future version of InnoDB supports the Cheetah format and that an ib-file set has been used with that version.

**Table 5.8. InnoDB Data File Compatibility and Related InnoDB Parameters**

innodb file format check	innodb file format	Highest file format used in ib-file set	Highest file format supported by InnoDB	Result
OFF	Antelope or Barracuda	Barracuda	Barracuda	Database can be opened; tables can be created which require Antelope or Barracuda file format
OFF	Antelope or Barracuda	Cheetah	Barracuda	Database can be opened with a warning, since the database contains files in a “too new” format; tables can be created in Antelope or Barracuda file format; tables in Cheetah format cannot be accessed
OFF	Cheetah	Barracuda	Barracuda	Database cannot be opened; <code>innodb_file_format</code> [1716] cannot be set to Cheetah
ON	Antelope or Barracuda	Barracuda	Barracuda	Database can be opened; tables can be created in Antelope or Barracuda file format
ON	Antelope or Barracuda	Cheetah	Barracuda	Database cannot be opened, since the database contains files in a “too new” format (Cheetah)

innodb file format check	innodb file format	Highest file format used in ib-file set	Highest file format supported by InnoDB	Result
ON	Cheetah	Barracuda	Barracuda	Database cannot be opened; <code>innodb_file_format [1716]</code> cannot be set to Cheetah

#### 5.4.7.2.2. Compatibility Check When a Table Is Opened

When a table is first accessed, InnoDB (including some releases prior to InnoDB 1.0) checks that the file format of the tablespace in which the table is stored is fully supported. This check prevents crashes or corruptions that would otherwise occur when tables using a “too new” data structure are encountered.

All tables using any file format supported by a release can be read or written (assuming the user has sufficient privileges). The setting of the system configuration parameter `innodb_file_format [1716]` can prevent creating a new table that uses specific file formats, even if they are supported by a given release. Such a setting might be used to preserve backward compatibility, but it does not prevent accessing any table that uses any supported format.

As noted in [Named File Formats](#), versions of MySQL older than 5.0.21 cannot reliably use database files created by newer versions if a new file format was used when a table was created. To prevent various error conditions or corruptions, InnoDB checks file format compatibility when it opens a file (for example, upon first access to a table). If the currently running version of InnoDB does not support the file format identified by the table type in the InnoDB data dictionary, MySQL reports the following error:

```
ERROR 1146 (42S02): Table 'test.t1' doesn't exist
```

InnoDB also writes a message to the error log:

```
InnoDB: table test/t1: unknown table type 33
```

The table type should be equal to the tablespace flags, which contains the file format version as discussed in [Section 5.4.7.3, “Identifying the File Format in Use”](#).

Versions of InnoDB prior to MySQL 4.1 did not include table format identifiers in the database files, and versions prior to MySQL 5.0.21 did not include a table format compatibility check. Therefore, there is no way to ensure proper operations if a table in a “too new” format is used with versions of InnoDB prior to 5.0.21.

The file format management capability in InnoDB 1.0 and higher (tablespace tagging and run-time checks) allows InnoDB to verify as soon as possible that the running version of software can properly process the tables existing in the database.

If you permit InnoDB to open a database containing files in a format it does not support (by setting the parameter `innodb_file_format_check [1717]` to `OFF`), the table-level checking described in this section still applies.

Users are *strongly* urged not to use database files that contain Barracuda file format tables with releases of InnoDB older than the MySQL 5.1 with the InnoDB Plugin. It is possible to “downgrade” such tables to the Antelope format with the procedure described in [Section 5.4.7.4, “Downgrading the File Format”](#).

#### 5.4.7.3. Identifying the File Format in Use

After you enable a given `innodb_file_format [1716]`, this change applies only to newly created tables rather than existing ones. If you do create a new table, the tablespace containing the table is tagged with the “earliest” or “simplest” file format that is required for the table’s features. For example, if you enable file format Barracuda, and create a new table that is not compressed and does not use `ROW_FORMAT=DYNAMIC`, the new tablespace that contains the table is tagged as using file format Antelope.



It is easy to identify the file format used by a given tablespace or table. The table uses the Barracuda format if the `Row_format` reported by `SHOW CREATE TABLE` or `INFORMATION_SCHEMA.TABLES` is one of 'Compressed' or 'Dynamic'. (The `Row_format` is a separate column; ignore the contents of the `Create_options` column, which may contain the string `ROW_FORMAT`.) If the table in a tablespace uses neither of those features, the file uses the format supported by prior releases of InnoDB, now called file format Antelope. Then, the `Row_format` is one of 'Redundant' or 'Compact'.

## Internal Details

The file format identifier is written as part of the tablespace flags (a 32-bit number) in the `*.ibd` file in the 4 bytes starting at position 54 of the file, most significant byte first. (The first byte of the file is byte zero.) On some systems, you can display these bytes in hexadecimal with the command `od -t x1 -j 54 -N 4 tablename.ibd`. If all bytes are zero, the tablespace uses the Antelope file format (which is the format used by the standard InnoDB storage engine up to version 5.1). Otherwise, the least significant bit should be set in the tablespace flags, and the file format identifier is written in the bits 5 through 11. (Divide the tablespace flags by 32 and take the remainder after dividing the integer part of the result by 128.)

### 5.4.7.4. Downgrading the File Format

Each InnoDB tablespace file (with a name matching `*.ibd`) is tagged with the file format used to create its table and indexes. The way to downgrade the tablespace is to re-create the table and its indexes. The easiest way to recreate a table and its indexes is to use the command:

```
ALTER TABLE t ROW_FORMAT=COMPACT;
```

on each table that you want to downgrade. The `COMPACT` row format uses the file format Antelope. It was introduced in MySQL 5.0.3.

### 5.4.7.5. Future InnoDB File Formats

The file format used by the standard built-in InnoDB in MySQL 5.1 is the `Antelope` format. The file format introduced with InnoDB Plugin 1.0 is the `Barracuda` format. Although no new features have been announced that would require additional new file formats, the InnoDB file format mechanism allows for future enhancements.

For the sake of completeness, these are the file format names that might be used for future file formats: Antelope, Barracuda, Cheetah, Dragon, Elk, Fox, Gazelle, Hornet, Impala, Jaguar, Kangaroo, Leopard, Moose, Nautilus, Ocelot, Porpoise, Quail, Rabbit, Shark, Tiger, Urchin, Viper, Whale, Xenops, Yak and Zebra. These file formats correspond to the internal identifiers 0..25.

## 5.4.8. How InnoDB Stores Variable-Length Columns

This section discusses how certain InnoDB features, such as table `compression` and off-page storage of long columns, are controlled by the `ROW_FORMAT` clause of the `CREATE TABLE` statement. It discusses considerations for choosing the right row format and compatibility of row formats between MySQL releases.

### 5.4.8.1. Overview of InnoDB Row Storage

The storage for rows and associated columns affects performance for queries and DML operations. As more rows fit into a single disk `page`, queries and index lookups can work faster, less cache memory is required in the InnoDB buffer pool, and less I/O is required to write out updated values for the numeric and short string columns.

The data in each InnoDB table is divided into `pages`. The pages that make up each table are arranged in a tree data structure called a `B-tree index`. Table data and `secondary indexes` both use this type of structure. The B-tree index that represents an entire table is known as the `clustered index`, which is organized according to the `primary key` columns. The nodes of the index data structure contain the values of all the columns in that row (for the clustered index) or the index columns and the primary key columns (for secondary indexes).

Variable-length columns are an exception to this rule. Columns such as `BLOB` and `VARCHAR` that are too long to fit on a B-tree page are stored on separately allocated disk pages called [overflow pages](#). We call such columns [off-page columns](#). The values of these columns are stored in singly-linked lists of overflow pages, and each such column has its own list of one or more overflow pages. In some cases, all or a prefix of the long column value is stored in the B-tree, to avoid wasting storage and eliminating the need to read a separate page.

This section describes the clauses you can use with the `CREATE TABLE` and `ALTER TABLE` statements to control how these variable-length columns are represented: `ROW_FORMAT` and `KEY_BLOCK_SIZE`. To use these clauses, you might also need to change the settings for the `innodb_file_per_table` [1718] and `innodb_file_format` [1716] configuration options.

#### 5.4.8.2. Specifying the Row Format for a Table

You specify the row format for a table with the `ROW_FORMAT` clause of the `CREATE TABLE` and `ALTER TABLE` statements.

#### 5.4.8.3. DYNAMIC and COMPRESSED Row Formats

This section discusses the `DYNAMIC` and `COMPRESSED` row formats for InnoDB tables. You can only create these kinds of tables when the `innodb_file_format` [1716] configuration option is set to `Barracuda`. (The Barracuda file format also allows the `COMPACT` and `REDUNDANT` row formats.)

When a table is created with `ROW_FORMAT=DYNAMIC` or `ROW_FORMAT=COMPRESSED`, long column values are stored fully off-page, and the clustered index record contains only a 20-byte pointer to the overflow page.

Whether any columns are stored off-page depends on the page size and the total size of the row. When the row is too long, InnoDB chooses the longest columns for off-page storage until the clustered index record fits on the B-tree page.

The `DYNAMIC` row format maintains the efficiency of storing the entire row in the index node if it fits (as do the `COMPACT` and `REDUNDANT` formats), but this new format avoids the problem of filling B-tree nodes with a large number of data bytes of long columns. The `DYNAMIC` format is based on the idea that if a portion of a long data value is stored off-page, it is usually most efficient to store all of the value off-page. With `DYNAMIC` format, shorter columns are likely to remain in the B-tree node, minimizing the number of overflow pages needed for any given row.

The `COMPRESSED` row format uses similar internal details for off-page storage as the `DYNAMIC` row format, with additional storage and performance considerations from the table and index data being compressed and using smaller page sizes. With the `COMPRESSED` row format, the option `KEY_BLOCK_SIZE` controls how much column data is stored in the clustered index, and how much is placed on overflow pages. For full details about the `COMPRESSED` row format, see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#).

#### 5.4.8.4. COMPACT and REDUNDANT Row Formats

Early versions of InnoDB used an unnamed file format (now called `Antelope`) for database files. With that file format, tables are defined with `ROW_FORMAT=COMPACT` or `ROW_FORMAT=REDUNDANT`. InnoDB stores up to the first 768 bytes of variable-length columns (such as `BLOB` and `VARCHAR`) in the index record within the B-tree node, with the remainder stored on the overflow pages.

To preserve compatibility with those prior versions, tables created with the newest InnoDB default to the `COMPACT` row format. See [Section 5.4.8.3, “DYNAMIC and COMPRESSED Row Formats”](#) for information about the newer `DYNAMIC` and `COMPRESSED` row formats.

With the Antelope file format, if the value of a column is 768 bytes or less, no overflow page is needed, and some savings in I/O may result, since the value is in the B-tree node. This works well for relatively short `BLOBS`, but may cause B-tree nodes to fill with data rather than key values, reducing their efficiency. Tables with many `BLOB` columns could cause B-tree nodes to become too full of data,



and contain too few rows, making the entire index less efficient than if the rows were shorter or if the column values were stored off-page.

## 5.5. Online DDL for InnoDB Tables

You can perform several kinds of [online DDL](#) operations on [InnoDB](#) tables: that is, allowing [DML](#) operations and [queries](#) on the table while the [DDL](#) is in progress, performing the operation “in-place” without rebuilding the entire table, or both. This enhancement has the following benefits:

- It improves responsiveness and availability in busy production environments, where making a table unavailable for minutes or hours whenever you modify its indexes or column definitions is not practical.
- It lets you adjust the balance between performance and concurrency during the DDL operation, by choosing whether to block access to the table entirely ([LOCK=EXCLUSIVE](#) clause), allow queries but not DML ([LOCK=SHARED](#) clause), or allow full query and DML access to the table ([LOCK=NONE](#) clause). When you omit the [LOCK](#) clause or specify [LOCK=DEFAULT](#), MySQL allows as much concurrency as possible depending on the type of operation.
- By doing the changes in-place where possible, rather than creating a new copy of the table, it avoids temporary increases in disk space usage and the I/O overhead of copying the table and reconstructing all the secondary indexes.

### 5.5.1. Overview of Online DDL

Historically, many [DDL](#) operations on [InnoDB](#) tables were expensive. Many [ALTER TABLE](#) operations worked by creating a new, empty table defined with the requested table options and indexes, then copying the existing rows to the new table one-by-one, updating the indexes as the rows were inserted. After all rows from the original table were copied, the old table was dropped and the copy was renamed with the name of the original table.

MySQL 5.5, and MySQL 5.1 with the InnoDB Plugin, optimized [CREATE INDEX](#) and [DROP INDEX](#) to avoid the table-copying behavior. That feature was known as Fast Index Creation. MySQL 5.6 enhances many other types of [ALTER TABLE](#) operations to avoid copying the table. Another enhancement allows [SELECT](#) queries and [INSERT](#), [UPDATE](#), and [DELETE](#) ([DML](#)) statements to proceed while the table is being altered. This combination of features is now known as [online DDL](#).

This new mechanism also means that you can generally speed the overall process of creating and loading a table and associated indexes by creating the table with without any secondary indexes, then adding the secondary indexes after the data is loaded.

Although no syntax changes are required in the [CREATE INDEX](#) or [DROP INDEX](#) commands, some factors affect the performance, space usage, and semantics of this operation (see [Section 5.5.9](#), “[Limitations of Online DDL](#)”).

The online DDL enhancements in MySQL 5.6 improve many DDL operations that formerly required a table copy, blocked DML operations on the table, or both. [Table 5.9](#), “[Summary of Online Status for DDL Operations](#)” shows the variations of the [ALTER TABLE](#) statement and shows how the online DDL feature applies to each one.

- The “In-Place?” column shows which operations allow the [ALGORITHM=INPLACE](#) clause; the preferred value is “Yes”.
- The “Copies Table?” column shows which operations are able to avoid the expensive table-copying operation; the preferred value is “No”. This column is mostly the reverse of the “In-Place?” column, except that a few operations allow [ALGORITHM=INPLACE](#) but still involve some amount of table copying.
- The “Allows Concurrent DML?” column shows which operations can be performed fully online; the preferred value is “Yes”. You can specify [LOCK=NONE](#) to assert that full concurrency is allowed

during the DDL, but MySQL automatically allows this level of concurrency when possible. When concurrent DML is allowed, concurrent queries are also always allowed.

- The “Allows Concurrent Queries?” column shows which DDL operations allow queries on the table while the operation is in progress; the preferred value is “Yes”. Concurrent query is allowed during all online DDL operations. It is shown with “Yes” listed for all cells, for reference purposes. You can specify `LOCK=SHARED` to assert that concurrent queries are allowed during the DDL, but MySQL automatically allows this level of concurrency when possible.
- The “Notes” column explains any exceptions to the “yes/no” values of the other columns, such as when the answer depends on the setting of a configuration option or some other clause in the DDL statement. The values “Yes\*” and “No\*” indicate that an answer depends on these additional notes.

**Table 5.9. Summary of Online Status for DDL Operations**

Operation	In-Place?	Copies Table?	Allows Concurrent DML?	Allows Concurrent Query?	Notes
<code>CREATE INDEX, ADD INDEX</code>	Yes*	No*	Yes	Yes	Some restrictions for <code>FULLTEXT</code> index; see next row. Currently, the operation is not in-place (that is, it copies the table) if the same index being created was also dropped by an earlier clause in the same <code>ALTER TABLE</code> statement.
<code>ADD FULLTEXT INDEX</code>	Yes	No*	No	Yes	Creating the first <code>FULLTEXT</code> index for a table involves a table copy, unless there is a user-supplied <code>FTS_DOC_ID</code> column. Subsequent <code>FULLTEXT</code> indexes on the same table can be created in-place.
<code>DROP INDEX</code>	Yes	No	Yes	Yes	
Set default value for a column	Yes	No	Yes	Yes	Modifies <code>.frm</code> file only, not the data file.
Change <code>auto-increment</code> value for a column	Yes	No	Yes	Yes	Modifies a value stored in memory, not the data file.
Add a <code>foreign key constraint</code>	Yes*	No*	Yes	Yes	To avoid copying the table, disable <code>foreign_key_checks</code> [483] during constraint creation.
Drop a <code>foreign key constraint</code>	Yes	No	Yes	Yes	The <code>foreign_key_checks</code> [483] option can be enabled or disabled.
Rename a column	Yes*	No*	Yes*	Yes	To allow concurrent DML, keep the same data type and only change the column name.
Add a column	Yes	Yes	Yes*	Yes	Concurrent DML is not allowed when adding an <code>auto-increment</code> column. Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.

Operation	In-Place?	Copies Table?	Allows Concurrent DML?	Allows Concurrent Query?	Notes
Drop a column	Yes	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Reorder columns	Yes	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Change <code>ROW_FORMAT</code> property	Yes	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Change <code>KEY_BLOCK_SIZE</code> property	Yes	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Make column <code>NULL</code>	Yes	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Make column <code>NOT NULL</code>	Yes*	Yes	Yes	Yes	When <code>SQL_MODE [545]</code> includes <code>strict_all_tables</code> or <code>strict_all_tables</code> , the operation fails if the column contains any nulls. Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation.
Change data type of column	No	Yes	No	Yes	
Add <code>primary key</code>	Yes*	Yes	Yes	Yes	Although <code>ALGORITHM=INPLACE</code> is allowed, the data is reorganized substantially, so it is still an expensive operation. <code>ALGORITHM=INPLACE</code> is not allowed under certain conditions if columns have to be converted to <code>NOT NULL</code> . See <a href="#">Example 5.9, "Creating and Dropping the Primary Key"</a> .
Drop <code>primary key</code> and add another	Yes	Yes	Yes	Yes	<code>ALGORITHM=INPLACE</code> is only allowed when you add a new primary key in the same <code>ALTER TABLE</code> ; the data is reorganized substantially, so it is still an expensive operation.
Drop <code>primary key</code>	No	Yes	No	Yes	Restrictions apply when you drop a primary key without adding a new one in the same <code>ALTER TABLE</code> statement.

Operation	In-Place?	Copies Table?	Allows Concurrent DML?	Allows Concurrent Query?	Notes
Convert character set	No	Yes	No	Yes	Rebuilds the table if the new character encoding is different.
Specify character set	No	Yes	No	Yes	Rebuilds the table if the new character encoding is different.
Rebuild with <code>FORCE</code> option	No	Yes	No	Yes	Acts like the <code>ALGORITHM=COPY</code> clause or the setting <code>old_alter_table=1</code> [517].

The following sections shows the basic syntax, and usage notes related to online DDL, for each of the major operations that can be performed with concurrent DML, in-place, or both:

## Secondary Indexes

- Create [secondary indexes](#): `CREATE INDEX name ON table (col_list)` or `ALTER TABLE table ADD INDEX name (col_list)`. (Creating a `FULLTEXT` index still requires locking the table.)
- Drop [secondary indexes](#): `DROP INDEX name ON table;` or `ALTER TABLE table DROP INDEX name`

Creating and dropping secondary indexes on [InnoDB](#) tables skips the table-copying behavior, the same as in MySQL 5.5 and MySQL 5.1 with the [InnoDB Plugin](#).

In MySQL 5.6 and higher, the table remains available for read and write operations while the index is being created or dropped. The `CREATE INDEX` or `DROP INDEX` statement only finishes after all transactions that are accessing the table are completed, so that the initial state of the index reflects the most recent contents of the table. Previously, modifying the table while an index was being created or dropped typically resulted in a [deadlock](#) that cancelled the `INSERT`, `UPDATE`, or `DELETE` statement on the table.

## Column Properties

- Set a default value for a column: `ALTER TABLE tbl ALTER COLUMN col SET DEFAULT literal` or `ALTER TABLE tbl ALTER COLUMN col DROP DEFAULT`

The default values for columns are stored in the `.frm` file for the table, not the [InnoDB data dictionary](#).

- Changing the [auto-increment](#) value for a column: `ALTER TABLE table AUTO_INCREMENT=next_value;`

Especially in a distributed system using replication or sharding, you sometimes reset the auto-increment counter for a table to a specific value. The next row inserted into the table uses the specified value for its auto-increment column. You might also use this technique in a data warehousing environment where you periodically empty all the tables and reload them, and you can restart the auto-increment sequence from 1.

- Renaming a column: `ALTER TABLE tbl CHANGE old_col_name new_col_name datatype`

When you keep the same data type and `[NOT] NULL` attribute, only changing the column name, this operation can always be performed online.

As part of this enhancement, you can now rename a column that is part of a foreign key constraint, which was not allowed before. The foreign key definition is automatically updated to use the new column name. Renaming a column participating in a foreign key only works with the in-place mode

of `ALTER TABLE`. If you use the `ALGORITHM=COPY` clause, or some other condition causes the command to use `ALGORITHM=COPY` behind the scenes, the `ALTER TABLE` statement will fail.

## Foreign Keys

- Adding or dropping a [foreign key constraint](#):

```
ALTER TABLE tbl1 ADD CONSTRAINT fk_name FOREIGN KEY index (col1) REFERENCES tbl2(col2) referential_action
ALTER TABLE tbl DROP FOREIGN KEY fk_name;
```

Dropping a foreign key can be performed online with the `foreign_key_checks [483]` option enabled or disabled. Creating a foreign key online requires `foreign_key_checks [483]` to be disabled.

If you do not know the names of the foreign key constraints on a particular table, issue the following statement and find the constraint name in the `CONSTRAINT` clause for each foreign key:

```
show create table table\G
```

Or, query the `information_schema.table_constraints` table and use the `constraint_name` and `constraint_type` columns to identify the foreign key names.

As a consequence of this enhancement, you can now also drop a foreign key and its associated index in a single statement, which previously required separate statements in a strict order:

```
ALTER TABLE table DROP FOREIGN KEY constraint, DROP INDEX index;
```

If [foreign keys](#) are already present in the table being altered (that is, it is a [child table](#) containing any `FOREIGN KEY ... REFERENCE` clauses), additional restrictions apply to online DDL operations, even those not directly involving the foreign key columns:

- Concurrent DML is disallowed during online DDL operations on such child tables. (This restriction is being evaluated as a bug and might be lifted.)
- An `ALTER TABLE` on the child table could also wait for another transaction to commit, if a change to the parent table caused associated changes in the child table through an `ON UPDATE` or `ON DELETE` clause using the `CASCADE` or `SET NULL` parameters.

In the same way, if a table is the [parent table](#) in a foreign key relationship, even though it does not contain any `FOREIGN KEY` clauses, it could wait for the `ALTER TABLE` to complete if an `INSERT`, `UPDATE`, or `DELETE` statement caused an `ON UPDATE` or `ON DELETE` action in the child table.

## Notes on `ALGORITHM=COPY`

Any `ALTER TABLE` operation run with the `ALGORITHM=COPY` clause prevents concurrent DML operations. Concurrent queries are still allowed. That is, a table-copying operation always includes at least the concurrency restrictions of `LOCK=SHARED` (allow queries but not DML). You can further restrict concurrency for such operations by specifying `LOCK=EXCLUSIVE` (prevent DML and queries).

## Concurrent DML but Table Copy Still Required

Some other `ALTER TABLE` operations allow concurrent DML, and are faster than MySQL 5.5 and prior: the table-copying operation is optimized, even though a table copy is still required:

- Adding, dropping, or reordering columns.
- Adding or dropping a [primary key](#).
- Changing the `ROW_FORMAT` or `KEY_BLOCK_SIZE` properties for a table.
- Changing the nullable status for a column.



### Note

As your database schema evolves with new columns, data types, constraints, indexes, and so on, keep your `CREATE TABLE` statements up to date with the latest table definitions. Even with the performance improvements of online DDL, it is more efficient to create stable database structures at the beginning, rather than creating part of the schema and then issuing `ALTER TABLE` statements afterward.

The main exception to this guideline is for [secondary indexes](#) on tables with large numbers of rows. It is typically most efficient to create the table with all details specified except the secondary indexes, load the data, then create the secondary indexes. You can use the same technique with [foreign keys](#) (load the data first, then set up the foreign keys) if you know the initial data is clean and do not need consistency checks during the loading process.

Whatever sequence of `CREATE TABLE`, `CREATE INDEX`, `ALTER TABLE`, and similar statements went into putting a table together, you can capture the SQL needed to reconstruct the current form of the table by issuing the statement `SHOW CREATE TABLE table\G` (uppercase `\G` required for tidy formatting). This output shows clauses such as numeric precision, `NOT NULL`, and `CHARACTER SET` that are sometimes added behind the scenes, and you might otherwise leave out when cloning the table on a new system or setting up foreign key columns with identical type.

## 5.5.2. Performance and Concurrency Considerations for Online DDL

[Online DDL](#) improves several aspects of MySQL operation, such as performance, concurrency, availability, and scalability:

- Because queries and [DML](#) operations on the table can proceed while the DDL is in progress, applications that access the table are more responsive. Reduced [locking](#) and waiting for other resources all throughout the MySQL server leads to greater scalability, even for operations not involving the table being altered.
- For in-place operations, by avoiding the disk I/O and CPU cycles to rebuild the table, you minimize the overall load on the database and maintain good performance and high throughput during the DDL operation.
- For in-place operations, because less data is read into the [buffer pool](#) than if all the data was copied, you avoid purging frequently accessed data from memory, which formerly could cause a temporary performance dip after a DDL operation.

### Locking Options for Online DDL

While an InnoDB table is being changed by a DDL operation, the table may or may not be [locked](#), depending on the internal workings of that operation and the `LOCK` clause of the `ALTER TABLE` statement. By default, MySQL uses as little locking as possible during a DDL operation; you specify the clause either to make the locking more restrictive than it normally would be (thus limiting concurrent DML, or DML and queries), or to ensure that some expected degree of locking is allowed for an operation. If the `LOCK` clause specifies a level of locking that is not available for that specific kind of DDL operation, such as `LOCK=SHARED` or `LOCK=NONE` while creating or dropping a primary key, the clause works like an assertion, causing the statement to fail with an error. The following list shows the different possibilities for the `LOCK` clause, from the most permissive to the most restrictive:

- For DDL operations with `LOCK=NONE`, both queries and concurrent DML are allowed. This clause makes the `ALTER TABLE` fail if the kind of DDL operation cannot be performed with the requested type of locking, so specify `LOCK=NONE` if keeping the table fully available is vital and it is OK to cancel the DDL if that is not possible. For example, you might use this clause in DDLs for tables



involving customer signups or purchases, to avoid making those tables unavailable by mistakenly issuing an expensive `ALTER TABLE` statement.

- For DDL operations with `LOCK=SHARED`, any writes to the table (that is, DML operations) are blocked, but the data in the table can be read. This clause makes the `ALTER TABLE` fail if the kind of DDL operation cannot be performed with the requested type of locking, so specify `LOCK=SHARED` if keeping the table available for queries is vital and it is OK to cancel the DDL if that is not possible. For example, you might use this clause in DDLs for tables in a data warehouse, where it is OK to delay data load operations until the DDL is finished, but queries cannot be delayed for long periods.
- For DDL operations with `LOCK=DEFAULT`, or with the `LOCK` clause omitted, MySQL uses the lowest level of locking that is available for that kind of operation, allowing concurrent queries, DML, or both wherever possible. This is the setting to use when making pre-planned, pre-tested changes that you know will not cause any availability problems based on the workload for that table.
- For DDL operations with `LOCK=EXCLUSIVE`, both queries and DML operations are blocked. This clause makes the `ALTER TABLE` fail if the kind of DDL operation cannot be performed with the requested type of locking, so specify `LOCK=EXCLUSIVE` if the primary concern is finishing the DDL in the shortest time possible, and it is OK to make applications wait when they try to access the table. You might also use `LOCK=EXCLUSIVE` if the server is supposed to be idle, to avoid unexpected accesses to the table.

An online DDL statement for an InnoDB table always waits for currently executing transactions that are accessing the table to `commit` or `roll back`, because it requires exclusive access to the table for a brief period while the DDL statement is being prepared. Likewise, it requires exclusive access to the table for a brief time before finishing. Thus, an online DDL statement waits for any transactions that are started while the DDL is in progress, and query or modify the table, to `commit` or `roll back` before the DDL completes.

Because there is some processing work involved with recording the changes made by concurrent DML operations, then applying those changes at the end, an online DDL operation could take longer overall than the old-style mechanism that blocks table access from other sessions. The reduction in raw performance is balanced against better responsiveness for applications that use the table. When evaluating the ideal techniques for changing table structure, consider end-user perception of performance, based on factors such as load times for web pages.

A newly created InnoDB secondary index contains only the committed data in the table at the time the `CREATE INDEX` or `ALTER TABLE` statement finishes executing. It does not contain any uncommitted values, old versions of values, or values marked for deletion but not yet removed from the old index.

## Performance of In-Place versus Table-Copying DDL Operations

The raw performance of an online DDL operation is largely determined by whether the operation is performed in-place, or requires copying and rebuilding the entire table. See [Table 5.9, “Summary of Online Status for DDL Operations”](#) to see what kinds of operations can be performed in-place, and any requirements for avoiding table-copy operations.

The performance speedup from in-place DDL applies to operations on secondary indexes, not to the primary key index. The rows of an InnoDB table are stored in a `clustered index` organized based on the `primary key`, forming what some database systems call an “index-organized table”. Because the table structure is so closely tied to the primary key, redefining the primary key still requires copying the data.

When an operation on the primary key uses `ALGORITHM=INPLACE`, even though the data is still copied, it is more efficient than using `ALGORITHM=COPY` because:

- No undo logging or associated redo logging is required for `ALGORITHM=INPLACE`. These operations add overhead to DDL statements that use `ALGORITHM=COPY`.
- The secondary index entries are pre-sorted, and so can be loaded in order.

- The change buffer is not used, because there are no random-access inserts into the secondary indexes.

To judge the relative performance of online DDL operations, you can run such operations on a big [InnoDB](#) table using current and earlier versions of MySQL. You can also run all the performance tests under the latest MySQL version, simulating the previous DDL behavior for the “before” results, by setting the `old_alter_table` [\[517\]](#) system variable. Issue the statement `set old_alter_table=1` in the session, and measure DDL performance to record the “before” figures. Then `set old_alter_table=0` to re-enable the newer, faster behavior, and run the DDL operations again to record the “after” figures.

For a basic idea of whether a DDL operation does its changes in-place or performs a table copy, look at the “rows affected” value displayed after the command finishes. For example, here are lines you might see after doing different types of DDL operations:

- Changing the default value of a column (super-fast, does not affect the table data at all):

```
Query OK, 0 rows affected (0.07 sec)
```

- Adding an index (takes time, but `0 rows affected` shows that the table is not copied):

```
Query OK, 0 rows affected (21.42 sec)
```

- Changing the data type of a column (takes substantial time and does require rebuilding all the rows of the table):

```
Query OK, 1671168 rows affected (1 min 35.54 sec)
```

For example, before running a DDL operation on a big table, you might check whether the operation will be fast or slow as follows:

1. Clone the table structure.
2. Populate the cloned table with a tiny amount of data.
3. Run the DDL operation on the cloned table.
4. Check whether the “rows affected” value is zero or not. A non-zero value means the operation will require rebuilding the entire table, which might require special planning. For example, you might do the DDL operation during a period of scheduled downtime, or on each replication slave server one at a time.

For a deeper understanding of the reduction in MySQL processing, examine the [performance\\_schema](#) and [INFORMATION\\_SCHEMA](#) tables related to [InnoDB](#) before and after DDL operations, to see the number of physical reads, writes, memory allocations, and so on.

### 5.5.3. SQL Syntax for Online DDL

Typically, you do not need to do anything special to enable [online DDL](#) when using the `ALTER TABLE` statement for [InnoDB](#) tables. See [Table 5.9, “Summary of Online Status for DDL Operations”](#) for the kinds of DDL operations that can be performed in-place, allowing concurrent DML, or both. Some variations require particular combinations of configuration settings or `ALTER TABLE` clauses.

You can control the various aspects of a particular online DDL operation by using the `LOCK` and `ALGORITHM` clauses of the `ALTER TABLE` statement. These clauses come at the end of the statement, separated from the table and column specifications by commas. The `LOCK` clause is useful for fine-tuning the degree of concurrent access to the table. The `ALGORITHM` clause is primarily intended for performance comparisons and as a fallback to the older table-copying behavior in case you encounter any issues with existing DDL code. For example:

- To avoid accidentally making the table unavailable for reads, writes, or both, you could specify a clause on the `ALTER TABLE` statement such as `LOCK=NONE` (allow both reads and writes) or



`LOCK=SHARED` (allow reads). The operation halts immediately if the requested level of concurrency is not available.

- To compare performance, you could run one statement with `ALGORITHM=INPLACE` and another with `ALGORITHM=COPY`, as an alternative to setting the `old_alter_table` [517] configuration option.
- To avoid the chance of tying up the server by running an `ALTER TABLE` that copied the table, you could include `ALGORITHM=INPLACE` so the statement halts immediately if it cannot use the in-place mechanism. See Table 5.9, “Summary of Online Status for DDL Operations” for a list of the DDL operations that can or cannot be performed in-place.

See Section 5.5.2, “Performance and Concurrency Considerations for Online DDL” for more details about the `LOCK` clause. For full examples of using online DDL, see Section 5.5.5, “Examples of Online DDL”.

## 5.5.4. Combining or Separating DDL Statements

Before the introduction of [online DDL](#), it was common practice to combine many DDL operations into a single `ALTER TABLE` statement. Because each `ALTER TABLE` statement involved copying and rebuilding the table, it was more efficient to make several changes to the same table at once, since those changes could all be done with a single rebuild operation for the table. The downside was that SQL code involving DDL operations was harder to maintain and to reuse in different scripts. If the specific changes were different each time, you might have to construct a new complex `ALTER TABLE` for each slightly different scenario.

For DDL operations that can be done in-place, as shown in Table 5.9, “Summary of Online Status for DDL Operations”, now you can separate them into individual `ALTER TABLE` statements for easier scripting and maintenance, without sacrificing efficiency. For example, you might take a complicated statement such as:

```
alter table t1 add index i1(c1), add unique index i2(c2), change c4_old_name c4_new_name integer unsigned;
```

and break it down into simpler parts that can be tested and performed independently, such as:

```
alter table t1 add index i1(c1);
alter table t1 add unique index i2(c2);
alter table t1 change c4_old_name c4_new_name integer unsigned not null;
```

You might still use multi-part `ALTER TABLE` statements for:

- Operations that must be performed in a specific sequence, such as creating an index followed by a foreign key constraint that uses that index.
- Operations all using the same specific `LOCK` clause, that you want to either succeed or fail as a group.
- Operations that cannot be performed in-place, that is, that still copy and rebuild the table.
- Operations for which you specify `ALGORITHM=COPY` or `old_alter_table=1`, to force the table-copying behavior if needed for precise backward-compatibility in specialized scenarios.

## 5.5.5. Examples of Online DDL

Here are code examples showing some operations whose performance, concurrency, and scalability are improved by the latest [online DDL](#) enhancements.

- [Example 5.1, “Schema Setup Code for Online DDL Experiments”](#) sets up tables named `BIG_TABLE` and `SMALL_TABLE` used in the subsequent examples.
- [Example 5.2, “Speed and Efficiency of CREATE INDEX and DROP INDEX”](#) illustrates the performance aspects of creating and dropping indexes.

- [Example 5.3, “Concurrent DML During CREATE INDEX and DROP INDEX”](#) shows queries and DML statements running during a `DROP INDEX` operation.
- [Example 5.4, “Renaming a Column”](#) demonstrates the speed improvement for renaming a column, and shows the care needed to keep the data type precisely the same when doing the rename operation.
- [Example 5.5, “Dropping Foreign Keys”](#) demonstrates how foreign keys work with online DDL. Because two tables are involved in foreign key operations, there are extra locking considerations. Thus, tables with foreign keys sometimes have restrictions for online DDL operations.
- [Example 5.6, “Changing Auto-Increment Value”](#) demonstrates how auto-increment columns work with online DDL. Tables with auto-increment columns sometimes have restrictions for online DDL operations.
- [Example 5.7, “Controlling Concurrency with the LOCK Clause”](#) demonstrates the options to permit or restrict concurrent queries and DML operations while an online DDL operation is in progress. It shows the situations when the DDL statement might wait, or the concurrent transaction might wait, or the concurrent transaction might cancel a DML statement due to a deadlock error.
- [Example 5.8, “Schema Setup Code for Online DDL Experiments”](#) demonstrates creating and dropping multiple indexes in a single statement, which can be more efficient than using a separate statement for each index operation.
- [Example 5.9, “Creating and Dropping the Primary Key”](#) demonstrates how it is more efficient to define a `primary key` when creating the table, and relatively expensive to add one later.

### Example 5.1. Schema Setup Code for Online DDL Experiments

Here is the code that sets up the initial tables used in these demonstrations:

```

/*
Setup code for the online DDL demonstration:
- Set up some config variables.
- Create 2 tables that are clones of one of the INFORMATION_SCHEMA tables
  that always has some data. The "small" table has a couple of thousand rows.
  For the "big" table, keep doubling the data until it reaches over a million rows.
- Set up a primary key for the sample tables, since we are demonstrating InnoDB aspects.
*/

set autocommit = 0;
set foreign_key_checks = 1;
set global innodb_file_per_table = 1;
set old_alter_table=0;
prompt mysql:

use test;

\! echo "Setting up 'small' table:"
drop table if exists small_table;
create table small_table as select * from information_schema.columns;
alter table small_table add id int unsigned not null primary key auto_increment;
select count(id) from small_table;

\! echo "Setting up 'big' table:"
drop table if exists big_table;
create table big_table as select * from information_schema.columns;
show create table big_table\G

insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;
insert into big_table select * from big_table;

```

```
insert into big_table select * from big_table;
commit;

alter table big_table add id int unsigned not null primary key auto_increment;
select count(id) from big_table;
```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```
Setting up 'small' table:
Query OK, 0 rows affected (0.01 sec)

Query OK, 1678 rows affected (0.13 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Query OK, 1678 rows affected (0.07 sec)
Records: 1678 Duplicates: 0 Warnings: 0

+-----+
| count(id) |
+-----+
|      1678 |
+-----+
1 row in set (0.00 sec)

Setting up 'big' table:
Query OK, 0 rows affected (0.16 sec)

Query OK, 1678 rows affected (0.17 sec)
Records: 1678 Duplicates: 0 Warnings: 0

***** 1. row *****
      Table: big_table
Create Table: CREATE TABLE `big_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
  `DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
  `COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT ''
) ENGINE=InnoDB DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

Query OK, 1678 rows affected (0.09 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Query OK, 3356 rows affected (0.07 sec)
Records: 3356 Duplicates: 0 Warnings: 0

Query OK, 6712 rows affected (0.17 sec)
Records: 6712 Duplicates: 0 Warnings: 0

Query OK, 13424 rows affected (0.44 sec)
Records: 13424 Duplicates: 0 Warnings: 0

Query OK, 26848 rows affected (0.63 sec)
Records: 26848 Duplicates: 0 Warnings: 0

Query OK, 53696 rows affected (1.72 sec)
```

```

Records: 53696 Duplicates: 0 Warnings: 0

Query OK, 107392 rows affected (3.02 sec)
Records: 107392 Duplicates: 0 Warnings: 0

Query OK, 214784 rows affected (6.28 sec)
Records: 214784 Duplicates: 0 Warnings: 0

Query OK, 429568 rows affected (13.25 sec)
Records: 429568 Duplicates: 0 Warnings: 0

Query OK, 859136 rows affected (28.16 sec)
Records: 859136 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.03 sec)

Query OK, 1718272 rows affected (1 min 9.22 sec)
Records: 1718272 Duplicates: 0 Warnings: 0

+-----+
| count(id) |
+-----+
|   1718272 |
+-----+
1 row in set (1.75 sec)

```

### Example 5.2. Speed and Efficiency of CREATE INDEX and DROP INDEX

Here is a sequence of statements demonstrating the relative speed of `CREATE INDEX` and `DROP INDEX` statements. For a small table, the elapsed time is less than a second whether we use the fast or slow technique, so we look at the “rows affected” output to verify which operations can avoid the table rebuild. For a large table, the difference in efficiency is obvious because skipping the table rebuild saves substantial time.

```

\! clear

\! echo "=== Create and drop index (small table, new/fast technique) ==="
\! echo
\! echo "Data size (kilobytes) before index created: "
\! du -k data/test/small_table.ibd
create index i_dtyp_small on small_table (data_type), algorithm=inplace;
\! echo "Data size after index created: "
\! du -k data/test/small_table.ibd
drop index i_dtyp_small on small_table, algorithm=inplace;

-- Compare against the older slower DDL.

\! echo "=== Create and drop index (small table, old/slow technique) ==="
\! echo
\! echo "Data size (kilobytes) before index created: "
\! du -k data/test/small_table.ibd
create index i_dtyp_small on small_table (data_type), algorithm=copy;
\! echo "Data size after index created: "
\! du -k data/test/small_table.ibd
drop index i_dtyp_small on small_table, algorithm=copy;

-- In the above example, we examined the "rows affected" number,
-- ideally looking for a zero figure. Let's try again with a larger
-- sample size, where we'll see that the actual time taken can
-- vary significantly.

\! echo "=== Create and drop index (big table, new/fast technique) ==="
\! echo
\! echo "Data size (kilobytes) before index created: "
\! du -k data/test/big_table.ibd
create index i_dtyp_big on big_table (data_type), algorithm=inplace;
\! echo "Data size after index created: "
\! du -k data/test/big_table.ibd
drop index i_dtyp_big on big_table, algorithm=inplace;

\! echo "=== Create and drop index (big table, old/slow technique) ==="

```

```

\! echo
\! echo "Data size (kilobytes) before index created: "
\! du -k data/test/big_table.ibd
create index i_dtyp_big on big_table (data_type), algorithm=copy;
\! echo "Data size after index created: "
\! du -k data/test/big_table.ibd
drop index i_dtyp_big on big_table, algorithm=copy;

```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```

Query OK, 0 rows affected (0.00 sec)

=== Create and drop index (small table, new/fast technique) ===

Data size (kilobytes) before index created:
384 data/test/small_table.ibd
Query OK, 0 rows affected (0.04 sec)
Records: 0 Duplicates: 0 Warnings: 0

Data size after index created:
432 data/test/small_table.ibd
Query OK, 0 rows affected (0.02 sec)
Records: 0 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)

=== Create and drop index (small table, old/slow technique) ===

Data size (kilobytes) before index created:
432 data/test/small_table.ibd
Query OK, 1678 rows affected (0.12 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Data size after index created:
448 data/test/small_table.ibd
Query OK, 1678 rows affected (0.10 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)

=== Create and drop index (big table, new/fast technique) ===

Data size (kilobytes) before index created:
315392 data/test/big_table.ibd
Query OK, 0 rows affected (33.32 sec)
Records: 0 Duplicates: 0 Warnings: 0

Data size after index created:
335872 data/test/big_table.ibd
Query OK, 0 rows affected (0.02 sec)
Records: 0 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)

=== Create and drop index (big table, old/slow technique) ===

Data size (kilobytes) before index created:
335872 data/test/big_table.ibd
Query OK, 1718272 rows affected (1 min 5.01 sec)
Records: 1718272 Duplicates: 0 Warnings: 0

Data size after index created:
348160 data/test/big_table.ibd
Query OK, 1718272 rows affected (46.59 sec)
Records: 1718272 Duplicates: 0 Warnings: 0

```

### Example 5.3. Concurrent DML During CREATE INDEX and DROP INDEX

Here are some snippets of code that I ran in separate `mysql` sessions connected to the same database, to illustrate DML statements (insert, update, or delete) running at the same time as `CREATE INDEX` and `DROP INDEX`.

```

/*
CREATE INDEX statement to run against a table while
insert/update/delete statements are modifying the
column being indexed.
*/

-- We'll run this script in one session, while simultaneously creating and dropping
-- an index on test/big_table.table_name in another session.

use test;
create index i_concurrent on big_table(table_name);

```

```

/*
DROP INDEX statement to run against a table while
insert/update/delete statements are modifying the
column being indexed.
*/

-- We'll run this script in one session, while simultaneously creating and dropping
-- an index on test/big_table.table_name in another session.

use test;
drop index i_concurrent on big_table;

```

```

/*
Some queries and insert/update/delete statements to run against a table
while an index is being created or dropped. Previously, these operations
would have stalled during the index create/drop period and possibly
timed out or deadlocked.
*/

-- We'll run this script in one session, while simultaneously creating and dropping
-- an index on test/big_table.table_name in another session.

-- In our test instance, that column has about 1.7M rows, with 136 different values.
-- Sample values: COLUMNS (20480), ENGINES (6144), EVENTS (24576), FILES (38912), TABLES (21504), VIEWS

set autocommit = 0;
use test;

select distinct character_set_name from big_table where table_name = 'FILES';
delete from big_table where table_name = 'FILES';
select distinct character_set_name from big_table where table_name = 'FILES';

-- I'll issue the final rollback interactively, not via script,
-- the better to control the timing.
-- rollback;

```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```

mysql: source concurrent_ddl_create.sql
Database changed
Query OK, 0 rows affected (1 min 25.15 sec)
Records: 0 Duplicates: 0 Warnings: 0

mysql: source concurrent_ddl_drop.sql
Database changed
Query OK, 0 rows affected (24.98 sec)
Records: 0 Duplicates: 0 Warnings: 0

mysql: source concurrent_dml.sql
Query OK, 0 rows affected (0.00 sec)

Database changed
+-----+
| character_set_name |
+-----+
| NULL               |
| utf8               |
+-----+
2 rows in set (0.32 sec)

```

```
Query OK, 38912 rows affected (1.84 sec)

Empty set (0.01 sec)

mysql: rollback;
Query OK, 0 rows affected (1.05 sec)
```

### Example 5.4. Renaming a Column

Here is a demonstration of using `ALTER TABLE` to rename a column. We use the new, fast DDL mechanism to change the name, then the old, slow DDL mechanism (with `old_alter_table=1`) to restore the original column name.

Notes:

- Because the syntax for renaming a column also involves re-specifying the data type, be very careful to specify exactly the same data type to avoid a costly table rebuild. In this case, we checked the output of `show create table table\G` and copied any clauses such as `CHARACTER SET` and `NOT NULL` from the original column definition.
- Again, renaming a column for a small table is fast enough that we need to examine the “rows affected” number to verify that the new DDL mechanism is more efficient than the old one. With a big table, the difference in elapsed time makes the improvement obvious.

```
/*
Run through a sequence of 'rename column' statements.
Because this operation involves only metadata, not table data,
it is fast for big and small tables, with new or old DDL mechanisms.
*/

\! clear

\! echo "Rename column (fast technique, small table):"
alter table small_table change `IS_NULLABLE` `NULLABLE` varchar(3) character set utf8 not null, algorithm=innodb;
\! echo "Rename back to original name (slow technique):"
alter table small_table change `NULLABLE` `IS_NULLABLE` varchar(3) character set utf8 not null, algorithm=copy;

\! echo "Rename column (fast technique, big table):"
alter table big_table change `IS_NULLABLE` `NULLABLE` varchar(3) character set utf8 not null, algorithm=innodb;
\! echo "Rename back to original name (slow technique):"
alter table big_table change `NULLABLE` `IS_NULLABLE` varchar(3) character set utf8 not null, algorithm=copy;
```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```
Rename column (fast technique, small table):
Query OK, 0 rows affected (0.05 sec)

Query OK, 0 rows affected (0.13 sec)
Records: 0 Duplicates: 0 Warnings: 0

Rename back to original name (slow technique):
Query OK, 0 rows affected (0.00 sec)

Query OK, 1678 rows affected (0.35 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Rename column (fast technique, big table):
Query OK, 0 rows affected (0.00 sec)

Query OK, 0 rows affected (0.11 sec)
Records: 0 Duplicates: 0 Warnings: 0

Rename back to original name (slow technique):
Query OK, 0 rows affected (0.00 sec)

Query OK, 1718272 rows affected (1 min 0.00 sec)
Records: 1718272 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)
```

**Example 5.5. Dropping Foreign Keys**

Here is a demonstration of foreign keys, including improvement to the speed of dropping a foreign key constraint.

```

/*
Demonstrate aspects of foreign keys that are or aren't affected by the DDL improvements.
- Create a new table with only a few values to serve as the parent table.
- Set up the 'small' and 'big' tables as child tables using a foreign key.
- Verify that the ON DELETE CASCADE clause makes changes ripple from parent to child tables.
- Drop the foreign key constraints, and optionally associated indexes. (This is the operation that is s
*/

\! clear

-- Make sure foreign keys are being enforced, and allow
-- rollback after doing some DELETES that affect both
-- parent and child tables.
set foreign_key_checks = 1;
set autocommit = 0;

-- Create a parent table, containing values that we know are already present
-- in the child tables.
drop table if exists schema_names;
create table schema_names (id int unsigned not null primary key auto_increment, schema_name varchar(64))

show create table schema_names\G
show create table small_table\G
show create table big_table\G

-- Creating the foreign key constraint still involves a table rebuild when foreign_key_checks=1,
-- as illustrated by the "rows affected" figure.
alter table small_table add constraint small_fk foreign key i_table_schema (table_schema) references schema_names (id) on delete cascade;
alter table big_table add constraint big_fk foreign key i_table_schema (table_schema) references schema_names (id) on delete cascade;

show create table small_table\G
show create table big_table\G

select schema_name from schema_names order by schema_name;
select count(table_schema) howmany, table_schema from small_table group by table_schema;
select count(table_schema) howmany, table_schema from big_table group by table_schema;

-- big_table is the parent table.
-- schema_names is the parent table.
-- big_table is the child table.
-- (One row in the parent table can have many "children" in the child table.)
-- Changes to the parent table can ripple through to the child table.
-- For example, removing the value 'test' from schema_names.schema_name will
-- result in the removal of 20K or so rows from big_table.

delete from schema_names where schema_name = 'test';

select schema_name from schema_names order by schema_name;
select count(table_schema) howmany, table_schema from small_table group by table_schema;
select count(table_schema) howmany, table_schema from big_table group by table_schema;

-- Because we've turned off autocommit, we can still get back those deleted rows
-- if the DELETE was issued by mistake.
rollback;

select schema_name from schema_names order by schema_name;
select count(table_schema) howmany, table_schema from small_table group by table_schema;
select count(table_schema) howmany, table_schema from big_table group by table_schema;

-- All of the cross-checking between parent and child tables would be
-- deadly slow if there wasn't the requirement for the corresponding
-- columns to be indexed!

-- But we can get rid of the foreign key using a fast operation
-- that doesn't rebuild the table.
-- If we didn't specify a constraint name when setting up the foreign key, we would
-- have to find the auto-generated name such as 'big_table_ibfk_1' in the

```



```
-- output from 'show create table'.

-- For the small table, we'll drop the foreign key and the associated index.
-- Having an index on a small table is less critical.

\! echo "DROP FOREIGN KEY and INDEX from small_table:"
alter table small_table drop foreign key small_fk, drop index small_fk;

-- For the big table, we'll drop the foreign key and leave the associated index.
-- If we are still doing queries that reference the indexed column, the index is
-- very important to avoid a full table scan of the big table.
\! echo "DROP FOREIGN KEY from big_table:"
alter table big_table drop foreign key big_fk;

show create table small_table\G
show create table big_table\G
```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```
Query OK, 0 rows affected (0.00 sec)

Query OK, 0 rows affected (0.00 sec)

Query OK, 0 rows affected (0.01 sec)

Query OK, 4 rows affected (0.03 sec)
Records: 4 Duplicates: 0 Warnings: 0

***** 1. row *****
      Table: schema_names
Create Table: CREATE TABLE `schema_names` (
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  `schema_name` varchar(64) CHARACTER SET utf8 NOT NULL,
  PRIMARY KEY (`id`),
  KEY `i_schema` (`schema_name`)
) ENGINE=InnoDB AUTO_INCREMENT=8 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

***** 1. row *****
      Table: small_table
Create Table: CREATE TABLE `small_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
  `DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
  `COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB AUTO_INCREMENT=1679 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

***** 1. row *****
      Table: big_table
Create Table: CREATE TABLE `big_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
```

```

`TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
`COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
`ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
`COLUMN_DEFAULT` longtext CHARACTER SET utf8,
`IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
`DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
`CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
`CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
`NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
`NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
`DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
`CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
`COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
`COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
`COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
`EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
`PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
`COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
`id` int(10) unsigned NOT NULL AUTO_INCREMENT,
PRIMARY KEY (`id`),
KEY `big_fk` (`TABLE_SCHEMA`)
) ENGINE=InnoDB AUTO_INCREMENT=1718273 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

Query OK, 1678 rows affected (0.10 sec)
Records: 1678 Duplicates: 0 Warnings: 0

Query OK, 1718272 rows affected (1 min 14.54 sec)
Records: 1718272 Duplicates: 0 Warnings: 0

***** 1. row *****
Table: small_table
Create Table: CREATE TABLE `small_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
  `DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
  `COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
PRIMARY KEY (`id`),
KEY `small_fk` (`TABLE_SCHEMA`),
CONSTRAINT `small_fk` FOREIGN KEY (`TABLE_SCHEMA`) REFERENCES `schema_names` (`schema_name`) ON DELETE NO ACTION
) ENGINE=InnoDB AUTO_INCREMENT=1679 DEFAULT CHARSET=latin1
1 row in set (0.12 sec)

***** 1. row *****
Table: big_table
Create Table: CREATE TABLE `big_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,

```

```

`CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
`NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
`NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
`DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
`CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
`COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
`COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
`COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
`EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
`PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
`COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
`id` int(10) unsigned NOT NULL AUTO_INCREMENT,
PRIMARY KEY (`id`),
KEY `big_fk` (`TABLE_SCHEMA`),
CONSTRAINT `big_fk` FOREIGN KEY (`TABLE_SCHEMA`) REFERENCES `schema_names` (`schema_name`) ON DELETE CASCADE
) ENGINE=InnoDB AUTO_INCREMENT=1718273 DEFAULT CHARSET=latin1
1 row in set (0.01 sec)

+-----+
| schema_name |
+-----+
| information_schema |
| mysql |
| performance_schema |
| test |
+-----+
4 rows in set (0.00 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 563 | information_schema |
| 286 | mysql |
| 786 | performance_schema |
| 43 | test |
+-----+-----+
4 rows in set (0.01 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 576512 | information_schema |
| 292864 | mysql |
| 804864 | performance_schema |
| 44032 | test |
+-----+-----+
4 rows in set (2.10 sec)

Query OK, 1 row affected (1.52 sec)

+-----+
| schema_name |
+-----+
| information_schema |
| mysql |
| performance_schema |
+-----+
3 rows in set (0.00 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 563 | information_schema |
| 286 | mysql |
| 786 | performance_schema |
+-----+-----+
3 rows in set (0.00 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 576512 | information_schema |
+-----+-----+

```

```

| 292864 | mysql |
| 804864 | performance_schema |
+-----+-----+
3 rows in set (1.74 sec)

Query OK, 0 rows affected (0.60 sec)

+-----+-----+
| schema_name |
+-----+-----+
| information_schema |
| mysql |
| performance_schema |
| test |
+-----+-----+
4 rows in set (0.00 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 563 | information_schema |
| 286 | mysql |
| 786 | performance_schema |
| 43 | test |
+-----+-----+
4 rows in set (0.01 sec)

+-----+-----+
| howmany | table_schema |
+-----+-----+
| 576512 | information_schema |
| 292864 | mysql |
| 804864 | performance_schema |
| 44032 | test |
+-----+-----+
4 rows in set (1.59 sec)

DROP FOREIGN KEY and INDEX from small_table:
Query OK, 0 rows affected (0.02 sec)
Records: 0 Duplicates: 0 Warnings: 0

DROP FOREIGN KEY from big_table:
Query OK, 0 rows affected (0.02 sec)
Records: 0 Duplicates: 0 Warnings: 0

***** 1. row *****
Table: small_table
Create Table: CREATE TABLE `small_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
  `DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
  `COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB AUTO_INCREMENT=1679 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

```

```

***** 1. row *****
      Table: big_table
Create Table: CREATE TABLE `big_table` (
  `TABLE_CATALOG` varchar(512) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_SCHEMA` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `TABLE_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_NAME` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `ORDINAL_POSITION` bigint(21) unsigned NOT NULL DEFAULT '0',
  `COLUMN_DEFAULT` longtext CHARACTER SET utf8,
  `IS_NULLABLE` varchar(3) CHARACTER SET utf8 NOT NULL,
  `DATA_TYPE` varchar(64) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `CHARACTER_MAXIMUM_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_OCTET_LENGTH` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `NUMERIC_SCALE` bigint(21) unsigned DEFAULT NULL,
  `DATETIME_PRECISION` bigint(21) unsigned DEFAULT NULL,
  `CHARACTER_SET_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLLATION_NAME` varchar(32) CHARACTER SET utf8 DEFAULT NULL,
  `COLUMN_TYPE` longtext CHARACTER SET utf8 NOT NULL,
  `COLUMN_KEY` varchar(3) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `EXTRA` varchar(30) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `PRIVILEGES` varchar(80) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `COLUMN_COMMENT` varchar(1024) CHARACTER SET utf8 NOT NULL DEFAULT '',
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  PRIMARY KEY (`id`),
  KEY `big_fk` (`TABLE_SCHEMA`)
) ENGINE=InnoDB AUTO_INCREMENT=1718273 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

```

### Example 5.6. Changing Auto-Increment Value

Here is an illustration of increasing the [auto-increment](#) lower limit for a table column, demonstrating how this operation now avoids a table rebuild, plus some other fun facts about [InnoDB](#) auto-increment columns.

```

/*
If this script is run after foreign_key.sql, the schema_names table is
already set up. But to allow this script to run multiple times without
running into duplicate ID errors, we set up the schema_names table
all over again.
*/

\! clear

\! echo "=== Adjusting the Auto-Increment Limit for a Table ==="
\! echo

drop table if exists schema_names;
create table schema_names (id int unsigned not null primary key auto_increment,
  schema_name varchar(64) character set utf8 not null, index i_schema (schema_name))
as select distinct table_schema schema_name from small_table;

\! echo "Initial state of schema_names table. AUTO_INCREMENT is included in SHOW CREATE TABLE output."
\! echo "Note how MySQL reserved a block of IDs, but only needed 4 of them in this transaction, so the next"
show create table schema_names\G
select * from schema_names order by id;

\! echo "Inserting even a tiny amount of data can produce gaps in the ID sequence."
insert into schema_names (schema_name) values ('eight'), ('nine');

\! echo "Bumping auto-increment lower limit to 20 (fast mechanism):"
alter table schema_names auto_increment=20, algorithm=inplace;

\! echo "Inserting 2 rows that should get IDs 20 and 21:"
insert into schema_names (schema_name) values ('foo'), ('bar');
commit;

\! echo "Bumping auto-increment lower limit to 30 (slow mechanism):"
alter table schema_names auto_increment=30, algorithm=copy;

```

```

\! echo "Inserting 2 rows that should get IDs 30 and 31:"
insert into schema_names (schema_name) values ('bletch'),('baz');
commit;

select * from schema_names order by id;

\! echo "Final state of schema_names table. AUTO_INCREMENT value shows the next inserted row would get
show create table schema_names\G

```

Running this code gives this output, condensed for brevity and with the most important points bolded:

```

=== Adjusting the Auto-Increment Limit for a Table ===

Query OK, 0 rows affected (0.01 sec)

Query OK, 4 rows affected (0.02 sec)
Records: 4 Duplicates: 0 Warnings: 0

Initial state of schema_names table. AUTO_INCREMENT is included in SHOW CREATE TABLE output.
Note how MySQL reserved a block of IDs, but only needed 4 of them in this transaction, so the next inserted row
***** 1. row *****

Table: schema_names
Create Table: CREATE TABLE `schema_names` (
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  `schema_name` varchar(64) CHARACTER SET utf8 NOT NULL,
  PRIMARY KEY (`id`),
  KEY `i_schema` (`schema_name`)
) ENGINE=InnoDB AUTO_INCREMENT=8 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

+----+-----+
| id | schema_name |
+----+-----+
|  1 | information_schema |
|  2 | mysql |
|  3 | performance_schema |
|  4 | test |
+----+-----+
4 rows in set (0.00 sec)

Inserting even a tiny amount of data can produce gaps in the ID sequence.
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)

Bumping auto-increment lower limit to 20 (fast mechanism):
Query OK, 0 rows affected (0.01 sec)
Records: 0 Duplicates: 0 Warnings: 0

Inserting 2 rows that should get IDs 20 and 21:
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.00 sec)

Query OK, 0 rows affected (0.00 sec)

Bumping auto-increment lower limit to 30 (slow mechanism):
Query OK, 8 rows affected (0.02 sec)
Records: 8 Duplicates: 0 Warnings: 0

Inserting 2 rows that should get IDs 30 and 31:
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

Query OK, 0 rows affected (0.01 sec)

+----+-----+
| id | schema_name |
+----+-----+
|  1 | information_schema |

```

```

| 2 | mysql
| 3 | performance_schema
| 4 | test
| 8 | eight
| 9 | nine
| 20 | foo
| 21 | bar
| 30 | bletch
| 31 | baz
+-----+
10 rows in set (0.00 sec)

Query OK, 0 rows affected (0.00 sec)

Final state of schema_names table. AUTO_INCREMENT value shows the next inserted row would get ID=32.
***** 1. row *****
      Table: schema_names
Create Table: CREATE TABLE `schema_names` (
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  `schema_name` varchar(64) CHARACTER SET utf8 NOT NULL,
  PRIMARY KEY (`id`),
  KEY `i_schema` (`schema_name`)
) ENGINE=InnoDB AUTO_INCREMENT=32 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

```

### Example 5.7. Controlling Concurrency with the `LOCK` Clause

This example shows how to use the `LOCK` clause of the `ALTER TABLE` statement to allow or deny concurrent access to the table while an online DDL operation is in progress. The clause has settings that allow queries and DML statements (`LOCK=NONE`), just queries (`LOCK=SHARED`), or no concurrent access at all (`LOCK=EXCLUSIVE`).

In one session, we run a succession of `ALTER TABLE` statements to create and drop an index, using different values for the `LOCK` clause to see what happens with waiting or deadlocking in either session. We are using the same `BIG_TABLE` table as in previous examples, starting with approximately 1.7 million rows. For illustration purposes, we will index and query the `IS_NULLABLE` column. (Although in real life it would be silly to make an index for a tiny column with only 2 distinct values.)

```

mysql: desc big_table;
+-----+-----+-----+-----+-----+-----+
| Field          | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| TABLE_CATALOG | varchar(512)  | NO   |     |         |       |
| TABLE_SCHEMA | varchar(64)   | NO   |     |         |       |
| TABLE_NAME    | varchar(64)   | NO   |     |         |       |
| COLUMN_NAME    | varchar(64)   | NO   |     |         |       |
| ORDINAL_POSITION | bigint(21) unsigned | NO   |     | 0       |       |
| COLUMN_DEFAULT | longtext      | YES  |     | NULL    |       |
| IS_NULLABLE    | varchar(3)    | NO   |     |         |       |
...
+-----+-----+-----+-----+-----+-----+
21 rows in set (0.14 sec)

mysql: alter table big_table add index i1(is_nullable);
Query OK, 0 rows affected (20.71 sec)

mysql: alter table big_table drop index i1;
Query OK, 0 rows affected (0.02 sec)

mysql: alter table big_table add index i1(is_nullable), lock=exclusive;
Query OK, 0 rows affected (19.44 sec)

mysql: alter table big_table drop index i1;
Query OK, 0 rows affected (0.03 sec)

mysql: alter table big_table add index i1(is_nullable), lock=shared;
Query OK, 0 rows affected (16.71 sec)

mysql: alter table big_table drop index i1;
Query OK, 0 rows affected (0.05 sec)

```

```
mysql: alter table big_table add index i1(is_nullable), lock=none;
Query OK, 0 rows affected (12.26 sec)

mysql: alter table big_table drop index i1;
Query OK, 0 rows affected (0.01 sec)

... repeat statements like the above while running queries ...
... and DML statements at the same time in another session ...
```

Nothing dramatic happens in the session running the DDL statements. Sometimes, an `ALTER TABLE` takes unusually long because it is waiting for another transaction to finish, when that transaction modified the table during the DDL or queried the table before the DDL:

```
mysql: alter table big_table add index i1(is_nullable), lock=none;
Query OK, 0 rows affected (59.27 sec)

mysql: -- The previous ALTER took so long because it was waiting for all the concurrent
mysql: -- transactions to commit or roll back.

mysql: alter table big_table drop index i1;
Query OK, 0 rows affected (41.05 sec)

mysql: -- Even doing a SELECT on the table in the other session first causes
mysql: -- the ALTER TABLE above to stall until the transaction
mysql: -- surrounding the SELECT is committed or rolled back.
```

Here is the log from another session running concurrently, where we issue queries and DML statements against the table before, during, and after the DDL operations shown in the previous listings. This first listing shows queries only. We expect the queries to be allowed during DDL operations using `LOCK=NONE` or `LOCK=SHARED`, and for the query to wait until the DDL is finished if the `ALTER TABLE` statement includes `LOCK=EXCLUSIVE`.

```
mysql: show variables like 'autocommit';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| autocommit    | ON    |
+-----+-----+
1 row in set (0.01 sec)

mysql: -- A trial query before any ADD INDEX in the other session:
mysql: -- Note: because autocommit is enabled, each
mysql: -- transaction finishes immediately after the query.
mysql: select distinct is_nullable from big_table;
+-----+
| is_nullable |
+-----+
| NO          |
| YES        |
+-----+
2 rows in set (4.49 sec)

mysql: -- Index is being created with LOCK=EXCLUSIVE on the ALTER statement.
mysql: -- The query waits until the DDL is finished before proceeding.
mysql: select distinct is_nullable from big_table;
+-----+
| is_nullable |
+-----+
| NO          |
| YES        |
+-----+
2 rows in set (17.26 sec)

mysql: -- Index is being created with LOCK=SHARED on the ALTER statement.
mysql: -- The query returns its results while the DDL is in progress.
mysql: -- The same thing happens with LOCK=NONE on the ALTER statement.
mysql: select distinct is_nullable from big_table;
+-----+
| is_nullable |
+-----+
```



```

| NO          |
| YES        |
+-----+
2 rows in set (3.11 sec)

mysql: -- Once the index is created, and with no DDL in progress,
mysql: -- queries referencing the indexed column are very fast:
mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
|  411648 |
+-----+
1 row in set (0.20 sec)

mysql: select distinct is_nullable from big_table;
+-----+
| is_nullable |
+-----+
| NO          |
| YES        |
+-----+
2 rows in set (0.00 sec)

```

Now in this concurrent session, we run some transactions including DML statements, or a combination of DML statements and queries. We use `DELETE` statements to illustrate predictable, verifiable changes to the table. Because the transactions in this part can span multiple statements, we run these tests with `autocommit` [464] turned off.

```

mysql: set global autocommit = off;
Query OK, 0 rows affected (0.00 sec)

mysql: -- Count the rows that will be involved in our DELETE statements:
mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
|  411648 |
+-----+
1 row in set (0.95 sec)

mysql: -- After this point, any DDL statements back in the other session
mysql: -- stall until we commit or roll back.

mysql: delete from big_table where is_nullable = 'YES' limit 11648;
Query OK, 11648 rows affected (0.14 sec)

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
|  400000 |
+-----+
1 row in set (1.04 sec)

mysql: rollback;
Query OK, 0 rows affected (0.09 sec)

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
|  411648 |
+-----+
1 row in set (0.93 sec)

mysql: -- OK, now we're going to try that during index creation with LOCK=NONE.
mysql: delete from big_table where is_nullable = 'YES' limit 11648;
Query OK, 11648 rows affected (0.21 sec)

mysql: -- We expect that now there will be 400000 'YES' rows left:

```

```

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 400000 |
+-----+
1 row in set (1.25 sec)

mysql: -- In the other session, the ALTER TABLE is waiting before finishing,
mysql: -- because _this_ transaction hasn't committed or rolled back yet.
mysql: rollback;
Query OK, 0 rows affected (0.11 sec)

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 411648 |
+-----+
1 row in set (0.19 sec)

mysql: -- The ROLLBACK left the table in the same state we originally found it.
mysql: -- Now let's make a permanent change while the index is being created,
mysql: -- again with ALTER TABLE ... , LOCK=NONE.
mysql: -- First, commit so the DROP INDEX in the other shell can finish;
mysql: -- the previous SELECT started a transaction that accessed the table.
mysql: commit;
Query OK, 0 rows affected (0.00 sec)

mysql: -- Now we add the index back in the other shell, then issue DML in this one
mysql: -- while the DDL is running.
mysql: delete from big_table where is_nullable = 'YES' limit 11648;
Query OK, 11648 rows affected (0.23 sec)

mysql: commit;
Query OK, 0 rows affected (0.01 sec)

mysql: -- In the other shell, the ADD INDEX has finished.
mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 400000 |
+-----+
1 row in set (0.19 sec)

mysql: -- At the point the new index is finished being created, it contains entries
mysql: -- only for the 400000 'YES' rows left when all concurrent transactions are finished.
mysql:
mysql: -- Now we will run a similar test, while ALTER TABLE ... , LOCK=SHARED is running.
mysql: -- We expect a query to complete during the ALTER TABLE, but for the DELETE
mysql: -- to run into some kind of issue.
mysql: commit;
Query OK, 0 rows affected (0.00 sec)

mysql: -- As expected, the query returns results while the LOCK=SHARED DDL is running:
mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 400000 |
+-----+
1 row in set (2.07 sec)

mysql: -- The DDL in the other session is not going to finish until this transaction
mysql: -- is committed or rolled back. If we tried a DELETE now and it waited because
mysql: -- of LOCK=SHARED on the DDL, both transactions would wait forever (deadlock).
mysql: -- MySQL detects this condition and cancels the attempted DML statement.
mysql: delete from big_table where is_nullable = 'YES' limit 100000;
ERROR 1213 (40001): Deadlock found when trying to get lock; try restarting transaction
mysql: -- The transaction here is still going, so in the other shell, the ADD INDEX operation
mysql: -- is waiting for this transaction to commit or roll back.

```

```

mysql: rollback;
Query OK, 0 rows affected (0.00 sec)

mysql: -- Now let's try issuing a query and some DML, on one line, while running
mysql: -- ALTER TABLE ... , LOCK=EXCLUSIVE in the other shell.
mysql: -- Notice how even the query is held up until the DDL is finished.
mysql: -- By the time the DELETE is issued, there is no conflicting access
mysql: -- to the table and we avoid the deadlock error.
mysql: select count(*) from big_table where is_nullable = 'YES'; delete from big_table where is_nullable =
+-----+
| count(*) |
+-----+
| 400000 |
+-----+
1 row in set (15.98 sec)

Query OK, 100000 rows affected (2.81 sec)

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 300000 |
+-----+
1 row in set (0.17 sec)

mysql: rollback;
Query OK, 0 rows affected (1.36 sec)

mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 400000 |
+-----+
1 row in set (0.19 sec)

mysql: commit;
Query OK, 0 rows affected (0.00 sec)

mysql: -- Next, we try ALTER TABLE ... , LOCK=EXCLUSIVE in the other session
mysql: -- and only issue DML, not any query, in the concurrent transaction here.
mysql: delete from big_table where is_nullable = 'YES' limit 100000;
Query OK, 100000 rows affected (16.37 sec)

mysql: -- That was OK because the ALTER TABLE did not have to wait for the transaction
mysql: -- here to complete. The DELETE in this session waited until the index was ready.
mysql: select count(*) from big_table where is_nullable = 'YES';
+-----+
| count(*) |
+-----+
| 300000 |
+-----+
1 row in set (0.16 sec)

mysql: commit;
Query OK, 0 rows affected (0.00 sec)

```

In the preceding example listings, we learned that:

- The `LOCK` clause for `ALTER TABLE` is set off from the rest of the statement by a comma.
- Online DDL operations might wait before starting, until any prior transactions that access the table are committed or rolled back.
- Online DDL operations might wait before completing, until any concurrent transactions that access the table are committed or rolled back.
- While an online DDL operation is running, concurrent queries are relatively straightforward, as long as the `ALTER TABLE` statement uses `LOCK=NONE` or `LOCK=SHARED`.

- Pay attention to whether `autocommit` [464] is turned on or off. If it is turned off, be careful to end transactions in other sessions (even just queries) before performing DDL operations on the table.
- With `LOCK=SHARED`, concurrent transactions that mix queries and DML could encounter deadlock errors and have to be restarted after the DDL is finished.
- With `LOCK=NONE`, concurrent transactions can freely mix queries and DML. The DDL operation waits until the concurrent transactions are committed or rolled back.
- With `LOCK=NONE`, concurrent transactions can freely mix queries and DML, but those transactions wait until the DDL operation is finished before they can access the table.

### Example 5.8. Schema Setup Code for Online DDL Experiments

You can create multiple indexes on a table with one `ALTER TABLE` statement. This is relatively efficient, because the clustered index of the table needs to be scanned only once (although the data is sorted separately for each new index). For example:

```
CREATE TABLE T1(A INT PRIMARY KEY, B INT, C CHAR(1)) ENGINE=InnoDB;
INSERT INTO T1 VALUES (1,2,'a'), (2,3,'b'), (3,2,'c'), (4,3,'d'), (5,2,'e');
COMMIT;
ALTER TABLE T1 ADD INDEX (B), ADD UNIQUE INDEX (C);
```

The above statements create table `T1` with the primary key on column `A`, insert several rows, then build two new indexes on columns `B` and `C`. If there were many rows inserted into `T1` before the `ALTER TABLE` statement, this approach is much more efficient than creating all the secondary indexes before loading the data.

Because dropping InnoDB secondary indexes also does not require any copying of table data, it is equally efficient to drop multiple indexes with a single `ALTER TABLE` statement or multiple `DROP INDEX` statements:

```
ALTER TABLE T1 DROP INDEX B, DROP INDEX C;
```

or:

```
DROP INDEX B ON T1;
DROP INDEX C ON T1;
```

### Example 5.9. Creating and Dropping the Primary Key

Restructuring the `clustered index` for an InnoDB table always requires copying the table data. Thus, it is best to define the `primary key` when you create a table, rather than issuing `ALTER TABLE ... ADD PRIMARY KEY` later, to avoid rebuilding the table.

Defining a `PRIMARY KEY` later causes the data to be copied, as in the following example:

```
CREATE TABLE T2 (A INT, B INT);
INSERT INTO T2 VALUES (NULL, 1);
ALTER TABLE T2 ADD PRIMARY KEY (B);
```

When you create a `UNIQUE` or `PRIMARY KEY` index, MySQL must do some extra work. For `UNIQUE` indexes, MySQL checks that the table contains no duplicate values for the key. For a `PRIMARY KEY` index, MySQL also checks that none of the `PRIMARY KEY` columns contains a `NULL`.

When you add a primary key using the `ALGORITHM=COPY` clause, MySQL actually converts `NULL` values in the associated columns to default values: 0 for numbers, the empty string for character-based columns and BLOBs, and January 1, 1975 for dates. This is a non-standard behavior that Oracle recommends you not rely on. Adding a primary key using `ALGORITHM=INPLACE` is only allowed when the `SQL_MODE` [545] setting includes the `strict_trans_tables` or `strict_all_tables` flags; when the `SQL_MODE` setting is strict, `ADD PRIMARY KEY ... , ALGORITHM=INPLACE` is allowed,

but the statement can still fail if the requested primary key columns contain any `NULL` values. The `ALGORITHM=INPLACE` behavior is more standard-compliant.

The following example shows the different possibilities for the `ADD PRIMARY KEY` clause. With the `ALGORITHM=COPY` clause, the operation succeeds despite the presence of `NULL` values in the primary key columns; the data is silently changed, which could cause problems. With the `ALGORITHM=INPLACE` clause, the operation could fail for different reasons, because this setting considers data integrity a high priority: the statement gives an error if the `SQL_MODE` [545] setting is not “strict” enough, or if the primary key columns contain any `NULL` values. Once we address both of those requirements, the `ALTER TABLE` operation succeeds.

```
CREATE TABLE add_pk_via_copy (c1 INT, c2 VARCHAR(10), c3 DATETIME);
INSERT INTO add_pk_via_copy VALUES (1,'a','...'),(NULL,NULL,NULL);
ALTER TABLE add_pk_via_copy ADD PRIMARY KEY (c1,c2,c3), ALGORITHM=COPY;
SELECT * FROM add_pk_via_copy;

CREATE TABLE add_pk_via_inplace (c1 INT, c2 VARCHAR(10), c3 DATETIME);
INSERT INTO add_pk_via_inplace VALUES (1,'a','...'),(NULL,NULL,NULL);
SET sql_mode = 'strict_trans_tables';
ALTER TABLE add_pk_via_inplace ADD PRIMARY KEY (c1,c2,c3), ALGORITHM=COPY;
SET sql_mode = '';
ALTER TABLE add_pk_via_inplace ADD PRIMARY KEY (c1,c2,c3), ALGORITHM=COPY;
DELETE FROM add_pk_via_inplace WHERE c1 IS NULL OR c2 IS NULL OR c3 IS NULL;
ALTER TABLE add_pk_via_inplace ADD PRIMARY KEY (c1,c2,c3), ALGORITHM=COPY;
SELECT * FROM add_pk_via_inplace;
```

If you create a table without a primary key, InnoDB chooses one for you, which can be the first `UNIQUE` key defined on `NOT NULL` columns, or a system-generated key. To avoid any uncertainty and the potential space requirement for an extra hidden column, specify the `PRIMARY KEY` clause as part of the `CREATE TABLE` statement.

## 5.5.6. Implementation Details of Online DDL

Each `ALTER TABLE` operation for an InnoDB table is governed by several aspects:

- Whether there is any change to the physical representation of the table, or whether it purely a change to metadata that can be done without touching the table itself.
- Whether the volume of data in the table stays the same, increases, or decreases.
- Whether a change in table data involves the clustered index, secondary indexes, or both.
- Whether there are any `foreign key` relationships between the table being altered and some other table. The mechanics differ depending on whether the `foreign_key_checks` [483] configuration option is enabled or disabled.
- Whether the table is partitioned. Partitioning clauses of `ALTER TABLE` are turned into low-level operations involving one or more tables, and those operations follow the regular rules for online DDL.
- Whether the table data must be copied, whether the table can be reorganized “in-place”, or a combination of both.
- Whether the table contains any `auto-increment` columns.
- What degree of `locking` is required, either by the nature of the underlying database operations, or a `LOCK` clause that you specify in the `ALTER TABLE` statement.

This section explains how these factors affect the different kinds of `ALTER TABLE` operations on InnoDB tables.

## Error Conditions for Online DDL

Here are the primary reasons why an online DDL operation could fail:

- If a `LOCK` clause specifies a low degree of locking (`SHARED` or `NONE`) that is not compatible with the particular type of DDL operation.
- If a timeout occurs while waiting to get an `exclusive lock` on the table, which is needed briefly during the initial and final phases of the DDL operation.
- If the `tmpdir` [559] file system runs out of disk space, while MySQL writes temporary sort files on disk during index creation.
- If the `ALTER TABLE` takes so long, and concurrent DML modifies the table so much, that the size of the temporary online log exceeds the value of the `innodb_online_alter_log_max_size` [1739] configuration option. This condition causes a `DB_ONLINE_LOG_TOO_BIG` error.
- If concurrent DML makes changes to the table that are allowed with the original table definition, but not with the new one. The operation only fails at the very end, when MySQL tries to apply all the changes from concurrent DML statements. For example, you might insert duplicate values into a column while a unique index is being created, or you might insert `NULL` values into a column while creating a `primary key` index on that column. The changes made by the concurrent DML take precedence, and the `ALTER TABLE` operation is effectively `rolled back`.

Although the configuration option `innodb_file_per_table` [1718] has a dramatic effect on the representation for an `InnoDB` table, all online DDL operations work equally well whether that option is enabled or disabled, and whether the table is physically located in its own `.ibd file` or inside the `system tablespace`.

`InnoDB` has two types of indexes: the `clustered index` representing all the data in the table, and optional `secondary indexes` to speed up queries. Since the clustered index contains the data values in its B-tree nodes, adding or dropping a clustered index does involve copying the data, and creating a new copy of the table. A secondary index, however, contains only the index key and the value of the primary key. This type of index can be created or dropped without copying the data in the clustered index. Because each secondary index contains copies of the primary key values (used to access the clustered index when needed), when you change the definition of the primary key, all secondary indexes are recreated as well.

Dropping a secondary index is simple. Only the internal `InnoDB` system tables and the MySQL data dictionary tables are updated to reflect the fact that the index no longer exists. `InnoDB` returns the storage used for the index to the tablespace that contained it, so that new indexes or additional table rows can use the space.

To add a secondary index to an existing table, `InnoDB` scans the table, and sorts the rows using memory buffers and temporary files in order by the values of the secondary index key columns. The B-tree is then built in key-value order, which is more efficient than inserting rows into an index in random order. Because the B-tree nodes are split when they fill, building the index in this way results in a higher fill-factor for the index, making it more efficient for subsequent access.

## Primary Key and Secondary Key Indexes

Historically, the MySQL server and `InnoDB` have each kept their own metadata about table and index structures. The MySQL server stores this information in `.frm files` that are not protected by a transactional mechanism, while `InnoDB` has its own `data dictionary` as part of the `system tablespace`. If a DDL operation was interrupted by a crash or other unexpected event partway through, the metadata could be left inconsistent between these two locations, causing problems such as startup errors or inability to access the table that was being altered. Now that `InnoDB` is the default storage engine, addressing such issues is a high priority. These enhancements to DDL operations reduce the window of opportunity for such issues to occur.

### 5.5.7. How Crash Recovery Works with Online DDL

Although no data is lost if the server crashes while an `ALTER TABLE` statement is executing, the `crash recovery` process is different for `clustered indexes` and `secondary indexes`.

If the server crashes while creating an InnoDB secondary index, upon recovery, MySQL drops any partially created indexes. You must re-run the `ALTER TABLE` or `CREATE INDEX` statement.

When a crash occurs during the creation of an InnoDB clustered index, recovery is more complicated, because the data in the table must be copied to an entirely new clustered index. Remember that all InnoDB tables are stored as clustered indexes. In the following discussion, we use the word table and clustered index interchangeably.

MySQL creates the new clustered index by copying the existing data from the original InnoDB table to a temporary table that has the desired index structure. Once the data is completely copied to this temporary table, the original table is renamed with a different temporary table name. The temporary table comprising the new clustered index is renamed with the name of the original table, and the original table is dropped from the database.

If a system crash occurs while creating a new clustered index, no data is lost, but you must complete the recovery process using the temporary tables that exist during the process. Since it is rare to re-create a clustered index or re-define primary keys on large tables, or to encounter a system crash during this operation, this manual does not provide information on recovering from this scenario.

### 5.5.8. DDL for Partitioned Tables

The mechanism used for DDL operations on partitioned tables is somewhat different from the online DDL discussed throughout this section. The partitioning clauses of the `ALTER TABLE` statement are only allowed in conjunction with `ALGORITHM=DEFAULT` and `LOCK=DEFAULT`. Although these operations do not go through the same internal online DDL API as for non-partitioned tables, MySQL still attempts to minimize data copying and locking where possible:

- `ADD PARTITION` and `DROP PARTITION` for tables partitioned by `RANGE` or `LIST` do not copy any existing data.
- `TRUNCATE PARTITION` does not copy any existing data, for all types of partitioned tables.
- Concurrent queries are allowed during `ADD PARTITION` and `COALESCE PARTITION` for tables partitioned by `HASH` or `LIST`. MySQL copies the data while holding a shared lock.
- For `REORGANIZE PARTITION`, `REBUILD PARTITION`, or `ADD PARTITION` or `COALESCE PARTITION` for a table partitioned by `LINEAR HASH` or `LIST`, concurrent queries are allowed. Data from the affected partitions is copied while holding a shared lock.

### 5.5.9. Limitations of Online DDL

Take the following considerations into account when creating or dropping InnoDB indexes:

- During an online DDL operation that copies the table, files are written to the temporary directory (`$TMPDIR` on Unix, `%TEMP%` on Windows, or the directory specified by the `--tmpdir` [559] configuration variable). Each temporary file is large enough to hold one column in the new table or index, and each one is removed as soon as it is merged into the final table or index.
- The table is copied, rather than using Fast Index Creation when you create an index on a `TEMPORARY TABLE`. This has been reported as MySQL Bug #39833.
- InnoDB handles error cases when users attempt to drop indexes needed for foreign keys. See section [Section 14.2.5.9, “Better Error Handling when Dropping Indexes”](#) for details.
- The `ALTER TABLE` clause `LOCK=NONE` is not allowed if there are `ON . . . CASCADE` or `ON . . . SET NULL` constraints on the table.
- During each online DDL `ALTER TABLE` statement, regardless of the `LOCK` clause, there are brief periods at the beginning and end requiring an **exclusive lock** on the table (the same kind of lock specified by the `LOCK=EXCLUSIVE` clause). Thus, an online DDL operation might wait before starting



if there is a long-running transaction performing inserts, updates, deletes, or `SELECT ... FOR UPDATE` on that table; and an online DDL operation might wait before finishing if a similar long-running transaction was started while the `ALTER TABLE` was in progress.

## 5.6. Running Multiple MySQL Instances on One Machine

In some cases, you might want to run multiple instances of MySQL on a single machine. You might want to test a new MySQL release while leaving an existing production setup undisturbed. Or you might want to give different users access to different `mysqld` servers that they manage themselves. (For example, you might be an Internet Service Provider that wants to provide independent MySQL installations for different customers.)

It is possible to use a different MySQL server binary per instance, or use the same binary for multiple instances, or any combination of the two approaches. For example, you might run a server from MySQL 5.5 and one from MySQL 5.6, to see how different versions handle a given workload. Or you might run multiple instances of the current production version, each managing a different set of databases.

Whether or not you use distinct server binaries, each instance that you run must be configured with unique values for several operating parameters. This eliminates the potential for conflict between instances. Parameters can be set on the command line, in option files, or by setting environment variables. See [Section 4.2.3, “Specifying Program Options”](#). To see the values used by a given instance, connect to it and execute a `SHOW VARIABLES` statement.

The primary resource managed by a MySQL instance is the data directory. Each instance should use a different data directory, the location of which is specified using the `--datadir=path` [418] option. For methods of configuring each instance with its own data directory, and warnings about the dangers of failing to do so, see [Section 5.6.1, “Setting Up Multiple Data Directories”](#).

In addition to using different data directories, several other options must have different values for each server instance:

- `--port=port_num` [437]

`--port` [437] controls the port number for TCP/IP connections. Alternatively, if the host has multiple network addresses, you can use `--bind-address` [414] to cause each server to listen to a different address.

- `--socket=path` [442]

`--socket` [442] controls the Unix socket file path on Unix or the named pipe name on Windows. On Windows, it is necessary to specify distinct pipe names only for those servers configured to permit named-pipe connections.

- `--shared-memory-base-name=name` [439]

This option is used only on Windows. It designates the shared-memory name used by a Windows server to permit clients to connect using shared memory. It is necessary to specify distinct shared-memory names only for those servers configured to permit shared-memory connections.

- `--pid-file=file_name` [435]

This option indicates the path name of the file in which the server writes its process ID.

If you use the following log file options, their values must differ for each server:

- `--general_log_file=file_name` [486]
- `--log-bin[=file_name]` [2002]
- `--slow_query_log_file=file_name` [541]



- `--log-error[=file_name]` [427]

For further discussion of log file options, see [Section 5.2, “MySQL Server Logs”](#).

To achieve better performance, you can specify the following option differently for each server, to spread the load between several physical disks:

- `--tmpdir=path` [446]

Having different temporary directories also makes it easier to determine which MySQL server created any given temporary file.

If you have multiple MySQL installations in different locations, you can specify the base directory for each installation with the `--basedir=path` [414] option. This causes each instance to automatically use a different data directory, log files, and PID file because the default for each of those parameters is relative to the base directory. In that case, the only other options you need to specify are the `--socket` [442] and `--port` [437] options. Suppose that you install different versions of MySQL using `tar` file binary distributions. These install in different locations, so you can start the server for each installation using the command `bin/mysqld_safe` under its corresponding base directory. `mysqld_safe` determines the proper `--basedir` [414] option to pass to `mysqld`, and you need specify only the `--socket` [229] and `--port` [228] options to `mysqld_safe`.

As discussed in the following sections, it is possible to start additional servers by specifying appropriate command options or by setting environment variables. However, if you need to run multiple servers on a more permanent basis, it is more convenient to use option files to specify for each server those option values that must be unique to it. The `--defaults-file` [219] option is useful for this purpose.

### 5.6.1. Setting Up Multiple Data Directories

Each MySQL Instance on a machine should have its own data directory. The location is specified using the `--datadir=path` [418] option.

There are different methods of setting up a data directory for a new instance:

- Create a new data directory.
- Copy an existing data directory.

The following discussion provides more detail about each method.



#### Warning

Normally, you should never have two servers that update data in the same databases. This may lead to unpleasant surprises if your operating system does not support fault-free system locking. If (despite this warning) you run multiple servers using the same data directory and they have logging enabled, you must use the appropriate options to specify log file names that are unique to each server. Otherwise, the servers try to log to the same files.

Even when the preceding precautions are observed, this kind of setup works only with `MyISAM` and `MERGE` tables, and not with any of the other storage engines. Also, this warning against sharing a data directory among servers always applies in an NFS environment. Permitting multiple MySQL servers to access a common data directory over NFS is a *very bad idea*. The primary problem is that NFS is the speed bottleneck. It is not meant for such use. Another risk with NFS is that you must devise a way to ensure that two or more servers do not interfere with each other. Usually NFS file locking is handled by the `lockd` daemon, but at the moment there is no platform that performs locking 100% reliably in every situation.

## Create a New Data Directory

With this method, the data directory will be in the same state as when you first install MySQL. It will have the default set of MySQL accounts and no user data.

On Unix, initialize the data directory by running `mysql_install_db`. See [Section 2.10.1, “Unix Postinstallation Procedures”](#).

On Windows, the data directory is included in the MySQL distribution:

- MySQL Zip archive distributions for Windows contain an unmodified data directory. You can unpack such a distribution into a temporary location, then copy it `data` directory to where you are setting up the new instance.
- Windows MSI package installers create and set up the data directory that the installed server will use, but also create a pristine “template” data directory named `data` under the installation directory. After an installation has been performed using an MSI package, the template data directory can be copied to set up additional MySQL instances.

## Copy an Existing Data Directory

With this method, any MySQL accounts or user data present in the data directory are carried over to the new data directory.

1. Stop the existing MySQL instance using the data directory. This must be a clean shutdown so that the instance flushes any pending changes to disk.
2. Copy the data directory to the location where the new data directory should be.
3. Copy the `my.cnf` or `my.ini` option file used by the existing instance. This serves as a basis for the new instance.
4. Modify the new option file so that any pathnames referring to the original data directory refer to the new data directory. Also, modify any other options that must be unique per instance, such as the TCP/IP port number and the log files. For a list of parameters that must be unique per instance, see [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).
5. Start the new instance, telling it to use the new option file.

## 5.6.2. Running Multiple MySQL Instances on Windows

You can run multiple servers on Windows by starting them manually from the command line, each with appropriate operating parameters, or by installing several servers as Windows services and running them that way. General instructions for running MySQL from the command line or as a service are given in [Section 2.3, “Installing MySQL on Microsoft Windows”](#). The following sections describe how to start each server with different values for those options that must be unique per server, such as the data directory. These options are listed in [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#).

### 5.6.2.1. Starting Multiple MySQL Instances at the Windows Command Line

The procedure for starting a single MySQL server manually from the command line is described in [Section 2.3.4.5, “Starting MySQL from the Windows Command Line”](#). To start multiple servers this way, you can specify the appropriate options on the command line or in an option file. It is more convenient to place the options in an option file, but it is necessary to make sure that each server gets its own set of options. To do this, create an option file for each server and tell the server the file name with a `--defaults-file [219]` option when you run it.

Suppose that you want to run `mysqld` on port 3307 with a data directory of `C:\mydata1`, and `mysqld-debug` on port 3308 with a data directory of `C:\mydata2`. Use this procedure:

1. Make sure that each data directory exists, including its own copy of the `mysql` database that contains the grant tables.
2. Create two option files. For example, create one file named `C:\my-opts1.cnf` that looks like this:

```
[mysqld]
datadir = C:/mydata1
port = 3307
```

Create a second file named `C:\my-opts2.cnf` that looks like this:

```
[mysqld]
datadir = C:/mydata2
port = 3308
```

3. Use the `--defaults-file` [219] option to start each server with its own option file:

```
C:\> C:\mysql\bin\mysqld --defaults-file=C:\my-opts1.cnf
C:\> C:\mysql\bin\mysqld-debug --defaults-file=C:\my-opts2.cnf
```

Each server starts in the foreground (no new prompt appears until the server exits later), so you will need to issue those two commands in separate console windows.

To shut down the servers, connect to each using the appropriate port number:

```
C:\> C:\mysql\bin\mysqladmin --port=3307 shutdown
C:\> C:\mysql\bin\mysqladmin --port=3308 shutdown
```

Servers configured as just described permit clients to connect over TCP/IP. If your version of Windows supports named pipes and you also want to permit named-pipe connections, use the `mysqld` or `mysqld-debug` server and specify options that enable the named pipe and specify its name. Each server that supports named-pipe connections must use a unique pipe name. For example, the `C:\my-opts1.cnf` file might be written like this:

```
[mysqld]
datadir = C:/mydata1
port = 3307
enable-named-pipe
socket = mypipe1
```

Modify `C:\my-opts2.cnf` similarly for use by the second server. Then start the servers as described previously.

A similar procedure applies for servers that you want to permit shared-memory connections. Enable such connections with the `--shared-memory` [439] option and specify a unique shared-memory name for each server with the `--shared-memory-base-name` [439] option.

### 5.6.2.2. Starting Multiple MySQL Instances as Windows Services

On Windows, a MySQL server can run as a Windows service. The procedures for installing, controlling, and removing a single MySQL service are described in [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

To set up multiple MySQL services, you must make sure that each instance uses a different service name in addition to the other parameters that must be unique per instance.

For the following instructions, suppose that you want to run the `mysqld` server from two different versions of MySQL that are installed at `C:\mysql-5.5.9` and `C:\mysql-5.6.12`, respectively. (This might be the case if you are running 5.5.9 as your production server, but also want to conduct tests using 5.6.12.)

To install MySQL as a Windows service, use the `--install` or `--install-manual` option. For information about these options, see [Section 2.3.4.8, “Starting MySQL as a Windows Service”](#).

Based on the preceding information, you have several ways to set up multiple services. The following instructions describe some examples. Before trying any of them, shut down and remove any existing MySQL services.

- **Approach 1:** Specify the options for all services in one of the standard option files. To do this, use a different service name for each server. Suppose that you want to run the 5.5.9 `mysqld` using the service name of `mysqld1` and the 5.6.12 `mysqld` using the service name `mysqld2`. In this case, you can use the `[mysqld1]` group for 5.5.9 and the `[mysqld2]` group for 5.6.12. For example, you can set up `C:\my.cnf` like this:

```
# options for mysqld1 service
[mysqld1]
basedir = C:/mysql-5.5.9
port = 3307
enable-named-pipe
socket = mypipe1

# options for mysqld2 service
[mysqld2]
basedir = C:/mysql-5.6.12
port = 3308
enable-named-pipe
socket = mypipe2
```

Install the services as follows, using the full server path names to ensure that Windows registers the correct executable program for each service:

```
C:\> C:\mysql-5.5.9\bin\mysqld --install mysqld1
C:\> C:\mysql-5.6.12\bin\mysqld --install mysqld2
```

To start the services, use the services manager, or use `NET START` with the appropriate service names:

```
C:\> NET START mysqld1
C:\> NET START mysqld2
```

To stop the services, use the services manager, or use `NET STOP` with the appropriate service names:

```
C:\> NET STOP mysqld1
C:\> NET STOP mysqld2
```

- **Approach 2:** Specify options for each server in separate files and use `--defaults-file` [219] when you install the services to tell each server what file to use. In this case, each file should list options using a `[mysqld]` group.

With this approach, to specify options for the 5.5.9 `mysqld`, create a file `C:\my-opts1.cnf` that looks like this:

```
[mysqld]
basedir = C:/mysql-5.5.9
port = 3307
enable-named-pipe
socket = mypipe1
```

For the 5.6.12 `mysqld`, create a file `C:\my-opts2.cnf` that looks like this:

```
[mysqld]
basedir = C:/mysql-5.6.12
port = 3308
enable-named-pipe
socket = mypipe2
```

Install the services as follows (enter each command on a single line):

```
C:\> C:\mysql-5.5.9\bin\mysqld --install mysqld1
--defaults-file=C:\my-opts1.cnf
```

```
C:\> C:\mysql-5.6.12\bin\mysqld --install mysqld2
      --defaults-file=C:\my-opts2.cnf
```

When you install a MySQL server as a service and use a `--defaults-file` [219] option, the service name must precede the option.

After installing the services, start and stop them the same way as in the preceding example.

To remove multiple services, use `mysqld --remove` for each one, specifying a service name following the `--remove` [437] option. If the service name is the default (MySQL), you can omit it.

### 5.6.3. Running Multiple MySQL Instances on Unix

One way to run multiple MySQL instances on Unix is to compile different servers with different default TCP/IP ports and Unix socket files so that each one listens on different network interfaces. Compiling in different base directories for each installation also results automatically in a separate, compiled-in data directory, log file, and PID file location for each server.

Assume that an existing 5.5 server is configured for the default TCP/IP port number (3306) and Unix socket file (`/tmp/mysql.sock`). To configure a new 5.6.12 server to have different operating parameters, use a CMake command something like this:

```
shell> cmake . -DMYSQL_TCP_PORT=port_number \
           -DMYSQL_UNIX_ADDR=file_name \
           -DCMAKE_INSTALL_PREFIX=/usr/local/mysql-5.6.12
```

Here, `port_number` and `file_name` must be different from the default TCP/IP port number and Unix socket file path name, and the `CMAKE_INSTALL_PREFIX` [132] value should specify an installation directory different from the one under which the existing MySQL installation is located.

If you have a MySQL server listening on a given port number, you can use the following command to find out what operating parameters it is using for several important configurable variables, including the base directory and Unix socket file name:

```
shell> mysqladmin --host=host_name --port=port_number variables
```

With the information displayed by that command, you can tell what option values *not* to use when configuring an additional server.

If you specify `localhost` as the host name, `mysqladmin` defaults to using a Unix socket file connection rather than TCP/IP. To explicitly specify the connection protocol, use the `--protocol={TCP|SOCKET|PIPE|MEMORY}` [211] option.

You need not compile a new MySQL server just to start with a different Unix socket file and TCP/IP port number. It is also possible to use the same server binary and start each invocation of it with different parameter values at runtime. One way to do so is by using command-line options:

```
shell> mysqld_safe --socket=file_name --port=port_number
```

To start a second server, provide different `--socket` [442] and `--port` [437] option values, and pass a `--datadir=path` [418] option to `mysqld_safe` so that the server uses a different data directory.

Alternatively, put the options for each server in a different option file, then start each server using a `--defaults-file` [219] option that specifies the path to the appropriate option file. For example, if the option files for two server instances are named `/usr/local/mysql/my.cnf` and `/usr/local/mysql/my.cnf2`, start the servers like this: command:

```
shell> mysqld_safe --defaults-file=/usr/local/mysql/my.cnf
shell> mysqld_safe --defaults-file=/usr/local/mysql/my.cnf2
```

Another way to achieve a similar effect is to use environment variables to set the Unix socket file name and TCP/IP port number:

```

shell> MYSQL_UNIX_PORT=/tmp/mysql-new.sock
shell> MYSQL_TCP_PORT=3307
shell> export MYSQL_UNIX_PORT MYSQL_TCP_PORT
shell> mysql_install_db --user=mysql
shell> mysqld_safe --datadir=/path/to/datadir &

```

This is a quick way of starting a second server to use for testing. The nice thing about this method is that the environment variable settings apply to any client programs that you invoke from the same shell. Thus, connections for those clients are automatically directed to the second server.

[Section 2.12, “Environment Variables”](#), includes a list of other environment variables you can use to affect MySQL programs.

On Unix, the `mysqld_multi` script provides another way to start multiple servers. See [Section 4.3.4, “mysqld\\_multi — Manage Multiple MySQL Servers”](#).

## 5.6.4. Using Client Programs in a Multiple-Server Environment

To connect with a client program to a MySQL server that is listening to different network interfaces from those compiled into your client, you can use one of the following methods:

- Start the client with `--host=host_name [211] --port=port_number [211]` to connect using TCP/IP to a remote server, with `--host=127.0.0.1 [211] --port=port_number [211]` to connect using TCP/IP to a local server, or with `--host=localhost [211] --socket=file_name [211]` to connect to a local server using a Unix socket file or a Windows named pipe.
- Start the client with `--protocol=TCP [211]` to connect using TCP/IP, `--protocol=SOCKET [211]` to connect using a Unix socket file, `--protocol=PIPE [211]` to connect using a named pipe, or `--protocol=MEMORY [211]` to connect using shared memory. For TCP/IP connections, you may also need to specify `--host [211]` and `--port [211]` options. For the other types of connections, you may need to specify a `--socket [211]` option to specify a Unix socket file or Windows named-pipe name, or a `--shared-memory-base-name [211]` option to specify the shared-memory name. Shared-memory connections are supported only on Windows.
- On Unix, set the `MYSQL_UNIX_PORT` and `MYSQL_TCP_PORT` environment variables to point to the Unix socket file and TCP/IP port number before you start your clients. If you normally use a specific socket file or port number, you can place commands to set these environment variables in your `.login` file so that they apply each time you log in. See [Section 2.12, “Environment Variables”](#).
- Specify the default Unix socket file and TCP/IP port number in the `[client]` group of an option file. For example, you can use `C:\my.cnf` on Windows, or the `.my.cnf` file in your home directory on Unix. See [Section 4.2.3.3, “Using Option Files”](#).
- In a C program, you can specify the socket file or port number arguments in the `mysql_real_connect()` call. You can also have the program read option files by calling `mysql_options()`. See [Section 21.8.6, “C API Function Descriptions”](#).
- If you are using the Perl `DBD: :mysql` module, you can read options from MySQL option files. For example:

```

$dsn = "DBI:mysql:test;mysql_read_default_group=client;"
      . "mysql_read_default_file=/usr/local/mysql/data/my.cnf";
$dbh = DBI->connect($dsn, $user, $password);

```

See [Section 21.10, “MySQL Perl API”](#).

Other programming interfaces may provide similar capabilities for reading option files.

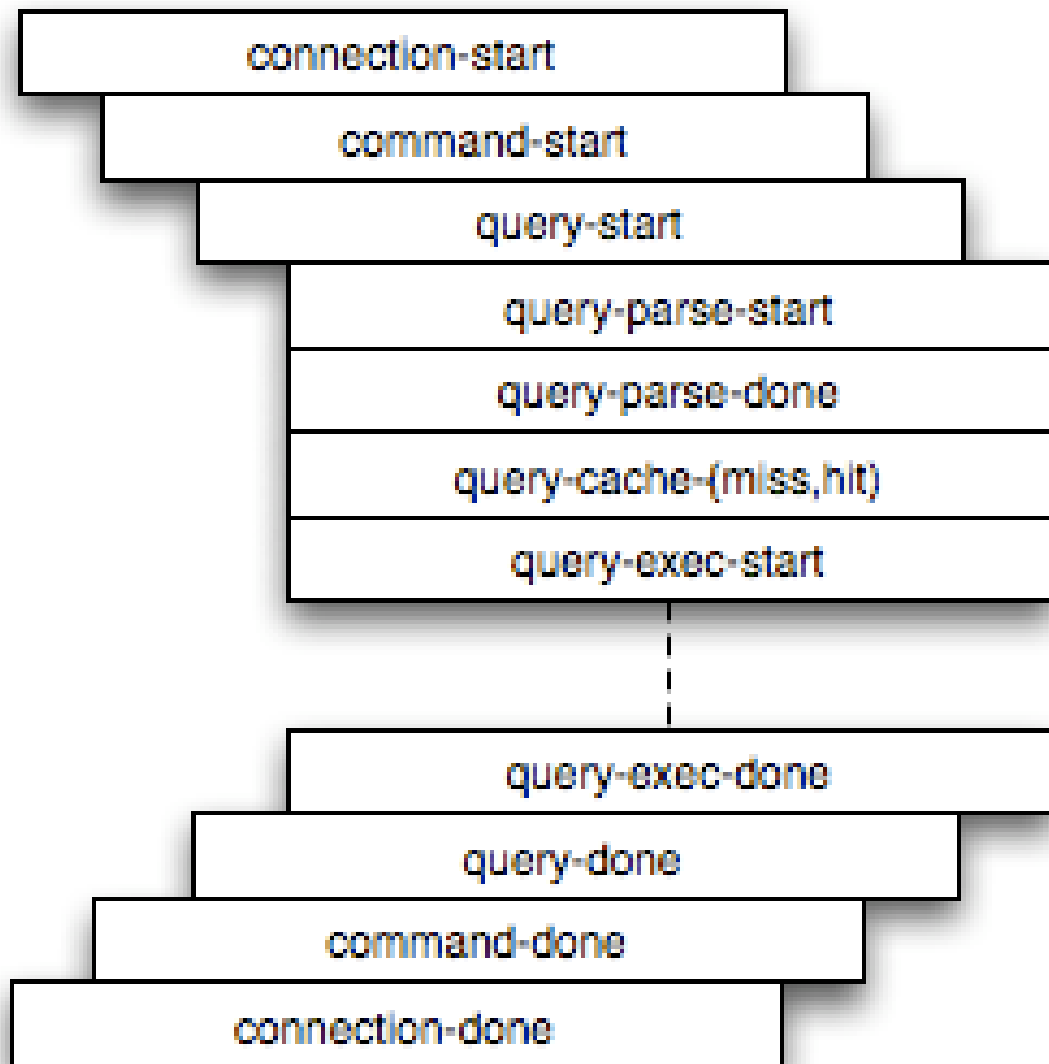
## 5.7. Tracing `mysqld` Using DTrace

The DTrace probes in the MySQL server are designed to provide information about the execution of queries within MySQL and the different areas of the system being utilized during that process. The

organization and triggering of the probes means that the execution of an entire query can be monitored with one level of probes (`query-start` and `query-done`) but by monitoring other probes you can get successively more detailed information about the execution of the query in terms of the locks used, sort methods and even row-by-row and storage-engine level execution information.

The DTrace probes are organized so that you can follow the entire query process, from the point of connection from a client, through the query execution, row-level operations, and back out again. You can think of the probes as being fired within a specific sequence during a typical client connect/execute/disconnect sequence, as shown in the following figure.

**Figure 5.1. The MySQL Architecture Using Pluggable Storage Engines**



Global information is provided in the arguments to the DTrace probes at various levels. Global information, that is, the connection ID and user/host and where relevant the query string, is provided at key levels (`connection-start`, `command-start`, `query-start`, and `query-exec-start`). As you go deeper into the probes, it is assumed either you are only interested in the individual executions (row-level probes provide information on the database and table name only), or that you will combine the row-level probes with the notional parent probes to provide the information about a specific query. Examples of this will be given as the format and arguments of each probe are provided.



For more information on DTrace and writing DTrace scripts, read the [DTrace User Guide](#).

MySQL 5.6 includes support for DTrace probes on Solaris 10 Update 5 (Solaris 5/08) on SPARC, x86 and x86\_64 platforms. Probes are also supported on Mac OS X 10.4 and higher. Enabling the probes should be automatic on these platforms. To explicitly enable or disable the probes during building, use the `-DENABLE_DTRACE=1` [135] or `-DENABLE_DTRACE=0` [135] option to CMake.

### 5.7.1. mysqlld DTrace Probe Reference

MySQL supports the following static probes, organized into groups of functionality.

**Table 5.10. MySQL DTrace Probes**

Group	Probes	Introduced
Connection	<code>connection-start</code> , <code>connection-done</code>	5.4.0
Command	<code>command-start</code> , <code>command-done</code>	5.4.0
Query	<code>query-start</code> , <code>query-done</code>	5.4.0
Query Parsing	<code>query-parse-start</code> , <code>query-parse-done</code>	5.4.0
Query Cache	<code>query-cache-hit</code> , <code>query-cache-miss</code>	5.4.0
Query Execution	<code>query-exec-start</code> , <code>query-exec-done</code>	5.4.0
Row Level	<code>insert-row-start</code> , <code>insert-row-done</code>	5.4.0
	<code>update-row-start</code> , <code>update-row-done</code>	5.4.0
	<code>delete-row-start</code> , <code>delete-row-done</code>	5.4.0
Row Reads	<code>read-row-start</code> , <code>read-row-done</code>	5.4.0
Index Reads	<code>index-read-row-start</code> , <code>index-read-row-done</code>	5.4.0
Lock	<code>handler-rdlock-start</code> , <code>handler-rdlock-done</code>	5.4.0
	<code>handler-wrlock-start</code> , <code>handler-wrlock-done</code>	5.4.0
	<code>handler-unlock-start</code> , <code>handler-unlock-done</code>	5.4.0
Filesort	<code>filesort-start</code> , <code>filesort-done</code>	5.4.0
Statement	<code>select-start</code> , <code>select-done</code>	5.4.0
	<code>insert-start</code> , <code>insert-done</code>	5.4.0
	<code>insert-select-start</code> , <code>insert-select-done</code>	5.4.0
	<code>update-start</code> , <code>update-done</code>	5.4.0
	<code>multi-update-start</code> , <code>multi-update-done</code>	5.4.0
	<code>delete-start</code> , <code>delete-done</code>	5.4.0
	<code>multi-delete-start</code> , <code>multi-delete-done</code>	5.4.0
Network	<code>net-read-start</code> , <code>net-read-done</code> , <code>net-write-start</code> , <code>net-write-done</code>	5.4.0
Keycache	<code>keycache-read-start</code> , <code>keycache-read-block</code> , <code>keycache-read-done</code> , <code>keycache-read-hit</code> , <code>keycache-read-miss</code> , <code>keycache-write-start</code> , <code>keycache-write-block</code> , <code>keycache-write-done</code>	5.4.0



#### Note

When extracting the argument data from the probes, each argument is available as `argN`, starting with `arg0`. To identify each argument within the definitions they are provided with a descriptive name, but you must access the information using the corresponding `argN` parameter.



### 5.7.1.1. Connection Probes

The `connection-start` and `connection-done` probes enclose a connection from a client, regardless of whether the connection is through a socket or network connection.

```
connection-start(connectionid, user, host)
connection-done(status, connectionid)
```

- `connection-start`: Triggered after a connection and successful login/authentication have been completed by a client. The arguments contain the connection information:
  - `connectionid`: An `unsigned long` containing the connection ID. This is the same as the process ID shown as the `Id` value in the output from `SHOW PROCESSLIST`.
  - `user`: The username used when authenticating. The value will be blank for the anonymous user.
  - `host`: The host of the client connection. For a connection made using UNIX sockets, the value will be blank.
- `connection-done`: Triggered just as the connection to the client has been closed. The arguments are:
  - `status`: The status of the connection when it was closed. A logout operation will have a value of 0; any other termination of the connection has a nonzero value.
  - `connectionid`: The connection ID of the connection that was closed.

The following D script will quantify and summarize the average duration of individual connections, and provide a count, dumping the information every 60 seconds:

```
#!/usr/sbin/dtrace -s

mysql*:::connection-start
{
    self->start = timestamp;
}

mysql*:::connection-done
/self->start/
{
    @ = quantize(((timestamp - self->start)/1000000));
    self->start = 0;
}

tick-60s
{
    printa(@);
}
```

When executed on a server with a large number of clients you might see output similar to this:

```
1 57413          :tick-60s

      value  ----- Distribution ----- count
      -1 | 0
      0 | @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@ 30011
      1 | 59
      2 | 5
      4 | 20
      8 | 29
     16 | 18
     32 | 27
     64 | 30
    128 | 11
    256 | 10
    512 | 1
```

1024	6
2048	8
4096	9
8192	8
16384	2
32768	1
65536	1
131072	0
262144	1
524288	0

### 5.7.1.2. Command Probes

The command probes are executed before and after a client command is executed, including any SQL statement that might be executed during that period. Commands include operations such as the initialization of the DB, use of the `COM_CHANGE_USER` operation (supported by the MySQL protocol), and manipulation of prepared statements. Many of these commands are used only by the MySQL client API from various connectors such as PHP and Java.

```
command-start(connectionid, command, user, host)
command-done(status)
```

- `command-start`: Triggered when a command is submitted to the server.
- `connectionid`: The connection ID of the client executing the command.
- `command`: An integer representing the command that was executed. Possible values are shown in the following table.

Value	Name	Description
00	COM_SLEEP	Internal thread state
01	COM_QUIT	Close connection
02	COM_INIT_DB	Select database ( <code>USE ...</code> )
03	COM_QUERY	Execute a query
04	COM_FIELD_LIST	Get a list of fields
05	COM_CREATE_DB	Create a database (deprecated)
06	COM_DROP_DB	Drop a database (deprecated)
07	COM_REFRESH	Refresh connection
08	COM_SHUTDOWN	Shutdown server
09	COM_STATISTICS	Get statistics
10	COM_PROCESS_INFO	Get processes ( <code>SHOW PROCESSLIST</code> )
11	COM_CONNECT	Initialize connection
12	COM_PROCESS_KILL	Kill process
13	COM_DEBUG	Get debug information
14	COM_PING	Ping
15	COM_TIME	Internal thread state
16	COM_DELAYED_INSERT	Internal thread state
17	COM_CHANGE_USER	Change user
18	COM_BINLOG_DUMP	Used by a replication slave or <code>mysqlbinlog</code> to initiate a binary log read
19	COM_TABLE_DUMP	Used by a replication slave to get the master table information
20	COM_CONNECT_OUT	Used by a replication slave to log a connection to the server
21	COM_REGISTER_SLAVE	Used by a replication slave during registration

Value	Name	Description
22	COM_STMT_PREPARE	Prepare a statement
23	COM_STMT_EXECUTE	Execute a statement
24	COM_STMT_SEND_LONG_DATA	Used when requesting extended data
25	COM_STMT_CLOSE	Close a prepared statement
26	COM_STMT_RESET	Reset a prepared statement
27	COM_SET_OPTIONS	Set a server option
28	COM_STMT_FETCH	Fetch a prepared statement

- `user`: The user executing the command.
- `host`: The client host.
- `command-done`: Triggered when the command execution completes. The `status` argument contains 0 if the command executed successfully, or 1 if the statement was terminated before normal completion.

The `command-start` and `command-done` probes are best used when combined with the statement probes to get an idea of overall execution time.

### 5.7.1.3. Query Probes

The `query-start` and `query-done` probes are triggered when a specific query is received by the server and when the query has been completed and the information has been successfully sent to the client.

```
query-start(query, connectionid, database, user, host)
query-done(status)
```

- `query-start`: Triggered after the query string has been received from the client. The arguments are:
  - `query`: The full text of the submitted query.
  - `connectionid`: The connection ID of the client that submitted the query. The connection ID equals the connection ID returned when the client first connects and the `Id` value in the output from `SHOW PROCESSLIST`.
  - `database`: The database name on which the query is being executed.
  - `user`: The username used to connect to the server.
  - `host`: The hostname of the client.
- `query-done`: Triggered once the query has been executed and the information has been returned to the client. The probe includes a single argument, `status`, which returns 0 when the query is successfully executed and 1 if there was an error.

You can get a simple report of the execution time for each query using the following D script:

```
#!/usr/sbin/dtrace -s
#pragma D option quiet
dtrace:::BEGIN
{
    printf("%-20s %-20s %-40s %-9s\n", "Who", "Database", "Query", "Time(ms)");
}
mysql*:::query-start
```

```

{
    self->query = copyinstr(arg0);
    self->connid = arg1;
    self->db     = copyinstr(arg2);
    self->who    = strjoin(copyinstr(arg3), strjoin("@", copyinstr(arg4)));
    self->querystart = timestamp;
}

mysql*:::query-done
{
    printf("%-20s %-20s %-40s %-9d\n", self->who, self->db, self->query,
        (timestamp - self->querystart) / 1000000);
}

```

When executing the above script you should get a basic idea of the execution time of your queries:

```

shell> ./query.d
Who           Database      Query                                               Time(ms)
root@localhost test         select * from t1 order by i limit 10              0
root@localhost test         set global query_cache_size=0                    0
root@localhost test         select * from t1 order by i limit 10              776
root@localhost test         select * from t1 order by i limit 10              773
root@localhost test         select * from t1 order by i desc limit 10        795

```

#### 5.7.1.4. Query Parsing Probes

The query parsing probes are triggered before the original SQL statement is parsed and when the parsing of the statement and determination of the execution model required to process the statement has been completed:

```

query-parse-start(query)
query-parse-done(status)

```

- **query-parse-start**: Triggered just before the statement is parsed by the MySQL query parser. The single argument, **query**, is a string containing the full text of the original query.
- **query-parse-done**: Triggered when the parsing of the original statement has been completed. The **status** is an integer describing the status of the operation. A **0** indicates that the query was successfully parsed. A **1** indicates that the parsing of the query failed.

For example, you could monitor the execution time for parsing a given query using the following D script:

```

#!/usr/sbin/dtrace -s

#pragma D option quiet

mysql*:::query-parse-start
{
    self->parsestart = timestamp;
    self->parsequery = copyinstr(arg0);
}

mysql*:::query-parse-done
/arg0 == 0/
{
    printf("Parsing %s: %d microseconds\n", self->parsequery, ((timestamp - self->parsestart)/1000));
}

mysql*:::query-parse-done
/arg0 != 0/
{
    printf("Error parsing %s: %d microseconds\n", self->parsequery, ((timestamp - self->parsestart)/1000));
}

```

In the above script a predicate is used on **query-parse-done** so that different output is generated based on the status value of the probe.

When running the script and monitoring the execution:

```
shell> ./query-parsing.d
Error parsing select from t1 join (t2) on (t1.i = t2.i) order by t1.s,t1.i limit 10: 36 ms
Parsing select * from t1 join (t2) on (t1.i = t2.i) order by t1.s,t1.i limit 10: 176 ms
```

### 5.7.1.5. Query Cache Probes

The query cache probes are fired when executing any query. The `query-cache-hit` query is triggered when a query exists in the query cache and can be used to return the query cache information. The arguments contain the original query text and the number of rows returned from the query cache for the query. If the query is not within the query cache, or the query cache is not enabled, then the `query-cache-miss` probe is triggered instead.

```
query-cache-hit(query, rows)
query-cache-miss(query)
```

- `query-cache-hit`: Triggered when the query has been found within the query cache. The first argument, `query`, contains the original text of the query. The second argument, `rows`, is an integer containing the number of rows in the cached query.
- `query-cache-miss`: Triggered when the query is not found within the query cache. The first argument, `query`, contains the original text of the query.

The query cache probes are best combined with a probe on the main query so that you can determine the differences in times between using or not using the query cache for specified queries. For example, in the following D script, the query and query cache information are combined into the information output during monitoring:

```
#!/usr/sbin/dtrace -s

#pragma D option quiet

dtrace::BEGIN
{
    printf("%-20s %-20s %-40s %2s %-9s\n", "Who", "Database", "Query", "QC", "Time(ms)");
}

mysql*:::query-start
{
    self->query = copyinstr(arg0);
    self->connid = arg1;
    self->db = copyinstr(arg2);
    self->who = strjoin(copyinstr(arg3),strjoin("@",copyinstr(arg4)));
    self->querystart = timestamp;
    self->qc = 0;
}

mysql*:::query-cache-hit
{
    self->qc = 1;
}

mysql*:::query-cache-miss
{
    self->qc = 0;
}

mysql*:::query-done
{
    printf("%-20s %-20s %-40s %-2s %-9d\n",self->who,self->db,self->query,(self->qc ? "Y" : "N"),
        (timestamp - self->querystart) / 1000000);
}
```

When executing the script you can see the effects of the query cache. Initially the query cache is disabled. If you set the query cache size and then execute the query multiple times you should see that the query cache is being used to return the query data:

```
shell> ./query-cache.d
root@localhost      test                select * from t1 order by i limit 10      N 1072
```

```

root@localhost          set global query_cache_size=262144      N  0
root@localhost test      select * from t1 order by i limit 10      N 781
root@localhost test      select * from t1 order by i limit 10      Y  0

```

### 5.7.1.6. Query Execution Probes

The query execution probe is triggered when the actual execution of the query starts, after the parsing and checking the query cache but before any privilege checks or optimization. By comparing the difference between the start and done probes you can monitor the time actually spent servicing the query (instead of just handling the parsing and other elements of the query).

```

query-exec-start(query, connectionid, database, user, host, exec_type)
query-exec-done(status)

```



#### Note

The information provided in the arguments for `query-start` and `query-exec-start` are almost identical and designed so that you can choose to monitor either the entire query process (using `query-start`) or only the execution (using `query-exec-start`) while exposing the core information about the user, client, and query being executed.

- `query-exec-start`: Triggered when the execution of a individual query is started. The arguments are:
  - `query`: The full text of the submitted query.
  - `connectionid`: The connection ID of the client that submitted the query. The connection ID equals the connection ID returned when the client first connects and the `Id` value in the output from `SHOW PROCESSLIST`.
  - `database`: The database name on which the query is being executed.
  - `user`: The username used to connect to the server.
  - `host`: The hostname of the client.
  - `exec_type`: The type of execution. Execution types are determined based on the contents of the query and where it was submitted. The values for each type are shown in the following table.

Value	Description
0	Executed query from <code>sql_parse</code> , top-level query.
1	Executed prepared statement
2	Executed cursor statement
3	Executed query in stored procedure

- `query-exec-done`: Triggered when the execution of the query has completed. The probe includes a single argument, `status`, which returns 0 when the query is successfully executed and 1 if there was an error.

### 5.7.1.7. Row-Level Probes

The `*row-{start,done}` probes are triggered each time a row operation is pushed down to a storage engine. For example, if you execute an `INSERT` statement with 100 rows of data, then the `insert-row-start` and `insert-row-done` probes will be triggered 100 times each, for each row insert.

```

insert-row-start(database, table)
insert-row-done(status)

```

```

update-row-start(database, table)
update-row-done(status)

delete-row-start(database, table)
delete-row-done(status)

```

- `insert-row-start`: Triggered before a row is inserted into a table.
- `insert-row-done`: Triggered after a row is inserted into a table.
- `update-row-start`: Triggered before a row is updated in a table.
- `update-row-done`: Triggered before a row is updated in a table.
- `delete-row-start`: Triggered before a row is deleted from a table.
- `delete-row-done`: Triggered before a row is deleted from a table.

The arguments supported by the probes are consistent for the corresponding `start` and `done` probes in each case:

- `database`: The database name.
- `table`: The table name.
- `status`: The status; 0 for success or 1 for failure.

Because the row-level probes are triggered for each individual row access, these probes can be triggered many thousands of times each second, which may have a detrimental effect on both the monitoring script and MySQL. The DTrace environment should limit the triggering on these probes to prevent the performance being adversely affected. Either use the probes sparingly, or use counter or aggregation functions to report on these probes and then provide a summary when the script terminates or as part of a `query-done` or `query-exec-done` probes.

The following example script summarizes the duration of each row operation within a larger query:

```

#!/usr/sbin/dtrace -s

#pragma D option quiet

dtrace::BEGIN
{
    printf("%-2s %-10s %-10s %9s %9s %-s \n",
          "St", "Who", "DB", "ConnID", "Dur ms", "Query");
}

mysql*:::query-start
{
    self->query = copyinstr(arg0);
    self->who   = strjoin(copyinstr(arg3),strjoin("@",copyinstr(arg4)));
    self->db    = copyinstr(arg2);
    self->connid = arg1;
    self->querystart = timestamp;
    self->rowdur = 0;
}

mysql*:::query-done
{
    this->elapsed = (timestamp - self->querystart) /1000000;
    printf("%2d %-10s %-10s %9d %9d %s\n",
          arg0, self->who, self->db,
          self->connid, this->elapsed, self->query);
}

mysql*:::query-done
/ self->rowdur /
{
    printf("%34s %9d %s\n", "", (self->rowdur/1000000), "-> Row ops");
}

```

```

}

mysql*:::insert-row-start
{
    self->rowstart = timestamp;
}

mysql*:::delete-row-start
{
    self->rowstart = timestamp;
}

mysql*:::update-row-start
{
    self->rowstart = timestamp;
}

mysql*:::insert-row-done
{
    self->rowdur += (timestamp-self->rowstart);
}

mysql*:::delete-row-done
{
    self->rowdur += (timestamp-self->rowstart);
}

mysql*:::update-row-done
{
    self->rowdur += (timestamp-self->rowstart);
}

```

Running the above script with a query that inserts data into a table, you can monitor the exact time spent performing the raw row insertion:

St	Who	DB	ConnID	Dur ms	Query
0	@localhost	test	13	20767	insert into t1(select * from t2)
				4827	-> Row ops

### 5.7.1.8. Read Row Probes

The read row probes are triggered at a storage engine level each time a row read operation occurs. These probes are specified within each storage engine (as opposed to the `*row-start` probes which are in the storage engine interface). These probes can therefore be used to monitor individual storage engine row-level operations and performance. Because these probes are triggered around the storage engine row read interface, they may be hit a significant number of times during a basic query.

```

read-row-start(database, table, scan_flag)
read-row-done(status)

```

- `read-row-start`: Triggered when a row is read by the storage engine from the specified `database` and `table`. The `scan_flag` is set to 1 (true) when the read is part of a table scan (that is, a sequential read), or 0 (false) when the read is of a specific record.
- `read-row-done`: Triggered when a row read operation within a storage engine completes. The `status` returns 0 on success, or a positive value on failure.

### 5.7.1.9. Index Probes

The index probes are triggered each time a a row is read using one of the indexes for the specified table. The probe is triggered within the corresponding storage engine for the table.

```

index-read-row-start(database, table)
index-read-row-done(status)

```

- `index-read-row-start`: Triggered when a row is read by the storage engine from the specified `database` and `table`.



- `index-read-row-done`: Triggered when an indexed row read operation within a storage engine completes. The `status` returns 0 on success, or a positive value on failure.

### 5.7.1.10. Lock Probes

The lock probes are called whenever an external lock is requested by MySQL for a table using the corresponding lock mechanism on the table as defined by the table's engine type. There are three different types of lock, the read lock, write lock, and unlock operations. Using the probes you can determine the duration of the external locking routine (that is, the time taken by the storage engine to implement the lock, including any time waiting for another lock to become free) and the total duration of the lock/unlock process.

```
handler-rdlock-start(database, table)
handler-rdlock-done(status)

handler-wrlock-start(database, table)
handler-wrlock-done(status)

handler-unlock-start(database, table)
handler-unlock-done(status)
```

- `handler-rdlock-start`: Triggered when a read lock is requested on the specified `database` and `table`.
- `handler-wrlock-start`: Triggered when a write lock is requested on the specified `database` and `table`.
- `handler-unlock-start`: Triggered when an unlock request is made on the specified `database` and `table`.
- `handler-rdlock-done`: Triggered when a read lock request completes. The `status` is 0 if the lock operation succeeded, or `>0` on failure.
- `handler-wrlock-done`: Triggered when a write lock request completes. The `status` is 0 if the lock operation succeeded, or `>0` on failure.
- `handler-unlock-done`: Triggered when an unlock request completes. The `status` is 0 if the unlock operation succeeded, or `>0` on failure.

You can use arrays to monitor the locking and unlocking of individual tables and then calculate the duration of the entire table lock using the following script:

```
#!/usr/sbin/dtrace -s

#pragma D option quiet

mysql*:::handler-rdlock-start
{
    self->rdlockstart = timestamp;
    this->lockref = strjoin(copyinstr(arg0), strjoin("@", copyinstr(arg1)));
    self->lockmap[this->lockref] = self->rdlockstart;
    printf("Start: Lock->Read   %s.%s\n", copyinstr(arg0), copyinstr(arg1));
}

mysql*:::handler-wrlock-start
{
    self->wrlockstart = timestamp;
    this->lockref = strjoin(copyinstr(arg0), strjoin("@", copyinstr(arg1)));
    self->lockmap[this->lockref] = self->rdlockstart;
    printf("Start: Lock->Write  %s.%s\n", copyinstr(arg0), copyinstr(arg1));
}

mysql*:::handler-unlock-start
{
    self->unlockstart = timestamp;
    this->lockref = strjoin(copyinstr(arg0), strjoin("@", copyinstr(arg1)));
    printf("Start: Lock->Unlock %s.%s (%d ms lock duration)\n",
```

```

        copyinstr(arg0),copyinstr(arg1),
        (timestamp - self->lockmap[this->lockref])/1000000);
    }

mysql*:::handler-rdlock-done
{
    printf("End:   Lock->Read   %d ms\n",
        (timestamp - self->rdlockstart)/1000000);
}

mysql*:::handler-wrlock-done
{
    printf("End:   Lock->Write  %d ms\n",
        (timestamp - self->wrlockstart)/1000000);
}

mysql*:::handler-unlock-done
{
    printf("End:   Lock->Unlock %d ms\n",
        (timestamp - self->unlockstart)/1000000);
}

```

When executed, you should get information both about the duration of the locking process itself, and of the locks on a specific table:

```

Start: Lock->Read   test.t2
End:   Lock->Read   0 ms
Start: Lock->Unlock test.t2 (25743 ms lock duration)
End:   Lock->Unlock 0 ms
Start: Lock->Read   test.t2
End:   Lock->Read   0 ms
Start: Lock->Unlock test.t2 (1 ms lock duration)
End:   Lock->Unlock 0 ms
Start: Lock->Read   test.t2
End:   Lock->Read   0 ms
Start: Lock->Unlock test.t2 (1 ms lock duration)
End:   Lock->Unlock 0 ms
Start: Lock->Read   test.t2
End:   Lock->Read   0 ms

```

### 5.7.1.11. Filesort Probes

The filesort probes are triggered whenever a filesort operation is applied to a table. For more information on filesort and the conditions under which it occurs, see [Section 8.13.13, "ORDER BY Optimization"](#).

```

filestort-start(database, table)
filestort-done(status, rows)

```

- **filestort-start**: Triggered when the filesort operation starts on a table. The two arguments to the probe, **database** and **table**, will identify the table being sorted.
- **filestort-done**: Triggered when the filesort operation completes. Two arguments are supplied, the **status** (0 for success, 1 for failure), and the number of rows sorted during the filesort process.

An example of this is in the following script, which tracks the duration of the filesort process in addition to the duration of the main query:

```

#!/usr/sbin/dtrace -s

#pragma D option quiet

dtrace:::BEGIN
{
    printf("%-2s %-10s %-10s %9s %18s %-s \n",
        "St", "Who", "DB", "ConnID", "Dur microsec", "Query");
}

mysql*:::query-start

```

```

{
    self->query = copyinstr(arg0);
    self->who   = strjoin(copyinstr(arg3),strjoin("@",copyinstr(arg4)));
    self->db    = copyinstr(arg2);
    self->connid = arg1;
    self->querystart = timestamp;
    self->filesort = 0;
    self->fsdb = "";
    self->fstable = "";
}

mysql*:::filesort-start
{
    self->filesort = timestamp;
    self->fsdb = copyinstr(arg0);
    self->fstable = copyinstr(arg1);
}

mysql*:::filesort-done
{
    this->elapsed = (timestamp - self->filesort) /1000;
    printf("%2d %-10s %-10s %9d %18d Filesort on %s\n",
        arg0, self->who, self->fsdb,
        self->connid, this->elapsed, self->fstable);
}

mysql*:::query-done
{
    this->elapsed = (timestamp - self->querystart) /1000;
    printf("%2d %-10s %-10s %9d %18d %s\n",
        arg0, self->who, self->db,
        self->connid, this->elapsed, self->query);
}

```

Executing a query on a large table with an `ORDER BY` clause that triggers a filesort, and then creating an index on the table and then repeating the same query, you can see the difference in execution speed:

St	Who	DB	ConnID	Dur	microsec	Query
0	@localhost	test	14	11335469		Filesort on t1
0	@localhost	test	14	11335787		select * from t1 order by i limit 100
0	@localhost	test	14	466734378		create index t1a on t1 (i)
0	@localhost	test	14	26472		select * from t1 order by i limit 100

### 5.7.1.12. Statement Probes

The individual statement probes are provided to give specific information about different statement types. For the start probes the string of the query is provided as the only argument. Depending on the statement type, the information provided by the corresponding done probe will differ. For all done probes the status of the operation (0 for success, >0 for failure) is provided. For `SELECT`, `INSERT`, `INSERT ... (SELECT FROM ...)`, `DELETE`, and `DELETE FROM t1,t2` operations the number of rows affected is returned.

For `UPDATE` and `UPDATE t1,t2 ...` statements the number of rows matched and the number of rows actually changed is provided. This is because the number of rows actually matched by the corresponding `WHERE` clause, and the number of rows changed can differ. MySQL does not update the value of a row if the value already matches the new setting.

```

select-start(query)
select-done(status,rows)

insert-start(query)
insert-done(status,rows)

insert-select-start(query)
insert-select-done(status,rows)

update-start(query)
update-done(status,rowsmatched,rowschanged)

```

```

multi-update-start(query)
multi-update-done(status, rowsmatched, rowschanged)

delete-start(query)
delete-done(status, rows)

multi-delete-start(query)
multi-delete-done(status, rows)

```

- `select-start`: Triggered before a `SELECT` statement.
- `select-done`: Triggered at the end of a `SELECT` statement.
- `insert-start`: Triggered before a `INSERT` statement.
- `insert-done`: Triggered at the end of an `INSERT` statement.
- `insert-select-start`: Triggered before an `INSERT ... SELECT` statement.
- `insert-select-done`: Triggered at the end of an `INSERT ... SELECT` statement.
- `update-start`: Triggered before an `UPDATE` statement.
- `update-done`: Triggered at the end of an `UPDATE` statement.
- `multi-update-start`: Triggered before an `UPDATE` statement involving multiple tables.
- `multi-update-done`: Triggered at the end of an `UPDATE` statement involving multiple tables.
- `delete-start`: Triggered before a `DELETE` statement.
- `delete-done`: Triggered at the end of a `DELETE` statement.
- `multi-delete-start`: Triggered before a `DELETE` statement involving multiple tables.
- `multi-delete-done`: Triggered at the end of a `DELETE` statement involving multiple tables.

The arguments for the statement probes are:

- `query`: The query string.
- `status`: The status of the query. `0` for success, and `>0` for failure.
- `rows`: The number of rows affected by the statement. This returns the number rows found for `SELECT`, the number of rows deleted for `DELETE`, and the number of rows successfully inserted for `INSERT`.
- `rowsmatched`: The number of rows matched by the `WHERE` clause of an `UPDATE` operation.
- `rowschanged`: The number of rows actually changed during an `UPDATE` operation.

You use these probes to monitor the execution of these statement types without having to monitor the user or client executing the statements. A simple example of this is to track the execution times:

```

#!/usr/sbin/dtrace -s

#pragma D option quiet

dtrace::BEGIN
{
    printf("%-60s %-8s %-8s %-8s\n", "Query", "RowsU", "RowsM", "Dur (ms)");
}

mysql*:::update-start, mysql*:::insert-start,
mysql*:::delete-start, mysql*:::multi-delete-start,
mysql*:::multi-delete-done, mysql*:::select-start,

```

```

mysql*:::insert-select-start, mysql*:::multi-update-start
{
    self->query = copyinstr(arg0);
    self->querystart = timestamp;
}

mysql*:::insert-done, mysql*:::select-done,
mysql*:::delete-done, mysql*:::multi-delete-done, mysql*:::insert-select-done
/ self->querystart /
{
    this->elapsed = ((timestamp - self->querystart)/1000000);
    printf("%-60s %-8d %-8d %d\n",
        self->query,
        0,
        arg1,
        this->elapsed);
    self->querystart = 0;
}

mysql*:::update-done, mysql*:::multi-update-done
/ self->querystart /
{
    this->elapsed = ((timestamp - self->querystart)/1000000);
    printf("%-60s %-8d %-8d %d\n",
        self->query,
        arg1,
        arg2,
        this->elapsed);
    self->querystart = 0;
}

```

When executed you can see the basic execution times and rows matches:

Query	RowsU	RowsM	Dur (ms)
select * from t2	0	275	0
insert into t2 (select * from t2)	0	275	9
update t2 set i=5 where i > 75	110	110	8
update t2 set i=5 where i < 25	254	134	12
delete from t2 where i < 5	0	0	0

Another alternative is to use the aggregation functions in DTrace to aggregate the execution time of individual statements together:

```

#!/usr/sbin/dtrace -s

#pragma D option quiet

mysql*:::update-start, mysql*:::insert-start,
mysql*:::delete-start, mysql*:::multi-delete-start,
mysql*:::multi-delete-done, mysql*:::select-start,
mysql*:::insert-select-start, mysql*:::multi-update-start
{
    self->querystart = timestamp;
}

mysql*:::select-done
{
    @statements["select"] = sum(((timestamp - self->querystart)/1000000));
}

mysql*:::insert-done, mysql*:::insert-select-done
{
    @statements["insert"] = sum(((timestamp - self->querystart)/1000000));
}

mysql*:::update-done, mysql*:::multi-update-done
{
    @statements["update"] = sum(((timestamp - self->querystart)/1000000));
}

mysql*:::delete-done, mysql*:::multi-delete-done

```

```

{
    @statements["delete"] = sum(((timestamp - self->querystart)/1000000));
}

tick-30s
{
    printa(@statements);
}

```

The script just shown aggregates the times spent doing each operation, which could be used to help benchmark a standard suite of tests.

```

delete           0
update           0
insert          23
select         2484

delete           0
update           0
insert          39
select        10744

delete           0
update          26
insert          56
select        10944

delete           0
update          26
insert         2287
select        15985

```

### 5.7.1.13. Network Probes

The network probes monitor the transfer of information from the MySQL server and clients of all types over the network. The probes are defined as follows:

```

net-read-start()
net-read-done(status, bytes)
net-write-start(bytes)
net-write-done(status)

```

- **net-read-start**: Triggered when a network read operation is started.
- **net-read-done**: Triggered when the network read operation completes. The `status` is an `integer` representing the return status for the operation, `0` for success and `1` for failure. The `bytes` argument is an integer specifying the number of bytes read during the process.
- **net-start-bytes**: Triggered when data is written to a network socket. The single argument, `bytes`, specifies the number of bytes written to the network socket.
- **net-write-done**: Triggered when the network write operation has completed. The single argument, `status`, is an integer representing the return status for the operation, `0` for success and `1` for failure.

You can use the network probes to monitor the time spent reading from and writing to network clients during execution. The following D script provides an example of this. Both the cumulative time for the read or write is calculated, and the number of bytes. Note that the dynamic variable size has been increased (using the `dynvarsize` option) to cope with the rapid firing of the individual probes for the network reads/writes.

```

#!/usr/sbin/dtrace -s

#pragma D option quiet
#pragma D option dynvarsize=4m

dtrace:::BEGIN

```

```

{
    printf("%-2s %-30s %-10s %9s %18s %-s \n",
           "St", "Who", "DB", "ConnID", "Dur microsec", "Query");
}

mysql*:::query-start
{
    self->query = copyinstr(arg0);
    self->who   = strjoin(copyinstr(arg3),strjoin("@",copyinstr(arg4)));
    self->db    = copyinstr(arg2);
    self->connid = arg1;
    self->querystart = timestamp;
    self->netwrite = 0;
    self->netwritecum = 0;
    self->netwritebase = 0;
    self->netread = 0;
    self->netreadcum = 0;
    self->netreadbase = 0;
}

mysql*:::net-write-start
{
    self->netwrite += arg0;
    self->netwritebase = timestamp;
}

mysql*:::net-write-done
{
    self->netwritecum += (timestamp - self->netwritebase);
    self->netwritebase = 0;
}

mysql*:::net-read-start
{
    self->netreadbase = timestamp;
}

mysql*:::net-read-done
{
    self->netread += arg1;
    self->netreadcum += (timestamp - self->netreadbase);
    self->netreadbase = 0;
}

mysql*:::query-done
{
    this->elapsed = (timestamp - self->querystart) /1000000;
    printf("%2d %-30s %-10s %9d %18d %s\n",
           arg0, self->who, self->db,
           self->connid, this->elapsed, self->query);
    printf("Net read: %d bytes (%d ms) write: %d bytes (%d ms)\n",
           self->netread, (self->netreadcum/1000000),
           self->netwrite, (self->netwritecum/1000000));
}

```

When executing the above script on a machine with a remote client, you can see that approximately a third of the time spent executing the query is related to writing the query results back to the client.

St	Who	DB	ConnID	Dur microsec	Query
0	root@::ffff:192.168.0.108	test	31	3495	select * from t1 limit 100000
Net read: 0 bytes (0 ms) write: 10000075 bytes (1220 ms)					

### 5.7.1.14. Keycache Probes

The keycache probes are triggered when using the index key cache used with the MyISAM storage engine. Probes exist to monitor when data is read into the keycache, cached key data is written from the cache into a cached file, or when accessing the keycache.

Keycache usage indicates when data is read or written from the index files into the cache, and can be used to monitor how efficient the memory allocated to the keycache is being used. A high number of

keycache reads across a range of queries may indicate that the keycache is too small for size of data being accessed.

```
keycache-read-start(filepath, bytes, mem_used, mem_free)
keycache-read-block(bytes)
keycache-read-hit()
keycache-read-miss()
keycache-read-done(mem_used, mem_free)
keycache-write-start(filepath, bytes, mem_used, mem_free)
keycache-write-block(bytes)
keycache-write-done(mem_used, mem_free)
```

When reading data from the index files into the keycache, the process first initializes the read operation (indicated by `keycache-read-start`), then loads blocks of data (`keycache-read-block`), and then the read block is either matches the data being identified (`keycache-read-hit`) or more data needs to be read (`keycache-read-miss`). Once the read operation has completed, reading stops with the `keycache-read-done`.

Data will be read from the index file into the keycache only when the specified key is not already within the keycache.

- `keycache-read-start`: Triggered when the keycache read operation is started. Data is read from the specified `filepath`, reading the specified number of `bytes`. The `mem_used` and `mem_avail` indicate memory currently used by the keycache and the amount of memory available within the keycache.
- `keycache-read-block`: Triggered when the keycache reads a block of data, of the specified number of `bytes`, from the index file into the keycache.
- `keycache-read-hit`: Triggered when the block of data read from the index file matches the key data requested.
- `keycache-read-miss`: Triggered when the block of data read from the index file does not match the key data needed.
- `keycache-read-done`: Triggered when the keycache read operation has completed. The `mem_used` and `mem_avail` indicate memory currently used by the keycache and the amount of memory available within the keycache.

Keycache writes occur when the index information is updated during an `INSERT`, `UPDATE`, or `DELETE` operation, and the cached key information is flushed back to the index file.

- `keycache-write-start`: Triggered when the keycache write operation is started. Data is written to the specified `filepath`, reading the specified number of `bytes`. The `mem_used` and `mem_avail` indicate memory currently used by the keycache and the amount of memory available within the keycache.
- `keycache-write-block`: Triggered when the keycache writes a block of data, of the specified number of `bytes`, to the index file from the keycache.
- `keycache-write-done`: Triggered when the keycache write operation has completed. The `mem_used` and `mem_avail` indicate memory currently used by the keycache and the amount of memory available within the keycache.





---

## Chapter 6. Security

### Table of Contents

6.1. General Security Issues .....	732
6.1.1. Security Guidelines .....	732
6.1.2. Keeping Passwords Secure .....	733
6.1.3. Making MySQL Secure Against Attackers .....	746
6.1.4. Security-Related <code>mysqld</code> Options and Variables .....	748
6.1.5. How to Run MySQL as a Normal User .....	748
6.1.6. Security Issues with <code>LOAD DATA LOCAL</code> .....	749
6.1.7. Client Programming Security Guidelines .....	750
6.2. The MySQL Access Privilege System .....	751
6.2.1. Privileges Provided by MySQL .....	752
6.2.2. Privilege System Grant Tables .....	756
6.2.3. Specifying Account Names .....	762
6.2.4. Access Control, Stage 1: Connection Verification .....	763
6.2.5. Access Control, Stage 2: Request Verification .....	766
6.2.6. When Privilege Changes Take Effect .....	768
6.2.7. Causes of Access-Denied Errors .....	768
6.3. MySQL User Account Management .....	773
6.3.1. User Names and Passwords .....	773
6.3.2. Adding User Accounts .....	774
6.3.3. Removing User Accounts .....	778
6.3.4. Setting Account Resource Limits .....	778
6.3.5. Assigning Account Passwords .....	780
6.3.6. Password Expiration and Sandbox Mode .....	781
6.3.7. Pluggable Authentication .....	784
6.3.8. Proxy Users .....	803
6.3.9. Using SSL for Secure Connections .....	806
6.3.10. Connecting to MySQL Remotely from Windows with SSH .....	817
6.3.11. MySQL Enterprise Audit Log Plugin .....	817
6.3.12. SQL-Based MySQL Account Activity Auditing .....	827

When thinking about security within a MySQL installation, you should consider a wide range of possible topics and how they affect the security of your MySQL server and related applications:

- General factors that affect security. These include choosing good passwords, not granting unnecessary privileges to users, ensuring application security by preventing SQL injections and data corruption, and others. See [Section 6.1, “General Security Issues”](#).
- Security of the installation itself. The data files, log files, and the all the application files of your installation should be protected to ensure that they are not readable or writable by unauthorized parties. For more information, see [Section 2.10, “Postinstallation Setup and Testing”](#).
- Access control and security within the database system itself, including the users and databases granted with access to the databases, views and stored programs in use within the database. For more information, see [Section 6.2, “The MySQL Access Privilege System”](#), and [Section 6.3, “MySQL User Account Management”](#).
- Network security of MySQL and your system. The security is related to the grants for individual users, but you may also wish to restrict MySQL so that it is available only locally on the MySQL server host, or to a limited set of other hosts.
- Ensure that you have adequate and appropriate backups of your database files, configuration and log files. Also be sure that you have a recovery solution in place and test that you are able to successfully recover the information from your backups. See [Chapter 7, Backup and Recovery](#).

## 6.1. General Security Issues

This section describes general security issues to be aware of and what you can do to make your MySQL installation more secure against attack or misuse. For information specifically about the access control system that MySQL uses for setting up user accounts and checking database access, see [Section 6.2, “The MySQL Access Privilege System”](#).

For answers to some questions that are often asked about MySQL Server security issues, see [Section B.9, “MySQL 5.6 FAQ: Security”](#).

### 6.1.1. Security Guidelines

Anyone using MySQL on a computer connected to the Internet should read this section to avoid the most common security mistakes.

In discussing security, it is necessary to consider fully protecting the entire server host (not just the MySQL server) against all types of applicable attacks: eavesdropping, altering, playback, and denial of service. We do not cover all aspects of availability and fault tolerance here.

MySQL uses security based on Access Control Lists (ACLs) for all connections, queries, and other operations that users can attempt to perform. There is also support for SSL-encrypted connections between MySQL clients and servers. Many of the concepts discussed here are not specific to MySQL at all; the same general ideas apply to almost all applications.

When running MySQL, follow these guidelines:

- **Do not ever give anyone (except MySQL `root` accounts) access to the `user` table in the `mysql` database!** This is critical.
- Learn how the MySQL access privilege system works (see [Section 6.2, “The MySQL Access Privilege System”](#)). Use the `GRANT` and `REVOKE` statements to control access to MySQL. Do not grant more privileges than necessary. Never grant privileges to all hosts.

Checklist:

- Try `mysql -u root`. If you are able to connect successfully to the server without being asked for a password, anyone can connect to your MySQL server as the MySQL `root` user with full privileges! Review the MySQL installation instructions, paying particular attention to the information about setting a `root` password. See [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).
- Use the `SHOW GRANTS` statement to check which accounts have access to what. Then use the `REVOKE` statement to remove those privileges that are not necessary.
- Do not store cleartext passwords in your database. If your computer becomes compromised, the intruder can take the full list of passwords and use them. Instead, use `SHA2 ( )` [1264], `SHA1 ( )` [1264], `MD5 ( )` [1263], or some other one-way hashing function and store the hash value.
- Do not choose passwords from dictionaries. Special programs exist to break passwords. Even passwords like “xfish98” are very bad. Much better is “duag98” which contains the same word “fish” but typed one key to the left on a standard QWERTY keyboard. Another method is to use a password that is taken from the first characters of each word in a sentence (for example, “Four score and seven years ago” results in a password of “Fsasya”). The password is easy to remember and type, but difficult to guess for someone who does not know the sentence. In this case, you can additionally substitute digits for the number words to obtain the phrase “4 score and 7 years ago”, yielding the password “4sa7ya” which is even more difficult to guess.
- Invest in a firewall. This protects you from at least 50% of all types of exploits in any software. Put MySQL behind the firewall or in a demilitarized zone (DMZ).

Checklist:

- Try to scan your ports from the Internet using a tool such as [nmap](#). MySQL uses port 3306 by default. This port should not be accessible from untrusted hosts. As a simple way to check whether your MySQL port is open, try the following command from some remote machine, where `server_host` is the host name or IP address of the host on which your MySQL server runs:

```
shell> telnet server_host 3306
```

If `telnet` hangs or the connection is refused, the port is blocked, which is how you want it to be. If you get a connection and some garbage characters, the port is open, and should be closed on your firewall or router, unless you really have a good reason to keep it open.

- Applications that access MySQL should not trust any data entered by users, and should be written using proper defensive programming techniques. See [Section 6.1.7, “Client Programming Security Guidelines”](#).
- Do not transmit plain (unencrypted) data over the Internet. This information is accessible to everyone who has the time and ability to intercept it and use it for their own purposes. Instead, use an encrypted protocol such as SSL or SSH. MySQL supports internal SSL connections. Another technique is to use SSH port-forwarding to create an encrypted (and compressed) tunnel for the communication.
- Learn to use the `tcpdump` and `strings` utilities. In most cases, you can check whether MySQL data streams are unencrypted by issuing a command like the following:

```
shell> tcpdump -l -i eth0 -w - src or dst port 3306 | strings
```

This works under Linux and should work with small modifications under other systems.



### Warning

If you do not see cleartext data, this does not always mean that the information actually is encrypted. If you need high security, consult with a security expert.

## 6.1.2. Keeping Passwords Secure

Passwords occur in several contexts within MySQL. The following sections provide guidelines that enable end users and administrators to keep these passwords secure and avoid exposing them. There is also a discussion of how MySQL uses password hashing internally and of a plugin that you can use to enforce stricter passwords.

### 6.1.2.1. End-User Guidelines for Password Security

MySQL users should use the following guidelines to keep passwords secure.

When you run a client program to connect to the MySQL server, it is inadvisable to specify your password in a way that exposes it to discovery by other users. The methods you can use to specify your password when you run client programs are listed here, along with an assessment of the risks of each method. In short, the safest methods are to have the client program prompt for the password or to specify the password in a properly protected option file.

- Use the `mysql_config_editor` utility, which enables you to store authentication credentials in an encrypted login file named `.mylogin.cnf`. The file can be read later by MySQL client programs to obtain authentication credentials for connecting to MySQL Server. See [Section 4.6.6, “mysql\\_config\\_editor — MySQL Configuration Utility”](#).
- Use a `-pyour_pass` or `--password=your_pass` option on the command line. For example:

```
shell> mysql -u francis -pfrank db_name
```

This is convenient *but insecure*. On some systems, your password becomes visible to system status programs such as `ps` that may be invoked by other users to display command lines. MySQL clients typically overwrite the command-line password argument with zeros during their initialization sequence. However, there is still a brief interval during which the value is visible. Also, on some systems this overwriting strategy is ineffective and the password remains visible to `ps`. (SystemV Unix systems and perhaps others are subject to this problem.)

If your operating environment is set up to display your current command in the title bar of your terminal window, the password remains visible as long as the command is running, even if the command has scrolled out of view in the window content area.

- Use the `-p` or `--password` option on the command line with no password value specified. In this case, the client program solicits the password interactively:

```
shell> mysql -u francis -p db_name
Enter password: *****
```

The “\*” characters indicate where you enter your password. The password is not displayed as you enter it.

It is more secure to enter your password this way than to specify it on the command line because it is not visible to other users. However, this method of entering a password is suitable only for programs that you run interactively. If you want to invoke a client from a script that runs noninteractively, there is no opportunity to enter the password from the keyboard. On some systems, you may even find that the first line of your script is read and interpreted (incorrectly) as your password.

- Store your password in an option file. For example, on Unix, you can list your password in the `[client]` section of the `.my.cnf` file in your home directory:

```
[client]
password=your_pass
```

To keep the password safe, the file should not be accessible to anyone but yourself. To ensure this, set the file access mode to `400` or `600`. For example:

```
shell> chmod 600 .my.cnf
```

To name from the command line a specific option file containing the password, use the `--defaults-file=file_name` [219] option, where `file_name` is the full path name to the file. For example:

```
shell> mysql --defaults-file=/home/francis/mysql-opts
```

[Section 4.2.3.3, “Using Option Files”](#), discusses option files in more detail.

- Store your password in the `MYSQL_PWD` environment variable. See [Section 2.12, “Environment Variables”](#).

This method of specifying your MySQL password must be considered *extremely insecure* and should not be used. Some versions of `ps` include an option to display the environment of running processes. On some systems, if you set `MYSQL_PWD`, your password is exposed to any other user who runs `ps`. Even on systems without such a version of `ps`, it is unwise to assume that there are no other methods by which users can examine process environments.

On Unix, the `mysql` client writes a record of executed statements to a history file (see [Section 4.5.1.3, “mysql Logging”](#)). By default, this file is named `.mysql_history` and is created in your home directory. Passwords can be written as plain text in SQL statements such as `CREATE USER`, `GRANT`, and `SET PASSWORD`, so if you use these statements, they are logged in the history file. To keep this file safe, use a restrictive access mode, the same way as described earlier for the `.my.cnf` file.

If your command interpreter is configured to maintain a history, any file in which the commands are saved will contain MySQL passwords entered on the command line. For example, `bash` uses `~/.bash_history`. Any such file should have a restrictive access mode.

### 6.1.2.2. Administrator Guidelines for Password Security

Database administrators should use the following guidelines to keep passwords secure.

MySQL stores passwords for user accounts in the `mysql.user` table. Access to this table should never be granted to any nonadministrative accounts.

Account passwords can be expired so that users must reset them. See [Section 6.3.6, “Password Expiration and Sandbox Mode”](#).

The `validate_password` plugin can be used to enforce a policy on acceptable password. See [Section 6.1.2.6, “The Password Validation Plugin”](#).

A user who has access to modify the plugin directory (the value of the `plugin_dir` [525] system variable) or the `my.cnf` file that specifies the location of the plugin directory can replace plugins and modify the capabilities provided by plugins, including authentication plugins.

Files such as log files to which passwords might be written should be protected. See [Section 6.1.2.3, “Passwords and Logging”](#).

### 6.1.2.3. Passwords and Logging

Passwords can be written as plain text in SQL statements such as `CREATE USER`, `GRANT`, and `SET PASSWORD`, or statements that invoke the `PASSWORD()` [1263] function. If these statements are logged by the MySQL server as written, such passwords become available to anyone with access to the logs.

Beginning with MySQL 5.6.3, statement logging is modified so that passwords do not appear in plain text for the following statements:

```
CREATE USER ... IDENTIFIED BY ...
GRANT ... IDENTIFIED BY ...
SET PASSWORD ...
SLAVE START ... PASSWORD = ... (as of 5.6.4)
CREATE SERVER ... OPTIONS(... PASSWORD ...) (as of 5.6.9)
ALTER SERVER ... OPTIONS(... PASSWORD ...) (as of 5.6.9)
```

Passwords in those statements are rewritten not to appear literally in statement text, for the general query log, slow query log, and binary log. Rewriting does not apply to other statements.

For the general query log, password rewriting can be suppressed by starting the server with the `--log-raw` [428] option. This option may be useful for diagnostic purposes, to see the exact text of statements as received by the server, but for security reasons is not recommended for production use.

Passwords can be written as plain text in SQL statements in the audit log plugin log file (see [Section 6.3.11, “MySQL Enterprise Audit Log Plugin”](#)).

To guard log files against unwarranted exposure, they should be located in a directory that restricts access to only the server and the database administrator. If you log to tables in the `mysql` database, access to those tables should never be granted to any nonadministrative accounts.

Replication slaves store the password for the replication master in the `master.info` file. Restrict this file to be accessible only to the database administrator. As of MySQL 5.6.4, an alternative to storing the password in a file is to use the `START SLAVE` statement to specify credentials for connecting to the master.

Database backups that include tables or log files containing passwords should be protected using a restricted access mode.

### 6.1.2.4. Password Hashing in MySQL

MySQL lists user accounts in the `user` table of the `mysql` database. Each MySQL account can be assigned a password, although the `user` table does not store the cleartext version of the password, but a hash value computed from it.

MySQL uses passwords in two phases of client/server communication:

- When a client attempts to connect to the server, there is an initial authentication step in which the client must present a password that has a hash value matching the hash value stored in the `user` table for the account the client wants to use.
- After the client connects, it can (if it has sufficient privileges) set or change the password hash for accounts listed in the `user` table. The client can do this by using the `PASSWORD( )` [1263] function to generate a password hash, or by using a password-generating statement (`CREATE USER`, `GRANT`, or `SET PASSWORD`).

In other words, the server *checks* hash values during authentication when a client first attempts to connect. The server *generates* hash values if a connected client invokes the `PASSWORD( )` [1263] function or uses a password-generating statement to set or change a password.

Password hashing methods in MySQL have the history described following. These changes are illustrated by changes in the result from the `PASSWORD( )` [1263] function that computes password hash values and in the structure of the `user` table where passwords are stored.

#### The Original (Pre-4.1) Hashing Method

The original hashing method produced a 16-byte string. Such hashes look like this:

```
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| 6f8c114b58f2ce9e |
+-----+
```

To store account passwords, the `Password` column of the `user` table was at this point 16 bytes long.

#### The 4.1 Hashing Method

MySQL 4.1 introduced password hashing that provides better security and reduces the risk of passwords being intercepted. There were several aspects to this change:

- Different `PASSWORD( )` [1263] function result format
- Widening of the `Password` column
- Control over the default hashing method
- Control over the permitted hashing methods for clients attempting to connect to the server

The changes in MySQL 4.1 took place in two stages:

- MySQL 4.1.0 used a preliminary version of the 4.1 hashing method. Because this method was so short lived, the following discussion says no more about it.
- In MySQL 4.1.1, the hashing method was modified to produce a longer 41-byte hash value:

```
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| *6C8989366EAF75BB670AD8EA7A7FC1176A95CEF4 |
+-----+
```



The longer password hash format has better cryptographic properties, and client authentication based on long hashes is more secure than that based on the older short hashes.

To accommodate longer password hashes, the `Password` column in the `user` table was changed at this point to be 41 bytes, its current length.

A widened `Password` column can store password hashes in both the pre-4.1 and 4.1 formats. The format of any given hash value can be determined two ways:

- The length: 4.1 and pre-4.1 hashes are 41 and 16 bytes, respectively.
- Password hashes in the 4.1 format always begin with a “\*” character, whereas passwords in the pre-4.1 format never do.

To permit explicit generation of pre-4.1 password hashes, two additional changes were made:

- The `OLD_PASSWORD()` [1263] function was added, which returns hash values in the 16-byte format.
- For compatibility purposes, the `old_passwords` [517] system variable was added, to enable DBAs and applications control over the hashing method. The default `old_passwords` [517] value of 0 causes hashing to use the 4.1 method (41-byte hash values), but setting `old_passwords=1` [517] causes hashing to use the pre-4.1 method. In this case, `PASSWORD()` [1263] produces 16-byte values and is equivalent to `OLD_PASSWORD()` [1263]

To permit DBAs control over how clients are permitted to connect, the `secure_auth` [538] system variable was added. Starting the server with this variable disabled or enabled permits or prohibits clients to connect using the older pre-4.1 password hashing method. Before MySQL 5.6.5, `secure_auth` [538] is disabled by default. As of 5.6.5, `secure_auth` [538] is enabled by default to promote a more secure default configuration. (DBAs can disable it at their discretion, but this is not recommended.)

In addition, the `mysql` client supports a `--secure-auth` [253] option that is analogous to `secure_auth` [538], but from the client side. It can be used to prevent connections to less secure accounts that use pre-4.1 password hashing. This option is disabled by default before MySQL 5.6.7, enabled thereafter.

## Compatibility Issues Related to Hashing Methods

The widening of the `Password` column in MySQL 4.1 from 16 bytes to 41 bytes affects installation or upgrade operations as follows:

- If you perform a new installation of MySQL, the `Password` column is made 41 bytes long automatically.
- Upgrades from MySQL 4.1 or later to current versions of MySQL should not give rise to any issues in regard to the `Password` column because both versions use the same column length and password hashing method.
- For upgrades from a pre-4.1 release to 4.1 or later, you must upgrade the system tables after upgrading. (See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).)

The 4.1 hashing method is understood only by MySQL 4.1 (and newer) servers and clients, which can result in some compatibility problems. A 4.1 or newer client can connect to a pre-4.1 server, because the client understands both the pre-4.1 and 4.1 password hashing methods. However, a pre-4.1 client that attempts to connect to a 4.1 or newer server may run into difficulties. For example, a 4.0 `mysql` client may fail with the following error message:

```
shell> mysql -h localhost -u root
Client does not support authentication protocol requested
by server; consider upgrading MySQL client
```



This phenomenon also occurs for attempts to use the older PHP `mysql` extension after upgrading to MySQL 4.1 or newer. (See [Section 21.9.12, “Common Problems with MySQL and PHP”](#).)

The following discussion describes the differences between the pre-4.1 and 4.1 hashing methods, and what you should do if you upgrade your server but need to maintain backward compatibility with pre-4.1 clients. (However, permitting connections by old clients is not recommended and should be avoided if possible.) Additional information can be found in [Section C.5.2.4, “Client does not support authentication protocol”](#). This information is of particular importance to PHP programmers migrating MySQL databases from versions older than 4.1 to 4.1 or higher.

The differences between short and long password hashes are relevant both for how the server uses passwords during authentication and for how it generates password hashes for connected clients that perform password-changing operations.

The way in which the server uses password hashes during authentication is affected by the width of the `Password` column:

- If the column is short, only short-hash authentication is used.
- If the column is long, it can hold either short or long hashes, and the server can use either format:
  - Pre-4.1 clients can connect, but because they know only about the pre-4.1 hashing method, they can authenticate only using accounts that have short hashes.
  - 4.1 and later clients can authenticate using accounts that have short or long hashes.

Even for short-hash accounts, the authentication process is actually a bit more secure for 4.1 and later clients than for older clients. In terms of security, the gradient from least to most secure is:

- Pre-4.1 client authenticating with short password hash
- 4.1 or later client authenticating with short password hash
- 4.1 or later client authenticating with long password hash

The way in which the server generates password hashes for connected clients is affected by the width of the `Password` column and by the `old_passwords [517]` system variable. A 4.1 or later server generates long hashes only if certain conditions are met: The `Password` column must be wide enough to hold long values and `old_passwords [517]` must not be set to 1.

Those conditions apply as follows:

- The `Password` column must be wide enough to hold long hashes (41 bytes). If the column has not been updated and still has the pre-4.1 width of 16 bytes, the server notices that long hashes cannot fit into it and generates only short hashes when a client performs password-changing operations using the `PASSWORD( ) [1263]` function or a password-generating statement. This is the behavior that occurs if you have upgraded from a version of MySQL older than 4.1 to 4.1 or later but have not yet run the `mysql_upgrade` program to widen the `Password` column.
- If the `Password` column is wide, it can store either short or long password hashes. In this case, the `PASSWORD( ) [1263]` function and password-generating statements generate long hashes unless the server was started with the `old_passwords [517]` system variable set to 1 to force the server to generate short password hashes instead.

The purpose of the `old_passwords [517]` system variable is to permit backward compatibility with pre-4.1 clients under circumstances where the server would otherwise generate long password hashes. The option does not affect authentication (4.1 and later clients can still use accounts that have long password hashes), but it does prevent creation of a long password hash in the `user` table as the result of a password-changing operation. Were that permitted to occur, the account could no longer be used by pre-4.1 clients. With `old_passwords [517]` disabled, the following undesirable scenario is possible:

- An old pre-4.1 client connects to an account that has a short password hash.

- The client changes its own password. With `old_passwords` [517] disabled, this results in the account having a long password hash.
- The next time the old client attempts to connect to the account, it cannot, because the account has a long password hash that requires the 4.1 hashing method during authentication. (Once an account has a long password hash in the user table, only 4.1 and later clients can authenticate for it because pre-4.1 clients do not understand long hashes.)

This scenario illustrates that, if you must support older pre-4.1 clients, it is problematic to run a 4.1 or newer server without `old_passwords` [517] set to 1. By running the server with `old_passwords=1` [517], password-changing operations do not generate long password hashes and thus do not cause accounts to become inaccessible to older clients. (Those clients cannot inadvertently lock themselves out by changing their password and ending up with a long password hash.)

The downside of `old_passwords=1` [517] is that any passwords created or changed use short hashes, even for 4.1 or later clients. Thus, you lose the additional security provided by long password hashes. To create an account that has a long hash (for example, for use by 4.1 clients) or to change an existing account to use a long password hash, an administrator can set the session value of `old_passwords` [517] set to 0 while leaving the global value set to 1:

```
mysql> SET @@session.old_passwords = 0;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @@session.old_passwords, @@global.old_passwords;
+-----+-----+
| @@session.old_passwords | @@global.old_passwords |
+-----+-----+
| 0 | 1 |
+-----+-----+
1 row in set (0.00 sec)

mysql> CREATE USER 'newuser'@'localhost' IDENTIFIED BY 'newpass';
Query OK, 0 rows affected (0.03 sec)

mysql> SET PASSWORD FOR 'existinguser'@'localhost' = PASSWORD('existingpass');
Query OK, 0 rows affected (0.00 sec)
```

The following scenarios are possible in MySQL 4.1 or later. The factors are whether the `Password` column is short or long, and, if long, whether the server is started with `old_passwords` [517] enabled or disabled.

**Scenario 1:** Short `Password` column in user table:

- Only short hashes can be stored in the `Password` column.
- The server uses only short hashes during client authentication.
- For connected clients, password hash-generating operations involving the `PASSWORD( )` [1263] function or password-generating statements use short hashes exclusively. Any change to an account's password results in that account having a short password hash.
- The value of `old_passwords` [517] is irrelevant because with a short `Password` column, the server generates only short password hashes anyway.

This scenario occurs when a pre-4.1 MySQL installation has been upgraded to 4.1 or later but `mysql_upgrade` has not been run to upgrade the system tables in the `mysql` database. (This is not a recommended configuration because it does not permit use of more secure 4.1 password hashing.)

**Scenario 2:** Long `Password` column; server started with `old_passwords=1` [517]:

- Short or long hashes can be stored in the `Password` column.
- 4.1 and later clients can authenticate for accounts that have short or long hashes.
- Pre-4.1 clients can authenticate only for accounts that have short hashes.

- For connected clients, password hash-generating operations involving the `PASSWORD()` [1263] function or password-generating statements use short hashes exclusively. Any change to an account's password results in that account having a short password hash.

In this scenario, newly created accounts have short password hashes because `old_passwords=1` [517] prevents generation of long hashes. Also, if you create an account with a long hash before setting `old_passwords` [517] to 1, changing the account's password while `old_passwords=1` [517] results in the account being given a short password, causing it to lose the security benefits of a longer hash.

To create a new account that has a long password hash, or to change the password of any existing account to use a long hash, first set the session value of `old_passwords` [517] set to 0 while leaving the global value set to 1, as described previously.

In this scenario, the server has an up to date `Password` column, but is running with the default password hashing method set to generate pre-4.1 hash values. This is not a recommended configuration but may be useful during a transitional period in which pre-4.1 clients and passwords are upgraded to 4.1 or later. When that has been done, it is preferable to run the server with `old_passwords=0` [517] and `secure_auth=1` [538].

**Scenario 3:** Long `Password` column; server started with `old_passwords=0` [517]:

- Short or long hashes can be stored in the `Password` column.
- 4.1 and later clients can authenticate using accounts that have short or long hashes.
- Pre-4.1 clients can authenticate only using accounts that have short hashes.
- For connected clients, password hash-generating operations involving the `PASSWORD()` [1263] function or password-generating statements use short hashes exclusively. A change to an account's password results in that account having a long password hash.

As indicated earlier, a danger in this scenario is that it is possible for accounts that have a short password hash to become inaccessible to pre-4.1 clients. A change to such an account's password made using the `PASSWORD()` [1263] function or a password-generating statement results in the account being given a long password hash. From that point on, no pre-4.1 client can connect to the server using that account. The client must upgrade to 4.1 or later.

If this is a problem, you can change a password in a special way. For example, normally you use `SET PASSWORD` as follows to change an account password:

```
SET PASSWORD FOR 'some_user'@'some_host' = PASSWORD('mypass');
```

To change the password but create a short hash, use the `OLD_PASSWORD()` [1263] function instead:

```
SET PASSWORD FOR 'some_user'@'some_host' = OLD_PASSWORD('mypass');
```

`OLD_PASSWORD()` [1263] is useful for situations in which you explicitly want to generate a short hash.

The disadvantages for each of the preceding scenarios may be summarized as follows:

In scenario 1, you cannot take advantage of longer hashes that provide more secure authentication.

In scenario 2, `old_passwords=1` [517] prevents accounts with short hashes from becoming inaccessible, but password-changing operations cause accounts with long hashes to revert to short hashes unless you take care to change the session value of `old_passwords` [517] to 0 first.

In scenario 3, accounts with short hashes become inaccessible to pre-4.1 clients if you change their passwords without explicitly using `OLD_PASSWORD()` [1263].

The best way to avoid compatibility problems related to short password hashes is to not use them:

- Upgrade all client programs to MySQL 4.1 or later.
- Run the server with `old_passwords=0` [517].

- Reset the password for any account with a short password hash to use a long password hash.
- For additional security, run the server with `secure_auth=1` [538].

### 6.1.2.5. Implications of Password Hashing Changes in MySQL 4.1 for Application Programs

An upgrade to MySQL version 4.1 or later can cause compatibility issues for applications that use `PASSWORD()` [1263] to generate passwords for their own purposes. Applications really should not do this, because `PASSWORD()` [1263] should be used only to manage passwords for MySQL accounts. But some applications use `PASSWORD()` [1263] for their own purposes anyway.

If you upgrade to 4.1 or later from a pre-4.1 version of MySQL and run the server under conditions where it generates long password hashes, an application using `PASSWORD()` [1263] for its own passwords breaks. The recommended course of action in such cases is to modify the application to use another function, such as `SHA2()` [1264], `SHA1()` [1264], or `MD5()` [1263], to produce hashed values. If that is not possible, you can use the `OLD_PASSWORD()` [1263] function, which is provided for generate short hashes in the old format. However, you should note that `OLD_PASSWORD()` [1263] may one day no longer be supported.

If the server is running with `old_passwords=1` [517], it generates short hashes and `OLD_PASSWORD()` [1263] is equivalent to `PASSWORD()` [1263].

PHP programmers migrating their MySQL databases from version 4.0 or lower to version 4.1 or higher should see [Section 21.9, “MySQL PHP API”](#).

### 6.1.2.6. The Password Validation Plugin

The `validate_password` plugin (available as of MySQL 5.6.6) can be used to test passwords and improve security. This plugin implements two capabilities:

- In statements that assign a password supplied as a cleartext value, the value is checked against the current password policy and rejected if it is weak (the statement returns an `ER_NOT_VALID_PASSWORD` [3714] error). This affects the `CREATE USER`, `GRANT`, and `SET PASSWORD` statements. Passwords given as arguments to the `PASSWORD()` [1263] and `OLD_PASSWORD()` [1263] functions are checked as well.
- The strength of potential passwords can be assessed using the `VALIDATE_PASSWORD_STRENGTH()` [1265] SQL function, which takes a password argument and returns an integer from 0 (weak) to 100 (strong).

For example, the cleartext password in the following statement is checked. Under the default password policy, which requires passwords to be at least 8 characters long, the password is weak and the statement produces an error:

```
mysql> SET PASSWORD = PASSWORD('abc');
ERROR 1819 (HY000): Your password does not satisfy the current policy
requirements
```

Passwords specified as already hashed values are not checked because the original password value is not available:

```
mysql> SET PASSWORD = '*0D3CED9BEC10A777AEC23CCC353A8C08A633045E';
Query OK, 0 rows affected (0.01 sec)
```

The parameters that control password checking are available as the values of the system variables having names of the form `validate_password_xxx`. These variables can be modified to configure password checking; see [Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”](#).

The three levels of password checking are `LOW`, `MEDIUM`, and `STRONG`. The default is `MEDIUM`; to change this, modify the value of `validate_password_policy` [745]. The policies implement increasingly strict password tests. The following descriptions refer to default parameter values; these can be modified by changing the appropriate system variables.

- **LOW** policy tests password length only. Passwords must be at least 8 characters long.
- **MEDIUM** policy adds the conditions that passwords must contain at least 1 numeric character, 1 lowercase and uppercase character, and 1 special (nonalphanumeric) character.
- **STRONG** policy adds the condition that password substrings of length 4 or longer must not match words in the dictionary file, if one has been specified.

If the `validate_password` plugin is not installed, the `validate_password_xxx` system variables are not available, passwords in statements are not checked, and `VALIDATE_PASSWORD_STRENGTH()` [1265] always returns 0. For example, accounts can be assigned passwords shorter than 8 characters.

### 6.1.2.6.1. Password Validation Plugin Installation

The password-validation plugin is named `validate_password`. To be usable by the server, the plugin library object file must be located in the MySQL plugin directory (the directory named by the `plugin_dir` [525] system variable). If necessary, set the value of `plugin_dir` [525] at server startup to tell the server the location of the plugin directory.

To load the plugin at server startup, use the `--plugin-load` [435] option to name the object file that contains the plugin. With this plugin-loading method, the option must be given each time the server starts. For example, put these lines in your `my.cnf` file:

```
[mysqld]
plugin-load=validate_password.so
```

If object files have a suffix different from `.so` on your system, substitute the correct suffix (for example, `.dll` on Windows).

Alternatively, to register the plugin at runtime, use this statement (changing the extension as necessary):

```
mysql> INSTALL PLUGIN validate_password SONAME 'validate_password.so';
```

`INSTALL PLUGIN` loads the plugin, and also registers it in the `mysql.plugins` table to cause the plugin to be loaded for each subsequent normal server startup.

If the plugin has been previously registered with `INSTALL PLUGIN` or is loaded with `--plugin-load` [435], you can use the `--validate-password` option at server startup to control plugin activation. For example, to load the plugin and prevent it from being removed at runtime, use these options:

```
[mysqld]
plugin-load=validate_password.so
validate-password=FORCE_PLUS_PERMANENT
```

If it is desired to prevent the server from running without the password-validation plugin, use `--validate-password` [742] with a value of `FORCE` or `FORCE_PLUS_PERMANENT` to force server startup to fail if the plugin does not initialize successfully.

For general information about installing plugins, see [Section 5.1.8, “Server Plugins”](#). To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement. See [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).

### 6.1.2.6.2. Password Validation Plugin Options and Variables

To control the activation of the `validate_password` plugin, use this option:

- `--validate-password[=value]` [742]

<b>Version Introduced</b>	5.6.6
---------------------------	-------

<b>Command-Line Format</b>	<code>--validate-password[=value]</code>
<b>Option-File Format</b>	<code>validate-password[=value]</code>
	<b>Permitted Values</b>
<b>Type</b>	enumeration
<b>Default</b>	ON
<b>Valid Values</b>	ON OFF FORCE FORCE_PLUS_PERMANENT

This option controls how the server loads the `validate_password` plugin at startup. The value should be one of those available for plugin-loading options, as described in [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#). For example, `--validate-password=FORCE_PLUS_PERMANENT` [742] tells the server to load the plugin and prevent it from being removed while the server is running.

This option is available only if the `validate_password` plugin has been previously registered with `INSTALL PLUGIN` or is loaded with `--plugin-load` [435]. See [Section 6.1.2.6.1, “Password Validation Plugin Installation”](#).

If the `validate_password` plugin is installed, it exposes several system variables that indicate the parameters that control password checking:

```
mysql> SHOW VARIABLES LIKE 'validate_password%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| validate_password_dictionary_file |      |
| validate_password_length | 8 |
| validate_password_mixed_case_count | 1 |
| validate_password_number_count | 1 |
| validate_password_policy | MEDIUM |
| validate_password_special_char_count | 1 |
+-----+-----+
```

To change how passwords are checked, you can set any of these variables at server startup, and most of them at runtime. The following list describes the meaning of each variable.

- `validate_password_dictionary_file` [743]

<b>Version Introduced</b>	5.6.6
<b>Variable Name</b>	<code>validate_password_dictionary_file</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	file name

The path name of the dictionary file used by the `validate_password` plugin for checking passwords. This variable is unavailable unless that plugin is installed.

By default, this variable has an empty value and dictionary checks are not performed. To enable dictionary checks, you must set this variable to a nonempty value. If the file is named as a relative path, it is interpreted relative to the server data directory. Its contents should be lowercase, one word per line. Contents are treated as having a character set of `utf8`. The maximum permitted file size is 1MB.



For the dictionary file to be used during password checking, the password policy must be set to 2 (`STRONG`); see the description of the `validate_password_policy` [745] system variable. Assuming that is true, each substring of the password of length 4 up to 100 is compared to the words in the dictionary file. Any match causes the password to be rejected. Comparisons are not case sensitive.

For `VALIDATE_PASSWORD_STRENGTH()` [1265] the password is checked against all policies, including `STRONG`, so the strength assessment includes the dictionary check regardless of the `validate_password_policy` [745] value.

Changes to the dictionary file while the server is running require a restart for the server to recognize the changes.

- `validate_password_length` [744]

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>validate_password_length</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	8
	<b>Min Value</b>	0

The minimum number of characters that passwords checked by the `validate_password` plugin must have. This variable is unavailable unless that plugin is installed.

The `validate_password_length` [744] minimum value is a function of several other related system variables. As of MySQL 5.6.10, the server will not set the value less than the value of this expression:

```
validate_password_number_count
+ validate_password_special_char_count
+ (2 * validate_password_mixed_case_count)
```

If the `validate_password` plugin adjusts the value of `validate_password_length` [744] due to the preceding constraint, it writes a message to the error log.

- `validate_password_mixed_case_count` [744]

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>validate_password_mixed_case_count</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	1
	<b>Min Value</b>	0

The minimum number of lowercase and uppercase characters that passwords checked by the `validate_password` plugin must have if the password policy is `MEDIUM` or stronger. This variable is unavailable unless that plugin is installed.

- [validate\\_password\\_number\\_count \[745\]](#)

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>validate_password_number_count</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>
	<b>Min Value</b>	<code>0</code>

The minimum number of numeric (digit) characters that passwords checked by the `validate_password` plugin must have if the password policy is `MEDIUM` or stronger. This variable is unavailable unless that plugin is installed.

- [validate\\_password\\_policy \[745\]](#)

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>validate_password_policy</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>1</code>
	<b>Valid Values</b>	<code>0</code> <code>1</code> <code>2</code>

The password policy enforced by the `validate_password` plugin. This variable is unavailable unless that plugin is installed.

The `validate_password_policy [745]` value can be specified using numeric values 0, 1, 2, or the corresponding symbolic values `LOW`, `MEDIUM`, `STRONG`. The following table describes the tests performed for each policy. For the length test, the required length is the value of the `validate_password_length [744]` system variable. Similarly, the required values for the other tests are given by other `validate_password_xxx` variables.

Policy	Tests Performed
<code>0</code> or <code>LOW</code>	Length
<code>1</code> or <code>MEDIUM</code>	Length; numeric, lowercase/uppercase, and special characters
<code>2</code> or <code>STRONG</code>	Length; numeric, lowercase/uppercase, and special characters; dictionary file



**Note**

Before MySQL 5.6.10, `validate_password_policy` was named `validate_password_policy_number`.

- [validate\\_password\\_special\\_char\\_count \[745\]](#)



<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>validate_password_special_char_count</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>
	<b>Min Value</b>	<code>0</code>

The minimum number of nonalphanumeric characters that passwords checked by the `validate_password` plugin must have if the password policy is `MEDIUM` or stronger. This variable is unavailable unless that plugin is installed.

### 6.1.3. Making MySQL Secure Against Attackers

When you connect to a MySQL server, you should use a password. The password is not transmitted in clear text over the connection. Password handling during the client connection sequence was upgraded in MySQL 4.1.1 to be very secure. If you are still using pre-4.1.1-style passwords, the encryption algorithm is not as strong as the newer algorithm. With some effort, a clever attacker who can sniff the traffic between the client and the server can crack the password. (See [Section 6.1.2.4, “Password Hashing in MySQL”](#), for a discussion of the different password handling methods.)

All other information is transferred as text, and can be read by anyone who is able to watch the connection. If the connection between the client and the server goes through an untrusted network, and you are concerned about this, you can use the compressed protocol to make traffic much more difficult to decipher. You can also use MySQL's internal SSL support to make the connection even more secure. See [Section 6.3.9, “Using SSL for Secure Connections”](#). Alternatively, use SSH to get an encrypted TCP/IP connection between a MySQL server and a MySQL client. You can find an Open Source SSH client at <http://www.openssh.org/>, and a commercial SSH client at <http://www.ssh.com/>.

To make a MySQL system secure, you should strongly consider the following suggestions:

- Require all MySQL accounts to have a password. A client program does not necessarily know the identity of the person running it. It is common for client/server applications that the user can specify any user name to the client program. For example, anyone can use the `mysql` program to connect as any other person simply by invoking it as `mysql -u other_user db_name` if `other_user` has no password. If all accounts have a password, connecting using another user's account becomes much more difficult.

For a discussion of methods for setting passwords, see [Section 6.3.5, “Assigning Account Passwords”](#).

- Make sure that the only Unix user account with read or write privileges in the database directories is the account that is used for running `mysqld`.
- Never run the MySQL server as the Unix `root` user. This is extremely dangerous, because any user with the `FILE [754]` privilege is able to cause the server to create files as `root` (for example, `~root/.bashrc`). To prevent this, `mysqld` refuses to run as `root` unless that is specified explicitly using the `--user=root [446]` option.

`mysqld` can (and should) be run as an ordinary, unprivileged user instead. You can create a separate Unix account named `mysql` to make everything even more secure. Use this account only for administering MySQL. To start `mysqld` as a different Unix user, add a `user` option that specifies the user name in the `[mysqld]` group of the `my.cnf` option file where you specify server options. For example:

```
[mysqld]
user=mysql
```

This causes the server to start as the designated user whether you start it manually or by using `mysqld_safe` or `mysql.server`. For more details, see [Section 6.1.5, “How to Run MySQL as a Normal User”](#).

Running `mysqld` as a Unix user other than `root` does not mean that you need to change the `root` user name in the `user` table. *User names for MySQL accounts have nothing to do with user names for Unix accounts.*

- Do not grant the [FILE \[754\]](#) privilege to nonadministrative users. Any user that has this privilege can write a file anywhere in the file system with the privileges of the `mysqld` daemon. This includes the server's data directory containing the files that implement the privilege tables. To make [FILE \[754\]](#)-privilege operations a bit safer, files generated with `SELECT ... INTO OUTFILE` do not overwrite existing files and are writable by everyone.

The [FILE \[754\]](#) privilege may also be used to read any file that is world-readable or accessible to the Unix user that the server runs as. With this privilege, you can read any file into a database table. This could be abused, for example, by using `LOAD DATA` to load `/etc/passwd` into a table, which then can be displayed with `SELECT`.

To limit the location in which files can be read and written, set the [secure\\_file\\_priv \[538\]](#) system to a specific directory. See [Section 5.1.4, “Server System Variables”](#).

- Do not grant the [PROCESS \[755\]](#) or [SUPER \[756\]](#) privilege to nonadministrative users. The output of `mysqladmin processlist` and `SHOW PROCESSLIST` shows the text of any statements currently being executed, so any user who is permitted to see the server process list might be able to see statements issued by other users such as `UPDATE user SET password=PASSWORD('not_secure')`.

`mysqld` reserves an extra connection for users who have the [SUPER \[756\]](#) privilege, so that a MySQL `root` user can log in and check server activity even if all normal connections are in use.

The [SUPER \[756\]](#) privilege can be used to terminate client connections, change server operation by changing the value of system variables, and control replication servers.

- Do not permit the use of symlinks to tables. (This capability can be disabled with the `--skip-symbolic-links [441]` option.) This is especially important if you run `mysqld` as `root`, because anyone that has write access to the server's data directory then could delete any file in the system! See [Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”](#).
- Stored programs and views should be written using the security guidelines discussed in [Section 18.6, “Access Control for Stored Programs and Views”](#).
- If you do not trust your DNS, you should use IP addresses rather than host names in the grant tables. In any case, you should be very careful about creating grant table entries using host name values that contain wildcards.
- If you want to restrict the number of connections permitted to a single account, you can do so by setting the [max\\_user\\_connections \[509\]](#) variable in `mysqld`. The `GRANT` statement also supports resource control options for limiting the extent of server use permitted to an account. See [Section 13.7.1.4, “GRANT Syntax”](#).
- If the plugin directory is writable by the server, it may be possible for a user to write executable code to a file in the directory using `SELECT ... INTO DUMPFILE`. This can be prevented by making [plugin\\_dir \[525\]](#) read only to the server or by setting `--secure-file-priv [438]` to a directory where `SELECT` writes can be made safely.

## 6.1.4. Security-Related `mysqld` Options and Variables

The following table shows `mysqld` options and system variables that affect security. For descriptions of each of these, see [Section 5.1.3, “Server Command Options”](#), and [Section 5.1.4, “Server System Variables”](#).

**Table 6.1. Security Option/Variable Summary**

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">allow-suspicious-udfs [413]</a>	Yes	Yes				
<a href="#">automatic_sp_privileges [464]</a>			Yes		Global	Yes
<a href="#">chroot [417]</a>	Yes	Yes				
<a href="#">des-key-file [421]</a>	Yes	Yes				
<a href="#">local_infile [496]</a>			Yes		Global	Yes
<a href="#">old_passwords [517]</a>			Yes		Both	Yes
<a href="#">safe-user-create [438]</a>	Yes	Yes				
<a href="#">secure-auth [438]</a>	Yes	Yes			Global	Yes
- Variable: secure_auth			Yes		Global	Yes
<a href="#">secure-file-priv [438]</a>	Yes	Yes			Global	No
- Variable: secure_file_priv			Yes		Global	No
<a href="#">skip-grant-tables [439]</a>	Yes	Yes				
<a href="#">skip-name-resolve [440]</a>	Yes	Yes			Global	No
- Variable: skip_name_resolve			Yes		Global	No
<a href="#">skip-networking [440]</a>	Yes	Yes			Global	No
- Variable: skip_networking			Yes		Global	No
<a href="#">skip-show-database [441]</a>	Yes	Yes			Global	No
- Variable: skip_show_database			Yes		Global	No

## 6.1.5. How to Run MySQL as a Normal User

On Windows, you can run the server as a Windows service using a normal user account.

On Unix, the MySQL server `mysqld` can be started and run by any user. However, you should avoid running the server as the Unix `root` user for security reasons. To change `mysqld` to run as a normal unprivileged Unix user `user_name`, you must do the following:

1. Stop the server if it is running (use `mysqladmin shutdown`).
2. Change the database directories and files so that `user_name` has privileges to read and write files in them (you might need to do this as the Unix `root` user):

```
shell> chown -R user_name /path/to/mysql/datadir
```

If you do not do this, the server will not be able to access databases or tables when it runs as `user_name`.

If directories or files within the MySQL data directory are symbolic links, `chown -R` might not follow symbolic links for you. If it does not, you will also need to follow those links and change the directories and files they point to.

3. Start the server as user `user_name`. Another alternative is to start `mysqld` as the Unix `root` user and use the `--user=user_name` [446] option. `mysqld` starts up, then switches to run as the Unix user `user_name` before accepting any connections.
4. To start the server as the given user automatically at system startup time, specify the user name by adding a `user` option to the `[mysqld]` group of the `/etc/my.cnf` option file or the `my.cnf` option file in the server's data directory. For example:

```
[mysqld]
user=user_name
```

If your Unix machine itself is not secured, you should assign passwords to the MySQL `root` accounts in the grant tables. Otherwise, any user with a login account on that machine can run the `mysql` client with a `--user=root` [254] option and perform any operation. (It is a good idea to assign passwords to MySQL accounts in any case, but especially so when other login accounts exist on the server host.) See [Section 2.10, "Postinstallation Setup and Testing"](#).

### 6.1.6. Security Issues with `LOAD DATA LOCAL`

The `LOAD DATA` statement can load a file that is located on the server host, or it can load a file that is located on the client host when the `LOCAL` keyword is specified.

There are two potential security issues with supporting the `LOCAL` version of `LOAD DATA` statements:

- The transfer of the file from the client host to the server host is initiated by the MySQL server. In theory, a patched server could be built that would tell the client program to transfer a file of the server's choosing rather than the file named by the client in the `LOAD DATA` statement. Such a server could access any file on the client host to which the client user has read access.
- In a Web environment where the clients are connecting from a Web server, a user could use `LOAD DATA LOCAL` to read any files that the Web server process has read access to (assuming that a user could run any command against the SQL server). In this environment, the client with respect to the MySQL server actually is the Web server, not the remote program being run by the user who connects to the Web server.

To deal with these problems, we changed how `LOAD DATA LOCAL` is handled as of MySQL 3.23.49 and MySQL 4.0.2 (4.0.13 on Windows):

- By default, all MySQL clients and libraries in binary distributions are compiled with the `-DENABLED_LOCAL_INFILE=1` [135] option, to be compatible with MySQL 3.23.48 and before.
- If you build MySQL from source but do not invoke `CMake` with the `-DENABLED_LOCAL_INFILE=1` [135] option, `LOAD DATA LOCAL` cannot be used by any client unless it is written explicitly to invoke `mysql_options(... MYSQL_OPT_LOCAL_INFILE, 0)`. See [Section 21.8.6.49, "mysql\\_options\(\)"](#).
- You can disable all `LOAD DATA LOCAL` statements from the server side by starting `mysqld` with the `--local-infile=0` [496] option.
- For the `mysql` command-line client, enable `LOAD DATA LOCAL` by specifying the `--local-infile[=1]` [250] option, or disable it with the `--local-infile=0` [250] option. For

`mysqlimport`, local data file loading is off by default; enable it with the `--local [300]` or `-L` option. In any case, successful use of a local load operation requires that the server permits it.

- If you use `LOAD DATA LOCAL` in Perl scripts or other programs that read the `[client]` group from option files, you can add the `local-infile=1` option to that group. However, to keep this from causing problems for programs that do not understand `local-infile`, specify it using the `loose-` prefix:

```
[client]
loose-local-infile=1
```

- If `LOAD DATA LOCAL` is disabled, either in the server or the client, a client that attempts to issue such a statement receives the following error message:

```
ERROR 1148: The used command is not allowed with this MySQL version
```

## 6.1.7. Client Programming Security Guidelines

Applications that access MySQL should not trust any data entered by users, who can try to trick your code by entering special or escaped character sequences in Web forms, URLs, or whatever application you have built. Be sure that your application remains secure if a user enters something like “`;  
DROP DATABASE mysql;`”. This is an extreme example, but large security leaks and data loss might occur as a result of hackers using similar techniques, if you do not prepare for them.

A common mistake is to protect only string data values. Remember to check numeric data as well. If an application generates a query such as `SELECT * FROM table WHERE ID=234` when a user enters the value `234`, the user can enter the value `234 OR 1=1` to cause the application to generate the query `SELECT * FROM table WHERE ID=234 OR 1=1`. As a result, the server retrieves every row in the table. This exposes every row and causes excessive server load. The simplest way to protect from this type of attack is to use single quotation marks around the numeric constants: `SELECT * FROM table WHERE ID='234'`. If the user enters extra information, it all becomes part of the string. In a numeric context, MySQL automatically converts this string to a number and strips any trailing nonnumeric characters from it.

Sometimes people think that if a database contains only publicly available data, it need not be protected. This is incorrect. Even if it is permissible to display any row in the database, you should still protect against denial of service attacks (for example, those that are based on the technique in the preceding paragraph that causes the server to waste resources). Otherwise, your server becomes unresponsive to legitimate users.

Checklist:

- Enable strict SQL mode to tell the server to be more restrictive of what data values it accepts. See [Section 5.1.7, “Server SQL Modes”](#).
- Try to enter single and double quotation marks (“’” and “””) in all of your Web forms. If you get any kind of MySQL error, investigate the problem right away.
- Try to modify dynamic URLs by adding `%22` (“”), `%23` (“#”), and `%27` (“’”) to them.
- Try to modify data types in dynamic URLs from numeric to character types using the characters shown in the previous examples. Your application should be safe against these and similar attacks.
- Try to enter characters, spaces, and special symbols rather than numbers in numeric fields. Your application should remove them before passing them to MySQL or else generate an error. Passing unchecked values to MySQL is very dangerous!
- Check the size of data before passing it to MySQL.
- Have your application connect to the database using a user name different from the one you use for administrative purposes. Do not give your applications any access privileges they do not need.

Many application programming interfaces provide a means of escaping special characters in data values. Properly used, this prevents application users from entering values that cause the application to generate statements that have a different effect than you intend:

- MySQL C API: Use the `mysql_real_escape_string()` API call.
- MySQL++: Use the `escape` and `quote` modifiers for query streams.
- PHP: Use either the `mysqli` or `pdo_mysql` extensions, and not the older `ext/mysql` extension. The preferred API's support the improved MySQL authentication protocol and passwords, as well as prepared statements with placeholders. See also [Section 21.9.1.3, "Choosing an API"](#).

If the older `ext/mysql` extension must be used, then for escaping use the `mysql_real_escape_string()` function and not `mysql_escape_string()` or `addslashes()` because only `mysql_real_escape_string()` is character set-aware; the other functions can be "bypassed" when using (invalid) multi-byte character sets.

- Perl DBI: Use placeholders or the `quote()` method.
- Ruby DBI: Use placeholders or the `quote()` method.
- Java JDBC: Use a `PreparedStatement` object and placeholders.

Other programming interfaces might have similar capabilities.

## 6.2. The MySQL Access Privilege System

The primary function of the MySQL privilege system is to authenticate a user who connects from a given host and to associate that user with privileges on a database such as `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. Additional functionality includes the ability to have anonymous users and to grant privileges for MySQL-specific functions such as `LOAD DATA INFILE` and administrative operations.

There are some things that you cannot do with the MySQL privilege system:

- You cannot explicitly specify that a given user should be denied access. That is, you cannot explicitly match a user and then refuse the connection.
- You cannot specify that a user has privileges to create or drop tables in a database but not to create or drop the database itself.
- A password applies globally to an account. You cannot associate a password with a specific object such as a database, table, or routine.

The user interface to the MySQL privilege system consists of SQL statements such as `CREATE USER`, `GRANT`, and `REVOKE`. See [Section 13.7.1, "Account Management Statements"](#).

Internally, the server stores privilege information in the grant tables of the `mysql` database (that is, in the database named `mysql`). The MySQL server reads the contents of these tables into memory when it starts and bases access-control decisions on the in-memory copies of the grant tables.

The MySQL privilege system ensures that all users may perform only the operations permitted to them. As a user, when you connect to a MySQL server, your identity is determined by *the host from which you connect* and *the user name you specify*. When you issue requests after connecting, the system grants privileges according to your identity and *what you want to do*.

MySQL considers both your host name and user name in identifying you because there is no reason to assume that a given user name belongs to the same person on all hosts. For example, the user `joe` who connects from `office.example.com` need not be the same person as the user `joe` who connects from `home.example.com`. MySQL handles this by enabling you to distinguish users on different hosts that happen to have the same name: You can grant one set of privileges for connections by `joe` from `office.example.com`, and a different set of privileges for connections by `joe` from



`home.example.com`. To see what privileges a given account has, use the `SHOW GRANTS` statement. For example:

```
SHOW GRANTS FOR 'joe'@'office.example.com';
SHOW GRANTS FOR 'joe'@'home.example.com';
```

MySQL access control involves two stages when you run a client program that connects to the server:

**Stage 1:** The server accepts or rejects the connection based on your identity and whether you can verify your identity by supplying the correct password.

**Stage 2:** Assuming that you can connect, the server checks each statement you issue to determine whether you have sufficient privileges to perform it. For example, if you try to select rows from a table in a database or drop a table from the database, the server verifies that you have the `SELECT` [755] privilege for the table or the `DROP` [754] privilege for the database.

For a more detailed description of what happens during each stage, see [Section 6.2.4, “Access Control, Stage 1: Connection Verification”](#), and [Section 6.2.5, “Access Control, Stage 2: Request Verification”](#).

If your privileges are changed (either by yourself or someone else) while you are connected, those changes do not necessarily take effect immediately for the next statement that you issue. For details about the conditions under which the server reloads the grant tables, see [Section 6.2.6, “When Privilege Changes Take Effect”](#).

For general security-related advice, see [Section 6.1, “General Security Issues”](#). For help in diagnosing privilege-related problems, see [Section 6.2.7, “Causes of Access-Denied Errors”](#).

## 6.2.1. Privileges Provided by MySQL

MySQL provides privileges that apply in different contexts and at different levels of operation:

- Administrative privileges enable users to manage operation of the MySQL server. These privileges are global because they are not specific to a particular database.
- Database privileges apply to a database and to all objects within it. These privileges can be granted for specific databases, or globally so that they apply to all databases.
- Privileges for database objects such as tables, indexes, views, and stored routines can be granted for specific objects within a database, for all objects of a given type within a database (for example, all tables in a database), or globally for all objects of a given type in all databases).

Information about account privileges is stored in the `user`, `db`, `tables_priv`, `columns_priv`, and `procs_priv` tables in the `mysql` database (see [Section 6.2.2, “Privilege System Grant Tables”](#)). The MySQL server reads the contents of these tables into memory when it starts and reloads them under the circumstances indicated in [Section 6.2.6, “When Privilege Changes Take Effect”](#). Access-control decisions are based on the in-memory copies of the grant tables.

Some releases of MySQL introduce changes to the structure of the grant tables to add new access privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

The following table shows the privilege names used at the SQL level in the `GRANT` and `REVOKE` statements, along with the column name associated with each privilege in the grant tables and the context in which the privilege applies.

**Table 6.2. Permissible Privileges for `GRANT` and `REVOKE`**

Privilege	Column	Context
<code>CREATE</code> [754]	<code>Create_priv</code>	databases, tables, or indexes

Privilege	Column	Context
DROP [754]	Drop_priv	databases, tables, or views
GRANT OPTION [755]	Grant_priv	databases, tables, or stored routines
LOCK TABLES [755]	Lock_tables_priv	databases
REFERENCES [755]	References_priv	databases or tables
EVENT [754]	Event_priv	databases
ALTER [753]	Alter_priv	tables
DELETE [754]	Delete_priv	tables
INDEX [755]	Index_priv	tables
INSERT [755]	Insert_priv	tables or columns
SELECT [755]	Select_priv	tables or columns
UPDATE [756]	Update_priv	tables or columns
CREATE TEMPORARY TABLES [754]	Create_tmp_table_priv	tables
TRIGGER [756]	Trigger_priv	tables
CREATE VIEW [754]	Create_view_priv	views
SHOW VIEW [756]	Show_view_priv	views
ALTER ROUTINE [754]	Alter_routine_priv	stored routines
CREATE ROUTINE [754]	Create_routine_priv	stored routines
EXECUTE [754]	Execute_priv	stored routines
FILE [754]	File_priv	file access on server host
CREATE TABLESPACE [754]	Create_tablespace_priv	server administration
CREATE USER [754]	Create_user_priv	server administration
PROCESS [755]	Process_priv	server administration
PROXY [755]	see proxies_priv table	server administration
RELOAD [755]	Reload_priv	server administration
REPLICATION CLIENT [755]	Repl_client_priv	server administration
REPLICATION SLAVE [755]	Repl_slave_priv	server administration
SHOW DATABASES [756]	Show_db_priv	server administration
SHUTDOWN [756]	Shutdown_priv	server administration
SUPER [756]	Super_priv	server administration
ALL [PRIVILEGES] [753]		server administration
USAGE [756]		server administration

The following list provides a general description of each privilege available in MySQL. Particular SQL statements might have more specific privilege requirements than indicated here. If so, the description for the statement in question provides the details.

- The `ALL [753]` or `ALL PRIVILEGES [753]` privilege specifier is shorthand. It stands for “all privileges available at a given privilege level” (except `GRANT OPTION [755]`). For example, granting `ALL [753]` at the global or table level grants all global privileges or all table-level privileges.
- The `ALTER [753]` privilege enables use of `ALTER TABLE` to change the structure of tables. `ALTER TABLE` also requires the `CREATE [754]` and `INSERT [755]` privileges. Renaming a table



requires `ALTER` [753] and `DROP` [754] on the old table, `ALTER` [753], `CREATE` [754], and `INSERT` [755] on the new table.

- The `ALTER ROUTINE` [754] privilege is needed to alter or drop stored routines (procedures and functions).
- The `CREATE` [754] privilege enables creation of new databases and tables.
- The `CREATE ROUTINE` [754] privilege is needed to create stored routines (procedures and functions).
- The `CREATE TABLESPACE` [754] privilege is needed to create, alter, or drop tablespaces and log file groups.
- The `CREATE TEMPORARY TABLES` [754] privilege enables the creation of temporary tables using the `CREATE TEMPORARY TABLE` statement.

As of MySQL 5.6.3, after a session has created a temporary table, the server performs no further privilege checks on the table. The creating session can perform any operation on the table, such as `DROP TABLE`, `INSERT`, `UPDATE`, or `SELECT`.

One implication of this behavior is that a session can manipulate its temporary tables even if the current user has no privilege to create them. Suppose that the current user does not have the `CREATE TEMPORARY TABLES` [754] privilege but is able to execute a `DEFINER`-context stored procedure that executes with the privileges of a user who does have `CREATE TEMPORARY TABLES` [754] and that creates a temporary table. While the procedure executes, the session uses the privileges of the defining user. After the procedure returns, the effective privileges revert to those of the current user, which can still see the temporary table and perform any operation on it.

Before MySQL 5.6.3, other operations on a temporary table, such as `INSERT`, `UPDATE`, or `SELECT`, require additional privileges for those operations for the database containing the temporary table, or for the nontemporary table of the same name.

To keep privileges for temporary and nontemporary tables separate, a common workaround for this situation is to create a database dedicated to the use of temporary tables. Then for that database, a user can be granted the `CREATE TEMPORARY TABLES` [754] privilege, along with any other privileges required for temporary table operations done by that user.

- The `CREATE USER` [754] privilege enables use of `CREATE USER`, `DROP USER`, `RENAME USER`, and `REVOKE ALL PRIVILEGES`.
- The `CREATE VIEW` [754] privilege enables use of `CREATE VIEW`.
- The `DELETE` [754] privilege enables rows to be deleted from tables in a database.
- The `DROP` [754] privilege enables you to drop (remove) existing databases, tables, and views. The `DROP` [754] privilege is required in order to use the statement `ALTER TABLE ... DROP PARTITION` on a partitioned table. The `DROP` [754] privilege is also required for `TRUNCATE TABLE`. *If you grant the `DROP` [754] privilege for the `mysql` database to a user, that user can drop the database in which the MySQL access privileges are stored.*
- The `EVENT` [754] privilege is required to create, alter, drop, or see events for the Event Scheduler.
- The `EXECUTE` [754] privilege is required to execute stored routines (procedures and functions).
- The `FILE` [754] privilege gives you permission to read and write files on the server host using the `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` statements and the `LOAD_FILE()` [1181] function. A user who has the `FILE` [754] privilege can read any file on the server host that is either world-readable or readable by the MySQL server. (This implies the user can read any file in any database directory, because the server can access any of those files.) The `FILE` [754] privilege also enables the user to create new files in any directory where the MySQL server has write access.

This includes the server's data directory containing the files that implement the privilege tables. As a security measure, the server will not overwrite existing files.

To limit the location in which files can be read and written, set the `secure_file_priv` [538] system to a specific directory. See [Section 5.1.4, “Server System Variables”](#).

- The `GRANT OPTION` [755] privilege enables you to give to other users or remove from other users those privileges that you yourself possess.
- The `INDEX` [755] privilege enables you to create or drop (remove) indexes. `INDEX` [755] applies to existing tables. If you have the `CREATE` [754] privilege for a table, you can include index definitions in the `CREATE TABLE` statement.
- The `INSERT` [755] privilege enables rows to be inserted into tables in a database. `INSERT` [755] is also required for the `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` table-maintenance statements.
- The `LOCK TABLES` [755] privilege enables the use of explicit `LOCK TABLES` statements to lock tables for which you have the `SELECT` [755] privilege. This includes the use of write locks, which prevents other sessions from reading the locked table.
- The `PROCESS` [755] privilege pertains to display of information about the threads executing within the server (that is, information about the statements being executed by sessions). The privilege enables use of `SHOW PROCESSLIST` or `mysqladmin processlist` to see threads belonging to other accounts; you can always see your own threads.
- The `PROXY` [755] privilege enables a user to impersonate or become known as another user. See [Section 6.3.8, “Proxy Users”](#).
- The `REFERENCES` [755] privilege currently is unused.
- The `RELOAD` [755] privilege enables use of the `FLUSH` statement. It also enables `mysqladmin` commands that are equivalent to `FLUSH` operations: `flush-hosts`, `flush-logs`, `flush-privileges`, `flush-status`, `flush-tables`, `flush-threads`, `refresh`, and `reload`.

The `reload` command tells the server to reload the grant tables into memory. `flush-privileges` is a synonym for `reload`. The `refresh` command closes and reopens the log files and flushes all tables. The other `flush-xxx` commands perform functions similar to `refresh`, but are more specific and may be preferable in some instances. For example, if you want to flush just the log files, `flush-logs` is a better choice than `refresh`.

- The `REPLICATION CLIENT` [755] privilege enables the use of `SHOW MASTER STATUS` and `SHOW SLAVE STATUS`. In MySQL 5.6.6 and later, it also enables the use of the `SHOW BINARY LOGS` statement.
- The `REPLICATION SLAVE` [755] privilege should be granted to accounts that are used by slave servers to connect to the current server as their master. Without this privilege, the slave cannot request updates that have been made to databases on the master server.
- The `SELECT` [755] privilege enables you to select rows from tables in a database. `SELECT` statements require the `SELECT` [755] privilege only if they actually retrieve rows from a table. Some `SELECT` statements do not access tables and can be executed without permission for any database. For example, you can use `SELECT` as a simple calculator to evaluate expressions that make no reference to tables:

```
SELECT 1+1;
SELECT PI()*2;
```

The `SELECT` [755] privilege is also needed for other statements that read column values. For example, `SELECT` [755] is needed for columns referenced on the right hand side of `col_name=expr` assignment in `UPDATE` statements or for columns named in the `WHERE` clause of `DELETE` or `UPDATE` statements.

- The `SHOW DATABASES` [756] privilege enables the account to see database names by issuing the `SHOW DATABASE` statement. Accounts that do not have this privilege see only databases for which they have some privileges, and cannot use the statement at all if the server was started with the `--skip-show-database` [441] option. Note that *any* global privilege is a privilege for the database.
- The `SHOW VIEW` [756] privilege enables use of `SHOW CREATE VIEW`.
- The `SHUTDOWN` [756] privilege enables use of the `mysqladmin shutdown` command. There is no corresponding SQL statement.
- The `SUPER` [756] privilege enables an account to use `CHANGE MASTER TO`, `KILL` or `mysqladmin kill` to kill threads belonging to other accounts (you can always kill your own threads), `PURGE BINARY LOGS`, configuration changes using `SET GLOBAL` to modify global system variables, the `mysqladmin debug` command, enabling or disabling logging, performing updates even if the `read_only` [532] system variable is enabled, starting and stopping replication on slave servers, specification of any account in the `DEFINER` attribute of stored programs and views, and enables you to connect (once) even if the connection limit controlled by the `max_connections` [503] system variable is reached.

To create or alter stored functions if binary logging is enabled, you may also need the `SUPER` [756] privilege, as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

- The `TRIGGER` [756] privilege enables trigger operations. You must have this privilege for a table to create, drop, or execute triggers for that table.
- The `UPDATE` [756] privilege enables rows to be updated in tables in a database.
- The `USAGE` [756] privilege specifier stands for “no privileges.” It is used at the global level with `GRANT` to modify account attributes such as resource limits or SSL characteristics without affecting existing account privileges.

It is a good idea to grant to an account only those privileges that it needs. You should exercise particular caution in granting the `FILE` [754] and administrative privileges:

- The `FILE` [754] privilege can be abused to read into a database table any files that the MySQL server can read on the server host. This includes all world-readable files and files in the server's data directory. The table can then be accessed using `SELECT` to transfer its contents to the client host.
- The `GRANT OPTION` [755] privilege enables users to give their privileges to other users. Two users that have different privileges and with the `GRANT OPTION` [755] privilege are able to combine privileges.
- The `ALTER` [753] privilege may be used to subvert the privilege system by renaming tables.
- The `SHUTDOWN` [756] privilege can be abused to deny service to other users entirely by terminating the server.
- The `PROCESS` [755] privilege can be used to view the plain text of currently executing statements, including statements that set or change passwords.
- The `SUPER` [756] privilege can be used to terminate other sessions or change how the server operates.
- Privileges granted for the `mysql` database itself can be used to change passwords and other access privilege information. Passwords are stored encrypted, so a malicious user cannot simply read them to know the plain text password. However, a user with write access to the `user` table `Password` column can change an account's password, and then connect to the MySQL server using that account.

## 6.2.2. Privilege System Grant Tables

Normally, you manipulate the contents of the grant tables in the `mysql` database indirectly by using statements such as `GRANT` and `REVOKE` to set up accounts and control the privileges available to

each one. See [Section 13.7.1, “Account Management Statements”](#). The discussion here describes the underlying structure of the grant tables and how the server uses their contents when interacting with clients.

These `mysql` database tables contain grant information:

- `user`: Contains user accounts, global privileges, and other non-privilege columns.
- `db`: Contains database-level privileges.
- `host`: Obsolete. New MySQL installations no longer create this table as of MySQL 5.6.7.
- `tables_priv`: Contains table-level privileges.
- `columns_priv`: Contains column-level privileges.
- `procs_priv`: Contains stored procedure and function privileges.
- `proxies_priv`: Contains proxy-user privileges.

Other tables in the `mysql` database do not hold grant information and are discussed elsewhere:

- `event`: Contains information about Event Scheduler events: See [Section 18.4, “Using the Event Scheduler”](#).
- `func`: Contains information about user-defined functions: See [Section 22.3, “Adding New Functions to MySQL”](#).
- `help_xxx`: These tables are used for server-side help: See [Section 5.1.10, “Server-Side Help”](#).
- `plugin`: Contains information about server plugins: See [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#), and [Section 22.2, “The MySQL Plugin API”](#).
- `proc`: Contains information about stored procedures and functions: See [Section 18.2, “Using Stored Routines \(Procedures and Functions\)”](#).
- `servers`: Used by the `FEDERATED` storage engine: See [Section 14.9.2.2, “Creating a FEDERATED Table Using CREATE SERVER”](#).
- `time_zone_xxx`: These tables contain time zone information: See [Section 10.6, “MySQL Server Time Zone Support”](#).
- Tables with `_log` in their name are used for logging: See [Section 5.2, “MySQL Server Logs”](#).

Each grant table contains scope columns and privilege columns:

- Scope columns determine the scope of each row (entry) in the tables; that is, the context in which the row applies. For example, a `user` table row with `Host` and `User` values of `'thomas.loc.gov'` and `'bob'` would be used for authenticating connections made to the server from the host `thomas.loc.gov` by a client that specifies a user name of `bob`. Similarly, a `db` table row with `Host`, `User`, and `Db` column values of `'thomas.loc.gov'`, `'bob'` and `'reports'` would be used when `bob` connects from the host `thomas.loc.gov` to access the `reports` database. The `tables_priv` and `columns_priv` tables contain scope columns indicating tables or table/column combinations to which each row applies. The `procs_priv` scope columns indicate the stored routine to which each row applies.
- Privilege columns indicate which privileges are granted by a table row; that is, what operations can be performed. The server combines the information in the various grant tables to form a complete description of a user's privileges. [Section 6.2.5, “Access Control, Stage 2: Request Verification”](#), describes the rules that are used to do this.

The server uses the grant tables in the following manner:

- The `user` table scope columns determine whether to reject or permit incoming connections. For permitted connections, any privileges granted in the `user` table indicate the user's global privileges. Any privilege granted in this table applies to *all* databases on the server.



**Note**

Because any global privilege is considered a privilege for all databases, any global privilege enables a user to see all database names with `SHOW DATABASES` or by examining the `SCHEMATA` table of `INFORMATION_SCHEMA`.

- The `db` table scope columns determine which users can access which databases from which hosts. The privilege columns determine which operations are permitted. A privilege granted at the database level applies to the database and to all objects in the database, such as tables and stored programs.
- The `tables_priv` and `columns_priv` tables are similar to the `db` table, but are more fine-grained: They apply at the table and column levels rather than at the database level. A privilege granted at the table level applies to the table and to all its columns. A privilege granted at the column level applies only to a specific column.
- The `procs_priv` table applies to stored routines. A privilege granted at the routine level applies only to a single routine.
- The `proxies_priv` table indicates which users can act as proxies for other users and whether proxy users can grant the `PROXY [755]` privilege to other users.

The server uses the `user` and `db` tables in the `mysql` database at both the first and second stages of access control (see Section 6.2, “The MySQL Access Privilege System”). The columns in the `user` and `db` tables are shown here.

**Table 6.3. `user` and `db` Table Columns**

Table Name	<code>user</code>	<code>db</code>
<b>Scope columns</b>	Host	Host
	User	Db
	Password	User
<b>Privilege columns</b>	Select_priv	Select_priv
	Insert_priv	Insert_priv
	Update_priv	Update_priv
	Delete_priv	Delete_priv
	Index_priv	Index_priv
	Alter_priv	Alter_priv
	Create_priv	Create_priv
	Drop_priv	Drop_priv
	Grant_priv	Grant_priv
	Create_view_priv	Create_view_priv
	Show_view_priv	Show_view_priv
	Create_routine_priv	Create_routine_priv
	Alter_routine_priv	Alter_routine_priv
	Execute_priv	Execute_priv
	Trigger_priv	Trigger_priv
	Event_priv	Event_priv
	Create_tmp_table_priv	Create_tmp_table_priv

Table Name	user	db
	Lock_tables_priv	Lock_tables_priv
	References_priv	References_priv
	Reload_priv	
	Shutdown_priv	
	Process_priv	
	File_priv	
	Show_db_priv	
	Super_priv	
	Repl_slave_priv	
	Repl_client_priv	
	Create_user_priv	
	Create_tablespace_priv	
<b>Security columns</b>	ssl_type	
	ssl_cipher	
	x509_issuer	
	x509_subject	
	plugin	
	authentication_string	
	password_expired	
<b>Resource control columns</b>	max_questions	
	max_updates	
	max_connections	
	max_user_connections	

The `mysql.user` table `plugin` and `authentication_string` columns store authentication plugin information.

If the `plugin` column for an account row is empty, the server uses native authentication for connection attempts for the account: Clients must match the password in the `Password` column of the account row.

If an account row names a plugin in the `plugin` column, the server uses it to authenticate connection attempts for the account. Whether the plugin uses the value in the `Password` column is up to the plugin.

The `password_expired` column was added in MySQL 5.6.6 to permit DBAs to expire account passwords and require users to reset their password. The default `password_expired` value is 'N', but can be set to 'Y' with the `ALTER USER` statement. After an account's password has been expired, all operations performed by the account in subsequent connections to the server result in an error until the user issues a `SET PASSWORD` statement to establish a new account password. See [Section 13.7.1.1, "ALTER USER Syntax"](#).

It is possible after password expiration to "reset" a password by using `SET PASSWORD` to set it to its current value. As a matter of good policy, it is preferable to choose a different password.



**Caution**

In MySQL 5.6.6, `ALTER USER` also sets the `Password` column to the empty string, so do not use this statement until 5.6.7.

During the second stage of access control, the server performs request verification to make sure that each client has sufficient privileges for each request that it issues. In addition to the `user` and `db` grant tables, the server may also consult the `tables_priv` and `columns_priv` tables for requests that involve tables. The latter tables provide finer privilege control at the table and column levels. They have the columns shown in the following table.

**Table 6.4. `tables_priv` and `columns_priv` Table Columns**

Table Name	<code>tables_priv</code>	<code>columns_priv</code>
<b>Scope columns</b>	<code>Host</code>	<code>Host</code>
	<code>Db</code>	<code>Db</code>
	<code>User</code>	<code>User</code>
	<code>Table_name</code>	<code>Table_name</code>
		<code>Column_name</code>
<b>Privilege columns</b>	<code>Table_priv</code>	<code>Column_priv</code>
	<code>Column_priv</code>	
<b>Other columns</b>	<code>Timestamp</code>	<code>Timestamp</code>
	<code>Grantor</code>	

The `Timestamp` and `Grantor` columns are set to the current timestamp and the `CURRENT_USER` [1267] value, respectively. However, they are unused and are discussed no further here.

For verification of requests that involve stored routines, the server may consult the `procs_priv` table, which has the columns shown in the following table.

**Table 6.5. `procs_priv` Table Columns**

Table Name	<code>procs_priv</code>
<b>Scope columns</b>	<code>Host</code>
	<code>Db</code>
	<code>User</code>
	<code>Routine_name</code>
	<code>Routine_type</code>
<b>Privilege columns</b>	<code>Proc_priv</code>
<b>Other columns</b>	<code>Timestamp</code>
	<code>Grantor</code>

The `Routine_type` column is an `ENUM` column with values of `'FUNCTION'` or `'PROCEDURE'` to indicate the type of routine the row refers to. This column enables privileges to be granted separately for a function and a procedure with the same name.

The `Timestamp` and `Grantor` columns currently are unused and are discussed no further here.

The `proxies_priv` table records information about proxy users. It has these columns:

- `Host`, `User`: These columns indicate the user account that has the `PROXY` [755] privilege for the proxied account.
- `Proxied_host`, `Proxied_user`: These columns indicate the account of the proxied user.
- `Grantor`: Currently unused.
- `Timestamp`: Currently unused.



- `With_grant`: This column indicates whether the proxy account can grant the `PROXY` [755] privilege to other accounts.

Scope columns in the grant tables contain strings. They are declared as shown here; the default value for each is the empty string.

**Table 6.6. Grant Table Scope Column Types**

Column Name	Type
<code>Host, Proxied_host</code>	<code>CHAR(60)</code>
<code>User, Proxied_user</code>	<code>CHAR(16)</code>
<code>Password</code>	<code>CHAR(41)</code>
<code>Db</code>	<code>CHAR(64)</code>
<code>Table_name</code>	<code>CHAR(64)</code>
<code>Column_name</code>	<code>CHAR(64)</code>
<code>Routine_name</code>	<code>CHAR(64)</code>

For access-checking purposes, comparisons of `User`, `Proxied_user`, `Password`, `Db`, and `Table_name` values are case sensitive. Comparisons of `Host`, `Proxied_host`, `Column_name`, and `Routine_name` values are not case sensitive.

In the `user` and `db` tables, each privilege is listed in a separate column that is declared as `ENUM('N', 'Y') DEFAULT 'N'`. In other words, each privilege can be disabled or enabled, with the default being disabled.

In the `tables_priv`, `columns_priv`, and `procs_priv` tables, the privilege columns are declared as `SET` columns. Values in these columns can contain any combination of the privileges controlled by the table. Only those privileges listed in the column value are enabled.

**Table 6.7. Set-Type Privilege Column Values**

Table Name	Column Name	Possible Set Elements
<code>tables_priv</code>	<code>Table_priv</code>	<code>'Select', 'Insert', 'Update', 'Delete', 'Create', 'Drop', 'Grant', 'References', 'Index', 'Alter', 'Create View', 'Show view', 'Trigger'</code>
<code>tables_priv</code>	<code>Column_priv</code>	<code>'Select', 'Insert', 'Update', 'References'</code>
<code>columns_priv</code>	<code>Column_priv</code>	<code>'Select', 'Insert', 'Update', 'References'</code>
<code>procs_priv</code>	<code>Proc_priv</code>	<code>'Execute', 'Alter Routine', 'Grant'</code>

Administrative privileges (such as `RELOAD` [755] or `SHUTDOWN` [756]) are specified only in the `user` table. Administrative operations are operations on the server itself and are not database-specific, so there is no reason to list these privileges in the other grant tables. Consequently, to determine whether you can perform an administrative operation, the server need consult only the `user` table.

The `FILE` [754] privilege also is specified only in the `user` table. It is not an administrative privilege as such, but your ability to read or write files on the server host is independent of the database you are accessing.

The `mysqld` server reads the contents of the grant tables into memory when it starts. You can tell it to reload the tables by issuing a `FLUSH PRIVILEGES` statement or executing a `mysqladmin flush-privileges` or `mysqladmin reload` command. Changes to the grant tables take effect as indicated in Section 6.2.6, “When Privilege Changes Take Effect”.



When you modify an account's privileges, it is a good idea to verify that the changes set up privileges the way you want. To check the privileges for a given account, use the `SHOW GRANTS` statement (see [Section 13.7.5.22, "SHOW GRANTS Syntax"](#)). For example, to determine the privileges that are granted to an account with user name and host name values of `bob` and `pc84.example.com`, use this statement:

```
SHOW GRANTS FOR 'bob'@'pc84.example.com';
```

### 6.2.3. Specifying Account Names

MySQL account names consist of a user name and a host name. This enables creation of accounts for users with the same name who can connect from different hosts. This section describes how to write account names, including special values and wildcard rules.

In SQL statements such as `CREATE USER`, `GRANT`, and `SET PASSWORD`, write account names using the following rules:

- Syntax for account names is `'user_name'@'host_name'`.
- An account name consisting only of a user name is equivalent to `'user_name'@'%'`. For example, `'me'` is equivalent to `'me'@'%'`.
- The user name and host name need not be quoted if they are legal as unquoted identifiers. Quotes are necessary to specify a `user_name` string containing special characters (such as `"-`"), or a `host_name` string containing special characters or wildcard characters (such as `"%"`); for example, `'test-user'@'%.com'`.
- Quote user names and host names as identifiers or as strings, using either backticks (`"`"`), single quotation marks (`"'"`), or double quotation marks (`""`).
- The user name and host name parts, if quoted, must be quoted separately. That is, write `'me'@'localhost'`, not `'me@localhost'`; the latter is interpreted as `'me@localhost'@'%'`.
- A reference to the `CURRENT_USER [1267]` or `CURRENT_USER() [1267]` function is equivalent to specifying the current client's user name and host name literally.

MySQL stores account names in grant tables in the `mysql` database using separate columns for the user name and host name parts:

- The `user` table contains one row for each account. The `User` and `Host` columns store the user name and host name. This table also indicates which global privileges the account has.
- Other grant tables indicate privileges an account has for databases and objects within databases. These tables have `User` and `Host` columns to store the account name. Each row in these tables associates with the account in the `user` table that has the same `User` and `Host` values.

For additional detail about grant table structure, see [Section 6.2.2, "Privilege System Grant Tables"](#).

User names and host names have certain special values or wildcard conventions, as described following.

A user name is either a nonblank value that literally matches the user name for incoming connection attempts, or a blank value (empty string) that matches any user name. An account with a blank user name is an anonymous user. To specify an anonymous user in SQL statements, use a quoted empty user name part, such as `'@'localhost'`.

The host name part of an account name can take many forms, and wildcards are permitted:

- A host value can be a host name or an IP address (IPv4 or IPv6). The name `'localhost'` indicates the local host. The IP address `'127.0.0.1'` indicates the IPv4 loopback interface. The IP address `:::1` indicates the IPv6 loopback interface.
- You can use the wildcard characters `"%"` and `"_"` in host name or IP address values. These have the same meaning as for pattern-matching operations performed with the `LIKE [1190]` operator.

For example, a host value of `'%'` matches any host name, whereas a value of `'%.mysql.com'` matches any host in the `mysql.com` domain. `'192.168.1.%'` matches any host in the 192.168.1 class C network.

Because you can use IP wildcard values in host values (for example, `'192.168.1.%'` to match every host on a subnet), someone could try to exploit this capability by naming a host `192.168.1.somewhere.com`. To foil such attempts, MySQL disallows matching on host names that start with digits and a dot. Thus, if you have a host named something like `1.2.example.com`, its name never matches the host part of account names. An IP wildcard value can match only IP addresses, not host names.

- For a host value specified as an IPv4 address, you can specify a netmask indicating how many address bits to use for the network number. Netmask notation cannot be used for IPv6 addresses.

The syntax is `host_ip/netmask`. For example:

```
CREATE USER 'david'@'192.58.197.0/255.255.255.0';
```

This enables `david` to connect from any client host having an IP address `client_ip` for which the following condition is true:

```
client_ip & netmask = host_ip
```

That is, for the `CREATE USER` statement just shown:

```
client_ip & 255.255.255.0 = 192.58.197.0
```

IP addresses that satisfy this condition and can connect to the MySQL server are those in the range from `192.58.197.0` to `192.58.197.255`.

The netmask can only be used to tell the server to use 8, 16, 24, or 32 bits of the address. Examples:

- `192.0.0.0/255.0.0.0`: Any host on the 192 class A network
- `192.168.0.0/255.255.0.0`: Any host on the 192.168 class B network
- `192.168.1.0/255.255.255.0`: Any host on the 192.168.1 class C network
- `192.168.1.1`: Only the host with this specific IP address

The following netmask will not work because it masks 28 bits, and 28 is not a multiple of 8:

```
192.168.0.1/255.255.255.240
```

The server performs matching of host values in account names against the client host using the value returned by the system DNS resolver for the client host name or IP address. Except in the case that the account host value is specified using netmask notation, this comparison is performed as a string match, even for an account host value given as an IP address. This means that you should specify account host values in the same format used by DNS. Here are examples of problems to watch out for:

- Suppose that a host on the local network has a fully qualified name of `host1.example.com`. If DNS returns name lookups for this host as `host1.example.com`, use that name in account host values. But if DNS returns just `host1`, use `host1` instead.
- If DNS returns the IP address for a given host as `192.168.1.2`, that will match an account host value of `192.168.1.2` but not `192.168.01.2`. Similarly, it will match an account host pattern like `192.168.1.%` but not `192.168.01.%`.

To avoid problems like this, it is advisable to check the format in which your DNS returns host names and addresses, and use values in the same format in MySQL account names.

## 6.2.4. Access Control, Stage 1: Connection Verification

When you attempt to connect to a MySQL server, the server accepts or rejects the connection based on your identity and whether you can verify your identity by supplying the correct password. If not, the server denies access to you completely. Otherwise, the server accepts the connection, and then enters Stage 2 and waits for requests.

Your identity is based on two pieces of information:

- The client host from which you connect
- Your MySQL user name

Identity checking is performed using the three `user` table scope columns (`Host`, `User`, and `Password`). The server accepts the connection only if the `Host` and `User` columns in some `user` table row match the client host name and user name and the client supplies the password specified in that row. The rules for permissible `Host` and `User` values are given in [Section 6.2.3, “Specifying Account Names”](#).

If the `User` column value is nonblank, the user name in an incoming connection must match exactly. If the `User` value is blank, it matches any user name. If the `user` table row that matches an incoming connection has a blank user name, the user is considered to be an anonymous user with no name, not a user with the name that the client actually specified. This means that a blank user name is used for all further access checking for the duration of the connection (that is, during Stage 2).

The `Password` column can be blank. This is not a wildcard and does not mean that any password matches. It means that the user must connect without specifying a password. If the server authenticates a client using a plugin, the authentication method that the plugin implements may or may not use the password in the `Password` column. In this case, it is possible that an external password is also used to authenticate to the MySQL server.

Nonblank `Password` values in the `user` table represent encrypted passwords. MySQL does not store passwords in plaintext form for anyone to see. Rather, the password supplied by a user who is attempting to connect is encrypted (using the `PASSWORD()` [1263] function). The encrypted password then is used during the connection process when checking whether the password is correct. This is done without the encrypted password ever traveling over the connection. See [Section 6.3.1, “User Names and Passwords”](#).

From MySQL's point of view, the encrypted password is the *real* password, so you should never give anyone access to it. In particular, *do not give nonadministrative users read access to tables in the `mysql` database*.

The following table shows how various combinations of `Host` and `User` values in the `user` table apply to incoming connections.

Host Value	User Value	Permissible Connections
'thomas.loc.gov'	'fred'	fred, connecting from thomas.loc.gov
'thomas.loc.gov'	' '	Any user, connecting from thomas.loc.gov
'%'	'fred'	fred, connecting from any host
'%'	' '	Any user, connecting from any host
'%.loc.gov'	'fred'	fred, connecting from any host in the loc.gov domain
'x.y.%'	'fred'	fred, connecting from x.y.net, x.y.com, x.y.edu, and so on; this is probably not useful
'144.155.166.177'	'fred'	fred, connecting from the host with IP address 144.155.166.177
'144.155.166.%'	'fred'	fred, connecting from any host in the 144.155.166 class C subnet
'144.155.166.0/255.255.255.0'	'fred'	Same as previous example

It is possible for the client host name and user name of an incoming connection to match more than one row in the `user` table. The preceding set of examples demonstrates this: Several of the entries shown match a connection from `thomas.loc.gov` by `fred`.

When multiple matches are possible, the server must determine which of them to use. It resolves this issue as follows:

- Whenever the server reads the `user` table into memory, it sorts the rows.
- When a client attempts to connect, the server looks through the rows in sorted order.
- The server uses the first row that matches the client host name and user name.

The server uses sorting rules that order rows with the most-specific `Host` values first. Literal host names and IP addresses are the most specific. (The specificity of a literal IP address is not affected by whether it has a netmask, so `192.168.1.13` and `192.168.1.0/255.255.255.0` are considered equally specific.) The pattern `'%'` means “any host” and is least specific. The empty string `' '` also means “any host” but sorts after `'%'`. Rows with the same `Host` value are ordered with the most-specific `User` values first (a blank `User` value means “any user” and is least specific).

To see how this works, suppose that the `user` table looks like this:

```
+-----+-----+
| Host      | User      | ...
+-----+-----+
| %         | root      | ...
| %         | jeffrey   | ...
| localhost | root      | ...
| localhost |           | ...
+-----+-----+
```

When the server reads the table into memory, it sorts the rows using the rules just described. The result after sorting looks like this:

```
+-----+-----+
| Host      | User      | ...
+-----+-----+
| localhost | root      | ...
| localhost |           | ...
| %         | jeffrey   | ...
| %         | root      | ...
+-----+-----+
```

When a client attempts to connect, the server looks through the sorted rows and uses the first match found. For a connection from `localhost` by `jeffrey`, two of the rows from the table match: the one with `Host` and `User` values of `'localhost'` and `' '`, and the one with values of `'%'` and `'jeffrey'`. The `'localhost'` row appears first in sorted order, so that is the one the server uses.

Here is another example. Suppose that the `user` table looks like this:

```
+-----+-----+
| Host      | User      | ...
+-----+-----+
| %         | jeffrey   | ...
| thomas.loc.gov |         | ...
+-----+-----+
```

The sorted table looks like this:

```
+-----+-----+
| Host      | User      | ...
+-----+-----+
| thomas.loc.gov |         | ...
| %         | jeffrey   | ...
+-----+-----+
```

A connection by `jeffrey` from `thomas.loc.gov` is matched by the first row, whereas a connection by `jeffrey` from any host is matched by the second.



#### Note

It is a common misconception to think that, for a given user name, all rows that explicitly name that user are used first when the server attempts to find a match for the connection. This is not true. The preceding example illustrates this, where a connection from `thomas.loc.gov` by `jeffrey` is first matched not by the row containing `'jeffrey'` as the `User` column value, but by the row with no user name. As a result, `jeffrey` is authenticated as an anonymous user, even though he specified a user name when connecting.

If you are able to connect to the server, but your privileges are not what you expect, you probably are being authenticated as some other account. To find out what account the server used to authenticate you, use the `CURRENT_USER()` [1267] function. (See [Section 12.14, “Information Functions”](#).) It returns a value in `user_name@host_name` format that indicates the `User` and `Host` values from the matching `user` table row. Suppose that `jeffrey` connects and issues the following query:

```
mysql> SELECT CURRENT_USER();
+-----+
| CURRENT_USER() |
+-----+
| @localhost     |
+-----+
```

The result shown here indicates that the matching `user` table row had a blank `User` column value. In other words, the server is treating `jeffrey` as an anonymous user.

Another way to diagnose authentication problems is to print out the `user` table and sort it by hand to see where the first match is being made.

## 6.2.5. Access Control, Stage 2: Request Verification

After you establish a connection, the server enters Stage 2 of access control. For each request that you issue through that connection, the server determines what operation you want to perform, then checks whether you have sufficient privileges to do so. This is where the privilege columns in the grant tables come into play. These privileges can come from any of the `user`, `db`, `tables_priv`, `columns_priv`, or `procs_priv` tables. (You may find it helpful to refer to [Section 6.2.2, “Privilege System Grant Tables”](#), which lists the columns present in each of the grant tables.)

The `user` table grants privileges that are assigned to you on a global basis and that apply no matter what the default database is. For example, if the `user` table grants you the `DELETE` [754] privilege, you can delete rows from any table in any database on the server host! It is wise to grant privileges in the `user` table only to people who need them, such as database administrators. For other users, you should leave all privileges in the `user` table set to `'N'` and grant privileges at more specific levels only. You can grant privileges for particular databases, tables, columns, or routines.

The `db` table grants database-specific privileges. Values in the scope columns of this table can take the following forms:

- A blank `User` value matches the anonymous user. A nonblank value matches literally; there are no wildcards in user names.
- The wildcard characters `“%”` and `“_”` can be used in the `Host` and `Db` columns. These have the same meaning as for pattern-matching operations performed with the `LIKE` [1190] operator. If you want to use either character literally when granting privileges, you must escape it with a backslash. For example, to include the underscore character (`“_”`) as part of a database name, specify it as `“\_”` in the `GRANT` statement.
- A `'%'` or blank `Host` value means “any host.”
- A `'%'` or blank `Db` value means “any database.”

The server reads the `db` table into memory and sorts it at the same time that it reads the `user` table. The server sorts the `db` table based on the `Host`, `Db`, and `User` scope columns. As with the `user` table, sorting puts the most-specific values first and least-specific values last, and when the server looks for matching entries, it uses the first match that it finds.

The `tables_priv`, `columns_priv`, and `procs_priv` tables grant table-specific, column-specific, and routine-specific privileges. Values in the scope columns of these tables can take the following forms:

- The wildcard characters “%” and “\_” can be used in the `Host` column. These have the same meaning as for pattern-matching operations performed with the `LIKE` [1190] operator.
- A ‘%’ or blank `Host` value means “any host.”
- The `Db`, `Table_name`, `Column_name`, and `Routine_name` columns cannot contain wildcards or be blank.

The server sorts the `tables_priv`, `columns_priv`, and `procs_priv` tables based on the `Host`, `Db`, and `User` columns. This is similar to `db` table sorting, but simpler because only the `Host` column can contain wildcards.

The server uses the sorted tables to verify each request that it receives. For requests that require administrative privileges such as `SHUTDOWN` [756] or `RELOAD` [755], the server checks only the `user` table row because that is the only table that specifies administrative privileges. The server grants access if the row permits the requested operation and denies access otherwise. For example, if you want to execute `mysqladmin shutdown` but your `user` table row does not grant the `SHUTDOWN` [756] privilege to you, the server denies access without even checking the `db` table. (It contains no `Shutdown_priv` column, so there is no need to do so.)

For database-related requests (`INSERT`, `UPDATE`, and so on), the server first checks the user's global privileges by looking in the `user` table row. If the row permits the requested operation, access is granted. If the global privileges in the `user` table are insufficient, the server determines the user's database-specific privileges by checking the `db` table:

The server looks in the `db` table for a match on the `Host`, `Db`, and `User` columns. The `Host` and `User` columns are matched to the connecting user's host name and MySQL user name. The `Db` column is matched to the database that the user wants to access. If there is no row for the `Host` and `User`, access is denied.

After determining the database-specific privileges granted by the `db` table entries, the server adds them to the global privileges granted by the `user` table. If the result permits the requested operation, access is granted. Otherwise, the server successively checks the user's table and column privileges in the `tables_priv` and `columns_priv` tables, adds those to the user's privileges, and permits or denies access based on the result. For stored-routine operations, the server uses the `procs_priv` table rather than `tables_priv` and `columns_priv`.

Expressed in boolean terms, the preceding description of how a user's privileges are calculated may be summarized like this:

```
global privileges
OR (database privileges AND host privileges)
OR table privileges
OR column privileges
OR routine privileges
```

It may not be apparent why, if the global `user` row privileges are initially found to be insufficient for the requested operation, the server adds those privileges to the database, table, and column privileges later. The reason is that a request might require more than one type of privilege. For example, if you execute an `INSERT INTO ... SELECT` statement, you need both the `INSERT` [755] and the `SELECT` [755] privileges. Your privileges might be such that the `user` table row grants one privilege and the `db` table row grants the other. In this case, you have the necessary privileges to perform the



request, but the server cannot tell that from either table by itself; the privileges granted by the entries in both tables must be combined.

## 6.2.6. When Privilege Changes Take Effect

When `mysqld` starts, it reads all grant table contents into memory. The in-memory tables become effective for access control at that point.

If you modify the grant tables indirectly using account-management statements such as `GRANT`, `REVOKE`, `SET PASSWORD`, or `RENAME USER`, the server notices these changes and loads the grant tables into memory again immediately.

If you modify the grant tables directly using statements such as `INSERT`, `UPDATE`, or `DELETE`, your changes have no effect on privilege checking until you either restart the server or tell it to reload the tables. If you change the grant tables directly but forget to reload them, your changes have *no effect* until you restart the server. This may leave you wondering why your changes seem to make no difference!

To tell the server to reload the grant tables, perform a flush-privileges operation. This can be done by issuing a `FLUSH PRIVILEGES` statement or by executing a `mysqladmin flush-privileges` or `mysqladmin reload` command.

A grant table reload affects privileges for each existing client connection as follows:

- Table and column privilege changes take effect with the client's next request.
- Database privilege changes take effect the next time the client executes a `USE db_name` statement.



### Note

Client applications may cache the database name; thus, this effect may not be visible to them without actually changing to a different database or flushing the privileges.

- Global privileges and passwords are unaffected for a connected client. These changes take effect only for subsequent connections.

If the server is started with the `--skip-grant-tables` [439] option, it does not read the grant tables or implement any access control. Anyone can connect and do anything, *which is insecure*. To cause a server thus started to read the tables and enable access checking, flush the privileges.

## 6.2.7. Causes of Access-Denied Errors

If you encounter problems when you try to connect to the MySQL server, the following items describe some courses of action you can take to correct the problem.

- Make sure that the server is running. If it is not, clients cannot connect to it. For example, if an attempt to connect to the server fails with a message such as one of those following, one cause might be that the server is not running:

```
shell> mysql
ERROR 2003: Can't connect to MySQL server on 'host_name' (111)
shell> mysql
ERROR 2002: Can't connect to local MySQL server through socket
'/tmp/mysql.sock' (111)
```

- It might be that the server is running, but you are trying to connect using a TCP/IP port, named pipe, or Unix socket file different from the one on which the server is listening. To correct this when you invoke a client program, specify a `--port` [211] option to indicate the proper port number, or a `--socket` [211] option to indicate the proper named pipe or Unix socket file. To find out where the socket file is, you can use this command:

```
shell> netstat -ln | grep mysql
```

- Make sure that the server has not been configured to ignore network connections or (if you are attempting to connect remotely) that it has not been configured to listen only locally on its network interfaces. If the server was started with `--skip-networking` [440], it will not accept TCP/IP connections at all. If the server was started with `--bind-address=127.0.0.1` [414], it will listen for TCP/IP connections only locally on the loopback interface and will not accept remote connections.
- Check to make sure that there is no firewall blocking access to MySQL. Your firewall may be configured on the basis of the application being executed, or the port number used by MySQL for communication (3306 by default). Under Linux or Unix, check your IP tables (or similar) configuration to ensure that the port has not been blocked. Under Windows, applications such as ZoneAlarm or the Windows XP personal firewall may need to be configured not to block the MySQL port.
- The grant tables must be properly set up so that the server can use them for access control. For some distribution types (such as binary distributions on Windows, or RPM distributions on Linux), the installation process initializes the `mysql` database containing the grant tables. For distributions that do not do this, you must initialize the grant tables manually by running the `mysql_install_db` script. For details, see [Section 2.10.1, “Unix Postinstallation Procedures”](#).

To determine whether you need to initialize the grant tables, look for a `mysql` directory under the data directory. (The data directory normally is named `data` or `var` and is located under your MySQL installation directory.) Make sure that you have a file named `user.MYD` in the `mysql` database directory. If not, execute the `mysql_install_db` script. After running this script and starting the server, test the initial privileges by executing this command:

```
shell> mysql -u root test
```

The server should let you connect without error.

- After a fresh installation, you should connect to the server and set up your users and their access permissions:

```
shell> mysql -u root mysql
```

The server should let you connect because the MySQL `root` user has no password initially. That is also a security risk, so setting the password for the `root` accounts is something you should do while you're setting up your other MySQL accounts. For instructions on setting the initial passwords, see [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

- If you have updated an existing MySQL installation to a newer version, did you run the `mysql_upgrade` script? If not, do so. The structure of the grant tables changes occasionally when new capabilities are added, so after an upgrade you should always make sure that your tables have the current structure. For instructions, see [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).
- If a client program receives the following error message when it tries to connect, it means that the server expects passwords in a newer format than the client is capable of generating:

```
shell> mysql
Client does not support authentication protocol requested
by server; consider upgrading MySQL client
```

For information on how to deal with this, see [Section 6.1.2.4, “Password Hashing in MySQL”](#), and [Section C.5.2.4, “Client does not support authentication protocol”](#).

- Remember that client programs use connection parameters specified in option files or environment variables. If a client program seems to be sending incorrect default connection parameters when you have not specified them on the command line, check any applicable option files and your environment. For example, if you get `Access denied` when you run a client without any options, make sure that you have not specified an old password in any of your option files!

You can suppress the use of option files by a client program by invoking it with the `--no-defaults` [220] option. For example:



```
shell> mysqladmin --no-defaults -u root version
```

The option files that clients use are listed in [Section 4.2.3.3, “Using Option Files”](#). Environment variables are listed in [Section 2.12, “Environment Variables”](#).

- If you get the following error, it means that you are using an incorrect `root` password:

```
shell> mysqladmin -u root -pxxxx ver
Access denied for user 'root'@'localhost' (using password: YES)
```

If the preceding error occurs even when you have not specified a password, it means that you have an incorrect password listed in some option file. Try the `--no-defaults` [\[220\]](#) option as described in the previous item.

For information on changing passwords, see [Section 6.3.5, “Assigning Account Passwords”](#).

If you have lost or forgotten the `root` password, see [Section C.5.4.1, “How to Reset the Root Password”](#).

- If you change a password by using `SET PASSWORD`, `INSERT`, or `UPDATE`, you must encrypt the password using the `PASSWORD()` [\[1263\]](#) function. If you do not use `PASSWORD()` [\[1263\]](#) for these statements, the password will not work. For example, the following statement assigns a password, but fails to encrypt it, so the user is not able to connect afterward:

```
SET PASSWORD FOR 'abe'@'host_name' = 'eagle';
```

Instead, set the password like this:

```
SET PASSWORD FOR 'abe'@'host_name' = PASSWORD('eagle');
```

The `PASSWORD()` [\[1263\]](#) function is unnecessary when you specify a password using the `CREATE USER` or `GRANT` statements or the `mysqladmin password` command. Each of those automatically uses `PASSWORD()` [\[1263\]](#) to encrypt the password. See [Section 6.3.5, “Assigning Account Passwords”](#), and [Section 13.7.1.2, “CREATE USER Syntax”](#).

- `localhost` is a synonym for your local host name, and is also the default host to which clients try to connect if you specify no host explicitly.

To avoid this problem on such systems, you can use a `--host=127.0.0.1` [\[211\]](#) option to name the server host explicitly. This will make a TCP/IP connection to the local `mysqld` server. You can also use TCP/IP by specifying a `--host` [\[211\]](#) option that uses the actual host name of the local host. In this case, the host name must be specified in a `user` table row on the server host, even though you are running the client program on the same host as the server.

- The `Access denied` error message tells you who you are trying to log in as, the client host from which you are trying to connect, and whether you were using a password. Normally, you should have one row in the `user` table that exactly matches the host name and user name that were given in the error message. For example, if you get an error message that contains `using password: NO`, it means that you tried to log in without a password.
- If you get an `Access denied` error when trying to connect to the database with `mysql -u user_name`, you may have a problem with the `user` table. Check this by executing `mysql -u root mysql` and issuing this SQL statement:

```
SELECT * FROM user;
```

The result should include a row with the `Host` and `User` columns matching your client's host name and your MySQL user name.

- If the following error occurs when you try to connect from a host other than the one on which the MySQL server is running, it means that there is no row in the `user` table with a `Host` value that matches the client host:

```
Host ... is not allowed to connect to this MySQL server
```

You can fix this by setting up an account for the combination of client host name and user name that you are using when trying to connect.

If you do not know the IP address or host name of the machine from which you are connecting, you should put a row with '%' as the `Host` column value in the `user` table. After trying to connect from the client machine, use a `SELECT USER()` query to see how you really did connect. Then change the '%' in the `user` table row to the actual host name that shows up in the log. Otherwise, your system is left insecure because it permits connections from any host for the given user name.

On Linux, another reason that this error might occur is that you are using a binary MySQL version that is compiled with a different version of the `glibc` library than the one you are using. In this case, you should either upgrade your operating system or `glibc`, or download a source distribution of MySQL version and compile it yourself. A source RPM is normally trivial to compile and install, so this is not a big problem.

- If you specify a host name when trying to connect, but get an error message where the host name is not shown or is an IP address, it means that the MySQL server got an error when trying to resolve the IP address of the client host to a name:

```
shell> mysqladmin -u root -pXXXX -h some_hostname ver
Access denied for user 'root'@' ' (using password: YES)
```

If you try to connect as `root` and get the following error, it means that you do not have a row in the `user` table with a `User` column value of 'root' and that `mysqld` cannot resolve the host name for your client:

```
Access denied for user ''@'unknown'
```

These errors indicate a DNS problem. To fix it, execute `mysqladmin flush-hosts` to reset the internal DNS host cache. See [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#).

Some permanent solutions are:

- Determine what is wrong with your DNS server and fix it.
- Specify IP addresses rather than host names in the MySQL grant tables.
- Put an entry for the client machine name in `/etc/hosts` on Unix or `\windows\hosts` on Windows.
- Start `mysqld` with the `--skip-name-resolve` [\[440\]](#) option.
- Start `mysqld` with the `--skip-host-cache` [\[439\]](#) option.
- On Unix, if you are running the server and the client on the same machine, connect to `localhost`. Unix connections to `localhost` use a Unix socket file rather than TCP/IP.
- On Windows, if you are running the server and the client on the same machine and the server supports named pipe connections, connect to the host name `.` (period). Connections to `.` use a named pipe rather than TCP/IP.
- If `mysql -u root test` works but `mysql -h your_hostname -u root test` results in `Access denied` (where `your_hostname` is the actual host name of the local host), you may not have the correct name for your host in the `user` table. A common problem here is that the `Host` value in the `user` table row specifies an unqualified host name, but your system's name resolution routines return a fully qualified domain name (or vice versa). For example, if you have an entry with host 'pluto' in the `user` table, but your DNS tells MySQL that your host name is 'pluto.example.com', the entry does not work. Try adding an entry to the `user` table that contains the IP address of your host as the `Host` column value. (Alternatively, you could add an

entry to the `user` table with a `Host` value that contains a wildcard; for example, `'pluto.%'`. However, use of `Host` values ending with `"%"` is *insecure* and is *not* recommended!)

- If `mysql -u user_name test` works but `mysql -u user_name other_db` does not, you have not granted access to the given user for the database named `other_db`.
- If `mysql -u user_name` works when executed on the server host, but `mysql -h host_name -u user_name` does not work when executed on a remote client host, you have not enabled access to the server for the given user name from the remote host.
- If you cannot figure out why you get `Access denied`, remove from the `user` table all entries that have `Host` values containing wildcards (entries that contain `'%'` or `'_'` characters). A very common error is to insert a new entry with `Host='%'` and `User='some_user'`, thinking that this enables you to specify `localhost` to connect from the same machine. The reason that this does not work is that the default privileges include an entry with `Host='localhost'` and `User=''`. Because that entry has a `Host` value `'localhost'` that is more specific than `'%'`, it is used in preference to the new entry when connecting from `localhost`! The correct procedure is to insert a second entry with `Host='localhost'` and `User='some_user'`, or to delete the entry with `Host='localhost'` and `User=''`. After deleting the entry, remember to issue a `FLUSH PRIVILEGES` statement to reload the grant tables. See also [Section 6.2.4, "Access Control, Stage 1: Connection Verification"](#).
- If you are able to connect to the MySQL server, but get an `Access denied` message whenever you issue a `SELECT ... INTO OUTFILE` or `LOAD DATA INFILE` statement, your entry in the `user` table does not have the `FILE [754]` privilege enabled.
- If you change the grant tables directly (for example, by using `INSERT`, `UPDATE`, or `DELETE` statements) and your changes seem to be ignored, remember that you must execute a `FLUSH PRIVILEGES` statement or a `mysqladmin flush-privileges` command to cause the server to reload the privilege tables. Otherwise, your changes have no effect until the next time the server is restarted. Remember that after you change the `root` password with an `UPDATE` statement, you will not need to specify the new password until after you flush the privileges, because the server will not know you've changed the password yet!
- If your privileges seem to have changed in the middle of a session, it may be that a MySQL administrator has changed them. Reloading the grant tables affects new client connections, but it also affects existing connections as indicated in [Section 6.2.6, "When Privilege Changes Take Effect"](#).
- If you have access problems with a Perl, PHP, Python, or ODBC program, try to connect to the server with `mysql -u user_name db_name` or `mysql -u user_name -p your_pass db_name`. If you are able to connect using the `mysql` client, the problem lies with your program, not with the access privileges. (There is no space between `-p` and the password; you can also use the `--password=your_pass [211]` syntax to specify the password. If you use the `-p` or `--password [211]` option with no password value, MySQL prompts you for the password.)
- For testing purposes, start the `mysqld` server with the `--skip-grant-tables [439]` option. Then you can change the MySQL grant tables and use the `mysqlaccess` script to check whether your modifications have the desired effect. When you are satisfied with your changes, execute `mysqladmin flush-privileges` to tell the `mysqld` server to reload the privileges. This enables you to begin using the new grant table contents without stopping and restarting the server.
- If everything else fails, start the `mysqld` server with a debugging option (for example, `--debug=d,general,query [418]`). This prints host and user information about attempted connections, as well as information about each command issued. See [MySQL Internals: Porting to Other Systems](#).
- If you have any other problems with the MySQL grant tables and feel you must post the problem to the mailing list, always provide a dump of the MySQL grant tables. You can dump the tables with the `mysqldump mysql` command. To file a bug report, see the instructions at [Section 1.7, "How to Report Bugs or Problems"](#). In some cases, you may need to restart `mysqld` with `--skip-grant-tables [439]` to run `mysqldump`.

## 6.3. MySQL User Account Management

This section describes how to set up accounts for clients of your MySQL server. It discusses the following topics:

- The meaning of account names and passwords as used in MySQL and how that compares to names and passwords used by your operating system
- How to set up new accounts and remove existing accounts
- How to change passwords
- Guidelines for using passwords securely
- How to use secure connections with SSL

See also [Section 13.7.1, “Account Management Statements”](#), which describes the syntax and use for all user-management SQL statements.

### 6.3.1. User Names and Passwords

MySQL stores accounts in the `user` table of the `mysql` database. An account is defined in terms of a user name and the client host or hosts from which the user can connect to the server. The account may also have a password. For information about account representation in the `user` table, see [Section 6.2.2, “Privilege System Grant Tables”](#). MySQL 5.6 supports authentication plugins, so it is possible that an account authenticates using some external authentication method. See [Section 6.3.7, “Pluggable Authentication”](#).

There are several distinctions between the way user names and passwords are used by MySQL and the way they are used by your operating system:

- User names, as used by MySQL for authentication purposes, have nothing to do with user names (login names) as used by Windows or Unix. On Unix, most MySQL clients by default try to log in using the current Unix user name as the MySQL user name, but that is for convenience only. The default can be overridden easily, because client programs permit any user name to be specified with a `-u` or `--user` option. Because this means that anyone can attempt to connect to the server using any user name, you cannot make a database secure in any way unless all MySQL accounts have passwords. Anyone who specifies a user name for an account that has no password is able to connect successfully to the server.
- MySQL user names can be up to 16 characters long. Operating system user names, because they are completely unrelated to MySQL user names, may be of a different maximum length. For example, Unix user names typically are limited to eight characters.



#### Warning

The limit on MySQL user name length is hard-coded in the MySQL servers and clients, and trying to circumvent it by modifying the definitions of the tables in the `mysql` database *does not work*.

You should never alter any of the tables in the `mysql` database in any manner whatsoever except by means of the procedure that is described in [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#). Attempting to redefine MySQL's system tables in any other fashion results in undefined (and unsupported!) behavior.

- The server uses MySQL passwords stored in the `user` table to authenticate client connections using MySQL native authentication (against passwords stored in the `mysql.user` table). These passwords have nothing to do with passwords for logging in to your operating system. There is no necessary connection between the “external” password you use to log in to a Windows or Unix machine and the password you use to access the MySQL server on that machine.

If the server authenticates a client using a plugin, the authentication method that the plugin implements may or may not use the password in the `user` table. In this case, it is possible that an external password is also used to authenticate to the MySQL server.

- MySQL encrypts passwords stored in the `user` table using its own algorithm. This encryption is the same as that implemented by the `PASSWORD()` [1263] SQL function but differs from that used during the Unix login process. Unix password encryption is the same as that implemented by the `ENCRYPT()` [1262] SQL function. See the descriptions of the `PASSWORD()` [1263] and `ENCRYPT()` [1262] functions in [Section 12.13, “Encryption and Compression Functions”](#).

From version 4.1 on, MySQL employs a stronger authentication method that has better password protection during the connection process than in earlier versions. It is secure even if TCP/IP packets are sniffed or the `mysql` database is captured. (In earlier versions, even though passwords are stored in encrypted form in the `user` table, knowledge of the encrypted password value could be used to connect to the MySQL server.) [Section 6.1.2.4, “Password Hashing in MySQL”](#), discusses password encryption further.

- It is possible to connect to the server regardless of character set settings if the user name and password contain only ASCII characters. To connect when the user name or password contain non-ASCII characters, the client should call the `mysql_options()` C API function with the `MYSQL_SET_CHARSET_NAME` option and appropriate character set name as arguments. This causes authentication to take place using the specified character set. Otherwise, authentication will fail unless the server default character set is the same as the encoding in the authentication defaults.

Standard MySQL client programs support a `--default-character-set` option that causes `mysql_options()` to be called as just described. In addition, character set autodetection is supported as described in [Section 10.1.4, “Connection Character Sets and Collations”](#). For programs that use a connector that is not based on the C API, the connector may provide an equivalent to `mysql_options()` that can be used instead. Check the connector documentation.

The preceding notes do not apply for `ucs2`, `utf16`, and `utf32`, which are not permitted as client character sets.

When you install MySQL, the grant tables are populated with an initial set of accounts. The names and access privileges for these accounts are described in [Section 2.10.2, “Securing the Initial MySQL Accounts”](#), which also discusses how to assign passwords to them. Thereafter, you normally set up, modify, and remove MySQL accounts using statements such as `CREATE USER`, `GRANT`, and `REVOKE`. See [Section 13.7.1, “Account Management Statements”](#).

When you connect to a MySQL server with a command-line client, specify the user name and password as necessary for the account that you want to use:

```
shell> mysql --user=monty --password=password db_name
```

If you prefer short options, the command looks like this:

```
shell> mysql -u monty -ppassword db_name
```

There must be *no space* between the `-p` option and the following password value.

If you omit the `password` value following the `--password` [211] or `-p` option on the command line, the client prompts for one.

Specifying a password on the command line should be considered insecure. See [Section 6.1.2.1, “End-User Guidelines for Password Security”](#). You can use an option file to avoid giving the password on the command line.

For additional information about specifying user names, passwords, and other connection parameters, see [Section 4.2.2, “Connecting to the MySQL Server”](#).

## 6.3.2. Adding User Accounts



You can create MySQL accounts in two ways:

- By using statements intended for creating accounts, such as `CREATE USER` or `GRANT`. These statements cause the server to make appropriate modifications to the grant tables.
- By manipulating the MySQL grant tables directly with statements such as `INSERT`, `UPDATE`, or `DELETE`.

The preferred method is to use account-creation statements because they are more concise and less error-prone than manipulating the grant tables directly. `CREATE USER` and `GRANT` are described in [Section 13.7.1, “Account Management Statements”](#).

Another option for creating accounts is to use the GUI tool MySQL Workbench. Or one of several available third-party programs that offer capabilities for MySQL account administration. `phpMyAdmin` is one such program.

The following examples show how to use the `mysql` client program to set up new accounts. These examples assume that privileges have been set up according to the defaults described in [Section 2.10.2, “Securing the Initial MySQL Accounts”](#). This means that to make changes, you must connect to the MySQL server as the MySQL `root` user, and the `root` account must have the `INSERT [755]` privilege for the `mysql` database and the `RELOAD [755]` administrative privilege.

As noted in the examples where appropriate, some of the statements will fail if the server's SQL mode has been set to enable certain restrictions. In particular, strict mode (`STRICT_TRANS_TABLES [607]`, `STRICT_ALL_TABLES [607]`) and `NO_AUTO_CREATE_USER [604]` will prevent the server from accepting some of the statements. Workarounds are indicated for these cases. For more information about SQL modes and their effect on grant table manipulation, see [Section 5.1.7, “Server SQL Modes”](#), and [Section 13.7.1.4, “GRANT Syntax”](#).

First, use the `mysql` program to connect to the server as the MySQL `root` user:

```
shell> mysql --user=root mysql
```

If you have assigned a password to the `root` account, you will also need to supply a `--password` or `-p` option, both for this `mysql` command and for those later in this section.

After connecting to the server as `root`, you can add new accounts. The following statements use `GRANT` to set up four new accounts:

```
mysql> CREATE USER 'monty'@'localhost' IDENTIFIED BY 'some_pass';
mysql> GRANT ALL PRIVILEGES ON *.* TO 'monty'@'localhost'
-> WITH GRANT OPTION;
mysql> CREATE USER 'monty'@'%' IDENTIFIED BY 'some_pass';
mysql> GRANT ALL PRIVILEGES ON *.* TO 'monty'@'%'
-> WITH GRANT OPTION;
mysql> CREATE USER 'admin'@'localhost';
mysql> GRANT RELOAD,PROCESS ON *.* TO 'admin'@'localhost';
mysql> CREATE USER 'dummy'@'localhost';
```

The accounts created by these statements have the following properties:

- Two of the accounts have a user name of `monty` and a password of `some_pass`. Both accounts are superuser accounts with full privileges to do anything. The `'monty'@'localhost'` account can be used only when connecting from the local host. The `'monty'@'%'` account uses the `'%'` wildcard for the host part, so it can be used to connect from any host.

It is necessary to have both accounts for `monty` to be able to connect from anywhere as `monty`. Without the `localhost` account, the anonymous-user account for `localhost` that is created by `mysql_install_db` would take precedence when `monty` connects from the local host. As a result, `monty` would be treated as an anonymous user. The reason for this is that the anonymous-user account has a more specific `Host` column value than the `'monty'@'%'` account and thus comes earlier in the `user` table sort order. (`user` table sorting is discussed in [Section 6.2.4, “Access Control, Stage 1: Connection Verification”](#).)

- The `'admin'@'localhost'` account has no password. This account can be used only by `admin` to connect from the local host. It is granted the `RELOAD [755]` and `PROCESS [755]` administrative privileges. These privileges enable the `admin` user to execute the `mysqladmin reload`, `mysqladmin refresh`, and `mysqladmin flush-xxx` commands, as well as `mysqladmin processlist`. No privileges are granted for accessing any databases. You could add such privileges later by issuing other `GRANT` statements.
- The `'dummy'@'localhost'` account has no password. This account can be used only to connect from the local host. No privileges are granted. It is assumed that you will grant specific privileges to the account later.

The statements that create accounts with no password will fail if the `NO_AUTO_CREATE_USER [604]` SQL mode is enabled. To deal with this, use an `IDENTIFIED BY` clause that specifies a nonempty password.

To check the privileges for an account, use `SHOW GRANTS`:

```
mysql> SHOW GRANTS FOR 'admin'@'localhost';
+-----+
| Grants for admin@localhost |
+-----+
| GRANT RELOAD, PROCESS ON *.* TO 'admin'@'localhost' |
+-----+
```

As an alternative to `CREATE USER` and `GRANT`, you can create the same accounts directly by issuing `INSERT` statements and then telling the server to reload the grant tables using `FLUSH PRIVILEGES`:

```
shell> mysql --user=root mysql
mysql> INSERT INTO user
-> VALUES('localhost','monty',PASSWORD('some_pass'),
-> 'Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO user
-> VALUES('%','monty',PASSWORD('some_pass'),
-> 'Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y',
-> 'Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y',
-> '','','',0,0,0,0);
mysql> INSERT INTO user SET Host='localhost',User='admin',
-> Reload_priv='Y', Process_priv='Y';
mysql> INSERT INTO user (Host,User>Password)
-> VALUES('localhost','dummy','');
mysql> FLUSH PRIVILEGES;
```

When you create accounts with `INSERT`, it is necessary to use `FLUSH PRIVILEGES` to tell the server to reload the grant tables. Otherwise, the changes go unnoticed until you restart the server. With `CREATE USER`, `FLUSH PRIVILEGES` is unnecessary.

The reason for using the `PASSWORD( ) [1263]` function with `INSERT` is to encrypt the password. The `CREATE USER` statement encrypts the password for you, so `PASSWORD( ) [1263]` is unnecessary.

The `'Y'` values enable privileges for the accounts. Depending on your MySQL version, you may have to use a different number of `'Y'` values in the first two `INSERT` statements. The `INSERT` statement for the `admin` account employs the more readable extended `INSERT` syntax using `SET`.

In the `INSERT` statement for the `dummy` account, only the `Host`, `User`, and `Password` columns in the `user` table row are assigned values. None of the privilege columns are set explicitly, so MySQL assigns them all the default value of `'N'`. This is equivalent to what `CREATE USER` does.

If strict SQL mode is enabled, all columns that have no default value must have a value specified. In this case, `INSERT` statements must explicitly specify values for the `ssl_cipher`, `x509_issuer`, and `x509_subject` columns.

To set up a superuser account, it is necessary only to insert a `user` table row with all privilege columns set to `'Y'`. The `user` table privileges are global, so no entries in any of the other grant tables are needed.

The next examples create three accounts and give them access to specific databases. Each of them has a user name of `custom` and password of `obscure`.

To create the accounts with `CREATE USER` and `GRANT`, use the following statements:

```
shell> mysql --user=root mysql
mysql> CREATE USER 'custom'@'localhost' IDENTIFIED BY 'obscure';
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON bankaccount.*
-> TO 'custom'@'localhost';
mysql> CREATE USER 'custom'@'host47.example.com' IDENTIFIED BY 'obscure';
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON expenses.*
-> TO 'custom'@'host47.example.com';
mysql> CREATE USER 'custom'@'server.domain' IDENTIFIED BY 'obscure';
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON customer.*
-> TO 'custom'@'server.domain';
```

The three accounts can be used as follows:

- The first account can access the `bankaccount` database, but only from the local host.
- The second account can access the `expenses` database, but only from the host `host47.example.com`.
- The third account can access the `customer` database, but only from the host `server.domain`.

To set up the `custom` accounts without `GRANT`, use `INSERT` statements as follows to modify the grant tables directly:

```
shell> mysql --user=root mysql
mysql> INSERT INTO user (Host,User,Password)
-> VALUES('localhost','custom',PASSWORD('obscure'));
mysql> INSERT INTO user (Host,User,Password)
-> VALUES('host47.example.com','custom',PASSWORD('obscure'));
mysql> INSERT INTO user (Host,User,Password)
-> VALUES('server.domain','custom',PASSWORD('obscure'));
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES('localhost','bankaccount','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES('host47.example.com','expenses','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES('server.domain','customer','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> FLUSH PRIVILEGES;
```

The first three `INSERT` statements add `user` table entries that permit the user `custom` to connect from the various hosts with the given password, but grant no global privileges (all privileges are set to the default value of `'N'`). The next three `INSERT` statements add `db` table entries that grant privileges to `custom` for the `bankaccount`, `expenses`, and `customer` databases, but only when accessed from the proper hosts. As usual when you modify the grant tables directly, you must tell the server to reload them with `FLUSH PRIVILEGES` so that the privilege changes take effect.

To create a user who has access from all machines in a given domain (for example, `mydomain.com`), you can use the `"%"` wildcard character in the host part of the account name:

```
mysql> CREATE USER 'myname'@'%.mydomain.com' IDENTIFIED BY 'mypass';
```

To do the same thing by modifying the grant tables directly, do this:



```
mysql> INSERT INTO user (Host,User,Password,...)
-> VALUES ('%.mydomain.com','myname',PASSWORD('mypass'),...);
mysql> FLUSH PRIVILEGES;
```

### 6.3.3. Removing User Accounts

To remove an account, use the `DROP USER` statement, which is described in [Section 13.7.1.3](#), “`DROP USER Syntax`”.

### 6.3.4. Setting Account Resource Limits

One means of limiting use of MySQL server resources is to set the global `max_user_connections` [509] system variable to a nonzero value. This limits the number of simultaneous connections that can be made by any given account, but places no limits on what a client can do once connected. In addition, setting `max_user_connections` [509] does not enable management of individual accounts. Both types of control are of interest to many MySQL administrators, particularly those working for Internet Service Providers.

In MySQL 5.6, you can limit use of the following server resources for individual accounts:

- The number of queries that an account can issue per hour
- The number of updates that an account can issue per hour
- The number of times an account can connect to the server per hour
- The number of simultaneous connections to the server by an account

Any statement that a client can issue counts against the query limit (unless its results are served from the query cache). Only statements that modify databases or tables count against the update limit.

An “account” in this context corresponds to a row in the `mysql.user` table. That is, a connection is assessed against the `User` and `Host` values in the `user` table row that applies to the connection. For example, an account `'usera'@'%.example.com'` corresponds to a row in the `user` table that has `User` and `Host` values of `usera` and `%.example.com`, to permit `usera` to connect from any host in the `example.com` domain. In this case, the server applies resource limits in this row collectively to all connections by `usera` from any host in the `example.com` domain because all such connections use the same account.

Before MySQL 5.0.3, an “account” was assessed against the actual host from which a user connects. This older method accounting may be selected by starting the server with the `--old-style-user-limits` [433] option. In this case, if `usera` connects simultaneously from `host1.example.com` and `host2.example.com`, the server applies the account resource limits separately to each connection. If `usera` connects again from `host1.example.com`, the server applies the limits for that connection together with the existing connection from that host.

To set resource limits for an account, use the `GRANT` statement (see [Section 13.7.1.4](#), “`GRANT Syntax`”). Provide a `WITH` clause that names each resource to be limited. The default value for each limit is zero (no limit). For example, to create a new account that can access the `customer` database, but only in a limited fashion, issue these statements:

```
mysql> CREATE USER 'francis'@'localhost' IDENTIFIED BY 'frank';
mysql> GRANT ALL ON customer.* TO 'francis'@'localhost'
-> WITH MAX_QUERIES_PER_HOUR 20
-> MAX_UPDATES_PER_HOUR 10
-> MAX_CONNECTIONS_PER_HOUR 5
-> MAX_USER_CONNECTIONS 2;
```

The limit types need not all be named in the `WITH` clause, but those named can be present in any order. The value for each per-hour limit should be an integer representing a count per hour. For `MAX_USER_CONNECTIONS`, the limit is an integer representing the maximum number of simultaneous connections by the account. If this limit is set to zero, the global

`max_user_connections` [509] system variable value determines the number of simultaneous connections. If `max_user_connections` [509] is also zero, there is no limit for the account.

To modify existing limits for an account, use a `GRANT USAGE` statement at the global level (`ON *.*`). The following statement changes the query limit for `francis` to 100:

```
mysql> GRANT USAGE ON *.* TO 'francis'@'localhost'
->     WITH MAX_QUERIES_PER_HOUR 100;
```

The statement modifies only the limit value specified and leaves the account otherwise unchanged.

To remove a limit, set its value to zero. For example, to remove the limit on how many times per hour `francis` can connect, use this statement:

```
mysql> GRANT USAGE ON *.* TO 'francis'@'localhost'
->     WITH MAX_CONNECTIONS_PER_HOUR 0;
```

As mentioned previously, the simultaneous-connection limit for an account is determined from the `MAX_USER_CONNECTIONS` limit and the `max_user_connections` [509] system variable. Suppose that the global `max_user_connections` [509] value is 10 and three accounts have resource limits specified with `GRANT`:

```
GRANT ... TO 'user1'@'localhost' WITH MAX_USER_CONNECTIONS 0;
GRANT ... TO 'user2'@'localhost' WITH MAX_USER_CONNECTIONS 5;
GRANT ... TO 'user3'@'localhost' WITH MAX_USER_CONNECTIONS 20;
```

`user1` has a connection limit of 10 (the global `max_user_connections` [509] value) because it has a zero `MAX_USER_CONNECTIONS` limit). `user2` and `user3` have connection limits of 5 and 20, respectively, because they have nonzero `MAX_USER_CONNECTIONS` limits.

The server stores resource limits for an account in the `user` table row corresponding to the account. The `max_questions`, `max_updates`, and `max_connections` columns store the per-hour limits, and the `max_user_connections` column stores the `MAX_USER_CONNECTIONS` limit. (See [Section 6.2.2, "Privilege System Grant Tables"](#).)

Resource-use counting takes place when any account has a nonzero limit placed on its use of any of the resources.

As the server runs, it counts the number of times each account uses resources. If an account reaches its limit on number of connections within the last hour, further connections for the account are rejected until that hour is up. Similarly, if the account reaches its limit on the number of queries or updates, further queries or updates are rejected until the hour is up. In all such cases, an appropriate error message is issued.

Resource counting is done per account, not per client. For example, if your account has a query limit of 50, you cannot increase your limit to 100 by making two simultaneous client connections to the server. Queries issued on both connections are counted together.

The current per-hour resource-use counts can be reset globally for all accounts, or individually for a given account:

- To reset the current counts to zero for all accounts, issue a `FLUSH USER_RESOURCES` statement. The counts also can be reset by reloading the grant tables (for example, with a `FLUSH PRIVILEGES` statement or a `mysqladmin reload` command).
- The counts for an individual account can be set to zero by re-granting it any of its limits. To do this, use `GRANT USAGE` as described earlier and specify a limit value equal to the value that the account currently has.

Counter resets do not affect the `MAX_USER_CONNECTIONS` limit.

All counts begin at zero when the server starts; counts are not carried over through a restart.

For the `MAX_USER_CONNECTIONS` limit, an edge case can occur if the account currently has open the maximum number of connections permitted to it: A disconnect followed quickly by a connect can result in an error (`ER_TOO_MANY_USER_CONNECTIONS` [3673] or `ER_USER_LIMIT_REACHED` [3675]) if the server has not fully processed the disconnect by the time the connect occurs. When the server finishes disconnect processing, another connection will once more be permitted.

### 6.3.5. Assigning Account Passwords

Required credentials for clients that connect to the MySQL server can include a password. This section describes how to assign passwords for MySQL accounts. In MySQL 5.6, it is also possible for clients to authenticate using plugins. For information, see [Section 6.3.7, “Pluggable Authentication”](#).

To assign a password when you create a new account with `CREATE USER`, include an `IDENTIFIED BY` clause:

```
mysql> CREATE USER 'jeffrey'@'localhost'
-> IDENTIFIED BY 'mypass';
```

To assign or change a password for an existing account, one way is to issue a `SET PASSWORD` statement:

```
mysql> SET PASSWORD FOR
-> 'jeffrey'@'localhost' = PASSWORD('mypass');
```

MySQL stores passwords in the `user` table in the `mysql` database. Only users such as `root` that have update access to the `mysql` database can change the password for other users. If you are not connected as an anonymous user, you can change your own password by omitting the `FOR` clause:

```
mysql> SET PASSWORD = PASSWORD('mypass');
```

The `old_passwords` [517] system variable value determines the hashing method used by `PASSWORD()` [1263]. If you specify the password using that function and `SET PASSWORD` rejects the password as not being in the correct format, it may be necessary to set `old_passwords` [517] to change the hashing method. For descriptions of the permitted values, see [Section 5.1.4, “Server System Variables”](#).

In MySQL 5.6, enabling the `read_only` [532] system variable prevents the use of the `SET PASSWORD` statement by any user not having the `SUPER` [756] privilege.

You can also use a `GRANT USAGE` statement at the global level (`ON *.*`) to assign a password to an account without affecting the account's current privileges:

```
mysql> GRANT USAGE ON *.* TO 'jeffrey'@'localhost'
-> IDENTIFIED BY 'mypass';
```

To assign a password from the command line, use the `mysqladmin` command:

```
shell> mysqladmin -u user_name -h host_name password "newpwd"
```

The account for which this command sets the password is the one with a `user` table row that matches `user_name` in the `User` column and the client host *from which you connect* in the `Host` column.

During authentication when a client connects to the server, MySQL treats the password in the `user` table as an encrypted hash value (the value that `PASSWORD()` [1263] would return for the password). When assigning a password to an account, it is important to store an encrypted value, not the plaintext password. Use the following guidelines:

- When you assign a password using `CREATE USER`, `GRANT` with an `IDENTIFIED BY` clause, or the `mysqladmin password` command, they encrypt the password for you. Specify the literal plaintext password:

```
mysql> CREATE USER 'jeffrey'@'localhost'
```

```
-> IDENTIFIED BY 'mypass';
```

- For `CREATE USER` or `GRANT`, you can avoid sending the plaintext password if you know the hash value that `PASSWORD()` [1263] would return for the password. Specify the hash value preceded by the keyword `PASSWORD`:

```
mysql> CREATE USER 'jeffrey'@'localhost'
-> IDENTIFIED BY PASSWORD '*90E462C37378CED12064BB3388827D2BA3A9B689';
```

- When you assign an account a nonempty password using `SET PASSWORD`, you must use the `PASSWORD()` [1263] function to encrypt the password, otherwise the password is stored as plaintext. Suppose that you assign a password like this:

```
mysql> SET PASSWORD FOR
-> 'jeffrey'@'localhost' = 'mypass';
```

The result is that the literal value `'mypass'` is stored as the password in the `user` table, not the encrypted value. When `jeffrey` attempts to connect to the server using this password, the value is encrypted and compared to the value stored in the `user` table. However, the stored value is the literal string `'mypass'`, so the comparison fails and the server rejects the connection with an `Access denied` error.



#### Note

`PASSWORD()` [1263] encryption differs from Unix password encryption. See [Section 6.3.1, “User Names and Passwords”](#).

It is preferable to assign passwords using `SET PASSWORD`, `GRANT`, or `mysqladmin`, but it is also possible to modify the `user` table directly. In this case, you must also use `FLUSH PRIVILEGES` to cause the server to reread the grant tables. Otherwise, the change remains unnoticed by the server until you restart it.

- To establish a password for a new account, provide a value for the `Password` column:

```
mysql> INSERT INTO mysql.user (Host,User,Password)
-> VALUES('localhost','jeffrey',PASSWORD('mypass'));
mysql> FLUSH PRIVILEGES;
```

- To change the password for an existing account, use `UPDATE` to set the `Password` column value:

```
mysql> UPDATE mysql.user SET Password = PASSWORD('bagel')
-> WHERE Host = 'localhost' AND User = 'francis';
mysql> FLUSH PRIVILEGES;
```

## 6.3.6. Password Expiration and Sandbox Mode

MySQL 5.6 introduces password-expiration capability, to enable database administrators to expire account passwords and require users to reset their password. The immediately following discussion describes how password expiration works currently. Later, the development of this capability is detailed as it occurred over several versions, as background to help you understand what features are available when. However, to ensure that you can take advantage of all features and fixes, you should use the most recent available version of MySQL 5.6 if possible.

### How Password Expiration Works

To expire an account password, use the `ALTER USER` statement. For example:

```
ALTER USER 'myuser'@'localhost' PASSWORD EXPIRE;
```

This statement modifies the row of the `mysql.user` table associated with the named account, setting the `password_expired` column to `'Y'`. This does not affect any current connections the account has open. For each subsequent connection that uses the account, the server either disconnects the client or handles the client in “sandbox mode,” in which the server permits the client only those operations

necessary to reset the expired password. (The action taken by the server depends on both client and server settings.)

If the server disconnects the client, it returns an `ER_MUST_CHANGE_PASSWORD_LOGIN` [3717] error:

```
shell> mysql -u myuser -p
Password: *****
ERROR 1862 (HY000): Your password has expired. To log in you must
change it using a client that supports expired passwords.
```

If the server puts the client in sandbox mode, these operations are permitted within the client session:

- The client can reset the account password with `SET PASSWORD`. This modifies the row of the `mysql.user` table associated with the current account, setting the `password_expired` column to 'N'. After the password has been reset, the server restores normal access for the session, as well as for subsequent connections that use the account.

It is possible to “reset” a password by setting it to its current value. As a matter of good policy, it is preferable to choose a different password.

- The client can use `SET` statements. This might be necessary prior to resetting the password; for example, if the account password uses a hashing format that requires the `old_passwords` [517] system variable to be set to a value different from its default.

For any operation not permitted within the session, the server returns an `ER_MUST_CHANGE_PASSWORD` [3714] error:

```
mysql> USE test;
ERROR 1820 (HY000): You must SET PASSWORD before executing this statement
```

As mentioned previously, whether the server disconnects an expired-password client or puts it in sandbox mode depends on a combination of client and server settings. The following discussion describes the relevant settings and how they interact.

On the client side, a given client indicates whether it can handle sandbox mode for expired passwords. For clients that use the C client library, there are two ways to do this:

- Pass the `MYSQL_OPT_CAN_HANDLE_EXPIRED_PASSWORDS` flag to `mysql_options()` prior to connecting:

```
arg = 1;
result = mysql_options(mysql,
                        MYSQL_OPT_CAN_HANDLE_EXPIRED_PASSWORDS, &arg);
```

- Pass the `CLIENT_CAN_HANDLE_EXPIRED_PASSWORDS` flag to `mysql_real_connect()` at connection time:

```
mysql = mysql_real_connect(mysql,
                           host, user, password, "test",
                           port, unix_socket,
                           CLIENT_CAN_HANDLE_EXPIRED_PASSWORDS);
```

Other MySQL Connectors have their own conventions for indicating readiness to handle sandbox mode. See the relevant Connector documentation.

On the server side, if a client indicates that it can handle expired passwords, the server puts it in sandbox mode.

If a client does not indicate that it can handle expired passwords (or uses an older version of the client library that cannot so indicate), the server action depends on the value of the `disconnect_on_expired_passwords` [477] system variable:

- If `disconnect_on_expired_passwords` [477] is enabled (the default), the server disconnects the client with an `ER_MUST_CHANGE_PASSWORD_LOGIN` [3717] error.

- If `disconnect_on_expired_passwords` [477] is disabled, the server puts the client in sandbox mode.

The preceding client and server settings apply only for accounts with expired passwords. If a client connects using a nonexpired password, the server handles the client normally.

## Development of Password-Expiration Capability

The following timeline describes the versions in which various password-expiration features were added.

MySQL 5.6.6: Initial implementation of password expiration

The `password_expired` column is introduced in the `mysql.user` table to enable DBAs to expire account passwords. The column default value is 'N' (not expired).

The `ALTER USER` statement is introduced as the SQL interface for setting the `password_expired` column to 'Y'.

Connections that use an account with an expired password enter “sandbox mode” that permits only `SET PASSWORD` statements. For other statements, the server returns an `ER_MUST_CHANGE_PASSWORD` [3714] error. The intent is to force the client to reset the password before the server permits any other operations. `SET PASSWORD` resets the account password and sets `password_expired` to 'N'.

A bug in the initial implementation is that `ALTER USER` sets the `Password` column in the `mysql.user` table to the empty string. The implication is that users should wait until MySQL 5.6.7 to use this statement.

MySQL 5.6.7:

`ALTER USER` is fixed to not set the `Password` column to the empty string.

MySQL 5.6.8:

`ALTER USER` can be used as a prepared statement.

`mysqladmin password` is made capable of setting passwords for accounts with expired native or old-native passwords.

Sandbox mode is changed to permit clients to execute `SET` statements in addition to `SET PASSWORD`. Prohibiting `SET` prevented clients that needed to set `old_passwords` from resetting their password. It also broke some Connectors, which use `SET` extensively at connect time to initialize the session environment.

MySQL 5.6.9:

Sandbox mode is changed to permit `SET PASSWORD` only if the account named in the statement matches the account the client authenticated as.

MySQL 5.6.10:

Sandbox mode is changed to permit better control over how the server handles client connections for accounts with expired passwords, and to permit clients to signal whether they are capable of handling expired passwords:

- The `disconnect_on_expired_passwords` [477] system variable is added, which controls how the server treats expired-password accounts.
- Two flags are added to the C API client library: `MYSQL_OPT_CAN_HANDLE_EXPIRED_PASSWORDS` for `mysql_options()` and `CLIENT_CAN_HANDLE_EXPIRED_PASSWORDS` for



`mysql_real_connect()`. Each flag enables a client program to indicate whether it can handle sandbox mode for accounts with expired passwords.

`MYSQL_OPT_CAN_HANDLE_EXPIRED_PASSWORDS` is enabled for `mysqltest` unconditionally, for `mysql` in interactive mode, and for `mysqladmin` if the first command is `password`.

- The `ER_MUST_CHANGE_PASSWORD_LOGIN` [3717] error is added. The server returns this error when it disconnects a client.

Concurrent with these changes to sandbox mode in MySQL Server and the C API client library, work begins to modify Connectors for conformance to the changes.

## 6.3.7. Pluggable Authentication

When a client connects to the MySQL server, the server uses the user name provided by the client and the client host to select the appropriate account row from the `mysql.user` table. It then uses this row to authenticate the client.

In MySQL 5.6, the server authenticates clients using plugins, as follows:

- The server determines from the account row which authentication plugin applies for the client:
  - If the account row specifies no plugin name, the server uses native authentication; that is, authentication against the password stored in the `Password` column of the account row. This is the same authentication method provided by MySQL servers older than 5.5.7, before pluggable authentication was implemented, but now is implemented using two plugins that are built in and cannot be disabled.
  - If the account row specifies a plugin, the server invokes it to authenticate the user. If the server cannot find the plugin, an error occurs.
- The plugin returns a status to the server indicating whether the user is permitted to connect.

Pluggable authentication enables two important capabilities:

- **External authentication:** Pluggable authentication makes it possible for clients to connect to the MySQL server with credentials that are appropriate for authentication methods other than native authentication based on passwords stored in the `mysql.user` table. For example, plugins can be created to use external authentication methods such as PAM, Windows login IDs, LDAP, or Kerberos.
- **Proxy users:** If a user is permitted to connect, an authentication plugin can return to the server a user name different from the name of the connecting user, to indicate that the connecting user is a proxy for another user. While the connection lasts, the proxy user is treated, for purposes of access control, as having the privileges of a different user. In effect, one user impersonates another. For more information, see [Section 6.3.8, “Proxy Users”](#).

Several authentication plugins are available in MySQL. The following sections provide details about specific plugins.

- Plugins that perform native authentication that matches the password against the `Password` column of the account row. See [Section 6.3.7.1, “The Native Authentication Plugins”](#). Native authentication is the default for accounts that have no plugin named explicitly in their account row.
- A plugin that performs authentication using SHA-256 password hashing. This plugin matches the password against the `authentication_string` column of the account row. This is stronger encryption than that available with native authentication. See [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).
- A plugin that performs external authentication against PAM (Pluggable Authentication Modules), enabling MySQL Server to use PAM to authenticate MySQL users. This plugin supports proxy users as well. See [Section 6.3.7.3, “The PAM Authentication Plugin”](#).

- A plugin that performs external authentication on Windows, enabling MySQL Server to use native Windows services to authenticate client connections. Users who have logged in to Windows can connect from MySQL client programs to the server based on the information in their environment without specifying an additional password. This plugin supports proxy users as well. See [Section 6.3.7.4, “The Windows Native Authentication Plugin”](#).
- A client-side plugin that sends the password to the server without hashing or encryption. This plugin can be used by server-side plugins that require access to the password exactly as provided by the client user. See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).
- A plugin that authenticates clients that connect from the local host through the Unix socket file. See [Section 6.3.7.6, “The Socket Peer-Credential Authentication Plugin”](#).
- A test plugin that authenticates using MySQL native authentication. This plugin is intended for testing and development purposes, and as an example of how to write an authentication plugin. See [Section 6.3.7.7, “The Test Authentication Plugin”](#).

**Note**

For information about current restrictions on the use of pluggable authentication, including which connectors support which plugins, see [Section E.9, “Restrictions on Pluggable Authentication”](#).

Third-party connector developers should read that section to determine the extent to which a connector can take advantage of pluggable authentication capabilities and what steps to take to become more compliant.

If you are interested in writing your own authentication plugins, see [Section 22.2.4.9, “Writing Authentication Plugins”](#).

In general, pluggable authentication uses corresponding plugins on the server and client sides, so you use a given authentication method like this:

- On the server host, install the appropriate library containing the server plugin, if necessary, so that the server can use it to authenticate client connections. Similarly, on each client host, install the appropriate library containing the client plugin for use by client programs.
- Create MySQL accounts that specify use of the plugin for authentication.
- When a client connects, the server plugin tells the client program which client plugin to use for authentication.

The remainder of this section provides general instructions for installing and using authentication plugins. The instructions use an example authentication plugin included in MySQL distributions (see [Section 6.3.7.7, “The Test Authentication Plugin”](#)). The procedure is similar for other authentication plugins; substitute the appropriate plugin and file names.

The example authentication plugin has these characteristics:

- The server-side plugin name is `test_plugin_server`.
- The client-side plugin name is `auth_test_plugin`.
- Both plugins are located in the shared library object file named `auth_test_plugin.so` in the plugin directory (the directory named by the `plugin_dir` [525] system variable). The file name suffix might differ on your system.

Install and use the example authentication plugin as follows:

1. Make sure that the plugin library is installed on the server and client hosts.
2. Install the server-side test plugin at server startup or at runtime:



- To install the plugin at startup, use the `--plugin-load` [435] option. With this plugin-loading method, the option must be given each time you start the server. For example, use these lines in a `my.cnf` option file:

```
[mysqld]
plugin-load=test_plugin_server=auth_test_plugin.so
```

- To install the plugin at runtime, use the `INSTALL PLUGIN` statement:

```
mysql> INSTALL PLUGIN test_plugin_server SONAME 'auth_test_plugin.so';
```

This installs the plugin permanently and need be done only once.

PAM authentication, when not done through proxy users or groups, requires the MySQL account to have the same user name as the Unix account. Because MySQL user names are limited to 16 characters (see [Section 6.2.2, “Privilege System Grant Tables”](#)), this limits PAM nonproxy authentication to Unix accounts with names of at most 16 characters.

3. Verify that the plugin is installed. For example, use `SHOW PLUGINS`:

```
mysql> SHOW PLUGINS\G
...
***** 21. row *****
Name: test_plugin_server
Status: ACTIVE
Type: AUTHENTICATION
Library: auth_test_plugin.so
License: GPL
```

For other ways to check the plugin, see [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).

4. To specify that a MySQL user must be authenticated using the plugin, name it in the `IDENTIFIED WITH` clause of the `CREATE USER` statement that creates the user:

```
CREATE USER 'testuser'@'localhost' IDENTIFIED WITH test_plugin_server;
```

5. Connect to the server using a client program. The test plugin authenticates the same way as native MySQL authentication, so provide the usual `--user` [212] and `--password` [211] options that you normally use to connect to the server. For example:

```
shell> mysql --user=your_name --password=your_pass
```

For connections by `testuser`, the server sees that the account must be authenticated using the server-side plugin named `test_plugin_server` and communicates to the client program which client-side plugin it must use—in this case, `auth_test_plugin`.

In the case that the account uses the authentication method that is the default for both the server and the client program, the server need not communicate to the client which plugin to use, and a round trip in client/server negotiation can be avoided. Currently this is true for accounts that use native MySQL authentication (`mysql_native_password`).

The `--default-auth=plugin_name` [249] option can be specified on the `mysql` command line to make explicit which client-side plugin the program can expect to use, although the server will override this if the user account requires a different plugin.

If `mysql` does not find the plugin, specify a `--plugin-dir=dir_name` [252] option to indicate where the plugin is located.



#### Note

If you start the server with the `--skip-grant-tables` [439] option, authentication plugins are not used even if loaded because the server performs no client authentication and permits any client to connect. Because this is

insecure, you might want to use `--skip-grant-tables` [439] in conjunction with `--skip-networking` [440] to prevent remote clients from connecting.

### 6.3.7.1. The Native Authentication Plugins

MySQL includes two plugins that implement the same kind of native authentication that older servers provide; that is, authentication against passwords stored in the `Password` column of the `mysql.user` table:

- The `mysql_native_password` authentication plugin implements the same default authentication against the `mysql.user` table as used prior to the implementation of pluggable authentication.
- The `mysql_old_password` plugin implements authentication as used before MySQL 4.1.1 that is based on shorter password hash values. For information about this authentication method, see [Section 6.1.2.4, “Password Hashing in MySQL”](#).

The native authentication plugins are backward compatible. Clients older than MySQL 5.5.7 do not support authentication plugins but use native authentication, so they can connect to servers from 5.5.7 and up.

The following tables show the plugin names. Both are considered to implement native authentication even though only one has “native” in the name.

**Table 6.8. MySQL Native Password Authentication Plugin**

Server-side plugin name	<code>mysql_native_password</code>
Client-side plugin name	<code>mysql_native_password</code>
Library object file name	None (plugins are built in)

**Table 6.9. MySQL Native Old-Password Authentication Plugin**

Server-side plugin name	<code>mysql_old_password</code>
Client-side plugin name	<code>mysql_old_password</code>
Library object file name	None (plugins are built in)

Each plugin exists in both client and server form. MySQL client programs use `mysql_native_password` by default. The `--default-auth` [249] option can be used to specify either plugin explicitly:

```
shell> mysql --default-auth=mysql_native_password ...
shell> mysql --default-auth=mysql_old_password ...
```

The server-side plugins are built into the server, need not be loaded explicitly, and cannot be disabled by unloading them. The client-side plugins are built into the `libmysqlclient` client library as of MySQL 5.5.7 and available to any program linked against `libmysqlclient` from that version or newer.

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#).

### 6.3.7.2. The SHA-256 Authentication Plugin

As of MySQL 5.6.6, MySQL provides an authentication plugin that implements SHA-256 hashing for user account passwords. The following table shows the plugin names on the server and client sides.

**Table 6.10. MySQL SHA-256 Authentication Plugin**

Server-side plugin name	<code>sha256_password</code>
Client-side plugin name	<code>sha256_password</code>
Library object file name	None (plugins are built in)

The server-side `sha256_password` plugin is built into the server, need not be loaded explicitly, and cannot be disabled by unloading it. Similarly, clients need not specify the location of the client-side plugin.



### Note

Use of the `sha256_password` plugin requires that MySQL be built with SSL capabilities. See [Section 6.3.9, “Using SSL for Secure Connections”](#).

To set up an account that uses SHA-256 password hashing, use the following procedure.

1. Create the account and specify that it authenticates using the `sha256_password` plugin:

```
CREATE USER 'sha256user'@'localhost' IDENTIFIED WITH sha256_password;
```

2. Set the `old_passwords` system variable to 2 to select SHA-256 hashing of password strings by the `PASSWORD()` function;

```
SET old_passwords = 2;
```

3. Set the account password:

```
SET PASSWORD FOR 'sha256user'@'localhost' = PASSWORD('sha256P@ss');
```

Alternatively, start the server with the default authentication plugin set to `sha256_password`. For example, put these lines in the server option file:

```
[mysqld]
default-authentication-plugin=sha256_password
```

That causes the `sha256_password` plugin to be used by default for new accounts and sets `old_passwords` to 2. As a result, it is possible to set the password at account-creation time with the `CREATE USER` statement:

```
mysql> CREATE USER 'sha256user2'@'localhost' IDENTIFIED BY 'sha256P@ss2';
Query OK, 0 rows affected (0.06 sec)
```

In this case, the server assigns the `sha256_password` plugin to the account and encrypts the password using SHA-256. (Another consequence is that to create an account that uses a different authentication plugin, you must specify that plugin in the `CREATE USER` statement, then set `old_passwords` appropriately for the plugin before using `SET PASSWORD` to set the account password.)

If `old_passwords` has a value other than 2, an error occurs for attempts to set the password for an account that requires a SHA-256 password:

```
mysql> SET old_passwords = 0;
mysql> SET PASSWORD FOR 'sha256user'@'localhost' = PASSWORD('sha256P@ss');
ERROR 1827 (HY000): The password hash doesn't have the expected format.
Check if the correct password algorithm is being used with the
PASSWORD() function.
```

For more information about `old_passwords` [\[517\]](#) and `PASSWORD()` [\[1263\]](#), see [Section 5.1.4, “Server System Variables”](#), and [Section 12.13, “Encryption and Compression Functions”](#).

Accounts in the `mysql.user` table that use SHA-256 passwords can be identified as rows with `'sha256_password'` in the `plugin` column and a SHA-256 password hash in the `authentication_string` column.

MySQL can be built with either yaSSL or OpenSSL and the `sha256_password` plugin works with distributions built using either package. The default is to use yaSSL. If MySQL is built using OpenSSL instead, RSA encryption is available and `sha256_password` implements the additional capabilities in the following list. (To enable these capabilities, you must also follow the RSA configuration procedure given later in this section.)

- It is possible for the client to transmit passwords to the server using RSA encryption during the client connection process, as described later.
- The server exposes two additional system variables, `sha256_password_private_key_path` [539] and `sha256_password_public_key_path` [539]. It is intended that the database administrator will set these to the names of the RSA private and public key files at server startup.
- The server exposes a status variable, `Rsa_public_key` [597], that displays the RSA public key value.
- The `mysql` and `mysqltest` client programs support a `--server-public-key-path` [253] option for specifying an RSA public key file explicitly. (This option was added in MySQL 5.6.6 under the name `--server-public-key` and renamed in 5.6.7 to `--server-public-key-path`.)

For clients that use the `sha256_password` plugin, passwords are never exposed as cleartext when connecting to the server. How password transmission occurs depends on whether an SSL connection is used and whether RSA encryption is available:

- If an SSL connection is used, the password is sent as cleartext but cannot be snooped because the connection is encrypted using SSL.
- If an SSL connection is not used but RSA encryption is available, the password is sent within an unencrypted connection, but the password is RSA-encrypted to prevent snooping. When the server receives the password, it decrypts it. A scramble is used in the encryption to prevent repeat attacks.
- If an SSL connection is not used and RSA encryption is not available, the `sha256_password` plugin causes the connection attempt to fail because the password cannot be sent without being exposed as cleartext.

As mentioned previously, RSA password encryption is available only if MySQL was built using OpenSSL. The implication for MySQL distributions built using yaSSL is that SHA-256 passwords can be used only when clients access the server using an SSL connection. For information about connecting to the server using SSL, see [Section 6.3.9, “Using SSL for Secure Connections”](#).

Assuming that MySQL has been built with OpenSSL, the following procedure describes how to enable RSA encryption of passwords during the client connection process:

1. Create the RSA private and public key files. Run these commands while logged into the system account used to run the MySQL server so the files will be owned by that account:

```
openssl genrsa -out mykey.pem 1024
openssl rsa -in mykey.pem -pubout > mykey.pub
```

2. Set the access modes for the key files. The private key should be readable only by the server:

```
chmod 400 mykey.pem
```

The public key can be freely distributed to client users:

```
chmod 444 mykey.pub
```

3. In the server option file, configure the appropriate system variables with the names of the key files. If you place the files in the server data directory, you need not specify their full path names:

```
[mysqld]
sha256_password_private_key_path=mykey.pem
sha256_password_public_key_path=mykey.pub
```

If the files are not in the data directory, or to make their locations explicit in the option values, use full path names:

```
[mysqld]
```

```
sha256_password_private_key_path=/usr/local/mysql/mykey.pem
sha256_password_public_key_path=/usr/local/mysql/mykey.pub
```

- Restart the server, then connect to it and check the `Rsa_public_key` [597] status variable value. The value will differ from that shown here, but should be nonempty:

```
mysql> SHOW STATUS LIKE 'Rsa_public_key'\G
***** 1. row *****
Variable_name: Rsa_public_key
Value: -----BEGIN PUBLIC KEY-----
MIGfMA0GC斯qGSIB3DQEBAQUAA4GNADCBiQKBgQDO9nRUdd+KvSZgY7cNBZMNpwX6
MvE1PbJFXO7u18nJ91wc99Du/E71w6CVXw7VKrXPehbVQUzGyUNkf45Nz/ckaaJa
aLgJOBcIDmNVnyU54OT/1lcs2xiyfaDMe8fCJ64ZwTnKbY2gkt1IMjUAB5Ogd5kJ
g8aV7EtKwyhHb0c30QIDAQAB
-----END PUBLIC KEY-----
```

If the value is empty, the server found some problem with the key files. Check the error log for diagnostic information.

After the server has been configured with the RSA key files, clients can connect to it using accounts that authenticate with the `sha256_password` plugin. As mentioned previously, such accounts can use either an SSL connection (in which case RSA is not used) or a plain connection that encrypts the password using RSA. Assume for the following discussion that SSL is not used. Connecting to the server involves no special preparation on the client side. For example:

```
shell> mysql -u sha256user -p
Enter password: sha256P@ss
```

For connection attempts by `sha256user`, the server determines that `sha256_password` is the appropriate authentication plugin and invokes it. The plugin finds that the connection does not use SSL and thus requires the password to be transmitted using RSA encryption. It sends the RSA public key to the client, which uses it to encrypt the password and returns the result to the server. The plugin uses the RSA key on the server side to decrypt the password and accepts or rejects the connection based on whether the password is correct.

The server sends the public key to the client as needed, but if a copy of the RSA public key is available on the client host, the client can use it to save a round trip in the client/server protocol:

```
shell> mysql -u sha256user -p --server-public-key-path=file_name
```

The public key value in the file named by the `--server-public-key-path` [253] option should be the same as the key value in the server-side file named by the `sha256_password_public_key_path` [539] system variable. If the key file contains a valid public key value but the value is incorrect, an access-denied error occurs. If the key file does not contain a valid public key, the client program cannot use it. In this case, the server sends the public key to the client as if no `--server-public-key-path` [253] option had been specified.

Client users can get the RSA public key two ways:

- The database administrator can provide a copy of the public key file.
- A client user who can connect to the server some other way can use a `SHOW STATUS LIKE 'Rsa_public_key'` statement and save the returned key value in a file.

### 6.3.7.3. The PAM Authentication Plugin



#### Note

The PAM authentication plugin is a commercial extension. To learn more about commercial products (MySQL Enterprise Edition), see <http://www.mysql.com/products/>.

As of MySQL 5.6.10, commercial distributions of MySQL include an authentication plugin that enables MySQL Server to use PAM (Pluggable Authentication Modules) to authenticate MySQL users. PAM

enables a system to use a standard interface to access various kinds of authentication methods, such as Unix passwords or an LDAP directory.

The PAM plugin uses the information passed to it by MySQL Server (such as user name, host name, password, and authentication string), plus whatever method is available for PAM lookup. The plugin checks the user credentials against PAM and returns `'Authentication succeeded, Username is user_name'` or `'Authentication failed'`.

The PAM authentication plugin provides these capabilities:

- External authentication: The plugin enables MySQL Server to accept connections from users defined outside the MySQL grant tables.
- Proxy user support: The plugin can return to MySQL a user name different from the login user, based on the groups the external user is in and the authentication string provided. This means that the plugin can return the MySQL user that defines the privileges the external PAM-authenticated user should have. For example, a PAM user named `joe` can connect and have the privileges of the MySQL user named `developer`.

The following table shows the plugin and library file names. The file name suffix might be different on your system. The file location must be the directory named by the `plugin_dir [525]` system variable. For installation information, see [Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”](#).

**Table 6.11. MySQL PAM Authentication Plugin**

Server-side plugin name	<code>authentication_pam</code>
Client-side plugin name	<code>mysql_clear_password</code>
Library object file name	<code>authentication_pam.so</code>

The library file includes only the server-side plugin. The client-side plugin is built into the `libmysqlclient` client library. See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).

The server-side PAM authentication plugin is included only in commercial distributions. It is not included in MySQL community distributions. The client-side clear-text plugin that communicates with the server-side plugin is built into the MySQL client library and is included in all distributions, including community distributions. This permits clients from any 5.6.10 or newer distribution to connect to a server that has the server-side plugin loaded.

The PAM authentication plugin has been tested on Linux and Mac OS X. It requires MySQL Server 5.6.10 or newer.

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#). For proxy user information, see [Section 6.3.8, “Proxy Users”](#).

### 6.3.7.3.1. Installing the PAM Authentication Plugin

The PAM authentication plugin must be located in the MySQL plugin directory (the directory named by the `plugin_dir [525]` system variable). If necessary, set the value of `plugin_dir [525]` at server startup to tell the server the location of the plugin directory.

To enable the plugin, start the server with the `--plugin-load [435]` option. For example, put the following lines in your `my.cnf` file. If object files have a suffix different from `.so` on your system, substitute the correct suffix.

```
[mysqld]
plugin-load=authentication_pam.so
```

Use the plugin name `authentication_pam` in the `IDENTIFIED WITH` clause of `CREATE USER` or `GRANT` statements for MySQL accounts that should be authenticated with this plugin.

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement. See [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).



### 6.3.7.3.2. Using the PAM Authentication Plugin

This section describes how to use the PAM authentication plugin to connect from MySQL client programs to the server. It is assumed that the server-side plugin is enabled and that client programs are recent enough to include the client-side plugin.



#### Note

The client-side plugin with which the PAM plugin communicates simply sends the password to the server in clear text so it can be passed to PAM. This may be a security problem in some configurations, but is necessary to use the server-side PAM library. To avoid problems if there is any possibility that the password would be intercepted, clients should connect to MySQL Server using SSL. See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).

To refer to the PAM authentication plugin in the `IDENTIFIED WITH` clause of a `CREATE USER` or `GRANT` statement, use the name `authentication_pam`. For example:

```
CREATE USER user
  IDENTIFIED WITH authentication_pam
  AS 'authentication_string';
```

The authentication string specifies the following types of information:

- PAM supports the notion of “service name,” which is a name that the system administrator can use to configure the authentication method for a particular application. There can be several such “applications” associated with a single database server instance, so the choice of service name is left to the SQL application developer. When you define an account that should authenticate using PAM, specify the service name in the authentication string.
- PAM provides a way for a PAM module to return to the server a MySQL user name other than the login name supplied at login time. Use the authentication string to control the mapping between login name and MySQL user name. If you want to take advantage of proxy user capabilities, the authentication string must include this kind of mapping.

For example, if the service name is `mysql` and users in the `root` and `users` PAM groups should be mapped to the `developer` and `data_entry` users, respectively, use a statement like this:

```
CREATE USER user
  IDENTIFIED WITH authentication_pam
  AS 'mysql, root=developer, users=data_entry';
```

Authentication string syntax for the PAM authentication plugin follows these rules:

- The string consists of a PAM service name, optionally followed by a group mapping list consisting of one or more keyword/value pairs each specifying a group name and a SQL user name:

```
pam_service_name[,group_name=sql_user_name]...
```

- Each `group_name=sql_user_name` pair must be preceded by a comma.
- Leading and trailing spaces not inside double quotation marks are ignored.
- Unquoted `pam_service_name`, `group_name`, and `sql_user_name` values can contain anything except equal sign, comma, or space.
- If a `pam_service_name`, `group_name`, or `sql_user_name` value is quoted with double quotation marks, everything between the quotation marks is part of the value. This is necessary, for example, if the value contains space characters. All characters are legal except double quotation mark and backslash (`'\'`). To include either character, escape it with a backslash.

The plugin parses the authentication string on each login check. To minimize overhead, keep the string as short as possible.

If the plugin successfully authenticates a login name, it looks for a group mapping list in the authentication string and uses it to return a different user name to the MySQL server based on the groups the external user is a member of:

- If the authentication string contains no group mapping list, the plugin returns the login name.
- If the authentication string does contain a group mapping list, the plugin examines each `group_name=sql_user_name` pair in the list from left to right and tries to find a match for the `group_name` value in a non-MySQL directory of the groups assigned to the authenticated user and returns `sql_user_name` for the first match it finds. If the plugin finds no match for any group, it returns the login name. If the plugin is not capable of looking up a group in a directory, it ignores the group mapping list and returns the login name.

The following sections describe how to set up several authentication scenarios that use the PAM authentication plugin:

- No proxy users. This uses PAM only to check login names and passwords. Every external user permitted to connect to MySQL Server should have a matching MySQL account that is defined to use external PAM authentication. Authentication can be performed by various PAM-supported methods. The discussion shows how to use traditional Unix passwords and LDAP.

PAM authentication, when not done through proxy users or groups, requires the MySQL account to have the same user name as the Unix account. Because MySQL user names are limited to 16 characters (see [Section 6.2.2, “Privilege System Grant Tables”](#)), this limits PAM nonproxy authentication to Unix accounts with names of at most 16 characters.

- Proxy login only and group mapping. For this scenario, create a few MySQL accounts that define different sets of privileges. (Ideally, nobody should log in through these directly.) Then define a default user authenticating through PAM that uses some mapping scheme (usually by the external groups the users are in) to map all the external logins to the few MySQL accounts holding the privilege sets. Any user that logs in is mapped to one of the MySQL accounts and uses its privileges. The discussion shows how to set this up using Unix passwords, but other PAM methods such as LDAP could be used instead.

Variations on these scenarios are possible. For example, you can permit some users to log in directly but require others to connect through proxy users.

The examples make the following assumptions. You might need to make some adjustments if your system is set up differently.

- The PAM configuration directory is `/etc/pam.d`.
- The PAM service name is `mysql`, which means that you must set up a PAM file named `mysql` in the PAM configuration directory (creating the file if it does not exist). If you use a different service name, the file name will be different and you must use a different name in the `AS` clause of `CREATE USER` and `GRANT` statements.
- The examples use a login name of `antonio` and password of `verysecret`. Change these to correspond to the users you want to authenticate.

The PAM authentication plugin checks at initialization time whether the `AUTHENTICATION_PAM_LOG` environment value is set. If so, the plugin enables logging of diagnostic messages to the standard output. These messages may be helpful for debugging PAM-related problems that occur when the plugin performs authentication. For more information, see [Section 6.3.7.3.3, “PAM Authentication Plugin Debugging”](#).

#### 6.3.7.3.2.1. Unix Password Authentication without Proxy Users

This authentication scenario uses PAM only to check Unix user login names and passwords. Every external user permitted to connect to MySQL Server should have a matching MySQL account that is defined to use external PAM authentication.



1. Verify that Unix authentication in PAM permits you to log in as `antonio` with password `verysecret`.
2. Set up PAM to authenticate the `mysql` service. Put the following in `/etc/pam.d/mysql`:

```

#%PAM-1.0
auth            include      password-auth
account         include      password-auth

```

3. Create a MySQL account with the same user name as the Unix login name and define it to authenticate using the PAM plugin:

```

CREATE USER 'antonio'@'localhost'
  IDENTIFIED WITH authentication_pam AS 'mysql';
GRANT ALL PRIVILEGES ON mydb.* TO 'antonio'@'localhost';

```

4. Try to connect to the MySQL server using the `mysql` command-line client. For example:

```
mysql --user=antonio --password=verysecret mydb
```

The server should permit the connection and the following query should return output as shown:

```

mysql> SELECT USER(), CURRENT_USER(), @@proxy_user;
+-----+-----+-----+
| USER()          | CURRENT_USER() | @@proxy_user |
+-----+-----+-----+
| antonio@localhost | antonio@localhost | NULL         |
+-----+-----+-----+

```

This shows that `antonio` uses the privileges granted to the `antonio` MySQL account, and that no proxying has occurred.

#### 6.3.7.3.2.2. LDAP Authentication without Proxy Users

This authentication scenario uses PAM only to check LDAP user login names and passwords. Every external user permitted to connect to MySQL Server should have a matching MySQL account that is defined to use external PAM authentication.

1. Verify that LDAP authentication in PAM permits you to log in as `antonio` with password `verysecret`.
2. Set up PAM to authenticate the `mysql` service through LDAP. Put the following in `/etc/pam.d/mysql`:

```

#%PAM-1.0
auth            required     pam_ldap.so
account         required     pam_ldap.so

```

If PAM object files have a suffix different from `.so` on your system, substitute the correct suffix.

3. MySQL account creation and connecting to the server is the same as previously described in [Section 6.3.7.3.2.1, “Unix Password Authentication without Proxy Users”](#).

#### 6.3.7.3.2.3. Unix Password Authentication with Proxy Users and Group Mapping

This authentication scheme uses proxying and group mapping to map users who connect to the MySQL server through PAM onto a few MySQL accounts that define different sets of privileges. Users do not connect directly through the accounts that define the privileges. Instead, they connect through a default proxy user authenticating through PAM that uses a mapping scheme to map all the external logins to the few MySQL accounts holding the privileges. Any user who connects is mapped to one of the MySQL accounts and uses its privileges.

The procedure shown here uses Unix password authentication. To use LDAP instead, see the early steps of [Section 6.3.7.3.2.2, “LDAP Authentication without Proxy Users”](#).

1. Verify that Unix authentication in PAM permits you to log in as `antonio` with password `verysecret` and that `antonio` is a member of the `root` or `users` group.
2. Set up PAM to authenticate the `mysql` service. Put the following in `/etc/pam.d/mysql`:

```

#%PAM-1.0
auth            include            password-auth
account        include            password-auth

```

3. Create the default proxy user that maps the external PAM users to the proxied accounts. It maps external users from the `root` PAM group to the `developer` MySQL account and the external users from the `users` PAM group to the `data_entry` MySQL account:

```

CREATE USER ''@'
  IDENTIFIED WITH authentication_pam
  AS 'mysql, root=developer, users=data_entry';

```

The mapping list following the service name is required when you set up proxy users. Otherwise, the plugin cannot tell how to map the name of PAM groups to the proper proxied user name.

4. Create the proxied accounts that will be used to access the databases:

```

CREATE USER 'developer'@'localhost' IDENTIFIED BY 'very secret password';
GRANT ALL PRIVILEGES ON mydevdb.* TO 'developer'@'localhost';
CREATE USER 'data_entry'@'localhost' IDENTIFIED BY 'very secret password';
GRANT ALL PRIVILEGES ON mydb.* TO 'data_entry'@'localhost';

```

If you do not let anyone know the passwords for these accounts, other users cannot use them to connect directly to the MySQL server. Instead, it is expected that users will authenticate using PAM and that they will use the `developer` or `data_entry` account by proxy based on their PAM group.

5. Grant the `PROXY [755]` privilege to the proxy account for the proxied accounts:

```

GRANT PROXY ON 'developer'@'localhost' TO ''@';
GRANT PROXY ON 'data_entry'@'localhost' TO ''@';

```

6. Try to connect to the MySQL server using the `mysql` command-line client. For example:

```
mysql --user=antonio --password=verysecret mydb
```

The server authenticates the connection using the `''@'` account. The privileges `antonio` will have depends on what PAM groups he is a member of. If `antonio` is a member of the `root` PAM group, the PAM plugin maps `root` to the `developer` MySQL user name and returns that name to the server. The server verifies that `''@'` has the `PROXY [755]` privilege for `developer` and permits the connection. the following query should return output as shown:

```

mysql> SELECT USER(), CURRENT_USER(), @@proxy_user;
+-----+-----+-----+
| USER()          | CURRENT_USER()    | @@proxy_user |
+-----+-----+-----+
| antonio@localhost | developer@localhost | ''@'         |
+-----+-----+-----+

```

This shows that `antonio` uses the privileges granted to the `developer` MySQL account, and that proxying occurred through the default proxy user account.

If `antonio` is not a member of the `root` PAM group but is a member of the `users` group, a similar process occurs, but the plugin maps `user` group membership to the `data_entry` MySQL user name and returns that name to the server. In this case, `antonio` uses the privileges of the `data_entry` MySQL account:

```

mysql> SELECT USER(), CURRENT_USER(), @@proxy_user;
+-----+-----+-----+
| USER()          | CURRENT_USER()    | @@proxy_user |
+-----+-----+-----+

```

```
| antonio@localhost | data_entry@localhost | '@' |
+-----+-----+-----+-----+
```

### 6.3.7.3.3. PAM Authentication Plugin Debugging

The PAM authentication plugin checks at initialization time whether the `AUTHENTICATION_PAM_LOG` environment value is set (the value does not matter). If so, the plugin enables logging of diagnostic messages to the standard output. These messages may be helpful for debugging PAM-related problems that occur when the plugin performs authentication.

Some messages include reference to PAM plugin source files and line numbers, which enables plugin actions to be tied more closely to the location in the code where they occur.

The following transcript demonstrates the kind of information produced by enabling logging. It resulted from a successful proxy authentication attempt.

```
entering auth_pam_server
entering auth_pam_next_token
auth_pam_next_token:reading at [cups,admin=writer,everyone=reader], sep=[,]
auth_pam_next_token:state=PRESPACE, ptr=[cups,admin=writer,everyone=reader],
out=[]
auth_pam_next_token:state=IDENT, ptr=[cups,admin=writer,everyone=reader],
out=[]
auth_pam_next_token:state=AFTERSPACE, ptr=[,admin=writer,everyone=reader],
out=[cups]
auth_pam_next_token:state=DELIMITER, ptr=[,admin=writer,everyone=reader],
out=[cups]
auth_pam_next_token:state=DONE, ptr=[,admin=writer,everyone=reader],
out=[cups]
leaving auth_pam_next_token on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/parser.c:19
auth_pam_server:password 12345qq received
auth_pam_server:pam_start rc=0
auth_pam_server:pam_set_item(PAM_RUSER,gkodinov) rc=0
auth_pam_server:pam_set_item(PAM_RHOST,localhost) rc=0
entering auth_pam_server_conv
auth_pam_server_conv:PAM_PROMPT_ECHO_OFF [Password:] received
leaving auth_pam_server_conv on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/authentication.c:19
auth_pam_server:pam_authenticate rc=0
auth_pam_server:pam_acct_mgmt rc=0
auth_pam_server:pam_setcred(PAM_ESTABLISH_CRED) rc=0
auth_pam_server:pam_get_item rc=0
auth_pam_server:pam_setcred(PAM_DELETE_CRED) rc=0
entering auth_pam_map_groups
entering auth_pam_walk_namevalue_list
auth_pam_walk_namevalue_list:reading at: [admin=writer,everyone=reader]
entering auth_pam_next_token
auth_pam_next_token:reading at [admin=writer,everyone=reader], sep=[,]
auth_pam_next_token:state=PRESPACE, ptr=[admin=writer,everyone=reader], out=[]
auth_pam_next_token:state=IDENT, ptr=[admin=writer,everyone=reader], out=[]
auth_pam_next_token:state=AFTERSPACE, ptr=[=writer,everyone=reader],
out=[admin]
auth_pam_next_token:state=DELIMITER, ptr=[=writer,everyone=reader],
out=[admin]
auth_pam_next_token:state=DONE, ptr=[=writer,everyone=reader], out=[admin]
leaving auth_pam_next_token on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/parser.c:19
auth_pam_walk_namevalue_list:name=[admin]
entering auth_pam_next_token
auth_pam_next_token:reading at [writer,everyone=reader], sep=[,]
auth_pam_next_token:state=PRESPACE, ptr=[writer,everyone=reader], out=[]
auth_pam_next_token:state=IDENT, ptr=[writer,everyone=reader], out=[]
auth_pam_next_token:state=AFTERSPACE, ptr=[,everyone=reader], out=[writer]
auth_pam_next_token:state=DELIMITER, ptr=[,everyone=reader], out=[writer]
auth_pam_next_token:state=DONE, ptr=[,everyone=reader], out=[writer]
leaving auth_pam_next_token on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/parser.c:19
walk, &error_namevalue_list:value=[writer]
entering auth_pam_map_group_to_user
```

```

auth_pam_map_group_to_user:pam_user=gkodinov, name=admin, value=writer
examining member root
examining member gkodinov
substitution was made to mysql user writer
leaving auth_pam_map_group_to_user on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/authent
auth_pam_walk_namevalue_list:found mapping
leaving auth_pam_walk_namevalue_list on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/parser.
auth_pam_walk_namevalue_list returned 0
leaving auth_pam_map_groups on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/authent
auth_pam_server:authenticated_as=writer
auth_pam_server: rc=0
leaving auth_pam_server on
/Users/gkodinov/mysql/work/x-5.5.16-release-basket/release/plugin/pam-authentication-plugin/src/authent

```

### 6.3.7.4. The Windows Native Authentication Plugin



#### Note

The Windows authentication plugin is a commercial extension. To learn more about commercial products (MySQL Enterprise Edition), see <http://www.mysql.com/products/>.

As of MySQL 5.6.10, commercial distributions of MySQL for Windows include an authentication plugin that performs external authentication on Windows, enabling MySQL Server to use native Windows services to authenticate client connections. Users who have logged in to Windows can connect from MySQL client programs to the server based on the information in their environment without specifying an additional password.

The client and server exchange data packets in the authentication handshake. As a result of this exchange, the server creates a security context object that represents the identity of the client in the Windows OS. This identity includes the name of the client account. The Windows authentication plugin uses the identity of the client to check whether it is a given account or a member of a group. By default, negotiation uses Kerberos to authenticate, then NTLM if Kerberos is unavailable.

The Windows authentication plugin provides these capabilities:

- External authentication: The plugin enables MySQL Server to accept connections from users defined outside the MySQL grant tables.
- Proxy user support: The plugin can return to MySQL a user name different from the client user. This means that the plugin can return the MySQL user that defines the privileges the external Windows-authenticated user should have. For example, a Windows user named `joe` can connect and have the privileges of the MySQL user named `developer`.

The following table shows the plugin and library file names. The file location must be the directory named by the `plugin_dir` [525] system variable. For installation information, see [Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”](#).

**Table 6.12. MySQL Windows Authentication Plugin**

Server-side plugin name	<code>authentication_windows</code>
Client-side plugin name	<code>authentication_windows_client</code>
Library object file name	<code>authentication_windows.dll</code>

The library file includes only the server-side plugin. The client-side plugin is built into the `libmysqlclient` client library.

The server-side Windows authentication plugin is included only in commercial distributions. It is not included in MySQL community distributions. The client-side plugin is included in all distributions,

including community distributions. This permits clients from any distribution to connect to a server that has the server-side plugin loaded.

The Windows authentication plugin should work on Windows 2000 Professional and up. It requires MySQL Server 5.6.10 or newer.

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#). For proxy user information, see [Section 6.3.8, “Proxy Users”](#).

#### 6.3.7.4.1. Installing the Windows Authentication Plugin

The Windows authentication plugin must be installed in the MySQL plugin directory (the directory named by the `plugin_dir` [525] system variable). If necessary, set the value of `plugin_dir` [525] at server startup to tell the server the location of the plugin directory.

To enable the plugin, start the server with the `--plugin-load` [435] option. For example, put these lines in your `my.ini` file:

```
[mysqld]
plugin-load=authentication_windows.dll
```

Use the plugin name `authentication_windows` in the `IDENTIFIED WITH` clause of `CREATE USER` or `GRANT` statements for MySQL accounts that should be authenticated with this plugin.

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement. See [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).

#### 6.3.7.4.2. Using the Windows Authentication Plugin

The Windows authentication plugin supports the use of MySQL accounts such that users who have logged in to Windows can connect to the MySQL server without having to specify an additional password. It is assumed that the server-side plugin is enabled and that client programs are recent enough to include the client-side plugin built into `libmysqlclient`. Once the DBA has enabled the server-side plugin and set up accounts to use it, clients can connect using those accounts with no other setup required on their part.

To refer to the Windows authentication plugin in the `IDENTIFIED WITH` clause of a `CREATE USER` or `GRANT` statement, use the name `authentication_windows`. Suppose that the Windows users `Rafal` and `Tasha` should be permitted to connect to MySQL, as well as any users in the `Administrators` or `Power Users` group. To set this up, create a MySQL account named `sql_admin` that uses the Windows plugin for authentication:

```
CREATE USER sql_admin
  IDENTIFIED WITH authentication_windows
  AS 'Rafal, Tasha, Administrators, "Power Users"');
```

The plugin name is `authentication_windows`. The string following the `AS` keyword is the authentication string. It specifies that the Windows users named `Rafal` or `Tasha` are permitted to authenticate to the server as the MySQL user `sql_admin`, as are any Windows users in the `Administrators` or `Power Users` group. The latter group name contains a space, so it must be quoted with double quote characters.

After you create the `sql_admin` account, a user who has logged in to Windows can attempt to connect to the server using that account:

```
C:\> mysql --user=sql_admin
```

No password is required here. The `authentication_windows` plugin uses the Windows security API to check which Windows user is connecting. If that user is named `Rafal` or `Tasha`, or is in the `Administrators` or `Power Users` group, the server grants access and the client is authenticated as `sql_admin` and has whatever privileges are granted to the `sql_admin` account. Otherwise, the server denies access.

Authentication string syntax for the Windows authentication plugin follows these rules:

- The string consists of one or more user mappings separated by commas.
- Each user mapping associates a Windows user or group name with a MySQL user name:

```
win_user_or_group_name=sql_user_name
win_user_or_group_name
```

For the latter syntax, with no *sql\_user\_name* value given, the implicit value is the MySQL user created by the `CREATE USER` statement. Thus, these statements are equivalent:

```
CREATE USER sql_admin
  IDENTIFIED WITH authentication_windows
  AS 'Rafal, Tasha, Administrators, "Power Users"';

CREATE USER sql_admin
  IDENTIFIED WITH authentication_windows
  AS 'Rafal=sql_admin, Tasha=sql_admin, Administrators=sql_admin,
    "Power Users"=sql_admin';
```

- Each backslash (`'\'`) in a value must be doubled because backslash is the escape character in MySQL strings.
- Leading and trailing spaces not inside double quotation marks are ignored.
- Unquoted *win\_user\_or\_group\_name* and *sql\_user\_name* values can contain anything except equal sign, comma, or space.
- If a *win\_user\_or\_group\_name* and or *sql\_user\_name* value is quoted with double quotation marks, everything between the quotation marks is part of the value. This is necessary, for example, if the name contains space characters. All characters within double quotes are legal except double quotation mark and backslash. To include either character, escape it with a backslash.
- *win\_user\_or\_group\_name* values use conventional syntax for Windows principals, either local or in a domain. Examples (note the doubling of backslashes):

```
domain\\user
.\user
domain\group
.\group
BUILTIN\WellKnownGroup
```

When invoked by the server to authenticate a client, the plugin scans the authentication string left to right for a user or group match to the Windows user. If there is a match, the plugin returns the corresponding *sql\_user\_name* to the MySQL server. If there is no match, authentication fails.

A user name match takes preference over a group name match. Suppose that the Windows user named *win\_user* is a member of *win\_group* and the authentication string looks like this:

```
'win_group = sql_user1, win_user = sql_user2'
```

When *win\_user* connects to the MySQL server, there is a match both to *win\_group* and to *win\_user*. The plugin authenticates the user as *sql\_user2* because the more-specific user match takes precedence over the group match, even though the group is listed first in the authentication string.

Windows authentication always works for connections from the same computer on which the server is running. For cross-computer connections, both computers must be registered with Windows Active Directory. If they are in the same Windows domain, it is unnecessary to specify a domain name. It is also possible to permit connections from a different domain, as in this example:

```
CREATE USER sql_accounting
  IDENTIFIED WITH authentication_windows
```



```
AS 'SomeDomain\\Accounting';
```

Here `SomeDomain` is the name of the other domain. The backslash character is doubled because it is the MySQL escape character within strings.

MySQL supports the concept of proxy users whereby a client can connect and authenticate to the MySQL server using one account but while connected has the privileges of another account (see [Section 6.3.8, “Proxy Users”](#)). Suppose that you want Windows users to connect using a single user name but be mapped based on their Windows user and group names onto specific MySQL accounts as follows:

- The `local_user` and `MyDomain\domain_user` local and domain Windows users should map to the `local_wlad` MySQL account.
- Users in the `MyDomain\Developers` domain group should map to the `local_dev` MySQL account.
- Local machine administrators should map to the `local_admin` MySQL account.

To set this up, create a proxy account for Windows users to connect to, and configure this account so that users and groups map to the appropriate MySQL accounts (`local_wlad`, `local_dev`, `local_admin`). In addition, grant the MySQL accounts the privileges appropriate to the operations they need to perform. The following instructions use `win_proxy` as the proxy account, and `local_wlad`, `local_dev`, and `local_admin` as the proxied accounts.

1. Create the proxy MySQL account:

```
CREATE USER win_proxy
  IDENTIFIED WITH authentication_windows
  AS 'local_user = local_wlad,
     MyDomain\\domain_user = local_wlad,
     MyDomain\\Developers = local_dev,
     BUILTIN\Administrators = local_admin';
```

2. For proxying to work, the proxied accounts must exist, so create them:

```
CREATE USER local_wlad IDENTIFIED BY 'wlad_pass';
CREATE USER local_dev IDENTIFIED BY 'dev_pass';
CREATE USER local_admin IDENTIFIED BY 'admin_pass';
```

If you do not let anyone know the passwords for these accounts, other users cannot use them to connect directly to the MySQL server.

You should also issue `GRANT` statements (not shown) that grant each proxied account the privileges it needs.

3. The proxy account must have the `PROXY [755]` privilege for each of the proxied accounts:

```
GRANT PROXY ON local_wlad TO win_proxy;
GRANT PROXY ON local_dev TO win_proxy;
GRANT PROXY ON local_admin TO win_proxy;
```

Now the Windows users `local_user` and `MyDomain\domain_user` can connect to the MySQL server as `win_proxy` and when authenticated have the privileges of the account given in the authentication string—in this case, `local_wlad`. A user in the `MyDomain\Developers` group who connects as `win_proxy` has the privileges of the `local_dev` account. A user in the `BUILTIN\Administrators` group has the privileges of the `local_admin` account.

To configure authentication so that all Windows users who do not have their own MySQL account go through a proxy account, substitute the default proxy user (`' '@ '`) for `win_proxy` in the preceding instructions. For information about the default proxy user, see [Section 6.3.8, “Proxy Users”](#).

To use the Windows authentication plugin with Connector/Net connection strings in Connector/Net 6.4.4 and higher, see [Section 21.2.5.5, “Using the Windows Native Authentication Plugin”](#).

Additional control over the Windows authentication plugin is provided by the `authentication_windows_use_principal_name` [464] and `authentication_windows_log_level` [463] system variables. See [Section 5.1.4, “Server System Variables”](#).

### 6.3.7.5. The Cleartext Client-Side Authentication Plugin

As of MySQL 5.6.2, a client-side authentication plugin is available that sends the password to the server without hashing or encryption. This plugin is built into the MySQL client library.

The following table shows the plugin name.

**Table 6.13. MySQL Cleartext Authentication Plugin**

Server-side plugin name	None, see discussion
Client-side plugin name	<code>mysql_clear_password</code>
Library object file name	None (plugin is built in)

With native MySQL authentication, the client performs one-way hashing on the password before sending it to the server. This enables the client to avoid sending the password in clear text. See [Section 6.1.2.4, “Password Hashing in MySQL”](#). However, because the hash algorithm is one way, the original password cannot be recovered on the server side.

One-way hashing cannot be done for authentication schemes that require the server to receive the password as entered on the client side. In such cases, the `mysql_clear_password` client-side plugin can be used to send the password to the server in clear text. There is no corresponding server-side plugin. Rather, the client-side plugin can be used by any server-side plugin that needs a clear text password.

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#).



#### Note

Sending passwords in clear text may be a security problem in some configurations. To avoid problems if there is any possibility that the password would be intercepted, clients should connect to MySQL Server using a method that protects the password. Possibilities include SSL (see [Section 6.3.9, “Using SSL for Secure Connections”](#)), IPsec, or a private network.

As of MySQL 5.6.7, to make inadvertent use of this plugin less likely, it is required that clients explicitly enable it. This can be done several ways:

- Set the `LIBMYSQL_ENABLE_CLEARTEXT_PLUGIN` environment variable to a value that begins with `1`, `Y`, or `y`. This enables the plugin for all client connections.
- The `mysql`, `mysqladmin`, and `mysqlslap` client programs support an `--enable-cleartext-plugin` option that enables the plugin on a per-invocation basis.
- The `mysql_options()` C API function supports a `MYSQL_ENABLE_CLEARTEXT_PLUGIN` option that enables the plugin on a per-connection basis. Also, any program that uses `libmysqlclient` and reads option files can enable the plugin by including an `enable-cleartext-plugin` option in an option group read by the client library.

### 6.3.7.6. The Socket Peer-Credential Authentication Plugin

As of MySQL 5.6.2, a server-side authentication plugin is available that authenticates clients that connect from the local host through the Unix socket file.

The source code for this plugin can be examined as a relatively simple example demonstrating how to write a loadable authentication plugin.



The following table shows the plugin and library file names. The file name suffix might differ on your system. The file location is the directory named by the `plugin_dir` [525] system variable. For installation information, see [Section 6.3.7, “Pluggable Authentication”](#).

**Table 6.14. MySQL Socket Peer-Credential Authentication Plugin**

Server-side plugin name	<code>auth_socket</code>
Client-side plugin name	None, see discussion
Library object file name	<code>auth_socket.so</code>

The `auth_socket` authentication plugin authenticates clients that connect from the local host through the Unix socket file. The plugin uses the `SO_PEERCRED` socket option to obtain information about the user running the client program. The plugin checks whether the user name matches the MySQL user name specified by the client program to the server, and permits the connection only if the names match. The plugin can be built only on systems that support the `SO_PEERCRED` option, such as Linux.

Suppose that a MySQL account is created for a user named `valerie` who is to be authenticated by the `auth_socket` plugin for connections from the local host through the socket file:

```
CREATE USER 'valerie'@'localhost' IDENTIFIED WITH auth_socket;
```

If a user on the local host with a login name of `stefanie` invokes `mysql` with the option `--user=valerie` to connect through the socket file, the server uses `auth_socket` to authenticate the client. The plugin determines that the `--user` option value (`valerie`) differs from the client user's name (`stefanie`) and refuses the connection. If a user named `valerie` tries the same thing, the plugin finds that the user name and the MySQL user name are both `valerie` and permits the connection. However, the plugin refuses the connection even for `valerie` if the connection is made using a different protocol, such as TCP/IP.

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#).

### 6.3.7.7. The Test Authentication Plugin

MySQL includes a test plugin that authenticates using MySQL native authentication, but is a loadable plugin (not built in) and must be installed prior to use. It can authenticate against either normal or older (shorter) password hash values.

This plugin is intended for testing and development purposes, and not for use in production environments. The test plugin source code is separate from the server source, unlike the built-in native plugin, so it can be examined as a relatively simple example demonstrating how to write a loadable authentication plugin.

The following table shows the plugin and library file names. The file name suffix might differ on your system. The file location is the directory named by the `plugin_dir` [525] system variable. For installation information, see [Section 6.3.7, “Pluggable Authentication”](#).

**Table 6.15. MySQL Test Authentication Plugin**

Server-side plugin name	<code>test_plugin_server</code>
Client-side plugin name	<code>auth_test_plugin</code>
Library object file name	<code>auth_test_plugin.so</code>

Because the test plugin authenticates the same way as native MySQL authentication, provide the usual `--user` [212] and `--password` [211] options that you normally use for accounts that use native authentication when you connect to the server. For example:

```
shell> mysql --user=your_name --password=your_pass
```

For general information about pluggable authentication in MySQL, see [Section 6.3.7, “Pluggable Authentication”](#).

### 6.3.8. Proxy Users

When authentication to the MySQL server occurs by means of an authentication plugin, the plugin may request that the connecting (external) user be treated as a different user for privilege-checking purposes. This enables the external user to be a proxy for the second user; that is, to have the privileges of the second user. In other words, the external user is a “proxy user” (a user who can impersonate or become known as another user) and the second user is a “proxied user” (a user whose identity can be taken on by a proxy user).

This section describes how the proxy user capability works. For general information about authentication plugins, see [Section 6.3.7, “Pluggable Authentication”](#). If you are interested in writing your own authentication plugins that support proxy users, see [Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”](#).

For proxying to occur, these conditions must be satisfied:

- When a connecting client should be treated as a proxy user, the plugin must return a different name, to indicate the proxied user name.
- A proxy user account must be set up to be authenticated by the plugin. Use the `CREATE USER` or `GRANT` statement to associate an account with a plugin.
- A proxy user account must have the `PROXY [755]` privilege for the proxied account. Use the `GRANT` statement for this.

Consider the following definitions:

```
CREATE USER 'empl_external'@'localhost'
  IDENTIFIED WITH auth_plugin AS 'auth_string';
CREATE USER 'employee'@'localhost'
  IDENTIFIED BY 'employee_pass';
GRANT PROXY
  ON 'employee'@'localhost'
  TO 'empl_external'@'localhost';
```

When a client connects as `empl_external` from the local host, MySQL uses `auth_plugin` to perform authentication. If `auth_plugin` returns the `employee` user name to the server (based on the content of `'auth_string'` and perhaps by consulting some external authentication system), that serves as a request to the server to treat this client, for purposes of privilege checking, as the `employee` local user.

In this case, `empl_external` is the proxy user and `employee` is the proxied user.

The server verifies that proxy authentication for `employee` is possible for the `empl_external` user by checking whether `empl_external` has the `PROXY [755]` privilege for `employee`. (If this privilege had not been granted, an error would occur.)

When proxying occurs, the `USER()` and `CURRENT_USER()` functions can be used to see the difference between the connecting user and the account whose privileges apply during the current session. For the example just described, those functions return these values:

```
mysql> SELECT USER(), CURRENT_USER();
+-----+-----+
| USER()          | CURRENT_USER() |
+-----+-----+
| empl_external@localhost | employee@localhost |
+-----+-----+
```

The `IDENTIFIED WITH` clause that names the authentication plugin may be followed by an `AS` clause specifying a string that the server passes to the plugin when the user connects. It is up to each plugin

whether the `AS` clause is required. If it is required, the format of the authentication string depends on how the plugin intends to use it. Consult the documentation for a given plugin for information about the authentication string values it accepts.

## Granting the Proxy Privilege

A special `PROXY [755]` privilege is needed to enable an external user to connect as and have the privileges of another user. To grant this privilege, use the `GRANT` statement. For example:

```
GRANT PROXY ON 'proxied_user' TO 'proxy_user';
```

`proxy_user` must represent a valid externally authenticated MySQL user at connection time or connection attempts fail. `proxied_user` must represent a valid locally authenticated user at connection time or connection attempts fail.

The corresponding `REVOKE` syntax is:

```
REVOKE PROXY ON 'proxied_user' FROM 'proxy_user';
```

MySQL `GRANT` and `REVOKE` syntax extensions work as usual. For example:

```
GRANT PROXY ON 'a' TO 'b', 'c', 'd';
GRANT PROXY ON 'a' TO 'd' IDENTIFIED BY ...;
GRANT PROXY ON 'a' TO 'd' WITH GRANT OPTION;
GRANT PROXY ON 'a' TO '@';
REVOKE PROXY ON 'a' FROM 'b', 'c', 'd';
```

In the preceding example, `'@'` is the default proxy user and means “any user.” The default proxy user is discussed later in this section.

The `PROXY [755]` privilege can be granted in these cases:

- By `proxied_user` for itself: The value of `USER()` must exactly match `CURRENT_USER()` and `proxied_user`, for both the user name and host name parts of the account name.
- By a user that has `GRANT PROXY ... WITH GRANT OPTION` for `proxied_user`.

The `root` account created by default during MySQL installation has the `PROXY ... WITH GRANT OPTION [755]` privilege for `'@'`, that is, for all users. This enables `root` to set up proxy users, as well as to delegate to other accounts the authority to set up proxy users. For example, `root` can do this:

```
CREATE USER 'admin'@'localhost' IDENTIFIED BY 'test';
GRANT PROXY ON '@' TO 'admin'@'localhost' WITH GRANT OPTION;
```

Now the `admin` user can manage all the specific `GRANT PROXY` mappings. For example, `admin` can do this:

```
GRANT PROXY ON sally TO joe;
```

## Default Proxy Users

To specify that some or all users should connect using a given external plugin, create a “blank” MySQL user, set it up to use that plugin for authentication, and let the plugin return the real authenticated user name (if different from the blank user). For example, suppose that there exists a hypothetical plugin named `ldap_auth` that implements LDAP authentication:

```
CREATE USER '@' IDENTIFIED WITH ldap_auth AS 'O=Oracle, OU=MySQL';
CREATE USER 'developer'@'localhost' IDENTIFIED BY 'developer_pass';
CREATE USER 'manager'@'localhost' IDENTIFIED BY 'manager_pass';
GRANT PROXY ON 'manager'@'localhost' TO '@';
GRANT PROXY ON 'developer'@'localhost' TO '@';
```

Now assume that a client tries to connect as follows:

```
mysql --user=myuser --password='myuser_pass' ...
```

The server will not find `myuser` defined as a MySQL user. But because there is a blank user account (`' '@'`), that matches the client user name and host name, the server authenticates the client against that account: The server invokes `ldap_auth`, passing it `myuser` and `myuser_pass` as the user name and password.

If the `ldap_auth` plugin finds in the LDAP directory that `myuser_pass` is not the correct password for `myuser`, authentication fails and the server rejects the connection.

If the password is correct and `ldap_auth` finds that `myuser` is a developer, it returns the user name `developer` to the MySQL server, rather than `myuser`. The server verifies that `' '@'` can authenticate as `developer` (because it has the `PROXY [755]` privilege to do so) and accepts the connection. The session proceeds with `myuser` having the privileges of `developer`. (These privileges should be set up by the DBA using `GRANT` statements, not shown.) The `USER()` [1273] and `CURRENT_USER()` [1267] functions return these values:

```
mysql> SELECT USER(), CURRENT_USER();
+-----+-----+
| USER()          | CURRENT_USER() |
+-----+-----+
| myuser@localhost | developer@localhost |
+-----+-----+
```

If the plugin instead finds in the LDAP directory that `myuser` is a manager, it returns `manager` as the user name and the session proceeds with `myuser` having the privileges of `manager`.

```
mysql> SELECT USER(), CURRENT_USER();
+-----+-----+
| USER()          | CURRENT_USER() |
+-----+-----+
| myuser@localhost | manager@localhost |
+-----+-----+
```

For simplicity, external authentication cannot be multilevel: Neither the credentials for `developer` nor those for `manager` are taken into account in the preceding example. However, they are still used if a client tries to authenticate directly against the `developer` or `manager` account, which is why those accounts should be assigned passwords.

The default proxy account uses `' '@'` in the host part, which matches any host. If you set up a default proxy user, take care to also check for accounts with `' @%'` in the host part, because that also matches any host, but has precedence over `' '@'` by the rules that the server uses to sort account rows internally (see [Section 6.2.4, “Access Control, Stage 1: Connection Verification”](#)).

Suppose that a MySQL installation includes these two accounts:

```
CREATE USER '@' IDENTIFIED WITH some_plugin;
CREATE USER '@%' IDENTIFIED BY 'some_password';
```

The intent of the first account is to serve as the default proxy user, to be used to authenticate connections for users who do not otherwise match a more-specific account. The second account might have been created, for example, to enable users without their own account as the anonymous user.

However, in this configuration, the first account will never be used because the matching rules sort `' @%'` ahead of `' '@'`. For accounts that do not match any more-specific account, the server will attempt to authenticate them against `' @%'` rather than `' '@'`.

If you intend to create a default proxy user, check for other existing “match any user” accounts that will take precedence over the default proxy user and thus prevent that user from working as intended. It may be necessary to remove any such accounts.

## Proxy User System Variables

Two system variables help trace the proxy login process:

- `proxy_user` [526]: This value is `NULL` if proxying is not used. Otherwise, it indicates the proxy user account. For example, if a client authenticates through the default proxy account, this variable will be set as follows:

```
mysql> SELECT @@proxy_user;
+-----+
| @@proxy_user |
+-----+
| ''@''        |
+-----+
```

- `external_user` [482]: Sometimes the authentication plugin may use an external user to authenticate to the MySQL server. For example, when using Windows native authentication, a plugin that authenticates using the windows API does not need the login ID passed to it. However, it still uses an Windows user ID to authenticate. The plugin may return this external user ID (or the first 512 UTF-8 bytes of it) to the server using the `external_user` read-only session variable. If the plugin does not set this variable, its value is `NULL`.

### 6.3.9. Using SSL for Secure Connections

MySQL supports secure (encrypted) connections between MySQL clients and the server using the Secure Sockets Layer (SSL) protocol. This section discusses how to use SSL connections. For information on how to require users to use SSL connections, see the discussion of the `REQUIRE` clause of the `GRANT` statement in [Section 13.7.1.4, “GRANT Syntax”](#).

The standard configuration of MySQL is intended to be as fast as possible, so encrypted connections are not used by default. For applications that require the security provided by encrypted connections, the extra computation to encrypt the data is worthwhile.

MySQL enables encryption on a per-connection basis. You can choose an unencrypted connection or a secure encrypted SSL connection according to the requirements of individual applications.

Secure connections are based on the OpenSSL API and are available through the MySQL C API. Replication uses the C API, so secure connections can be used between master and slave servers. See [Section 16.3.7, “Setting Up Replication Using SSL”](#).

Another way to connect securely is from within an SSH connection to the MySQL server host. For an example, see [Section 6.3.10, “Connecting to MySQL Remotely from Windows with SSH”](#).

#### 6.3.9.1. Basic SSL Concepts

To understand how MySQL uses SSL, it is necessary to explain some basic SSL and X509 concepts. People who are familiar with these concepts can skip this part of the discussion.

By default, MySQL uses unencrypted connections between the client and the server. This means that someone with access to the network could watch all your traffic and look at the data being sent or received. They could even change the data while it is in transit between client and server. To improve security a little, you can compress client/server traffic by using the `--compress` option when invoking client programs. However, this does not foil a determined attacker.

When you need to move information over a network in a secure fashion, an unencrypted connection is unacceptable. Encryption is the way to make any kind of data unreadable. Encryption algorithms must include security elements to resist many kinds of known attacks such as changing the order of encrypted messages or replaying data twice.

SSL is a protocol that uses different encryption algorithms to ensure that data received over a public network can be trusted. It has mechanisms to detect any data change, loss, or replay. SSL also incorporates algorithms that provide identity verification using the X509 standard.

X509 makes it possible to identify someone on the Internet. It is most commonly used in e-commerce applications. In basic terms, there should be some entity called a “Certificate Authority” (or CA) that assigns electronic certificates to anyone who needs them. Certificates rely on asymmetric encryption

algorithms that have two encryption keys (a public key and a secret key). A certificate owner can show the certificate to another party as proof of identity. A certificate consists of its owner's public key. Any data encrypted with this public key can be decrypted only using the corresponding secret key, which is held by the owner of the certificate.

For more information about SSL, X509, encryption, or public-key cryptography, perform an Internet search for the keywords in which you are interested.

### 6.3.9.2. Configuring MySQL for SSL

To use SSL connections between the MySQL server and client programs, your system must support either OpenSSL or yaSSL, and your version of MySQL must be built with SSL support. To make it easier to use secure connections, MySQL is bundled with yaSSL, which uses the same licensing model as MySQL. (OpenSSL uses an Apache-style license.) yaSSL support is available on all MySQL platforms supported by Oracle Corporation.

To get secure connections to work with MySQL and SSL, you must do the following:

1. If you are not using a binary (precompiled) version of MySQL that has been built with SSL support, and you are going to use OpenSSL rather than the bundled yaSSL library, install OpenSSL if it has not already been installed. We have tested MySQL with OpenSSL 1.0.0. To obtain OpenSSL, visit <http://www.openssl.org>.

Building MySQL using OpenSSL requires a shared OpenSSL library, otherwise linker errors occur. Alternatively, build MySQL using yaSSL.

2. If you are not using a binary (precompiled) version of MySQL that has been built with SSL support, configure a MySQL source distribution to use SSL. When you configure MySQL, invoke `CMake` like this:

```
shell> cmake . -DWITH_SSL=bundled
```

That command configures the distribution to use the bundled yaSSL library. To use the system SSL library instead, specify the option like this instead:

```
shell> cmake . -DWITH_SSL=system
```

See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

Then compile and install the distribution.

On Unix platforms, yaSSL retrieves true random numbers from either either `/dev/urandom` or `/dev/random`. Bug#13164 lists workarounds for some very old platforms which do not support these devices.

3. To check whether a `mysqld` server supports SSL, examine the value of the `have_ssl` [487] system variable:

```
mysql> SHOW VARIABLES LIKE 'have_ssl';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_ssl      | YES   |
+-----+-----+
```

If the value is `YES`, the server supports SSL connections. If the value is `DISABLED`, the server is capable of supporting SSL connections but was not started with the appropriate `--ssl-xxx` options to enable them to be used; see [Section 6.3.9.3, “Using SSL Connections”](#).

### 6.3.9.3. Using SSL Connections

To enable SSL connections, your MySQL distribution must be built with SSL support, as described in [Section 6.3.9.2, “Configuring MySQL for SSL”](#). In addition, the proper SSL-related options must



be used to specify the appropriate certificate and key files. For a complete list of SSL options, see [Section 6.3.9.4, “SSL Command Options”](#).

To start the MySQL server so that it permits clients to connect using SSL, use the options that identify the certificate and key files the server uses when establishing a secure connection:

- `--ssl-ca` [810] identifies the Certificate Authority (CA) certificate.
- `--ssl-cert` [811] identifies the server public key certificate. This can be sent to the client and authenticated against the CA certificate that it has.
- `--ssl-key` [811] identifies the server private key.

For example, start the server like this:

```
shell> mysqld --ssl-ca=ca-cert.pem \
             --ssl-cert=server-cert.pem \
             --ssl-key=server-key.pem
```

Each option names a file in PEM format. For instructions on generating the required SSL certificate and key files, see [Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”](#). If you have a MySQL source distribution, you can also test your setup using the demonstration certificate and key files in the `mysql-test/std_data` directory of the distribution.

Similar options are used on the client side, although in this case, `--ssl-cert` [811] and `--ssl-key` [811] identify the client public and private key. Note that the Certificate Authority certificate, if specified, must be the same as used by the server.

To establish a secure connection to a MySQL server with SSL support, the options that a client must specify depend on the SSL requirements of the MySQL account used by the client. (See the discussion of the `REQUIRE` clause in [Section 13.7.1.4, “GRANT Syntax”](#).)

Suppose that you want to connect using an account that has no special SSL requirements or was created using a `GRANT` statement that includes the `REQUIRE SSL` option. As a recommended set of SSL options, start the server with at least `--ssl-cert` [811] and `--ssl-key` [811], and invoke the client with `--ssl-ca` [810]. A client can connect securely like this:

```
shell> mysql --ssl-ca=ca-cert.pem
```

To require that a client certificate also be specified, create the account using the `REQUIRE X509` option. Then the client must also specify the proper client key and certificate files or the server will reject the connection:

```
shell> mysql --ssl-ca=ca-cert.pem \
             --ssl-cert=client-cert.pem \
             --ssl-key=client-key.pem
```

A client can determine whether the current connection with the server uses SSL by checking the value of the `Ssl_cipher` [599] status variable. The value of `Ssl_cipher` [599] is nonempty if SSL is used, and empty otherwise. For example:

```
mysql> SHOW STATUS LIKE 'Ssl_cipher';
+-----+-----+
| Variable_name | Value                               |
+-----+-----+
| Ssl_cipher    | DHE-RSA-AES256-SHA                |
+-----+-----+
```

For the `mysql` client, an alternative is to use the `STATUS` or `\s` command and check the `SSL` line:

```
mysql> \s
...
SSL:                               Not in use
```

...

Or:

```
mysql> \s
...
SSL:                Cipher in use is DHE-RSA-AES256-SHA
...
```

The C API enables application programs to use SSL:

- To establish a secure connection, program, use the `mysql_ssl_set()` C API function to set the appropriate certificate options before calling `mysql_real_connect()`. See [Section 21.8.6.68](#), “`mysql_ssl_set()`”.
- To determine whether SSL is in use after the connection is established, use `mysql_get_ssl_cipher()`. A non-NULL return value indicates a secure connection and names the SSL cipher used for encryption. A NULL return value indicates that SSL is not being used. See [Section 21.8.6.33](#), “`mysql_get_ssl_cipher()`”.

Replication uses the C API, so secure connections can be used between master and slave servers. See [Section 16.3.7](#), “[Setting Up Replication Using SSL](#)”.

### 6.3.9.4. SSL Command Options

This section describes options that are used to specify whether to use SSL and the names of SSL certificate and key files. These options can be given on the command line or in an option file. They are not available unless MySQL has been built with SSL support. See [Section 6.3.9.2](#), “[Configuring MySQL for SSL](#)”.

**Table 6.16. SSL Option/Variable Summary**

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">have_openssl [487]</a>			Yes		Global	No
<a href="#">have_ssl [487]</a>			Yes		Global	No
<a href="#">skip-ssl [810]</a>	Yes	Yes				
<a href="#">ssl [810]</a>	Yes	Yes				
<a href="#">ssl-ca [810]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_ca</code>			Yes		Global	No
<a href="#">ssl-capath [810]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_capath</code>			Yes		Global	No
<a href="#">ssl-cert [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_cert</code>			Yes		Global	No
<a href="#">ssl-cipher [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_cipher</code>			Yes		Global	No
<a href="#">ssl-crl [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_crl</code>			Yes		Global	No
<a href="#">ssl-crlpath [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_crlpath</code>			Yes		Global	No



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">ssl-key [811]</a>	Yes	Yes			Global	No
- Variable: <code>ssl_key</code>			Yes		Global	No
<a href="#">ssl-verify-server-cert [811]</a>	Yes	Yes				

- [--ssl \[810\]](#)

For the server, this option specifies that the server permits SSL connections. For a client program, it permits the client to connect to the server using SSL, but this option is not sufficient in itself to cause an SSL connection to be used. As a recommended set of options to enable SSL connections, use at least [--ssl-cert \[811\]](#) and [--ssl-key \[811\]](#) on the server side and [--ssl-ca \[810\]](#) on the client side.

[--ssl \[810\]](#) is implied by other [--ssl-xxx](#) options as indicated in the descriptions for those options. For this reason, [--ssl \[810\]](#) is not usually specified explicitly. It is more often used explicitly in its opposite form to override other SSL options and indicate that SSL should *not* be used. To do this, specify the option as [--skip-ssl \[810\]](#) or [--ssl=0 \[810\]](#). For example, you might have SSL options specified in the `[client]` group of your option file to use SSL connections by default when you invoke MySQL client programs. To use an unencrypted connection instead, invoke the client program with [--skip-ssl \[810\]](#) on the command line to override the options in the option file.

Use of [--ssl \[810\]](#) does not *require* an SSL connection to be used, it only permits it. For example, if you specify this option for a client program but the server has not been configured to permit SSL connections, an unencrypted connection is used.

The secure way to require use of an SSL connection is to create a MySQL account that includes at least a `REQUIRE SSL` clause in the `GRANT` statement. In this case, connections for that account will be rejected unless MySQL supports SSL connections and the server and client have been started with the proper SSL options.

The `REQUIRE` clause permits other SSL-related restrictions as well. These can be used for stricter requirements than `REQUIRE SSL`. The description of `REQUIRE` in [Section 13.7.1.4, “GRANT Syntax”](#), provides additional detail about which SSL command options may or must be specified by clients that connect using accounts that are created using the various `REQUIRE` options.

- [--ssl-ca=file\\_name \[810\]](#)

The path to a file in PEM format that contains a list of trusted SSL certificate authorities. This option implies [--ssl \[810\]](#).

If you use SSL when establishing a client connection, you can tell the client not to authenticate the server certificate by specifying neither [--ssl-ca \[810\]](#) nor [--ssl-capath \[810\]](#). The server still verifies the client according to any applicable requirements established using `GRANT` statements for the client, and it still uses any [--ssl-ca \[810\]](#)/[--ssl-capath \[810\]](#) values that were passed to server at startup.

- [--ssl-capath=dir\\_name \[810\]](#)

The path to a directory that contains trusted SSL certificate authority certificates in PEM format. This option implies [--ssl \[810\]](#).

If you use SSL when establishing a client connection, you can tell the client not to authenticate the server certificate by specifying neither [--ssl-ca \[810\]](#) nor [--ssl-capath \[810\]](#). The server still verifies the client according to any applicable requirements established using `GRANT` statements for the client, and it still uses any [--ssl-ca \[810\]](#)/[--ssl-capath \[810\]](#) values that were passed to server at startup.

MySQL distributions built with OpenSSL support the `--ssl-capath [810]` option. Distributions built with yaSSL do not because yaSSL does not look in any directory and does not follow a chained certificate tree. yaSSL requires that all components of the CA certificate tree be contained within a single CA certificate tree and that each certificate in the file has a unique SubjectName value. To work around this yaSSL limitation, concatenate the individual certificate files comprising the certificate tree into a new file. Then specify the new file as the value of the `--ssl-capath [810]` option.

- `--ssl-cert=file_name [811]`

The name of the SSL certificate file in PEM format to use for establishing a secure connection. This option implies `--ssl [810]`.

- `--ssl-cipher=cipher_list [811]`

A list of permissible ciphers to use for SSL encryption. If no cipher in the list is supported, SSL connections will not work. This option implies `--ssl [810]`.

For greatest portability, `cipher_list` should be a list of one or more cipher names, separated by colons. This format is understood both by OpenSSL and yaSSL. Examples:

```
--ssl-cipher=AES128-SHA
--ssl-cipher=DHE-RSA-AES256-SHA:AES128-SHA
```

OpenSSL supports a more flexible syntax for specifying ciphers, as described in the OpenSSL documentation at <http://www.openssl.org/docs/apps/ciphers.html>. However, yaSSL does not, so attempts to use that extended syntax fails for a MySQL distribution built with yaSSL.

- `--ssl-crl=file_name [811]`

The path to a file containing certificate revocation lists in PEM format. This option implies `--ssl [810]`. If neither `--ssl-crl [811]` nor `--ssl-crlpath [811]` is given, no CRL checks are performed, even if the CA path contains certificate revocation lists.

MySQL distributions built with OpenSSL support the `--ssl-crl [811]` option. Distributions built with yaSSL do not because yaSSL revocation lists do not work with yaSSL.

This option was added in MySQL 5.6.3.

- `--ssl-crlpath=dir_name [811]`

The path to a directory that contains files containing certificate revocation lists in PEM format. This option implies `--ssl [810]`. If neither `--ssl-crl [811]` nor `--ssl-crlpath [811]` is given, no CRL checks are performed, even if the CA path contains certificate revocation lists.

MySQL distributions built with OpenSSL support the `--ssl-crlpath [811]` option. Distributions built with yaSSL do not because yaSSL revocation lists do not work with yaSSL.

This option was added in MySQL 5.6.3.

- `--ssl-key=file_name [811]`

The name of the SSL key file in PEM format to use for establishing a secure connection.

If the MySQL distribution was built using OpenSSL or (as of MySQL 5.6.3) yaSSL and the key file is protected by a passphrase, the program will prompt the user for the passphrase. The password must be given interactively; it cannot be stored in a file. If the passphrase is incorrect, the program continues as if it could not read the key. Before MySQL 5.6.3, if the MySQL distribution was built using yaSSL and the key file is protected by a passphrase, an error occurs.

- `--ssl-verify-server-cert [811]`

This option is available for client programs only, not the server. It causes the client to check the server's Common Name value in the certificate that the server sends to the client. The client verifies that name against the host name the client uses for connecting to the server, and the connection fails if there is a mismatch. This feature can be used to prevent man-in-the-middle attacks. Verification is disabled by default.

### 6.3.9.5. Setting Up SSL Certificates and Keys for MySQL

This section demonstrates how to set up SSL certificate and key files for use by MySQL servers and clients. The first example shows a simplified procedure such as you might use from the command line. The second shows a script that contains more detail. The first two examples are intended for use on Unix and both use the `openssl` command that is part of OpenSSL. The third example describes how to set up SSL files on Windows.

#### Example 1: Creating SSL Files from the Command Line on Unix

The following example shows a set of commands to create MySQL server and client certificate and key files. You will need to respond to several prompts by the `openssl` commands. To generate test files, you can press Enter to all prompts. To generate files for production use, you should provide nonempty responses.

```
# Create clean environment
shell> rm -rf newcerts
shell> mkdir newcerts && cd newcerts

# Create CA certificate
shell> openssl genrsa 2048 > ca-key.pem
shell> openssl req -new -x509 -nodes -days 3600 \
    -key ca-key.pem -out ca-cert.pem

# Create server certificate, remove passphrase, and sign it
# server-cert.pem = public key, server-key.pem = private key
shell> openssl req -newkey rsa:2048 -days 3600 \
    -nodes -keyout server-key.pem -out server-req.pem
shell> openssl rsa -in server-key.pem -out server-key.pem
shell> openssl x509 -req -in server-req.pem -days 3600 \
    -CA ca-cert.pem -CAkey ca-key.pem -set_serial 01 -out server-cert.pem

# Create client certificate, remove passphrase, and sign it
# client-cert.pem = public key, client-key.pem = private key
shell> openssl req -newkey rsa:2048 -days 3600 \
    -nodes -keyout client-key.pem -out client-req.pem
shell> openssl rsa -in client-key.pem -out client-key.pem
shell> openssl x509 -req -in client-req.pem -days 3600 \
    -CA ca-cert.pem -CAkey ca-key.pem -set_serial 01 -out client-cert.pem
```

After generating the certificates, verify them:

```
shell> openssl verify -CAfile ca-cert.pem server-cert.pem client-cert.pem
server-cert.pem: OK
client-cert.pem: OK
```

Now you have a set of files that can be used as follows:

- `ca-cert.pem`: Use this as the argument to `--ssl-ca` [810] on the server and client sides. (The CA certificate, if used, must be the same on both sides.)
- `server-cert.pem`, `server-key.pem`: Use these as the arguments to `--ssl-cert` [811] and `--ssl-key` [811] on the server side.
- `client-cert.pem`, `client-key.pem`: Use these as the arguments to `--ssl-cert` [811] and `--ssl-key` [811] on the client side.

To use the files to test SSL connections, see [Section 6.3.9.3, “Using SSL Connections”](#).

**Example 2: Creating SSL Files Using a Script on Unix**

Here is an example script that shows how to set up SSL certificate and key files for MySQL. After executing the script, use the files to test SSL connections as described in [Section 6.3.9.3, "Using SSL Connections"](#).

```

DIR=`pwd`/openssl
PRIV=$DIR/private

mkdir $DIR $PRIV $DIR/newcerts
cp /usr/share/ssl/openssl.cnf $DIR
replace ./demoCA $DIR -- $DIR/openssl.cnf

# Create necessary files: $database, $serial and $new_certs_dir
# directory (optional)

touch $DIR/index.txt
echo "01" > $DIR/serial

#
# Generation of Certificate Authority(CA)
#

openssl req -new -x509 -keyout $PRIV/cakey.pem -out $DIR/ca-cert.pem \
    -days 3600 -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# .....+++++
# .....+++++
# writing new private key to '/home/monty/openssl/private/cakey.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank
# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:
# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL admin
# Email Address []:

#
# Create server request and key
#

openssl req -new -keyout $DIR/server-key.pem -out \
    $DIR/server-req.pem -days 3600 -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# ..+++++
# .....+++++
# writing new private key to '/home/monty/openssl/server-key.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank

```

```

# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:
# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL server
# Email Address []:
#
# Please enter the following 'extra' attributes
# to be sent with your certificate request
# A challenge password []:
# An optional company name []:

#
# Remove the passphrase from the key
#
openssl rsa -in $DIR/server-key.pem -out $DIR/server-key.pem

#
# Sign server cert
#
openssl ca -cert $DIR/ca-cert.pem -policy policy_anything \
  -out $DIR/server-cert.pem -config $DIR/openssl.cnf \
  -infile $DIR/server-req.pem

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Enter PEM pass phrase:
# Check that the request matches the signature
# Signature ok
# The Subjects Distinguished Name is as follows
# countryName           :PRINTABLE:'FI'
# organizationName      :PRINTABLE:'MySQL AB'
# commonName            :PRINTABLE:'MySQL admin'
# Certificate is to be certified until Sep 13 14:22:46 2003 GMT
# (365 days)
# Sign the certificate? [y/n]:y
#
#
# 1 out of 1 certificate requests certified, commit? [y/n]y
# Write out database with 1 new entries
# Data Base Updated

#
# Create client request and key
#
openssl req -new -keyout $DIR/client-key.pem -out \
  $DIR/client-req.pem -days 3600 -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# .....++++++
# .....++++++
# writing new private key to '/home/monty/openssl/client-key.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank
# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:

```

```

# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL user
# Email Address []:
#
# Please enter the following 'extra' attributes
# to be sent with your certificate request
# A challenge password []:
# An optional company name []:

#
# Remove the passphrase from the key
#
openssl rsa -in $DIR/client-key.pem -out $DIR/client-key.pem

#
# Sign client cert
#

openssl ca -cert $DIR/ca-cert.pem -policy policy_anything \
  -out $DIR/client-cert.pem -config $DIR/openssl.cnf \
  -infiles $DIR/client-req.pem

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Enter PEM pass phrase:
# Check that the request matches the signature
# Signature ok
# The Subjects Distinguished Name is as follows
# countryName           :PRINTABLE:'FI'
# organizationName      :PRINTABLE:'MySQL AB'
# commonName            :PRINTABLE:'MySQL user'
# Certificate is to be certified until Sep 13 16:45:17 2003 GMT
# (365 days)
# Sign the certificate? [y/n]:y
#
#
# 1 out of 1 certificate requests certified, commit? [y/n]y
# Write out database with 1 new entries
# Data Base Updated

#
# Create a my.cnf file that you can use to test the certificates
#

cat <<EOF > $DIR/my.cnf
[client]
ssl-ca=$DIR/ca-cert.pem
ssl-cert=$DIR/client-cert.pem
ssl-key=$DIR/client-key.pem
[mysqld]
ssl-ca=$DIR/ca-cert.pem
ssl-cert=$DIR/server-cert.pem
ssl-key=$DIR/server-key.pem
EOF

```

### Example 3: Creating SSL Files on Windows

Download OpenSSL for Windows if it is not installed on your system. An overview of available packages can be seen here:

<http://www.slproweb.com/products/Win32OpenSSL.html>

Choose the Win32 OpenSSL Light or Win64 OpenSSL Light package, depending on your architecture (32-bit or 64-bit). The default installation location will be `C:\OpenSSL-Win32` or `C:\OpenSSL-Win64`, depending on which package you downloaded. The following instructions assume a default location of `C:\OpenSSL-Win32`. Modify this as necessary if you are using the 64-bit package.

If a message occurs during setup indicating '`...critical component is missing: Microsoft Visual C++ 2008 Redistributables`', cancel the setup and download one of the following packages as well, again depending on your architecture (32-bit or 64-bit):

- Visual C++ 2008 Redistributables (x86), available at:

<http://www.microsoft.com/downloads/details.aspx?familyid=9B2DA534-3E03-4391-8A4D-074B9F2BC1BF>

- Visual C++ 2008 Redistributables (x64), available at:

<http://www.microsoft.com/downloads/details.aspx?familyid=bd2a6171-e2d6-4230-b809-9a8d7548c1b6>

After installing the additional package, restart the OpenSSL setup procedure.

During installation, leave the default `C:\OpenSSL-Win32` as the install path, and also leave the default option '`Copy OpenSSL DLL files to the Windows system directory`' selected.

When the installation has finished, add `C:\OpenSSL-Win32\bin` to the Windows System Path variable of your server:

1. On the Windows desktop, right-click the My Computer icon, and select Properties.
2. Select the Advanced tab from the **System Properties** menu that appears, and click the **Environment Variables** button.
3. Under **System Variables**, select Path, then click the **Edit** button. The **Edit System Variable** dialogue should appear.
4. Add ' `;C:\OpenSSL-Win32\bin`' to the end (notice the semicolon).
5. Press OK 3 times.
6. Check that OpenSSL was correctly integrated into the Path variable by opening a new command console (`Start>Run>cmd.exe`) and verifying that OpenSSL is available:

```
Microsoft Windows [Version ...]
Copyright (c) 2006 Microsoft Corporation. All rights reserved.

C:\Windows\system32>cd \

C:\>openssl
OpenSSL> exit <<< If you see the OpenSSL prompt, installation was successful.

C:\>
```

Depending on your version of Windows, the preceding path-setting instructions might differ slightly.

After OpenSSL has been installed, use instructions similar to those from from Example 1 (shown earlier in this section), with the following changes:

- Change the following Unix commands:

```
# Create clean environment
shell> rm -rf newcerts
shell> mkdir newcerts && cd newcerts
```

On Windows, use these commands instead:

```
# Create clean environment
shell> md c:\newcerts
shell> cd c:\newcerts
```

- When a '`\`' character is shown at the end of a command line, this '`\`' character must be removed and the command lines entered all on a single line.

After generating the certificate and key files, to use them to test SSL connections, see [Section 6.3.9.3, "Using SSL Connections"](#).

### 6.3.10. Connecting to MySQL Remotely from Windows with SSH

This section describes how to get a secure connection to a remote MySQL server with SSH. The information was provided by David Carlson <[dcarlson@mplcomm.com](mailto:dcarlson@mplcomm.com)>.

1. Install an SSH client on your Windows machine. As a user, the best nonfree one I have found is from *SecureCRT* from <http://www.vandyke.com/>. Another option is *f-secure* from <http://www.f-secure.com/>. You can also find some free ones on *Google* at <http://directory.google.com/Top/Computers/Internet/Protocols/SSH/Clients/Windows/>.
2. Start your Windows SSH client. Set `Host_Name = yourmysqlserver_URL_or_IP`. Set `userid=your_userid` to log in to your server. This `userid` value might not be the same as the user name of your MySQL account.
3. Set up port forwarding. Either do a remote forward (Set `local_port: 3306, remote_host: yourmysqlservername_or_ip, remote_port: 3306`) or a local forward (Set `port: 3306, host: localhost, remote port: 3306`).
4. Save everything, otherwise you will have to redo it the next time.
5. Log in to your server with the SSH session you just created.
6. On your Windows machine, start some ODBC application (such as Access).
7. Create a new file in Windows and link to MySQL using the ODBC driver the same way you normally do, except type in `localhost` for the MySQL host server, not `yourmysqlservername`.

At this point, you should have an ODBC connection to MySQL, encrypted using SSH.

### 6.3.11. MySQL Enterprise Audit Log Plugin



#### Note

MySQL Enterprise Audit is a commercial extension. To learn more about commercial products (MySQL Enterprise Edition), see <http://www.mysql.com/products/>.

As of MySQL 5.6.10, MySQL Enterprise Edition includes MySQL Enterprise Audit, implemented using a server plugin named `audit_log`. MySQL Enterprise Audit uses the open MySQL Audit API to enable standard, policy-based monitoring and logging of connection and query activity executed on specific MySQL servers. Designed to meet the Oracle audit specification, MySQL Enterprise Audit provides an out of box, easy to use auditing and compliance solution for applications that are governed by both internal and external regulatory guidelines.

When installed, the audit plugin enables MySQL Server to produce a log file containing an audit record of server activity. The log contents include when clients connect and disconnect, and what actions they perform while connected, such as which databases and tables they access.

After you install the plugin (see [Section 6.3.11.1, “Installing the Audit Log Plugin”](#)), it writes an audit log file. By default, the file is named `audit.log` in the server data directory. To change the name of the file, set the `audit_log_file` [825] system variable at server startup.

The audit log file is written in XML format, with auditable events encoded as `<AUDIT_RECORD>` elements. For details on file format and contents, see [Section 6.3.11.2, “The Audit Log File”](#).

To control what information `audit_log` writes to its log file, set the `audit_log_policy` [826] system variable. By default, this variable is set to `ALL` (write all auditable events), but also permits values of `LOGINS` or `QUERIES` to log only login or query events, or `NONE` to disable logging.

For more information about controlling how logging occurs, see [Section 6.3.11.3, “Audit Log Plugin Logging Control”](#). For descriptions of the parameters used to configure the audit log plugin, see [Section 6.3.11.4, “Audit Log Plugin Options and Variables”](#).



The Performance Schema (see [Chapter 20, MySQL Performance Schema](#)) has instrumentation for the audit log plugin. To identify the relevant instruments, use this query:

```
SELECT NAME FROM performance_schema.setup_instruments
WHERE NAME LIKE '/alog/%';
```

### 6.3.11.1. Installing the Audit Log Plugin

The audit log plugin is named `audit_log`. To be usable by the server, the plugin library object file must be located in the MySQL plugin directory (the directory named by the `plugin_dir` [525] system variable). If necessary, set the value of `plugin_dir` [525] at server startup to tell the server the location of the plugin directory.

To load the plugin at server startup, use the `--plugin-load` [435] option to name the object file that contains the plugin. With this plugin-loading method, the option must be given each time the server starts. For example, put the following lines in your `my.cnf` file:

```
[mysqld]
plugin-load=audit_log.so
```

If object files have a suffix different from `.so` on your system, substitute the correct suffix (for example, `.dll` on Windows).

Alternatively, to register the plugin at runtime, use this statement (changing the suffix as necessary):

```
mysql> INSTALL PLUGIN audit_log SONAME 'audit_log.so';
```

`INSTALL PLUGIN` loads the plugin, and also registers it in the `mysql.plugins` table to cause the plugin to be loaded for each subsequent normal server startup.

If the plugin is loaded with `--plugin-load` [435] or has been previously registered with `INSTALL PLUGIN`, you can use the `--audit-log` option at server startup to control plugin activation. For example, to load the plugin and prevent it from being removed at runtime, use these options:

```
[mysqld]
plugin-load=audit_log.so
audit-log=FORCE_PLUS_PERMANENT
```

If it is desired to prevent the server from running without the audit plugin, use `--audit-log` [824] with a value of `FORCE` or `FORCE_PLUS_PERMANENT` to force server startup to fail if the plugin does not initialize successfully.

For general information about installing plugins, see [Section 5.1.8, “Server Plugins”](#). To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement. See [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).

For security reasons, the audit log file produced by `audit_log` should be written to a directory accessible only to the MySQL server and users with a legitimate reason to view the log. The default file is `audit.log` in the data directory. This can be changed by setting the `audit_log_file` [825] system variable at server startup.

For additional information about the parameters used to configure operation of the `audit_log` plugin, see [Section 6.3.11.4, “Audit Log Plugin Options and Variables”](#).

### 6.3.11.2. The Audit Log File

The audit log file is written as XML, using UTF-8 (up to 4 bytes per character). The contents are not encrypted. The root element is `<AUDIT>`. The closing `</AUDIT>` tag of the root element is written when the plugin terminates, so the tag is not present in the file while the audit log plugin is active.

The root element contains `<AUDIT_RECORD>` elements. Each `<AUDIT_RECORD>` element has an empty body; all audit record fields are represented by element attributes.

Here is a sample log file, reformatted for readability:

```
<?xml version="1.0" encoding="UTF-8"?>
```

```

<AUDIT>
  <AUDIT_RECORD
    TIMESTAMP="2012-08-02T14:52:12"
    NAME="Audit"
    SERVER_ID="1"
    VERSION="1"
    STARTUP_OPTIONS="--port=3306"
    OS_VERSION="i686-Linux"
    MYSQL_VERSION="5.6.10-debug-log" />
  <AUDIT_RECORD
    TIMESTAMP="2012-08-02T14:52:41"
    NAME="Connect"
    CONNECTION_ID="1"
    STATUS="0"
    USER="root"
    PRIV_USER="root"
    OS_LOGIN=""
    PROXY_USER=""
    HOST="localhost"
    IP="127.0.0.1"
    DB="" />
  <AUDIT_RECORD
    TIMESTAMP="2012-08-02T14:53:45"
    NAME="Query"
    CONNECTION_ID="1"
    STATUS="0"
    SQLTEXT="INSERT INTO t1 () VALUES()" />
  <AUDIT_RECORD
    TIMESTAMP="2012-08-02T14:53:51"
    NAME="Quit"
    CONNECTION_ID="1"
    STATUS="0" />
  <AUDIT_RECORD
    TIMESTAMP="2012-08-06T14:21:03"
    NAME="NoAudit"
    SERVER_ID="1" />
</AUDIT>

```

Attributes of `<AUDIT_RECORD>` elements have these characteristics:

- Some attributes appear in every element, but most are optional and do not necessarily appear in every element.
- Order of attributes within an element is not guaranteed.
- Attribute values are not fixed length. Long values may be truncated as indicated in the attribute descriptions given later.
- The `<`, `>`, `"`, and `&` characters are encoded as `&lt;`, `&gt;`, `&quot;`, and `&amp;`, respectively. Null bytes (U+00) are encoded as the `?` character.
- Characters not valid as XML characters are encoded using numeric character references. Valid XML characters are:

```
#x9 | #xA | #xD | [#x20-#xD7FF] | [#xE000-#xFFFF] | [#x10000-#x10FFFF]
```

Every `<AUDIT_RECORD>` element contains a set of mandatory attributes. Other optional attributes may appear depending on the audit record type.

The following attributes are mandatory in every `<AUDIT_RECORD>` element:

- **NAME**

A string representing the type of instruction that generated the audit event, such as a command that the server received from a client.

Example: `NAME="Query"`

Some common **NAME** values:

"Audit"	When auditing starts, which may be server startup time
"Connect"	When a client connects, also known as logging in
"Query"	An SQL statement (executed directly)
"Prepare"	Preparation of an SQL statement; usually followed by Execute
"Execute"	Execution of an SQL statement; usually follows Prepare
"Shutdown"	Server shutdown
"Quit"	When a client disconnects
"NoAudit"	Auditing has been turned off

The possible values are "Audit", "Binlog Dump", "Change user", "Close stmt", "Connect Out", "Connect", "Create DB", "Daemon", "Debug", "Delayed insert", "Drop DB", "Execute", "Fetch", "Field List", "Init DB", "Kill", "Long Data", "NoAudit", "Ping", "Prepare", "Processlist", "Query", "Quit", "Refresh", "Register Slave", "Reset stmt", "Set option", "Shutdown", "Sleep", "Statistics", "Table Dump", "Time".

With the exception of "Audit" and "NoAudit", these values correspond to the COM\_xxx command values listed in the `mysql_com.h` header file. For example, "Create DB" and "Shutdown" correspond to `COM_CREATE_DB` and `COM_SHUTDOWN`, respectively.

- **TIMESTAMP**

The date and time that the audit event was generated. For example, the event corresponding to execution of an SQL statement received from a client has a **TIMESTAMP** value occurring after the statement finishes, not when it is received. The value is UTC, in the format `yyyy-mm-ddThh:mm:ss` (with **T**, no decimals).

Example: `TIMESTAMP="2012-08-09T12:55:16"`

The following attributes are optional in `<AUDIT_RECORD>` elements. Many of them occur only for elements with specific values of the **NAME** attribute.

- **CONNECTION\_ID**

An unsigned integer representing the client connection identifier. This is the same as the `CONNECTION_ID()` [1267] function value within the session.

Example: `CONNECTION_ID="1"`

- **DB**

A string representing the default database name. This attribute appears only if the **NAME** value is "Connect" or "Change user".

- **HOST**

A string representing the client host name. This attribute appears only if the **NAME** value is "Connect" or "Change user".

Example: `HOST="localhost"`

- **IP**

A string representing the client IP address. This attribute appears only if the **NAME** value is "Connect" or "Change user".

Example: `IP="127.0.0.1"`

- **MYSQL\_VERSION**

A string representing the MySQL server version. This is the same as the value of the `VERSION()` [1273] function or `version` [563] system variable. This attribute appears only if the **NAME** value is "Audit".

Example: `MYSQL_VERSION="5.6.10-debug-log"`

- `OS_LOGIN`

A string representing the external user (empty if none). The value may differ from `USER`, for example, if the server authenticates the client using an external authentication method. This attribute appears only if the `NAME` value is `"Connect"` or `"Change user"`.

- `OS_VERSION`

A string representing the operating system on which the server was built or is running. This attribute appears only if the `NAME` value is `"Audit"`.

Example: `OS_VERSION="x86_64-Linux"`

- `PRIV_USER`

A string representing the user that the server authenticated the client as. This is the user name that the server uses for privilege checking, and may be different from the `USER` value. This attribute appears only if the `NAME` value is `"Connect"` or `"Change user"`.

- `PROXY_USER`

A string representing the proxy user. The value is empty if user proxying is not in effect. This attribute appears only if the `NAME` value is `"Connect"` or `"Change user"`.

- `SERVER_ID`

An unsigned integer representing the server ID. This is the same as the value of the `server_id` [539] system variable. This attribute appears only if the `NAME` value is `"Audit"` or `"NoAudit"`.

Example: `SERVER_ID="1"`

- `SQLTEXT`

A string representing the text of an SQL statement. The value can be empty. Long values may be truncated. This attribute appears only if the `NAME` value is `"Query"` or `"Execute"`.

The string, like the audit log file itself, is written using UTF-8 (up to 4 bytes per character), so the value may be the result of conversion. For example, the original statement might have been received from the client as an SJIS string.

Example: `SQLTEXT="DELETE FROM t1"`

- `STARTUP_OPTIONS`

A string representing the options that were given on the command line or in option files when the MySQL server was started. This attribute appears only if the `NAME` value is `"Audit"`.

Example: `STARTUP_OPTIONS="--port=3306 --log-output=FILE"`

- `STATUS`

An unsigned integer representing the command status. The value is zero for success, nonzero if an error occurred. This is the same as the value of the `mysql_errno()` C API function.

The audit log does not contain the `SQLSTATE` value or error message. To see the associations between error codes, `SQLSTATE` values, and messages, see [Section C.3, "Server Error Codes and Messages"](#).

Warnings are not logged.

Example: `STATUS="1051"`

- **USER**

A string representing the user name sent by the client. This may be different from the `PRIV_USER` value. This attribute appears only if the `NAME` value is `"Connect"` or `"Change user"`.

- **VERSION**

An unsigned integer representing the version of the audit log file format. This attribute appears only if the `NAME` value is `"Audit"`.

Example: `VERSION="1"`

### 6.3.11.3. Audit Log Plugin Logging Control

This section describes how the `audit_log` plugin performs logging and the system variables that control how logging occurs. It assumes familiarity with the log file format described in [Section 6.3.11.2, “The Audit Log File”](#).

When the audit log plugin opens its log file, it checks whether the XML declaration and opening `<AUDIT>` root element tag need to be written and writes them if so. When the audit log plugin terminates, it writes a closing `</AUDIT>` tag to the file.

If the log file exists at open time, the plugin checks whether the file ends with an `</AUDIT>` tag and truncates it if so before writing any `<AUDIT_RECORD>` elements. If the log file exists but does not end with `</AUDIT>` or the `</AUDIT>` tag cannot be truncated, the plugin considers the file malformed and fails to initialize. This can occur if the server crashes or is killed with the audit log plugin running. No logging occurs until the problem is rectified. Check the error log for diagnostic information:

```
[ERROR] Plugin 'audit_log' init function returned error.
```

To deal with this problem, you must either remove or rename the malformed log file and restart the server.

The MySQL server calls the audit log plugin to write an `<AUDIT_RECORD>` element whenever an auditable event occurs, such as when it completes execution of an SQL statement received from a client. Typically the first `<AUDIT_RECORD>` element written after server startup has the server description and startup options. Elements following that one represent events such as client connect and disconnect events, executed SQL statements, and so forth. Only top-level statements are logged, not statements within stored programs such as triggers or stored procedures. Contents of files referenced by statements such as `LOAD DATA INFILE` are not logged.

To permit control over how logging occurs, the `audit_log` plugin provides several system variables, described following. For more information, see [Section 6.3.11.4, “Audit Log Plugin Options and Variables”](#).

- `audit_log_file` [825]: The name of the log file. By default, the name is `audit.log` in the server data directory. For security reasons, the audit log file should be written to a directory accessible only to the MySQL server and users with a legitimate reason to view the log.
- `audit_log_policy` [826]: What kinds of information the plugin writes. By default, this variable is set to `ALL` (write all auditable events), but also permits values of `LOGINS` or `QUERIES` to log only login or query events, or `NONE` to disable logging.
- `audit_log_strategy` [827]: The method used for log writes. By default, the strategy value is `ASYNCHRONOUS` and the plugin logs asynchronously to a buffer, waiting if the buffer is full. It's possible to tell the plugin not to wait (`PERFORMANCE`) or to log synchronously, either using file system caching (`SEMISYNCHRONOUS`) or forcing output with a `sync()` call after each write request (`SYNCHRONOUS`).

Asynchronous logging strategy has these characteristics:

- Minimal impact on server performance and scalability.
- Blocking of threads that generate audit events for the shortest possible time; that is, time to allocate the buffer plus time to copy the event to the buffer.
- Output goes to the buffer. Writes from the buffer to the log file are handled by a separate thread.

A disadvantage of `PERFORMANCE` strategy is that it drops events when the buffer is full. For a heavily loaded server, it is more likely that the audit log will be missing events.

With asynchronous logging, the integrity of the log file may be compromised if a problem occurs during a write to the file or if the plugin does not shut down cleanly (for example, in the event that the server host crashes). To reduce this risk, set `audit_log_strategy` [827] use synchronous logging. Regardless of strategy, logging occurs on a best-effort basis, with no guarantee of consistency.

- `audit_log_buffer_size` [824]: The size of the buffer for asynchronous logging. The plugin uses a single buffer, which it allocates when it initializes and removes when it terminates. The plugin allocates this buffer only if logging is asynchronous.
- `audit_log_rotate_on_size` [826], `audit_log_flush` [825]: These variables permit audit log file rotation and flushing. The audit log file has the potential to grow very large and consume a lot of disk space. To manage the space used, either enable automatic log rotation, or manually rename the audit file and flush the log to open a new file. The renamed file can be removed or backed up as desired.

By default, `audit_log_rotate_on_size=0` [826] and there is no log rotation. In this case, the audit log plugin closes and reopens the log file when the `audit_log_flush` [825] value changes from disabled to enabled. Log file renaming must be done externally to the server. Suppose that you want to maintain the three most recent log files, which cycle through the names `audit.log.1` through `audit.log.3`. On Unix, perform rotation manually like this:

1. From the command line, rename the current log files:

```
shell> mv audit.log.2 audit.log.3
shell> mv audit.log.1 audit.log.2
shell> mv audit.log audit.log.1
```

At this point, the plugin is still writing to the current log file, which has been renamed to `audit.log.1`.

2. Connect to the server and flush the log file so the plugin closes it and reopens a new `audit.log` file:

```
mysql> SET GLOBAL audit_log_flush = ON;
```

If `audit_log_rotate_on_size` [826] is greater than 0, setting `audit_log_flush` [825] has no effect. In this case, the audit log plugin closes and reopens its log file whenever a write to the file causes its size to exceed the `audit_log_rotate_on_size` [826] value. The plugin renames the original file to have a timestamp extension. For example, `audit.log` might be renamed to `audit.log.13440033615657730`. The last 7 digits are a fractional second part. The first 10 digits are a Unix timestamp value that can be interpreted using the `FROM_UNIXTIME()` [1218] function:

```
mysql> SELECT FROM_UNIXTIME(1344003361);
+-----+
| FROM_UNIXTIME(1344003361) |
+-----+
| 2012-08-03 09:16:01      |
+-----+
```

### 6.3.11.4. Audit Log Plugin Options and Variables

This section describes the command options and system variables that control operation of the audit log plugin. If values specified at startup time are incorrect, the plugin may fail to initialize properly and the server does not load it. In this case, the server may also produce error messages for other audit log settings because it will not recognize them.

To control the activation of the `audit_log` plugin, use this option:

- `--audit-log[=value]` [824]

<b>Version Introduced</b>	5.6.10	
<b>Command-Line Format</b>	<code>--audit-log[=value]</code>	
<b>Option-File Format</b>	<code>audit-log[=value]</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	ON
	<b>Valid Values</b>	ON OFF FORCE FORCE_PLUS_PERMANENT

This option controls how the server loads the `audit_log` plugin at startup. It is available only if the audit log plugin has been previously registered with `INSTALL PLUGIN` or is loaded with `--plugin-load` [435]. See Section 6.3.11.1, “Installing the Audit Log Plugin”.

The option value should be one of those available for plugin-loading options, as described in Section 5.1.8.1, “Installing and Uninstalling Plugins”. For example, `--audit-log=FORCE_PLUS_PERMANENT` [824] tells the server to load the plugin and prevent it from being removed while the server is running.

This option was added in MySQL 5.6.10.

If the `audit_log` plugin is installed, it exposes several system variables that permit control over logging:

```
mysql> SHOW VARIABLES LIKE 'audit_log%';
+-----+-----+
| Variable_name | Value          |
+-----+-----+
| audit_log_buffer_size | 1048576       |
| audit_log_file      | audit.log     |
| audit_log_flush     | OFF           |
| audit_log_policy     | ALL           |
| audit_log_rotate_on_size | 0             |
| audit_log_strategy   | ASYNCHRONOUS |
+-----+-----+
```

You can set any of these variables at server startup, and some of them at runtime.

- `audit_log_buffer_size` [824]

<b>Version Introduced</b>	5.6.10
<b>Variable Name</b>	<code>audit_log_buffer_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

	<b>Permitted Values</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	4096 .. 4294967295
	<b>Permitted Values</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	1048576
<b>Range</b>	4096 .. 18446744073709547520

When the audit log plugin writes events to the log asynchronously, it uses a buffer to store event contents prior to writing them. This variable controls the size of that buffer, in bytes. The server adjusts the value to a multiple of 4096. The plugin uses a single buffer, which it allocates when it initializes and removes when it terminates. The plugin allocates this buffer only if logging is asynchronous.

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

- [audit\\_log\\_file \[825\]](#)

<b>Version Introduced</b>	5.6.10
<b>Variable Name</b>	<code>audit_log_file</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	string
<b>Default</b>	<code>audit.log</code>

The name of the file to which the audit log plugin writes events. The default value is `audit.log`. If the file name is a relative path, the server interprets it relative to the data directory. For security reasons, the audit log file should be written to a directory accessible only to the MySQL server and users with a legitimate reason to view the log.

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

- [audit\\_log\\_flush \[825\]](#)

<b>Version Introduced</b>	5.6.10
<b>Variable Name</b>	<code>audit_log_flush</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	OFF

When this variable is set to enabled (1 or ON), the audit log plugin closes and reopens its log file to flush it. (The value remains OFF so that you need not disable it explicitly before



enabling it again to perform another flush.) Enabling this variable has no effect unless `audit_log_rotate_on_size` [826] is 0.

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

- `audit_log_policy` [826]

<b>Version Introduced</b>	5.6.10	
<b>Variable Name</b>	<code>audit_log_policy</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>ALL</code>
	<b>Valid Values</b>	<code>ALL</code> <code>NONE</code> <code>LOGINS</code> <code>QUERIES</code>

The policy controlling the information written by the audit log plugin to its log file. The following table shows the permitted values.

Value	Description
<code>ALL</code>	Log all events
<code>NONE</code>	Log nothing (disable the audit stream)
<code>LOGINS</code>	Log only login events
<code>QUERIES</code>	Log only query events

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

- `audit_log_rotate_on_size` [826]

<b>Version Introduced</b>	5.6.10	
<b>Variable Name</b>	<code>audit_log_rotate_on_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

If the `audit_log_rotate_on_size` [826] value is greater than 0, the audit log plugin closes and reopens its log file if a write to the file causes its size to exceed this value. The original file is renamed to have a timestamp extension.

If the `audit_log_rotate_on_size` [826] value is 0, the plugin does not close and reopen its log based on size. Instead, use `audit_log_flush` [825] to close and reopen the log on demand. In this case, rename the file externally to the server before flushing it.

For more information about audit log file rotation and timestamp interpretation, see [Section 6.3.11.3, “Audit Log Plugin Logging Control”](#).

If you set this variable to a value that is not a multiple of 4096, it is truncated to the nearest multiple.

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

- `audit_log_strategy` [827]

<b>Version Introduced</b>	5.6.10	
<b>Variable Name</b>	<code>audit_log_strategy</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	ASYNCHRONOUS
	<b>Valid Values</b>	ASYNCHRONOUS PERFORMANCE SEMISYNCHRONOUS SYNCHRONOUS

The logging method used by the audit log plugin. The following table describes the permitted values.

**Table 6.17. Audit Log Strategies**

Value	Meaning
ASYNCHRONOUS	Log asynchronously, wait for space in output buffer
PERFORMANCE	Log asynchronously, drop request if insufficient space in output buffer
SEMISYNCHRONOUS	Log synchronously, permit caching by operating system
SYNCHRONOUS	Log synchronously, call <code>sync()</code> after each request

This variable was added in MySQL 5.6.10. It is available only if the `audit_log` plugin is enabled.

### 6.3.11.5. Audit Log Plugin Restrictions

The audit log plugin is subject to these restrictions:

- Only top-level statements are logged, not statements within stored programs such as triggers or stored procedures.
- Contents of files referenced by statements such as `LOAD DATA INFILE` are not logged.

## 6.3.12. SQL-Based MySQL Account Activity Auditing

Applications can use the following guidelines to perform SQL-based auditing that ties database activity to MySQL accounts.

MySQL accounts correspond to rows in the `mysql.user` table. When a client connects successfully, the server authenticates the client to a particular row in this table. The `User` and `Host` column values in this row uniquely identify the account and correspond to the `'user_name'@'host_name'` format in which account names are written in SQL statements.

The account used to authenticate a client determines which privileges the client has. Normally, the `CURRENT_USER()` [1267] function can be invoked to determine which account this is for the client user. Its value is constructed from the `User` and `Host` columns of the `user` table row for the account.

However, there are circumstances under which the `CURRENT_USER()` [1267] value corresponds not to the client user but to a different account. This occurs in contexts when privilege checking is not based the client's account:

- Stored routines (procedures and functions) defined with the `SQL SECURITY DEFINER` characteristic
- Views defined with the `SQL SECURITY DEFINER` characteristic
- Triggers and events

In those contexts, privilege checking is done against the `DEFINER` account and `CURRENT_USER()` [1267] refers to that account, not to the account for the client who invoked the stored routine or view or who caused the trigger to activate. To determine the invoking user, you can call the `USER()` [1273] function, which returns a value indicating the actual user name provided by the client and the host from which the client connected. However, this value does not necessarily correspond directly to an account in the `user` table, because the `USER()` [1273] value never contains wildcards, whereas account values (as returned by `CURRENT_USER()` [1267]) may contain user name and host name wildcards.

For example, a blank user name matches any user, so an account of `'@localhost'` enables clients to connect as an anonymous user from the local host with any user name. If this case, if a client connects as `user1` from the local host, `USER()` [1273] and `CURRENT_USER()` [1267] return different values:

```
mysql> SELECT USER(), CURRENT_USER();
+-----+-----+
| USER()          | CURRENT_USER() |
+-----+-----+
| user1@localhost | @localhost     |
+-----+-----+
```

The host name part of an account can contain wildcards, too. If the host name contains a `'%'` or `'_'` pattern character or uses netmask notation, the account can be used for clients connecting from multiple hosts and the `CURRENT_USER()` [1267] value will not indicate which one. For example, the account `'user2'@'%.example.com'` can be used by `user2` to connect from any host in the `example.com` domain. If `user2` connects from `remote.example.com`, `USER()` [1273] and `CURRENT_USER()` [1267] return different values:

```
mysql> SELECT USER(), CURRENT_USER();
+-----+-----+
| USER()          | CURRENT_USER() |
+-----+-----+
| user2@remote.example.com | user2@%.example.com |
+-----+-----+
```

If an application must invoke `USER()` [1273] for user auditing (for example, if it does auditing from within triggers) but must also be able to associate the `USER()` [1273] value with an account in the `user` table, it is necessary to avoid accounts that contain wildcards in the `User` or `Host` column. Specifically, do not permit `User` to be empty (which creates an anonymous-user account), and do not permit pattern characters or netmask notation in `Host` values. All accounts must have a nonempty `User` value and literal `Host` value.

With respect to the previous examples, the `'@localhost'` and `'user2'@'%.example.com'` accounts should be changed not to use wildcards:

```
RENAME USER '@localhost' TO 'user1@localhost';
RENAME USER 'user2'@'%.example.com' TO 'user2'@'remote.example.com';
```

If `user2` must be able to connect from several hosts in the `example.com` domain, there should be a separate account for each host.

To extract the user name or host name part from a `CURRENT_USER()` [1267] or `USER()` [1273] value, use the `SUBSTRING()` [1185] function:

```
mysql> SELECT SUBSTRING_INDEX(CURRENT_USER(), '@', 1);
+-----+
| SUBSTRING_INDEX(CURRENT_USER(), '@', 1) |
+-----+
| user1                                     |
+-----+

mysql> SELECT SUBSTRING_INDEX(CURRENT_USER(), '@', -1);
+-----+
| SUBSTRING_INDEX(CURRENT_USER(), '@', -1) |
+-----+
| localhost                                 |
+-----+
```



---

# Chapter 7. Backup and Recovery

## Table of Contents

7.1. Backup and Recovery Types .....	832
7.2. Database Backup Methods .....	835
7.3. Example Backup and Recovery Strategy .....	837
7.3.1. Establishing a Backup Policy .....	838
7.3.2. Using Backups for Recovery .....	839
7.3.3. Backup Strategy Summary .....	840
7.4. Using <code>mysqldump</code> for Backups .....	840
7.4.1. Dumping Data in SQL Format with <code>mysqldump</code> .....	841
7.4.2. Reloading SQL-Format Backups .....	841
7.4.3. Dumping Data in Delimited-Text Format with <code>mysqldump</code> .....	842
7.4.4. Reloading Delimited-Text Format Backups .....	843
7.4.5. <code>mysqldump</code> Tips .....	844
7.5. Point-in-Time (Incremental) Recovery Using the Binary Log .....	845
7.5.1. Point-in-Time Recovery Using Event Times .....	847
7.5.2. Point-in-Time Recovery Using Event Positions .....	847
7.6. <code>MyISAM</code> Table Maintenance and Crash Recovery .....	848
7.6.1. Using <code>myisamchk</code> for Crash Recovery .....	848
7.6.2. How to Check <code>MyISAM</code> Tables for Errors .....	849
7.6.3. How to Repair <code>MyISAM</code> Tables .....	850
7.6.4. <code>MyISAM</code> Table Optimization .....	852
7.6.5. Setting Up a <code>MyISAM</code> Table Maintenance Schedule .....	853

It is important to back up your databases so that you can recover your data and be up and running again in case problems occur, such as system crashes, hardware failures, or users deleting data by mistake. Backups are also essential as a safeguard before upgrading a MySQL installation, and they can be used to transfer a MySQL installation to another system or to set up replication slave servers.

MySQL offers a variety of backup strategies from which you can choose the methods that best suit the requirements for your installation. This chapter discusses several backup and recovery topics with which you should be familiar:

- Types of backups: Logical versus physical, full versus incremental, and so forth.
- Methods for creating backups.
- Recovery methods, including point-in-time recovery.
- Backup scheduling, compression, and encryption.
- Table maintenance, to enable recovery of corrupt tables.

## Additional Resources

Resources related to backup or to maintaining data availability include the following:

- Customers of MySQL Enterprise Edition can use the MySQL Enterprise Backup product for backups. For an overview of the MySQL Enterprise Backup product, see [Section 23.2, “MySQL Enterprise Backup”](#).
- A forum dedicated to backup issues is available at <http://forums.mysql.com/list.php?93>.
- Details for `mysqldump`, `mysqlhotcopy`, and other MySQL backup programs can be found in [Chapter 4, MySQL Programs](#).
- The syntax of the SQL statements described here is given in [Chapter 13, SQL Statement Syntax](#).

- For additional information about [InnoDB](#) backup procedures, see [Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”](#).
- Replication enables you to maintain identical data on multiple servers. This has several benefits, such as enabling client query load to be distributed over servers, availability of data even if a given server is taken offline or fails, and the ability to make backups with no impact on the master by using a slave server. See [Chapter 16, Replication](#).
- MySQL Cluster provides a high-availability, high-redundancy version of MySQL adapted for the distributed computing environment. See [MySQL Cluster NDB 6.1 - 7.1](#), which provides information about MySQL Cluster NDB 7.0 and 7.1 (based on MySQL 5.1 but containing the latest improvements and fixes for the [NDBCLUSTER](#) storage engine).

**Note**

The [NDBCLUSTER](#) storage engine is currently not supported in MySQL 5.6.

- Distributed Replicated Block Device (DRBD) is another high-availability solution. It works by replicating a block device from a primary server to a secondary server at the block level. See [Chapter 15, High Availability and Scalability](#)

## 7.1. Backup and Recovery Types

This section describes the characteristics of different types of backups.

### Physical (Raw) Versus Logical Backups

Physical backups consist of raw copies of the directories and files that store database contents. This type of backup is suitable for large, important databases that need to be recovered quickly when problems occur.

Logical backups save information represented as logical database structure ([CREATE DATABASE](#), [CREATE TABLE](#) statements) and content ([INSERT](#) statements or delimited-text files). This type of backup is suitable for smaller amounts of data where you might edit the data values or table structure, or recreate the data on a different machine architecture.

Physical backup methods have these characteristics:

- The backup consists of exact copies of database directories and files. Typically this is a copy of all or part of the MySQL data directory.
- Physical backup methods are faster than logical because they involve only file copying without conversion.
- Output is more compact than for logical backup.
- Because backup speed and compactness are important for busy, important databases, the MySQL Enterprise Backup product performs physical backups. For an overview of the MySQL Enterprise Backup product, see [Section 23.2, “MySQL Enterprise Backup”](#).
- Backup and restore granularity ranges from the level of the entire data directory down to the level of individual files. This may or may not provide for table-level granularity, depending on storage engine. For example, [InnoDB](#) tables can each be in a separate file, or share file storage with other [InnoDB](#) tables; each [MyISAM](#) table corresponds uniquely to a set of files.
- In addition to databases, the backup can include any related files such as log or configuration files.
- Data from [MEMORY](#) tables is tricky to back up this way because their contents are not stored on disk. (The MySQL Enterprise Backup product has a feature where you can retrieve data from [MEMORY](#) tables during a backup.)

- Backups are portable only to other machines that have identical or similar hardware characteristics.
- Backups can be performed while the MySQL server is not running. If the server is running, it is necessary to perform appropriate locking so that the server does not change database contents during the backup. MySQL Enterprise Backup does this locking automatically for tables that require it.
- Physical backup tools include the `mysqlbackup` of MySQL Enterprise Backup for InnoDB or any other tables, file system-level commands (such as `cp`, `scp`, `tar`, `rsync`), or `mysqlhotcopy` for MyISAM tables.
- For restore:
  - MySQL Enterprise Backup restores InnoDB and other tables that it backed up.
  - `ndb_restore` restores NDB tables.
  - Files copied at the file system level or with `mysqlhotcopy` can be copied back to their original locations with file system commands.

Logical backup methods have these characteristics:

- The backup is done by querying the MySQL server to obtain database structure and content information.
- Backup is slower than physical methods because the server must access database information and convert it to logical format. If the output is written on the client side, the server must also send it to the backup program.
- Output is larger than for physical backup, particularly when saved in text format.
- Backup and restore granularity is available at the server level (all databases), database level (all tables in a particular database), or table level. This is true regardless of storage engine.
- The backup does not include log or configuration files, or other database-related files that are not part of databases.
- Backups stored in logical format are machine independent and highly portable.
- Logical backups are performed with the MySQL server running. The server is not taken offline.
- Logical backup tools include the `mysqldump` program and the `SELECT ... INTO OUTFILE` statement. These work for any storage engine, even MEMORY.
- To restore logical backups, SQL-format dump files can be processed using the `mysql` client. To load delimited-text files, use the `LOAD DATA INFILE` statement or the `mysqlimport` client.

## Online Versus Offline Backups

Online backups take place while the MySQL server is running so that the database information can be obtained from the server. Offline backups take place while the server is stopped. This distinction can also be described as “hot” versus “cold” backups; a “warm” backup is one where the server remains running but locked against modifying data while you access database files externally.

Online backup methods have these characteristics:

- The backup is less intrusive to other clients, which can connect to the MySQL server during the backup and may be able to access data depending on what operations they need to perform.
- Care must be taken to impose appropriate locking so that data modifications do not take place that would compromise backup integrity. The MySQL Enterprise Backup product does such locking automatically.



Offline backup methods have these characteristics:

- Clients can be affected adversely because the server is unavailable during backup. For that reason, such backups are often taken from a replication slave server that can be taken offline without harming availability.
- The backup procedure is simpler because there is no possibility of interference from client activity.

A similar distinction between online and offline applies for recovery operations, and similar characteristics apply. However, it is more likely that clients will be affected for online recovery than for online backup because recovery requires stronger locking. During backup, clients might be able to read data while it is being backed up. Recovery modifies data and does not just read it, so clients must be prevented from accessing data while it is being restored.

## Local Versus Remote Backups

A local backup is performed on the same host where the MySQL server runs, whereas a remote backup is done from a different host. For some types of backups, the backup can be initiated from a remote host even if the output is written locally on the server. host.

- `mysqldump` can connect to local or remote servers. For SQL output (`CREATE` and `INSERT` statements), local or remote dumps can be done and generate output on the client. For delimited-text output (with the `--tab` [291] option), data files are created on the server host.
- `mysqlhotcopy` performs only local backups: It connects to the server to lock it against data modifications and then copies local table files.
- `SELECT ... INTO OUTFILE` can be initiated from a local or remote client host, but the output file is created on the server host.
- Physical backup methods typically are initiated locally on the MySQL server host so that the server can be taken offline, although the destination for copied files might be remote.

## Snapshot Backups

Some file system implementations enable “snapshots” to be taken. These provide logical copies of the file system at a given point in time, without requiring a physical copy of the entire file system. (For example, the implementation may use copy-on-write techniques so that only parts of the file system modified after the snapshot time need be copied.) MySQL itself does not provide the capability for taking file system snapshots. It is available through third-party solutions such as Veritas, LVM, or ZFS.

## Full Versus Incremental Backups

A full backup includes all data managed by a MySQL server at a given point in time. An incremental backup consists of the changes made to the data during a given time span (from one point in time to another). MySQL has different ways to perform full backups, such as those described earlier in this section. Incremental backups are made possible by enabling the server's binary log, which the server uses to record data changes.

## Full Versus Point-in-Time (Incremental) Recovery

A full recovery restores all data from a full backup. This restores the server instance to the state that it had when the backup was made. If that state is not sufficiently current, a full recovery can be followed by recovery of incremental backups made since the full backup, to bring the server to a more up-to-date state.

Incremental recovery is recovery of changes made during a given time span. This is also called point-in-time recovery because it makes a server's state current up to a given time. Point-in-time recovery is based on the binary log and typically follows a full recovery from the backup files that restores the server to its state when the backup was made. Then the data changes written in the binary log files are

applied as incremental recovery to redo data modifications and bring the server up to the desired point in time.

## Table Maintenance

Data integrity can be compromised if tables become corrupt. For [InnoDB](#) tables, this is not a typical issue. For programs to check [MyISAM](#) tables and repair them if problems are found, see [Section 7.6](#), “[MyISAM Table Maintenance and Crash Recovery](#)”.

## Backup Scheduling, Compression, and Encryption

Backup scheduling is valuable for automating backup procedures. Compression of backup output reduces space requirements, and encryption of the output provides better security against unauthorized access of backed-up data. MySQL itself does not provide these capabilities. The MySQL Enterprise Backup product can compress [InnoDB](#) backups, and compression or encryption of backup output can be achieved using file system utilities. Other third-party solutions may be available.

## 7.2. Database Backup Methods

This section summarizes some general methods for making backups.

### Making a Hot Backup with MySQL Enterprise Backup

Customers of MySQL Enterprise Edition can use the [MySQL Enterprise Backup](#) product to do [physical](#) backups of entire instances or selected databases, tables, or both. This product includes features for [incremental](#) and [compressed](#) backups. Backing up the physical database files makes restore much faster than logical techniques such as the `mysqldump` command. [InnoDB](#) tables are copied using a [hot backup](#) mechanism. (Ideally, the [InnoDB](#) tables should represent a substantial majority of the data.) Tables from other storage engines are copied using a [warm backup](#) mechanism. For an overview of the MySQL Enterprise Backup product, see [Section 23.2](#), “[MySQL Enterprise Backup](#)”.

### Making Backups with `mysqldump` or `mysqlhotcopy`

The `mysqldump` program and the `mysqlhotcopy` script can make backups. `mysqldump` is more general because it can back up all kinds of tables. `mysqlhotcopy` works only with some storage engines. (See [Section 7.4](#), “[Using `mysqldump` for Backups](#)”, and [Section 4.6.10](#), “[`mysqlhotcopy` — A Database Backup Program](#)”.)

For [InnoDB](#) tables, it is possible to perform an online backup that takes no locks on tables using the `--single-transaction` [\[295\]](#) option to `mysqldump`. See [Section 7.3.1](#), “[Establishing a Backup Policy](#)”.

### Making Backups by Copying Table Files

For storage engines that represent each table using its own files, tables can be backed up by copying those files. For example, [MyISAM](#) tables are stored as files, so it is easy to do a backup by copying files (`*.frm`, `*.MYD`, and `*.MYI` files). To get a consistent backup, stop the server or lock and flush the relevant tables:

```
FLUSH TABLES tbl_list WITH READ LOCK;
```

You need only a read lock; this enables other clients to continue to query the tables while you are making a copy of the files in the database directory. The flush is needed to ensure that all active index pages are written to disk before you start the backup. See [Section 13.3.5](#), “[LOCK TABLES and UNLOCK TABLES Syntax](#)”, and [Section 13.7.6.3](#), “[FLUSH Syntax](#)”.

You can also create a binary backup simply by copying all table files, as long as the server isn't updating anything. The `mysqlhotcopy` script uses this method. (But note that table file copying methods do not work if your database contains [InnoDB](#) tables. `mysqlhotcopy` does not work for

InnoDB tables because InnoDB does not necessarily store table contents in database directories. Also, even if the server is not actively updating data, InnoDB may still have modified data cached in memory and not flushed to disk.)

## Making Delimited-Text File Backups

To create a text file containing a table's data, you can use `SELECT * INTO OUTFILE 'file_name' FROM tbl_name`. The file is created on the MySQL server host, not the client host. For this statement, the output file cannot already exist because permitting files to be overwritten constitutes a security risk. See [Section 13.2.9, "SELECT Syntax"](#). This method works for any kind of data file, but saves only table data, not the table structure.

Another way to create text data files (along with files containing `CREATE TABLE` statements for the backed up tables) is to use `mysqldump` with the `--tab` [\[291\]](#) option. See [Section 7.4.3, "Dumping Data in Delimited-Text Format with mysqldump"](#).

To reload a delimited-text data file, use `LOAD DATA INFILE` or `mysqlimport`.

## Making Incremental Backups by Enabling the Binary Log

MySQL supports incremental backups: You must start the server with the `--log-bin` [\[2002\]](#) option to enable binary logging; see [Section 5.2.4, "The Binary Log"](#). The binary log files provide you with the information you need to replicate changes to the database that are made subsequent to the point at which you performed a backup. At the moment you want to make an incremental backup (containing all changes that happened since the last full or incremental backup), you should rotate the binary log by using `FLUSH LOGS`. This done, you need to copy to the backup location all binary logs which range from the one of the moment of the last full or incremental backup to the last but one. These binary logs are the incremental backup; at restore time, you apply them as explained in [Section 7.5, "Point-in-Time \(Incremental\) Recovery Using the Binary Log"](#). The next time you do a full backup, you should also rotate the binary log using `FLUSH LOGS`, `mysqldump --flush-logs`, or `mysqlhotcopy --flushlog`. See [Section 4.5.4, "mysqldump — A Database Backup Program"](#), and [Section 4.6.10, "mysqlhotcopy — A Database Backup Program"](#).

## Making Backups Using Replication Slaves

If you have performance problems with your master server while making backups, one strategy that can help is to set up replication and perform backups on the slave rather than on the master. See [Section 16.3.1, "Using Replication for Backups"](#).

If you are backing up a slave replication server, you should back up its `master.info` and `relay-log.info` files when you back up the slave's databases, regardless of the backup method you choose. These information files are always needed to resume replication after you restore the slave's data. If your slave is replicating `LOAD DATA INFILE` statements, you should also back up any `SQL_LOAD-*` files that exist in the directory that the slave uses for this purpose. The slave needs these files to resume replication of any interrupted `LOAD DATA INFILE` operations. The location of this directory is the value of the `--slave-load-tmpdir` [\[1985\]](#) option. If the server was not started with that option, the directory location is the value of the `tmpdir` [\[559\]](#) system variable.

## Recovering Corrupt Tables

If you have to restore MyISAM tables that have become corrupt, try to recover them using `REPAIR TABLE` or `myisamchk -r` first. That should work in 99.9% of all cases. If `myisamchk` fails, see [Section 7.6, "MyISAM Table Maintenance and Crash Recovery"](#).

## Making Backups Using a File System Snapshot

If you are using a Veritas file system, you can make a backup like this:

1. From a client program, execute `FLUSH TABLES WITH READ LOCK`.

2. From another shell, execute `mount vxfs snapshot`.
3. From the first client, execute `UNLOCK TABLES`.
4. Copy files from the snapshot.
5. Unmount the snapshot.

Similar snapshot capabilities may be available in other file systems, such as LVM or ZFS.

## 7.3. Example Backup and Recovery Strategy

This section discusses a procedure for performing backups that enables you to recover data after several types of crashes:

- Operating system crash
- Power failure
- File system crash
- Hardware problem (hard drive, motherboard, and so forth)

The example commands do not include options such as `--user [212]` and `--password [211]` for the `mysqldump` and `mysql` client programs. You should include such options as necessary to enable client programs to connect to the MySQL server.

Assume that data is stored in the `InnoDB` storage engine, which has support for transactions and automatic crash recovery. Assume also that the MySQL server is under load at the time of the crash. If it were not, no recovery would ever be needed.

For cases of operating system crashes or power failures, we can assume that MySQL's disk data is available after a restart. The `InnoDB` data files might not contain consistent data due to the crash, but `InnoDB` reads its logs and finds in them the list of pending committed and noncommitted transactions that have not been flushed to the data files. `InnoDB` automatically rolls back those transactions that were not committed, and flushes to its data files those that were committed. Information about this recovery process is conveyed to the user through the MySQL error log. The following is an example log excerpt:

```
InnoDB: Database was not shut down normally.
InnoDB: Starting recovery from log files...
InnoDB: Starting log scan based on checkpoint at
InnoDB: log sequence number 0 13674004
InnoDB: Doing recovery: scanned up to log sequence number 0 13739520
InnoDB: Doing recovery: scanned up to log sequence number 0 13805056
InnoDB: Doing recovery: scanned up to log sequence number 0 13870592
InnoDB: Doing recovery: scanned up to log sequence number 0 13936128
...
InnoDB: Doing recovery: scanned up to log sequence number 0 20555264
InnoDB: Doing recovery: scanned up to log sequence number 0 20620800
InnoDB: Doing recovery: scanned up to log sequence number 0 20664692
InnoDB: 1 uncommitted transaction(s) which must be rolled back
InnoDB: Starting rollback of uncommitted transactions
InnoDB: Rolling back trx no 16745
InnoDB: Rolling back of trx no 16745 completed
InnoDB: Rollback of uncommitted transactions completed
InnoDB: Starting an apply batch of log records to the database...
InnoDB: Apply batch completed
InnoDB: Started
mysqld: ready for connections
```

For the cases of file system crashes or hardware problems, we can assume that the MySQL disk data is *not* available after a restart. This means that MySQL fails to start successfully because some blocks of disk data are no longer readable. In this case, it is necessary to reformat the disk, install a new one,

or otherwise correct the underlying problem. Then it is necessary to recover our MySQL data from backups, which means that backups must already have been made. To make sure that is the case, design and implement a backup policy.

### 7.3.1. Establishing a Backup Policy

To be useful, backups must be scheduled regularly. A full backup (a snapshot of the data at a point in time) can be done in MySQL with several tools. For example, [MySQL Enterprise Backup](#) can perform a [physical backup](#) of an entire instance, with optimizations to minimize overhead and avoid disruption when backing up [InnoDB](#) data files; `mysqldump` provides online [logical backup](#). This discussion uses `mysqldump`.

Assume that we make a full backup of all our [InnoDB](#) tables in all databases using the following command on Sunday at 1 p.m., when load is low:

```
shell> mysqldump --single-transaction --all-databases > backup_sunday_1_PM.sql
```

The resulting `.sql` file produced by `mysqldump` contains a set of SQL `INSERT` statements that can be used to reload the dumped tables at a later time.

This backup operation acquires a global read lock on all tables at the beginning of the dump (using `FLUSH TABLES WITH READ LOCK`). As soon as this lock has been acquired, the binary log coordinates are read and the lock is released. If long updating statements are running when the `FLUSH` statement is issued, the backup operation may stall until those statements finish. After that, the dump becomes lock-free and does not disturb reads and writes on the tables.

It was assumed earlier that the tables to back up are [InnoDB](#) tables, so `--single-transaction` [295] uses a consistent read and guarantees that data seen by `mysqldump` does not change. (Changes made by other clients to [InnoDB](#) tables are not seen by the `mysqldump` process.) If the backup operation includes nontransactional tables, consistency requires that they do not change during the backup. For example, for the [MyISAM](#) tables in the `mysql` database, there must be no administrative changes to MySQL accounts during the backup.

Full backups are necessary, but it is not always convenient to create them. They produce large backup files and take time to generate. They are not optimal in the sense that each successive full backup includes all data, even that part that has not changed since the previous full backup. It is more efficient to make an initial full backup, and then to make incremental backups. The incremental backups are smaller and take less time to produce. The tradeoff is that, at recovery time, you cannot restore your data just by reloading the full backup. You must also process the incremental backups to recover the incremental changes.

To make incremental backups, we need to save the incremental changes. In MySQL, these changes are represented in the binary log, so the MySQL server should always be started with the `--log-bin` [2002] option to enable that log. With binary logging enabled, the server writes each data change into a file while it updates data. Looking at the data directory of a MySQL server that was started with the `--log-bin` [2002] option and that has been running for some days, we find these MySQL binary log files:

```
-rw-rw---- 1 guilhem guilhem 1277324 Nov 10 23:59 gbichot2-bin.000001
-rw-rw---- 1 guilhem guilhem 4 Nov 10 23:59 gbichot2-bin.000002
-rw-rw---- 1 guilhem guilhem 79 Nov 11 11:06 gbichot2-bin.000003
-rw-rw---- 1 guilhem guilhem 508 Nov 11 11:08 gbichot2-bin.000004
-rw-rw---- 1 guilhem guilhem 220047446 Nov 12 16:47 gbichot2-bin.000005
-rw-rw---- 1 guilhem guilhem 998412 Nov 14 10:08 gbichot2-bin.000006
-rw-rw---- 1 guilhem guilhem 361 Nov 14 10:07 gbichot2-bin.index
```

Each time it restarts, the MySQL server creates a new binary log file using the next number in the sequence. While the server is running, you can also tell it to close the current binary log file and begin a new one manually by issuing a `FLUSH LOGS` SQL statement or with a `mysqladmin flush-logs` command. `mysqldump` also has an option to flush the logs. The `.index` file in the data directory contains the list of all MySQL binary logs in the directory.

The MySQL binary logs are important for recovery because they form the set of incremental backups. If you make sure to flush the logs when you make your full backup, the binary log files created afterward contain all the data changes made since the backup. Let's modify the previous `mysqldump` command a bit so that it flushes the MySQL binary logs at the moment of the full backup, and so that the dump file contains the name of the new current binary log:

```
shell> mysqldump --single-transaction --flush-logs --master-data=2 \
--all-databases > backup_sunday_1_PM.sql
```

After executing this command, the data directory contains a new binary log file, `gbichot2-bin.000007`, because the `--flush-logs` [294] option causes the server to flush its logs. The `--master-data` [288] option causes `mysqldump` to write binary log information to its output, so the resulting `.sql` dump file includes these lines:

```
-- Position to start replication or point-in-time recovery from
-- CHANGE MASTER TO MASTER_LOG_FILE='gbichot2-bin.000007',MASTER_LOG_POS=4;
```

Because the `mysqldump` command made a full backup, those lines mean two things:

- The dump file contains all changes made before any changes written to the `gbichot2-bin.000007` binary log file or newer.
- All data changes logged after the backup are not present in the dump file, but are present in the `gbichot2-bin.000007` binary log file or newer.

On Monday at 1 p.m., we can create an incremental backup by flushing the logs to begin a new binary log file. For example, executing a `mysqladmin flush-logs` command creates `gbichot2-bin.000008`. All changes between the Sunday 1 p.m. full backup and Monday 1 p.m. will be in the `gbichot2-bin.000007` file. This incremental backup is important, so it is a good idea to copy it to a safe place. (For example, back it up on tape or DVD, or copy it to another machine.) On Tuesday at 1 p.m., execute another `mysqladmin flush-logs` command. All changes between Monday 1 p.m. and Tuesday 1 p.m. will be in the `gbichot2-bin.000008` file (which also should be copied somewhere safe).

The MySQL binary logs take up disk space. To free up space, purge them from time to time. One way to do this is by deleting the binary logs that are no longer needed, such as when we make a full backup:

```
shell> mysqldump --single-transaction --flush-logs --master-data=2 \
--all-databases --delete-master-logs > backup_sunday_1_PM.sql
```



#### Note

Deleting the MySQL binary logs with `mysqldump --delete-master-logs` can be dangerous if your server is a replication master server, because slave servers might not yet fully have processed the contents of the binary log. The description for the `PURGE BINARY LOGS` statement explains what should be verified before deleting the MySQL binary logs. See [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#).

## 7.3.2. Using Backups for Recovery

Now, suppose that we have a catastrophic crash on Wednesday at 8 a.m. that requires recovery from backups. To recover, first we restore the last full backup we have (the one from Sunday 1 p.m.). The full backup file is just a set of SQL statements, so restoring it is very easy:

```
shell> mysql < backup_sunday_1_PM.sql
```

At this point, the data is restored to its state as of Sunday 1 p.m.. To restore the changes made since then, we must use the incremental backups; that is, the `gbichot2-bin.000007` and `gbichot2-bin.000008` binary log files. Fetch the files if necessary from where they were backed up, and then process their contents like this:



```
shell> mysqlbinlog gbichot2-bin.000007 gbichot2-bin.000008 | mysql
```

We now have recovered the data to its state as of Tuesday 1 p.m., but still are missing the changes from that date to the date of the crash. To not lose them, we would have needed to have the MySQL server store its MySQL binary logs into a safe location (RAID disks, SAN, ...) different from the place where it stores its data files, so that these logs were not on the destroyed disk. (That is, we can start the server with a `--log-bin` [2002] option that specifies a location on a different physical device from the one on which the data directory resides. That way, the logs are safe even if the device containing the directory is lost.) If we had done this, we would have the `gbichot2-bin.000009` file (and any subsequent files) at hand, and we could apply them using `mysqlbinlog` and `mysql` to restore the most recent data changes with no loss up to the moment of the crash:

```
shell> mysqlbinlog gbichot2-bin.000009 ... | mysql
```

For more information about using `mysqlbinlog` to process binary log files, see [Section 7.5, “Point-in-Time \(Incremental\) Recovery Using the Binary Log”](#).

### 7.3.3. Backup Strategy Summary

In case of an operating system crash or power failure, `InnoDB` itself does all the job of recovering data. But to make sure that you can sleep well, observe the following guidelines:

- Always run the MySQL server with the `--log-bin` [2002] option, or even `--log-bin=log_name` [2002], where the log file name is located on some safe media different from the drive on which the data directory is located. If you have such safe media, this technique can also be good for disk load balancing (which results in a performance improvement).
- Make periodic full backups, using the `mysqldump` command shown earlier in [Section 7.3.1, “Establishing a Backup Policy”](#), that makes an online, nonblocking backup.
- Make periodic incremental backups by flushing the logs with `FLUSH LOGS` or `mysqladmin flush-logs`.

## 7.4. Using `mysqldump` for Backups

This section describes how to use `mysqldump` to produce dump files, and how to reload dump files. A dump file can be used in several ways:

- As a backup to enable data recovery in case of data loss.
- As a source of data for setting up replication slaves.
- As a source of data for experimentation:
  - To make a copy of a database that you can use without changing the original data.
  - To test potential upgrade incompatibilities.

`mysqldump` produces two types of output, depending on whether the `--tab` [291] option is given:

- Without `--tab` [291], `mysqldump` writes SQL statements to the standard output. This output consists of `CREATE` statements to create dumped objects (databases, tables, stored routines, and so forth), and `INSERT` statements to load data into tables. The output can be saved in a file and reloaded later using `mysql` to recreate the dumped objects. Options are available to modify the format of the SQL statements, and to control which objects are dumped.
- With `--tab` [291], `mysqldump` produces two output files for each dumped table. The server writes one file as tab-delimited text, one line per table row. This file is named `tbl_name.txt` in the output directory. The server also sends a `CREATE TABLE` statement for the table to `mysqldump`, which writes it as a file named `tbl_name.sql` in the output directory.

### 7.4.1. Dumping Data in SQL Format with `mysqldump`

This section describes how to use `mysqldump` to create SQL-format dump files. For information about reloading such dump files, see [Section 7.4.2, “Reloading SQL-Format Backups”](#).

By default, `mysqldump` writes information as SQL statements to the standard output. You can save the output in a file:

```
shell> mysqldump [arguments] > file_name
```

To dump all databases, invoke `mysqldump` with the `--all-databases` [\[292\]](#) option:

```
shell> mysqldump --all-databases > dump.sql
```

To dump only specific databases, name them on the command line and use the `--databases` [\[292\]](#) option:

```
shell> mysqldump --databases db1 db2 db3 > dump.sql
```

The `--databases` [\[292\]](#) option causes all names on the command line to be treated as database names. Without this option, `mysqldump` treats the first name as a database name and those following as table names.

With `--all-databases` [\[292\]](#) or `--databases` [\[292\]](#), `mysqldump` writes `CREATE DATABASE` and `USE` statements prior to the dump output for each database. This ensures that when the dump file is reloaded, it creates each database if it does not exist and makes it the default database so database contents are loaded into the same database from which they came. If you want to cause the dump file to force a drop of each database before recreating it, use the `--add-drop-database` [\[286\]](#) option as well. In this case, `mysqldump` writes a `DROP DATABASE` statement preceding each `CREATE DATABASE` statement.

To dump a single database, name it on the command line:

```
shell> mysqldump --databases test > dump.sql
```

In the single-database case, it is permissible to omit the `--databases` [\[292\]](#) option:

```
shell> mysqldump test > dump.sql
```

The difference between the two preceding commands is that without `--databases` [\[292\]](#), the dump output contains no `CREATE DATABASE` or `USE` statements. This has several implications:

- When you reload the dump file, you must specify a default database name so that the server knows which database to reload.
- For reloading, you can specify a database name different from the original name, which enables you to reload the data into a different database.
- If the database to be reloaded does not exist, you must create it first.
- Because the output will contain no `CREATE DATABASE` statement, the `--add-drop-database` [\[286\]](#) option has no effect. If you use it, it produces no `DROP DATABASE` statement.

To dump only specific tables from a database, name them on the command line following the database name:

```
shell> mysqldump test t1 t3 t7 > dump.sql
```

### 7.4.2. Reloading SQL-Format Backups

To reload a dump file written by `mysqldump` that consists of SQL statements, use it as input to the `mysql` client. If the dump file was created by `mysqldump` with the `--all-databases` [\[292\]](#) or `--`



`databases` [292] option, it contains `CREATE DATABASE` and `USE` statements and it is not necessary to specify a default database into which to load the data:

```
shell> mysql < dump.sql
```

Alternatively, from within `mysql`, use a `source` command:

```
mysql> source dump.sql
```

If the file is a single-database dump not containing `CREATE DATABASE` and `USE` statements, create the database first (if necessary):

```
shell> mysqladmin create db1
```

Then specify the database name when you load the dump file:

```
shell> mysql db1 < dump.sql
```

Alternatively, from within `mysql`, create the database, select it as the default database, and load the dump file:

```
mysql> CREATE DATABASE IF NOT EXISTS db1;
mysql> USE db1;
mysql> source dump.sql
```

### 7.4.3. Dumping Data in Delimited-Text Format with `mysqldump`

This section describes how to use `mysqldump` to create delimited-text dump files. For information about reloading such dump files, see [Section 7.4.4, “Reloading Delimited-Text Format Backups”](#).

If you invoke `mysqldump` with the `--tab=dir_name` [291] option, it uses `dir_name` as the output directory and dumps tables individually in that directory using two files for each table. The table name is the basename for these files. For a table named `t1`, the files are named `t1.sql` and `t1.txt`. The `.sql` file contains a `CREATE TABLE` statement for the table. The `.txt` file contains the table data, one line per table row.

The following command dumps the contents of the `db1` database to files in the `/tmp` database:

```
shell> mysqldump --tab=/tmp db1
```

The `.txt` files containing table data are written by the server, so they are owned by the system account used for running the server. The server uses `SELECT ... INTO OUTFILE` to write the files, so you must have the `FILE` [754] privilege to perform this operation, and an error occurs if a given `.txt` file already exists.

The server sends the `CREATE` definitions for dumped tables to `mysqldump`, which writes them to `.sql` files. These files therefore are owned by the user who executes `mysqldump`.

It is best that `--tab` [291] be used only for dumping a local server. If you use it with a remote server, the `--tab` [291] directory must exist on both the local and remote hosts, and the `.txt` files will be written by the server in the remote directory (on the server host), whereas the `.sql` files will be written by `mysqldump` in the local directory (on the client host).

For `mysqldump --tab`, the server by default writes table data to `.txt` files one line per row with tabs between column values, no quotation marks around column values, and newline as the line terminator. (These are the same defaults as for `SELECT ... INTO OUTFILE`.)

To enable data files to be written using a different format, `mysqldump` supports these options:

- `--fields-terminated-by=str` [290]

The string for separating column values (default: tab).

- `--fields-enclosed-by=char` [290]

The character within which to enclose column values (default: no character).

- `--fields-optionally-enclosed-by=char` [290]

The character within which to enclose non-numeric column values (default: no character).

- `--fields-escaped-by=char` [290]

The character for escaping special characters (default: no escaping).

- `--lines-terminated-by=str` [290]

The line-termination string (default: newline).

Depending on the value you specify for any of these options, it might be necessary on the command line to quote or escape the value appropriately for your command interpreter. Alternatively, specify the value using hex notation. Suppose that you want `mysqldump` to quote column values within double quotation marks. To do so, specify double quote as the value for the `--fields-enclosed-by` [290] option. But this character is often special to command interpreters and must be treated specially. For example, on Unix, you can quote the double quote like this:

```
--fields-enclosed-by='"'
```

On any platform, you can specify the value in hex:

```
--fields-enclosed-by=0x22
```

It is common to use several of the data-formatting options together. For example, to dump tables in comma-separated values format with lines terminated by carriage-return/newline pairs (`\r\n`), use this command (enter it on a single line):

```
shell> mysqldump --tab=/tmp --fields-terminated-by=,
--fields-enclosed-by='"' --lines-terminated-by=0x0d0a db1
```

Should you use any of the data-formatting options to dump table data, you will need to specify the same format when you reload data files later, to ensure proper interpretation of the file contents.

## 7.4.4. Reloading Delimited-Text Format Backups

For backups produced with `mysqldump --tab`, each table is represented in the output directory by an `.sql` file containing the `CREATE TABLE` statement for the table, and a `.txt` file containing the table data. To reload a table, first change location into the output directory. Then process the `.sql` file with `mysql` to create an empty table and process the `.txt` file to load the data into the table:

```
shell> mysql db1 < t1.sql
shell> mysqlimport db1 t1.txt
```

An alternative to using `mysqlimport` to load the data file is to use the `LOAD DATA INFILE` statement from within the `mysql` client:

```
mysql> USE db1;
mysql> LOAD DATA INFILE 't1.txt' INTO TABLE t1;
```

If you used any data-formatting options with `mysqldump` when you initially dumped the table, you must use the same options with `mysqlimport` or `LOAD DATA INFILE` to ensure proper interpretation of the data file contents:

```
shell> mysqlimport --fields-terminated-by=,
--fields-enclosed-by='"' --lines-terminated-by=0x0d0a db1 t1.txt
```

Or:

```
mysql> USE db1;
mysql> LOAD DATA INFILE 't1.txt' INTO TABLE t1
```

```
-> FIELDS TERMINATED BY ',' FIELDS ENCLOSED BY ''''
-> LINES TERMINATED BY '\r\n';
```

## 7.4.5. mysqldump Tips

This section surveys techniques that enable you to use `mysqldump` to solve specific problems:

- How to make a copy a database
- How to copy a database from one server to another
- How to dump stored programs (stored procedures and functions, triggers, and events)
- How to dump definitions and data separately

### 7.4.5.1. Making a Copy of a Database

```
shell> mysqldump db1 > dump.sql
shell> mysqladmin create db2
shell> mysql db2 < dump.sql
```

Do not use `--databases [292]` on the `mysqldump` command line because that causes `USE db1` to be included in the dump file, which overrides the effect of naming `db2` on the `mysql` command line.

### 7.4.5.2. Copy a Database from one Server to Another

On Server 1:

```
shell> mysqldump --databases db1 > dump.sql
```

Copy the dump file from Server 1 to Server 2.

On Server 2:

```
shell> mysql < dump.sql
```

Use of `--databases [292]` with the `mysqldump` command line causes the dump file to include `CREATE DATABASE` and `USE` statements that create the database if it does exist and make it the default database for the reloaded data.

Alternatively, you can omit `--databases [292]` from the `mysqldump` command. Then you will need to create the database on Server 2 (if necessary) and specify it as the default database when you reload the dump file.

On Server 1:

```
shell> mysqldump db1 > dump.sql
```

On Server 2:

```
shell> mysqladmin create db1
shell> mysql db1 < dump.sql
```

You can specify a different database name in this case, so omitting `--databases [292]` from the `mysqldump` command enables you to dump data from one database and load it into another.

### 7.4.5.3. Dumping Stored Programs

Several options control how `mysqldump` handles stored programs (stored procedures and functions, triggers, and events):

- `--events [293]`: Dump Event Scheduler events
- `--routines [293]`: Dump stored procedures and functions

- `--triggers` [293]: Dump triggers for tables

The `--triggers` [293] option is enabled by default so that when tables are dumped, they are accompanied by any triggers they have. The other options are disabled by default and must be specified explicitly to dump the corresponding objects. To disable any of these options explicitly, use its skip form: `--skip-events` [293], `--skip-routines` [293], or `--skip-triggers` [293].

#### 7.4.5.4. Dumping Table Definitions and Content Separately

The `--no-data` [293] option tells `mysqldump` not to dump table data, resulting in the dump file containing only statements to create the tables. Conversely, the `--no-create-info` [286] option tells `mysqldump` to suppress `CREATE` statements from the output, so that the dump file contains only table data.

For example, to dump table definitions and data separately for the `test` database, use these commands:

```
shell> mysqldump --no-data test > dump-defs.sql
shell> mysqldump --no-create-info test > dump-data.sql
```

For a definition-only dump, add the `--routines` [293] and `--events` [293] options to also include stored routine and event definitions:

```
shell> mysqldump --no-data --routines --events test > dump-defs.sql
```

#### 7.4.5.5. Using `mysqldump` to Test for Upgrade Incompatibilities

When contemplating a MySQL upgrade, it is prudent to install the newer version separately from your current production version. Then you can dump the database and database object definitions from the production server and load them into the new server to verify that they are handled properly. (This is also useful for testing downgrades.)

On the production server:

```
shell> mysqldump --all-databases --no-data --routines --events > dump-defs.sql
```

On the upgraded server:

```
shell> mysql < dump-defs.sql
```

Because the dump file does not contain table data, it can be processed quickly. This enables you to spot potential incompatibilities without waiting for lengthy data-loading operations. Look for warnings or errors while the dump file is being processed.

After you have verified that the definitions are handled properly, dump the data and try to load it into the upgraded server.

On the production server:

```
shell> mysqldump --all-databases --no-create-info > dump-data.sql
```

On the upgraded server:

```
shell> mysql < dump-data.sql
```

Now check the table contents and run some test queries.

## 7.5. Point-in-Time (Incremental) Recovery Using the Binary Log

Point-in-time recovery refers to recovery of data changes made since a given point in time. Typically, this type of recovery is performed after restoring a full backup that brings the server to its state as of the time the backup was made. (The full backup can be made in several ways, such as those listed

in [Section 7.2, “Database Backup Methods”](#).) Point-in-time recovery then brings the server up to date incrementally from the time of the full backup to a more recent time.

Point-in-time recovery is based on these principles:

- The source of information for point-in-time recovery is the set of incremental backups represented by the binary log files generated subsequent to the full backup operation. Therefore, the server must be started with the `--log-bin [2002]` option to enable binary logging (see [Section 5.2.4, “The Binary Log”](#)).

To restore data from the binary log, you must know the name and location of the current binary log files. By default, the server creates binary log files in the data directory, but a path name can be specified with the `--log-bin [2002]` option to place the files in a different location. [Section 5.2.4, “The Binary Log”](#).

To see a listing of all binary log files, use this statement:

```
mysql> SHOW BINARY LOGS;
```

To determine the name of the current binary log file, issue the following statement:

```
mysql> SHOW MASTER STATUS;
```

- The `mysqlbinlog` utility converts the events in the binary log files from binary format to text so that they can be executed or viewed. `mysqlbinlog` has options for selecting sections of the binary log based on event times or position of events within the log. See [Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#).
- Executing events from the binary log causes the data modifications they represent to be redone. This enables recovery of data changes for a given span of time. To execute events from the binary log, process `mysqlbinlog` output using the `mysql` client:

```
shell> mysqlbinlog binlog_files | mysql -u root -p
```

- Viewing log contents can be useful when you need to determine event times or positions to select partial log contents prior to executing events. To view events from the log, send `mysqlbinlog` output into a paging program:

```
shell> mysqlbinlog binlog_files | more
```

Alternatively, save the output in a file and view the file in a text editor:

```
shell> mysqlbinlog binlog_files > tmpfile
shell> ... edit tmpfile ...
```

- Saving the output in a file is useful as a preliminary to executing the log contents with certain events removed, such as an accidental `DROP DATABASE`. You can delete from the file any statements not to be executed before executing its contents. After editing the file, execute the contents as follows:

```
shell> mysql -u root -p < tmpfile
```

If you have more than one binary log to execute on the MySQL server, the safe method is to process them all using a single connection to the server. Here is an example that demonstrates what may be *unsafe*:

```
shell> mysqlbinlog binlog.000001 | mysql -u root -p # DANGER!!
shell> mysqlbinlog binlog.000002 | mysql -u root -p # DANGER!!
```

Processing binary logs this way using different connections to the server causes problems if the first log file contains a `CREATE TEMPORARY TABLE` statement and the second log contains a statement that uses the temporary table. When the first `mysql` process terminates, the server drops the temporary table. When the second `mysql` process attempts to use the table, the server reports “unknown table.”

To avoid problems like this, use a *single* connection to execute the contents of all binary logs that you want to process. Here is one way to do so:

```
shell> mysqlbinlog binlog.000001 binlog.000002 | mysql -u root -p
```

Another approach is to write all the logs to a single file and then process the file:

```
shell> mysqlbinlog binlog.000001 > /tmp/statements.sql
shell> mysqlbinlog binlog.000002 >> /tmp/statements.sql
shell> mysql -u root -p -e "source /tmp/statements.sql"
```

## 7.5.1. Point-in-Time Recovery Using Event Times

To indicate the start and end times for recovery, specify the `--start-datetime` [353] and `--stop-datetime` [353] options for `mysqlbinlog`, in `DATETIME` format. As an example, suppose that exactly at 10:00 a.m. on April 20, 2005 an SQL statement was executed that deleted a large table. To restore the table and data, you could restore the previous night's backup, and then execute the following command:

```
shell> mysqlbinlog --stop-datetime="2005-04-20 9:59:59" \
/var/log/mysql/bin.123456 | mysql -u root -p
```

This command recovers all of the data up until the date and time given by the `--stop-datetime` [353] option. If you did not detect the erroneous SQL statement that was entered until hours later, you will probably also want to recover the activity that occurred afterward. Based on this, you could run `mysqlbinlog` again with a start date and time, like so:

```
shell> mysqlbinlog --start-datetime="2005-04-20 10:01:00" \
/var/log/mysql/bin.123456 | mysql -u root -p
```

In this command, the SQL statements logged from 10:01 a.m. on will be re-executed. The combination of restoring of the previous night's dump file and the two `mysqlbinlog` commands restores everything up until one second before 10:00 a.m. and everything from 10:01 a.m. on.

To use this method of point-in-time recovery, you should examine the log to be sure of the exact times to specify for the commands. To display the log file contents without executing them, use this command:

```
shell> mysqlbinlog /var/log/mysql/bin.123456 > /tmp/mysql_restore.sql
```

Then open the `/tmp/mysql_restore.sql` file with a text editor to examine it.

Excluding specific changes by specifying times for `mysqlbinlog` does not work well if multiple statements executed at the same time as the one to be excluded.

## 7.5.2. Point-in-Time Recovery Using Event Positions

Instead of specifying dates and times, the `--start-position` [353] and `--stop-position` [354] options for `mysqlbinlog` can be used for specifying log positions. They work the same as the start and stop date options, except that you specify log position numbers rather than dates. Using positions may enable you to be more precise about which part of the log to recover, especially if many transactions occurred around the same time as a damaging SQL statement. To determine the position numbers, run `mysqlbinlog` for a range of times near the time when the unwanted transaction was executed, but redirect the results to a text file for examination. This can be done like so:

```
shell> mysqlbinlog --start-datetime="2005-04-20 9:55:00" \
--stop-datetime="2005-04-20 10:05:00" \
/var/log/mysql/bin.123456 > /tmp/mysql_restore.sql
```

This command creates a small text file in the `/tmp` directory that contains the SQL statements around the time that the deleterious SQL statement was executed. Open this file with a text editor and look for the statement that you do not want to repeat. Determine the positions in the binary log for stopping and resuming the recovery and make note of them. Positions are labeled as `log_pos` followed by a

number. After restoring the previous backup file, use the position numbers to process the binary log file. For example, you would use commands something like these:

```
shell> mysqlbinlog --stop-position=368312 /var/log/mysql/bin.123456 \  
| mysql -u root -p  
  
shell> mysqlbinlog --start-position=368315 /var/log/mysql/bin.123456 \  
| mysql -u root -p
```

The first command recovers all the transactions up until the stop position given. The second command recovers all transactions from the starting position given until the end of the binary log. Because the output of `mysqlbinlog` includes `SET TIMESTAMP` statements before each SQL statement recorded, the recovered data and related MySQL logs will reflect the original times at which the transactions were executed.

## 7.6. MyISAM Table Maintenance and Crash Recovery

This section discusses how to use `myisamchk` to check or repair MyISAM tables (tables that have `.MYD` and `.MYI` files for storing data and indexes). For general `myisamchk` background, see [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#). Other table-repair information can be found at [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).

You can use `myisamchk` to check, repair, or optimize database tables. The following sections describe how to perform these operations and how to set up a table maintenance schedule. For information about using `myisamchk` to get information about your tables, see [Section 4.6.3.5, “Obtaining Table Information with myisamchk”](#).

Even though table repair with `myisamchk` is quite secure, it is always a good idea to make a backup *before* doing a repair or any maintenance operation that could make a lot of changes to a table.

`myisamchk` operations that affect indexes can cause MyISAM FULLTEXT indexes to be rebuilt with full-text parameters that are incompatible with the values used by the MySQL server. To avoid this problem, follow the guidelines in [Section 4.6.3.1, “myisamchk General Options”](#).

MyISAM table maintenance can also be done using the SQL statements that perform operations similar to what `myisamchk` can do:

- To check MyISAM tables, use `CHECK TABLE`.
- To repair MyISAM tables, use `REPAIR TABLE`.
- To optimize MyISAM tables, use `OPTIMIZE TABLE`.
- To analyze MyISAM tables, use `ANALYZE TABLE`.

For additional information about these statements, see [Section 13.7.2, “Table Maintenance Statements”](#).

These statements can be used directly or by means of the `mysqlcheck` client program. One advantage of these statements over `myisamchk` is that the server does all the work. With `myisamchk`, you must make sure that the server does not use the tables at the same time so that there is no unwanted interaction between `myisamchk` and the server.

### 7.6.1. Using myisamchk for Crash Recovery

This section describes how to check for and deal with data corruption in MySQL databases. If your tables become corrupted frequently, you should try to find the reason why. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#).

For an explanation of how MyISAM tables can become corrupted, see [Section 14.3.4, “MyISAM Table Problems”](#).



If you run `mysqld` with external locking disabled (which is the default), you cannot reliably use `myisamchk` to check a table when `mysqld` is using the same table. If you can be certain that no one will access the tables through `mysqld` while you run `myisamchk`, you only have to execute `mysqladmin flush-tables` before you start checking the tables. If you cannot guarantee this, you must stop `mysqld` while you check the tables. If you run `myisamchk` to check tables that `mysqld` is updating at the same time, you may get a warning that a table is corrupt even when it is not.

If the server is run with external locking enabled, you can use `myisamchk` to check tables at any time. In this case, if the server tries to update a table that `myisamchk` is using, the server will wait for `myisamchk` to finish before it continues.

If you use `myisamchk` to repair or optimize tables, you *must* always ensure that the `mysqld` server is not using the table (this also applies if external locking is disabled). If you do not stop `mysqld`, you should at least do a `mysqladmin flush-tables` before you run `myisamchk`. Your tables *may become corrupted* if the server and `myisamchk` access the tables simultaneously.

When performing crash recovery, it is important to understand that each MyISAM table `tbl_name` in a database corresponds to the three files in the database directory shown in the following table.

File	Purpose
<code>tbl_name.frm</code>	Definition (format) file
<code>tbl_name.MYD</code>	Data file
<code>tbl_name.MYI</code>	Index file

Each of these three file types is subject to corruption in various ways, but problems occur most often in data files and index files.

`myisamchk` works by creating a copy of the `.MYD` data file row by row. It ends the repair stage by removing the old `.MYD` file and renaming the new file to the original file name. If you use `--quick [322]`, `myisamchk` does not create a temporary `.MYD` file, but instead assumes that the `.MYD` file is correct and generates only a new index file without touching the `.MYD` file. This is safe, because `myisamchk` automatically detects whether the `.MYD` file is corrupt and aborts the repair if it is. You can also specify the `--quick [322]` option twice to `myisamchk`. In this case, `myisamchk` does not abort on some errors (such as duplicate-key errors) but instead tries to resolve them by modifying the `.MYD` file. Normally the use of two `--quick [322]` options is useful only if you have too little free disk space to perform a normal repair. In this case, you should at least make a backup of the table before running `myisamchk`.

## 7.6.2. How to Check MyISAM Tables for Errors

To check a MyISAM table, use the following commands:

- `myisamchk tbl_name`

This finds 99.99% of all errors. What it cannot find is corruption that involves *only* the data file (which is very unusual). If you want to check a table, you should normally run `myisamchk` without options or with the `-s` (silent) option.

- `myisamchk -m tbl_name`

This finds 99.999% of all errors. It first checks all index entries for errors and then reads through all rows. It calculates a checksum for all key values in the rows and verifies that the checksum matches the checksum for the keys in the index tree.

- `myisamchk -e tbl_name`

This does a complete and thorough check of all data (`-e` means “extended check”). It does a check-read of every key for each row to verify that they indeed point to the correct row. This may take a



long time for a large table that has many indexes. Normally, `myisamchk` stops after the first error it finds. If you want to obtain more information, you can add the `-v` (verbose) option. This causes `myisamchk` to keep going, up through a maximum of 20 errors.

- `myisamchk -e -i tbl_name`

This is like the previous command, but the `-i` option tells `myisamchk` to print additional statistical information.

In most cases, a simple `myisamchk` command with no arguments other than the table name is sufficient to check a table.

### 7.6.3. How to Repair MyISAM Tables

The discussion in this section describes how to use `myisamchk` on MyISAM tables (extensions `.MYI` and `.MYD`).

You can also use the `CHECK TABLE` and `REPAIR TABLE` statements to check and repair MyISAM tables. See [Section 13.7.2.2, “CHECK TABLE Syntax”](#), and [Section 13.7.2.5, “REPAIR TABLE Syntax”](#).

Symptoms of corrupted tables include queries that abort unexpectedly and observable errors such as these:

- `tbl_name.frm` is locked against change
- Can't find file `tbl_name.MYI` (Errcode: `nnn`)
- Unexpected end of file
- Record file is crashed
- Got error `nnn` from table handler

To get more information about the error, run `pererror nnn`, where `nnn` is the error number. The following example shows how to use `pererror` to find the meanings for the most common error numbers that indicate a problem with a table:

```
shell> pererror 126 127 132 134 135 136 141 144 145
MySQL error code 126 = Index file is crashed
MySQL error code 127 = Record-file is crashed
MySQL error code 132 = Old database file
MySQL error code 134 = Record was already deleted (or record file crashed)
MySQL error code 135 = No more room in record file
MySQL error code 136 = No more room in index file
MySQL error code 141 = Duplicate unique key or constraint on write or update
MySQL error code 144 = Table is crashed and last repair failed
MySQL error code 145 = Table was marked as crashed and should be repaired
```

Note that error 135 (no more room in record file) and error 136 (no more room in index file) are not errors that can be fixed by a simple repair. In this case, you must use `ALTER TABLE` to increase the `MAX_ROWS` and `AVG_ROW_LENGTH` table option values:

```
ALTER TABLE tbl_name MAX_ROWS=xxx AVG_ROW_LENGTH=yyy;
```

If you do not know the current table option values, use `SHOW CREATE TABLE`.

For the other errors, you must repair your tables. `myisamchk` can usually detect and fix most problems that occur.

The repair process involves up to four stages, described here. Before you begin, you should change location to the database directory and check the permissions of the table files. On Unix, make sure that they are readable by the user that `mysqld` runs as (and to you, because you need to access the files you are checking). If it turns out you need to modify files, they must also be writable by you.

This section is for the cases where a table check fails (such as those described in [Section 7.6.2, “How to Check MyISAM Tables for Errors”](#)), or you want to use the extended features that `myisamchk` provides.

The `myisamchk` options used for table maintenance with are described in [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#). `myisamchk` also has variables that you can set to control memory allocation that may improve performance. See [Section 4.6.3.6, “myisamchk Memory Usage”](#).

If you are going to repair a table from the command line, you must first stop the `mysqld` server. Note that when you do `mysqladmin shutdown` on a remote server, the `mysqld` server is still available for a while after `mysqladmin` returns, until all statement-processing has stopped and all index changes have been flushed to disk.

### Stage 1: Checking your tables

Run `myisamchk *.MYI` or `myisamchk -e *.MYI` if you have more time. Use the `-s` (silent) option to suppress unnecessary information.

If the `mysqld` server is stopped, you should use the `--update-state [321]` option to tell `myisamchk` to mark the table as “checked.”

You have to repair only those tables for which `myisamchk` announces an error. For such tables, proceed to Stage 2.

If you get unexpected errors when checking (such as `out of memory` errors), or if `myisamchk` crashes, go to Stage 3.

### Stage 2: Easy safe repair

First, try `myisamchk -r -q tbl_name` (`-r -q` means “quick recovery mode”). This attempts to repair the index file without touching the data file. If the data file contains everything that it should and the delete links point at the correct locations within the data file, this should work, and the table is fixed. Start repairing the next table. Otherwise, use the following procedure:

1. Make a backup of the data file before continuing.
2. Use `myisamchk -r tbl_name` (`-r` means “recovery mode”). This removes incorrect rows and deleted rows from the data file and reconstructs the index file.
3. If the preceding step fails, use `myisamchk --safe-recover tbl_name`. Safe recovery mode uses an old recovery method that handles a few cases that regular recovery mode does not (but is slower).



#### Note

If you want a repair operation to go much faster, you should set the values of the `sort_buffer_size [542]` and `key_buffer_size [492]` variables each to about 25% of your available memory when running `myisamchk`.

If you get unexpected errors when repairing (such as `out of memory` errors), or if `myisamchk` crashes, go to Stage 3.

### Stage 3: Difficult repair

You should reach this stage only if the first 16KB block in the index file is destroyed or contains incorrect information, or if the index file is missing. In this case, it is necessary to create a new index file. Do so as follows:

1. Move the data file to a safe place.
2. Use the table description file to create new (empty) data and index files:

```
shell> mysql db_name
```

```
mysql> SET autocommit=1;
mysql> TRUNCATE TABLE tbl_name;
mysql> quit
```

- Copy the old data file back onto the newly created data file. (Do not just move the old file back onto the new file. You want to retain a copy in case something goes wrong.)



### Important

If you are using replication, you should stop it prior to performing the above procedure, since it involves file system operations, and these are not logged by MySQL.

Go back to Stage 2. `myisamchk -r -q` should work. (This should not be an endless loop.)

You can also use the `REPAIR TABLE tbl_name USE_FRM` SQL statement, which performs the whole procedure automatically. There is also no possibility of unwanted interaction between a utility and the server, because the server does all the work when you use `REPAIR TABLE`. See [Section 13.7.2.5, “REPAIR TABLE Syntax”](#).

### Stage 4: Very difficult repair

You should reach this stage only if the `.frm` description file has also crashed. That should never happen, because the description file is not changed after the table is created:

- Restore the description file from a backup and go back to Stage 3. You can also restore the index file and go back to Stage 2. In the latter case, you should start with `myisamchk -r`.
- If you do not have a backup but know exactly how the table was created, create a copy of the table in another database. Remove the new data file, and then move the `.frm` description and `.MYI` index files from the other database to your crashed database. This gives you new description and index files, but leaves the `.MYD` data file alone. Go back to Stage 2 and attempt to reconstruct the index file.

## 7.6.4. MyISAM Table Optimization

To coalesce fragmented rows and eliminate wasted space that results from deleting or updating rows, run `myisamchk` in recovery mode:

```
shell> myisamchk -r tbl_name
```

You can optimize a table in the same way by using the `OPTIMIZE TABLE` SQL statement. `OPTIMIZE TABLE` does a table repair and a key analysis, and also sorts the index tree so that key lookups are faster. There is also no possibility of unwanted interaction between a utility and the server, because the server does all the work when you use `OPTIMIZE TABLE`. See [Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#).

`myisamchk` has a number of other options that you can use to improve the performance of a table:

- `--analyze [323]` or `-a`: Perform key distribution analysis. This improves join performance by enabling the join optimizer to better choose the order in which to join the tables and which indexes it should use.
- `--sort-index [323]` or `-S`: Sort the index blocks. This optimizes seeks and makes table scans that use indexes faster.
- `--sort-records=index_num [324]` or `-R index_num`: Sort data rows according to a given index. This makes your data much more localized and may speed up range-based `SELECT` and `ORDER BY` operations that use this index.

For a full description of all available options, see [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

## 7.6.5. Setting Up a MyISAM Table Maintenance Schedule

It is a good idea to perform table checks on a regular basis rather than waiting for problems to occur. One way to check and repair MyISAM tables is with the `CHECK TABLE` and `REPAIR TABLE` statements. See [Section 13.7.2, “Table Maintenance Statements”](#).

Another way to check tables is to use `myisamchk`. For maintenance purposes, you can use `myisamchk -s`. The `-s` option (short for `--silent` [319]) causes `myisamchk` to run in silent mode, printing messages only when errors occur.

It is also a good idea to enable automatic MyISAM table checking. For example, whenever the machine has done a restart in the middle of an update, you usually need to check each table that could have been affected before it is used further. (These are “expected crashed tables.”) To cause the server to check MyISAM tables automatically, start it with the `--myisam-recover-options` [432] option. See [Section 5.1.3, “Server Command Options”](#).

You should also check your tables regularly during normal system operation. For example, you can run a `cron` job to check important tables once a week, using a line like this in a `crontab` file:

```
35 0 * * 0 /path/to/myisamchk --fast --silent /path/to/datadir/*/*.MYI
```

This prints out information about crashed tables so that you can examine and repair them as necessary.

To start with, execute `myisamchk -s` each night on all tables that have been updated during the last 24 hours. As you see that problems occur infrequently, you can back off the checking frequency to once a week or so.

Normally, MySQL tables need little maintenance. If you are performing many updates to MyISAM tables with dynamic-sized rows (tables with `VARCHAR`, `BLOB`, or `TEXT` columns) or have tables with many deleted rows you may want to defragment/reclaim space from the tables from time to time. You can do this by using `OPTIMIZE TABLE` on the tables in question. Alternatively, if you can stop the `mysqld` server for a while, change location into the data directory and use this command while the server is stopped:

```
shell> myisamchk -r -s --sort-index --myisam_sort_buffer_size=16M /*.MYI
```



---

# Chapter 8. Optimization

## Table of Contents

8.1. Optimization Overview .....	856
8.2. Optimizing SQL Statements .....	858
8.2.1. Optimizing <code>SELECT</code> Statements .....	858
8.2.2. Optimizing DML Statements .....	862
8.2.3. Optimizing Database Privileges .....	863
8.2.4. Optimizing <code>INFORMATION_SCHEMA</code> Queries .....	864
8.2.5. Other Optimization Tips .....	868
8.3. Optimization and Indexes .....	871
8.3.1. How MySQL Uses Indexes .....	871
8.3.2. Using Primary Keys .....	872
8.3.3. Using Foreign Keys .....	872
8.3.4. Column Indexes .....	872
8.3.5. Multiple-Column Indexes .....	873
8.3.6. Verifying Index Usage .....	875
8.3.7. <code>InnoDB</code> and <code>MyISAM</code> Index Statistics Collection .....	875
8.3.8. Comparison of B-Tree and Hash Indexes .....	876
8.4. Optimizing Database Structure .....	877
8.4.1. Optimizing Data Size .....	878
8.4.2. Optimizing MySQL Data Types .....	879
8.4.3. Optimizing for Many Tables .....	881
8.5. Optimizing for <code>InnoDB</code> Tables .....	883
8.5.1. Optimizing Storage Layout for <code>InnoDB</code> Tables .....	883
8.5.2. Optimizing <code>InnoDB</code> Transaction Management .....	884
8.5.3. Optimizing <code>InnoDB</code> Logging .....	885
8.5.4. Bulk Data Loading for <code>InnoDB</code> Tables .....	885
8.5.5. Optimizing <code>InnoDB</code> Queries .....	886
8.5.6. Optimizing <code>InnoDB</code> DDL Operations .....	886
8.5.7. Optimizing <code>InnoDB</code> Disk I/O .....	887
8.5.8. Optimizing <code>InnoDB</code> Configuration Variables .....	888
8.5.9. Optimizing <code>InnoDB</code> for Systems with Many Tables .....	889
8.6. Optimizing for <code>MyISAM</code> Tables .....	889
8.6.1. Optimizing <code>MyISAM</code> Queries .....	889
8.6.2. Bulk Data Loading for <code>MyISAM</code> Tables .....	891
8.6.3. Speed of <code>REPAIR TABLE</code> Statements .....	892
8.7. Optimizing for <code>MEMORY</code> Tables .....	893
8.8. Understanding the Query Execution Plan .....	894
8.8.1. Optimizing Queries with <code>EXPLAIN</code> .....	894
8.8.2. <code>EXPLAIN</code> Output Format .....	895
8.8.3. <code>EXPLAIN EXTENDED</code> Output Format .....	906
8.8.4. Estimating Query Performance .....	908
8.8.5. Controlling the Query Optimizer .....	908
8.9. Buffering and Caching .....	911
8.9.1. The <code>InnoDB</code> Buffer Pool .....	911
8.9.2. The <code>MyISAM</code> Key Cache .....	914
8.9.3. The MySQL Query Cache .....	918
8.9.4. Caching of Prepared Statements and Stored Programs .....	924
8.10. Optimizing Locking Operations .....	925
8.10.1. Internal Locking Methods .....	925
8.10.2. Table Locking Issues .....	927
8.10.3. Concurrent Inserts .....	929
8.10.4. Metadata Locking .....	929
8.10.5. External Locking .....	930

8.11. Optimizing the MySQL Server .....	931
8.11.1. System Factors and Startup Parameter Tuning .....	931
8.11.2. Tuning Server Parameters .....	932
8.11.3. Optimizing Disk I/O .....	939
8.11.4. Optimizing Memory Use .....	943
8.11.5. Optimizing Network Use .....	946
8.11.6. The Thread Pool Plugin .....	949
8.12. Measuring Performance (Benchmarking) .....	954
8.12.1. Measuring the Speed of Expressions and Functions .....	955
8.12.2. The MySQL Benchmark Suite .....	955
8.12.3. Using Your Own Benchmarks .....	956
8.12.4. Measuring Performance with <code>performance_schema</code> .....	956
8.12.5. Examining Thread Information .....	956
8.13. Internal Details of MySQL Optimizations .....	969
8.13.1. Range Optimization .....	970
8.13.2. Index Merge Optimization .....	973
8.13.3. Engine Condition Pushdown Optimization .....	976
8.13.4. Index Condition Pushdown Optimization .....	977
8.13.5. Use of Index Extensions .....	978
8.13.6. <code>IS NULL</code> Optimization .....	981
8.13.7. <code>LEFT JOIN</code> and <code>RIGHT JOIN</code> Optimization .....	981
8.13.8. Nested-Loop Join Algorithms .....	982
8.13.9. Nested Join Optimization .....	984
8.13.10. Outer Join Simplification .....	989
8.13.11. Multi-Range Read Optimization .....	991
8.13.12. Block Nested-Loop and Batched Key Access Joins .....	993
8.13.13. <code>ORDER BY</code> Optimization .....	995
8.13.14. <code>GROUP BY</code> Optimization .....	998
8.13.15. <code>DISTINCT</code> Optimization .....	1000
8.13.16. Subquery Optimization .....	1001

This chapter explains how to optimize MySQL performance and provides examples. Optimization involves configuring, tuning, and measuring performance, at several levels. Depending on your job role (developer, DBA, or a combination of both), you might optimize at the level of individual SQL statements, entire applications, a single database server, or multiple networked database servers. Sometimes you can be proactive and plan in advance for performance, while other times you might troubleshoot a configuration or code issue after a problem occurs. Optimizing CPU and memory usage can also improve scalability, allowing the database to handle more load without slowing down.

## 8.1. Optimization Overview

Database performance depends on several factors at the database level, such as tables, queries, and configuration settings. These software constructs result in CPU and I/O operations at the hardware level, which you must minimize and make as efficient as possible. As you work on database performance, you start by learning the high-level rules and guidelines for the software side, and measuring performance using wall-clock time. As you become an expert, you learn more about what happens internally, and start measuring things such as CPU cycles and I/O operations.

Typical users aim to get the best database performance out of their existing software and hardware configurations. Advanced users look for opportunities to improve the MySQL software itself, or develop their own storage engines and hardware appliances to expand the MySQL ecosystem.

### Optimizing at the Database Level

The most important factor in making a database application fast is its basic design:

- Are the tables structured properly? In particular, do the columns have the right data types, and does each table have the appropriate columns for the type of work? For example, applications that

perform frequent updates often have many tables with few columns, while applications that analyze large amounts of data often have few tables with many columns.

- Are the right [indexes](#) in place to make queries efficient?
- Are you using the appropriate storage engine for each table, and taking advantage of the strengths and features of each storage engine you use? In particular, the choice of a transactional storage engine such as [InnoDB](#) or a nontransactional one such as [MyISAM](#) can be very important for performance and scalability.



#### Note

In MySQL 5.5 and higher, [InnoDB](#) is the default storage engine for new tables. In practice, the advanced [InnoDB](#) performance features mean that [InnoDB](#) tables often outperform the simpler [MyISAM](#) tables, especially for a busy database.

- Does each table use an appropriate row format? This choice also depends on the storage engine used for the table. In particular, compressed tables use less disk space and so require less disk I/O to read and write the data. Compression is available for all kinds of workloads with [InnoDB](#) tables, and for read-only [MyISAM](#) tables.
- Does the application use an appropriate [locking strategy](#)? For example, by allowing shared access when possible so that database operations can run concurrently, and requesting exclusive access when appropriate so that critical operations get top priority. Again, the choice of storage engine is significant. The [InnoDB](#) storage engine handles most locking issues without involvement from you, allowing for better concurrency in the database and reducing the amount of experimentation and tuning for your code.
- Are all [memory areas used for caching](#) sized correctly? That is, large enough to hold frequently accessed data, but not so large that they overload physical memory and cause paging. The main memory areas to configure are the [InnoDB](#) buffer pool, the [MyISAM](#) key cache, and the MySQL query cache.

## Optimizing at the Hardware Level

Any database application eventually hits hardware limits as the database becomes more and more busy. A DBA must evaluate whether it is possible to tune the application or reconfigure the server to avoid these [bottlenecks](#), or whether more hardware resources are required. System bottlenecks typically arise from these sources:

- Disk seeks. It takes time for the disk to find a piece of data. With modern disks, the mean time for this is usually lower than 10ms, so we can in theory do about 100 seeks a second. This time improves slowly with new disks and is very hard to optimize for a single table. The way to optimize seek time is to distribute the data onto more than one disk.
- Disk reading and writing. When the disk is at the correct position, we need to read or write the data. With modern disks, one disk delivers at least 10–20MB/s throughput. This is easier to optimize than seeks because you can read in parallel from multiple disks.
- CPU cycles. When the data is in main memory, we must process it to get our result. Having small tables compared to the amount of memory is the most common limiting factor. But with small tables, speed is usually not the problem.
- Memory bandwidth. When the CPU needs more data than can fit in the CPU cache, main memory bandwidth becomes a bottleneck. This is an uncommon bottleneck for most systems, but one to be aware of.

## Balancing Portability and Performance



To use performance-oriented SQL extensions in a portable MySQL program, you can wrap MySQL-specific keywords in a statement within `/*! */` comment delimiters. Other SQL servers ignore the commented keywords. For information about writing comments, see [Section 9.6, “Comment Syntax”](#).

## 8.2. Optimizing SQL Statements

The core logic of a database application is performed through SQL statements, whether issued directly through an interpreter or submitted behind the scenes through an API. The tuning guidelines in this section help to speed up all kinds of MySQL applications. The guidelines cover SQL operations that read and write data, the behind-the-scenes overhead for SQL operations in general, and operations used in specific scenarios such as database monitoring.

### 8.2.1. Optimizing `SELECT` Statements

Queries, in the form of `SELECT` statements, perform all the lookup operations in the database. Tuning these statements is a top priority, whether to achieve sub-second response times for dynamic web pages, or to chop hours off the time to generate huge overnight reports.

#### 8.2.1.1. Speed of `SELECT` Statements

The main considerations for optimizing queries are:

- To make a slow `SELECT . . . WHERE` query faster, the first thing to check is whether you can add an [index](#). Set up indexes on columns used in the `WHERE` clause, to speed up evaluation, filtering, and the final retrieval of results. To avoid wasted disk space, construct a small set of indexes that speed up many related queries used in your application.

Indexes are especially important for queries that reference different tables, using features such as [joins](#) and [foreign keys](#). You can use the `EXPLAIN` statement to determine which indexes are used for a `SELECT`. See [Section 8.3.1, “How MySQL Uses Indexes”](#) and [Section 8.8.1, “Optimizing Queries with `EXPLAIN`”](#).

- Isolate and tune any part of the query, such as a function call, that takes excessive time. Depending on how the query is structured, a function could be called once for every row in the result set, or even once for every row in the table, greatly magnifying any inefficiency.
- Minimize the number of [full table scans](#) in your queries, particularly for big tables.
- Keep table statistics up to date by using the `ANALYZE TABLE` statement periodically, so the optimizer has the information needed to construct an efficient execution plan.
- Learn the tuning techniques, indexing techniques, and configuration parameters that are specific to the storage engine for each table. Both [InnoDB](#) and [MyISAM](#) have sets of guidelines for enabling and sustaining high performance in queries. For details, see [Section 8.5.5, “Optimizing InnoDB Queries”](#) and [Section 8.6.1, “Optimizing MyISAM Queries”](#).
- In particular, in MySQL 5.6.4 and higher, you can optimize single-query transactions for [InnoDB](#) tables, using the technique in [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#).
- Avoid transforming the query in ways that make it hard to understand, especially if the optimizer does some of the same transformations automatically.
- If a performance issue is not easily solved by one of the basic guidelines, investigate the internal details of the specific query by reading the `EXPLAIN` plan and adjusting your indexes, `WHERE` clauses, join clauses, and so on. (When you reach a certain level of expertise, reading the `EXPLAIN` plan might be your first step for every query.)
- Adjust the size and properties of the memory areas that MySQL uses for caching. With efficient use of the [InnoDB buffer pool](#), [MyISAM key cache](#), and the MySQL query cache, repeated queries run faster because the results are retrieved from memory the second and subsequent times.

- Even for a query that runs fast using the cache memory areas, you might still optimize further so that they require less cache memory, making your application more scalable. Scalability means that your application can handle more simultaneous users, larger requests, and so on without experiencing a big drop in performance.
- Deal with locking issues, where the speed of your query might be affected by other sessions accessing the tables at the same time.

### 8.2.1.2. How MySQL Optimizes `WHERE` Clauses

This section discusses optimizations that can be made for processing `WHERE` clauses. The examples use `SELECT` statements, but the same optimizations apply for `WHERE` clauses in `DELETE` and `UPDATE` statements.



#### Note

Because work on the MySQL optimizer is ongoing, not all of the optimizations that MySQL performs are documented here.

You might be tempted to rewrite your queries to make arithmetic operations faster, while sacrificing readability. Because MySQL does similar optimizations automatically, you can often avoid this work, and leave the query in a more understandable and maintainable form. Some of the optimizations performed by MySQL follow:

- Removal of unnecessary parentheses:

```
((a AND b) AND c OR (((a AND b) AND (c AND d))))
-> (a AND b AND c) OR (a AND b AND c AND d)
```

- Constant folding:

```
(a<b AND b=c) AND a=5
-> b>5 AND b=c AND a=5
```

- Constant condition removal (needed because of constant folding):

```
(B>=5 AND B=5) OR (B=6 AND 5=5) OR (B=7 AND 5=6)
-> B=5 OR B=6
```

- Constant expressions used by indexes are evaluated only once.
- `COUNT(*)` [1283] on a single table without a `WHERE` is retrieved directly from the table information for `MyISAM` and `MEMORY` tables. This is also done for any `NOT NULL` expression when used with only one table.
- Early detection of invalid constant expressions. MySQL quickly detects that some `SELECT` statements are impossible and returns no rows.
- `HAVING` is merged with `WHERE` if you do not use `GROUP BY` or aggregate functions (`COUNT()` [1283], `MIN()` [1284], and so on).
- For each table in a join, a simpler `WHERE` is constructed to get a fast `WHERE` evaluation for the table and also to skip rows as soon as possible.
- All constant tables are read first before any other tables in the query. A constant table is any of the following:
  - An empty table or a table with one row.
  - A table that is used with a `WHERE` clause on a `PRIMARY KEY` or a `UNIQUE` index, where all index parts are compared to constant expressions and are defined as `NOT NULL`.

All of the following tables are used as constant tables:

```
SELECT * FROM t WHERE primary_key=1;
SELECT * FROM t1,t2
WHERE t1.primary_key=1 AND t2.primary_key=t1.id;
```

- The best join combination for joining the tables is found by trying all possibilities. If all columns in `ORDER BY` and `GROUP BY` clauses come from the same table, that table is preferred first when joining.
- If there is an `ORDER BY` clause and a different `GROUP BY` clause, or if the `ORDER BY` or `GROUP BY` contains columns from tables other than the first table in the join queue, a temporary table is created.
- If you use the `SQL_SMALL_RESULT` option, MySQL uses an in-memory temporary table.
- Each table index is queried, and the best index is used unless the optimizer believes that it is more efficient to use a table scan. At one time, a scan was used based on whether the best index spanned more than 30% of the table, but a fixed percentage no longer determines the choice between using an index or a scan. The optimizer now is more complex and bases its estimate on additional factors such as table size, number of rows, and I/O block size.
- In some cases, MySQL can read rows from the index without even consulting the data file. If all columns used from the index are numeric, only the index tree is used to resolve the query.
- Before each row is output, those that do not match the `HAVING` clause are skipped.

Some examples of queries that are very fast:

```
SELECT COUNT(*) FROM tbl_name;

SELECT MIN(key_part1),MAX(key_part1) FROM tbl_name;

SELECT MAX(key_part2) FROM tbl_name
WHERE key_part1=constant;

SELECT ... FROM tbl_name
ORDER BY key_part1,key_part2,... LIMIT 10;

SELECT ... FROM tbl_name
ORDER BY key_part1 DESC, key_part2 DESC, ... LIMIT 10;
```

MySQL resolves the following queries using only the index tree, assuming that the indexed columns are numeric:

```
SELECT key_part1,key_part2 FROM tbl_name WHERE key_part1=val;

SELECT COUNT(*) FROM tbl_name
WHERE key_part1=val1 AND key_part2=val2;

SELECT key_part2 FROM tbl_name GROUP BY key_part1;
```

The following queries use indexing to retrieve the rows in sorted order without a separate sorting pass:

```
SELECT ... FROM tbl_name
ORDER BY key_part1,key_part2,... ;

SELECT ... FROM tbl_name
ORDER BY key_part1 DESC, key_part2 DESC, ... ;
```

### 8.2.1.3. Optimizing `LIMIT` Queries

If you need only a specified number of rows from a result set, use a `LIMIT` clause in the query, rather than fetching the whole result set and throwing away the extra data.

MySQL sometimes optimizes a query that has a `LIMIT row_count` clause and no `HAVING` clause:

- If you select only a few rows with `LIMIT`, MySQL uses indexes in some cases when normally it would prefer to do a full table scan.

- If you use `LIMIT row_count` with `ORDER BY`, MySQL ends the sorting as soon as it has found the first `row_count` rows of the sorted result, rather than sorting the entire result. If ordering is done by using an index, this is very fast. If a filesort must be done, all rows that match the query without the `LIMIT` clause are selected, and most or all of them are sorted, before the first `row_count` are found. After the initial rows have been found, MySQL does not sort any remainder of the result set.
- When combining `LIMIT row_count` with `DISTINCT`, MySQL stops as soon as it finds `row_count` unique rows.
- In some cases, a `GROUP BY` can be resolved by reading the key in order (or doing a sort on the key) and then calculating summaries until the key value changes. In this case, `LIMIT row_count` does not calculate any unnecessary `GROUP BY` values.
- As soon as MySQL has sent the required number of rows to the client, it aborts the query unless you are using `SQL_CALC_FOUND_ROWS`.
- `LIMIT 0` quickly returns an empty set. This can be useful for checking the validity of a query. When using one of the MySQL APIs, it can also be employed for obtaining the types of the result columns. (This trick does not work in the MySQL Monitor (the `mysql` program), which merely displays `Empty set` in such cases; instead, use `SHOW COLUMNS` or `DESCRIBE` for this purpose.)
- When the server uses temporary tables to resolve the query, it uses the `LIMIT row_count` clause to calculate how much space is required.

As of MySQL 5.6.2, the optimizer more efficiently handles queries (and subqueries) of the following form:

```
SELECT ... FROM single_table ... ORDER BY non_index_column [DESC] LIMIT [M,]N;
```

That type of query is common in web applications that display only a few rows from a larger result set. For example:

```
SELECT col1, ... FROM t1 ... ORDER BY name LIMIT 10;
SELECT col1, ... FROM t1 ... ORDER BY RAND() LIMIT 15;
```

The sort buffer has a size of `sort_buffer_size` [542]. If the sort elements for `N` rows are small enough to fit in the sort buffer (`M+N` rows if `M` was specified), the server can avoid using a merge file and perform the sort entirely in memory by treating the sort buffer as a priority queue:

- Scan the table, inserting the select list columns from each selected row in sorted order in the queue. If the queue is full, bump out the last row in the sort order.
- Return the first `N` rows from the queue. (If `M` was specified, skip the first `M` rows and return the next `N` rows.)

Previously, the server performed this operation by using a merge file for the sort:

- Scan the table, repeating these steps through the end of the table:
  - Select rows until the sort buffer is filled.
  - Write the first `N` rows in the buffer (`M+N` rows if `M` was specified) to a merge file.
- Sort the merge file and return the first `N` rows. (If `M` was specified, skip the first `M` rows and return the next `N` rows.)

The cost of the table scan is the same for the queue and merge-file methods, so the optimizer chooses between methods based on other costs:

- The queue method involves more CPU for inserting rows into the queue in order
- The merge-file method has I/O costs to write and read the file and CPU cost to sort it

The optimizer considers the balance between these factors for particular values of `N` and the row size.

### 8.2.1.4. How to Avoid Full Table Scans

The output from `EXPLAIN` shows `ALL` in the `type` column when MySQL uses a `full table scan` to resolve a query. This usually happens under the following conditions:

- The table is so small that it is faster to perform a table scan than to bother with a key lookup. This is common for tables with fewer than 10 rows and a short row length.
- There are no usable restrictions in the `ON` or `WHERE` clause for indexed columns.
- You are comparing indexed columns with constant values and MySQL has calculated (based on the index tree) that the constants cover too large a part of the table and that a table scan would be faster. See [Section 8.2.1.2, “How MySQL Optimizes WHERE Clauses”](#).
- You are using a key with low cardinality (many rows match the key value) through another column. In this case, MySQL assumes that by using the key it probably will do many key lookups and that a table scan would be faster.

For small tables, a table scan often is appropriate and the performance impact is negligible. For large tables, try the following techniques to avoid having the optimizer incorrectly choose a table scan:

- Use `ANALYZE TABLE tbl_name` to update the key distributions for the scanned table. See [Section 13.7.2.1, “ANALYZE TABLE Syntax”](#).
- Use `FORCE INDEX` for the scanned table to tell MySQL that table scans are very expensive compared to using the given index:

```
SELECT * FROM t1, t2 FORCE INDEX (index_for_column)
WHERE t1.col_name=t2.col_name;
```

See [Section 13.2.9.3, “Index Hint Syntax”](#).

- Start `mysqld` with the `--max-seeks-for-key=1000 [507]` option or use `SET max_seeks_for_key=1000` to tell the optimizer to assume that no key scan causes more than 1,000 key seeks. See [Section 5.1.4, “Server System Variables”](#).

## 8.2.2. Optimizing DML Statements

This section explains how to speed up the data manipulation language (DML) statements, `INSERT`, `UPDATE`, and `DELETE`. Traditional OLTP applications and modern web applications typically do many small DML operations, where concurrency is vital. Data analysis and reporting applications typically run DML operations that affect many rows at once, where the main considerations is the I/O to write large amounts of data and keep indexes up-to-date. For inserting and updating large volumes of data (known in the industry as ETL, for “extract-transform-load”), sometimes you use other SQL statements or external commands, that mimic the effects of `INSERT`, `UPDATE`, and `DELETE` statements.

### 8.2.2.1. Speed of INSERT Statements

To optimize insert speed, combine many small operations into a single large operation. Ideally, you make a single connection, send the data for many new rows at once, and delay all index updates and consistency checking until the very end.

The time required for inserting a row is determined by the following factors, where the numbers indicate approximate proportions:

- Connecting: (3)
- Sending query to server: (2)
- Parsing query: (2)
- Inserting row: (1 × size of row)

- Inserting indexes: (1 × number of indexes)
- Closing: (1)

This does not take into consideration the initial overhead to open tables, which is done once for each concurrently running query.

The size of the table slows down the insertion of indexes by  $\log N$ , assuming B-tree indexes.

You can use the following methods to speed up inserts:

- If you are inserting many rows from the same client at the same time, use `INSERT` statements with multiple `VALUES` lists to insert several rows at a time. This is considerably faster (many times faster in some cases) than using separate single-row `INSERT` statements. If you are adding data to a nonempty table, you can tune the `bulk_insert_buffer_size` [466] variable to make data insertion even faster. See [Section 5.1.4, “Server System Variables”](#).
- When loading a table from a text file, use `LOAD DATA INFILE`. This is usually 20 times faster than using `INSERT` statements. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).
- Take advantage of the fact that columns have default values. Insert values explicitly only when the value to be inserted differs from the default. This reduces the parsing that MySQL must do and improves the insert speed.
- See [Section 8.5.4, “Bulk Data Loading for InnoDB Tables”](#) for tips specific to InnoDB tables.
- See [Section 8.6.2, “Bulk Data Loading for MyISAM Tables”](#) for tips specific to MyISAM tables.

### 8.2.2.2. Speed of `UPDATE` Statements

An update statement is optimized like a `SELECT` query with the additional overhead of a write. The speed of the write depends on the amount of data being updated and the number of indexes that are updated. Indexes that are not changed do not get updated.

Another way to get fast updates is to delay updates and then do many updates in a row later. Performing multiple updates together is much quicker than doing one at a time if you lock the table.

For a MyISAM table that uses dynamic row format, updating a row to a longer total length may split the row. If you do this often, it is very important to use `OPTIMIZE TABLE` occasionally. See [Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#).

### 8.2.2.3. Speed of `DELETE` Statements

The time required to delete individual rows in a MyISAM table is exactly proportional to the number of indexes. To delete rows more quickly, you can increase the size of the key cache by increasing the `key_buffer_size` [492] system variable. See [Section 8.11.2, “Tuning Server Parameters”](#).

To delete all rows from a MyISAM table, `TRUNCATE TABLE tbl_name` is faster than `DELETE FROM tbl_name`. Truncate operations are not transaction-safe; an error occurs when attempting one in the course of an active transaction or active table lock. See [Section 13.1.27, “TRUNCATE TABLE Syntax”](#).

## 8.2.3. Optimizing Database Privileges

The more complex your privilege setup, the more overhead applies to all SQL statements. Simplifying the privileges established by `GRANT` statements enables MySQL to reduce permission-checking overhead when clients execute statements. For example, if you do not grant any table-level or column-level privileges, the server need not ever check the contents of the `tables_priv` and `columns_priv` tables. Similarly, if you place no resource limits on any accounts, the server does not have to perform resource counting. If you have a very high statement-processing load, consider using a simplified grant structure to reduce permission-checking overhead.



## 8.2.4. Optimizing `INFORMATION_SCHEMA` Queries

Applications that monitor the database can make frequent use of the `INFORMATION_SCHEMA` tables. Certain types of queries for `INFORMATION_SCHEMA` tables can be optimized to execute more quickly. The goal is to minimize file operations (for example, scanning a directory or opening a table file) to collect the information that makes up these dynamic tables. These optimizations do have an effect on how collations are used for searches in `INFORMATION_SCHEMA` tables. For more information, see [Section 10.1.7.9, “Collation and `INFORMATION\_SCHEMA` Searches”](#).

### 1) Try to use constant lookup values for database and table names in the `WHERE` clause

You can take advantage of this principle as follows:

- To look up databases or tables, use expressions that evaluate to a constant, such as literal values, functions that return a constant, or scalar subqueries.
- Avoid queries that use a nonconstant database name lookup value (or no lookup value) because they require a scan of the data directory to find matching database directory names.
- Within a database, avoid queries that use a nonconstant table name lookup value (or no lookup value) because they require a scan of the database directory to find matching table files.

This principle applies to the `INFORMATION_SCHEMA` tables shown in the following table, which shows the columns for which a constant lookup value enables the server to avoid a directory scan. For example, if you are selecting from `TABLES`, using a constant lookup value for `TABLE_SCHEMA` in the `WHERE` clause enables a data directory scan to be avoided.

Table	Column to specify to avoid data directory scan	Column to specify to avoid database directory scan
<code>COLUMNS</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>KEY_COLUMN_USAGE</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>PARTITIONS</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>REFERENTIAL_CONSTRAINTS</code>	<code>CONSTRAINT_SCHEMA</code>	<code>TABLE_NAME</code>
<code>STATISTICS</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>TABLES</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>TABLE_CONSTRAINTS</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>
<code>TRIGGERS</code>	<code>EVENT_OBJECT_SCHEMA</code>	<code>EVENT_OBJECT_TABLE</code>
<code>VIEWS</code>	<code>TABLE_SCHEMA</code>	<code>TABLE_NAME</code>

The benefit of a query that is limited to a specific constant database name is that checks need be made only for the named database directory. Example:

```
SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA = 'test';
```

Use of the literal database name `test` enables the server to check only the `test` database directory, regardless of how many databases there might be. By contrast, the following query is less efficient because it requires a scan of the data directory to determine which database names match the pattern `'test%'`:

```
SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA LIKE 'test%';
```

For a query that is limited to a specific constant table name, checks need be made only for the named table within the corresponding database directory. Example:

```
SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA = 'test' AND TABLE_NAME = 't1';
```

Use of the literal table name `t1` enables the server to check only the files for the `t1` table, regardless of how many tables there might be in the `test` database. By contrast, the following query requires a scan of the `test` database directory to determine which table names match the pattern `'t%'`:

```
SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA = 'test' AND TABLE_NAME LIKE 't%';
```

The following query requires a scan of the database directory to determine matching database names for the pattern `'test%'`, and for each matching database, it requires a scan of the database directory to determine matching table names for the pattern `'t%'`:

```
SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA = 'test%' AND TABLE_NAME LIKE 't%';
```

## 2) Write queries that minimize the number of table files that must be opened

For queries that refer to certain INFORMATION\_SCHEMA table columns, several optimizations are available that minimize the number of table files that must be opened. Example:

```
SELECT TABLE_NAME, ENGINE FROM INFORMATION_SCHEMA.TABLES
WHERE TABLE_SCHEMA = 'test';
```

In this case, after the server has scanned the database directory to determine the names of the tables in the database, those names become available with no further file system lookups. Thus, `TABLE_NAME` requires no files to be opened. The `ENGINE` (storage engine) value can be determined by opening the table's `.frm` file, without touching other table files such as the `.MYD` or `.MYI` file.

Some values, such as `INDEX_LENGTH` for `MyISAM` tables, require opening the `.MYD` or `.MYI` file as well.

The file-opening optimization types are denoted thus:

- `SKIP_OPEN_TABLE`: Table files do not need to be opened. The information has already become available within the query by scanning the database directory.
- `OPEN_FRM_ONLY`: Only the table's `.frm` file need be opened.
- `OPEN_TRIGGER_ONLY`: Only the table's `.TRG` file need be opened.
- `OPEN_FULL_TABLE`: The unoptimized information lookup. The `.frm`, `.MYD`, and `.MYI` files must be opened.

The following list indicates how the preceding optimization types apply to INFORMATION\_SCHEMA table columns. For tables and columns not named, none of the optimizations apply.

- `COLUMNS`: `OPEN_FRM_ONLY` applies to all columns
- `KEY_COLUMN_USAGE`: `OPEN_FULL_TABLE` applies to all columns
- `PARTITIONS`: `OPEN_FULL_TABLE` applies to all columns
- `REFERENTIAL_CONSTRAINTS`: `OPEN_FULL_TABLE` applies to all columns
- `STATISTICS`:

Column	Optimization type
<code>TABLE_CATALOG</code>	<code>OPEN_FRM_ONLY</code>
<code>TABLE_SCHEMA</code>	<code>OPEN_FRM_ONLY</code>
<code>TABLE_NAME</code>	<code>OPEN_FRM_ONLY</code>
<code>NON_UNIQUE</code>	<code>OPEN_FRM_ONLY</code>
<code>INDEX_SCHEMA</code>	<code>OPEN_FRM_ONLY</code>
<code>INDEX_NAME</code>	<code>OPEN_FRM_ONLY</code>



Column	Optimization type
SEQ_IN_INDEX	OPEN_FRM_ONLY
COLUMN_NAME	OPEN_FRM_ONLY
COLLATION	OPEN_FRM_ONLY
CARDINALITY	OPEN_FULL_TABLE
SUB_PART	OPEN_FRM_ONLY
PACKED	OPEN_FRM_ONLY
NULLABLE	OPEN_FRM_ONLY
INDEX_TYPE	OPEN_FULL_TABLE
COMMENT	OPEN_FRM_ONLY

- TABLES:

Column	Optimization type
TABLE_CATALOG	SKIP_OPEN_TABLE
TABLE_SCHEMA	SKIP_OPEN_TABLE
TABLE_NAME	SKIP_OPEN_TABLE
TABLE_TYPE	OPEN_FRM_ONLY
ENGINE	OPEN_FRM_ONLY
VERSION	OPEN_FRM_ONLY
ROW_FORMAT	OPEN_FULL_TABLE
TABLE_ROWS	OPEN_FULL_TABLE
AVG_ROW_LENGTH	OPEN_FULL_TABLE
DATA_LENGTH	OPEN_FULL_TABLE
MAX_DATA_LENGTH	OPEN_FULL_TABLE
INDEX_LENGTH	OPEN_FULL_TABLE
DATA_FREE	OPEN_FULL_TABLE
AUTO_INCREMENT	OPEN_FULL_TABLE
CREATE_TIME	OPEN_FULL_TABLE
UPDATE_TIME	OPEN_FULL_TABLE
CHECK_TIME	OPEN_FULL_TABLE
TABLE_COLLATION	OPEN_FRM_ONLY
CHECKSUM	OPEN_FULL_TABLE
CREATE_OPTIONS	OPEN_FRM_ONLY
TABLE_COMMENT	OPEN_FRM_ONLY

- TABLE\_CONSTRAINTS: OPEN\_FULL\_TABLE applies to all columns

- TRIGGERS: OPEN\_TRIGGER\_ONLY applies to all columns

- VIEWS:

Column	Optimization type
TABLE_CATALOG	OPEN_FRM_ONLY
TABLE_SCHEMA	OPEN_FRM_ONLY
TABLE_NAME	OPEN_FRM_ONLY

Column	Optimization type
VIEW_DEFINITION	OPEN_FRM_ONLY
CHECK_OPTION	OPEN_FRM_ONLY
IS_UPDATABLE	OPEN_FULL_TABLE
DEFINER	OPEN_FRM_ONLY
SECURITY_TYPE	OPEN_FRM_ONLY
CHARACTER_SET_CLIENT	OPEN_FRM_ONLY
COLLATION_CONNECTION	OPEN_FRM_ONLY

### 3) Use `EXPLAIN` to determine whether the server can use `INFORMATION_SCHEMA` optimizations for a query

This applies particularly for `INFORMATION_SCHEMA` queries that search for information from more than one database, which might take a long time and impact performance. The `Extra` value in `EXPLAIN` output indicates which, if any, of the optimizations described earlier the server can use to evaluate `INFORMATION_SCHEMA` queries. The following examples demonstrate the kinds of information you can expect to see in the `Extra` value.

```
mysql> EXPLAIN SELECT TABLE_NAME FROM INFORMATION_SCHEMA.VIEWS WHERE
-> TABLE_SCHEMA = 'test' AND TABLE_NAME = 'v1'\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
        table: VIEWS
        type: ALL
possible_keys: NULL
      key: TABLE_SCHEMA, TABLE_NAME
      key_len: NULL
       ref: NULL
      rows: NULL
      Extra: Using where; Open_frm_only; Scanned 0 databases
```

Use of constant database and table lookup values enables the server to avoid directory scans. For references to `VIEWS.TABLE_NAME`, only the `.frm` file need be opened.

```
mysql> EXPLAIN SELECT TABLE_NAME, ROW_FORMAT FROM INFORMATION_SCHEMA.TABLES\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
        table: TABLES
        type: ALL
possible_keys: NULL
      key: NULL
      key_len: NULL
       ref: NULL
      rows: NULL
      Extra: Open_full_table; Scanned all databases
```

No lookup values are provided (there is no `WHERE` clause), so the server must scan the data directory and each database directory. For each table thus identified, the table name and row format are selected. `TABLE_NAME` requires no further table files to be opened (the `SKIP_OPEN_TABLE` optimization applies). `ROW_FORMAT` requires all table files to be opened (`OPEN_FULL_TABLE` applies). `EXPLAIN` reports `OPEN_FULL_TABLE` because it is more expensive than `SKIP_OPEN_TABLE`.

```
mysql> EXPLAIN SELECT TABLE_NAME, TABLE_TYPE FROM INFORMATION_SCHEMA.TABLES
-> WHERE TABLE_SCHEMA = 'test'\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
        table: TABLES
        type: ALL
possible_keys: NULL
      key: TABLE_SCHEMA
      key_len: NULL
```

```

      ref: NULL
      rows: NULL
      Extra: Using where; Open_frm_only; Scanned 1 database

```

No table name lookup value is provided, so the server must scan the `test` database directory. For the `TABLE_NAME` and `TABLE_TYPE` columns, the `SKIP_OPEN_TABLE` and `OPEN_FRM_ONLY` optimizations apply, respectively. `EXPLAIN` reports `OPEN_FRM_ONLY` because it is more expensive.

```

mysql> EXPLAIN SELECT B.TABLE_NAME
-> FROM INFORMATION_SCHEMA.TABLES AS A, INFORMATION_SCHEMA.COLUMNS AS B
-> WHERE A.TABLE_SCHEMA = 'test'
-> AND A.TABLE_NAME = 't1'
-> AND B.TABLE_NAME = A.TABLE_NAME\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
      table: A
      type: ALL
possible_keys: NULL
      key: TABLE_SCHEMA, TABLE_NAME
      key_len: NULL
      ref: NULL
      rows: NULL
      Extra: Using where; Skip_open_table; Scanned 0 databases
***** 2. row *****
      id: 1
      select_type: SIMPLE
      table: B
      type: ALL
possible_keys: NULL
      key: NULL
      key_len: NULL
      ref: NULL
      rows: NULL
      Extra: Using where; Open_frm_only; Scanned all databases;
            Using join buffer

```

For the first `EXPLAIN` output row: Constant database and table lookup values enable the server to avoid directory scans for `TABLES` values. References to `TABLES . TABLE_NAME` require no further table files.

For the second `EXPLAIN` output row: All `COLUMNS` table values are `OPEN_FRM_ONLY` lookups, so `COLUMNS . TABLE_NAME` requires the `.frm` file to be opened.

```

mysql> EXPLAIN SELECT * FROM INFORMATION_SCHEMA.COLLATIONS\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
      table: COLLATIONS
      type: ALL
possible_keys: NULL
      key: NULL
      key_len: NULL
      ref: NULL
      rows: NULL
      Extra:

```

In this case, no optimizations apply because `COLLATIONS` is not one of the `INFORMATION_SCHEMA` tables for which optimizations are available.

## 8.2.5. Other Optimization Tips

This section lists a number of miscellaneous tips for improving query processing speed:

- Use persistent connections to the database to avoid connection overhead. If you cannot use persistent connections and you are initiating many new connections to the database, you may want to change the value of the `thread_cache_size` [552] variable. See Section 8.11.2, “Tuning Server Parameters”.

- Always check whether all your queries really use the indexes that you have created in the tables. In MySQL, you can do this with the `EXPLAIN` statement. See [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#).
- Try to avoid complex `SELECT` queries on `MyISAM` tables that are updated frequently, to avoid problems with table locking that occur due to contention between readers and writers.
- `MyISAM` supports concurrent inserts: If a table has no free blocks in the middle of the data file, you can `INSERT` new rows into it at the same time that other threads are reading from the table. If it is important to be able to do this, consider using the table in ways that avoid deleting rows. Another possibility is to run `OPTIMIZE TABLE` to defragment the table after you have deleted a lot of rows from it. This behavior is altered by setting the `concurrent_insert` [471] variable. You can force new rows to be appended (and therefore permit concurrent inserts), even in tables that have deleted rows. See [Section 8.10.3, “Concurrent Inserts”](#).
- To fix any compression issues that may have occurred with `ARCHIVE` tables, you can use `OPTIMIZE TABLE`. See [Section 14.6, “The ARCHIVE Storage Engine”](#).
- Use `ALTER TABLE ... ORDER BY expr1, expr2, ...` if you usually retrieve rows in `expr1, expr2, ...` order. By using this option after extensive changes to the table, you may be able to get higher performance.
- In some cases, it may make sense to introduce a column that is “hashed” based on information from other columns. If this column is short, reasonably unique, and indexed, it may be much faster than a “wide” index on many columns. In MySQL, it is very easy to use this extra column:

```
SELECT * FROM tbl_name
WHERE hash_col=MD5(CONCAT(col1,col2))
AND col1='constant' AND col2='constant';
```

- For `MyISAM` tables that change frequently, try to avoid all variable-length columns (`VARCHAR`, `BLOB`, and `TEXT`). The table uses dynamic row format if it includes even a single variable-length column. See [Chapter 14, Storage Engines](#).
- It is normally not useful to split a table into different tables just because the rows become large. In accessing a row, the biggest performance hit is the disk seek needed to find the first byte of the row. After finding the data, most modern disks can read the entire row fast enough for most applications. The only cases where splitting up a table makes an appreciable difference is if it is a `MyISAM` table using dynamic row format that you can change to a fixed row size, or if you very often need to scan the table but do not need most of the columns. See [Chapter 14, Storage Engines](#).
- If you often need to calculate results such as counts based on information from a lot of rows, it may be preferable to introduce a new table and update the counter in real time. An update of the following form is very fast:

```
UPDATE tbl_name SET count_col=count_col+1 WHERE key_col=constant;
```

This is very important when you use MySQL storage engines such as `MyISAM` that has only table-level locking (multiple readers with single writers). This also gives better performance with most database systems, because the row locking manager in this case has less to do.

- If you need to collect statistics from large log tables, use summary tables instead of scanning the entire log table. Maintaining the summaries should be much faster than trying to calculate statistics “live.” Regenerating new summary tables from the logs when things change (depending on business decisions) is faster than changing the running application.
- If possible, classify reports as “live” or as “statistical,” where data needed for statistical reports is created only from summary tables that are generated periodically from the live data.
- Take advantage of the fact that columns have default values. Insert values explicitly only when the value to be inserted differs from the default. This reduces the parsing that MySQL must do and improves the insert speed.

- In some cases, it is convenient to pack and store data into a `BLOB` column. In this case, you must provide code in your application to pack and unpack information, but this may save a lot of accesses at some stage. This is practical when you have data that does not conform well to a rows-and-columns table structure.
- Normally, try to keep all data nonredundant (observing what is referred to in database theory as *third normal form*). However, there may be situations in which it can be advantageous to duplicate information or create summary tables to gain more speed.
- Stored routines or UDFs (user-defined functions) may be a good way to gain performance for some tasks. See [Section 18.2, “Using Stored Routines \(Procedures and Functions\)”](#), and [Section 22.3, “Adding New Functions to MySQL”](#), for more information.
- You can increase performance by caching queries or answers in your application and then executing many inserts or updates together. If your database system supports table locks (as does MySQL), this should help to ensure that the index cache is only flushed once after all updates. You can also take advantage of MySQL's query cache to achieve similar results; see [Section 8.9.3, “The MySQL Query Cache”](#).
- Use multiple-row `INSERT` statements to store many rows with one SQL statement. (This is a relatively portable technique.)
- Use `LOAD DATA INFILE` to load large amounts of data. This is faster than using `INSERT` statements.
- Use `AUTO_INCREMENT` columns so that each row in a table can be identified by a single unique value.
- Use `OPTIMIZE TABLE` once in a while to avoid fragmentation with dynamic-format `MyISAM` tables. See [Section 14.3.3, “MyISAM Table Storage Formats”](#).
- Use `MEMORY` tables when possible to get more speed. See [Section 14.4, “The MEMORY Storage Engine”](#). `MEMORY` tables are useful for noncritical data that is accessed often, such as information about the last displayed banner for users who don't have cookies enabled in their Web browser. User sessions are another alternative available in many Web application environments for handling volatile state data.
- With Web servers, images and other binary assets should normally be stored as files. That is, store only a reference to the file rather than the file itself in the database. Most Web servers are better at caching files than database contents, so using files is generally faster.
- Columns with identical information in different tables should be declared to have identical data types so that joins based on the corresponding columns will be faster.
- Try to keep column names simple. For example, in a table named `customer`, use a column name of `name` instead of `customer_name`. To make your names portable to other SQL servers, consider keeping them shorter than 18 characters.
- If you need really high speed, look at the low-level interfaces for data storage that the different SQL servers support. For example, by accessing the MySQL `MyISAM` storage engine directly, you could get a speed increase of two to five times compared to using the SQL interface. To be able to do this, the data must be on the same server as the application, and usually it should only be accessed by one process (because external file locking is really slow). One could eliminate these problems by introducing low-level `MyISAM` commands in the MySQL server (this could be one easy way to get more performance if needed). By carefully designing the database interface, it should be quite easy to support this type of optimization.
- If you are using numeric data, it is faster in many cases to access information from a database (using a live connection) than to access a text file. Information in the database is likely to be stored in a more compact format than in the text file, so accessing it involves fewer disk accesses. You also save code in your application because you need not parse your text files to find line and column boundaries.

- Replication can provide a performance benefit for some operations. You can distribute client retrievals among replication servers to split up the load. To avoid slowing down the master while making backups, you can make backups using a slave server. See [Chapter 16, Replication](#).
- Declaring a `MyISAM` table with the `DELAY_KEY_WRITE=1` table option makes index updates faster because they are not flushed to disk until the table is closed. The downside is that if something kills the server while such a table is open, you must ensure that the table is okay by running the server with the `--myisam-recover-options` [432] option, or by running `myisamchk` before restarting the server. (However, even in this case, you should not lose anything by using `DELAY_KEY_WRITE`, because the key information can always be generated from the data rows.)
- Use `INSERT LOW_PRIORITY` for supported nontransactional tables when you want to give `SELECT` statements higher priority than your inserts.
- Use `SELECT HIGH_PRIORITY` for supported nontransactional tables to get retrievals that jump the queue. That is, the `SELECT` is executed even if there is another client waiting to do a write.

`LOW_PRIORITY` and `HIGH_PRIORITY` have an effect only for nontransactional storage engines that use only table-level locking.

## 8.3. Optimization and Indexes

The best way to improve the performance of `SELECT` operations is to create indexes on one or more of the columns that are tested in the query. The index entries act like pointers to the table rows, allowing the query to quickly determine which rows match a condition in the `WHERE` clause, and retrieve the other column values for those rows. All MySQL data types can be indexed.

Although it can be tempting to create an indexes for every possible column used in a query, unnecessary indexes waste space and waste time for MySQL to determine which indexes to use. You must find the right balance to achieve fast queries using the optimal set of indexes.

### 8.3.1. How MySQL Uses Indexes

Indexes are used to find rows with specific column values quickly. Without an index, MySQL must begin with the first row and then read through the entire table to find the relevant rows. The larger the table, the more this costs. If the table has an index for the columns in question, MySQL can quickly determine the position to seek to in the middle of the data file without having to look at all the data. If a table has 1,000 rows, this is at least 100 times faster than reading sequentially.

Most MySQL indexes (`PRIMARY KEY`, `UNIQUE`, `INDEX`, and `FULLTEXT`) are stored in [B-trees](#). Exceptions are that indexes on spatial data types use R-trees, and that `MEMORY` tables also support [hash indexes](#).

In general, indexes are used as described in the following discussion. Characteristics specific to hash indexes (as used in `MEMORY` tables) are described at the end of this section.

MySQL uses indexes for these operations:

- To find the rows matching a `WHERE` clause quickly.
- To eliminate rows from consideration. If there is a choice between multiple indexes, MySQL normally uses the index that finds the smallest number of rows (the most [selective](#) index).
- To retrieve rows from other tables when performing joins. MySQL can use indexes on columns more efficiently if they are declared as the same type and size. In this context, `VARCHAR` and `CHAR` are considered the same if they are declared as the same size. For example, `VARCHAR(10)` and `CHAR(10)` are the same size, but `VARCHAR(10)` and `CHAR(15)` are not.

Comparison of dissimilar columns may prevent use of indexes if values cannot be compared directly without conversion. Suppose that a numeric column is compared to a string column. For a given value such as `1` in the numeric column, it might compare equal to any number of values in the string

column such as '1', ' 1', '00001', or '01.e1'. This rules out use of any indexes for the string column.

- To find the `MIN()` [1284] or `MAX()` [1284] value for a specific indexed column `key_col`. This is optimized by a preprocessor that checks whether you are using `WHERE key_part_N = constant` on all key parts that occur before `key_col` in the index. In this case, MySQL does a single key lookup for each `MIN()` [1284] or `MAX()` [1284] expression and replaces it with a constant. If all expressions are replaced with constants, the query returns at once. For example:

```
SELECT MIN(key_part2),MAX(key_part2)
FROM tbl_name WHERE key_part1=10;
```

- To sort or group a table if the sorting or grouping is done on a leftmost prefix of a usable key (for example, `ORDER BY key_part1, key_part2`). If all key parts are followed by `DESC`, the key is read in reverse order. See [Section 8.13.13, “ORDER BY Optimization”](#), and [Section 8.13.14, “GROUP BY Optimization”](#).
- In some cases, a query can be optimized to retrieve values without consulting the data rows. (An index that provides all the necessary results for a query is called a [covering index](#).) If a query uses only columns from a table that are numeric and that form a leftmost prefix for some key, the selected values can be retrieved from the index tree for greater speed:

```
SELECT key_part3 FROM tbl_name
WHERE key_part1=1
```

Indexes are less important for queries on small tables, or big tables where report queries process most or all of the rows. When a query needs to access most of the rows, reading sequentially is faster than working through an index. Sequential reads minimize disk seeks, even if not all the rows are needed for the query. See [Section 8.2.1.4, “How to Avoid Full Table Scans”](#) for details.

### 8.3.2. Using Primary Keys

The primary key for a table represents the column or set of columns that you use in your most vital queries. It has an associated index, for fast query performance. Query performance benefits from the `NOT NULL` optimization, because it cannot include any `NULL` values. With the `InnoDB` storage engine, the table data is physically organized to do ultra-fast lookups and sorts based on the primary key column or columns.

If your table is big and important, but does not have an obvious column or set of columns to use as a primary key, you might create a separate column with auto-increment values to use as the primary key. These unique IDs can serve as pointers to corresponding rows in other tables when you join tables using foreign keys.

### 8.3.3. Using Foreign Keys

If a table has many columns, and you query many different combinations of columns, it might be efficient to split the less-frequently used data into separate tables with a few columns each, and relate them back to the main table by duplicating the numeric ID column from the main table. That way, each small table can have a primary key for fast lookups of its data, and you can query just the set of columns that you need using a join operation. Depending on how the data is distributed, the queries might perform less I/O and take up less cache memory because the relevant columns are packed together on disk. (To maximize performance, queries try to read as few data blocks as possible from disk; tables with only a few columns can fit more rows in each data block.)

### 8.3.4. Column Indexes

The most common type of index involves a single column, storing copies of the values from that column in a data structure, allowing fast lookups for the rows with the corresponding column values. The B-tree data structure lets the index quickly find a specific value, a set of values, or a range of values, corresponding to operators such as `=`, `>`, `≤`, `BETWEEN`, `IN`, and so on, in a `WHERE` clause.



The maximum number of indexes per table and the maximum index length is defined per storage engine. See [Chapter 14, Storage Engines](#). All storage engines support at least 16 indexes per table and a total index length of at least 256 bytes. Most storage engines have higher limits.

## Prefix Indexes

With `col_name(N)` syntax in an index specification, you can create an index that uses only the first *N* characters of a string column. Indexing only a prefix of column values in this way can make the index file much smaller. When you index a `BLOB` or `TEXT` column, you *must* specify a prefix length for the index. For example:

```
CREATE TABLE test (blob_col BLOB, INDEX(blob_col(10)));
```

Prefixes can be up to 1000 bytes long (767 bytes for `InnoDB` tables). Note that prefix limits are measured in bytes, whereas the prefix length in `CREATE TABLE` statements is interpreted as number of characters. *Be sure to take this into account when specifying a prefix length for a column that uses a multi-byte character set.*

## FULLTEXT Indexes

You can also create `FULLTEXT` indexes. These are used for full-text searches. Only the `InnoDB` and `MyISAM` storage engines support `FULLTEXT` indexes and only for `CHAR`, `VARCHAR`, and `TEXT` columns. Indexing always takes place over the entire column and column prefix indexing is not supported. For details, see [Section 12.9, “Full-Text Search Functions”](#).

Optimizations are applied to certain kinds of `FULLTEXT` queries against single `InnoDB` tables. Queries with these characteristics are particularly efficient:

- `FULLTEXT` queries that only return the document ID, or the document ID and the search rank.
- `FULLTEXT` queries that sort the matching rows in descending order of score and apply a `LIMIT` clause to take the top *N* matching rows. For this optimization to apply, there must be no `WHERE` clauses and only a single `ORDER BY` clause in descending order.
- `FULLTEXT` queries that retrieve only the `COUNT(*)` value of rows matching a search term, with no additional `WHERE` clauses. Code the `WHERE` clause as `WHERE MATCH(text) AGAINST ('other_text')`, without any `> 0` comparison operator.

## Spatial Indexes

You can also create indexes on spatial data types. Currently, only `MyISAM` supports R-tree indexes on spatial types. Other storage engines use B-trees for indexing spatial types (except for `ARCHIVE`, which does not support spatial type indexing).

## Indexes in the MEMORY Storage Engine

The `MEMORY` storage engine uses `HASH` indexes by default, but also supports `BTREE` indexes.

### 8.3.5. Multiple-Column Indexes

MySQL can create composite indexes (that is, indexes on multiple columns). An index may consist of up to 16 columns. For certain data types, you can index a prefix of the column (see [Section 8.3.4, “Column Indexes”](#)).

MySQL can use multiple-column indexes for queries that test all the columns in the index, or queries that test just the first column, the first two columns, the first three columns, and so on. If you specify the columns in the right order in the index definition, a single composite index can speed up several kinds of queries on the same table.

A multiple-column index can be considered a sorted array, the rows of which contain values that are created by concatenating the values of the indexed columns.



**Note**

As an alternative to a composite index, you can introduce a column that is “hashed” based on information from other columns. If this column is short, reasonably unique, and indexed, it might be faster than a “wide” index on many columns. In MySQL, it is very easy to use this extra column:

```
SELECT * FROM tbl_name
  WHERE hash_col=MD5(CONCAT(val1,val2))
  AND col1=val1 AND col2=val2;
```

Suppose that a table has the following specification:

```
CREATE TABLE test (
  id          INT NOT NULL,
  last_name   CHAR(30) NOT NULL,
  first_name  CHAR(30) NOT NULL,
  PRIMARY KEY (id),
  INDEX name (last_name,first_name)
);
```

The `name` index is an index over the `last_name` and `first_name` columns. The index can be used for lookups in queries that specify values in a known range for combinations of `last_name` and `first_name` values. It can also be used for queries that specify just a `last_name` value because that column is a leftmost prefix of the index (see [Section 8.3.5, “Multiple-Column Indexes”](#)). Therefore, the `name` index is used for lookups in the following queries:

```
SELECT * FROM test WHERE last_name='Widenius';

SELECT * FROM test
  WHERE last_name='Widenius' AND first_name='Michael';

SELECT * FROM test
  WHERE last_name='Widenius'
  AND (first_name='Michael' OR first_name='Monty');

SELECT * FROM test
  WHERE last_name='Widenius'
  AND first_name >='M' AND first_name < 'N';
```

However, the `name` index is *not* used for lookups in the following queries:

```
SELECT * FROM test WHERE first_name='Michael';

SELECT * FROM test
  WHERE last_name='Widenius' OR first_name='Michael';
```

Suppose that you issue the following `SELECT` statement:

```
mysql> SELECT * FROM tbl_name WHERE col1=val1 AND col2=val2;
```

If a multiple-column index exists on `col1` and `col2`, the appropriate rows can be fetched directly. If separate single-column indexes exist on `col1` and `col2`, the optimizer attempts to use the Index Merge optimization (see [Section 8.13.2, “Index Merge Optimization”](#)), or attempts to find the most restrictive index by deciding which index excludes more rows and using that index to fetch the rows.

If the table has a multiple-column index, any leftmost prefix of the index can be used by the optimizer to find rows. For example, if you have a three-column index on (`col1`, `col2`, `col3`), you have indexed search capabilities on (`col1`), (`col1`, `col2`), and (`col1`, `col2`, `col3`).

MySQL cannot use the index to perform lookups if the columns do not form a leftmost prefix of the index. Suppose that you have the `SELECT` statements shown here:

```
SELECT * FROM tbl_name WHERE col1=val1;
SELECT * FROM tbl_name WHERE col1=val1 AND col2=val2;

SELECT * FROM tbl_name WHERE col2=val2;
SELECT * FROM tbl_name WHERE col2=val2 AND col3=val3;
```

If an index exists on `(col1, col2, col3)`, only the first two queries use the index. The third and fourth queries do involve indexed columns, but `(col2)` and `(col2, col3)` are not leftmost prefixes of `(col1, col2, col3)`.

### 8.3.6. Verifying Index Usage

Always check whether all your queries really use the indexes that you have created in the tables. Use the `EXPLAIN` statement, as described in [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#).

### 8.3.7. InnoDB and MyISAM Index Statistics Collection

Storage engines collect statistics about tables for use by the optimizer. Table statistics are based on value groups, where a value group is a set of rows with the same key prefix value. For optimizer purposes, an important statistic is the average value group size.

MySQL uses the average value group size in the following ways:

- To estimate how many rows must be read for each [ref \[898\]](#) access
- To estimate how many row a partial join will produce; that is, the number of rows that an operation of this form will produce:

```
(...) JOIN tbl_name ON tbl_name.key = expr
```

As the average value group size for an index increases, the index is less useful for those two purposes because the average number of rows per lookup increases: For the index to be good for optimization purposes, it is best that each index value target a small number of rows in the table. When a given index value yields a large number of rows, the index is less useful and MySQL is less likely to use it.

The average value group size is related to table cardinality, which is the number of value groups. The `SHOW INDEX` statement displays a cardinality value based on  $N/S$ , where  $N$  is the number of rows in the table and  $S$  is the average value group size. That ratio yields an approximate number of value groups in the table.

For a join based on the `<=>` comparison operator, `NULL` is not treated differently from any other value: `NULL <=> NULL`, just as `N <=> N` for any other  $N$ .

However, for a join based on the `=` operator, `NULL` is different from non-`NULL` values: `expr1 = expr2` is not true when `expr1` or `expr2` (or both) are `NULL`. This affects [ref \[898\]](#) accesses for comparisons of the form `tbl_name.key = expr`: MySQL will not access the table if the current value of `expr` is `NULL`, because the comparison cannot be true.

For `=` comparisons, it does not matter how many `NULL` values are in the table. For optimization purposes, the relevant value is the average size of the non-`NULL` value groups. However, MySQL does not currently enable that average size to be collected or used.

For `InnoDB` and `MyISAM` tables, you have some control over collection of table statistics by means of the `innodb_stats_method [1746]` and `myisam_stats_method [513]` system variables, respectively. These variables have three possible values, which differ as follows:

- When the variable is set to `nulls_equal`, all `NULL` values are treated as identical (that is, they all form a single value group).

If the `NULL` value group size is much higher than the average non-`NULL` value group size, this method skews the average value group size upward. This makes index appear to the optimizer to be less useful than it really is for joins that look for non-`NULL` values. Consequently, the `nulls_equal` method may cause the optimizer not to use the index for [ref \[898\]](#) accesses when it should.

- When the variable is set to `nulls_unequal`, `NULL` values are not considered the same. Instead, each `NULL` value forms a separate value group of size 1.

If you have many `NULL` values, this method skews the average value group size downward. If the average non-`NULL` value group size is large, counting `NULL` values each as a group of size

1 causes the optimizer to overestimate the value of the index for joins that look for non-NULL values. Consequently, the `nulls_unequal` method may cause the optimizer to use this index for [ref \[898\]](#) lookups when other methods may be better.

- When the variable is set to `nulls_ignored`, NULL values are ignored.

If you tend to use many joins that use `<=>` rather than `=`, NULL values are not special in comparisons and one NULL is equal to another. In this case, `nulls_equal` is the appropriate statistics method.

The `innodb_stats_method [1746]` system variable has a global value; the `myisam_stats_method [513]` system variable has both global and session values. Setting the global value affects statistics collection for tables from the corresponding storage engine. Setting the session value affects statistics collection only for the current client connection. This means that you can force a table's statistics to be regenerated with a given method without affecting other clients by setting the session value of `myisam_stats_method [513]`.

To regenerate table statistics, you can use any of the following methods:

- Execute `myisamchk --stats_method=method_name --analyze`
- Change the table to cause its statistics to go out of date (for example, insert a row and then delete it), and then set `myisam_stats_method [513]` and issue an `ANALYZE TABLE` statement

Some caveats regarding the use of `innodb_stats_method [1746]` and `myisam_stats_method [513]`:

- You can force table statistics to be collected explicitly, as just described. However, MySQL may also collect statistics automatically. For example, if during the course of executing statements for a table, some of those statements modify the table, MySQL may collect statistics. (This may occur for bulk inserts or deletes, or some `ALTER TABLE` statements, for example.) If this happens, the statistics are collected using whatever value `innodb_stats_method [1746]` or `myisam_stats_method [513]` has at the time. Thus, if you collect statistics using one method, but the system variable is set to the other method when a table's statistics are collected automatically later, the other method will be used.
- There is no way to tell which method was used to generate statistics for a given table.
- These variables apply only to InnoDB and MyISAM tables. Other storage engines have only one method for collecting table statistics. Usually it is closer to the `nulls_equal` method.

### 8.3.8. Comparison of B-Tree and Hash Indexes

Understanding the B-tree and hash data structures can help predict how different queries perform on different storage engines that use these data structures in their indexes, particularly for the `MEMORY` storage engine that lets you choose B-tree or hash indexes.

#### B-Tree Index Characteristics

A B-tree index can be used for column comparisons in expressions that use the `= [1166]`, `> [1167]`, `>= [1167]`, `< [1167]`, `<= [1167]`, or `BETWEEN [1168]` operators. The index also can be used for `LIKE [1190]` comparisons if the argument to `LIKE [1190]` is a constant string that does not start with a wildcard character. For example, the following `SELECT` statements use indexes:

```
SELECT * FROM tbl_name WHERE key_col LIKE 'Patrick%';
SELECT * FROM tbl_name WHERE key_col LIKE 'Pat%ck%';
```

In the first statement, only rows with `'Patrick' <= key_col < 'Patricl'` are considered. In the second statement, only rows with `'Pat' <= key_col < 'Pau'` are considered.

The following `SELECT` statements do not use indexes:

```
SELECT * FROM tbl_name WHERE key_col LIKE '%Patrick%';
SELECT * FROM tbl_name WHERE key_col LIKE other_col;
```

In the first statement, the `LIKE [1190]` value begins with a wildcard character. In the second statement, the `LIKE [1190]` value is not a constant.

If you use `... LIKE '%string%'` and `string` is longer than three characters, MySQL uses the *Turbo Boyer-Moore algorithm* to initialize the pattern for the string and then uses this pattern to perform the search more quickly.

A search using `col_name IS NULL` employs indexes if `col_name` is indexed.

Any index that does not span all `AND [1171]` levels in the `WHERE` clause is not used to optimize the query. In other words, to be able to use an index, a prefix of the index must be used in every `AND [1171]` group.

The following `WHERE` clauses use indexes:

```
... WHERE index_part1=1 AND index_part2=2 AND other_column=3
/* index = 1 OR index = 2 */
... WHERE index=1 OR A=10 AND index=2
/* optimized like "index_part1='hello'" */
... WHERE index_part1='hello' AND index_part3=5
/* Can use index on index1 but not on index2 or index3 */
... WHERE index1=1 AND index2=2 OR index1=3 AND index3=3;
```

These `WHERE` clauses do *not* use indexes:

```
/* index_part1 is not used */
... WHERE index_part2=1 AND index_part3=2

/* Index is not used in both parts of the WHERE clause */
... WHERE index=1 OR A=10

/* No index spans all rows */
... WHERE index_part1=1 OR index_part2=10
```

Sometimes MySQL does not use an index, even if one is available. One circumstance under which this occurs is when the optimizer estimates that using the index would require MySQL to access a very large percentage of the rows in the table. (In this case, a table scan is likely to be much faster because it requires fewer seeks.) However, if such a query uses `LIMIT` to retrieve only some of the rows, MySQL uses an index anyway, because it can much more quickly find the few rows to return in the result.

## Hash Index Characteristics

Hash indexes have somewhat different characteristics from those just discussed:

- They are used only for equality comparisons that use the `=` or `<=>` operators (but are *very* fast). They are not used for comparison operators such as `<` that find a range of values. Systems that rely on this type of single-value lookup are known as “key-value stores”; to use MySQL for such applications, use hash indexes wherever possible.
- The optimizer cannot use a hash index to speed up `ORDER BY` operations. (This type of index cannot be used to search for the next entry in order.)
- MySQL cannot determine approximately how many rows there are between two values (this is used by the range optimizer to decide which index to use). This may affect some queries if you change a `MyISAM` table to a hash-indexed `MEMORY` table.
- Only whole keys can be used to search for a row. (With a B-tree index, any leftmost prefix of the key can be used to find rows.)

## 8.4. Optimizing Database Structure

In your role as a database designer, look for the most efficient way to organize your schemas, tables, and columns. As when tuning application code, you minimize I/O, keep related items together, and plan

ahead so that performance stays high as the data volume increases. Starting with an efficient database design makes it easier for team members to write high-performing application code, and makes the database likely to endure as applications evolve and are rewritten.

### 8.4.1. Optimizing Data Size

Design your tables to minimize their space on the disk. This can result in huge improvements by reducing the amount of data written to and read from disk. Smaller tables normally require less main memory while their contents are being actively processed during query execution. Any space reduction for table data also results in smaller indexes that can be processed faster.

MySQL supports many different storage engines (table types) and row formats. For each table, you can decide which storage and indexing method to use. Choosing the proper table format for your application can give you a big performance gain. See [Chapter 14, Storage Engines](#).

You can get better performance for a table and minimize storage space by using the techniques listed here:

#### Table Columns

- Use the most efficient (smallest) data types possible. MySQL has many specialized types that save disk space and memory. For example, use the smaller integer types if possible to get smaller tables. `MEDIUMINT` is often a better choice than `INT` because a `MEDIUMINT` column uses 25% less space.
- Declare columns to be `NOT NULL` if possible. It makes SQL operations faster, by enabling better use of indexes and eliminating overhead for testing whether each value is `NULL`. You also save some storage space, one bit per column. If you really need `NULL` values in your tables, use them. Just avoid the default setting that allows `NULL` values in every column.

#### Row Format

- `InnoDB` tables use a compact storage format. In versions of MySQL earlier than 5.0.3, `InnoDB` rows contain some redundant information, such as the number of columns and the length of each column, even for fixed-size columns. By default, tables are created in the compact format (`ROW_FORMAT=COMPACT`). If you wish to downgrade to older versions of MySQL, you can request the old format with `ROW_FORMAT=REDUNDANT`.

The presence of the compact row format decreases row storage space by about 20% at the cost of increasing CPU use for some operations. If your workload is a typical one that is limited by cache hit rates and disk speed it is likely to be faster. If it is a rare case that is limited by CPU speed, it might be slower.

The compact `InnoDB` format also changes how `CHAR` columns containing UTF-8 data are stored. With `ROW_FORMAT=REDUNDANT`, a UTF-8 `CHAR(N)` occupies  $3 \times N$  bytes, given that the maximum length of a UTF-8 encoded character is three bytes. Many languages can be written primarily using single-byte UTF-8 characters, so a fixed storage length often wastes space. With `ROW_FORMAT=COMPACT` format, `InnoDB` allocates a variable amount of storage in the range from  $N$  to  $3 \times N$  bytes for these columns by stripping trailing spaces if necessary. The minimum storage length is kept as  $N$  bytes to facilitate in-place updates in typical cases.

- To minimize space even further by storing table data in compressed form, specify `ROW_FORMAT=COMPRESSED` when creating `InnoDB` tables, or run the `myisampack` command on an existing `MyISAM` table. (`InnoDB` tables compressed tables are readable and writable, while `MyISAM` compressed tables are read-only.)
- For `MyISAM` tables, if you do not have any variable-length columns (`VARCHAR`, `TEXT`, or `BLOB` columns), a fixed-size row format is used. This is faster but may waste some space. See [Section 14.3.3, “MyISAM Table Storage Formats”](#). You can hint that you want to have fixed length rows even if you have `VARCHAR` columns with the `CREATE TABLE` option `ROW_FORMAT=FIXED`.

## Indexes

- The primary index of a table should be as short as possible. This makes identification of each row easy and efficient. For `InnoDB` tables, the primary key columns are duplicated in each secondary index entry, so a short primary key saves considerable space if you have many secondary indexes.
- Create only the indexes that you need to improve query performance. Indexes are good for retrieval, but slow down insert and update operations. If you access a table mostly by searching on a combination of columns, create a single composite index on them rather than a separate index for each column. The first part of the index should be the column most used. If you *always* use many columns when selecting from the table, the first column in the index should be the one with the most duplicates, to obtain better compression of the index.
- If it is very likely that a long string column has a unique prefix on the first number of characters, it is better to index only this prefix, using MySQL's support for creating an index on the leftmost part of the column (see [Section 13.1.11, "CREATE INDEX Syntax"](#)). Shorter indexes are faster, not only because they require less disk space, but because they also give you more hits in the index cache, and thus fewer disk seeks. See [Section 8.11.2, "Tuning Server Parameters"](#).

## Joins

- In some circumstances, it can be beneficial to split into two a table that is scanned very often. This is especially true if it is a dynamic-format table and it is possible to use a smaller static format table that can be used to find the relevant rows when scanning the table.
- Declare columns with identical information in different tables with identical data types, to speed up joins based on the corresponding columns.
- Keep column names simple, so that you can use the same name across different tables and simplify join queries. For example, in a table named `customer`, use a column name of `name` instead of `customer_name`. To make your names portable to other SQL servers, consider keeping them shorter than 18 characters.

## Normalization

- Normally, try to keep all data nonredundant (observing what is referred to in database theory as *third normal form*). Instead of repeating lengthy values such as names and addresses, assign them unique IDs, repeat these IDs as needed across multiple smaller tables, and join the tables in queries by referencing the IDs in the join clause.
- If speed is more important than disk space and the maintenance costs of keeping multiple copies of data, for example in a business intelligence scenario where you analyze all the data from large tables, you can relax the normalization rules, duplicating information or creating summary tables to gain more speed.

## 8.4.2. Optimizing MySQL Data Types

### 8.4.2.1. Optimizing for Numeric Data

- For unique IDs or other values that can be represented as either strings or numbers, prefer numeric columns to string columns. Since large numeric values can be stored in fewer bytes than the corresponding strings, it is faster and takes less memory to transfer and compare them.
- If you are using numeric data, it is faster in many cases to access information from a database (using a live connection) than to access a text file. Information in the database is likely to be stored in a more compact format than in the text file, so accessing it involves fewer disk accesses. You also save code in your application because you can avoid parsing the text file to find line and column boundaries.

### 8.4.2.2. Optimizing for Character and String Types



For character and string columns, follow these guidelines:

- Use binary collation order for fast comparison and sort operations, when you do not need language-specific collation features. You can use the [BINARY](#) operator to use binary collation within a particular query.
- When comparing values from different columns, declare those columns with the same character set and collation wherever possible, to avoid string conversions while running the query.
- For column values less than 8KB in size, use binary [VARCHAR](#) instead of [BLOB](#). The [GROUP BY](#) and [ORDER BY](#) clauses can generate temporary tables, and these temporary tables can use the [MEMORY](#) storage engine if the original table does not contain any [BLOB](#) columns.
- If a table contains string columns such as name and address, but many queries do not retrieve those columns, consider splitting the string columns into a separate table and using join queries with a foreign key when necessary. When MySQL retrieves any value from a row, it reads a data block containing all the columns of that row (and possibly other adjacent rows). Keeping each row small, with only the most frequently used columns, allows more rows to fit in each data block. Such compact tables reduce disk I/O and memory usage for common queries.
- When you use a randomly generated value as a primary key in an [InnoDB](#) table, prefix it with an ascending value such as the current date and time if possible. When consecutive primary values are physically stored near each other, [InnoDB](#) can insert and retrieve them faster.
- See [Section 8.4.2.1, “Optimizing for Numeric Data”](#) for reasons why a numeric column is usually preferable to an equivalent string column.

### 8.4.2.3. Optimizing for BLOB Types

- When storing a large blob containing textual data, consider compressing it first. Do not use this technique when the entire table is compressed by [InnoDB](#) or [MyISAM](#).
- For a table with several columns, to reduce memory requirements for queries that do not use the BLOB column, consider splitting the BLOB column into a separate table and referencing it with a join query when needed.
- Since the performance requirements to retrieve and display a BLOB value might be very different from other data types, you could put the BLOB-specific table on a different storage device or even a separate database instance. For example, to retrieve a BLOB might require a large sequential disk read that is better suited to a traditional hard drive than to an [SSD device](#).
- See [Section 8.4.2.2, “Optimizing for Character and String Types”](#) for reasons why a binary [VARCHAR](#) column is sometimes preferable to an equivalent BLOB column.
- Rather than testing for equality against a very long text string, you can store a hash of the column value in a separate column, index that column, and test the hashed value in queries. (Use the [MD5\(\)](#) or [CRC32\(\)](#) function to produce the hash value.) Since hash functions can produce duplicate results for different inputs, you still include a clause `AND blob_column = long_string_value` in the query to guard against false matches; the performance benefit comes from the smaller, easily scanned index for the hashed values.

### 8.4.2.4. Using [PROCEDURE ANALYSE](#)

```
ANALYSE([max_elements[,max_memory]])
```

[ANALYSE\(\)](#) examines the result from a query and returns an analysis of the results that suggests optimal data types for each column that may help reduce table sizes. To obtain this analysis, append [PROCEDURE ANALYSE](#) to the end of a [SELECT](#) statement:

```
SELECT ... FROM ... WHERE ... PROCEDURE ANALYSE([max_elements,[max_memory]])
```

For example:



```
SELECT col1, col2 FROM table1 PROCEDURE ANALYSE(10, 2000);
```

The results show some statistics for the values returned by the query, and propose an optimal data type for the columns. This can be helpful for checking your existing tables, or after importing new data. You may need to try different settings for the arguments so that `PROCEDURE ANALYSE()` does not suggest the `ENUM` data type when it is not appropriate.

The arguments are optional and are used as follows:

- `max_elements` (default 256) is the maximum number of distinct values that `ANALYSE()` notices per column. This is used by `ANALYSE()` to check whether the optimal data type should be of type `ENUM`; if there are more than `max_elements` distinct values, then `ENUM` is not a suggested type.
- `max_memory` (default 8192) is the maximum amount of memory that `ANALYSE()` should allocate per column while trying to find all distinct values.

### 8.4.3. Optimizing for Many Tables

Some techniques for keeping individual queries fast involve splitting data across many tables. When the number of tables runs into the thousands or even millions, the overhead of dealing with all these tables becomes a new performance consideration.

#### 8.4.3.1. How MySQL Opens and Closes Tables

When you execute a `mysqladmin status` command, you should see something like this:

```
Uptime: 426 Running threads: 1 Questions: 11082
Reloads: 1 Open tables: 12
```

The `Open tables` value of 12 can be somewhat puzzling if you have only six tables.

MySQL is multi-threaded, so there may be many clients issuing queries for a given table simultaneously. To minimize the problem with multiple client sessions having different states on the same table, the table is opened independently by each concurrent session. This uses additional memory but normally increases performance. With `MyISAM` tables, one extra file descriptor is required for the data file for each client that has the table open. (By contrast, the index file descriptor is shared between all sessions.)

The `table_open_cache` [551] and `max_connections` [503] system variables affect the maximum number of files the server keeps open. If you increase one or both of these values, you may run up against a limit imposed by your operating system on the per-process number of open file descriptors. Many operating systems permit you to increase the open-files limit, although the method varies widely from system to system. Consult your operating system documentation to determine whether it is possible to increase the limit and how to do so.

`table_open_cache` [551] is related to `max_connections` [503]. For example, for 200 concurrent running connections, specify a table cache size of at least  $200 * N$ , where  $N$  is the maximum number of tables per join in any of the queries which you execute. You must also reserve some extra file descriptors for temporary tables and files.

Make sure that your operating system can handle the number of open file descriptors implied by the `table_open_cache` [551] setting. If `table_open_cache` [551] is set too high, MySQL may run out of file descriptors and refuse connections, fail to perform queries, and be very unreliable. You also have to take into account that the `MyISAM` storage engine needs two file descriptors for each unique open table. You can increase the number of file descriptors available to MySQL using the `--open-files-limit` [434] startup option to `mysqld`. See [Section C.5.2.18, “'File' Not Found and Similar Errors”](#).

The cache of open tables is kept at a level of `table_open_cache` [551] entries. The default value is 400; this can be changed with the `--table_open_cache` [551] option to `mysqld`. Note that MySQL may temporarily open more tables than this to execute queries.

MySQL closes an unused table and removes it from the table cache under the following circumstances:

- When the cache is full and a thread tries to open a table that is not in the cache.
- When the cache contains more than `table_open_cache` [551] entries and a table in the cache is no longer being used by any threads.
- When a table flushing operation occurs. This happens when someone issues a `FLUSH TABLES` statement or executes a `mysqladmin flush-tables` or `mysqladmin refresh` command.

When the table cache fills up, the server uses the following procedure to locate a cache entry to use:

- Tables that are not currently in use are released, beginning with the table least recently used.
- If a new table needs to be opened, but the cache is full and no tables can be released, the cache is temporarily extended as necessary. When the cache is in a temporarily extended state and a table goes from a used to unused state, the table is closed and released from the cache.

A `MyISAM` table is opened for each concurrent access. This means the table needs to be opened twice if two threads access the same table or if a thread accesses the table twice in the same query (for example, by joining the table to itself). Each concurrent open requires an entry in the table cache. The first open of any `MyISAM` table takes two file descriptors: one for the data file and one for the index file. Each additional use of the table takes only one file descriptor for the data file. The index file descriptor is shared among all threads.

If you are opening a table with the `HANDLER tbl_name OPEN` statement, a dedicated table object is allocated for the thread. This table object is not shared by other threads and is not closed until the thread calls `HANDLER tbl_name CLOSE` or the thread terminates. When this happens, the table is put back in the table cache (if the cache is not full). See [Section 13.2.4, “HANDLER Syntax”](#).

You can determine whether your table cache is too small by checking the `mysqld` status variable `Opened_tables` [595], which indicates the number of table-opening operations since the server started:

```
mysql> SHOW GLOBAL STATUS LIKE 'Opened_tables';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Opened_tables | 2741  |
+-----+-----+
```

If the value is very large or increases rapidly, even when you have not issued many `FLUSH TABLES` statements, increase the table cache size. See [Section 5.1.4, “Server System Variables”](#), and [Section 5.1.6, “Server Status Variables”](#).

### 8.4.3.2. Disadvantages of Creating Many Tables in the Same Database

If you have many `MyISAM` tables in the same database directory, open, close, and create operations are slow. If you execute `SELECT` statements on many different tables, there is a little overhead when the table cache is full, because for every table that has to be opened, another must be closed. You can reduce this overhead by increasing the number of entries permitted in the table cache.

### 8.4.3.3. How MySQL Uses Internal Temporary Tables

In some cases, the server creates internal temporary tables while processing queries. Such a table can be held in memory and processed by the `MEMORY` storage engine, or stored on disk and processed by the `MyISAM` storage engine. The server may create a temporary table initially as an in-memory table, then convert it to an on-disk table if it becomes too large. Users have no direct control over when the server creates an internal temporary table or which storage engine the server uses to manage it.

Temporary tables can be created under conditions such as these:

- If there is an `ORDER BY` clause and a different `GROUP BY` clause, or if the `ORDER BY` or `GROUP BY` contains columns from tables other than the first table in the join queue, a temporary table is created.

- `DISTINCT` combined with `ORDER BY` may require a temporary table.
- If you use the `SQL_SMALL_RESULT` option, MySQL uses an in-memory temporary table, unless the query also contains elements (described later) that require on-disk storage.
- Derived tables (subqueries in the `FROM` clause).
- Tables created for subquery or semi-join materialization.

To determine whether a query requires a temporary table, use `EXPLAIN` and check the `Extra` column to see whether it says `Using temporary`. See Section 8.8.1, “Optimizing Queries with `EXPLAIN`”. `EXPLAIN` will not necessarily say `Using temporary` for derived or materialized temporary tables.

Some conditions prevent the use of an in-memory temporary table, in which case the server uses an on-disk table instead:

- Presence of a `BLOB` or `TEXT` column in the table
- Presence of any column in a `GROUP BY` or `DISTINCT` clause larger than 512 bytes
- Presence of any column larger than 512 bytes in the `SELECT` list, if `UNION` or `UNION ALL` is used

If an internal temporary table is created initially as an in-memory table but becomes too large, MySQL automatically converts it to an on-disk table. The maximum size for in-memory temporary tables is the minimum of the `tmp_table_size` [558] and `max_heap_table_size` [504] values. This differs from `MEMORY` tables explicitly created with `CREATE TABLE`: For such tables, the `max_heap_table_size` [504] system variable determines how large the table is permitted to grow and there is no conversion to on-disk format.

When the server creates an internal temporary table (either in memory or on disk), it increments the `Created_tmp_tables` [588] status variable. If the server creates the table on disk (either initially or by converting an in-memory table) it increments the `Created_tmp_disk_tables` [587] status variable.

## 8.5. Optimizing for InnoDB Tables

`InnoDB` is the storage engine that MySQL customers typically use in production databases where reliability and concurrency are important. Because `InnoDB` is the default storage engine in MySQL 5.5 and higher, you can expect to see `InnoDB` tables more often than before. This section explains how to optimize database operations for `InnoDB` tables.

### 8.5.1. Optimizing Storage Layout for InnoDB Tables

- Once your data reaches a stable size, or a growing table has increased by tens or some hundreds of megabytes, consider using the `OPTIMIZE TABLE` statement to reorganize the table and compact any wasted space. The reorganized tables require less disk I/O to perform full table scans. This is a straightforward technique that can improve performance when other techniques such as improving index usage or tuning application code are not practical.

`OPTIMIZE TABLE` copies the data part of the table and rebuilds the indexes. The benefits come from improved packing of data within indexes, and reduced fragmentation within the tablespaces and on disk. The benefits vary depending on the data in each table. You may find that there are significant gains for some and not for others, or that the gains decrease over time until you next optimize the table. This operation can be slow if the table is large or if the indexes being rebuilt don't fit into the buffer pool. The first run after adding a lot of data to a table is often much slower than later runs.

- In `InnoDB`, having a long `PRIMARY KEY` (either a single column with a lengthy value, or several columns that form a long composite value) wastes a lot of disk space. The primary key value for a row is duplicated in all the secondary index records that point to the same row. (See Section 14.2.3.12, “`InnoDB` Table and Index Structures”.) Create an `AUTO_INCREMENT` column as

the primary key if your primary key is long, or index a prefix of a long `VARCHAR` column instead of the entire column.

- Use the `VARCHAR` data type instead of `CHAR` to store variable-length strings or for columns with many `NULL` values. A `CHAR(N)` column always takes `N` characters to store data, even if the string is shorter or its value is `NULL`. Smaller tables fit better in the buffer pool and reduce disk I/O.

When using `COMPACT` row format (the default InnoDB format in MySQL 5.6) and variable-length character sets, such as `utf8` or `sjis`, `CHAR(N)` columns occupy a variable amount of space, but still at least `N` bytes.

- For tables that are big, or contain lots of repetitive text or numeric data, consider using `COMPRESSED` row format. Less disk I/O is required to bring data into the buffer pool, or to perform full table scans. Before making a permanent decision, measure the amount of compression you can achieve by using `COMPRESSED` versus `COMPACT` row format.

## 8.5.2. Optimizing InnoDB Transaction Management

To optimize InnoDB transaction processing, find the ideal balance between the performance overhead of transactional features and the workload of your server. For example, an application might encounter performance issues if it commits thousands of times per second, and different performance issues if it commits only every 2-3 hours.

- The default MySQL setting `AUTOCOMMIT=1` can impose performance limitations on a busy database server. Where practical, wrap several related DML operations into a single transaction, by issuing `SET AUTOCOMMIT=0` or a `START TRANSACTION` statement, followed by a `COMMIT` statement after making all the changes.

InnoDB must flush the log to disk at each transaction commit if that transaction made modifications to the database. When each change is followed by a commit (as with the default autocommit setting), the I/O throughput of the storage device puts a cap on the number of potential operations per second.

- Alternatively, for transactions that consist only of a single `SELECT` statement, turning on `AUTOCOMMIT` helps InnoDB to recognize read-only transactions and optimize them. See [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#) for requirements.
- Avoid performing rollbacks after inserting, updating, or deleting huge numbers of rows. If a big transaction is slowing down server performance, rolling it back can make the problem worse, potentially taking several times as long to perform as the original DML operations. Killing the database process does not help, because the rollback starts again on server startup.

To minimize the chance of this issue occurring: increase the size of the `buffer pool` so that all the DML changes can be cached rather than immediately written to disk; set `innodb_change_buffering=all` [1709] so that update and delete operations are buffered in addition to inserts; and consider issuing `COMMIT` statements periodically during the big DML operation, possibly breaking a single delete or update into multiple statements that operate on smaller numbers of rows.

To get rid of a runaway rollback once it occurs, increase the buffer pool so that the rollback becomes CPU-bound and runs fast, or kill the server and restart with `innodb_force_recovery=3` [1722], as explained in [Section 14.2.3.13, “The InnoDB Recovery Process”](#).

This issue is expected to be less prominent in MySQL 5.5 and up, or in MySQL 5.1 with the InnoDB Plugin, because the default setting `innodb_change_buffering=all` [1709] allows update and delete operations to be cached in memory, making them faster to perform in the first place, and also faster to roll back if needed. Make sure to use this parameter setting on servers that process long-running transactions with many inserts, updates, or deletes.

- If you can afford the loss of some of the latest committed transactions if a crash occurs, you can set the `innodb_flush_log_at_trx_commit` [1718] parameter to 0. InnoDB tries to

flush the log once per second anyway, although the flush is not guaranteed. Also, set the value of `innodb_support_xa` [1749] to 0, which will reduce the number of disk flushes due to synchronizing on disk data and the binary log.

- When rows are modified or deleted, the rows and associated [undo logs](#) are not physically removed immediately, or even immediately after the transaction commits. The old data is preserved until transactions that started earlier or concurrently are finished, so that those transactions can access the previous state of modified or deleted rows. Thus, a long-running transaction can prevent [InnoDB](#) from purging data that was changed by a different transaction.
- When rows are modified or deleted within a long-running transaction, other transactions using the [READ COMMITTED](#) [1469] and [REPEATABLE READ](#) [1469] isolation levels have to do more work to reconstruct the older data if they read those same rows.
- When a long-running transaction modifies a table, queries against that table from other transactions do not make use of the [covering index](#) technique. Queries that normally could retrieve all the result columns from a secondary index, instead look up the appropriate values from the table data.

### 8.5.3. Optimizing InnoDB Logging

- Make your log files big, even as big as the buffer pool. When [InnoDB](#) has written the log files full, it must write the modified contents of the buffer pool to disk in a checkpoint. Small log files cause many unnecessary disk writes. Although historically big log files caused lengthy recovery times, recovery is now much faster and you can confidently use large log files.
- Make the log buffer quite large as well (on the order of 8MB).

### 8.5.4. Bulk Data Loading for InnoDB Tables

These performance tips supplement the general guidelines for fast inserts in [Section 8.2.2.1, “Speed of INSERT Statements”](#).

- When importing data into [InnoDB](#), turn off autocommit mode, because it performs a log flush to disk for every insert. To disable autocommit during your import operation, surround it with `SET autocommit` and `COMMIT` statements:

```
SET autocommit=0;... SQL import statements ...
COMMIT;
```

The `mysqldump` option `--opt` [294] creates dump files that are fast to import into an [InnoDB](#) table, even without wrapping them with the `SET autocommit` and `COMMIT` statements.

- If you have [UNIQUE](#) constraints on secondary keys, you can speed up table imports by temporarily turning off the uniqueness checks during the import session:

```
SET unique_checks=0;... SQL import statements ...
SET unique_checks=1;
```

For big tables, this saves a lot of disk I/O because [InnoDB](#) can use its insert buffer to write secondary index records in a batch. Be certain that the data contains no duplicate keys.

- If you have [FOREIGN KEY](#) constraints in your tables, you can speed up table imports by turning off the foreign key checks for the duration of the import session:

```
SET foreign_key_checks=0;... SQL import statements ...
SET foreign_key_checks=1;
```

For big tables, this can save a lot of disk I/O.

- Use the multiple-row [INSERT](#) syntax to reduce communication overhead between the client and the server if you need to insert many rows:

```
INSERT INTO yourtable VALUES (1,2), (5,5), ...;
```

This tip is valid for inserts into any table, not just InnoDB tables.

- When doing bulk inserts into tables with auto-increment columns, set `innodb_autoinc_lock_mode` [1705] to 2 instead of the default value 1. See [Section 5.4.4.2, “Configurable InnoDB Auto-Increment Locking”](#) for details.
- For optimal performance when loading data into an InnoDB FULLTEXT index, follow this set of steps:
  - Define a column `FTS_DOC_ID` at table creation time, of type `BIGINT UNSIGNED NOT NULL`, with a unique index named `FTS_DOC_ID_INDEX`.
  - Load the data into the table.
  - Create the FULLTEXT index after the data is loaded.

### 8.5.5. Optimizing InnoDB Queries

To tune queries for InnoDB tables, create an appropriate set of indexes on each table. See [Section 8.3.1, “How MySQL Uses Indexes”](#) for details. Follow these guidelines for InnoDB indexes:

- Because each InnoDB table has a **primary key** (whether you request one or not), specify a set of primary key columns for each table, columns that are used in the most important and time-critical queries.
- Do not specify too many or too long columns in the primary key, because these column values are duplicated in each secondary index. When an index contains unnecessary data, the I/O to read this data and memory to cache it reduce the performance and scalability of the server.
- Do not create a separate **secondary index** for each column, because each query can only make use of one index. Indexes on rarely tested columns or columns with only a few different values might not be helpful for any queries. If you have many queries for the same table, testing different combinations of columns, try to create a small number of **concatenated indexes** rather than a large number of single-column indexes. If an index contains all the columns needed for the result set (known as a **covering index**), the query might be able to avoid reading the table data at all.
- If an indexed column cannot contain any `NULL` values, declare it as `NOT NULL` when you create the table. The optimizer can better determine which index is most effective to use for a query, when it knows whether each column contains `NULL` values or not.
- In MySQL 5.6.4 and higher, you can optimize single-query transactions for InnoDB tables, using the technique in [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#).
- If you often have recurring queries for tables that are not updated frequently, enable the query cache:

```
[mysqld]
query_cache_type = 1
query_cache_size = 10M
```

### 8.5.6. Optimizing InnoDB DDL Operations

- For DDL operations on tables and indexes (`CREATE`, `ALTER`, and `DROP` statements), the most significant aspect for InnoDB tables is that creating and dropping secondary indexes is much faster in MySQL 5.5 and higher, than in MySQL 5.1 and before. See [Fast Index Creation in the InnoDB Storage Engine](#) for details.
- “Fast index creation” makes it faster in some cases to drop an index before loading data into a table, then re-create the index after loading the data.
- Use `TRUNCATE TABLE` to empty a table, not `DELETE FROM tbl_name`. Foreign key constraints can make a `TRUNCATE` statement work like a regular `DELETE` statement, in which case a sequence of commands like `DROP TABLE` and `CREATE TABLE` might be fastest.



- Because the primary key is integral to the storage layout of each InnoDB table, and changing the definition of the primary key involves reorganizing the whole table, always set up the primary key as part of the CREATE TABLE statement, and plan ahead so that you do not need to ALTER or DROP the primary key afterward.

### 8.5.7. Optimizing InnoDB Disk I/O

If you follow the best practices for database design and the tuning techniques for SQL operations, but your database is still slowed by heavy disk I/O activity, explore these low-level techniques related to disk I/O. If the Unix `top` tool or the Windows Task Manager shows that the CPU usage percentage with your workload is less than 70%, your workload is probably disk-bound.

- When table data is cached in the InnoDB buffer pool, it can be processed over and over by queries without requiring any disk I/O. Specify the size of the buffer pool with the `innodb_buffer_pool_size` [1708] option. This memory area is important enough that busy databases often specify a size approximately 80% of the amount of physical memory. For more information, see Section 8.9.1, “The InnoDB Buffer Pool”.
- In some versions of GNU/Linux and Unix, flushing files to disk with the Unix `fsync()` call (which InnoDB uses by default) and similar methods is surprisingly slow. If database write performance is an issue, conduct benchmarks with the `innodb_flush_method` [1719] parameter set to `O_DSYNC`.
- When using the InnoDB storage engine on Solaris 10 for x86\_64 architecture (AMD Opteron), use direct I/O for InnoDB-related files, to avoid degradation of InnoDB performance. To use direct I/O for an entire UFS file system used for storing InnoDB-related files, mount it with the `forcedirectio` option; see `mount_ufs(1M)`. (The default on Solaris 10/x86\_64 is *not* to use this option.) To apply direct I/O only to InnoDB file operations rather than the whole file system, set `innodb_flush_method = O_DIRECT` [1719]. With this setting, InnoDB calls `directio()` instead of `fcntl()` for I/O to data files (not for I/O to log files).
- When using the InnoDB storage engine with a large `innodb_buffer_pool_size` [1708] value on any release of Solaris 2.6 and up and any platform (sparc/x86/x64/amd64), conduct benchmarks with InnoDB data files and log files on raw devices or on a separate direct I/O UFS file system, using the `forcedirectio` mount option as described earlier. (It is necessary to use the mount option rather than setting `innodb_flush_method` [1719] if you want direct I/O for the log files.) Users of the Veritas file system VxFS should use the `convosync=direct` mount option.

Do not place other MySQL data files, such as those for MyISAM tables, on a direct I/O file system. Executables or libraries *must not* be placed on a direct I/O file system.

- If you have additional storage devices available to set up a RAID configuration or symbolic links to different disks, Section 8.11.3, “Optimizing Disk I/O” for additional low-level I/O tips.
- If throughput drops periodically because of InnoDB checkpoint operations, consider increasing the value of the `innodb_io_capacity` [1727] configuration option. Higher values cause more frequent flushing, avoiding the backlog of work that can cause dips in throughput.
- If the system is not falling behind with InnoDB flushing operations, consider lowering the value of the `innodb_io_capacity` [1727] configuration option. Typically, you keep this option value as low as practical, but not so low that it causes periodic drops in throughput as mentioned in the preceding bullet. In a typical scenario where you could lower the option value, you might see a combination like this in the output from `SHOW ENGINE INNODB STATUS`:
  - History list length low, below a few thousand.
  - Insert buffer merges close to rows inserted.
  - Modified pages in buffer pool consistently well below `innodb_max_dirty_pages_pct` [1734] of the buffer pool. (Measure at a time when the server is not doing bulk inserts; it is normal during bulk inserts for the modified pages percentage to rise significantly.)



- `log_sequence_number` - Last checkpoint is at less than 7/8 or ideally less than 6/8 of the total size of the InnoDB log files.
- Other InnoDB configuration options to consider when tuning I/O-bound workloads include `innodb_adaptive_flushing` [1700], `innodb_change_buffer_max_size` [1709], `innodb_change_buffering` [1709], `innodb_flush_neighbors` [1721], `innodb_log_buffer_size` [1732], `innodb_log_file_size` [1732], `innodb_lru_scan_depth` [1734], `innodb_max_dirty_pages_pct` [1734], `innodb_max_purge_lag` [1735], `innodb_open_files` [1739], `innodb_page_size` [1740], `innodb_random_read_ahead` [1742], `innodb_read_ahead_threshold` [1742], `innodb_read_io_threads` [1743], `innodb_rollback_segments` [1744], `innodb_write_io_threads` [1754], and `sync_binlog` [2017].

## 8.5.8. Optimizing InnoDB Configuration Variables

Different settings work best for servers with light, predictable loads, versus servers that are running near full capacity all the time, or that experience spikes of high activity.

Because the InnoDB storage engine performs many of its optimizations automatically, many performance-tuning tasks involve monitoring to ensure that the database is performing well, and changing configuration options when performance drops. See [Section 14.2.4.2.26, “Integration with the MySQL Performance Schema”](#) for information about detailed InnoDB performance monitoring.

For information about the most important and most recent InnoDB performance features, see [Section 14.2.4.2, “InnoDB Performance and Scalability Enhancements”](#). Even if you have used InnoDB tables in prior versions, these features might be new to you, because they are from the “InnoDB Plugin”. The Plugin co-existed alongside the built-in InnoDB in MySQL 5.1, and becomes the default storage engine in MySQL 5.5 and higher.

The main configuration steps you can perform include:

- Enabling InnoDB to use high-performance memory allocators on systems that include them. See [Section 14.2.4.2.12, “Using Operating System Memory Allocators”](#).
- Controlling the types of DML operations for which InnoDB buffers the changed data, to avoid frequent small disk writes. See [Section 14.2.4.2.13, “Controlling InnoDB Change Buffering”](#). Because the default is to buffer all types of DML operations, only change this setting if you need to reduce the amount of buffering.
- Turning the adaptive hash indexing feature on and off. See [Section 14.2.4.2.14, “Controlling Adaptive Hash Indexing”](#). You might change this setting during periods of unusual activity, then restore it to its original setting.
- Setting a limit on the number of concurrent threads that InnoDB processes, if context switching is a bottleneck. See [Section 14.2.4.2.15, “Changes Regarding Thread Concurrency”](#).
- Controlling the amount of prefetching that InnoDB does with its read-ahead operations. When the system has unused I/O capacity, more read-ahead can improve the performance of queries. Too much read-ahead can cause periodic drops in performance on a heavily loaded system. See [Section 14.2.4.2.16, “Changes in the Read-Ahead Algorithm”](#).
- Increasing the number of background threads for read or write operations, if you have a high-end I/O subsystem that is not fully utilized by the default values. See [Section 14.2.4.2.17, “Multiple Background InnoDB I/O Threads”](#).
- Controlling how much I/O InnoDB performs in the background. See [Section 14.2.4.2.20, “Controlling the InnoDB Master Thread I/O Rate”](#). The amount of background I/O is higher than in MySQL 5.1, so you might scale back this setting if you observe periodic drops in performance.
- Controlling the algorithm that determines when InnoDB performs certain types of background writes. See [Section 14.2.4.2.21, “Controlling the Flushing Rate of Dirty Pages from the InnoDB Buffer Pool”](#).

The algorithm works for some types of workloads but not others, so might turn off this setting if you observe periodic drops in performance.

- Taking advantage of multicore processors and their cache memory configuration, to minimize delays in context switching. See [Section 14.2.4.2.23, “Control of Spin Lock Polling”](#).
- Preventing one-time operations such as table scans from interfering with the frequently accessed data stored in the InnoDB buffer cache. See [Section 14.2.4.2.24, “Making the Buffer Pool Scan Resistant”](#).
- Adjusting your log files to a size that makes sense for reliability and crash recovery. See [Section 14.2.4.2.25, “Improvements to Crash Recovery Performance”](#). Historically, people have kept their InnoDB log files small to avoid long startup times after a crash. Internal improvements in InnoDB make startup much faster, so the log file size is not such a performance factor anymore. If your log files are artificially small, increasing the size can help performance by reducing the I/O that occurs as redo log records are recycled.
- Configuring the size and number of instances for the InnoDB buffer pool, especially important for systems with multi-gigabyte buffer pools. See [Section 14.2.4.2.27, “Improvements to Performance from Multiple Buffer Pools”](#).
- Increasing the maximum number of concurrent transactions, which dramatically improves scalability for the busiest databases. See [Section 14.2.4.2.28, “Better Scalability with Multiple Rollback Segments”](#). Although this feature does not require any action during day-to-day operation, you must perform a [slow shutdown](#) during or after upgrading the database to MySQL 5.5 to enable the higher limit.
- Moving purge operations (a type of garbage collection) into a background thread. See [Section 14.2.4.2.29, “Better Scalability with Improved Purge Scheduling”](#). To effectively measure the results of this setting, tune the other I/O-related and thread-related configuration settings first.
- Reducing the amount of switching that InnoDB does between concurrent threads, so that SQL operations on a busy server do not queue up and form a “traffic jam”. Set a value for the `innodb_thread_concurrency [1751]` option, up to approximately 32 for a high-powered modern system. Increase the value for the `innodb_concurrency_tickets [1714]` option, typically to 5000 or so. This combination of options sets a cap on the number of threads that InnoDB processes at any one time, and allows each thread to do substantial work before being swapped out, so that the number of waiting threads stays low and operations can complete without excessive context switching.

## 8.5.9. Optimizing InnoDB for Systems with Many Tables

- InnoDB computes index [cardinality](#) values for a table the first time that table is accessed after startup, instead of storing such values in the table. This step can take significant time on systems that partition the data into many tables. Since this overhead only applies to the initial table open operation, to “warm up” a table for later use, access it immediately after startup by issuing a statement such as `SELECT 1 FROM tbl_name LIMIT 1`.

## 8.6. Optimizing for MyISAM Tables

The MyISAM storage engine performs best with read-mostly data or with low-concurrency operations, because table locks limit the ability to perform simultaneous updates. In MySQL 5.6, InnoDB is the default storage engine rather than MyISAM.

### 8.6.1. Optimizing MyISAM Queries

Some general tips for speeding up queries on MyISAM tables:

- To help MySQL better optimize queries, use `ANALYZE TABLE` or run `myisamchk --analyze` on a table after it has been loaded with data. This updates a value for each index part that indicates the average number of rows that have the same value. (For unique indexes, this is always 1.)

MySQL uses this to decide which index to choose when you join two tables based on a nonconstant expression. You can check the result from the table analysis by using `SHOW INDEX FROM tbl_name` and examining the `Cardinality` value. `myisamchk --description --verbose` shows index distribution information.

- To sort an index and data according to an index, use `myisamchk --sort-index --sort-records=1` (assuming that you want to sort on index 1). This is a good way to make queries faster if you have a unique index from which you want to read all rows in order according to the index. The first time you sort a large table this way, it may take a long time.
- Try to avoid complex `SELECT` queries on MyISAM tables that are updated frequently, to avoid problems with table locking that occur due to contention between readers and writers.
- MyISAM supports concurrent inserts: If a table has no free blocks in the middle of the data file, you can `INSERT` new rows into it at the same time that other threads are reading from the table. If it is important to be able to do this, consider using the table in ways that avoid deleting rows. Another possibility is to run `OPTIMIZE TABLE` to defragment the table after you have deleted a lot of rows from it. This behavior is altered by setting the `concurrent_insert` [471] variable. You can force new rows to be appended (and therefore permit concurrent inserts), even in tables that have deleted rows. See Section 8.10.3, “Concurrent Inserts”.
- For MyISAM tables that change frequently, try to avoid all variable-length columns (`VARCHAR`, `BLOB`, and `TEXT`). The table uses dynamic row format if it includes even a single variable-length column. See Chapter 14, *Storage Engines*.
- It is normally not useful to split a table into different tables just because the rows become large. In accessing a row, the biggest performance hit is the disk seek needed to find the first byte of the row. After finding the data, most modern disks can read the entire row fast enough for most applications. The only cases where splitting up a table makes an appreciable difference is if it is a MyISAM table using dynamic row format that you can change to a fixed row size, or if you very often need to scan the table but do not need most of the columns. See Chapter 14, *Storage Engines*.
- Use `ALTER TABLE ... ORDER BY expr1, expr2, ...` if you usually retrieve rows in `expr1, expr2, ...` order. By using this option after extensive changes to the table, you may be able to get higher performance.
- If you often need to calculate results such as counts based on information from a lot of rows, it may be preferable to introduce a new table and update the counter in real time. An update of the following form is very fast:

```
UPDATE tbl_name SET count_col=count_col+1 WHERE key_col=constant;
```

This is very important when you use MySQL storage engines such as MyISAM that has only table-level locking (multiple readers with single writers). This also gives better performance with most database systems, because the row locking manager in this case has less to do.

- Use `INSERT DELAYED` for MyISAM (or other supported nontransactional tables) when you do not need to know when your data is written. This reduces the overall insertion impact because many rows can be written with a single disk write.



#### Note

As of MySQL 5.6.6, `INSERT DELAYED` is deprecated, and will be removed in a future release. Use `INSERT` (without `DELAYED`) instead.

- Use `OPTIMIZE TABLE` periodically to avoid fragmentation with dynamic-format MyISAM tables. See Section 14.3.3, “MyISAM Table Storage Formats”.
- Declaring a MyISAM table with the `DELAY_KEY_WRITE=1` table option makes index updates faster because they are not flushed to disk until the table is closed. The downside is that if something kills the server while such a table is open, you must ensure that the table is okay by running the server with the `--myisam-recover-options` [432] option, or by running `myisamchk` before restarting

the server. (However, even in this case, you should not lose anything by using `DELAY_KEY_WRITE`, because the key information can always be generated from the data rows.)

- Strings are automatically prefix- and end-space compressed in MyISAM indexes. See [Section 13.1.11, “CREATE INDEX Syntax”](#).
- You can increase performance by caching queries or answers in your application and then executing many inserts or updates together. Locking the table during this operation ensures that the index cache is only flushed once after all updates. You can also take advantage of MySQL's query cache to achieve similar results; see [Section 8.9.3, “The MySQL Query Cache”](#).

## 8.6.2. Bulk Data Loading for MyISAM Tables

These performance tips supplement the general guidelines for fast inserts in [Section 8.2.2.1, “Speed of INSERT Statements”](#).

- To improve performance when multiple clients insert a lot of rows, use the `INSERT DELAYED` statement. See [Section 13.2.5.2, “INSERT DELAYED Syntax”](#). This technique works for MyISAM and some other storage engines, but not InnoDB.



### Note

As of MySQL 5.6.6, `INSERT DELAYED` is deprecated, and will be removed in a future release. Use `INSERT` (without `DELAYED`) instead.

- For a MyISAM table, you can use concurrent inserts to add rows at the same time that `SELECT` statements are running, if there are no deleted rows in middle of the data file. See [Section 8.10.3, “Concurrent Inserts”](#).
- With some extra work, it is possible to make `LOAD DATA INFILE` run even faster for a MyISAM table when the table has many indexes. Use the following procedure:
  1. Execute a `FLUSH TABLES` statement or a `mysqladmin flush-tables` command.
  2. Use `myisamchk --keys-used=0 -rq /path/to/db/tbl_name` to remove all use of indexes for the table.
  3. Insert data into the table with `LOAD DATA INFILE`. This does not update any indexes and therefore is very fast.
  4. If you intend only to read from the table in the future, use `myisampack` to compress it. See [Section 14.3.3.3, “Compressed Table Characteristics”](#).
  5. Re-create the indexes with `myisamchk -rq /path/to/db/tbl_name`. This creates the index tree in memory before writing it to disk, which is much faster than updating the index during `LOAD DATA INFILE` because it avoids lots of disk seeks. The resulting index tree is also perfectly balanced.
  6. Execute a `FLUSH TABLES` statement or a `mysqladmin flush-tables` command.

`LOAD DATA INFILE` performs the preceding optimization automatically if the MyISAM table into which you insert data is empty. The main difference between automatic optimization and using the procedure explicitly is that you can let `myisamchk` allocate much more temporary memory for the index creation than you might want the server to allocate for index re-creation when it executes the `LOAD DATA INFILE` statement.

You can also disable or enable the nonunique indexes for a MyISAM table by using the following statements rather than `myisamchk`. If you use these statements, you can skip the `FLUSH TABLE` operations:

```
ALTER TABLE tbl_name DISABLE KEYS;
```

```
ALTER TABLE tbl_name ENABLE KEYS;
```

- To speed up `INSERT` operations that are performed with multiple statements for nontransactional tables, lock your tables:

```
LOCK TABLES a WRITE;
INSERT INTO a VALUES (1,23),(2,34),(4,33);
INSERT INTO a VALUES (8,26),(6,29);
...
UNLOCK TABLES;
```

This benefits performance because the index buffer is flushed to disk only once, after all `INSERT` statements have completed. Normally, there would be as many index buffer flushes as there are `INSERT` statements. Explicit locking statements are not needed if you can insert all rows with a single `INSERT`.

Locking also lowers the total time for multiple-connection tests, although the maximum wait time for individual connections might go up because they wait for locks. Suppose that five clients attempt to perform inserts simultaneously as follows:

- Connection 1 does 1000 inserts
- Connections 2, 3, and 4 do 1 insert
- Connection 5 does 1000 inserts

If you do not use locking, connections 2, 3, and 4 finish before 1 and 5. If you use locking, connections 2, 3, and 4 probably do not finish before 1 or 5, but the total time should be about 40% faster.

`INSERT`, `UPDATE`, and `DELETE` operations are very fast in MySQL, but you can obtain better overall performance by adding locks around everything that does more than about five successive inserts or updates. If you do very many successive inserts, you could do a `LOCK TABLES` followed by an `UNLOCK TABLES` once in a while (each 1,000 rows or so) to permit other threads to access table. This would still result in a nice performance gain.

`INSERT` is still much slower for loading data than `LOAD DATA INFILE`, even when using the strategies just outlined.

- To increase performance for `MyISAM` tables, for both `LOAD DATA INFILE` and `INSERT`, enlarge the key cache by increasing the `key_buffer_size` [492] system variable. See [Section 8.11.2, “Tuning Server Parameters”](#).

### 8.6.3. Speed of `REPAIR TABLE` Statements

`REPAIR TABLE` for `MyISAM` tables is similar to using `myisamchk` for repair operations, and some of the same performance optimizations apply:

- `myisamchk` has variables that control memory allocation. You may be able to improve performance by setting these variables, as described in [Section 4.6.3.6, “myisamchk Memory Usage”](#).
- For `REPAIR TABLE`, the same principle applies, but because the repair is done by the server, you set server system variables instead of `myisamchk` variables. Also, in addition to setting memory-allocation variables, increasing the `myisam_max_sort_file_size` [511] system variable increases the likelihood that the repair will use the faster filesort method and avoid the slower repair by key cache method. Set the variable to the maximum file size for your system, after checking to be sure that there is enough free space to hold a copy of the table files. The free space must be available in the file system containing the original table files.

Suppose that a `myisamchk` table-repair operation is done using the following options to set its memory-allocation variables:

```
--key_buffer_size=128M --myisam_sort_buffer_size=256M
--read_buffer_size=64M --write_buffer_size=64M
```

Some of those `myisamchk` variables correspond to server system variables:

<code>myisamchk</code> Variable	System Variable
<code>key_buffer_size</code>	<code>key_buffer_size</code> [492]
<code>myisam_sort_buffer_size</code>	<code>myisam_sort_buffer_size</code> [513]
<code>read_buffer_size</code>	<code>read_buffer_size</code> [532]
<code>write_buffer_size</code>	none

Each of the server system variables can be set at runtime, and some of them (`myisam_sort_buffer_size` [513], `read_buffer_size` [532]) have a session value in addition to a global value. Setting a session value limits the effect of the change to your current session and does not affect other users. Changing a global-only variable (`key_buffer_size` [492], `myisam_max_sort_file_size` [511]) affects other users as well. For `key_buffer_size` [492], you must take into account that the buffer is shared with those users. For example, if you set the `myisamchk key_buffer_size` variable to 128MB, you could set the corresponding `key_buffer_size` [492] system variable larger than that (if it is not already set larger), to permit key buffer use by activity in other sessions. However, changing the global key buffer size invalidates the buffer, causing increased disk I/O and slowdown for other sessions. An alternative that avoids this problem is to use a separate key cache, assign to it the indexes from the table to be repaired, and deallocate it when the repair is complete. See [Section 8.9.2.2, “Multiple Key Caches”](#).

Based on the preceding remarks, a `REPAIR TABLE` operation can be done as follows to use settings similar to the `myisamchk` command. Here a separate 128MB key buffer is allocated and the file system is assumed to permit a file size of at least 100GB.

```
SET SESSION myisam_sort_buffer_size = 256*1024*1024;
SET SESSION read_buffer_size = 64*1024*1024;
SET GLOBAL myisam_max_sort_file_size = 100*1024*1024*1024;
SET GLOBAL repair_cache.key_buffer_size = 128*1024*1024;
CACHE INDEX tbl_name IN repair_cache;
LOAD INDEX INTO CACHE tbl_name;
REPAIR TABLE tbl_name ;
SET GLOBAL repair_cache.key_buffer_size = 0;
```

If you intend to change a global variable but want to do so only for the duration of a `REPAIR TABLE` operation to minimally affect other users, save its value in a user variable and restore it afterward. For example:

```
SET @old_myisam_sort_buffer_size = @@global.myisam_max_sort_file_size;
SET GLOBAL myisam_max_sort_file_size = 100*1024*1024*1024;
REPAIR TABLE tbl_name ;
SET GLOBAL myisam_max_sort_file_size = @old_myisam_max_sort_file_size;
```

The system variables that affect `REPAIR TABLE` can be set globally at server startup if you want the values to be in effect by default. For example, add these lines to the server `my.cnf` file:

```
[mysqld]
myisam_sort_buffer_size=256M
key_buffer_size=1G
myisam_max_sort_file_size=100G
```

These settings do not include `read_buffer_size` [532]. Setting `read_buffer_size` [532] globally to a large value does so for all sessions and can cause performance to suffer due to excessive memory allocation for a server with many simultaneous sessions.

## 8.7. Optimizing for MEMORY Tables

Consider using `MEMORY` tables for noncritical data that is accessed often, and is read-only or rarely updated. Benchmark your application against equivalent `InnoDB` or `MyISAM` tables under a realistic



workload, to confirm that any additional performance is worth the risk of losing data, or the overhead of copying data from a disk-based table at application start.

For best performance with `MEMORY` tables, examine the kinds of queries against each table, and specify the type to use for each associated index, either a B-tree index or a hash index. On the `CREATE INDEX` statement, use the clause `USING BTREE` or `USING HASH`. B-tree indexes are fast for queries that do greater-than or less-than comparisons through operators such as `>` or `BETWEEN`. Hash indexes are only fast for queries that look up single values through the `=` operator, or a restricted set of values through the `IN` operator. For why `USING BTREE` is often a better choice than the default `USING HASH`, see [Section 8.2.1.4, “How to Avoid Full Table Scans”](#). For implementation details of the different types of `MEMORY` indexes, see [Section 8.3.8, “Comparison of B-Tree and Hash Indexes”](#).

## 8.8. Understanding the Query Execution Plan

Depending on the details of your tables, columns, indexes, and the conditions in your `WHERE` clause, the MySQL optimizer considers many techniques to efficiently perform the lookups involved in an SQL query. A query on a huge table can be performed without reading all the rows; a join involving several tables can be performed without comparing every combination of rows. The set of operations that the optimizer chooses to perform the most efficient query is called the “query execution plan”, also known as the `EXPLAIN` plan. Your goals are to recognize the aspects of the `EXPLAIN` plan that indicate a query is optimized well, and to learn the SQL syntax and indexing techniques to improve the plan if you see some inefficient operations.

### 8.8.1. Optimizing Queries with `EXPLAIN`

The `EXPLAIN` statement can be used either as a way to obtain information about how MySQL executes a statement, or as a synonym for `DESCRIBE`:

- When you precede a statement with the keyword `EXPLAIN`, MySQL displays information from the optimizer about the query execution plan. That is, MySQL explains how it would process the statement, including information about how tables are joined and in which order. See [Section 8.8.2, “EXPLAIN Output Format”](#).
- `EXPLAIN EXTENDED` can be used to obtain additional information; see [Section 8.8.3, “EXPLAIN EXTENDED Output Format”](#).

As of MySQL 5.6.3, `EXPLAIN` provides information about `SELECT`, `DELETE`, `INSERT`, `REPLACE`, and `UPDATE` statements. Before MySQL 5.6.3, `EXPLAIN` provides information only about `SELECT` statements.

The following sections describe how to use `EXPLAIN` and `EXPLAIN EXTENDED` to obtain query execution plan information.

- `EXPLAIN PARTITIONS` is useful only when examining queries involving partitioned tables. For details, see [Section 17.3.5, “Obtaining Information About Partitions”](#).
- `EXPLAIN tbl_name` is synonymous with `DESCRIBE tbl_name` or `SHOW COLUMNS FROM tbl_name`. For information about `DESCRIBE` and `SHOW COLUMNS`, see [Section 13.8.1, “DESCRIBE Syntax”](#), and [Section 13.7.5.6, “SHOW COLUMNS Syntax”](#).

With the help of `EXPLAIN`, you can see where you should add indexes to tables so that the statement executes faster by using indexes to find rows. You can also use `EXPLAIN` to check whether the optimizer joins the tables in an optimal order. To give a hint to the optimizer to use a join order corresponding to the order in which the tables are named in a `SELECT` statement, begin the statement with `SELECT STRAIGHT_JOIN` rather than just `SELECT`. (See [Section 13.2.9, “SELECT Syntax”](#).)

The optimizer trace may sometimes provide information complementary to that of `EXPLAIN`. However, the optimizer trace format and content are subject to change between versions”. For details, see [MySQL Internals: Tracing the Optimizer](#).



If you have a problem with indexes not being used when you believe that they should be, run [ANALYZE TABLE](#) to update table statistics, such as cardinality of keys, that can affect the choices the optimizer makes. See [Section 13.7.2.1, “ANALYZE TABLE Syntax”](#).

## 8.8.2. EXPLAIN Output Format

The [EXPLAIN](#) statement provides information about the execution plan for a [SELECT](#) statement.

[EXPLAIN](#) returns a row of information for each table used in the [SELECT](#) statement. It lists the tables in the output in the order that MySQL would read them while processing the statement. MySQL resolves all joins using a nested-loop join method. This means that MySQL reads a row from the first table, and then finds a matching row in the second table, the third table, and so on. When all tables are processed, MySQL outputs the selected columns and backtracks through the table list until a table is found for which there are more matching rows. The next row is read from this table and the process continues with the next table.

When the [EXTENDED](#) keyword is used, [EXPLAIN](#) produces extra information that can be viewed by issuing a [SHOW WARNINGS](#) statement following the [EXPLAIN](#) statement. [EXPLAIN EXTENDED](#) also displays the [filtered](#) column. See [Section 8.8.3, “EXPLAIN EXTENDED Output Format”](#).



### Note

You cannot use the [EXTENDED](#) and [PARTITIONS](#) keywords together in the same [EXPLAIN](#) statement.

- [EXPLAIN Output Columns](#)
- [EXPLAIN Join Types](#)
- [EXPLAIN Extra Information](#)
- [EXPLAIN Output Interpretation](#)

### EXPLAIN Output Columns

This section describes the output columns produced by [EXPLAIN](#). Later sections provide additional information about the [type](#) and [Extra](#) columns.

Each output row from [EXPLAIN](#) provides information about one table. Each row contains the values summarized in [Table 8.1, “EXPLAIN Output Columns”](#), and described in more detail following the table.

**Table 8.1. EXPLAIN Output Columns**

Column	Meaning
<a href="#">id</a> [895]	The <a href="#">SELECT</a> identifier
<a href="#">select_type</a> [896]	The <a href="#">SELECT</a> type
<a href="#">table</a> [896]	The table for the output row
<a href="#">partitions</a> [896]	The matching partitions
<a href="#">type</a> [897]	The join type
<a href="#">possible_keys</a> [897]	The possible indexes to choose
<a href="#">key</a> [897]	The index actually chosen
<a href="#">key_len</a> [897]	The length of the chosen key
<a href="#">ref</a> [897]	The columns compared to the index
<a href="#">rows</a> [897]	Estimate of rows to be examined
<a href="#">filtered</a> [897]	Percentage of rows filtered by table condition
<a href="#">Extra</a> [898]	Additional information

- [id](#)

The `SELECT` identifier. This is the sequential number of the `SELECT` within the query. The value can be `NULL` if the row refers to the union result of other rows. In this case, the `table` column shows a value like `<unionM,N>` to indicate that the row refers to the union of the rows with `id` values of `M` and `N`.

- `select_type`

The type of `SELECT`, which can be any of those shown in the following table.

<code>select_type</code> Value	Meaning
<code>SIMPLE</code>	Simple <code>SELECT</code> (not using <code>UNION</code> or subqueries)
<code>PRIMARY</code>	Outermost <code>SELECT</code>
<code>UNION</code>	Second or later <code>SELECT</code> statement in a <code>UNION</code>
<code>DEPENDENT UNION</code>	Second or later <code>SELECT</code> statement in a <code>UNION</code> , dependent on outer query
<code>UNION RESULT</code>	Result of a <code>UNION</code> .
<code>SUBQUERY</code>	First <code>SELECT</code> in subquery
<code>DEPENDENT SUBQUERY</code>	First <code>SELECT</code> in subquery, dependent on outer query
<code>DERIVED</code>	Derived table <code>SELECT</code> (subquery in <code>FROM</code> clause)
<code>MATERIALIZED</code>	Materialized subquery
<code>UNCACHEABLE SUBQUERY</code>	A subquery for which the result cannot be cached and must be re-evaluated for each row of the outer query
<code>UNCACHEABLE UNION</code>	The second or later select in a <code>UNION</code> that belongs to an uncacheable subquery (see <code>UNCACHEABLE SUBQUERY</code> )

`DEPENDENT` typically signifies the use of a correlated subquery. See [Section 13.2.10.7, “Correlated Subqueries”](#).

`DEPENDENT SUBQUERY` evaluation differs from `UNCACHEABLE SUBQUERY` evaluation. For `DEPENDENT SUBQUERY`, the subquery is re-evaluated only once for each set of different values of the variables from its outer context. For `UNCACHEABLE SUBQUERY`, the subquery is re-evaluated for each row of the outer context.

Cacheability of subqueries differs from caching of query results in the query cache (which is described in [Section 8.9.3.1, “How the Query Cache Operates”](#)). Subquery caching occurs during query execution, whereas the query cache is used to store results only after query execution finishes.

- `table`

The name of the table to which the row of output refers. This can also be one of the following values:

- `<unionM,N>`: The row refers to the union of the rows with `id` values of `M` and `N`.
- `<derivedN>`: The row refers to the derived table result for the row with an `id` value of `N`. A derived table may result, for example, from a subquery in the `FROM` clause.
- `<subqueryN>`: The row refers to the result of a materialized subquery for the row with an `id` value of `N`. See [Section 8.13.16.2, “Optimizing Subqueries with Subquery Materialization”](#).

- `partitions`

The partitions from which records would be matched by the query. This column is displayed only if the `PARTITIONS` keyword is used. The value is `NULL` for nonpartitioned tables. See [Section 17.3.5, “Obtaining Information About Partitions”](#).

- `type`

The `type` join type. For descriptions of the different types, see [EXPLAIN Join Types](#).

- `possible_keys`

The `possible_keys` column indicates which indexes MySQL can choose from use to find the rows in this table. Note that this column is totally independent of the order of the tables as displayed in the output from `EXPLAIN`. That means that some of the keys in `possible_keys` might not be usable in practice with the generated table order.

If this column is `NULL`, there are no relevant indexes. In this case, you may be able to improve the performance of your query by examining the `WHERE` clause to check whether it refers to some column or columns that would be suitable for indexing. If so, create an appropriate index and check the query with `EXPLAIN` again. See [Section 13.1.6, “ALTER TABLE Syntax”](#).

To see what indexes a table has, use `SHOW INDEX FROM tbl_name`.

- `key`

The `key` column indicates the key (index) that MySQL actually decided to use. If MySQL decides to use one of the `possible_keys` indexes to look up rows, that index is listed as the key value.

It is possible that `key` will name an index that is not present in the `possible_keys` value. This can happen if none of the `possible_keys` indexes are suitable for looking up rows, but all the columns selected by the query are columns of some other index. That is, the named index covers the selected columns, so although it is not used to determine which rows to retrieve, an index scan is more efficient than a data row scan.

For `InnoDB`, a secondary index might cover the selected columns even if the query also selects the primary key because `InnoDB` stores the primary key value with each secondary index. If `key` is `NULL`, MySQL found no index to use for executing the query more efficiently.

To force MySQL to use or ignore an index listed in the `possible_keys` column, use `FORCE INDEX`, `USE INDEX`, or `IGNORE INDEX` in your query. See [Section 13.2.9.3, “Index Hint Syntax”](#).

For `MyISAM` tables, running `ANALYZE TABLE` helps the optimizer choose better indexes. For `MyISAM` tables, `myisamchk --analyze` does the same. See [Section 13.7.2.1, “ANALYZE TABLE Syntax”](#), and [Section 7.6, “MyISAM Table Maintenance and Crash Recovery”](#).

- `key_len`

The `key_len` column indicates the length of the key that MySQL decided to use. The length is `NULL` if the `key` column says `NULL`. Note that the value of `key_len` enables you to determine how many parts of a multiple-part key MySQL actually uses.

- `ref`

The `ref` column shows which columns or constants are compared to the index named in the `key` column to select rows from the table.

- `rows`

The `rows` column indicates the number of rows MySQL believes it must examine to execute the query.

For `InnoDB` tables, this number is an estimate, and may not always be exact.

- `filtered`

The `filtered` column indicates an estimated percentage of table rows that will be filtered by the table condition. That is, `rows` shows the estimated number of rows examined and `rows × filtered`

/ 100 shows the number of rows that will be joined with previous tables. This column is displayed if you use `EXPLAIN EXTENDED`.

- `Extra`

This column contains additional information about how MySQL resolves the query. For descriptions of the different values, see `EXPLAIN Extra Information`.

## EXPLAIN Join Types

The `type` column of `EXPLAIN` output describes how tables are joined. The following list describes the join types, ordered from the best type to the worst:

- `system` [898]

The table has only one row (= system table). This is a special case of the `const` [898] join type.

- `const` [898]

The table has at most one matching row, which is read at the start of the query. Because there is only one row, values from the column in this row can be regarded as constants by the rest of the optimizer. `const` [898] tables are very fast because they are read only once.

`const` [898] is used when you compare all parts of a `PRIMARY KEY` or `UNIQUE` index to constant values. In the following queries, `tbl_name` can be used as a `const` [898] table:

```
SELECT * FROM tbl_name WHERE primary_key=1;

SELECT * FROM tbl_name
WHERE primary_key_part1=1 AND primary_key_part2=2;
```

- `eq_ref` [898]

One row is read from this table for each combination of rows from the previous tables. Other than the `system` [898] and `const` [898] types, this is the best possible join type. It is used when all parts of an index are used by the join and the index is a `PRIMARY KEY` or `UNIQUE NOT NULL` index.

`eq_ref` [898] can be used for indexed columns that are compared using the `=` operator. The comparison value can be a constant or an expression that uses columns from tables that are read before this table. In the following examples, MySQL can use an `eq_ref` [898] join to process `ref_table`:

```
SELECT * FROM ref_table,other_table
WHERE ref_table.key_column=other_table.column;

SELECT * FROM ref_table,other_table
WHERE ref_table.key_column_part1=other_table.column
AND ref_table.key_column_part2=1;
```

- `ref` [898]

All rows with matching index values are read from this table for each combination of rows from the previous tables. `ref` [898] is used if the join uses only a leftmost prefix of the key or if the key is not a `PRIMARY KEY` or `UNIQUE` index (in other words, if the join cannot select a single row based on the key value). If the key that is used matches only a few rows, this is a good join type.

`ref` [898] can be used for indexed columns that are compared using the `=` or `<=>` operator. In the following examples, MySQL can use a `ref` [898] join to process `ref_table`:

```
SELECT * FROM ref_table WHERE key_column=expr;

SELECT * FROM ref_table,other_table
WHERE ref_table.key_column=other_table.column;

SELECT * FROM ref_table,other_table
```

```
WHERE ref_table.key_column_part1=other_table.column
AND ref_table.key_column_part2=1;
```

- [fulltext \[899\]](#)

The join is performed using a `FULLTEXT` index.

- [ref\\_or\\_null \[899\]](#)

This join type is like [ref \[898\]](#), but with the addition that MySQL does an extra search for rows that contain `NULL` values. This join type optimization is used most often in resolving subqueries. In the following examples, MySQL can use a [ref\\_or\\_null \[899\]](#) join to process *ref\_table*:

```
SELECT * FROM ref_table
WHERE key_column=expr OR key_column IS NULL;
```

See [Section 8.13.6, “IS NULL Optimization”](#).

- [index\\_merge \[899\]](#)

This join type indicates that the Index Merge optimization is used. In this case, the `key` column in the output row contains a list of indexes used, and `key_len` contains a list of the longest key parts for the indexes used. For more information, see [Section 8.13.2, “Index Merge Optimization”](#).

- [unique\\_subquery \[899\]](#)

This type replaces [ref \[898\]](#) for some `IN` subqueries of the following form:

```
value IN (SELECT primary_key FROM single_table WHERE some_expr)
```

[unique\\_subquery \[899\]](#) is just an index lookup function that replaces the subquery completely for better efficiency.

- [index\\_subquery \[899\]](#)

This join type is similar to [unique\\_subquery \[899\]](#). It replaces `IN` subqueries, but it works for nonunique indexes in subqueries of the following form:

```
value IN (SELECT key_column FROM single_table WHERE some_expr)
```

- [range \[899\]](#)

Only rows that are in a given range are retrieved, using an index to select the rows. The `key` column in the output row indicates which index is used. The `key_len` contains the longest key part that was used. The `ref` column is `NULL` for this type.

[range \[899\]](#) can be used when a key column is compared to a constant using any of the `= [1166]`, `<> [1167]`, `> [1167]`, `>= [1167]`, `< [1167]`, `<= [1167]`, `IS NULL [1168]`, `<=> [1167]`, `BETWEEN [1168]`, or `IN() [1169]` operators:

```
SELECT * FROM tbl_name
  WHERE key_column = 10;

SELECT * FROM tbl_name
  WHERE key_column BETWEEN 10 and 20;

SELECT * FROM tbl_name
  WHERE key_column IN (10,20,30);

SELECT * FROM tbl_name
  WHERE key_part1 = 10 AND key_part2 IN (10,20,30);
```

- [index \[899\]](#)

The `index` join type is the same as `ALL`, except that the index tree is scanned. This occurs two ways:

- If the index is a covering index for the queries and can be used to satisfy all data required from the table, only the index tree is scanned. In this case, the `Extra` column says `Using index`. An index-only scan usually is faster than `ALL` because the size of the index usually is smaller than the table data.
- A full table scan is performed using reads from the index to look up data rows in index order. `Uses index` does not appear in the `Extra` column.

MySQL can use this join type when the query uses only columns that are part of a single index.

- `ALL`

A full table scan is done for each combination of rows from the previous tables. This is normally not good if the table is the first table not marked `const` [898], and usually very bad in all other cases. Normally, you can avoid `ALL` by adding indexes that enable row retrieval from the table based on constant values or column values from earlier tables.

## EXPLAIN Extra Information

The `Extra` column of `EXPLAIN` output contains additional information about how MySQL resolves the query. The following list explains the values that can appear in this column. If you want to make your queries as fast as possible, look out for `Extra` values of `Using filesort` and `Using temporary`.

- `const row not found`

For a query such as `SELECT ... FROM tbl_name`, the table was empty.

- `Deleting all rows`

For `DELETE`, some storage engines (such as `MyISAM`) support a handler method that removes all table rows in a simple and fast way. This `Extra` value is displayed if the engine uses this optimization.

- `Distinct`

MySQL is looking for distinct values, so it stops searching for more rows for the current row combination after it has found the first matching row.

- `FirstMatch(tbl_name)`

The semi-join `FirstMatch` join shortcutting strategy is used for `tbl_name`.

- `Full scan on NULL key`

This occurs for subquery optimization as a fallback strategy when the optimizer cannot use an index-lookup access method.

- `Impossible HAVING`

The `HAVING` clause is always false and cannot select any rows.

- `Impossible WHERE`

The `WHERE` clause is always false and cannot select any rows.

- `Impossible WHERE noticed after reading const tables`

MySQL has read all `const` [898] (and `system` [898]) tables and notice that the `WHERE` clause is always false.

- `LooseScan(m..n)`

The semi-join LooseScan strategy is used. *m* and *n* are key part numbers.

- **Materialize, Scan**

Before MySQL 5.6.7, this indicates use of a single materialized temporary table. If `Scan` is present, no temporary table index is used for table reads. Otherwise, an index lookup is used. See also the `Start materialize` entry.

As of MySQL 5.6.7, materialization is indicated by rows with a `select_type` value of `MATERIALIZED` and rows with a `table` value of `<subqueryN>`.

- **No matching min/max row**

No row satisfies the condition for a query such as `SELECT MIN(...) FROM ... WHERE condition`.

- **no matching row in const table**

For a query with a join, there was an empty table or a table with no rows satisfying a unique index condition.

- **No matching rows after partition pruning**

For `DELETE` or `UPDATE`, the optimizer found nothing to delete or update after partition pruning. It is similar in meaning to `Impossible WHERE` for `SELECT` statements.

- **No tables used**

The query has no `FROM` clause, or has a `FROM DUAL` clause.

For `INSERT` or `REPLACE` statements, `EXPLAIN` displays this value when there is no `SELECT` part. For example, it appears for `EXPLAIN INSERT INTO t VALUES(10)` because that is equivalent to `EXPLAIN INSERT INTO t SELECT 10 FROM DUAL`.

- **Not exists**

MySQL was able to do a `LEFT JOIN` optimization on the query and does not examine more rows in this table for the previous row combination after it finds one row that matches the `LEFT JOIN` criteria. Here is an example of the type of query that can be optimized this way:

```
SELECT * FROM t1 LEFT JOIN t2 ON t1.id=t2.id
WHERE t2.id IS NULL;
```

Assume that `t2.id` is defined as `NOT NULL`. In this case, MySQL scans `t1` and looks up the rows in `t2` using the values of `t1.id`. If MySQL finds a matching row in `t2`, it knows that `t2.id` can never be `NULL`, and does not scan through the rest of the rows in `t2` that have the same `id` value. In other words, for each row in `t1`, MySQL needs to do only a single lookup in `t2`, regardless of how many rows actually match in `t2`.

- **Range checked for each record (index map: N)**

MySQL found no good index to use, but found that some of indexes might be used after column values from preceding tables are known. For each row combination in the preceding tables, MySQL checks whether it is possible to use a `range [899]` or `index_merge [899]` access method to retrieve rows. This is not very fast, but is faster than performing a join with no index at all. The applicability criteria are as described in [Section 8.13.1, “Range Optimization”](#), and [Section 8.13.2, “Index Merge Optimization”](#), with the exception that all column values for the preceding table are known and considered to be constants.

Indexes are numbered beginning with 1, in the same order as shown by `SHOW INDEX` for the table. The index map value *N* is a bitmask value that indicates which indexes are candidates. For example, a value of `0x19` (binary 11001) means that indexes 1, 4, and 5 will be considered.



- `Scanned N databases`

This indicates how many directory scans the server performs when processing a query for `INFORMATION_SCHEMA` tables, as described in [Section 8.2.4, “Optimizing `INFORMATION\_SCHEMA` Queries”](#). The value of `N` can be 0, 1, or `all`.

- `Select tables optimized away`

The query contained only aggregate functions (`MIN()` [1284], `MAX()` [1284]) that were all resolved using an index, or `COUNT(*)` [1283] for `MyISAM`, and no `GROUP BY` clause. The optimizer determined that only one row should be returned.

- `Skip_open_table`, `Open_frm_only`, `Open_trigger_only`, `Open_full_table`

These values indicate file-opening optimizations that apply to queries for `INFORMATION_SCHEMA` tables, as described in [Section 8.2.4, “Optimizing `INFORMATION\_SCHEMA` Queries”](#).

- `Skip_open_table`: Table files do not need to be opened. The information has already become available within the query by scanning the database directory.
- `Open_frm_only`: Only the table's `.frm` file need be opened.
- `Open_trigger_only`: Only the table's `.TRG` file need be opened.
- `Open_full_table`: The unoptimized information lookup. The `.frm`, `.MYD`, and `.MYI` files must be opened.

- `Start materialize`, `End materialize`, `Scan`

Before MySQL 5.6.7, this indicates use of multiple materialized temporary tables. If `Scan` is present, no temporary table index is used for table reads. Otherwise, an index lookup is used. See also the [Materialize](#) entry.

As of MySQL 5.6.7, materialization is indicated by rows with a `select_type` value of `MATERIALIZED` and rows with a `table` value of `<subqueryN>`.

- `Start temporary`, `End temporary`

This indicates temporary table use for the semi-join Duplicate Weedout strategy.

- `unique row not found`

For a query such as `SELECT ... FROM tbl_name`, no rows satisfy the condition for a `UNIQUE` index or `PRIMARY KEY` on the table.

- `Using filesort`

MySQL must do an extra pass to find out how to retrieve the rows in sorted order. The sort is done by going through all rows according to the join type and storing the sort key and pointer to the row for all rows that match the `WHERE` clause. The keys then are sorted and the rows are retrieved in sorted order. See [Section 8.13.13, “ORDER BY Optimization”](#).

- `Using index`

The column information is retrieved from the table using only information in the index tree without having to do an additional seek to read the actual row. This strategy can be used when the query uses only columns that are part of a single index.

If the `Extra` column also says `Using where`, it means the index is being used to perform lookups of key values. Without `Using where`, the optimizer may be reading the index to avoid reading data rows but not using it for lookups. For example, if the index is a covering index for the query, the optimizer may scan it without using it for lookups.

For InnoDB tables that have a user-defined clustered index, that index can be used even when `Using index` is absent from the `Extra` column. This is the case if `type` is `index` [899] and `key` is `PRIMARY`.

- `Using index condition`

Tables are read by accessing index tuples and testing them first to determine whether to read full table rows. In this way, index information is used to defer (“push down”) reading full table rows unless it is necessary. See [Section 8.13.4, “Index Condition Pushdown Optimization”](#).

- `Using index for group-by`

Similar to the `Using index` table access method, `Using index for group-by` indicates that MySQL found an index that can be used to retrieve all columns of a `GROUP BY` or `DISTINCT` query without any extra disk access to the actual table. Additionally, the index is used in the most efficient way so that for each group, only a few index entries are read. For details, see [Section 8.13.14, “GROUP BY Optimization”](#).

- `Using join buffer (Block Nested Loop)`, `Using join buffer (Batched Key Access)`

Tables from earlier joins are read in portions into the join buffer, and then their rows are used from the buffer to perform the join with the current table. `(Block Nested Loop)` indicates use of the Block Nested-Loop algorithm and `(Batched Key Access)` indicates use of the Batched Key Access algorithm. That is, the keys from the table on the preceding line of the `EXPLAIN` output will be buffered, and the matching rows will be fetched in batches from the table represented by the line in which `Using join buffer` appears.

- `Using MRR`

Tables are read using the Multi-Range Read optimization strategy. See [Section 8.13.11, “Multi-Range Read Optimization”](#).

- `Using sort_union(...)`, `Using union(...)`, `Using intersect(...)`

These indicate how index scans are merged for the `index_merge` [899] join type. See [Section 8.13.2, “Index Merge Optimization”](#).

- `Using temporary`

To resolve the query, MySQL needs to create a temporary table to hold the result. This typically happens if the query contains `GROUP BY` and `ORDER BY` clauses that list columns differently.

- `Using where`

A `WHERE` clause is used to restrict which rows to match against the next table or send to the client. Unless you specifically intend to fetch or examine all rows from the table, you may have something wrong in your query if the `Extra` value is not `Using where` and the table join type is `ALL` or `index` [899].

- `Using where with pushed condition`

This item applies to `NDB` tables *only*. It means that MySQL Cluster is using the Condition Pushdown optimization to improve the efficiency of a direct comparison between a nonindexed column and a constant. In such cases, the condition is “pushed down” to the cluster’s data nodes and is evaluated on all data nodes simultaneously. This eliminates the need to send nonmatching rows over the network, and can speed up such queries by a factor of 5 to 10 times over cases where Condition Pushdown could be but is not used. For more information, see [Section 8.13.3, “Engine Condition Pushdown Optimization”](#).

## EXPLAIN Output Interpretation

You can get a good indication of how good a join is by taking the product of the values in the `rows` column of the `EXPLAIN` output. This should tell you roughly how many rows MySQL must examine to execute the query. If you restrict queries with the `max_join_size` [505] system variable, this row product also is used to determine which multiple-table `SELECT` statements to execute and which to abort. See [Section 8.11.2, “Tuning Server Parameters”](#).

The following example shows how a multiple-table join can be optimized progressively based on the information provided by `EXPLAIN`.

Suppose that you have the `SELECT` statement shown here and that you plan to examine it using `EXPLAIN`:

```
EXPLAIN SELECT tt.TicketNumber, tt.TimeIn,
              tt.ProjectReference, tt.EstimatedShipDate,
              tt.ActualShipDate, tt.ClientID,
              tt.ServiceCodes, tt.RepetitiveID,
              tt.CurrentProcess, tt.CurrentDPPerson,
              tt.RecordVolume, tt.DPPrinted, et.COUNTRY,
              et_1.COUNTRY, do.CUSTNAME
FROM tt, et, et AS et_1, do
WHERE tt.SubmitTime IS NULL
      AND tt.ActualPC = et.EMPLOYID
      AND tt.AssignedPC = et_1.EMPLOYID
      AND tt.ClientID = do.CUSTNMBR;
```

For this example, make the following assumptions:

- The columns being compared have been declared as follows.

Table	Column	Data Type
tt	ActualPC	CHAR(10)
tt	AssignedPC	CHAR(10)
tt	ClientID	CHAR(10)
et	EMPLOYID	CHAR(15)
do	CUSTNMBR	CHAR(15)

- The tables have the following indexes.

Table	Index
tt	ActualPC
tt	AssignedPC
tt	ClientID
et	EMPLOYID (primary key)
do	CUSTNMBR (primary key)

- The `tt.ActualPC` values are not evenly distributed.

Initially, before any optimizations have been performed, the `EXPLAIN` statement produces the following information:

```
table type possible_keys key key_len ref rows Extra
et ALL PRIMARY NULL NULL NULL 74
do ALL PRIMARY NULL NULL NULL 2135
et_1 ALL PRIMARY NULL NULL NULL 74
tt ALL AssignedPC, NULL NULL NULL 3872
    ClientID,
    ActualPC
    Range checked for each record (index map: 0x23)
```

Because `type` is `ALL` for each table, this output indicates that MySQL is generating a Cartesian product of all the tables; that is, every combination of rows. This takes quite a long time, because the product of the number of rows in each table must be examined. For the case at hand, this product is  $74 \times 2135 \times 74 \times 3872 = 45,268,558,720$  rows. If the tables were bigger, you can only imagine how long it would take.

One problem here is that MySQL can use indexes on columns more efficiently if they are declared as the same type and size. In this context, `VARCHAR` and `CHAR` are considered the same if they are declared as the same size. `tt.ActualPC` is declared as `CHAR(10)` and `et.EMPLOYID` is `CHAR(15)`, so there is a length mismatch.

To fix this disparity between column lengths, use `ALTER TABLE` to lengthen `ActualPC` from 10 characters to 15 characters:

```
mysql> ALTER TABLE tt MODIFY ActualPC VARCHAR(15);
```

Now `tt.ActualPC` and `et.EMPLOYID` are both `VARCHAR(15)`. Executing the `EXPLAIN` statement again produces this result:

table	type	possible_keys	key	key_len	ref	rows	Extra	
tt	ALL	AssignedPC, ClientID, ActualPC	NULL	NULL	NULL	3872	Using where	
do	ALL	PRIMARY	NULL	NULL	NULL	2135		
		Range checked for each record (index map: 0x1)						
et_1	ALL	PRIMARY	NULL	NULL	NULL	74		
		Range checked for each record (index map: 0x1)						
et	eq_ref	PRIMARY	PRIMARY	15	tt.ActualPC	1		

This is not perfect, but is much better: The product of the `rows` values is less by a factor of 74. This version executes in a couple of seconds.

A second alteration can be made to eliminate the column length mismatches for the `tt.AssignedPC = et_1.EMPLOYID` and `tt.ClientID = do.CUSTNMBR` comparisons:

```
mysql> ALTER TABLE tt MODIFY AssignedPC VARCHAR(15),
->      MODIFY ClientID VARCHAR(15);
```

After that modification, `EXPLAIN` produces the output shown here:

table	type	possible_keys	key	key_len	ref	rows	Extra
et	ALL	PRIMARY	NULL	NULL	NULL	74	
tt	ref	AssignedPC, ClientID, ActualPC	ActualPC	15	et.EMPLOYID	52	Using where
et_1	eq_ref	PRIMARY	PRIMARY	15	tt.AssignedPC	1	
do	eq_ref	PRIMARY	PRIMARY	15	tt.ClientID	1	

At this point, the query is optimized almost as well as possible. The remaining problem is that, by default, MySQL assumes that values in the `tt.ActualPC` column are evenly distributed, and that is not the case for the `tt` table. Fortunately, it is easy to tell MySQL to analyze the key distribution:

```
mysql> ANALYZE TABLE tt;
```

With the additional index information, the join is perfect and `EXPLAIN` produces this result:

table	type	possible_keys	key	key_len	ref	rows	Extra
tt	ALL	AssignedPC, ClientID, ActualPC	NULL	NULL	NULL	3872	Using where
et	eq_ref	PRIMARY	PRIMARY	15	tt.ActualPC	1	
et_1	eq_ref	PRIMARY	PRIMARY	15	tt.AssignedPC	1	
do	eq_ref	PRIMARY	PRIMARY	15	tt.ClientID	1	

Note that the `rows` column in the output from `EXPLAIN` is an educated guess from the MySQL join optimizer. Check whether the numbers are even close to the truth by comparing the `rows` product with the actual number of rows that the query returns. If the numbers are quite different, you might get better

performance by using `STRAIGHT_JOIN` in your `SELECT` statement and trying to list the tables in a different order in the `FROM` clause.

It is possible in some cases to execute statements that modify data when `EXPLAIN SELECT` is used with a subquery; for more information, see [Section 13.2.10.8, “Subqueries in the FROM Clause”](#).

### 8.8.3. EXPLAIN EXTENDED Output Format

When `EXPLAIN` is used with the `EXTENDED` keyword, the output includes a `filtered` column not otherwise displayed. This column indicates the estimated percentage of table rows that will be filtered by the table condition. In addition, the statement produces extra information that can be viewed by issuing a `SHOW WARNINGS` statement following the `EXPLAIN` statement. The `Message` value in `SHOW WARNINGS` output displays how the optimizer qualifies table and column names in the `SELECT` statement, what the `SELECT` looks like after the application of rewriting and optimization rules, and possibly other notes about the optimization process. Here is an example:

```
mysql> EXPLAIN EXTENDED
-> SELECT t1.a, t1.a IN (SELECT t2.a FROM t2) FROM t1\G
***** 1. row *****
      id: 1
  select_type: PRIMARY
      table: t1
      type: index
possible_keys: NULL
      key: PRIMARY
     key_len: 4
      ref: NULL
      rows: 4
  filtered: 100.00
     Extra: Using index
***** 2. row *****
      id: 2
  select_type: SUBQUERY
      table: t2
      type: index
possible_keys: a
      key: a
     key_len: 5
      ref: NULL
      rows: 3
  filtered: 100.00
     Extra: Using index
2 rows in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Note
Code: 1003
Message: /* select#1 */ select `test`.`t1`.`a` AS `a`,
<in_optimizer>(`test`.`t1`.`a`,`test`.`t1`.`a` in
( <materialize> (/* select#2 */ select `test`.`t2`.`a`
from `test`.`t2` where 1 having 1 ),
<primary_index_lookup>(`test`.`t1`.`a` in
<temporary table> on <auto_key>
where ((`test`.`t1`.`a` = `materialized-subquery`.`a`)))) AS `t1.a`
IN (SELECT t2.a FROM t2)` from `test`.`t1`
1 row in set (0.00 sec)
```

As of MySQL 5.6.3, `EXPLAIN EXTENDED` can be used with `SELECT`, `DELETE`, `INSERT`, `REPLACE`, and `UPDATE` statements. However, the following `SHOW WARNINGS` statement displays a nonempty result only for `SELECT` statements. Before MySQL 5.6.3, `EXPLAIN EXTENDED` can be used only with `SELECT` statements.

Because the statement displayed by `SHOW WARNINGS` may contain special markers to provide information about query rewriting or optimizer actions, the statement is not necessarily valid SQL and is not intended to be executed. The output may also include rows with `Message` values that provide additional non-SQL explanatory notes about actions taken by the optimizer.

The following list describes special markers that can appear in `EXTENDED` output displayed by `SHOW WARNINGS`:

- `<auto_key>`  
An automatically generated key for a temporary table.
- `<cache>(expr)`  
The expression (such as a scalar subquery) is executed once and the resulting value is saved in memory for later use. For results consisting of multiple values, a temporary table may be created and you will see `<temporary table>` instead.
- `<exists>(query fragment)`  
The subquery predicate is converted to an `EXISTS` predicate and the subquery is transformed so that it can be used together with the `EXISTS` predicate.
- `<in_optimizer>(query fragment)`  
This is an internal optimizer object with no user significance.
- `<index_lookup>(query fragment)`  
The query fragment is processed using an index lookup to find qualifying rows.
- `<if>(condition, expr1, expr2)`  
If the condition is true, evaluate to `expr1`, otherwise `expr2`.
- `<is_not_null_test>(expr)`  
A test to verify that the expression does not evaluate to `NULL`.
- `<materialize>(query fragment)`  
Subquery materialization is used.
- ``materialized-subquery`.col_name, `materialized subselect`.col_name`  
A reference to the column `col_name` in an internal temporary table materialized to hold the result from evaluating a subquery.
- `<primary_index_lookup>(query fragment)`  
The query fragment is processed using a primary key lookup to find qualifying rows.
- `<ref_null_helper>(expr)`  
This is an internal optimizer object with no user significance.
- `/* select#N */ select_stmt`  
The `SELECT` is associated with the row in non-`EXTENDED EXPLAIN` output that has an `id` value of `N`.
- `outer_tables semi join (inner_tables)`  
A semi-join operation. `inner_tables` shows the tables that were not pulled out. See [Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”](#).
- `<temporary table>`  
This represents an internal temporary table created to cache an intermediate result.

## 8.8.4. Estimating Query Performance

In most cases, you can estimate query performance by counting disk seeks. For small tables, you can usually find a row in one disk seek (because the index is probably cached). For bigger tables, you can estimate that, using B-tree indexes, you need this many seeks to find a row:  $\log(\text{row\_count}) / \log(\text{index\_block\_length} / 3 * 2 / (\text{index\_length} + \text{data\_pointer\_length})) + 1$ .

In MySQL, an index block is usually 1,024 bytes and the data pointer is usually four bytes. For a 500,000-row table with a key value length of three bytes (the size of `MEDIUMINT`), the formula indicates  $\log(500,000) / \log(1024 / 3 * 2 / (3 + 4)) + 1 = 4$  seeks.

This index would require storage of about  $500,000 * 7 * 3/2 = 5.2\text{MB}$  (assuming a typical index buffer fill ratio of 2/3), so you probably have much of the index in memory and so need only one or two calls to read data to find the row.

For writes, however, you need four seek requests to find where to place a new index value and normally two seeks to update the index and write the row.

Note that the preceding discussion does not mean that your application performance slowly degenerates by  $\log N$ . As long as everything is cached by the OS or the MySQL server, things become only marginally slower as the table gets bigger. After the data gets too big to be cached, things start to go much slower until your applications are bound only by disk seeks (which increase by  $\log N$ ). To avoid this, increase the key cache size as the data grows. For `MyISAM` tables, the key cache size is controlled by the `key_buffer_size` [492] system variable. See Section 8.11.2, “Tuning Server Parameters”.

## 8.8.5. Controlling the Query Optimizer

MySQL provides optimizer control through system variables that affect how query plans are evaluated and which switchable optimizations are enabled.

### 8.8.5.1. Controlling Query Plan Evaluation

The task of the query optimizer is to find an optimal plan for executing an SQL query. Because the difference in performance between “good” and “bad” plans can be orders of magnitude (that is, seconds versus hours or even days), most query optimizers, including that of MySQL, perform a more or less exhaustive search for an optimal plan among all possible query evaluation plans. For join queries, the number of possible plans investigated by the MySQL optimizer grows exponentially with the number of tables referenced in a query. For small numbers of tables (typically less than 7 to 10) this is not a problem. However, when larger queries are submitted, the time spent in query optimization may easily become the major bottleneck in the server's performance.

A more flexible method for query optimization enables the user to control how exhaustive the optimizer is in its search for an optimal query evaluation plan. The general idea is that the fewer plans that are investigated by the optimizer, the less time it spends in compiling a query. On the other hand, because the optimizer skips some plans, it may miss finding an optimal plan.

The behavior of the optimizer with respect to the number of plans it evaluates can be controlled using two system variables:

- The `optimizer_prune_level` [519] variable tells the optimizer to skip certain plans based on estimates of the number of rows accessed for each table. Our experience shows that this kind of “educated guess” rarely misses optimal plans, and may dramatically reduce query compilation times. That is why this option is on (`optimizer_prune_level=1`) by default. However, if you believe that the optimizer missed a better query plan, this option can be switched off (`optimizer_prune_level=0`) with the risk that query compilation may take much longer. Note that, even with the use of this heuristic, the optimizer still explores a roughly exponential number of plans.
- The `optimizer_search_depth` [520] variable tells how far into the “future” of each incomplete plan the optimizer should look to evaluate whether it should be expanded further. Smaller values



of `optimizer_search_depth` [520] may result in orders of magnitude smaller query compilation times. For example, queries with 12, 13, or more tables may easily require hours and even days to compile if `optimizer_search_depth` [520] is close to the number of tables in the query. At the same time, if compiled with `optimizer_search_depth` [520] equal to 3 or 4, the optimizer may compile in less than a minute for the same query. If you are unsure of what a reasonable value is for `optimizer_search_depth` [520], this variable can be set to 0 to tell the optimizer to determine the value automatically.

### 8.8.5.2. Controlling Switchable Optimizations

The `optimizer_switch` [520] system variable enables control over optimizer behavior. Its value is a set of flags, each of which has a value of `on` or `off` to indicate whether the corresponding optimizer behavior is enabled or disabled. This variable has global and session values and can be changed at runtime. The global default can be set at server startup.

To see the current set of optimizer flags, select the variable value:

```
mysql> SELECT @@optimizer_switch\G
***** 1. row *****
@@optimizer_switch: index_merge=on,index_merge_union=on,
                    index_merge_sort_union=on,
                    index_merge_intersection=on,
                    engine_condition_pushdown=on,
                    index_condition_pushdown=on,
                    mrr=on,mrr_cost_based=on,
                    block_nested_loop=on,batched_key_access=off,
                    materialization=on,semi_join=on,loosescan=on,
                    firstmatch=on,
                    subquery_materialization_cost_based=on,
                    use_index_extensions=on
```

To change the value of `optimizer_switch` [520], assign a value consisting of a comma-separated list of one or more commands:

```
SET [GLOBAL|SESSION] optimizer_switch='command[,command]...';
```

Each `command` value should have one of the forms shown in the following table.

Command Syntax	Meaning
<code>default</code>	Reset every optimization to its default value
<code>opt_name=default</code>	Set the named optimization to its default value
<code>opt_name=off</code>	Disable the named optimization
<code>opt_name=on</code>	Enable the named optimization

The order of the commands in the value does not matter, although the `default` command is executed first if present. Setting an `opt_name` flag to `default` sets it to whichever of `on` or `off` is its default value. Specifying any given `opt_name` more than once in the value is not permitted and causes an error. Any errors in the value cause the assignment to fail with an error, leaving the value of `optimizer_switch` [520] unchanged.

The following table lists the permissible `opt_name` flag names, grouped by optimization strategy.

Optimization	Flag Name	Meaning
Batched Key Access	<code>batched_key_access</code>	Controls use of BKA join algorithm
Block Nested-Loop	<code>block_nested_loop</code>	Controls use of BNL join algorithm
Engine Condition Pushdown	<code>engine_condition_pushdown</code>	Controls engine condition pushdown
Index Condition Pushdown	<code>index_condition_pushdown</code>	Controls index condition pushdown

Optimization	Flag Name	Meaning
Index Extensions	<code>use_index_extensions</code>	Controls use of index extensions
Index Merge	<code>index_merge</code>	Controls all Index Merge optimizations
	<code>index_merge_intersection</code>	Controls the Index Merge Intersection Access optimization
	<code>index_merge_sort_union</code>	Controls the Index Merge Sort-Union Access optimization
	<code>index_merge_union</code>	Controls the Index Merge Union Access optimization
Multi-Range Read	<code>mrr</code>	Controls the Multi-Range Read strategy
	<code>mrr_cost_based</code>	Controls use of cost-based MRR if <code>mrr=on</code>
Semi-join	<code>semi_join</code>	Controls all semi-join strategies
	<code>firstmatch</code>	Controls the semi-join FirstMatch strategy
	<code>loosescan</code>	Controls the semi-join LooseScan strategy (not to be confused with LooseScan for <code>GROUP BY</code> )
Subquery materialization	<code>materialization</code>	Controls materialization (including semi-join materialization)
	<code>subquery_materialization_cost_based</code>	Used cost-based materialization choice

The `block_nested_loop` and `batched_key_access` flags were added in MySQL 5.6.3. For `batched_key_access` to have any effect when set to `on`, the `mrr` flag must also be `on`. Currently, the cost estimation for MRR is too pessimistic. Hence, it is also necessary for `mrr_cost_based` to be `off` for BKA to be used.

`semi_join`, `firstmatch`, `loosescan`, and `materialization` flags were added in MySQL 5.6.5 to enable control over semi-join and subquery materialization strategies. The `semi_join` flag controls whether semi-joins are used. If it is set to `on`, the `firstmatch` and `loosescan` flags enable finer control over the permitted semi-join strategies. The `materialization` flag controls whether subquery materialization is used. If `semi_join` and `materialization` are both `on`, semi-joins also use materialization where applicable. These flags are `on` by default.

The `subquery_materialization_cost_based` was added in MySQL 5.6.7 to enable control over the choice between subquery materialization and `IN -> EXISTS` subquery transformation. If the flag is `on` (the default), the optimizer performs a cost-based choice between subquery materialization and `IN -> EXISTS` subquery transformation if either method could be used. If the flag is `off`, the optimizer chooses subquery materialization over `IN -> EXISTS` subquery transformation, which was the previous behavior before MySQL 5.6.7.

For more information about individual optimization strategies, see the following sections:

- [Section 8.13.12, “Block Nested-Loop and Batched Key Access Joins”](#)
- [Section 8.13.3, “Engine Condition Pushdown Optimization”](#)
- [Section 8.13.5, “Use of Index Extensions”](#)
- [Section 8.13.4, “Index Condition Pushdown Optimization”](#)
- [Section 8.13.2, “Index Merge Optimization”](#)
- [Section 8.13.11, “Multi-Range Read Optimization”](#)
- [Section 8.13.16, “Subquery Optimization”](#)

When you assign a value to `optimizer_switch` [520], flags that are not mentioned keep their current values. This makes it possible to enable or disable specific optimizer behaviors in a single statement without affecting other behaviors. The statement does not depend on what other optimizer flags exist and what their values are. Suppose that all Index Merge optimizations are enabled:

```
mysql> SELECT @@optimizer_switch\G
***** 1. row *****
@@optimizer_switch: index_merge=on,index_merge_union=on,
                    index_merge_sort_union=on,
                    index_merge_intersection=on,
                    engine_condition_pushdown=on,
                    index_condition_pushdown=on,
                    mrr=on,mrr_cost_based=on,
                    block_nested_loop=on,batched_key_access=off,
                    materialization=on,semijoin=on,loosescan=on,
                    firstmatch=on,
                    subquery_materialization_cost_based=on,
                    use_index_extensions=on
```

If the server is using the Index Merge Union or Index Merge Sort-Union access methods for certain queries and you want to check whether the optimizer will perform better without them, set the variable value like this:

```
mysql> SET optimizer_switch='index_merge_union=off,index_merge_sort_union=off';

mysql> SELECT @@optimizer_switch\G
***** 1. row *****
@@optimizer_switch: index_merge=on,index_merge_union=off,
                    index_merge_sort_union=off,
                    index_merge_intersection=on,
                    engine_condition_pushdown=on,
                    index_condition_pushdown=on,
                    mrr=on,mrr_cost_based=on,
                    block_nested_loop=on,batched_key_access=off,
                    materialization=on,semijoin=on,loosescan=on,
                    firstmatch=on,
                    subquery_materialization_cost_based=on,
                    use_index_extensions=on
```

## 8.9. Buffering and Caching

MySQL uses several strategies that cache information in memory buffers to increase performance.

### 8.9.1. The InnoDB Buffer Pool

InnoDB maintains a storage area called the **buffer pool** for caching data and indexes in memory. Knowing how the InnoDB buffer pool works, and taking advantage of it to keep frequently accessed data in memory, is an important aspect of MySQL tuning.

#### Guidelines

Ideally, you set the size of the buffer pool to as large a value as practical, leaving enough memory for other processes on the server to run without excessive paging. The larger the buffer pool, the more InnoDB acts like an in-memory database, reading data from disk once and then accessing the data from memory during subsequent reads. The buffer pool even caches data changed by insert and update operations, so that disk writes can be grouped together for better performance.

Depending on the typical workload on your system, you might adjust the proportions of the parts within the buffer pool. You can tune the way the buffer pool chooses which blocks to cache once it fills up, to keep frequently accessed data in memory despite sudden spikes of activity for operations such as backups or reporting.

With 64-bit systems with large memory sizes, you can split the buffer pool into multiple parts, to minimize contention for the memory structures among concurrent operations. For details, see [Section 14.2.4.2.27, “Improvements to Performance from Multiple Buffer Pools”](#).

## Internal Details

`InnoDB` manages the pool as a list, using a variation of the least recently used (LRU) algorithm. When room is needed to add a new block to the pool, `InnoDB` evicts the least recently used block and adds the new block to the middle of the list. This “midpoint insertion strategy” treats the list as two sublists:

- At the head, a sublist of “new” (or “young”) blocks that were accessed recently.
- At the tail, a sublist of “old” blocks that were accessed less recently.

This algorithm keeps blocks that are heavily used by queries in the new sublist. The old sublist contains less-used blocks; these blocks are candidates for [eviction](#).

The LRU algorithm operates as follows by default:

- 3/8 of the buffer pool is devoted to the old sublist.
- The midpoint of the list is the boundary where the tail of the new sublist meets the head of the old sublist.
- When `InnoDB` reads a block into the buffer pool, it initially inserts it at the midpoint (the head of the old sublist). A block can be read in because it is required for a user-specified operation such as an SQL query, or as part of a [read-ahead](#) operation performed automatically by `InnoDB`.
- Accessing to a block in the old sublist makes it “young”, moving it to the head of the buffer pool (the head of the new sublist). If the block was read in because it was required, the first access occurs immediately and the block is made young. If the block was read in due to read-ahead, the first access does not occur immediately (and might not occur at all before the block is evicted).
- As the database operates, blocks in the buffer pool that are not accessed “age” by moving toward the tail of the list. Blocks in both the new and old sublists age as other blocks are made new. Blocks in the old sublist also age as blocks are inserted at the midpoint. Eventually, a block that remains unused for long enough reaches the tail of the old sublist and is evicted.

By default, blocks read by queries immediately move into the new sublist, meaning they will stay in the buffer pool for a long time. A table scan (such as performed for a `mysqldump` operation, or a `SELECT` statement with no `WHERE` clause) can bring a large amount of data into the buffer pool and evict an equivalent amount of older data, even if the new data is never used again. Similarly, blocks that are loaded by the read-ahead background thread and then accessed only once move to the head of the new list. These situations can push frequently used blocks to the old sublist, where they become subject to eviction.

## Configuration Options

Several `InnoDB` system variables control the size of the buffer pool and let you tune the LRU algorithm:

- [innodb\\_buffer\\_pool\\_size \[1708\]](#)

Specifies the size of the buffer pool. If your buffer pool is small and you have sufficient memory, making the pool larger can improve performance by reducing the amount of disk I/O needed as queries access `InnoDB` tables.

- [innodb\\_buffer\\_pool\\_instances \[1706\]](#)

Divides the buffer pool into a user-specified number of separate regions, each with its own LRU list and related data structures, to reduce contention during concurrent memory read and write operations. This option takes effect only when you set the [innodb\\_buffer\\_pool\\_size](#) to a size of 1 gigabyte or more. The total size you specify is divided among all the buffer pools. For best efficiency, specify a combination of [innodb\\_buffer\\_pool\\_instances \[1706\]](#) and [innodb\\_buffer\\_pool\\_size \[1708\]](#) so that each buffer pool instance is at least 1 gigabyte.

- [innodb\\_old\\_blocks\\_pct \[1738\]](#)

Specifies the approximate percentage of the buffer pool that InnoDB uses for the old block sublist. The range of values is 5 to 95. The default value is 37 (that is, 3/8 of the pool).

- [innodb\\_old\\_blocks\\_time \[1738\]](#)

Specifies how long in milliseconds (ms) a block inserted into the old sublist must stay there after its first access before it can be moved to the new sublist. The default value is 0: A block inserted into the old sublist moves immediately to the new sublist the first time it is accessed, no matter how soon after insertion the access occurs. If the value is greater than 0, blocks remain in the old sublist until an access occurs at least that many ms after the first access. For example, a value of 1000 causes blocks to stay in the old sublist for 1 second after the first access before they become eligible to move to the new sublist.

Setting [innodb\\_old\\_blocks\\_time \[1738\]](#) greater than 0 prevents one-time table scans from flooding the new sublist with blocks used only for the scan. Rows in a block read in for a scan are accessed many times in rapid succession, but the block is unused after that. If [innodb\\_old\\_blocks\\_time \[1738\]](#) is set to a value greater than time to process the block, the block remains in the “old” sublist and ages to the tail of the list to be evicted quickly. This way, blocks used only for a one-time scan do not act to the detriment of heavily used blocks in the new sublist.

[innodb\\_old\\_blocks\\_time \[1738\]](#) can be set at runtime, so you can change it temporarily while performing operations such as table scans and dumps:

```
SET GLOBAL innodb_old_blocks_time = 1000;... perform queries that scan tables ...
SET GLOBAL innodb_old_blocks_time = 0;
```

This strategy does not apply if your intent is to “warm up” the buffer pool by filling it with a table's content. For example, benchmark tests often perform a table or index scan at server startup, because that data would normally be in the buffer pool after a period of normal use. In this case, leave [innodb\\_old\\_blocks\\_time \[1738\]](#) set to 0, at least until the warmup phase is complete.

## Monitoring the Buffer Pool

The output from the InnoDB Standard Monitor contains several fields in the [BUFFER POOL AND MEMORY](#) section that pertain to operation of the buffer pool LRU algorithm:

- [Old database pages](#): The number of pages in the old sublist of the buffer pool.
- [Pages made young, not young](#): The number of old pages that were moved to the head of the buffer pool (the new sublist), and the number of pages that have remained in the old sublist without being made new.
- [young/s non-young/s](#): The number of accesses to old pages that have resulted in making them young or not. This metric differs from that of the previous item in two ways. First, it relates only to old pages. Second, it is based on number of accesses to pages and not the number of pages. (There can be multiple accesses to a given page, all of which are counted.)
- [young-making rate](#): Hits that cause blocks to move to the head of the buffer pool.
- [not](#): Hits that do not cause blocks to move to the head of the buffer pool (due to the delay not being met).

The [young-making](#) rate and [not](#) rate will not normally add up to the overall buffer pool hit rate. Hits for blocks in the old sublist cause them to move to the new sublist, but hits to blocks in the new sublist cause them to move to the head of the list only if they are a certain distance from the head.

The preceding information from the Monitor can help you make LRU tuning decisions:

- If you see very low [young/s](#) values when you do not have large scans going on, that indicates that you might need to either reduce the delay time, or increase the percentage of the buffer pool used for the old sublist. Increasing the percentage makes the old sublist larger, so blocks in that sublist

take longer to move to the tail and be evicted. This increases the likelihood that they will be accessed again and be made young.

- If you do not see a lot of `non-youngs/s` when you are doing large table scans (and lots of `youngs/s`), to tune your delay value to be larger.

For more information about InnoDB Monitors, see [Section 14.2.4.4, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”](#).

## 8.9.2. The MyISAM Key Cache

To minimize disk I/O, the MyISAM storage engine exploits a strategy that is used by many database management systems. It employs a cache mechanism to keep the most frequently accessed table blocks in memory:

- For index blocks, a special structure called the *key cache* (or *key buffer*) is maintained. The structure contains a number of block buffers where the most-used index blocks are placed.
- For data blocks, MySQL uses no special cache. Instead it relies on the native operating system file system cache.

This section first describes the basic operation of the MyISAM key cache. Then it discusses features that improve key cache performance and that enable you to better control cache operation:

- Multiple sessions can access the cache concurrently.
- You can set up multiple key caches and assign table indexes to specific caches.

To control the size of the key cache, use the `key_buffer_size` [492] system variable. If this variable is set equal to zero, no key cache is used. The key cache also is not used if the `key_buffer_size` [492] value is too small to allocate the minimal number of block buffers (8).

When the key cache is not operational, index files are accessed using only the native file system buffering provided by the operating system. (In other words, table index blocks are accessed using the same strategy as that employed for table data blocks.)

An index block is a contiguous unit of access to the MyISAM index files. Usually the size of an index block is equal to the size of nodes of the index B-tree. (Indexes are represented on disk using a B-tree data structure. Nodes at the bottom of the tree are leaf nodes. Nodes above the leaf nodes are nonleaf nodes.)

All block buffers in a key cache structure are the same size. This size can be equal to, greater than, or less than the size of a table index block. Usually one of these two values is a multiple of the other.

When data from any table index block must be accessed, the server first checks whether it is available in some block buffer of the key cache. If it is, the server accesses data in the key cache rather than on disk. That is, it reads from the cache or writes into it rather than reading from or writing to disk. Otherwise, the server chooses a cache block buffer containing a different table index block (or blocks) and replaces the data there by a copy of required table index block. As soon as the new index block is in the cache, the index data can be accessed.

If it happens that a block selected for replacement has been modified, the block is considered “dirty.” In this case, prior to being replaced, its contents are flushed to the table index from which it came.

Usually the server follows an *LRU* (*Least Recently Used*) strategy: When choosing a block for replacement, it selects the least recently used index block. To make this choice easier, the key cache module maintains all used blocks in a special list (*LRU chain*) ordered by time of use. When a block is accessed, it is the most recently used and is placed at the end of the list. When blocks need to be replaced, blocks at the beginning of the list are the least recently used and become the first candidates for eviction.

The InnoDB storage engine also uses an LRU algorithm, to manage its buffer pool. See [Section 8.9.1, “The InnoDB Buffer Pool”](#).



### 8.9.2.1. Shared Key Cache Access

Threads can access key cache buffers simultaneously, subject to the following conditions:

- A buffer that is not being updated can be accessed by multiple sessions.
- A buffer that is being updated causes sessions that need to use it to wait until the update is complete.
- Multiple sessions can initiate requests that result in cache block replacements, as long as they do not interfere with each other (that is, as long as they need different index blocks, and thus cause different cache blocks to be replaced).

Shared access to the key cache enables the server to improve throughput significantly.

### 8.9.2.2. Multiple Key Caches

Shared access to the key cache improves performance but does not eliminate contention among sessions entirely. They still compete for control structures that manage access to the key cache buffers. To reduce key cache access contention further, MySQL also provides multiple key caches. This feature enables you to assign different table indexes to different key caches.

Where there are multiple key caches, the server must know which cache to use when processing queries for a given `MyISAM` table. By default, all `MyISAM` table indexes are cached in the default key cache. To assign table indexes to a specific key cache, use the `CACHE INDEX` statement (see [Section 13.7.6.2, “CACHE INDEX Syntax”](#)). For example, the following statement assigns indexes from the tables `t1`, `t2`, and `t3` to the key cache named `hot_cache`:

```
mysql> CACHE INDEX t1, t2, t3 IN hot_cache;
```

Table	Op	Msg_type	Msg_text
test.t1	assign_to_keycache	status	OK
test.t2	assign_to_keycache	status	OK
test.t3	assign_to_keycache	status	OK

The key cache referred to in a `CACHE INDEX` statement can be created by setting its size with a `SET GLOBAL` parameter setting statement or by using server startup options. For example:

```
mysql> SET GLOBAL keycache1.key_buffer_size=128*1024;
```

To destroy a key cache, set its size to zero:

```
mysql> SET GLOBAL keycache1.key_buffer_size=0;
```

Note that you cannot destroy the default key cache. Any attempt to do this will be ignored:

```
mysql> SET GLOBAL key_buffer_size = 0;
```

```
mysql> SHOW VARIABLES LIKE 'key_buffer_size';
```

Variable_name	Value
key_buffer_size	8384512

Key cache variables are structured system variables that have a name and components. For `keycache1.key_buffer_size`, `keycache1` is the cache variable name and `key_buffer_size [492]` is the cache component. See [Section 5.1.5.1, “Structured System Variables”](#), for a description of the syntax used for referring to structured key cache system variables.

By default, table indexes are assigned to the main (default) key cache created at the server startup. When a key cache is destroyed, all indexes assigned to it are reassigned to the default key cache.



For a busy server, you can use a strategy that involves three key caches:

- A “hot” key cache that takes up 20% of the space allocated for all key caches. Use this for tables that are heavily used for searches but that are not updated.
- A “cold” key cache that takes up 20% of the space allocated for all key caches. Use this cache for medium-sized, intensively modified tables, such as temporary tables.
- A “warm” key cache that takes up 60% of the key cache space. Employ this as the default key cache, to be used by default for all other tables.

One reason the use of three key caches is beneficial is that access to one key cache structure does not block access to the others. Statements that access tables assigned to one cache do not compete with statements that access tables assigned to another cache. Performance gains occur for other reasons as well:

- The hot cache is used only for retrieval queries, so its contents are never modified. Consequently, whenever an index block needs to be pulled in from disk, the contents of the cache block chosen for replacement need not be flushed first.
- For an index assigned to the hot cache, if there are no queries requiring an index scan, there is a high probability that the index blocks corresponding to nonleaf nodes of the index B-tree remain in the cache.
- An update operation most frequently executed for temporary tables is performed much faster when the updated node is in the cache and need not be read in from disk first. If the size of the indexes of the temporary tables are comparable with the size of cold key cache, the probability is very high that the updated node is in the cache.

The `CACHE INDEX` statement sets up an association between a table and a key cache, but the association is lost each time the server restarts. If you want the association to take effect each time the server starts, one way to accomplish this is to use an option file: Include variable settings that configure your key caches, and an `init-file` option that names a file containing `CACHE INDEX` statements to be executed. For example:

```
key_buffer_size = 4G
hot_cache.key_buffer_size = 2G
cold_cache.key_buffer_size = 2G
init_file=/path/to/data-directory/mysql_d_init.sql
```

The statements in `mysql_d_init.sql` are executed each time the server starts. The file should contain one SQL statement per line. The following example assigns several tables each to `hot_cache` and `cold_cache`:

```
CACHE INDEX db1.t1, db1.t2, db2.t3 IN hot_cache
CACHE INDEX db1.t4, db2.t5, db2.t6 IN cold_cache
```

### 8.9.2.3. Midpoint Insertion Strategy

By default, the key cache management system uses a simple LRU strategy for choosing key cache blocks to be evicted, but it also supports a more sophisticated method called the *midpoint insertion strategy*.

When using the midpoint insertion strategy, the LRU chain is divided into two parts: a hot sublist and a warm sublist. The division point between two parts is not fixed, but the key cache management system takes care that the warm part is not “too short,” always containing at least `key_cache_division_limit` [494] percent of the key cache blocks. `key_cache_division_limit` [494] is a component of structured key cache variables, so its value is a parameter that can be set per cache.

When an index block is read from a table into the key cache, it is placed at the end of the warm sublist. After a certain number of hits (accesses of the block), it is promoted to the hot sublist. At present, the number of hits required to promote a block (3) is the same for all index blocks.

A block promoted into the hot sublist is placed at the end of the list. The block then circulates within this sublist. If the block stays at the beginning of the sublist for a long enough time, it is demoted to the warm sublist. This time is determined by the value of the `key_cache_age_threshold` [493] component of the key cache.

The threshold value prescribes that, for a key cache containing  $N$  blocks, the block at the beginning of the hot sublist not accessed within the last  $N * \text{key\_cache\_age\_threshold} / 100$  hits is to be moved to the beginning of the warm sublist. It then becomes the first candidate for eviction, because blocks for replacement always are taken from the beginning of the warm sublist.

The midpoint insertion strategy enables you to keep more-valued blocks always in the cache. If you prefer to use the plain LRU strategy, leave the `key_cache_division_limit` [494] value set to its default of 100.

The midpoint insertion strategy helps to improve performance when execution of a query that requires an index scan effectively pushes out of the cache all the index blocks corresponding to valuable high-level B-tree nodes. To avoid this, you must use a midpoint insertion strategy with the `key_cache_division_limit` [494] set to much less than 100. Then valuable frequently hit nodes are preserved in the hot sublist during an index scan operation as well.

#### 8.9.2.4. Index Preloading

If there are enough blocks in a key cache to hold blocks of an entire index, or at least the blocks corresponding to its nonleaf nodes, it makes sense to preload the key cache with index blocks before starting to use it. Preloading enables you to put the table index blocks into a key cache buffer in the most efficient way: by reading the index blocks from disk sequentially.

Without preloading, the blocks are still placed into the key cache as needed by queries. Although the blocks will stay in the cache, because there are enough buffers for all of them, they are fetched from disk in random order, and not sequentially.

To preload an index into a cache, use the `LOAD INDEX INTO CACHE` statement. For example, the following statement preloads nodes (index blocks) of indexes of the tables `t1` and `t2`:

```
mysql> LOAD INDEX INTO CACHE t1, t2 IGNORE LEAVES;
+-----+-----+-----+-----+
| Table | Op          | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.t1 | preload_keys | status   | OK       |
| test.t2 | preload_keys | status   | OK       |
+-----+-----+-----+-----+
```

The `IGNORE LEAVES` modifier causes only blocks for the nonleaf nodes of the index to be preloaded. Thus, the statement shown preloads all index blocks from `t1`, but only blocks for the nonleaf nodes from `t2`.

If an index has been assigned to a key cache using a `CACHE INDEX` statement, preloading places index blocks into that cache. Otherwise, the index is loaded into the default key cache.

#### 8.9.2.5. Key Cache Block Size

It is possible to specify the size of the block buffers for an individual key cache using the `key_cache_block_size` [493] variable. This permits tuning of the performance of I/O operations for index files.

The best performance for I/O operations is achieved when the size of read buffers is equal to the size of the native operating system I/O buffers. But setting the size of key nodes equal to the size of the I/O buffer does not always ensure the best overall performance. When reading the big leaf nodes, the server pulls in a lot of unnecessary data, effectively preventing reading other leaf nodes.

To control the size of blocks in the `.MYI` index file of MyISAM tables, use the `--myisam-block-size` [432] option at server startup.

### 8.9.2.6. Restructuring a Key Cache

A key cache can be restructured at any time by updating its parameter values. For example:

```
mysql> SET GLOBAL cold_cache.key_buffer_size=4*1024*1024;
```

If you assign to either the `key_buffer_size` [492] or `key_cache_block_size` [493] key cache component a value that differs from the component's current value, the server destroys the cache's old structure and creates a new one based on the new values. If the cache contains any dirty blocks, the server saves them to disk before destroying and re-creating the cache. Restructuring does not occur if you change other key cache parameters.

When restructuring a key cache, the server first flushes the contents of any dirty buffers to disk. After that, the cache contents become unavailable. However, restructuring does not block queries that need to use indexes assigned to the cache. Instead, the server directly accesses the table indexes using native file system caching. File system caching is not as efficient as using a key cache, so although queries execute, a slowdown can be anticipated. After the cache has been restructured, it becomes available again for caching indexes assigned to it, and the use of file system caching for the indexes ceases.

### 8.9.3. The MySQL Query Cache

The query cache stores the text of a `SELECT` statement together with the corresponding result that was sent to the client. If an identical statement is received later, the server retrieves the results from the query cache rather than parsing and executing the statement again. The query cache is shared among sessions, so a result set generated by one client can be sent in response to the same query issued by another client.

The query cache can be useful in an environment where you have tables that do not change very often and for which the server receives many identical queries. This is a typical situation for many Web servers that generate many dynamic pages based on database content.

The query cache does not return stale data. When tables are modified, any relevant entries in the query cache are flushed.



#### Note

The query cache does not work in an environment where you have multiple `mysqld` servers updating the same `MyISAM` tables.

The query cache is used for prepared statements under the conditions described in [Section 8.9.3.1, "How the Query Cache Operates"](#).



#### Note

As of MySQL 5.6.5, the query cache is not supported for partitioned tables, and is automatically disabled for queries involving partitioned tables. The query cache cannot be enabled for such queries. (Bug #53775)

Some performance data for the query cache follows. These results were generated by running the MySQL benchmark suite on a Linux Alpha 2x500MHz system with 2GB RAM and a 64MB query cache.

- If all the queries you are performing are simple (such as selecting a row from a table with one row), but still differ so that the queries cannot be cached, the overhead for having the query cache active is 13%. This could be regarded as the worst case scenario. In real life, queries tend to be much more complicated, so the overhead normally is significantly lower.
- Searches for a single row in a single-row table are 238% faster with the query cache than without it. This can be regarded as close to the minimum speedup to be expected for a query that is cached.

To disable the query cache at server startup, set the `query_cache_size` [528] system variable to 0. By disabling the query cache code, there is no noticeable overhead.

The query cache offers the potential for substantial performance improvement, but do not assume that it will do so under all circumstances. With some query cache configurations or server workloads, you might actually see a performance decrease:

- Be cautious about sizing the query cache excessively large, which increases the overhead required to maintain the cache, possibly beyond the benefit of enabling it. Sizes in tens of megabytes are usually beneficial. Sizes in the hundreds of megabytes might not be.
- Server workload has a significant effect on query cache efficiency. A query mix consisting almost entirely of a fixed set of `SELECT` statements is much more likely to benefit from enabling the cache than a mix in which frequent `INSERT` statements cause continual invalidation of results in the cache. In some cases, a workaround is to use the `SQL_NO_CACHE` option to prevent results from even entering the cache for `SELECT` statements that use frequently modified tables. (See [Section 8.9.3.2, “Query Cache `SELECT` Options”](#).)

To verify that enabling the query cache is beneficial, test the operation of your MySQL server with the cache enabled and disabled. Then retest periodically because query cache efficiency may change as server workload changes.

### 8.9.3.1. How the Query Cache Operates

This section describes how the query cache works when it is operational. [Section 8.9.3.3, “Query Cache Configuration”](#), describes how to control whether it is operational.

Incoming queries are compared to those in the query cache before parsing, so the following two queries are regarded as different by the query cache:

```
SELECT * FROM tbl_name
Select * from tbl_name
```

Queries must be *exactly* the same (byte for byte) to be seen as identical. In addition, query strings that are identical may be treated as different for other reasons. Queries that use different databases, different protocol versions, or different default character sets are considered different queries and are cached separately.

The cache is not used for queries of the following types:

- Queries that are a subquery of an outer query
- Queries executed within the body of a stored function, trigger, or event

Before a query result is fetched from the query cache, MySQL checks whether the user has `SELECT` privilege for all databases and tables involved. If this is not the case, the cached result is not used.

If a query result is returned from query cache, the server increments the `Qcache_hits` [595] status variable, not `Com_select`. See [Section 8.9.3.4, “Query Cache Status and Maintenance”](#).

If a table changes, all cached queries that use the table become invalid and are removed from the cache. This includes queries that use `MERGE` tables that map to the changed table. A table can be changed by many types of statements, such as `INSERT`, `UPDATE`, `DELETE`, `TRUNCATE TABLE`, `ALTER TABLE`, `DROP TABLE`, or `DROP DATABASE`.

The query cache also works within transactions when using `InnoDB` tables.

In MySQL 5.6, the result from a `SELECT` query on a view is cached.

The query cache works for `SELECT SQL_CALC_FOUND_ROWS . . .` queries and stores a value that is returned by a following `SELECT FOUND_ROWS()` query. `FOUND_ROWS()` [1268] returns the correct value even if the preceding query was fetched from the cache because the number of found rows is also stored in the cache. The `SELECT FOUND_ROWS()` query itself cannot be cached.

Prepared statements that are issued using the binary protocol using `mysql_stmt_prepare()` and `mysql_stmt_execute()` (see [Section 21.8.7, “C API Prepared Statements”](#)), are subject to

limitations on caching. Comparison with statements in the query cache is based on the text of the statement after expansion of ? parameter markers. The statement is compared only with other cached statements that were executed using the binary protocol. That is, for query cache purposes, prepared statements issued using the binary protocol are distinct from prepared statements issued using the text protocol (see [Section 13.5, “SQL Syntax for Prepared Statements”](#)).

A query cannot be cached if it contains any of the functions shown in the following table.

<a href="#">BENCHMARK() [1266]</a>	<a href="#">CONNECTION_ID() [1267]</a>	<a href="#">CONVERT_TZ() [1212]</a>
<a href="#">CURDATE() [1212]</a>	<a href="#">CURRENT_DATE() [1212]</a>	<a href="#">CURRENT_TIME() [1212]</a>
<a href="#">CURRENT_TIMESTAMP() [1213]</a>	<a href="#">CURTIME() [1213]</a>	<a href="#">DATABASE() [1268]</a>
<a href="#">ENCRYPT() [1262]</a> with one parameter	<a href="#">FOUND_ROWS() [1268]</a>	<a href="#">GET_LOCK() [1275]</a>
<a href="#">LAST_INSERT_ID() [1269]</a>	<a href="#">LOAD_FILE() [1181]</a>	<a href="#">MASTER_POS_WAIT() [1279]</a>
<a href="#">NOW() [1220]</a>	<a href="#">PASSWORD() [1263]</a>	<a href="#">RAND() [1206]</a>
<a href="#">RELEASE_LOCK() [1280]</a>	<a href="#">SLEEP() [1280]</a>	<a href="#">SYSDATE() [1223]</a>
<a href="#">UNIX_TIMESTAMP() [1226]</a> with no parameters	<a href="#">USER() [1273]</a>	<a href="#">UUID() [1280]</a>
<a href="#">UUID_SHORT() [1281]</a>		

A query also is not cached under these conditions:

- It refers to user-defined functions (UDFs) or stored functions.
- It refers to user variables or local stored program variables.
- It refers to tables in the `mysql`, `INFORMATION_SCHEMA`, or `performance_schema` database.
- (*MySQL 5.6.5 and later.*) It refers to any partitioned tables.
- It is of any of the following forms:

```
SELECT ... LOCK IN SHARE MODE
SELECT ... FOR UPDATE
SELECT ... INTO OUTFILE ...
SELECT ... INTO DUMPFILE ...
SELECT * FROM ... WHERE autoincrement_col IS NULL
```

The last form is not cached because it is used as the ODBC workaround for obtaining the last insert ID value. See the Connector/ODBC section of [Chapter 21, Connectors and APIs](#).

Statements within transactions that use `SERIALIZABLE [1470]` isolation level also cannot be cached because they use `LOCK IN SHARE MODE` locking.

- It uses `TEMPORARY` tables.
- It does not use any tables.
- It generates warnings.
- The user has a column-level privilege for any of the involved tables.

### 8.9.3.2. Query Cache `SELECT` Options

Two query cache-related options may be specified in `SELECT` statements:

- `SQL_CACHE`

The query result is cached if it is cacheable and the value of the `query_cache_type [529]` system variable is `ON` or `DEMAND`.

- [SQL\\_NO\\_CACHE](#)

The server does not use the query cache. It neither checks the query cache to see whether the result is already cached, nor does it cache the query result. (Due to a limitation in the parser, a space character must precede and follow the [SQL\\_NO\\_CACHE](#) keyword; a non-space such as a newline causes the server to check the query cache to see whether the result is already cached.)

Examples:

```
SELECT SQL_CACHE id, name FROM customer;
SELECT SQL_NO_CACHE id, name FROM customer;
```

### 8.9.3.3. Query Cache Configuration

The [have\\_query\\_cache](#) [487] server system variable indicates whether the query cache is available:

```
mysql> SHOW VARIABLES LIKE 'have_query_cache';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_query_cache | YES |
+-----+-----+
```

When using a standard MySQL binary, this value is always [YES](#), even if query caching is disabled.

Several other system variables control query cache operation. These can be set in an option file or on the command line when starting `mysqld`. The query cache system variables all have names that begin with [query\\_cache\\_](#). They are described briefly in [Section 5.1.4, “Server System Variables”](#), with additional configuration information given here.

To set the size of the query cache, set the [query\\_cache\\_size](#) [528] system variable. Setting it to 0 disables the query cache, as does setting [query\\_cache\\_type=0](#) [529]. By default, the query cache is disabled. This is achieved using a default size of 1M, with a default for [query\\_cache\\_type](#) of 0. (Before MySQL 5.6.8, the default size is 0, with a default [query\\_cache\\_type](#) of 1.)

To reduce overhead significantly, also start the server with [query\\_cache\\_type=0](#) [529] if you will not be using the query cache.



#### Note

When using the Windows Configuration Wizard to install or configure MySQL, the default value for [query\\_cache\\_size](#) [528] will be configured automatically for you based on the different configuration types available. When using the Windows Configuration Wizard, the query cache may be enabled (that is, set to a nonzero value) due to the selected configuration. The query cache is also controlled by the setting of the [query\\_cache\\_type](#) [529] variable. Check the values of these variables as set in your `my.ini` file after configuration has taken place.

When you set [query\\_cache\\_size](#) [528] to a nonzero value, keep in mind that the query cache needs a minimum size of about 40KB to allocate its structures. (The exact size depends on system architecture.) If you set the value too small, you'll get a warning, as in this example:

```
mysql> SET GLOBAL query_cache_size = 40000;
Query OK, 0 rows affected, 1 warning (0.00 sec)

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1282
Message: Query cache failed to set size 39936;
        new query cache size is 0

mysql> SET GLOBAL query_cache_size = 41984;
```



```
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE 'query_cache_size';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| query_cache_size | 41984 |
+-----+-----+
```

For the query cache to actually be able to hold any query results, its size must be set larger:

```
mysql> SET GLOBAL query_cache_size = 1000000;
Query OK, 0 rows affected (0.04 sec)

mysql> SHOW VARIABLES LIKE 'query_cache_size';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| query_cache_size | 999424 |
+-----+-----+
1 row in set (0.00 sec)
```

The `query_cache_size` [528] value is aligned to the nearest 1024 byte block. The value reported may therefore be different from the value that you assign.

If the query cache size is greater than 0, the `query_cache_type` [529] variable influences how it works. This variable can be set to the following values:

- A value of 0 or `OFF` prevents caching or retrieval of cached results.
- A value of 1 or `ON` enables caching except of those statements that begin with `SELECT SQL_NO_CACHE`.
- A value of 2 or `DEMAND` causes caching of only those statements that begin with `SELECT SQL_CACHE`.

If `query_cache_size` [528] is 0, you should also set `query_cache_type` [529] variable to 0. In this case, the server does not acquire the query cache mutex at all, which means that the query cache cannot be enabled at runtime and there is reduced overhead in query execution.

Setting the `GLOBAL query_cache_type` [529] value determines query cache behavior for all clients that connect after the change is made. Individual clients can control cache behavior for their own connection by setting the `SESSION query_cache_type` [529] value. For example, a client can disable use of the query cache for its own queries like this:

```
mysql> SET SESSION query_cache_type = OFF;
```

If you set `query_cache_type` [529] at server startup (rather than at runtime with a `SET` statement), only the numeric values are permitted.

To control the maximum size of individual query results that can be cached, set the `query_cache_limit` [527] system variable. The default value is 1MB.

Be careful not to set the size of the cache too large. Due to the need for threads to lock the cache during updates, you may see lock contention issues with a very large cache.



#### Note

You can set the maximum size that can be specified for the query cache at runtime with the `SET` statement by using the `--maximum-query_cache_size=32M` option on the command line or in the configuration file.

When a query is to be cached, its result (the data sent to the client) is stored in the query cache during result retrieval. Therefore the data usually is not handled in one big chunk. The query cache allocates blocks for storing this data on demand, so when one block is filled, a new block is allocated.



Because memory allocation operation is costly (timewise), the query cache allocates blocks with a minimum size given by the `query_cache_min_res_unit` [528] system variable. When a query is executed, the last result block is trimmed to the actual data size so that unused memory is freed. Depending on the types of queries your server executes, you might find it helpful to tune the value of `query_cache_min_res_unit` [528]:

- The default value of `query_cache_min_res_unit` [528] is 4KB. This should be adequate for most cases.
- If you have a lot of queries with small results, the default block size may lead to memory fragmentation, as indicated by a large number of free blocks. Fragmentation can force the query cache to prune (delete) queries from the cache due to lack of memory. In this case, decrease the value of `query_cache_min_res_unit` [528]. The number of free blocks and queries removed due to pruning are given by the values of the `Qcache_free_blocks` [595] and `Qcache_lowmem_prunes` [595] status variables.
- If most of your queries have large results (check the `Qcache_total_blocks` [596] and `Qcache_queries_in_cache` [596] status variables), you can increase performance by increasing `query_cache_min_res_unit` [528]. However, be careful to not make it too large (see the previous item).

### 8.9.3.4. Query Cache Status and Maintenance

To check whether the query cache is present in your MySQL server, use the following statement:

```
mysql> SHOW VARIABLES LIKE 'have_query_cache';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_query_cache | YES |
+-----+-----+
```

You can defragment the query cache to better utilize its memory with the `FLUSH QUERY CACHE` statement. The statement does not remove any queries from the cache.

The `RESET QUERY CACHE` statement removes all query results from the query cache. The `FLUSH TABLES` statement also does this.

To monitor query cache performance, use `SHOW STATUS` to view the cache status variables:

```
mysql> SHOW STATUS LIKE 'Qcache%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Qcache_free_blocks | 36 |
| Qcache_free_memory | 138488 |
| Qcache_hits | 79570 |
| Qcache_inserts | 27087 |
| Qcache_lowmem_prunes | 3114 |
| Qcache_not_cached | 22989 |
| Qcache_queries_in_cache | 415 |
| Qcache_total_blocks | 912 |
+-----+-----+
```

Descriptions of each of these variables are given in [Section 5.1.6, “Server Status Variables”](#). Some uses for them are described here.

The total number of `SELECT` queries is given by this formula:

```
Com_select
+ Qcache_hits
+ queries with errors found by parser
```

The `Com_select` value is given by this formula:

```
Qcache_inserts
+ Qcache_not_cached
```

```
+ queries with errors found during the column-privileges check
```

The query cache uses variable-length blocks, so `Qcache_total_blocks` [596] and `Qcache_free_blocks` [595] may indicate query cache memory fragmentation. After `FLUSH QUERY CACHE`, only a single free block remains.

Every cached query requires a minimum of two blocks (one for the query text and one or more for the query results). Also, every table that is used by a query requires one block. However, if two or more queries use the same table, only one table block needs to be allocated.

The information provided by the `Qcache_lowmem_prunes` [595] status variable can help you tune the query cache size. It counts the number of queries that have been removed from the cache to free up memory for caching new queries. The query cache uses a least recently used (LRU) strategy to decide which queries to remove from the cache. Tuning information is given in [Section 8.9.3.3, “Query Cache Configuration”](#).

## 8.9.4. Caching of Prepared Statements and Stored Programs

For certain statements that a client might execute multiple times during a session, the server converts the statement to an internal structure and caches that structure to be used during execution. Caching enables the server to perform more efficiently because it avoids the overhead of reconverting the statement should it be needed again during the session. Conversion and caching occurs for these statements:

- Prepared statements, both those processed at the SQL level (using the `PREPARE` statement) and those processed using the binary client/server protocol (using the `mysql_stmt_prepare()` C API function). The `max_prepared_stmt_count` [506] system variable controls the total number of statements the server caches. (The sum of the number of prepared statements across all sessions.)
- Stored programs (stored procedures and functions, triggers, and events). In this case, the server converts and caches the entire program body. The `stored_program_cache` [549] system variable indicates the approximate number of stored programs the the server caches per session.

The server maintains caches for prepared statements and stored programs on a per-session basis. Statements cached for one session are not accessible to other sessions. When a session ends, the server discards any statements cached for it.

When the server uses a cached internal statement structure, it must take care that the structure does not go out of date. Metadata changes can occur for an object used by the statement, causing a mismatch between the current object definition and the definition as represented in the internal statement structure. Metadata changes occur for DDL statements such as those that create, drop, alter, rename, or truncate tables, or that analyze, optimize, or repair tables. Table content changes (for example, with `INSERT` or `UPDATE`) do not change metadata, nor do `SELECT` statements.

Here is an illustration of the problem. Suppose that a client prepares this statement:

```
PREPARE s1 FROM 'SELECT * FROM t1';
```

The `SELECT *` expands in the internal structure to the list of columns in the table. If the set of columns in the table is modified with `ALTER TABLE`, the prepared statement goes out of date. If the server does not detect this change the next time the client executes `s1`, the prepared statement will return incorrect results.

To avoid problems caused by metadata changes to tables or views referred to by the prepared statement, the server detects these changes and automatically reprepares the statement when it is next executed. That is, the server reparses the statement and rebuilds the internal structure. Reparsing also occurs after referenced tables or views are flushed from the table definition cache, either implicitly to make room for new entries in the cache, or explicitly due to `FLUSH TABLES`.

Similarly, if changes occur to objects used by a stored program, the server reparses affected statements within the program. (Before MySQL 5.6.6, the server does not detect metadata changes affecting stored programs, so such changes can cause incorrect results or errors.)

The server also detects metadata changes for objects in expressions. These might be used in statements specific to stored programs, such as `DECLARE CURSOR` or flow-control statements such as `IF`, `CASE`, and `RETURN`.

To avoid reparsing entire stored programs, the server reparses affected statements or expressions within a program only as needed. Examples:

- Suppose that metadata for a table or view is changed. Reparsing occurs for a `SELECT *` within the program that accesses the table or view, but not for a `SELECT *` that does not access the table or view.
- When a statement is affected, the server reparses it only partially if possible. Consider this `CASE` statement:

```
CASE case_expr
  WHEN when_expr1 ...
  WHEN when_expr2 ...
  WHEN when_expr3 ...
  ...
END CASE
```

If a metadata change affects only `WHEN when_expr3`, that expression is reparsed. `case_expr` and the other `WHEN` expressions are not reparsed.

Reparsing uses the default database and SQL mode that were in effect for the original conversion to internal form.

The server attempts reparsing up to three times. An error occurs if all attempts fail.

Reparsing is automatic, but to the extent that it occurs, diminishes prepared statement and stored program performance.

For prepared statements, the `Com_stmt_reprepare` [586] status variable tracks the number of reparsings.

## 8.10. Optimizing Locking Operations

MySQL manages contention for table contents using [locking](#):

- Internal locking is performed within the MySQL server itself to manage contention for table contents by multiple threads. This type of locking is internal because it is performed entirely by the server and involves no other programs. See [Section 8.10.1, “Internal Locking Methods”](#).
- External locking occurs when the server and other programs lock `MyISAM` table files to coordinate among themselves which program can access the tables at which time. See [Section 8.10.5, “External Locking”](#).

### 8.10.1. Internal Locking Methods

This section discusses internal locking; that is, locking performed within the MySQL server itself to manage contention for table contents by multiple sessions. This type of locking is internal because it is performed entirely by the server and involves no other programs. For locking performed on MySQL files by other programs, see [Section 8.10.5, “External Locking”](#).

#### Row-Level Locking

MySQL uses [row-level locking](#) for `InnoDB` tables to support simultaneous write access by multiple sessions, making them suitable for multi-user, highly concurrent, and OLTP applications.

To avoid [deadlocks](#) when performing multiple concurrent write operations on a single `InnoDB` table, acquire necessary locks at the start of the transaction by issuing a `SELECT ... FOR UPDATE` statement for each group of rows expected to be modified, even if the `DML` statements come later in

the transaction. If transactions modify or lock more than one table, issue the applicable statements in the same order within each transaction. Deadlocks affect performance rather than representing a serious error, because [InnoDB](#) automatically [detects](#) deadlock conditions and rolls back one of the affected transactions.

Advantages of row-level locking:

- Fewer lock conflicts when different sessions access different rows.
- Fewer changes for rollbacks.
- Possible to lock a single row for a long time.

## Table-Level Locking

MySQL uses [table-level locking](#) for [MyISAM](#), [MEMORY](#), and [MERGE](#) tables, allowing only one session to update those tables at a time, making them more suitable for read-only, read-mostly, or single-user applications.

These storage engines avoid [deadlocks](#) by always requesting all needed locks at once at the beginning of a query and always locking the tables in the same order. The tradeoff is that this strategy reduces concurrency; other sessions that want to modify the table must wait until the current [DML](#) statement finishes.

MySQL grants table write locks as follows:

1. If there are no locks on the table, put a write lock on it.
2. Otherwise, put the lock request in the write lock queue.

MySQL grants table read locks as follows:

1. If there are no write locks on the table, put a read lock on it.
2. Otherwise, put the lock request in the read lock queue.

Table updates are given higher priority than table retrievals. Therefore, when a lock is released, the lock is made available to the requests in the write lock queue and then to the requests in the read lock queue. This ensures that updates to a table are not “starved” even if there is heavy [SELECT](#) activity for the table. However, if you have many updates for a table, [SELECT](#) statements wait until there are no more updates.

For information on altering the priority of reads and writes, see [Section 8.10.2, “Table Locking Issues”](#).

You can analyze the table lock contention on your system by checking the [Table\\_locks\\_immediate \[600\]](#) and [Table\\_locks\\_waited \[600\]](#) status variables, which indicate the number of times that requests for table locks could be granted immediately and the number that had to wait, respectively:

```
mysql> SHOW STATUS LIKE 'Table%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Table_locks_immediate | 1151552 |
| Table_locks_waited | 15324 |
+-----+-----+
```

The [MyISAM](#) storage engine supports concurrent inserts to reduce contention between readers and writers for a given table: If a [MyISAM](#) table has no free blocks in the middle of the data file, rows are always inserted at the end of the data file. In this case, you can freely mix concurrent [INSERT](#) and [SELECT](#) statements for a [MyISAM](#) table without locks. That is, you can insert rows into a [MyISAM](#) table at the same time other clients are reading from it. Holes can result from rows having been deleted from or updated in the middle of the table. If there are holes, concurrent inserts are disabled but are

enabled again automatically when all holes have been filled with new data.. This behavior is altered by the `concurrent_insert` [471] system variable. See [Section 8.10.3, “Concurrent Inserts”](#).

If you acquire a table lock explicitly with `LOCK TABLES`, you can request a `READ LOCAL` lock rather than a `READ` lock to enable other sessions to perform concurrent inserts while you have the table locked.

To perform many `INSERT` and `SELECT` operations on a table `real_table` when concurrent inserts are not possible, you can insert rows into a temporary table `temp_table` and update the real table with the rows from the temporary table periodically. This can be done with the following code:

```
mysql> LOCK TABLES real_table WRITE, temp_table WRITE;
mysql> INSERT INTO real_table SELECT * FROM temp_table;
mysql> DELETE FROM temp_table;
mysql> UNLOCK TABLES;
```

Advantages of table-level locking:

- Requires relatively little memory.
- Fast when used on a large part of the table because only a single lock is involved.
- Fast if you often do `GROUP BY` operations on a large part of the data or if you must scan the entire table frequently.

Generally, table locks are suited to the following cases:

- Most statements for the table are reads.
- Statements for the table are a mix of reads and writes, where writes are updates or deletes for a single row that can be fetched with one key read:

```
UPDATE tbl_name SET column=value WHERE unique_key_col=key_value;
DELETE FROM tbl_name WHERE unique_key_col=key_value;
```

- `SELECT` combined with concurrent `INSERT` statements, and very few `UPDATE` or `DELETE` statements.
- Many scans or `GROUP BY` operations on the entire table without any writers.

## 8.10.2. Table Locking Issues

`InnoDB` tables use row-level locking so that multiple sessions and applications can read from and write to the same table simultaneously, without making each other wait or producing inconsistent results. For this storage engine, avoid using the `LOCK TABLES` statement, because it does not offer any extra protection, but instead reduces concurrency. The automatic row-level locking makes these tables suitable for your busiest databases with your most important data, while also simplifying application logic since you do not need to lock and unlock tables. Consequently, the `InnoDB` storage engine is the default in MySQL 5.6.

MySQL uses table locking (instead of page, row, or column locking) for all storage engines except `InnoDB`. The locking operations themselves do not have much overhead. But because only one session can write to a table at any one time, for best performance with these other storage engines, use them primarily for tables that are queried often and rarely inserted into or updated.

### Performance Considerations Favoring `InnoDB`

When choosing whether to create a table using `InnoDB` or a different storage engine, keep in mind the following disadvantages of table locking:

- Table locking enables many sessions to read from a table at the same time, but if a session wants to write to a table, it must first get exclusive access, meaning it might have to wait for other sessions to finish with the table first. During the update, all other sessions that want to access this particular table must wait until the update is done.

- Table locking causes problems when a session is waiting because the disk is full and free space needs to become available before the session can proceed. In this case, all sessions that want to access the problem table are also put in a waiting state until more disk space is made available.
- A `SELECT` statement that takes a long time to run prevents other sessions from updating the table in the meantime, making the other sessions appear slow or unresponsive. While a session is waiting to get exclusive access to the table for updates, other sessions that issue `SELECT` statements will queue up behind it, reducing concurrency even for read-only sessions.

## Workarounds for Locking Performance Issues

The following items describe some ways to avoid or reduce contention caused by table locking:

- Consider switching the table to the `InnoDB` storage engine, either using `CREATE TABLE ... ENGINE=INNODB` during setup, or using `ALTER TABLE ... ENGINE=INNODB` for an existing table. See [Section 14.2, “The InnoDB Storage Engine”](#) for more details about this storage engine.
- Optimize `SELECT` statements to run faster so that they lock tables for a shorter time. You might have to create some summary tables to do this.
- Start `mysqld` with `--low-priority-updates [430]`. For storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`), this gives all statements that update (modify) a table lower priority than `SELECT` statements. In this case, the second `SELECT` statement in the preceding scenario would execute before the `UPDATE` statement, and would not wait for the first `SELECT` to finish.
- To specify that all updates issued in a specific connection should be done with low priority, set the `low_priority_updates [500]` server system variable equal to 1.
- To give a specific `INSERT`, `UPDATE`, or `DELETE` statement lower priority, use the `LOW_PRIORITY` attribute.
- To give a specific `SELECT` statement higher priority, use the `HIGH_PRIORITY` attribute. See [Section 13.2.9, “SELECT Syntax”](#).
- Start `mysqld` with a low value for the `max_write_lock_count [509]` system variable to force MySQL to temporarily elevate the priority of all `SELECT` statements that are waiting for a table after a specific number of inserts to the table occur. This permits `READ` locks after a certain number of `WRITE` locks.
- If you have problems with `INSERT` combined with `SELECT`, consider switching to `MyISAM` tables, which support concurrent `SELECT` and `INSERT` statements. (See [Section 8.10.3, “Concurrent Inserts”](#).)
- If you mix inserts and deletes on the same nontransactional table, `INSERT DELAYED` may help. See [Section 13.2.5.2, “INSERT DELAYED Syntax”](#).



### Note

As of MySQL 5.6.6, `INSERT DELAYED` is deprecated, and will be removed in a future release. Use `INSERT` (without `DELAYED`) instead.

- If you have problems with mixed `SELECT` and `DELETE` statements, the `LIMIT` option to `DELETE` may help. See [Section 13.2.2, “DELETE Syntax”](#).
- Using `SQL_BUFFER_RESULT` with `SELECT` statements can help to make the duration of table locks shorter. See [Section 13.2.9, “SELECT Syntax”](#).
- Splitting table contents into separate tables may help, by allowing queries to run against columns in one table, while updates are confined to columns in a different table.
- You could change the locking code in `mysys/thr_lock.c` to use a single queue. In this case, write locks and read locks would have the same priority, which might help some applications.



### 8.10.3. Concurrent Inserts

The `MyISAM` storage engine supports concurrent inserts to reduce contention between readers and writers for a given table: If a `MyISAM` table has no holes in the data file (deleted rows in the middle), an `INSERT` statement can be executed to add rows to the end of the table at the same time that `SELECT` statements are reading rows from the table. If there are multiple `INSERT` statements, they are queued and performed in sequence, concurrently with the `SELECT` statements. The results of a concurrent `INSERT` may not be visible immediately.

The `concurrent_insert` [471] system variable can be set to modify the concurrent-insert processing. By default, the variable is set to `AUTO` (or 1) and concurrent inserts are handled as just described. If `concurrent_insert` [471] is set to `NEVER` (or 0), concurrent inserts are disabled. If the variable is set to `ALWAYS` (or 2), concurrent inserts at the end of the table are permitted even for tables that have deleted rows. See also the description of the `concurrent_insert` [471] system variable.

Under circumstances where concurrent inserts can be used, there is seldom any need to use the `DELAYED` modifier for `INSERT` statements. See [Section 13.2.5.2, “INSERT DELAYED Syntax”](#).

If you are using the binary log, concurrent inserts are converted to normal inserts for `CREATE ... SELECT` or `INSERT ... SELECT` statements. This is done to ensure that you can re-create an exact copy of your tables by applying the log during a backup operation. See [Section 5.2.4, “The Binary Log”](#). In addition, for those statements a read lock is placed on the selected-from table such that inserts into that table are blocked. The effect is that concurrent inserts for that table must wait as well.

With `LOAD DATA INFILE`, if you specify `CONCURRENT` with a `MyISAM` table that satisfies the condition for concurrent inserts (that is, it contains no free blocks in the middle), other sessions can retrieve data from the table while `LOAD DATA` is executing. Use of the `CONCURRENT` option affects the performance of `LOAD DATA` a bit, even if no other session is using the table at the same time.

If you specify `HIGH_PRIORITY`, it overrides the effect of the `--low-priority-updates` [430] option if the server was started with that option. It also causes concurrent inserts not to be used.

For `LOCK TABLE`, the difference between `READ LOCAL` and `READ` is that `READ LOCAL` permits nonconflicting `INSERT` statements (concurrent inserts) to execute while the lock is held. However, this cannot be used if you are going to manipulate the database using processes external to the server while you hold the lock.

### 8.10.4. Metadata Locking

MySQL uses metadata locking to manage access to objects (tables, triggers, and so forth). Metadata locking is used to ensure data consistency but does involve some overhead, which increases as query volume increases. Metadata contention increases the more that multiple queries attempt to access the same objects.

Metadata locking is not a replacement for the table definition case, and its mutexes and locks differ from the `LOCK_open` mutex. The following discussion provides some information about how metadata locking works.

To ensure transaction serializability, the server must not permit one session to perform a data definition language (DDL) statement on a table that is used in an uncompleted transaction in another session. The server achieves this by acquiring metadata locks on tables used within a transaction and deferring release of those locks until the transaction ends. A metadata lock on a table prevents changes to the table's structure. This locking approach has the implication that a table that is being used by a transaction within one session cannot be used in DDL statements by other sessions until the transaction ends.

This principle applies not only to transactional tables, but also to nontransactional tables. Suppose that a session begins a transaction that uses transactional table `t` and nontransactional table `nt` as follows:

```
START TRANSACTION;
SELECT * FROM t;
SELECT * FROM nt;
```



Metadata locks are held on both `t` and `nt` until the transaction ends. If another session attempts a DDL operation on either table, it blocks until metadata lock release at transaction end. For example, a second session blocks if it attempts any of these operations:

```
DROP TABLE t;
ALTER TABLE t ...;
DROP TABLE nt;
ALTER TABLE nt ...;
```

If the server acquires metadata locks for a statement that is syntactically valid but fails during execution, it does not release the locks early. Lock release is still deferred to the end of the transaction because the failed statement is written to the binary log and the locks protect log consistency.

In autocommit mode, each statement is in effect a complete transaction, so metadata locks acquired for the statement are held only to the end of the statement.

Metadata locks acquired during a `PREPARE` statement are released once the statement has been prepared, even if preparation occurs within a multiple-statement transaction.

Before MySQL 5.5, when a transaction acquired the equivalent of a metadata lock for a table used within a statement, it released the lock at the end of the statement. This approach had the disadvantage that if a DDL statement occurred for a table that was being used by another session in an active transaction, statements could be written to the binary log in the wrong order.

## 8.10.5. External Locking

External locking is the use of file system locking to manage contention for `MyISAM` database tables by multiple processes. External locking is used in situations where a single process such as the MySQL server cannot be assumed to be the only process that requires access to tables. Here are some examples:

- If you run multiple servers that use the same database directory (not recommended), each server must have external locking enabled.
- If you use `myisamchk` to perform table maintenance operations on `MyISAM` tables, you must either ensure that the server is not running, or that the server has external locking enabled so that it locks table files as necessary to coordinate with `myisamchk` for access to the tables. The same is true for use of `myisampack` to pack `MyISAM` tables.

If the server is run with external locking enabled, you can use `myisamchk` at any time for read operations such as checking tables. In this case, if the server tries to update a table that `myisamchk` is using, the server will wait for `myisamchk` to finish before it continues.

If you use `myisamchk` for write operations such as repairing or optimizing tables, or if you use `myisampack` to pack tables, you *must* always ensure that the `mysqld` server is not using the table. If you don't stop `mysqld`, at least do a `mysqladmin flush-tables` before you run `myisamchk`. Your tables *may become corrupted* if the server and `myisamchk` access the tables simultaneously.

With external locking in effect, each process that requires access to a table acquires a file system lock for the table files before proceeding to access the table. If all necessary locks cannot be acquired, the process is blocked from accessing the table until the locks can be obtained (after the process that currently holds the locks releases them).

External locking affects server performance because the server must sometimes wait for other processes before it can access tables.

External locking is unnecessary if you run a single server to access a given data directory (which is the usual case) and if no other programs such as `myisamchk` need to modify tables while the server is running. If you only *read* tables with other programs, external locking is not required, although `myisamchk` might report warnings if the server changes tables while `myisamchk` is reading them.

With external locking disabled, to use `myisamchk`, you must either stop the server while `myisamchk` executes or else lock and flush the tables before running `myisamchk`. (See [Section 8.11.1](#), “System

Factors and Startup Parameter Tuning".) To avoid this requirement, use the `CHECK TABLE` and `REPAIR TABLE` statements to check and repair `MyISAM` tables.

For `mysqld`, external locking is controlled by the value of the `skip_external_locking` [540] system variable. When this variable is enabled, external locking is disabled, and vice versa. From MySQL 4.0 on, external locking is disabled by default.

Use of external locking can be controlled at server startup by using the `--external-locking` [422] or `--skip-external-locking` [439] option.

If you do use external locking option to enable updates to `MyISAM` tables from many MySQL processes, you must ensure that the following conditions are satisfied:

- Do not use the query cache for queries that use tables that are updated by another process.
- Do not start the server with the `--delay-key-write=ALL` [421] option or use the `DELAY_KEY_WRITE=1` table option for any shared tables. Otherwise, index corruption can occur.

The easiest way to satisfy these conditions is to always use `--external-locking` [422] together with `--delay-key-write=OFF` [421] and `--query-cache-size=0` [528]. (This is not done by default because in many setups it is useful to have a mixture of the preceding options.)

## 8.11. Optimizing the MySQL Server

This section discusses optimization techniques for the database server, primarily dealing with system configuration rather than tuning SQL statements. The information in this section is appropriate for DBAs who want to ensure performance and scalability across the servers they manage; for developers constructing installation scripts that include setting up the database; and people running MySQL themselves for development, testing, and so on who want to maximize their own productivity.

### 8.11.1. System Factors and Startup Parameter Tuning

We start with system-level factors, because some of these decisions must be made very early to achieve large performance gains. In other cases, a quick look at this section may suffice. However, it is always nice to have a sense of how much can be gained by changing factors that apply at this level.

The operating system to use is very important. To get the best use of multiple-CPU machines, you should use Solaris (because its threads implementation works well) or Linux (because the 2.4 and later kernels have good SMP support). Note that older Linux kernels have a 2GB filesize limit by default. If you have such a kernel and a need for files larger than 2GB, get the Large File Support (LFS) patch for the ext2 file system. Other file systems such as ReiserFS and XFS do not have this 2GB limitation.

Before using MySQL in production, we advise you to test it on your intended platform.

Other tips:

- If you have enough RAM, you could remove all swap devices. Some operating systems use a swap device in some contexts even if you have free memory.
- Avoid external locking for `MyISAM` tables. Since MySQL 4.0, the default has been for external locking to be disabled on all systems. The `--external-locking` [422] and `--skip-external-locking` [439] options explicitly enable and disable external locking.

Note that disabling external locking does not affect MySQL's functionality as long as you run only one server. Just remember to take down the server (or lock and flush the relevant tables) before you run `myisamchk`. On some systems it is mandatory to disable external locking because it does not work, anyway.

The only case in which you cannot disable external locking is when you run multiple MySQL servers (not clients) on the same data, or if you run `myisamchk` to check (not repair) a table without telling the server to flush and lock the tables first. Note that using multiple MySQL servers to access the same data concurrently is generally *not* recommended, except when using MySQL Cluster.

**Note**

MySQL Cluster is currently not supported in MySQL 5.6. Users wishing to upgrade a MySQL Cluster from MySQL 5.0 or 5.1 should instead migrate to MySQL Cluster NDB 7.0 or 7.1; these are based on MySQL 5.1 but contain the latest improvements and fixes for [NDB](#).

The [LOCK TABLES](#) and [UNLOCK TABLES](#) statements use internal locking, so you can use them even if external locking is disabled.

## 8.11.2. Tuning Server Parameters

You can determine the default buffer sizes used by the `mysqld` server using this command:

```
shell> mysqld --verbose --help
```

This command produces a list of all `mysqld` options and configurable system variables. The output includes the default variable values and looks something like this:

```

abort-slave-event-count          0
allow-suspicious-udfs           FALSE
archive                          ON
auto-increment-increment        1
auto-increment-offset           1
autocommit                       TRUE
automatic-sp-privileges         TRUE
back-log                         80
basedir                          /home/jon/bin/mysql-5.6/
big-tables                       FALSE
bind-address                     *
binlog-cache-size                32768
binlog-checksum                  CRC32
binlog-direct-non-transactional-updates FALSE
binlog-format                    STATEMENT
binlog-max-flush-queue-time      0
binlog-order-commits             TRUE
binlog-row-event-max-size        8192
binlog-row-image                 FULL
binlog-rows-query-log-events     FALSE
binlog-stmt-cache-size           32768
blackhole                        ON
bulk-insert-buffer-size          8388608
character-set-client-handshake   TRUE
character-set-filesystem         binary
character-set-server             latin1
character-sets-dir               /home/jon/bin/mysql-5.6/share/charsets/
chroot                          (No default value)
collation-server                 latin1_swedish_ci
completion-type                  NO_CHAIN
concurrent-insert                AUTO
connect-timeout                  10
console                          FALSE
datadir                         (No default value)
date-format                      %Y-%m-%d
datetime-format                  %Y-%m-%d %H:%i:%s
default-storage-engine           InnoDB
default-time-zone                (No default value)
default-tmp-storage-engine       InnoDB
default-week-format              0
delay-key-write                  ON
delayed-insert-limit             100
delayed-insert-timeout           300
delayed-queue-size               1000
des-key-file                      (No default value)
disconnect-on-expired-password   TRUE
disconnect-slave-event-count     0
div-precision-increment          4
end-markers-in-json              FALSE

```

## Tuning Server Parameters

enforce-gtid-consistency	FALSE
eq-range-index-dive-limit	10
event-scheduler	OFF
expire-logs-days	0
explicit-defaults-for-timestamp	FALSE
external-locking	FALSE
flush	FALSE
flush-time	0
ft-boolean-syntax	+ -><()~*:""&
ft-max-word-len	84
ft-min-word-len	4
ft-query-expansion-limit	20
ft-stopword-file	(No default value)
gdb	FALSE
general-log	FALSE
general-log-file	/home/jon/bin/mysql-5.6/data/havskatt.log
group-concat-max-len	1024
gtid-mode	OFF
help	TRUE
host-cache-size	279
ignore-builtin-innodb	FALSE
init-connect	
init-file	(No default value)
init-slave	
innodb	ON
innodb-adaptive-flushing	TRUE
innodb-adaptive-flushing-lwm	10
innodb-adaptive-hash-index	TRUE
innodb-adaptive-max-sleep-delay	150000
innodb-additional-mem-pool-size	8388608
innodb-api-bk-commit-interval	5
innodb-api-disable-rowlock	FALSE
innodb-api-enable-binlog	FALSE
innodb-api-enable-mdl	FALSE
innodb-api-trx-level	0
innodb-autoextend-increment	64
innodb-autoinc-lock-mode	1
innodb-buffer-page	ON
innodb-buffer-page-lru	ON
innodb-buffer-pool-dump-at-shutdown	FALSE
innodb-buffer-pool-dump-now	FALSE
innodb-buffer-pool-filename	ib_buffer_pool
innodb-buffer-pool-instances	0
innodb-buffer-pool-load-abort	FALSE
innodb-buffer-pool-load-at-startup	FALSE
innodb-buffer-pool-load-now	FALSE
innodb-buffer-pool-size	134217728
innodb-buffer-pool-stats	ON
innodb-change-buffer-max-size	25
innodb-change-buffering	all
innodb-checksum-algorithm	innodb
innodb-checksums	TRUE
innodb-cmp	ON
innodb-cmp-per-index	ON
innodb-cmp-per-index-enabled	FALSE
innodb-cmp-per-index-reset	ON
innodb-cmp-reset	ON
innodb-cmpmem	ON
innodb-cmpmem-reset	ON
innodb-commit-concurrency	0
innodb-compression-failure-threshold-pct	5
innodb-compression-level	6
innodb-compression-pad-pct-max	50
innodb-concurrency-tickets	5000
innodb-data-file-path	(No default value)
innodb-data-home-dir	(No default value)
innodb-disable-sort-file-cache	FALSE
innodb-doublewrite	TRUE
innodb-fast-shutdown	1
innodb-file-format	Antelope
innodb-file-format-check	TRUE
innodb-file-format-max	Antelope

## Tuning Server Parameters

innodb-file-io-threads	4
innodb-file-per-table	TRUE
innodb-flush-log-at-timeout	1
innodb-flush-log-at-trx-commit	1
innodb-flush-method	(No default value)
innodb-flush-neighbors	1
innodb-flushing-avg-loops	30
innodb-force-load-corrupted	FALSE
innodb-force-recovery	0
innodb-ft-aux-table	(No default value)
innodb-ft-being-deleted	ON
innodb-ft-cache-size	8000000
innodb-ft-config	ON
innodb-ft-default-stopword	ON
innodb-ft-deleted	ON
innodb-ft-enable-diag-print	FALSE
innodb-ft-enable-stopword	TRUE
innodb-ft-index-cache	ON
innodb-ft-index-table	ON
innodb-ft-inserted	ON
innodb-ft-max-token-size	84
innodb-ft-min-token-size	3
innodb-ft-num-word-optimize	2000
innodb-ft-server-stopword-table	(No default value)
innodb-ft-sort-pll-degree	2
innodb-ft-user-stopword-table	(No default value)
innodb-io-capacity	200
innodb-io-capacity-max	18446744073709551615
innodb-large-prefix	FALSE
innodb-lock-wait-timeout	50
innodb-lock-waits	ON
innodb-locks	ON
innodb-locks-unsafe-for-binlog	FALSE
innodb-log-buffer-size	8388608
innodb-log-compressed-pages	TRUE
innodb-log-file-size	50331648
innodb-log-files-in-group	2
innodb-log-group-home-dir	(No default value)
innodb-lru-scan-depth	1024
innodb-max-dirty-pages-pct	75
innodb-max-dirty-pages-pct-lwm	0
innodb-max-purge-lag	0
innodb-max-purge-lag-delay	0
innodb-metrics	ON
innodb-mirrored-log-groups	1
innodb-monitor-disable	(No default value)
innodb-monitor-enable	(No default value)
innodb-monitor-reset	(No default value)
innodb-monitor-reset-all	(No default value)
innodb-old-blocks-pct	37
innodb-old-blocks-time	1000
innodb-online-alter-log-max-size	134217728
innodb-open-files	0
innodb-optimize-fulltext-only	FALSE
innodb-page-size	16384
innodb-print-all-deadlocks	FALSE
innodb-purge-batch-size	300
innodb-purge-threads	1
innodb-random-read-ahead	FALSE
innodb-read-ahead-threshold	56
innodb-read-io-threads	4
innodb-read-only	FALSE
innodb-replication-delay	0
innodb-rollback-on-timeout	FALSE
innodb-rollback-segments	128
innodb-sort-buffer-size	1048576
innodb-spin-wait-delay	6
innodb-stats-auto-recalc	TRUE
innodb-stats-method	nulls_equal
innodb-stats-on-metadata	FALSE
innodb-stats-persistent	TRUE
innodb-stats-persistent-sample-pages	20

## Tuning Server Parameters

innodb-stats-sample-pages	8
innodb-stats-transient-sample-pages	8
innodb-status-file	FALSE
innodb-strict-mode	FALSE
innodb-support-xa	TRUE
innodb-sync-array-size	1
innodb-sync-spin-loops	30
innodb-sys-columns	ON
innodb-sys-datafiles	ON
innodb-sys-fields	ON
innodb-sys-foreign	ON
innodb-sys-foreign-cols	ON
innodb-sys-indexes	ON
innodb-sys-tables	ON
innodb-sys-tablespaces	ON
innodb-sys-tablestats	ON
innodb-table-locks	TRUE
innodb-thread-concurrency	0
innodb-thread-sleep-delay	10000
innodb-trx	ON
innodb-undo-directory	.
innodb-undo-logs	128
innodb-undo-tablespaces	0
innodb-use-native-aio	TRUE
innodb-use-sys-malloc	TRUE
innodb-write-io-threads	4
interactive-timeout	28800
join-buffer-size	262144
keep-files-on-create	FALSE
key-buffer-size	8388608
key-cache-age-threshold	300
key-cache-block-size	1024
key-cache-division-limit	100
language	/home/jon/bin/mysql-5.6/share/
large-pages	FALSE
lc-messages	en_US
lc-messages-dir	/home/jon/bin/mysql-5.6/share/
lc-time-names	en_US
local-infile	TRUE
lock-wait-timeout	31536000
log-bin	(No default value)
log-bin-index	(No default value)
log-bin-trust-function-creators	FALSE
log-bin-use-vl-row-events	FALSE
log-error	
log-isam	myisam.log
log-output	FILE
log-queries-not-using-indexes	FALSE
log-raw	FALSE
log-short-format	FALSE
log-slave-updates	FALSE
log-slow-admin-statements	FALSE
log-slow-slave-statements	FALSE
log-tc	tc.log
log-tc-size	24576
log-throttle-queries-not-using-indexes	0
log-warnings	1
long-query-time	10
low-priority-updates	FALSE
lower-case-table-names	0
master-info-file	master.info
master-info-repository	FILE
master-retry-count	86400
master-verify-checksum	FALSE
max-allowed-packet	4194304
max-binlog-cache-size	18446744073709547520
max-binlog-dump-events	0
max-binlog-size	1073741824
max-binlog-stmt-cache-size	18446744073709547520
max-connect-errors	100
max-connections	151
max-delayed-threads	20

## Tuning Server Parameters

max-error-count	64
max-heap-table-size	16777216
max-join-size	18446744073709551615
max-length-for-sort-data	1024
max-prepared-stmt-count	16382
max-relay-log-size	0
max-seeks-for-key	18446744073709551615
max-sort-length	1024
max-sp-recursion-depth	0
max-tmp-tables	32
max-user-connections	0
max-write-lock-count	18446744073709551615
memlock	FALSE
metadata-locks-cache-size	1024
metadata-locks-hash-instances	8
min-examined-row-limit	0
multi-range-count	256
mysam-block-size	1024
mysam-data-pointer-size	6
mysam-max-sort-file-size	9223372036853727232
mysam-mmap-size	18446744073709551615
mysam-recover-options	OFF
mysam-repair-threads	1
mysam-sort-buffer-size	8388608
mysam-stats-method	nulls_unequal
mysam-use-mmap	FALSE
net-buffer-length	16384
net-read-timeout	30
net-retry-count	10
net-write-timeout	60
new	FALSE
old	FALSE
old-alter-table	FALSE
old-passwords	0
old-style-user-limits	FALSE
open-files-limit	1024
optimizer-prune-level	1
optimizer-search-depth	62
optimizer-switch	index_merge=on,index_merge_union=on,index_merge
optimizer-trace	
optimizer-trace-features	greedy_search=on,range_optimizer=on,dynamic_rang
optimizer-trace-limit	1
optimizer-trace-max-mem-size	16384
optimizer-trace-offset	-1
partition	ON
performance-schema	TRUE
performance-schema-accounts-size	-1
performance-schema-consumer-events-stages-current	FALSE
performance-schema-consumer-events-stages-history	FALSE
performance-schema-consumer-events-stages-history-long	FALSE
performance-schema-consumer-events-statements-current	TRUE
performance-schema-consumer-events-statements-history	FALSE
performance-schema-consumer-events-statements-history-long	FALSE
performance-schema-consumer-events-waits-current	FALSE
performance-schema-consumer-events-waits-history	FALSE
performance-schema-consumer-events-waits-history-long	FALSE
performance-schema-consumer-global-instrumentation	TRUE
performance-schema-consumer-statements-digest	TRUE
performance-schema-consumer-thread-instrumentation	TRUE
performance-schema-digests-size	-1
performance-schema-events-stages-history-long-size	-1
performance-schema-events-stages-history-size	-1
performance-schema-events-statements-history-long-size	-1
performance-schema-events-statements-history-size	-1
performance-schema-events-waits-history-long-size	-1
performance-schema-events-waits-history-size	-1
performance-schema-hosts-size	-1
performance-schema-instrument	
performance-schema-max-cond-classes	80
performance-schema-max-cond-instances	-1
performance-schema-max-file-classes	50
performance-schema-max-file-handles	32768



## Tuning Server Parameters

performance-schema-max-file-instances	-1
performance-schema-max-mutex-classes	200
performance-schema-max-mutex-instances	-1
performance-schema-max-rwlock-classes	30
performance-schema-max-rwlock-instances	-1
performance-schema-max-socket-classes	10
performance-schema-max-socket-instances	-1
performance-schema-max-stage-classes	150
performance-schema-max-statement-classes	167
performance-schema-max-table-handles	-1
performance-schema-max-table-instances	-1
performance-schema-max-thread-classes	50
performance-schema-max-thread-instances	-1
performance-schema-session-connect-attrs-size	-1
performance-schema-setup-actors-size	100
performance-schema-setup-objects-size	100
performance-schema-users-size	-1
pid-file	/home/jon/bin/mysql-5.6/data/havskatt.pid
plugin-dir	/home/jon/bin/mysql-5.6/lib/plugin/
port	3306
port-open-timeout	0
preload-buffer-size	32768
profiling-history-size	15
query-alloc-block-size	8192
query-cache-limit	1048576
query-cache-min-res-unit	4096
query-cache-size	1048576
query-cache-type	OFF
query-cache-wlock-invalidate	FALSE
query-prealloc-size	8192
range-alloc-block-size	4096
read-buffer-size	131072
read-only	FALSE
read-rnd-buffer-size	262144
relay-log	(No default value)
relay-log-index	(No default value)
relay-log-info-file	relay-log.info
relay-log-info-repository	FILE
relay-log-purge	TRUE
relay-log-recovery	FALSE
relay-log-space-limit	0
replicate-same-server-id	FALSE
report-host	(No default value)
report-password	(No default value)
report-port	0
report-user	(No default value)
safe-user-create	FALSE
secure-auth	TRUE
secure-file-priv	(No default value)
server-id	0
server-id-bits	32
sha256-password-private-key-path	private_key.pem
sha256-password-public-key-path	public_key.pem
show-slave-auth-info	FALSE
skip-grant-tables	FALSE
skip-name-resolve	FALSE
skip-networking	FALSE
skip-show-database	FALSE
skip-slave-start	FALSE
slave-allow-batching	FALSE
slave-checkpoint-group	512
slave-checkpoint-period	300
slave-compressed-protocol	FALSE
slave-exec-mode	STRICT
slave-load-tmpdir	/tmp
slave-max-allowed-packet	1073741824
slave-net-timeout	3600
slave-parallel-workers	0
slave-pending-jobs-size-max	16777216
slave-rows-search-algorithms	TABLE_SCAN, INDEX_SCAN
slave-skip-errors	(No default value)
slave-sql-verify-checksum	TRUE

slave-transaction-retries	10
slave-type-conversions	
slow-launch-time	2
slow-query-log	FALSE
slow-query-log-file	/home/jon/bin/mysql-5.6/data/havskatt-slow.log
socket	/tmp/mysql.sock
sort-buffer-size	262144
sporadic-binlog-dump-fail	FALSE
sql-mode	NO_ENGINE_SUBSTITUTION
ssl	FALSE
ssl-ca	(No default value)
ssl-capath	(No default value)
ssl-cert	(No default value)
ssl-cipher	(No default value)
ssl-crl	(No default value)
ssl-crlpath	(No default value)
ssl-key	(No default value)
stored-program-cache	256
super-large-pages	FALSE
symbolic-links	TRUE
sync-binlog	0
sync_frm	TRUE
sync-master-info	10000
sync-relay-log	10000
sync-relay-log-info	10000
sysdate-is-now	FALSE
table-definition-cache	615
table-open-cache	431
table-open-cache-instances	1
tc-heuristic-recover	COMMIT
temp-pool	TRUE
thread-cache-size	9
thread-concurrency	10
thread-handling	one-thread-per-connection
thread-stack	262144
time-format	%H:%i:%s
timed-mutexes	FALSE
tmp-table-size	16777216
tmpdir	/tmp
transaction-alloc-block-size	8192
transaction-isolation	REPEATABLE-READ
transaction-prealloc-size	4096
transaction-read-only	FALSE
updatable-views-with-limit	YES
verbose	TRUE
wait-timeout	

For a `mysqld` server that is currently running, you can see the current values of its system variables by connecting to it and issuing this statement:

```
mysql> SHOW VARIABLES;
```

You can also see some statistical and status indicators for a running server by issuing this statement:

```
mysql> SHOW STATUS;
```

System variable and status information also can be obtained using `mysqladmin`:

```
shell> mysqladmin variables
shell> mysqladmin extended-status
```

For a full description of all system and status variables, see [Section 5.1.4, “Server System Variables”](#), and [Section 5.1.6, “Server Status Variables”](#).

MySQL uses algorithms that are very scalable, so you can usually run with very little memory. However, normally you get better performance by giving MySQL more memory.

When tuning a MySQL server, the two most important variables to configure are `key_buffer_size` [492] and `table_open_cache` [551]. You should first feel confident that you have these set appropriately before trying to change any other variables.

The following examples indicate some typical variable values for different runtime configurations.

- If you have at least 256MB of memory and many tables and want maximum performance with a moderate number of clients, use something like this:

```
shell> mysqld_safe --key_buffer_size=64M --table_open_cache=256 \  
--sort_buffer_size=4M --read_buffer_size=1M &
```

- If you have only 128MB of memory and only a few tables, but you still do a lot of sorting, you can use something like this:

```
shell> mysqld_safe --key_buffer_size=16M --sort_buffer_size=1M
```

If there are very many simultaneous connections, swapping problems may occur unless `mysqld` has been configured to use very little memory for each connection. `mysqld` performs better if you have enough memory for all connections.

- With little memory and lots of connections, use something like this:

```
shell> mysqld_safe --key_buffer_size=512K --sort_buffer_size=100K \  
--read_buffer_size=100K &
```

Or even this:

```
shell> mysqld_safe --key_buffer_size=512K --sort_buffer_size=16K \  
--table_open_cache=32 --read_buffer_size=8K \  
--net_buffer_length=1K &
```

If you are performing `GROUP BY` or `ORDER BY` operations on tables that are much larger than your available memory, increase the value of `read_rnd_buffer_size` [533] to speed up the reading of rows following sorting operations.

You can make use of the example option files included with your MySQL distribution; see [Section 5.1.2, “Server Configuration Defaults”](#).

If you specify an option on the command line for `mysqld` or `mysqld_safe`, it remains in effect only for that invocation of the server. To use the option every time the server runs, put it in an option file.

To see the effects of a parameter change, do something like this:

```
shell> mysqld --key_buffer_size=32M --verbose --help
```

The variable values are listed near the end of the output. Make sure that the `--verbose` [447] and `--help` [413] options are last. Otherwise, the effect of any options listed after them on the command line are not reflected in the output.

For information on tuning the `InnoDB` storage engine, see [Section 14.2.4.1, “InnoDB Performance Tuning Tips”](#).

### 8.11.3. Optimizing Disk I/O

- Disk seeks are a huge performance bottleneck. This problem becomes more apparent when the amount of data starts to grow so large that effective caching becomes impossible. For large databases where you access data more or less randomly, you can be sure that you need at least one disk seek to read and a couple of disk seeks to write things. To minimize this problem, use disks with low seek times.
- Increase the number of available disk spindles (and thereby reduce the seek overhead) by either symlinking files to different disks or striping the disks:
  - Using symbolic links

This means that, for `MyISAM` tables, you symlink the index file and data files from their usual location in the data directory to another disk (that may also be striped). This makes both the

seek and read times better, assuming that the disk is not used for other purposes as well. See [Section 8.11.3.1, “Using Symbolic Links”](#).

- **Striping**

Striping means that you have many disks and put the first block on the first disk, the second block on the second disk, and the  $N$ -th block on the  $(N \bmod \text{number\_of\_disks})$  disk, and so on. This means if your normal data size is less than the stripe size (or perfectly aligned), you get much better performance. Striping is very dependent on the operating system and the stripe size, so benchmark your application with different stripe sizes. See [Section 8.12.3, “Using Your Own Benchmarks”](#).

The speed difference for striping is very dependent on the parameters. Depending on how you set the striping parameters and number of disks, you may get differences measured in orders of magnitude. You have to choose to optimize for random or sequential access.

- For reliability, you may want to use RAID 0+1 (striping plus mirroring), but in this case, you need  $2 \times N$  drives to hold  $N$  drives of data. This is probably the best option if you have the money for it. However, you may also have to invest in some volume-management software to handle it efficiently.
- A good option is to vary the RAID level according to how critical a type of data is. For example, store semi-important data that can be regenerated on a RAID 0 disk, but store really important data such as host information and logs on a RAID 0+1 or RAID  $N$  disk. RAID  $N$  can be a problem if you have many writes, due to the time required to update the parity bits.
- On Linux, you can get much better performance by using `hdparm` to configure your disk's interface. (Up to 100% under load is not uncommon.) The following `hdparm` options should be quite good for MySQL, and probably for many other applications:

```
hdparm -m 16 -d 1
```

Note that performance and reliability when using this command depend on your hardware, so we strongly suggest that you test your system thoroughly after using `hdparm`. Please consult the `hdparm` manual page for more information. If `hdparm` is not used wisely, file system corruption may result, so back up everything before experimenting!

- You can also set the parameters for the file system that the database uses:

If you do not need to know when files were last accessed (which is not really useful on a database server), you can mount your file systems with the `-o noatime` option. That skips updates to the last access time in inodes on the file system, which avoids some disk seeks.

On many operating systems, you can set a file system to be updated asynchronously by mounting it with the `-o async` option. If your computer is reasonably stable, this should give you better performance without sacrificing too much reliability. (This flag is on by default on Linux.)

### 8.11.3.1. Using Symbolic Links

You can move databases or tables from the database directory to other locations and replace them with symbolic links to the new locations. You might want to do this, for example, to move a database to a file system with more free space or increase the speed of your system by spreading your tables to different disks.

For `InnoDB` tables, use the `DATA DIRECTORY` clause on the `CREATE TABLE` statement instead of symbolic links, as explained in [Section 5.4.1.2, “Specifying the Location of a Tablespace”](#). This new feature is a supported, cross-platform technique.

The recommended way to do this is to symlink entire database directories to a different disk. Symlink `MyISAM` tables only as a last resort.

To determine the location of your data directory, use this statement:

```
SHOW VARIABLES LIKE 'datadir';
```

### 8.11.3.1.1. Using Symbolic Links for Databases on Unix

On Unix, the way to symlink a database is first to create a directory on some disk where you have free space and then to create a soft link to it from the MySQL data directory.

```
shell> mkdir /dr1/databases/test
shell> ln -s /dr1/databases/test /path/to/datadir
```

MySQL does not support linking one directory to multiple databases. Replacing a database directory with a symbolic link works as long as you do not make a symbolic link between databases. Suppose that you have a database `db1` under the MySQL data directory, and then make a symlink `db2` that points to `db1`:

```
shell> cd /path/to/datadir
shell> ln -s db1 db2
```

The result is that, or any table `tbl_a` in `db1`, there also appears to be a table `tbl_a` in `db2`. If one client updates `db1.tbl_a` and another client updates `db2.tbl_a`, problems are likely to occur.

### 8.11.3.1.2. Using Symbolic Links for MyISAM Tables on Unix

Symlinks are fully supported only for `MyISAM` tables. For files used by tables for other storage engines, you may get strange problems if you try to use symbolic links. For `InnoDB` tables, use the alternative technique explained in [Section 5.4.1.2, “Specifying the Location of a Tablespace”](#) instead.

Do not symlink tables on systems that do not have a fully operational `realpath()` call. (Linux and Solaris support `realpath()`.) To determine whether your system supports symbolic links, check the value of the `have_symlink` [\[487\]](#) system variable using this statement:

```
SHOW VARIABLES LIKE 'have_symlink';
```

The handling of symbolic links for `MyISAM` tables works as follows:

- In the data directory, you always have the table format (`.frm`) file, the data (`.MYD`) file, and the index (`.MYI`) file. The data file and index file can be moved elsewhere and replaced in the data directory by symlinks. The format file cannot.
- You can symlink the data file and the index file independently to different directories.
- To instruct a running MySQL server to perform the symlinking, use the `DATA DIRECTORY` and `INDEX DIRECTORY` options to `CREATE TABLE`. See [Section 13.1.14, “CREATE TABLE Syntax”](#). Alternatively, if `mysqld` is not running, symlinking can be accomplished manually using `ln -s` from the command line.



#### Note

The path used with either or both of the `DATA DIRECTORY` and `INDEX DIRECTORY` options may not include the MySQL `data` directory. (Bug #32167)

- `myisamchk` does not replace a symlink with the data file or index file. It works directly on the file to which the symlink points. Any temporary files are created in the directory where the data file or index file is located. The same is true for the `ALTER TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements.



#### Note

When you drop a table that is using symlinks, *both the symlink and the file to which the symlink points are dropped*. This is an extremely good reason

not to run `mysqld` as the system `root` or permit system users to have write access to MySQL database directories.

- If you rename a table with `ALTER TABLE ... RENAME` or `RENAME TABLE` and you do not move the table to another database, the symlinks in the database directory are renamed to the new names and the data file and index file are renamed accordingly.
- If you use `ALTER TABLE ... RENAME` or `RENAME TABLE` to move a table to another database, the table is moved to the other database directory. If the table name changed, the symlinks in the new database directory are renamed to the new names and the data file and index file are renamed accordingly.
- If you are not using symlinks, start `mysqld` with the `--skip-symbolic-links` [441] option to ensure that no one can use `mysqld` to drop or rename a file outside of the data directory.

These table symlink operations are not supported:

- `ALTER TABLE` ignores the `DATA DIRECTORY` and `INDEX DIRECTORY` table options.
- As indicated previously, only the data and index files can be symbolic links. The `.frm` file must *never* be a symbolic link. Attempting to do this (for example, to make one table name a synonym for another) produces incorrect results. Suppose that you have a database `db1` under the MySQL data directory, a table `tbl1` in this database, and in the `db1` directory you make a symlink `tbl2` that points to `tbl1`:

```
shell> cd /path/to/datadir/db1
shell> ln -s tbl1.frm tbl2.frm
shell> ln -s tbl1.MYD tbl2.MYD
shell> ln -s tbl1.MYI tbl2.MYI
```

Problems result if one thread reads `db1.tbl1` and another thread updates `db1.tbl2`:

- The query cache is “fooled” (it has no way of knowing that `tbl1` has not been updated, so it returns outdated results).
- `ALTER` statements on `tbl2` fail.

### 8.11.3.1.3. Using Symbolic Links for Databases on Windows

On Windows, symbolic links can be used for database directories. This enables you to put a database directory at a different location (for example, on a different disk) by setting up a symbolic link to it. Use of database symlinks on Windows is similar to their use on Unix, although the procedure for setting up the link differs.

Suppose that you want to place the database directory for a database named `mydb` at `D:\data\mydb`. To do this, create a symbolic link in the MySQL data directory that points to `D:\data\mydb`. However, before creating the symbolic link, make sure that the `D:\data\mydb` directory exists by creating it if necessary. If you already have a database directory named `mydb` in the data directory, move it to `D:\data`. Otherwise, the symbolic link will be ineffective. To avoid problems, make sure that the server is not running when you move the database directory.

The procedure for creating the database symbolic link depends on your version of Windows.

Windows Vista, Windows Server 2008, or newer have native symbolic link support, so you can create a symlink using the `mklink` command. This command requires administrative privileges.

1. Change location into the data directory:

```
C:\> cd \path\to\datadir
```

2. In the data directory, create a symlink named `mydb` that points to the location of the database directory:

```
C:\> mklink /d mydb D:\data\mydb
```

After this, all tables created in the database `mydb` are created in `D:\data\mydb`.

Alternatively, on any version of Windows supported by MySQL, you can create a symbolic link to a MySQL database by creating a `.sym` file in the data directory that contains the path to the destination directory. The file should be named `db_name.sym`, where `db_name` is the database name.

Support for database symbolic links on Windows using `.sym` files is enabled by default. If you do not need `.sym` file symbolic links, you can disable support for them by starting `mysqld` with the `--skip-symbolic-links` [441] option. To determine whether your system supports `.sym` file symbolic links, check the value of the `have_symlink` [487] system variable using this statement:

```
SHOW VARIABLES LIKE 'have_symlink';
```

To create a `.sym` file symlink, use this procedure:

1. Change location into the data directory:

```
C:\> cd \path\to\datadir
```

2. In the data directory, create a text file named `mydb.sym` that contains this path name: `D:\data\mydb\`



#### Note

The path name to the new database and tables should be absolute. If you specify a relative path, the location will be relative to the `mydb.sym` file.

After this, all tables created in the database `mydb` are created in `D:\data\mydb`.



#### Note

Because support for `.sym` files is redundant with native symlink support available using `mklink`, use of `.sym` files is deprecated as of MySQL 5.6.9 and support for them will be removed in a future MySQL release.

The following limitations apply to the use of `.sym` files for database symbolic linking on Windows. These limitations do not apply for symlinks created using `mklink`.

- The symbolic link is not used if a directory with the same name as the database exists in the MySQL data directory.
- The `--innodb_file_per_table` [1718] option cannot be used.
- If you run `mysqld` as a service, you cannot use a mapped drive to a remote server as the destination of the symbolic link. As a workaround, you can use the full path (`\\servername\path\`).

## 8.11.4. Optimizing Memory Use

### 8.11.4.1. How MySQL Uses Memory

The following list indicates some of the ways that the `mysqld` server uses memory. Where applicable, the name of the system variable relevant to the memory use is given:

- All threads share the `MyISAM` key buffer; its size is determined by the `key_buffer_size` [492] variable. Other buffers used by the server are allocated as needed. See [Section 8.11.2, “Tuning Server Parameters”](#).
- Each thread that is used to manage client connections uses some thread-specific space. The following list indicates these and which variables control their size:



- A stack (variable `thread_stack` [556])
- A connection buffer (variable `net_buffer_length` [515])
- A result buffer (variable `net_buffer_length` [515])

The connection buffer and result buffer each begin with a size equal to `net_buffer_length` [515] bytes, but are dynamically enlarged up to `max_allowed_packet` [502] bytes as needed. The result buffer shrinks to `net_buffer_length` [515] bytes after each SQL statement. While a statement is running, a copy of the current statement string is also allocated.

- All threads share the same base memory.
- When a thread is no longer needed, the memory allocated to it is released and returned to the system unless the thread goes back into the thread cache. In that case, the memory remains allocated.
- The `myisam_use_mmap` [514] system variable can be set to 1 to enable memory-mapping for all `MyISAM` tables.
- Each request that performs a sequential scan of a table allocates a *read buffer* (variable `read_buffer_size` [532]).
- When reading rows in an arbitrary sequence (for example, following a sort), a *random-read buffer* (variable `read_rnd_buffer_size` [533]) may be allocated to avoid disk seeks.
- All joins are executed in a single pass, and most joins can be done without even using a temporary table. Most temporary tables are memory-based hash tables. Temporary tables with a large row length (calculated as the sum of all column lengths) or that contain `BLOB` columns are stored on disk.

If an internal in-memory temporary table becomes too large, MySQL handles this automatically by changing the table from in-memory to on-disk format, to be handled by the `MyISAM` storage engine. You can increase the permissible temporary table size as described in [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).

- Most requests that perform a sort allocate a sort buffer and zero to two temporary files depending on the result set size. See [Section C.5.4.4, “Where MySQL Stores Temporary Files”](#).
- Almost all parsing and calculating is done in thread-local and reusable memory pools. No memory overhead is needed for small items, so the normal slow memory allocation and freeing is avoided. Memory is allocated only for unexpectedly large strings.
- For each `MyISAM` table that is opened, the index file is opened once; the data file is opened once for each concurrently running thread. For each concurrent thread, a table structure, column structures for each column, and a buffer of size  $3 * N$  are allocated (where  $N$  is the maximum row length, not counting `BLOB` columns). A `BLOB` column requires five to eight bytes plus the length of the `BLOB` data. The `MyISAM` storage engine maintains one extra row buffer for internal use.
- For each table having `BLOB` columns, a buffer is enlarged dynamically to read in larger `BLOB` values. If you scan a table, a buffer as large as the largest `BLOB` value is allocated.
- Handler structures for all in-use tables are saved in a cache and managed as a FIFO. The initial cache size is taken from the value of the `table_open_cache` [551] system variable. If a table has been used by two running threads at the same time, the cache contains two entries for the table. See [Section 8.4.3.1, “How MySQL Opens and Closes Tables”](#).
- A `FLUSH TABLES` statement or `mysqladmin flush-tables` command closes all tables that are not in use at once and marks all in-use tables to be closed when the currently executing thread finishes. This effectively frees most in-use memory. `FLUSH TABLES` does not return until all tables have been closed.

- The server caches information in memory as a result of `GRANT`, `CREATE USER`, `CREATE SERVER`, and `INSTALL PLUGIN` statements. This memory is not released by the corresponding `REVOKE`, `DROP USER`, `DROP SERVER`, and `UNINSTALL PLUGIN` statements, so for a server that executes many instances of the statements that cause caching, there will be an increase in memory use. This cached memory can be freed with `FLUSH PRIVILEGES`.

`ps` and other system status programs may report that `mysqld` uses a lot of memory. This may be caused by thread stacks on different memory addresses. For example, the Solaris version of `ps` counts the unused memory between stacks as used memory. To verify this, check available swap with `swap -s`. We test `mysqld` with several memory-leakage detectors (both commercial and Open Source), so there should be no memory leaks.

#### 8.11.4.2. Enabling Large Page Support

Some hardware/operating system architectures support memory pages greater than the default (usually 4KB). The actual implementation of this support depends on the underlying hardware and operating system. Applications that perform a lot of memory accesses may obtain performance improvements by using large pages due to reduced Translation Lookaside Buffer (TLB) misses.

In MySQL, large pages can be used by InnoDB, to allocate memory for its buffer pool and additional memory pool.

Standard use of large pages in MySQL attempts to use the largest size supported, up to 4MB. Under Solaris, a “super large pages” feature enables uses of pages up to 256MB. This feature is available for recent SPARC platforms. It can be enabled or disabled by using the `--super-large-pages` [440] or `--skip-super-large-pages` [440] option.

MySQL also supports the Linux implementation of large page support (which is called HugeTLB in Linux).

Before large pages can be used on Linux, the kernel must be enabled to support them and it is necessary to configure the HugeTLB memory pool. For reference, the HugeTBL API is documented in the `Documentation/vm/hugetlbpage.txt` file of your Linux sources.

The kernel for some recent systems such as Red Hat Enterprise Linux appear to have the large pages feature enabled by default. To check whether this is true for your kernel, use the following command and look for output lines containing “huge”:

```
shell> cat /proc/meminfo | grep -i huge
HugePages_Total:      0
HugePages_Free:       0
HugePages_Rsvd:       0
HugePages_Surp:       0
Hugepagesize:         4096 kB
```

The nonempty command output indicates that large page support is present, but the zero values indicate that no pages are configured for use.

If your kernel needs to be reconfigured to support large pages, consult the `hugetlbpage.txt` file for instructions.

Assuming that your Linux kernel has large page support enabled, configure it for use by MySQL using the following commands. Normally, you put these in an `rc` file or equivalent startup file that is executed during the system boot sequence, so that the commands execute each time the system starts. The commands should execute early in the boot sequence, before the MySQL server starts. Be sure to change the allocation numbers and the group number as appropriate for your system.

```
# Set the number of pages to be used.
# Each page is normally 2MB, so a value of 20 = 40MB.
# This command actually allocates memory, so this much
# memory must be available.
```

```
echo 20 > /proc/sys/vm/nr_hugepages

# Set the group number that is permitted to access this
# memory (102 in this case). The mysql user must be a
# member of this group.
echo 102 > /proc/sys/vm/hugetlb_shm_group

# Increase the amount of shmem permitted per segment
# (12G in this case).
echo 1560281088 > /proc/sys/kernel/shmmax

# Increase total amount of shared memory. The value
# is the number of pages. At 4KB/page, 4194304 = 16GB.
echo 4194304 > /proc/sys/kernel/shmall
```

For MySQL usage, you normally want the value of `shmmax` to be close to the value of `shmall`.

To verify the large page configuration, check `/proc/meminfo` again as described previously. Now you should see some nonzero values:

```
shell> cat /proc/meminfo | grep -i huge
HugePages_Total:      20
HugePages_Free:       20
HugePages_Rsvd:       0
HugePages_Surp:       0
Hugepagesize:         4096 kB
```

The final step to make use of the `hugetlb_shm_group` is to give the `mysql` user an “unlimited” value for the memlock limit. This can be done either by editing `/etc/security/limits.conf` or by adding the following command to your `mysqld_safe` script:

```
ulimit -l unlimited
```

Adding the `ulimit` command to `mysqld_safe` causes the `root` user to set the memlock limit to `unlimited` before switching to the `mysql` user. (This assumes that `mysqld_safe` is started by `root`.)

Large page support in MySQL is disabled by default. To enable it, start the server with the `--large-pages [425]` option. For example, you can use the following lines in your server's `my.cnf` file:

```
[mysqld]
large-pages
```

With this option, `InnoDB` uses large pages automatically for its buffer pool and additional memory pool. If `InnoDB` cannot do this, it falls back to use of traditional memory and writes a warning to the error log: `Warning: Using conventional memory pool`

To verify that large pages are being used, check `/proc/meminfo` again:

```
shell> cat /proc/meminfo | grep -i huge
HugePages_Total:      20
HugePages_Free:       20
HugePages_Rsvd:       2
HugePages_Surp:       0
Hugepagesize:         4096 kB
```

## 8.11.5. Optimizing Network Use

### 8.11.5.1. How MySQL Uses Threads for Client Connections

Connection manager threads handle client connection requests on the network interfaces that the server listens to. On all platforms, one manager thread handles TCP/IP connection requests. On Unix, this manager thread also handles Unix socket file connection requests. On Windows, a manager thread

handles shared-memory connection requests, and another handles named-pipe connection requests. The server does not create threads to handle interfaces that it does not listen to. For example, a Windows server that does not have support for named-pipe connections enabled does not create a thread to handle them.

Connection manager threads associate each client connection with a thread dedicated to it that handles authentication and request processing for that connection. Manager threads create a new thread when necessary but try to avoid doing so by consulting the thread cache first to see whether it contains a thread that can be used for the connection. When a connection ends, its thread is returned to the thread cache if the cache is not full.

In this connection thread model, there are as many threads as there are clients currently connected, which has some disadvantages when server workload must scale to handle large numbers of connections. For example, thread creation and disposal becomes expensive. Also, each thread requires server and kernel resources, such as stack space. To accommodate a large number of simultaneous connections, the stack size per thread must be kept small, leading to a situation where it is either too small or the server consumes large amounts of memory. Exhaustion of other resources can occur as well, and scheduling overhead can become significant.

As of MySQL 5.6.10, commercial distributions of MySQL 5.6 include a thread pool plugin that provides an alternative thread-handling model designed to reduce overhead and improve performance. It implements a thread pool that increases server performance by efficiently managing statement execution threads for large numbers of client connections. See [Section 8.11.6, “The Thread Pool Plugin”](#).

To control and monitor how the server manages threads that handle client connections, several system and status variables are relevant. (See [Section 5.1.4, “Server System Variables”](#), and [Section 5.1.6, “Server Status Variables”](#).)

The thread cache has a size determined by the `thread_cache_size` [552] system variable. The default value is 0 (no caching), which causes a thread to be set up for each new connection and disposed of when the connection terminates. Set `thread_cache_size` [552] to *N* to enable *N* inactive connection threads to be cached. `thread_cache_size` [552] can be set at server startup or changed while the server runs. A connection thread becomes inactive when the client connection with which it was associated terminates.

To monitor the number of threads in the cache and how many threads have been created because a thread could not be taken from the cache, monitor the `Threads_cached` [601] and `Threads_created` [601] status variables.

You can set `max_connections` [503] at server startup or at runtime to control the maximum number of clients that can connect simultaneously.

When the thread stack is too small, this limits the complexity of the SQL statements which the server can handle, the recursion depth of stored procedures, and other memory-consuming actions. To set a stack size of *N* bytes for each thread, start the server with `--thread_stack=N` [556].

### 8.11.5.2. DNS Lookup Optimization and the Host Cache

The MySQL server maintains a host cache in memory that contains information about clients: IP address, host name, and error information. The server uses this cache for nonlocal TCP connections. It does not use the cache for TCP connections established using a loopback interface address (`127.0.0.1` or `:::1`), or for connections established using a Unix socket file, named pipe, or shared memory.

For each new client connection, the server uses the client IP address to check whether the client host name is in the host cache. If not, the server attempts to resolve the host name. First, it resolves the IP address to a host name and resolves that host name back to an IP address. Then it compares the result to the original IP address to ensure that they are the same. The server stores information

about the result of this operation in the host cache. If the cache is full, the least recently used entry is discarded.

The `host_cache` Performance Schema table exposes the contents of the host cache so that it can be examined using `SELECT` statements. This may help you diagnose the causes of connection problems. See [Section 20.9.9.1, “The `host\_cache` Table”](#).

The server handles entries in the host cache like this:

1. When the first TCP client connection reaches the server from a given IP address, a new entry is created to record the client IP, host name, and client lookup validation flag. Initially, the host name is set to `NULL` and the flag is false. This entry is also used for subsequent client connections from the same originating IP.
2. If the validation flag for the client IP entry is false, the server attempts an IP-to-host name DNS resolution. If that is successful, the host name is updated with the resolved host name and the validation flag is set to true. If resolution is unsuccessful, the action taken depends on whether the error is permanent or transient. For permanent failures, the host name remains `NULL` and the validation flag is set to true. For transient failures, the host name and validation flag remain unchanged. (Another DNS resolution attempt occurs the next time a client connects from this IP.)
3. If an error occurs while processing an incoming client connection from a given IP address, the server updates the corresponding error counters in the entry for that IP. For a description of the errors recorded, see [Section 20.9.9.1, “The `host\_cache` Table”](#).

The server performs host name resolution using the thread-safe `gethostbyaddr_r()` and `gethostbyname_r()` calls if the operating system supports them. Otherwise, the thread performing the lookup locks a mutex and calls `gethostbyaddr()` and `gethostbyname()` instead. In this case, no other thread can resolve host names that are not in the host cache until the thread holding the mutex lock releases it.

The server uses the host cache for several purposes:

- By caching the results of IP-to-host name lookups, the server avoids doing a DNS lookup for each client connection. Instead, for a given host, it needs to perform a lookup only for the first connection from that host.
- The cache contains information about errors that occur during the connection process. Some errors are considered “blocking.” If too many of these occur successively from a given host without a successful connection, the server blocks further connections from that host. The `max_connect_errors` [502] system variable determines the number of permitted errors before blocking occurs. See [Section C.5.2.6, “Host ‘`host\_name`’ is blocked”](#).

To unblock blocked hosts, flush the host cache by issuing a `FLUSH HOSTS` statement or executing a `mysqladmin flush-hosts` command.

It is possible for a blocked host to become unblocked even without `FLUSH HOSTS` if activity from other hosts has occurred since the last connection attempt from the blocked host. This can occur because the server discards the least recently used cache entry to make room for a new entry if the cache is full when a connection arrives from a client IP not in the cache. If the discarded entry is for a blocked host, that host becomes unblocked.

The host cache is enabled by default. To disable it, set the `host_cache_size` [487] system variable to 0, either at server startup or at runtime.

To disable DNS host name lookups, start the server with the `--skip-name-resolve` [440] option. In this case, the server uses only IP addresses and not host names to match connecting hosts to rows in the MySQL grant tables. Only accounts specified in those tables using IP addresses can be used.

If you have a very slow DNS and many hosts, you might be able to improve performance either by disabling DNS lookups with `--skip-name-resolve` [440] or by increasing the value of `host_cache_size` [487] to make the host cache larger.

To disallow TCP/IP connections entirely, start the server with the `--skip-networking` [440] option.

Some connection errors are not associated with TCP connections, occur very early in the connection process (even before an IP address is known), or are not specific to any particular IP address (such as out-of-memory conditions). For information about these errors, check the `Connection_errors_xxx` [587] status variables (see Section 5.1.6, “Server Status Variables”).

## 8.11.6. The Thread Pool Plugin



### Note

MySQL Thread Pool is a commercial extension. To learn more about commercial products (MySQL Enterprise Edition), see <http://www.mysql.com/products/>.

As of MySQL 5.6.10, commercial distributions of MySQL 5.6 include MySQL Thread Pool, implemented using a server plugin. The default thread-handling model in MySQL Server executes statements using one thread per client connection. As more clients connect to the server and execute statements, overall performance degrades. The thread pool plugin provides an alternative thread-handling model designed to reduce overhead and improve performance. The plugin implements a thread pool that increases server performance by efficiently managing statement execution threads for large numbers of client connections.

The thread pool addresses several problems of the one thread per connection model:

- Too many thread stacks make CPU caches almost useless in highly parallel execution workloads. The thread pool promotes thread stack reuse to minimize the CPU cache footprint.
- With too many threads executing in parallel, context switching overhead is high. This also presents a challenging task to the operating system scheduler. The thread pool controls the number of active threads to keep the parallelism within the MySQL server at a level that it can handle and that is appropriate for the server host on which MySQL is executing.
- Too many transactions executing in parallel increases resource contention. In `InnoDB`, this increases the time spent holding central mutexes. The thread pool controls when transactions start to ensure that not too many execute in parallel.

The thread pool plugin is a commercial feature. It is not included in MySQL community distributions.

On Windows, the thread pool plugin requires Windows Vista or newer. On Linux, the plugin requires kernel 2.6.9 or newer.

## Additional Resources

[Section B.14, “MySQL 5.6 FAQ: MySQL Enterprise Scalability Thread Pool”](#)

### 8.11.6.1. Thread Pool Components and Installation

The thread pool feature comprises these components:

- A plugin library object file contains a plugin for the thread pool code and plugins for several `INFORMATION_SCHEMA` tables.

For a detailed description of how the thread pool works, see [Section 8.11.6.2, “Thread Pool Operation”](#).

The `INFORMATION_SCHEMA` tables are named `TP_THREAD_STATE`, `TP_THREAD_GROUP_STATE`, and `TP_THREAD_GROUP_STATS`. These tables provide information about thread pool operation. For more information, see [Section 19.31, “Thread Pool `INFORMATION\_SCHEMA` Tables”](#).



- Several system variables are related to the thread pool. The `thread_handling` [553] system variable has a value of `loaded-dynamically` when the server successfully loads the thread pool plugin.

The other related variables are implemented by the thread pool plugin; they are not available unless it is enabled:

- `thread_pool_algorithm` [554]: The concurrency algorithm to use for scheduling.
- `thread_pool_high_priority_connection` [554]: How to schedule statement execution for a session.
- `thread_pool_prio_kickup_timer` [555]: How long before the thread pool moves a statement awaiting execution from the low-priority queue to the high-priority queue.
- `thread_pool_max_unused_threads` [555]: How many sleeping threads to permit.
- `thread_pool_size` [556]: The number of thread groups in the thread pool. This is the most important parameter controlling thread pool performance.
- `thread_pool_stall_limit` [556]: The time before an executing statement is considered to be stalled.

If any variable implemented by the plugin is set to an illegal value at startup, plugin initialization fails and the plugin does not load.

For information about setting thread pool parameters, see [Section 8.11.6.3, “Thread Pool Tuning”](#).

- The Performance Schema exposes information about the thread pool and may be used to investigate operational performance. For more information, see [Chapter 20, MySQL Performance Schema](#).

To be usable by the server, the thread pool library object file must be located in the MySQL plugin directory (the directory named by the `plugin_dir` [525] system variable). To enable thread pool capability, load the plugins to be used by starting the server with the `--plugin-load` [435] option. For example, if you name just the plugin object file, the server loads all plugins that it contains (that is, the thread pool plugin and all the `INFORMATION_SCHEMA` tables). To do this, put these lines in your `my.cnf` file:

```
[mysqld]
plugin-load=thread_pool.so
```

If object files have a suffix different from `.so` on your system, substitute the correct suffix (for example, `.dll` on Windows).

If necessary, set the value of the `plugin_dir` [525] system variable to tell the server the location of the plugin directory.

If desired, you can load individual plugins from the library file. To load the thread pool plugin but not the `INFORMATION_SCHEMA` tables, use an option like this:

```
[mysqld]
plugin-load=thread_pool=thread_pool.so
```

To load the thread pool plugin and only the `TP_THREAD_STATE INFORMATION_SCHEMA` table, use an option like this:

```
[mysqld]
plugin-load=thread_pool=thread_pool.so;TP_THREAD_STATE=thread_pool.so
```

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement. See [Section 5.1.8.2, “Obtaining Server Plugin Information”](#).



If the server loads the thread plugin successfully, it sets the `thread_handling` system variable to `dynamically-loaded`. If the plugin fails to load, the server writes a message to the error log.

### 8.11.6.2. Thread Pool Operation

The thread pool consists of a number of thread groups, each of which manages a set of client connections. As connections are established, the thread pool assigns them to thread groups in round-robin fashion.

The number of thread groups is configurable using the `thread_pool_size` [556] system variable. The default number of groups is 16. For guidelines on setting this variable, see [Section 8.11.6.3, “Thread Pool Tuning”](#).

The maximum number of threads per group is 4096 (or 4095 on some systems where one thread is used internally).

The thread pool separates connections and threads, so there is no fixed relationship between connections and the threads that execute statements received from those connections. This differs from the default thread-handling model that associates one thread with one connection such that the thread executes all statements from the connection.

The thread pool tries to ensure a maximum of one thread executing in each group at any time, but sometimes permits more threads to execute temporarily for best performance. The algorithm works in the following manner:

- Each thread group has a listener thread that listens for incoming statements from the connections assigned to the group. When a statement arrives, the thread group either begins executing it immediately or queues it for later execution:
  - Immediate execution occurs if the statement is the only one received and no statements are queued or currently executing.
  - Queuing occurs if the statement cannot begin executing immediately.
- If immediate execution occurs, execution is performed by the listener thread. (This means that temporarily no thread in the group is listening.) If the statement finishes quickly, the executing thread returns to listening for statements. Otherwise, the thread pool considers the statement stalled and starts another thread as a listener thread (creating it if necessary). To ensure that no thread group becomes blocked by stalled statements, the thread pool has a background thread that regularly monitors thread group states.

By using the listening thread to execute a statement that can begin immediately, there is no need to create an additional thread if the statement finishes quickly. This ensures the most efficient execution possible in the case of a low number of concurrent threads.

When the thread pool plugin starts, it creates one thread per group (the listener thread), plus the background thread. Additional threads are created as necessary to execute statements.

- The value of the `thread_pool_stall_limit` [556] system variable determines the meaning of “finishes quickly” in the previous item. The default time before threads are considered stalled is 60ms but can be set to a maximum of 6s. This parameter is configurable to enable you to strike a balance appropriate for the server work load. Short wait values permit threads to start more quickly. Short values are also better for avoiding deadlock situations. Long wait values are useful for workloads that include long-running statements, to avoid starting too many new statements while the current ones execute.
- The thread pool focuses on limiting the number of concurrent short-running statements. Before an executing statement reaches the stall time, it prevents other statements from beginning to execute. If the statement executes past the stall time, it is permitted to continue but no longer prevents other statements from starting. In this way, the thread pool tries to ensure that in each thread group there

is never more than one short-running statement, although there might be multiple long-running statements. It is undesirable to let long-running statements prevent other statements from executing because there is no limit on the amount of waiting that might be necessary. For example, on a replication master, a thread that is sending binary log events to a slave effectively runs forever.

- A statement becomes blocked if it encounters a disk I/O operation or a user level lock (row lock or table lock). The block would cause the thread group to become unused, so there are callbacks to the thread pool to ensure that the thread pool can immediately start a new thread in this group to execute another statement. When a blocked thread returns, the thread pool permits it to restart immediately.
- There are two queues, a high-priority queue and a low-priority queue. The first statement in a transaction goes to the low-priority queue. Any following statements for the transaction go to the high-priority queue if the transaction is ongoing (statements for it have begun executing), or to the low-priority queue otherwise. Queue assignment can be affected by enabling the `thread_pool_high_priority_connection` [554] system variable, which causes all queued statements for a session to go into the high-priority queue.

Statements for a nontransactional storage engine, or a transactional engine if `autocommit` [464] is enabled, are treated as low-priority statements because in this case each statement is a transaction. Thus, given a mix of statements for `InnoDB` and `MyISAM` tables, the thread pool prioritizes those for `InnoDB` over those for `MyISAM` unless `autocommit` [464] is enabled. With `autocommit` [464] enabled, all statements will be low priority.

- When the thread group selects a queued statement for execution, it first looks in the high-priority queue, then in the low-priority queue. If a statement is found, it is removed from its queue and begins to execute.
- If a statement stays in the low-priority queue too long, the thread pool moves to the high-priority queue. The value of the `thread_pool_prio_kickup_timer` [555] system variable controls the time before movement. For each thread group, a maximum of one statement per 10ms or 100 per second will be moved from the low-priority queue to the high-priority queue.
- The thread pool reuses the most active threads to obtain a much better use of CPU caches. This is a small adjustment that has a great impact on performance.
- While a thread executes a statement from a user connection, Performance Schema instrumentation accounts thread activity to the user connection. Otherwise, Performance Schema accounts activity to the thread pool.

Here are examples of conditions under which a thread group might have multiple threads started to execute statements:

- One thread begins executing a statement, but runs long enough to be considered stalled. The thread group permits another thread to begin executing another statement even though the first thread is still executing.
- One thread begins executing a statement, then becomes blocked and reports this back to the thread pool. The thread group permits another thread to begin executing another statement.
- One thread begins executing a statement, becomes blocked, but does not report back that it is blocked because the block does not occur in code that has been instrumented with thread pool callbacks. In this case, the thread appears to the thread group to be still running. If the block lasts long enough for the statement to be considered stalled, the group permits another thread to begin executing another statement.

The thread pool is designed to be scalable across an increasing number of connections. It is also designed to avoid deadlocks that can arise from limiting the number of actively executing statements. It is important that threads that do not report back to the thread pool do not prevent other statements from executing and thus cause the thread pool to become deadlocked. Examples of such statements follow:

- Long-running statements. These would lead to all resources used by only a few statements and they could prevent all others from accessing the server.
- Binary log dump threads that read the binary log and send it to slaves. This is a kind of long-running “statement” that runs for a very long time, and that should not prevent other statements from executing.
- Statements blocked on a row lock, table lock, sleep, or any other blocking activity that has not been reported back to the thread pool by MySQL Server or a storage engine.

In each case, to prevent deadlock, the statement is moved to the stalled category when it does not complete quickly, so that the thread group can permit another statement to begin executing. With this design, when a thread executes or becomes blocked for an extended time, the thread pool moves the thread to the stalled category and for the rest of the statement's execution, it does not prevent other statements from executing.

The maximum number of threads that can occur is the sum of `max_connections` [503] and `thread_pool_size` [556]. This can happen in a situation where all connections are in execution mode and an extra thread is created per group to listen for more statements. This is not necessarily a state that happens often, but it is theoretically possible.

### 8.11.6.3. Thread Pool Tuning

This section provides guidelines on setting thread pool system variables for best performance, measured using a metric such as transactions per second.

`thread_pool_size` [556] is the most important parameter controlling thread pool performance. It can be set only at server startup. Our experience in testing the thread pool indicates the following:

- If the primary storage engine is `InnoDB`, the optimal `thread_pool_size` [556] setting is likely to be between 16 and 36, with the most common optimal values tending to be from 24 to 36. We have not seen any situation where the setting has been optimal beyond 36. There may be special cases where a value smaller than 16 is optimal.

For workloads such as DBT2 and Sysbench, the optimum for `InnoDB` seems to be usually around 36. For very write-intensive workloads, the optimal setting can sometimes be lower.

- If the primary storage engine is `MyISAM`, the `thread_pool_size` [556] setting should be fairly low. We tend to get optimal performance for values from 4 to 8. Higher values tend to have a slightly negative but not dramatic impact on performance.

Another system variable, `thread_pool_stall_limit` [556], is important for handling of blocked and long-running statements. If all calls that block the MySQL Server are reported to the thread pool, it would always know when execution threads are blocked. However, this may not always be true. For example, blocks could occur in code that has not been instrumented with thread pool callbacks. For such cases, the thread pool must be able to identify threads that appear to be blocked. This is done by means of a timeout, the length of which can be tuned using the `thread_pool_stall_limit` [556] system variable. This parameter ensures that the server does not become completely blocked. The value of `thread_pool_stall_limit` [556] has an upper limit of 6 seconds to prevent the risk of a deadlocked server.

`thread_pool_stall_limit` [556] also enables the thread pool to handle long-running statements. If a long-running statement was permitted to block a thread group, all other connections assigned to the group would be blocked and unable to start execution until the long-running statement completed. In the worst case, this could take hours or even days.

The value of `thread_pool_stall_limit` [556] should be chosen such that statements that execute longer than its value are considered stalled. Stalled statements generate a lot of extra overhead since they involve extra context switches and in some cases even extra thread creations. On the other hand, setting the `thread_pool_stall_limit` [556] parameter too high means that long-

running statements will block a number of short-running statements for longer than necessary. Short wait values permit threads to start more quickly. Short values are also better for avoiding deadlock situations. Long wait values are useful for workloads that include long-running statements, to avoid starting too many new statements while the current ones execute.

Suppose a server executes a workload where 99.9% of the statements complete within 100ms even when the server is loaded, and the remaining statements take between 100ms and 2 hours fairly evenly spread. In this case, it would make sense to set `thread_pool_stall_limit` [556] to 10 (meaning 100ms). The default value of 60ms is okay for servers that primarily execute very simple statements.

The `thread_pool_stall_limit` [556] parameter can be changed at runtime to enable you to strike a balance appropriate for the server work load. Assuming that the `TP_THREAD_GROUP_STATS` table is enabled, you can use the following query to determine the fraction of executed statements that stalled:

```
SELECT SUM(STALLED_QUERIES_EXECUTED) / SUM(QUERIES_EXECUTED)
FROM information_schema.TP_THREAD_GROUP_STATS;
```

This number should be as low as possible. To decrease the likelihood of statements stalling, increase the value of `thread_pool_stall_limit` [556].

When a statement arrives, what is the maximum time it can be delayed before it actually starts executing? Suppose that the following conditions apply:

- There are 200 statements queued in the low-priority queue.
- There are 10 statements queued in the high-priority queue.
- `thread_pool_prio_kickup_timer` [555] is set to 10000 (10 seconds).
- `thread_pool_stall_limit` [556] is set to 100 (1 second).

In the worst case, the 10 high-priority statements represent 10 transactions that continue executing for a long time. Thus, in the worst case, no statements will be moved to the high-priority queue because it will always already contain statements awaiting execution. After 10 seconds, the new statement is eligible to be moved to the high-priority queue. However, before it can be moved, all the statements before it must be moved as well. This could take another 2 seconds because a maximum of 100 statements per second are moved to the high-priority queue. Now when the statement reaches the high-priority queue, there could potentially be many long-running statements ahead of it. In the worst case, every one of those will become stalled and it will take 1 second for each statement before the next statement is retrieved from the high-priority queue. Thus, in this scenario, it will take 222 seconds before the new statement starts executing.

This example shows a worst case for an application. How to handle it depends on the application. If the application has high requirements for the response time, it should most likely throttle users at a higher level itself. Otherwise, it can use the thread pool configuration parameters to set some kind of a maximum waiting time.

## 8.12. Measuring Performance (Benchmarking)

To measure performance, consider the following factors:

- Whether you are measuring the speed of a single operation on a quiet system, or how a set of operations (a “workload”) works over a period of time. With simple tests, you usually test how changing one aspect (a configuration setting, the set of indexes on a table, the SQL clauses in a query) affects performance. Benchmarks are typically long-running and elaborate performance tests, where the results could dictate high-level choices such as hardware and storage configuration, or how soon to upgrade to a new MySQL version.
- For benchmarking, sometimes you must simulate a heavy database workload to get an accurate picture.

- Performance can vary depending on so many different factors that a difference of a few percentage points might not be a decisive victory. The results might shift the opposite way when you test in a different environment.
- Certain MySQL features help or do not help performance depending on the workload. For completeness, always test performance with those features turned on and turned off. The two most important features to try with each workload are the [MySQL query cache](#), and the [adaptive hash index](#) for [InnoDB](#) tables.

This section progresses from simple and direct measurement techniques that a single developer can do, to more complicated ones that require additional expertise to perform and interpret the results.

### 8.12.1. Measuring the Speed of Expressions and Functions

To measure the speed of a specific MySQL expression or function, invoke the `BENCHMARK()` [1266] function using the `mysql` client program. Its syntax is `BENCHMARK(loop_count, expression)` [1266]. The return value is always zero, but `mysql` prints a line displaying approximately how long the statement took to execute. For example:

```
mysql> SELECT BENCHMARK(1000000,1+1);
+-----+
| BENCHMARK(1000000,1+1) |
+-----+
|                          0 |
+-----+
1 row in set (0.32 sec)
```

This result was obtained on a Pentium II 400MHz system. It shows that MySQL can execute 1,000,000 simple addition expressions in 0.32 seconds on that system.

The built-in MySQL functions are typically highly optimized, but there may be some exceptions. `BENCHMARK()` [1266] is an excellent tool for finding out if some function is a problem for your queries.

### 8.12.2. The MySQL Benchmark Suite

This benchmark suite is meant to tell any user what operations a given SQL implementation performs well or poorly. You can get a good idea for how the benchmarks work by looking at the code and results in the `sql-bench` directory in any MySQL source distribution.

Note that this benchmark is single-threaded, so it measures the minimum time for the operations performed. We plan to add multi-threaded tests to the benchmark suite in the future.

To use the benchmark suite, the following requirements must be satisfied:

- The benchmark suite is provided with MySQL source distributions. You can either download a released distribution from <http://dev.mysql.com/downloads/>, or use the current development source tree. (See [Section 2.9.3, “Installing MySQL from a Development Source Tree”](#).)
- The benchmark scripts are written in Perl and use the Perl DBI module to access database servers, so DBI must be installed. You also need the server-specific DBD drivers for each of the servers you want to test. For example, to test MySQL, PostgreSQL, and DB2, you must have the `DBD::mysql`, `DBD::Pg`, and `DBD::DB2` modules installed. See [Section 2.13, “Perl Installation Notes”](#).

After you obtain a MySQL source distribution, you can find the benchmark suite located in its `sql-bench` directory. To run the benchmark tests, build MySQL, and then change location into the `sql-bench` directory and execute the `run-all-tests` script:

```
shell> cd sql-bench
shell> perl run-all-tests --server=server_name
```

`server_name` should be the name of one of the supported servers. To get a list of all options and supported servers, invoke this command:

```
shell> perl run-all-tests --help
```

The `crash-me` script also is located in the `sql-bench` directory. `crash-me` tries to determine what features a database system supports and what its capabilities and limitations are by actually running queries. For example, it determines:

- What data types are supported
- How many indexes are supported
- What functions are supported
- How big a query can be
- How big a `VARCHAR` column can be

For more information about benchmark results, visit <http://www.mysql.com/why-mysql/benchmarks/>.

### 8.12.3. Using Your Own Benchmarks

Benchmark your application and database to find out where the bottlenecks are. After fixing one bottleneck (or by replacing it with a “dummy” module), you can proceed to identify the next bottleneck. Even if the overall performance for your application currently is acceptable, you should at least make a plan for each bottleneck and decide how to solve it if someday you really need the extra performance.

For examples of portable benchmark programs, look at those in the MySQL benchmark suite. See [Section 8.12.2, “The MySQL Benchmark Suite”](#). You can take any program from this suite and modify it for your own needs. By doing this, you can try different solutions to your problem and test which really is fastest for you.

Another free benchmark suite is the Open Source Database Benchmark, available at <http://osdb.sourceforge.net/>.

It is very common for a problem to occur only when the system is very heavily loaded. We have had many customers who contact us when they have a (tested) system in production and have encountered load problems. In most cases, performance problems turn out to be due to issues of basic database design (for example, table scans are not good under high load) or problems with the operating system or libraries. Most of the time, these problems would be much easier to fix if the systems were not already in production.

To avoid problems like this, benchmark your whole application under the worst possible load:

- The `mysqlslap` program can be helpful for simulating a high load produced by multiple clients issuing queries simultaneously. See [Section 4.5.7, “mysqlslap — Load Emulation Client”](#).
- You can also try benchmarking packages such as SysBench and DBT2, available at <http://sourceforge.net/projects/sysbench/>, and <http://osldb.sourceforge.net/#dbt2>.

These programs or packages can bring a system to its knees, so be sure to use them only on your development systems.

### 8.12.4. Measuring Performance with `performance_schema`

You can query the tables in the `performance_schema` database to see real-time information about the performance characteristics of your server and the applications it is running. See [Chapter 20, \*MySQL Performance Schema\*](#) for details.

### 8.12.5. Examining Thread Information



When you are attempting to ascertain what your MySQL server is doing, it can be helpful to examine the process list, which is the set of threads currently executing within the server. Process list information is available from these sources:

- The `SHOW [FULL] PROCESSLIST` statement: [Section 13.7.5.30, “SHOW PROCESSLIST Syntax”](#)
- The `SHOW PROFILE` statement: [Section 13.7.5.32, “SHOW PROFILES Syntax”](#)
- The `INFORMATION_SCHEMA.PROCESSLIST` table: [Section 19.16, “The INFORMATION\\_SCHEMA PROCESSLIST Table”](#)
- The `mysqladmin processlist` command: [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#)
- The `performance_schema.threads` table: [Section 20.9.9, “Performance Schema Miscellaneous Tables”](#)

Access to `threads` does not require a mutex and has minimal impact on server performance. `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` have negative performance consequences because they require a mutex. `threads` also shows information about background threads, which `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` do not. This means that `threads` can be used to monitor activity the other thread information sources cannot.

You can always view information about your own threads. To view information about threads being executed for other accounts, you must have the `PROCESS [755]` privilege.

Each process list entry contains several pieces of information:

- `Id` is the connection identifier for the client associated with the thread.
- `User` and `Host` indicate the account associated with the thread.
- `db` is the default database for the thread, or `NULL` if none is selected.
- `Command` and `State` indicate what the thread is doing.

Most states correspond to very quick operations. If a thread stays in a given state for many seconds, there might be a problem that needs to be investigated.

- `Time` indicates how long the thread has been in its current state. The thread's notion of the current time may be altered in some cases: The thread can change the time with `SET TIMESTAMP = value`. For a thread running on a slave that is processing events from the master, the thread time is set to the time found in the events and thus reflects current time on the master and not the slave.
- `Info` contains the text of the statement being executed by the thread, or `NULL` if it is not executing one. By default, this value contains only the first 100 characters of the statement. To see the complete statements, use `SHOW FULL PROCESSLIST`.

The following sections list the possible `Command` values, and `State` values grouped by category. The meaning for some of these values is self-evident. For others, additional description is provided.

### 8.12.5.1. Thread Command Values

A thread can have any of the following `Command` values:

- `Binlog Dump`
- `Change user`

This is a thread on a master server for sending binary log contents to a slave server.

The thread is executing a change-user operation.



- `Close stmt`  
The thread is closing a prepared statement.
- `Connect`  
A replication slave is connected to its master.
- `Connect Out`  
A replication slave is connecting to its master.
- `Create DB`  
The thread is executing a create-database operation.
- `Daemon`  
This thread is internal to the server, not a thread that services a client connection.
- `Debug`  
The thread is generating debugging information.
- `Delayed insert`  
The thread is a delayed-insert handler.
- `Drop DB`  
The thread is executing a drop-database operation.
- `Error`
- `Execute`  
The thread is executing a prepared statement.
- `Fetch`  
The thread is fetching the results from executing a prepared statement.
- `Field List`  
The thread is retrieving information for table columns.
- `Init DB`  
The thread is selecting a default database.
- `Kill`  
The thread is killing another thread.
- `Long Data`  
The thread is retrieving long data in the result of executing a prepared statement.
- `Ping`  
The thread is handling a server-ping request.
- `Prepare`  
The thread is preparing a prepared statement.

- `Processlist`  
The thread is producing information about server threads.
- `Query`  
The thread is executing a statement.
- `Quit`  
The thread is terminating.
- `Refresh`  
The thread is flushing table, logs, or caches, or resetting status variable or replication server information.
- `Register Slave`  
The thread is registering a slave server.
- `Reset stmt`  
The thread is resetting a prepared statement.
- `Set option`  
The thread is setting or resetting a client statement-execution option.
- `Shutdown`  
The thread is shutting down the server.
- `Sleep`  
The thread is waiting for the client to send a new statement to it.
- `Statistics`  
The thread is producing server-status information.
- `Table Dump`  
The thread is sending table contents to a slave server.
- `Time`  
Unused.

### 8.12.5.2. General Thread States

The following list describes thread `State` values that are associated with general query processing and not more specialized activities such as replication. Many of these are useful only for finding bugs in the server.

- `After create`  
This occurs when the thread creates a table (including internal temporary tables), at the end of the function that creates the table. This state is used even if the table could not be created due to some error.
- `altering table`  
The server is in the process of executing an in-place `ALTER TABLE`.

- `Analyzing`

The thread is calculating a `MyISAM` table key distributions (for example, for `ANALYZE TABLE`).

- `checking permissions`

The thread is checking whether the server has the required privileges to execute the statement.

- `Checking table`

The thread is performing a table check operation.

- `cleaning up`

The thread has processed one command and is preparing to free memory and reset certain state variables.

- `closing tables`

The thread is flushing the changed table data to disk and closing the used tables. This should be a fast operation. If not, verify that you do not have a full disk and that the disk is not in very heavy use.

- `committing alter table to storage engine`

The server has finished an in-place `ALTER TABLE` and is committing the result.

- `converting HEAP to MyISAM`

The thread is converting an internal temporary table from a `MEMORY` table to an on-disk `MyISAM` table.

- `copy to tmp table`

The thread is processing an `ALTER TABLE` statement. This state occurs after the table with the new structure has been created but before rows are copied into it.

- `Copying to group table`

If a statement has different `ORDER BY` and `GROUP BY` criteria, the rows are sorted by group and copied to a temporary table.

- `Copying to tmp table`

The server is copying to a temporary table in memory.

- `Copying to tmp table on disk`

The server is copying to a temporary table on disk. The temporary result set has become too large (see [Section 8.4.3.3, "How MySQL Uses Internal Temporary Tables"](#)). Consequently, the thread is changing the temporary table from in-memory to disk-based format to save memory.

- `Creating index`

The thread is processing `ALTER TABLE ... ENABLE KEYS` for a `MyISAM` table.

- `Creating sort index`

The thread is processing a `SELECT` that is resolved using an internal temporary table.

- `creating table`

The thread is creating a table. This includes creation of temporary tables.

- `Creating tmp table`

The thread is creating a temporary table in memory or on disk. If the table is created in memory but later is converted to an on-disk table, the state during that operation will be `Copying to tmp table on disk`.

- `deleting from main table`

The server is executing the first part of a multiple-table delete. It is deleting only from the first table, and saving columns and offsets to be used for deleting from the other (reference) tables.

- `deleting from reference tables`

The server is executing the second part of a multiple-table delete and deleting the matched rows from the other tables.

- `discard_or_import_tablespace`

The thread is processing an `ALTER TABLE ... DISCARD TABLESPACE` or `ALTER TABLE ... IMPORT TABLESPACE` statement.

- `end`

This occurs at the end but before the cleanup of `ALTER TABLE`, `CREATE VIEW`, `DELETE`, `INSERT`, `SELECT`, or `UPDATE` statements.

- `executing`

The thread has begun executing a statement.

- `Execution of init_command`

The thread is executing statements in the value of the `init_command` system variable.

- `freeing items`

The thread has executed a command. Some freeing of items done during this state involves the query cache. This state is usually followed by `cleaning up`.

- `Flushing tables`

The thread is executing `FLUSH TABLES` and is waiting for all threads to close their tables.

- `FULLTEXT initialization`

The server is preparing to perform a natural-language full-text search.

- `init`

This occurs before the initialization of `ALTER TABLE`, `DELETE`, `INSERT`, `SELECT`, or `UPDATE` statements. Actions taken by the server in this state include flushing the binary log, the `InnoDB` log, and some query cache cleanup operations.

For the `end` state, the following operations could be happening:

- Removing query cache entries after data in a table is changed
- Writing an event to the binary log
- Freeing memory buffers, including for blobs
- `Killed`

Someone has sent a `KILL` statement to the thread and it should abort next time it checks the kill flag. The flag is checked in each major loop in MySQL, but in some cases it might still take a short time

for the thread to die. If the thread is locked by some other thread, the kill takes effect as soon as the other thread releases its lock.

- `logging slow query`

The thread is writing a statement to the slow-query log.

- `NULL`

This state is used for the `SHOW PROCESSLIST` state.

- `login`

The initial state for a connection thread until the client has been authenticated successfully.

- `manage keys`

The server is enabling or disabling a table index.

- `Opening tables, Opening table`

The thread is trying to open a table. This is should be very fast procedure, unless something prevents opening. For example, an `ALTER TABLE` or a `LOCK TABLE` statement can prevent opening a table until the statement is finished. It is also worth checking that your `table_open_cache` [551] value is large enough.

- `optimizing`

The server is performing initial optimizations for a query.

- `preparing`

This state occurs during query optimization.

- `preparing for alter table`

The server is preparing to execute an in-place `ALTER TABLE`.

- `Purging old relay logs`

The thread is removing unneeded relay log files.

- `query end`

This state occurs after processing a query but before the `freeing items` state.

- `Reading from net`

The server is reading a packet from the network.

- `Removing duplicates`

The query was using `SELECT DISTINCT` in such a way that MySQL could not optimize away the distinct operation at an early stage. Because of this, MySQL requires an extra stage to remove all duplicated rows before sending the result to the client.

- `removing tmp table`

The thread is removing an internal temporary table after processing a `SELECT` statement. This state is not used if no temporary table was created.

- `rename`

The thread is renaming a table.

- `rename result table`

The thread is processing an `ALTER TABLE` statement, has created the new table, and is renaming it to replace the original table.

- `Reopen tables`

The thread got a lock for the table, but noticed after getting the lock that the underlying table structure changed. It has freed the lock, closed the table, and is trying to reopen it.

- `Repair by sorting`

The repair code is using a sort to create indexes.

- `Repair done`

The thread has completed a multi-threaded repair for a `MyISAM` table.

- `Repair with keycache`

The repair code is using creating keys one by one through the key cache. This is much slower than `Repair by sorting`.

- `Rolling back`

The thread is rolling back a transaction.

- `Saving state`

For `MyISAM` table operations such as repair or analysis, the thread is saving the new table state to the `.MYI` file header. State includes information such as number of rows, the `AUTO_INCREMENT` counter, and key distributions.

- `Searching rows for update`

The thread is doing a first phase to find all matching rows before updating them. This has to be done if the `UPDATE` is changing the index that is used to find the involved rows.

- `Sending data`

The thread is reading and processing rows for a `SELECT` statement, and sending data to the client. Because operations occurring during this state tend to perform large amounts of disk access (reads), it is often the longest-running state over the lifetime of a given query.

- `setup`

The thread is beginning an `ALTER TABLE` operation.

- `Sorting for group`

The thread is doing a sort to satisfy a `GROUP BY`.

- `Sorting for order`

The thread is doing a sort to satisfy a `ORDER BY`.

- `Sorting index`

The thread is sorting index pages for more efficient access during a `MyISAM` table optimization operation.

- `Sorting result`

For a `SELECT` statement, this is similar to `Creating sort index`, but for nontemporary tables.

- `statistics`

The server is calculating statistics to develop a query execution plan. If a thread is in this state for a long time, the server is probably disk-bound performing other work.

- `System lock`

The thread is going to request or is waiting for an internal or external system lock for the table. If this state is being caused by requests for external locks and you are not using multiple `mysqld` servers that are accessing the same `MyISAM` tables, you can disable external system locks with the `--skip-external-locking` [439] option. However, external locking is disabled by default, so it is likely that this option will have no effect. For `SHOW PROFILE`, this state means the thread is requesting the lock (not waiting for it).

- `update`

The thread is getting ready to start updating the table.

- `Updating`

The thread is searching for rows to update and is updating them.

- `updating main table`

The server is executing the first part of a multiple-table update. It is updating only the first table, and saving columns and offsets to be used for updating the other (reference) tables.

- `updating reference tables`

The server is executing the second part of a multiple-table update and updating the matched rows from the other tables.

- `User lock`

The thread is going to request or is waiting for an advisory lock requested with a `GET_LOCK()` [1275] call. For `SHOW PROFILE`, this state means the thread is requesting the lock (not waiting for it).

- `User sleep`

The thread has invoked a `SLEEP()` [1280] call.

- `Waiting for commit lock`

`FLUSH TABLES WITH READ LOCK`) is waiting for a commit lock.

- `Waiting for global read lock`

`FLUSH TABLES WITH READ LOCK`) is waiting for a global read lock.

- `Waiting for tables,Waiting for table flush`

The thread got a notification that the underlying structure for a table has changed and it needs to reopen the table to get the new structure. However, to reopen the table, it must wait until all other threads have closed the table in question.

This notification takes place if another thread has used `FLUSH TABLES` or one of the following statements on the table in question: `FLUSH TABLES tbl_name`, `ALTER TABLE`, `RENAME TABLE`, `REPAIR TABLE`, `ANALYZE TABLE`, or `OPTIMIZE TABLE`.

- `Waiting for lock_type lock`

The server is waiting to acquire a lock, where `lock_type` indicates the type of lock:



- `Waiting for event metadata lock`
- `Waiting for global read lock`
- `Waiting for schema metadata lock`
- `Waiting for stored function metadata lock`
- `Waiting for stored procedure metadata lock`
- `Waiting for table level lock`
- `Waiting for table metadata lock`
- `Waiting for trigger metadata lock`
- `Waiting on cond`

A generic state in which the thread is waiting for a condition to become true. No specific state information is available.

- `Writing to net`

The server is writing a packet to the network.

### 8.12.5.3. Delayed-Insert Thread States

These thread states are associated with processing for `DELAYED` inserts (see [Section 13.2.5.2, “INSERT DELAYED Syntax”](#)). Some states are associated with connection threads that process `INSERT DELAYED` statements from clients. Other states are associated with delayed-insert handler threads that insert the rows. There is a delayed-insert handler thread for each table for which `INSERT DELAYED` statements are issued.

States associated with a connection thread that processes an `INSERT DELAYED` statement from the client:

- `allocating local table`

The thread is preparing to feed rows to the delayed-insert handler thread.

- `Creating delayed handler`

The thread is creating a handler for `DELAYED` inserts.

- `got handler lock`

This occurs before the `allocating local table` state and after the `waiting for handler lock` state, when the connection thread gets access to the delayed-insert handler thread.

- `got old table`

This occurs after the `waiting for handler open` state. The delayed-insert handler thread has signaled that it has ended its initialization phase, which includes opening the table for delayed inserts.

- `storing row into queue`

The thread is adding a new row to the list of rows that the delayed-insert handler thread must insert.

- `waiting for delay_list`

This occurs during the initialization phase when the thread is trying to find the delayed-insert handler thread for the table, and before attempting to gain access to the list of delayed-insert threads.

- `waiting for handler insert`

An `INSERT DELAYED` handler has processed all pending inserts and is waiting for new ones.

- `waiting for handler lock`

This occurs before the `allocating local table` state when the connection thread waits for access to the delayed-insert handler thread.

- `waiting for handler open`

This occurs after the `Creating delayed handler` state and before the `got old table` state. The delayed-insert handler thread has just been started, and the connection thread is waiting for it to initialize.

States associated with a delayed-insert handler thread that inserts the rows:

- `insert`

The state that occurs just before inserting rows into the table.

- `reschedule`

After inserting a number of rows, the delayed-insert thread sleeps to let other threads do work.

- `upgrading lock`

A delayed-insert handler is trying to get a lock for the table to insert rows.

- `Waiting for INSERT`

A delayed-insert handler is waiting for a connection thread to add rows to the queue (see `storing row into queue`).

#### 8.12.5.4. Query Cache Thread States

These thread states are associated with the query cache (see [Section 8.9.3, “The MySQL Query Cache”](#)).

- `checking privileges on cached query`

The server is checking whether the user has privileges to access a cached query result.

- `checking query cache for query`

The server is checking whether the current query is present in the query cache.

- `invalidating query cache entries`

Query cache entries are being marked invalid because the underlying tables have changed.

- `sending cached result to client`

The server is taking the result of a query from the query cache and sending it to the client.

- `storing result in query cache`

The server is storing the result of a query in the query cache.

- `Waiting for query cache lock`

This state occurs while a session is waiting to take the query cache lock. This can happen for any statement that needs to perform some query cache operation, such as an `INSERT` or `DELETE` that

invalidates the query cache, a `SELECT` that looks for a cached entry, `RESET QUERY CACHE`, and so forth.

### 8.12.5.5. Replication Master Thread States

The following list shows the most common states you may see in the `State` column for the master's `Binlog Dump` thread. If you see no `Binlog Dump` threads on a master server, this means that replication is not running—that is, that no slaves are currently connected.

- `Sending binlog event to slave`

Binary logs consist of *events*, where an event is usually an update plus some other information. The thread has read an event from the binary log and is now sending it to the slave.

- `Finished reading one binlog; switching to next binlog`

The thread has finished reading a binary log file and is opening the next one to send to the slave.

- `Master has sent all binlog to slave; waiting for binlog to be updated`

The thread has read all outstanding updates from the binary logs and sent them to the slave. The thread is now idle, waiting for new events to appear in the binary log resulting from new updates occurring on the master.

- `Waiting to finalize termination`

A very brief state that occurs as the thread is stopping.

### 8.12.5.6. Replication Slave I/O Thread States

The following list shows the most common states you see in the `State` column for a slave server I/O thread. This state also appears in the `Slave_IO_State` column displayed by `SHOW SLAVE STATUS`, so you can get a good view of what is happening by using that statement.

- `Waiting for master update`

The initial state before `Connecting to master`.

- `Connecting to master`

The thread is attempting to connect to the master.

- `Checking master version`

A state that occurs very briefly, after the connection to the master is established.

- `Registering slave on master`

A state that occurs very briefly after the connection to the master is established.

- `Requesting binlog dump`

A state that occurs very briefly, after the connection to the master is established. The thread sends to the master a request for the contents of its binary logs, starting from the requested binary log file name and position.

- `Waiting to reconnect after a failed binlog dump request`

If the binary log dump request failed (due to disconnection), the thread goes into this state while it sleeps, then tries to reconnect periodically. The interval between retries can be specified using the `CHANGE MASTER TO` statement.

- `Reconnecting after a failed binlog dump request`

The thread is trying to reconnect to the master.

- `Waiting for master to send event`

The thread has connected to the master and is waiting for binary log events to arrive. This can last for a long time if the master is idle. If the wait lasts for `slave_net_timeout` [1993] seconds, a timeout occurs. At that point, the thread considers the connection to be broken and makes an attempt to reconnect.

- `Queueing master event to the relay log`

The thread has read an event and is copying it to the relay log so that the SQL thread can process it.

- `Waiting to reconnect after a failed master event read`

An error occurred while reading (due to disconnection). The thread is sleeping for the number of seconds set by the `CHANGE MASTER TO` statement (default 60) before attempting to reconnect.

- `Reconnecting after a failed master event read`

The thread is trying to reconnect to the master. When connection is established again, the state becomes `Waiting for master to send event`.

- `Waiting for the slave SQL thread to free enough relay log space`

You are using a nonzero `relay_log_space_limit` [534] value, and the relay logs have grown large enough that their combined size exceeds this value. The I/O thread is waiting until the SQL thread frees enough space by processing relay log contents so that it can delete some relay log files.

- `Waiting for slave mutex on exit`

A state that occurs briefly as the thread is stopping.

### 8.12.5.7. Replication Slave SQL Thread States

The following list shows the most common states you may see in the `State` column for a slave server SQL thread:

- `Waiting for the next event in relay log`

The initial state before `Reading event from the relay log`.

- `Reading event from the relay log`

The thread has read an event from the relay log so that the event can be processed.

- `Making temp file`

The thread is executing a `LOAD DATA INFILE` statement and is creating a temporary file containing the data from which the slave will read rows.

- `Slave has read all relay log; waiting for the slave I/O thread to update it`

The thread has processed all events in the relay log files, and is now waiting for the I/O thread to write new events to the relay log.

- `Waiting for slave mutex on exit`

A very brief state that occurs as the thread is stopping.

- `Waiting until MASTER_DELAY seconds after master executed event`

The SQL thread has read an event but is waiting for the slave delay to lapse. This delay is set with the `MASTER_DELAY` option of `CHANGE MASTER TO`.

The `State` column for the I/O thread may also show the text of a statement. This indicates that the thread has read an event from the relay log, extracted the statement from it, and is executing it.

### 8.12.5.8. Replication Slave Connection Thread States

These thread states occur on a replication slave but are associated with connection threads, not with the I/O or SQL threads.

- `Changing master`

The thread is processing a `CHANGE MASTER TO` statement.

- `Killing slave`

The thread is processing a `STOP SLAVE` statement.

- `Opening master dump table`

This state occurs after `Creating table from master dump`.

- `Reading master dump table data`

This state occurs after `Opening master dump table`.

- `Rebuilding the index on master dump table`

This state occurs after `Reading master dump table data`.

### 8.12.5.9. Event Scheduler Thread States

These states occur for the Event Scheduler thread, threads that are created to execute scheduled events, or threads that terminate the scheduler.

- `Clearing`

The scheduler thread or a thread that was executing an event is terminating and is about to end.

- `Initialized`

The scheduler thread or a thread that will execute an event has been initialized.

- `Waiting for next activation`

The scheduler has a nonempty event queue but the next activation is in the future.

- `Waiting for scheduler to stop`

The thread issued `SET GLOBAL event_scheduler=OFF` and is waiting for the scheduler to stop.

- `Waiting on empty queue`

The scheduler's event queue is empty and it is sleeping.

## 8.13. Internal Details of MySQL Optimizations

This background information helps you to understand some of the terms you see in the `EXPLAIN` plan output. If you are rewriting queries to make them more efficient, or writing your own application logic for lookups, joining, or sorting, use this information to determine which optimizations you write yourself and which you can rely on MySQL to perform.

## 8.13.1. Range Optimization

The [range \[899\]](#) access method uses a single index to retrieve a subset of table rows that are contained within one or several index value intervals. It can be used for a single-part or multiple-part index. The following sections give a detailed description of how intervals are extracted from the [WHERE](#) clause.

### 8.13.1.1. The Range Access Method for Single-Part Indexes

For a single-part index, index value intervals can be conveniently represented by corresponding conditions in the [WHERE](#) clause, so we speak of *range conditions* rather than “intervals.”

The definition of a range condition for a single-part index is as follows:

- For both [BTREE](#) and [HASH](#) indexes, comparison of a key part with a constant value is a range condition when using the [= \[1166\]](#), [<=> \[1167\]](#), [IN\(\) \[1169\]](#), [IS NULL \[1168\]](#), or [IS NOT NULL \[1168\]](#) operators.
- Additionally, for [BTREE](#) indexes, comparison of a key part with a constant value is a range condition when using the [> \[1167\]](#), [< \[1167\]](#), [>= \[1167\]](#), [<= \[1167\]](#), [BETWEEN \[1168\]](#), [!= \[1167\]](#), or [<> \[1167\]](#) operators, or [LIKE \[1190\]](#) comparisons if the argument to [LIKE \[1190\]](#) is a constant string that does not start with a wildcard character.
- For all types of indexes, multiple range conditions combined with [OR \[1172\]](#) or [AND \[1171\]](#) form a range condition.

“Constant value” in the preceding descriptions means one of the following:

- A constant from the query string
- A column of a [const \[898\]](#) or [system \[898\]](#) table from the same join
- The result of an uncorrelated subquery
- Any expression composed entirely from subexpressions of the preceding types

Here are some examples of queries with range conditions in the [WHERE](#) clause:

```
SELECT * FROM t1
  WHERE key_col > 1
     AND key_col < 10;

SELECT * FROM t1
  WHERE key_col = 1
     OR key_col IN (15,18,20);

SELECT * FROM t1
  WHERE key_col LIKE 'ab%'
     OR key_col BETWEEN 'bar' AND 'foo';
```

Note that some nonconstant values may be converted to constants during the constant propagation phase.

MySQL tries to extract range conditions from the [WHERE](#) clause for each of the possible indexes. During the extraction process, conditions that cannot be used for constructing the range condition are dropped, conditions that produce overlapping ranges are combined, and conditions that produce empty ranges are removed.

Consider the following statement, where [key1](#) is an indexed column and [nonkey](#) is not indexed:

```
SELECT * FROM t1 WHERE
(key1 < 'abc' AND (key1 LIKE 'abcde%' OR key1 LIKE '%b')) OR
(key1 < 'bar' AND nonkey = 4) OR
```

```
(key1 < 'uux' AND key1 > 'z');
```

The extraction process for key `key1` is as follows:

1. Start with original `WHERE` clause:

```
(key1 < 'abc' AND (key1 LIKE 'abcde%' OR key1 LIKE '%b')) OR
(key1 < 'bar' AND nonkey = 4) OR
(key1 < 'uux' AND key1 > 'z')
```

2. Remove `nonkey = 4` and `key1 LIKE '%b'` because they cannot be used for a range scan. The correct way to remove them is to replace them with `TRUE`, so that we do not miss any matching rows when doing the range scan. Having replaced them with `TRUE`, we get:

```
(key1 < 'abc' AND (key1 LIKE 'abcde%' OR TRUE)) OR
(key1 < 'bar' AND TRUE) OR
(key1 < 'uux' AND key1 > 'z')
```

3. Collapse conditions that are always true or false:

- `(key1 LIKE 'abcde%' OR TRUE)` is always true
- `(key1 < 'uux' AND key1 > 'z')` is always false

Replacing these conditions with constants, we get:

```
(key1 < 'abc' AND TRUE) OR (key1 < 'bar' AND TRUE) OR (FALSE)
```

Removing unnecessary `TRUE` and `FALSE` constants, we obtain:

```
(key1 < 'abc') OR (key1 < 'bar')
```

4. Combining overlapping intervals into one yields the final condition to be used for the range scan:

```
(key1 < 'bar')
```

In general (and as demonstrated by the preceding example), the condition used for a range scan is less restrictive than the `WHERE` clause. MySQL performs an additional check to filter out rows that satisfy the range condition but not the full `WHERE` clause.

The range condition extraction algorithm can handle nested `AND` [1171]/`OR` [1172] constructs of arbitrary depth, and its output does not depend on the order in which conditions appear in `WHERE` clause.

Currently, MySQL does not support merging multiple ranges for the `range` [899] access method for spatial indexes. To work around this limitation, you can use a `UNION` with identical `SELECT` statements, except that you put each spatial predicate in a different `SELECT`.

### 8.13.1.2. The Range Access Method for Multiple-Part Indexes

Range conditions on a multiple-part index are an extension of range conditions for a single-part index. A range condition on a multiple-part index restricts index rows to lie within one or several key tuple intervals. Key tuple intervals are defined over a set of key tuples, using ordering from the index.

For example, consider a multiple-part index defined as `key1(key_part1, key_part2, key_part3)`, and the following set of key tuples listed in key order:

<code>key_part1</code>	<code>key_part2</code>	<code>key_part3</code>
NULL	1	'abc'
NULL	1	'xyz'
NULL	2	'foo'
1	1	'abc'
1	1	'xyz'
1	2	'abc'



```
2      1      'aaa'
```

The condition `key_part1 = 1` defines this interval:

```
(1,-inf,-inf) <= (key_part1,key_part2,key_part3) < (1,+inf,+inf)
```

The interval covers the 4th, 5th, and 6th tuples in the preceding data set and can be used by the range access method.

By contrast, the condition `key_part3 = 'abc'` does not define a single interval and cannot be used by the range access method.

The following descriptions indicate how range conditions work for multiple-part indexes in greater detail.

- For [HASH](#) indexes, each interval containing identical values can be used. This means that the interval can be produced only for conditions in the following form:

```
key_part1 cmp const1
AND key_part2 cmp const2
AND ...
AND key_partN cmp constN;
```

Here, `const1`, `const2`, ... are constants, `cmp` is one of the `=` [1166], `<=>` [1167], or `IS NULL` [1168] comparison operators, and the conditions cover all index parts. (That is, there are `N` conditions, one for each part of an `N`-part index.) For example, the following is a range condition for a three-part [HASH](#) index:

```
key_part1 = 1 AND key_part2 IS NULL AND key_part3 = 'foo'
```

For the definition of what is considered to be a constant, see [Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”](#).

- For a [BTREE](#) index, an interval might be usable for conditions combined with `AND` [1171], where each condition compares a key part with a constant value using `=` [1166], `<=>` [1167], `IS NULL` [1168], `>` [1167], `<` [1167], `>=` [1167], `<=` [1167], `!=` [1167], `<>` [1167], `BETWEEN` [1168], or `LIKE 'pattern'` [1190] (where `'pattern'` does not start with a wildcard). An interval can be used as long as it is possible to determine a single key tuple containing all rows that match the condition (or two intervals if `<>` [1167] or `!=` [1167] is used).

The optimizer attempts to use additional key parts to determine the interval as long as the comparison operator is `=` [1166], `<=>` [1167], or `IS NULL` [1168]. If the operator is `>` [1167], `<` [1167], `>=` [1167], `<=` [1167], `!=` [1167], `<>` [1167], `BETWEEN` [1168], or `LIKE` [1190], the optimizer uses it but considers no more key parts. For the following expression, the optimizer uses `=` [1166] from the first comparison. It also uses `>=` [1167] from the second comparison but considers no further key parts and does not use the third comparison for interval construction:

```
key_part1 = 'foo' AND key_part2 >= 10 AND key_part3 > 10
```

The single interval is:

```
('foo',10,-inf) < (key_part1,key_part2,key_part3) < ('foo',+inf,+inf)
```

It is possible that the created interval contains more rows than the initial condition. For example, the preceding interval includes the value `('foo', 11, 0)`, which does not satisfy the original condition.

- If conditions that cover sets of rows contained within intervals are combined with `OR` [1172], they form a condition that covers a set of rows contained within the union of their intervals. If the conditions are combined with `AND` [1171], they form a condition that covers a set of rows contained within the intersection of their intervals. For example, for this condition on a two-part index:

```
(key_part1 = 1 AND key_part2 < 2) OR (key_part1 > 5)
```

The intervals are:

```
(1,-inf) < (key_part1,key_part2) < (1,2)
(5,-inf) < (key_part1,key_part2)
```

In this example, the interval on the first line uses one key part for the left bound and two key parts for the right bound. The interval on the second line uses only one key part. The `key_len` column in the `EXPLAIN` output indicates the maximum length of the key prefix used.

In some cases, `key_len` may indicate that a key part was used, but that might be not what you would expect. Suppose that `key_part1` and `key_part2` can be `NULL`. Then the `key_len` column displays two key part lengths for the following condition:

```
key_part1 >= 1 AND key_part2 < 2
```

But, in fact, the condition is converted to this:

```
key_part1 >= 1 AND key_part2 IS NOT NULL
```

Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”, describes how optimizations are performed to combine or eliminate intervals for range conditions on a single-part index. Analogous steps are performed for range conditions on multiple-part indexes.

### 8.13.1.3. Equality Range Optimization of Many-Valued Comparisons

Consider these expressions, where `col_name` is an indexed column:

```
col_name IN(val1, ..., valN)
col_name = val1 OR ... OR col_name = valN
```

Each expression is true if `col_name` is equal to any of several values. These comparisons are equality range comparisons (where the “range” is a single value). The optimizer estimates the cost of reading qualifying rows for equality range comparisons as follows:

- If there is a unique index on `col_name`, the row estimate for each range is 1 because at most one row can have the given value.
- Otherwise, the optimizer can estimate the row count for each range using dives into the index or index statistics.

With index dives, the optimizer makes a dive at each end of a range and uses the number of rows in the range as the estimate. For example, the expression `col_name IN (10, 20, 30)` has three equality ranges and the optimizer makes two dives per range to generate a row estimate. Each pair of dives yields an estimate of the number of rows that have the given value.

Index dives provide accurate row estimates, but as the number of comparison values in the expression increases, the optimizer takes longer to generate a row estimate. Use of index statistics is less accurate than index dives but permits faster row estimation for large value lists.

The `eq_range_index_dive_limit` [479] system variable enables you to configure the number of values at which the optimizer switches from one row estimation strategy to the other. To disable use of statistics and always use index dives, set `eq_range_index_dive_limit` [479] to 0. To permit use of index dives for comparisons of up to *N* equality ranges, set `eq_range_index_dive_limit` [479] to *N* + 1.

`eq_range_index_dive_limit` [479] is available as of MySQL 5.6.5. Before 5.6.5, the optimizer uses index dives, which is equivalent to `eq_range_index_dive_limit=0` [479].

To update table index statistics for best estimates, use `ANALYZE TABLE`.

## 8.13.2. Index Merge Optimization

The *Index Merge* method is used to retrieve rows with several [range \[899\]](#) scans and to merge their results into one. The merge can produce unions, intersections, or unions-of-intersections of its underlying scans. This access method merges index scans from a single table; it does not merge scans across multiple tables.

In `EXPLAIN` output, the Index Merge method appears as `index_merge [899]` in the `type` column. In this case, the `key` column contains a list of indexes used, and `key_len` contains a list of the longest key parts for those indexes.

Examples:

```
SELECT * FROM tbl_name WHERE key1 = 10 OR key2 = 20;

SELECT * FROM tbl_name
  WHERE (key1 = 10 OR key2 = 20) AND non_key=30;

SELECT * FROM t1, t2
  WHERE (t1.key1 IN (1,2) OR t1.key2 LIKE 'value%')
     AND t2.key1=t1.some_col;

SELECT * FROM t1, t2
  WHERE t1.key1=1
     AND (t2.key1=t1.some_col OR t2.key2=t1.some_col2);
```

The Index Merge method has several access algorithms (seen in the `Extra` field of `EXPLAIN` output):

- `Using intersect(...)`
- `Using union(...)`
- `Using sort_union(...)`

The following sections describe these methods in greater detail.



#### Note

The Index Merge optimization algorithm has the following known deficiencies:

- If your query has a complex `WHERE` clause with deep `AND [1171]/OR [1172]` nesting and MySQL doesn't choose the optimal plan, try distributing terms using the following identity laws:

```
(x AND y) OR z = (x OR z) AND (y OR z)
(x OR y) AND z = (x AND z) OR (y AND z)
```

- Index Merge is not applicable to full-text indexes. We plan to extend it to cover these in a future MySQL release.
- Before MySQL 5.6.6, if a range scan is possible on some key, the optimizer will not consider using Index Merge Union or Index Merge Sort-Union algorithms. For example, consider this query:

```
SELECT * FROM t1 WHERE (goodkey1 < 10 OR goodkey2 < 20) AND badkey < 30;
```

For this query, two plans are possible:

- An Index Merge scan using the `(goodkey1 < 10 OR goodkey2 < 20)` condition.
- A range scan using the `badkey < 30` condition.

However, the optimizer considers only the second plan.

The choice between different possible variants of the Index Merge access method and other access methods is based on cost estimates of various available options.

### 8.13.2.1. The Index Merge Intersection Access Algorithm

This access algorithm can be employed when a `WHERE` clause was converted to several range conditions on different keys combined with `AND` [1171], and each condition is one of the following:

- In this form, where the index has exactly  $N$  parts (that is, all index parts are covered):

```
key_part1=const1 AND key_part2=const2 ... AND key_partN=constN
```

- Any range condition over a primary key of an `InnoDB` table.

Examples:

```
SELECT * FROM innodb_table WHERE primary_key < 10 AND key_coll=20;
```

```
SELECT * FROM tbl_name
WHERE (key1_part1=1 AND key1_part2=2) AND key2=2;
```

The Index Merge intersection algorithm performs simultaneous scans on all used indexes and produces the intersection of row sequences that it receives from the merged index scans.

If all columns used in the query are covered by the used indexes, full table rows are not retrieved (`EXPLAIN` output contains `Using index` in `Extra` field in this case). Here is an example of such a query:

```
SELECT COUNT(*) FROM t1 WHERE key1=1 AND key2=1;
```

If the used indexes don't cover all columns used in the query, full rows are retrieved only when the range conditions for all used keys are satisfied.

If one of the merged conditions is a condition over a primary key of an `InnoDB` table, it is not used for row retrieval, but is used to filter out rows retrieved using other conditions.

### 8.13.2.2. The Index Merge Union Access Algorithm

The applicability criteria for this algorithm are similar to those for the Index Merge method intersection algorithm. The algorithm can be employed when the table's `WHERE` clause was converted to several range conditions on different keys combined with `OR` [1172], and each condition is one of the following:

- In this form, where the index has exactly  $N$  parts (that is, all index parts are covered):

```
key_part1=const1 AND key_part2=const2 ... AND key_partN=constN
```

- Any range condition over a primary key of an `InnoDB` table.
- A condition for which the Index Merge method intersection algorithm is applicable.

Examples:

```
SELECT * FROM t1 WHERE key1=1 OR key2=2 OR key3=3;
```

```
SELECT * FROM innodb_table WHERE (key1=1 AND key2=2) OR
(key3='foo' AND key4='bar') AND key5=5;
```

### 8.13.2.3. The Index Merge Sort-Union Access Algorithm

This access algorithm is employed when the `WHERE` clause was converted to several range conditions combined by `OR` [1172], but for which the Index Merge method union algorithm is not applicable.

Examples:

```
SELECT * FROM tbl_name WHERE key_coll < 10 OR key_coll2 < 20;
```

```
SELECT * FROM tbl_name
WHERE (key_coll > 10 OR key_coll2 = 20) AND nonkey_coll=30;
```

The difference between the sort-union algorithm and the union algorithm is that the sort-union algorithm must first fetch row IDs for all rows and sort them before returning any rows.

### 8.13.3. Engine Condition Pushdown Optimization

This optimization improves the efficiency of direct comparisons between a nonindexed column and a constant. In such cases, the condition is “pushed down” to the storage engine for evaluation. This optimization can be used only by the [NDB](#) storage engine.



**Note**

The [NDB](#) storage engine is currently not available in MySQL 5.6. If you are interested in using MySQL Cluster, see [MySQL Cluster NDB 7.2](#), which provides information about MySQL Cluster NDB 7.2, which is based on MySQL 5.5 but contains the latest improvements and fixes for [NDB](#).

For MySQL Cluster, this optimization can eliminate the need to send nonmatching rows over the network between the cluster's data nodes and the MySQL Server that issued the query, and can speed up queries where it is used by a factor of 5 to 10 times over cases where condition pushdown could be but is not used.

Suppose that a MySQL Cluster table is defined as follows:

```
CREATE TABLE t1 (
  a INT,
  b INT,
  KEY(a)
) ENGINE=NDB;
```

Condition pushdown can be used with queries such as the one shown here, which includes a comparison between a nonindexed column and a constant:

```
SELECT a, b FROM t1 WHERE b = 10;
```

The use of condition pushdown can be seen in the output of [EXPLAIN](#):

```
mysql> EXPLAIN SELECT a,b FROM t1 WHERE b = 10\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: t1
         type: ALL
possible_keys: NULL
          key: NULL
      key_len: NULL
         ref: NULL
        rows: 10
     Extra: Using where with pushed condition
```

However, condition pushdown *cannot* be used with either of these two queries:

```
SELECT a,b FROM t1 WHERE a = 10;
SELECT a,b FROM t1 WHERE b + 1 = 10;
```

Condition pushdown is not applicable to the first query because an index exists on column [a](#). (An index access method would be more efficient and so would be chosen in preference to condition pushdown.) Condition pushdown cannot be employed for the second query because the comparison involving the nonindexed column [b](#) is indirect. (However, condition pushdown could be applied if you were to reduce  $b + 1 = 10$  to  $b = 9$  in the [WHERE](#) clause.)

Condition pushdown may also be employed when an indexed column is compared with a constant using a  $>$  or  $<$  operator:

```
mysql> EXPLAIN SELECT a, b FROM t1 WHERE a < 2\G
***** 1. row *****
      id: 1
```

```

select_type: SIMPLE
  table: t1
  type: range
possible_keys: a
  key: a
  key_len: 5
  ref: NULL
  rows: 2
Extra: Using where with pushed condition

```

Other supported comparisons for condition pushdown include the following:

- `column [NOT] LIKE pattern`

`pattern` must be a string literal containing the pattern to be matched; for syntax, see [Section 12.5.1, “String Comparison Functions”](#).

- `column IS [NOT] NULL`
- `column IN (value_list)`

Each item in the `value_list` must be a constant, literal value.

- `column BETWEEN constant1 AND constant2`

`constant1` and `constant2` must each be a constant, literal value.

In all of the cases in the preceding list, it is possible for the condition to be converted into the form of one or more direct comparisons between a column and a constant.

Engine condition pushdown is enabled by default. To disable it at server startup, set the `optimizer_switch [520]` system variable. For example, in a `my.cnf` file, use these lines:

```

[mysqld]
optimizer_switch=engine_condition_pushdown=off

```

At runtime, enable condition pushdown like this:

```

SET optimizer_switch='engine_condition_pushdown=off';

```

**Limitations.** Engine condition pushdown is subject to the following limitations:

- Condition pushdown is supported only by the [NDB](#) storage engine.
- Columns may be compared with constants only; however, this includes expressions which evaluate to constant values.
- Columns used in comparisons cannot be of any of the [BLOB](#) or [TEXT](#) types.
- A string value to be compared with a column must use the same collation as the column.
- Joins are not directly supported; conditions involving multiple tables are pushed separately where possible. Use [EXPLAIN EXTENDED](#) to determine which conditions are actually pushed down.

## 8.13.4. Index Condition Pushdown Optimization

Index Condition Pushdown (ICP) is an optimization for the case where MySQL retrieves rows from a table using an index. Without ICP, the storage engine traverses the index to locate rows in the base table and returns them to the MySQL server which evaluates the [WHERE](#) condition for the rows. With ICP enabled, and if parts of the [WHERE](#) condition can be evaluated by using only fields from the index, the MySQL server pushes this part of the [WHERE](#) condition down to the storage engine. The storage engine then evaluates the pushed index condition by using the index entry and only if this is satisfied is the row read from the table. ICP can reduce the number of times the storage engine must access the base table and the number of times the MySQL server must access the storage engine.

Index Condition Pushdown optimization is used for the [range \[899\]](#), [ref \[898\]](#), [eq\\_ref \[898\]](#), and [ref\\_or\\_null \[899\]](#) access methods when there is a need to access full table rows. This strategy can be used for [InnoDB](#) and [MyISAM](#) tables.

To see how this optimization works, consider first how an index scan proceeds when Index Condition Pushdown is not used:

1. Get the next row, first by reading the index tuple, and then by using the index tuple to locate and read the full table row.
2. Test the part of the `WHERE` condition that applies to this table. Accept or reject the row based on the test result.

When Index Condition Pushdown is used, the scan proceeds like this instead:

1. Get the next row's index tuple (but not the full table row).
2. Test the part of the `WHERE` condition that applies to this table and can be checked using only index columns. If the condition is not satisfied, proceed to the index tuple for the next row.
3. If the condition is satisfied, use the index tuple to locate and read the full table row.
4. Test the remaining part of the `WHERE` condition that applies to this table. Accept or reject the row based on the test result.

When Index Condition Pushdown is used, the `Extra` column in `EXPLAIN` output shows `Using index condition`. It will not show `Index only` because that does not apply when full table rows must be read.

Suppose that we have a table containing information about people and their addresses and that the table has an index defined as `INDEX (zipcode, lastname, firstname)`. If we know a person's `zipcode` value but are not sure about the last name, we can search like this:

```
SELECT * FROM people
WHERE zipcode='95054'
AND lastname LIKE '%etrunia%'
AND address LIKE '%Main Street%';
```

MySQL can use the index to scan through people with `zipcode='95054'`. The second part (`lastname LIKE '%etrunia%'`) cannot be used to limit the number of rows that must be scanned, so without Index Condition Pushdown, this query must retrieve full table rows for all the people who have `zipcode='95054'`.

With Index Condition Pushdown, MySQL will check the `lastname LIKE '%etrunia%'` part before reading the full table row. This avoids reading full rows corresponding to all index tuples that do not match the `lastname` condition.

Index Condition Pushdown is enabled by default; it can be controlled with the [optimizer\\_switch \[520\]](#) system variable by setting the `index_condition_pushdown` flag. See [Section 8.8.5.2, "Controlling Switchable Optimizations"](#).

## 8.13.5. Use of Index Extensions

[InnoDB](#) automatically extends each secondary index by appending the primary key columns to it. Consider this table definition:

```
CREATE TABLE t1 (
  i1 INT NOT NULL DEFAULT 0,
  i2 INT NOT NULL DEFAULT 0,
  d DATE DEFAULT NULL,
  PRIMARY KEY (i1, i2),
  INDEX k_d (d)
```



```
) ENGINE = InnoDB;
```

This table defines the primary key on columns (*i1*, *i2*). It also defines a secondary index *k\_d* on column (*d*), but internally *InnoDB* extends this index and treats it as columns (*d*, *i1*, *i2*).

Before MySQL 5.6.9, the optimizer does not take into account the primary key columns of the extended secondary index when determining how and whether to use that index. As of 5.6.9, the optimizer takes the primary key columns into account, which can result in more efficient query execution plans and better performance.

The optimizer can use extended secondary indexes for *ref*, *range*, and *index\_merge* index access, for loose index scans, for join and sorting optimization, and for *MIN()* [1284]/*MAX()* [1284] optimization.

The following example shows how execution plans are affected by whether the optimizer uses extended secondary indexes. Suppose that *t1* is populated with these rows:

```
INSERT INTO t1 VALUES
(1, 1, '1998-01-01'), (1, 2, '1999-01-01'),
(1, 3, '2000-01-01'), (1, 4, '2001-01-01'),
(1, 5, '2002-01-01'), (2, 1, '1998-01-01'),
(2, 2, '1999-01-01'), (2, 3, '2000-01-01'),
(2, 4, '2001-01-01'), (2, 5, '2002-01-01'),
(3, 1, '1998-01-01'), (3, 2, '1999-01-01'),
(3, 3, '2000-01-01'), (3, 4, '2001-01-01'),
(3, 5, '2002-01-01'), (4, 1, '1998-01-01'),
(4, 2, '1999-01-01'), (4, 3, '2000-01-01'),
(4, 4, '2001-01-01'), (4, 5, '2002-01-01'),
(5, 1, '1998-01-01'), (5, 2, '1999-01-01'),
(5, 3, '2000-01-01'), (5, 4, '2001-01-01'),
(5, 5, '2002-01-01');
```

Now consider this query:

```
EXPLAIN SELECT COUNT(*) FROM t1 WHERE i1 = 3 AND d = '2000-01-01'
```

The optimizer cannot use the primary key in this case because that comprises columns (*i1*, *i2*) and the query does not refer to *i2*. Instead, the optimizer can use the secondary index *k\_d* on (*d*), and the execution plan depends on whether the extended index is used.

When the optimizer does not consider index extensions, it treats the index *k\_d* as only (*d*). *EXPLAIN* for the query produces this result:

```
mysql> EXPLAIN SELECT COUNT(*) FROM t1 WHERE i1 = 3 AND d = '2000-01-01'\G
***** 1. row *****
      id: 1
    select_type: SIMPLE
      table: t1
       type: ref
possible_keys: PRIMARY,k_d
         key: k_d
      key_len: 4
         ref: const
        rows: 5
      Extra: Using where; Using index
```

When the optimizer takes index extensions into account, it treats *k\_d* as (*d*, *i1*, *i2*). In this case, it can use the leftmost index prefix (*d*, *i1*) to produce a better execution plan:

```
mysql> EXPLAIN SELECT COUNT(*) FROM t1 WHERE i1 = 3 AND d = '2000-01-01'\G
***** 1. row *****
      id: 1
    select_type: SIMPLE
      table: t1
       type: ref
possible_keys: PRIMARY,k_d
         key: k_d
```

```

key_len: 8
  ref: const,const
  rows: 1
  Extra: Using index

```

In both cases, `key` indicates that the optimizer will use secondary index `k_d` but the `EXPLAIN` output shows these improvements from using the extended index:

- `key_len` goes from 4 bytes to 8 bytes, indicating that key lookups use columns `d` and `i1`, not just `d`.
- The `ref` value changes from `const` to `const,const` because the key lookup uses two key parts, not one.
- The `rows` count decreases from 5 to 1, indicating that `InnoDB` should need to examine fewer rows to produce the result.
- The `Extra` value changes from `Using where; Using index` to `Using index`. This means that rows can be read using only the index, without consulting columns in the data row.

Differences in optimizer behavior for use of extended indexes can also be seen with `SHOW STATUS`:

```

FLUSH TABLE t1;
FLUSH STATUS;
SELECT COUNT(*) FROM t1 WHERE i1 = 3 AND d = '2000-01-01';
SHOW STATUS LIKE 'handler_read%'

```

The preceding statements include `FLUSH TABLE` and `FLUSH STATUS` to flush the table cache and clear the status counters.

Without index extensions, `SHOW STATUS` produces this result:

```

+-----+-----+
| Variable_name | Value |
+-----+-----+
| Handler_read_first | 0 |
| Handler_read_key | 1 |
| Handler_read_last | 0 |
| Handler_read_next | 5 |
| Handler_read_prev | 0 |
| Handler_read_rnd | 0 |
| Handler_read_rnd_next | 0 |
+-----+-----+

```

With index extensions, `SHOW STATUS` produces this result. The `Handler_read_next` [589] value decreases from 5 to 1, indicating more efficient use of the index:

```

+-----+-----+
| Variable_name | Value |
+-----+-----+
| Handler_read_first | 0 |
| Handler_read_key | 1 |
| Handler_read_last | 0 |
| Handler_read_next | 1 |
| Handler_read_prev | 0 |
| Handler_read_rnd | 0 |
| Handler_read_rnd_next | 0 |
+-----+-----+

```

The `use_index_extensions` flag of the `optimizer_switch` [520] system variable permits control over whether the optimizer takes the primary key columns into account when determining how to use an `InnoDB` table's secondary indexes. By default, `use_index_extensions` is enabled. To check whether disabling use of index extensions will improve performance, use this statement:

```

SET optimizer_switch = 'use_index_extensions=off';

```

Use of index extensions by the optimizer is subject to the usual limits on the number of key parts in an index (16) and the maximum key length (3072 bytes).

### 8.13.6. IS NULL Optimization

MySQL can perform the same optimization on `col_name IS NULL` [1168] that it can use for `col_name = constant_value`. For example, MySQL can use indexes and ranges to search for `NULL` with `IS NULL` [1168].

Examples:

```
SELECT * FROM tbl_name WHERE key_col IS NULL;

SELECT * FROM tbl_name WHERE key_col <=> NULL;

SELECT * FROM tbl_name
WHERE key_col=const1 OR key_col=const2 OR key_col IS NULL;
```

If a `WHERE` clause includes a `col_name IS NULL` [1168] condition for a column that is declared as `NOT NULL`, that expression is optimized away. This optimization does not occur in cases when the column might produce `NULL` anyway; for example, if it comes from a table on the right side of a `LEFT JOIN`.

MySQL can also optimize the combination `col_name = expr OR col_name IS NULL`, a form that is common in resolved subqueries. `EXPLAIN` shows `ref_or_null` [899] when this optimization is used.

This optimization can handle one `IS NULL` [1168] for any key part.

Some examples of queries that are optimized, assuming that there is an index on columns `a` and `b` of table `t2`:

```
SELECT * FROM t1 WHERE t1.a=expr OR t1.a IS NULL;

SELECT * FROM t1, t2 WHERE t1.a=t2.a OR t2.a IS NULL;

SELECT * FROM t1, t2
WHERE (t1.a=t2.a OR t2.a IS NULL) AND t2.b=t1.b;

SELECT * FROM t1, t2
WHERE t1.a=t2.a AND (t2.b=t1.b OR t2.b IS NULL);

SELECT * FROM t1, t2
WHERE (t1.a=t2.a AND t2.a IS NULL AND ...)
OR (t1.a=t2.a AND t2.a IS NULL AND ...);
```

`ref_or_null` [899] works by first doing a read on the reference key, and then a separate search for rows with a `NULL` key value.

Note that the optimization can handle only one `IS NULL` [1168] level. In the following query, MySQL uses key lookups only on the expression `(t1.a=t2.a AND t2.a IS NULL)` and is not able to use the key part on `b`:

```
SELECT * FROM t1, t2
WHERE (t1.a=t2.a AND t2.a IS NULL)
OR (t1.b=t2.b AND t2.b IS NULL);
```

### 8.13.7. LEFT JOIN and RIGHT JOIN Optimization

MySQL implements an `A LEFT JOIN B join_condition` as follows:

- Table `B` is set to depend on table `A` and all tables on which `A` depends.
- Table `A` is set to depend on all tables (except `B`) that are used in the `LEFT JOIN` condition.
- The `LEFT JOIN` condition is used to decide how to retrieve rows from table `B`. (In other words, any condition in the `WHERE` clause is not used.)

- All standard join optimizations are performed, with the exception that a table is always read after all tables on which it depends. If there is a circular dependence, MySQL issues an error.
- All standard `WHERE` optimizations are performed.
- If there is a row in `A` that matches the `WHERE` clause, but there is no row in `B` that matches the `ON` condition, an extra `B` row is generated with all columns set to `NULL`.
- If you use `LEFT JOIN` to find rows that do not exist in some table and you have the following test: `col_name IS NULL` in the `WHERE` part, where `col_name` is a column that is declared as `NOT NULL`, MySQL stops searching for more rows (for a particular key combination) after it has found one row that matches the `LEFT JOIN` condition.

The implementation of `RIGHT JOIN` is analogous to that of `LEFT JOIN` with the roles of the tables reversed.

The join optimizer calculates the order in which tables should be joined. The table read order forced by `LEFT JOIN` or `STRAIGHT_JOIN` helps the join optimizer do its work much more quickly, because there are fewer table permutations to check. Note that this means that if you do a query of the following type, MySQL does a full scan on `b` because the `LEFT JOIN` forces it to be read before `d`:

```
SELECT *
FROM a JOIN b LEFT JOIN c ON (c.key=a.key)
LEFT JOIN d ON (d.key=a.key)
WHERE b.key=d.key;
```

The fix in this case is reverse the order in which `a` and `b` are listed in the `FROM` clause:

```
SELECT *
FROM b JOIN a LEFT JOIN c ON (c.key=a.key)
LEFT JOIN d ON (d.key=a.key)
WHERE b.key=d.key;
```

For a `LEFT JOIN`, if the `WHERE` condition is always false for the generated `NULL` row, the `LEFT JOIN` is changed to a normal join. For example, the `WHERE` clause would be false in the following query if `t2.column1` were `NULL`:

```
SELECT * FROM t1 LEFT JOIN t2 ON (column1) WHERE t2.column2=5;
```

Therefore, it is safe to convert the query to a normal join:

```
SELECT * FROM t1, t2 WHERE t2.column2=5 AND t1.column1=t2.column1;
```

This can be made faster because MySQL can use table `t2` before table `t1` if doing so would result in a better query plan. To provide a hint about the table join order, use `STRAIGHT_JOIN`. (See [Section 13.2.9, "SELECT Syntax"](#).)

### 8.13.8. Nested-Loop Join Algorithms

MySQL executes joins between tables using a nested-loop algorithm or variations on it.

#### Nested-Loop Join Algorithm

A simple nested-loop join (NLJ) algorithm reads rows from the first table in a loop one at a time, passing each row to a nested loop that processes the next table in the join. This process is repeated as many times as there remain tables to be joined.

Assume that a join between three tables `t1`, `t2`, and `t3` is to be executed using the following join types:

Table	Join Type
t1	range
t2	ref

t3      ALL

If a simple NLJ algorithm is used, the join is processed like this:

```
for each row in t1 matching range {
  for each row in t2 matching reference key {
    for each row in t3 {
      if row satisfies join conditions,
        send to client
    }
  }
}
```

Because the NLJ algorithm passes rows one at a time from outer loops to inner loops, it typically reads tables processed in the inner loops many times.

### Block Nested-Loop Join Algorithm

A Block Nested-Loop (BNL) join algorithm uses buffering of rows read in outer loops to reduce the number of times that tables in inner loops must be read. For example, if 10 rows are read into a buffer and the buffer is passed to the next inner loop, each row read in the inner loop can be compared against all 10 rows in the buffer. This reduces the number of times the inner table must be read by an order of magnitude.

MySQL uses join buffering under these conditions:

- The `join_buffer_size` [490] system variable determines the size of each join buffer.
- Join buffering can be used when the join is of type `ALL` [900] or `index` [899] (in other words, when no possible keys can be used, and a full scan is done, of either the data or index rows, respectively), or `range` [899]. In MySQL 5.6, use of buffering is extended to be applicable to outer joins, as described in [Section 8.13.12, “Block Nested-Loop and Batched Key Access Joins”](#).
- One buffer is allocated for each join that can be buffered, so a given query might be processed using multiple join buffers.
- A join buffer is never allocated for the first nonconst table, even if it would be of type `ALL` [900] or `index` [899].
- A join buffer is allocated prior to executing the join and freed after the query is done.
- Only columns of interest to the join are stored in the join buffer, not whole rows.

For the example join described previously for the NLJ algorithm (without buffering), the join is done as follows using join buffering:

```
for each row in t1 matching range {
  for each row in t2 matching reference key {
    store used columns from t1, t2 in join buffer
    if buffer is full {
      for each row in t3 {
        for each t1, t2 combination in join buffer {
          if row satisfies join conditions,
            send to client
        }
      }
      empty buffer
    }
  }
}

if buffer is not empty {
  for each row in t3 {
    for each t1, t2 combination in join buffer {
      if row satisfies join conditions,
        send to client
    }
  }
}
```

```
}
}
```

If  $S$  is the size of each stored  $t_1$ ,  $t_2$  combination is the join buffer and  $C$  is the number of combinations in the buffer, the number of times table  $t_3$  is scanned is:

```
(S * C) / join_buffer_size + 1
```

The number of  $t_3$  scans decreases as the value of `join_buffer_size` [490] increases, up to the point when `join_buffer_size` [490] is large enough to hold all previous row combinations. At that point, there is no speed to be gained by making it larger.

### 8.13.9. Nested Join Optimization

The syntax for expressing joins permits nested joins. The following discussion refers to the join syntax described in [Section 13.2.9.2, “JOIN Syntax”](#).

The syntax of *table\_factor* is extended in comparison with the SQL Standard. The latter accepts only *table\_reference*, not a list of them inside a pair of parentheses. This is a conservative extension if we consider each comma in a list of *table\_reference* items as equivalent to an inner join. For example:

```
SELECT * FROM t1 LEFT JOIN (t2, t3, t4)
      ON (t2.a=t1.a AND t3.b=t1.b AND t4.c=t1.c)
```

is equivalent to:

```
SELECT * FROM t1 LEFT JOIN (t2 CROSS JOIN t3 CROSS JOIN t4)
      ON (t2.a=t1.a AND t3.b=t1.b AND t4.c=t1.c)
```

In MySQL, `CROSS JOIN` is a syntactic equivalent to `INNER JOIN` (they can replace each other). In standard SQL, they are not equivalent. `INNER JOIN` is used with an `ON` clause; `CROSS JOIN` is used otherwise.

In general, parentheses can be ignored in join expressions containing only inner join operations. After removing parentheses and grouping operations to the left, the join expression:

```
t1 LEFT JOIN (t2 LEFT JOIN t3 ON t2.b=t3.b OR t2.b IS NULL)
      ON t1.a=t2.a
```

transforms into the expression:

```
(t1 LEFT JOIN t2 ON t1.a=t2.a) LEFT JOIN t3
      ON t2.b=t3.b OR t2.b IS NULL
```

Yet, the two expressions are not equivalent. To see this, suppose that the tables  $t_1$ ,  $t_2$ , and  $t_3$  have the following state:

- Table  $t_1$  contains rows (1), (2)
- Table  $t_2$  contains row (1,101)
- Table  $t_3$  contains row (101)

In this case, the first expression returns a result set including the rows (1,1,101,101), (2,NULL,NULL,NULL), whereas the second expression returns the rows (1,1,101,101), (2,NULL,NULL,101):

```
mysql> SELECT *
-> FROM t1
-> LEFT JOIN
-> (t2 LEFT JOIN t3 ON t2.b=t3.b OR t2.b IS NULL)
-> ON t1.a=t2.a;
+-----+-----+-----+-----+
| a | a | b | b |
+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+
| 1 | 1 | 101 | 101 |
| 2 | NULL | NULL | NULL |
+-----+-----+-----+-----+

mysql> SELECT *
-> FROM (t1 LEFT JOIN t2 ON t1.a=t2.a)
-> LEFT JOIN t3
-> ON t2.b=t3.b OR t2.b IS NULL;
+-----+-----+-----+-----+
| a | a | b | b |
+-----+-----+-----+-----+
| 1 | 1 | 101 | 101 |
| 2 | NULL | NULL | 101 |
+-----+-----+-----+-----+

```

In the following example, an outer join operation is used together with an inner join operation:

```
t1 LEFT JOIN (t2, t3) ON t1.a=t2.a
```

That expression cannot be transformed into the following expression:

```
t1 LEFT JOIN t2 ON t1.a=t2.a, t3.
```

For the given table states, the two expressions return different sets of rows:

```

mysql> SELECT *
-> FROM t1 LEFT JOIN (t2, t3) ON t1.a=t2.a;
+-----+-----+-----+-----+
| a | a | b | b |
+-----+-----+-----+-----+
| 1 | 1 | 101 | 101 |
| 2 | NULL | NULL | NULL |
+-----+-----+-----+-----+

mysql> SELECT *
-> FROM t1 LEFT JOIN t2 ON t1.a=t2.a, t3;
+-----+-----+-----+-----+
| a | a | b | b |
+-----+-----+-----+-----+
| 1 | 1 | 101 | 101 |
| 2 | NULL | NULL | 101 |
+-----+-----+-----+-----+

```

Therefore, if we omit parentheses in a join expression with outer join operators, we might change the result set for the original expression.

More exactly, we cannot ignore parentheses in the right operand of the left outer join operation and in the left operand of a right join operation. In other words, we cannot ignore parentheses for the inner table expressions of outer join operations. Parentheses for the other operand (operand for the outer table) can be ignored.

The following expression:

```
(t1,t2) LEFT JOIN t3 ON P(t2.b,t3.b)
```

is equivalent to this expression:

```
t1, t2 LEFT JOIN t3 ON P(t2.b,t3.b)
```

for any tables `t1, t2, t3` and any condition `P` over attributes `t2.b` and `t3.b`.

Whenever the order of execution of the join operations in a join expression (*join\_table*) is not from left to right, we talk about nested joins. Consider the following queries:

```

SELECT * FROM t1 LEFT JOIN (t2 LEFT JOIN t3 ON t2.b=t3.b) ON t1.a=t2.a
WHERE t1.a > 1

SELECT * FROM t1 LEFT JOIN (t2, t3) ON t1.a=t2.a

```



```
WHERE (t2.b=t3.b OR t2.b IS NULL) AND t1.a > 1
```

Those queries are considered to contain these nested joins:

```
t2 LEFT JOIN t3 ON t2.b=t3.b
t2, t3
```

The nested join is formed in the first query with a left join operation, whereas in the second query it is formed with an inner join operation.

In the first query, the parentheses can be omitted: The grammatical structure of the join expression will dictate the same order of execution for join operations. For the second query, the parentheses cannot be omitted, although the join expression here can be interpreted unambiguously without them. (In our extended syntax the parentheses in `(t2, t3)` of the second query are required, although theoretically the query could be parsed without them: We still would have unambiguous syntactical structure for the query because `LEFT JOIN` and `ON` would play the role of the left and right delimiters for the expression `(t2,t3)`.)

The preceding examples demonstrate these points:

- For join expressions involving only inner joins (and not outer joins), parentheses can be removed. You can remove parentheses and evaluate left to right (or, in fact, you can evaluate the tables in any order).
- The same is not true, in general, for outer joins or for outer joins mixed with inner joins. Removal of parentheses may change the result.

Queries with nested outer joins are executed in the same pipeline manner as queries with inner joins. More exactly, a variation of the nested-loop join algorithm is exploited. Recall by what algorithmic schema the nested-loop join executes a query. Suppose that we have a join query over 3 tables `T1, T2, T3` of the form:

```
SELECT * FROM T1 INNER JOIN T2 ON P1(T1,T2)
           INNER JOIN T3 ON P2(T2,T3)
WHERE P(T1,T2,T3).
```

Here, `P1(T1,T2)` and `P2(T2,T3)` are some join conditions (on expressions), whereas `P(t1,t2,t3)` is a condition over columns of tables `T1, T2, T3`.

The nested-loop join algorithm would execute this query in the following manner:

```
FOR each row t1 in T1 {
  FOR each row t2 in T2 such that P1(t1,t2) {
    FOR each row t3 in T3 such that P2(t2,t3) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
    }
  }
}
```

The notation `t1 || t2 || t3` means “a row constructed by concatenating the columns of rows `t1`, `t2`, and `t3`.” In some of the following examples, `NULL` where a row name appears means that `NULL` is used for each column of that row. For example, `t1 || t2 || NULL` means “a row constructed by concatenating the columns of rows `t1` and `t2`, and `NULL` for each column of `t3`.”

Now let's consider a query with nested outer joins:

```
SELECT * FROM T1 LEFT JOIN
           (T2 LEFT JOIN T3 ON P2(T2,T3))
           ON P1(T1,T2)
WHERE P(T1,T2,T3).
```

For this query, we modify the nested-loop pattern to get:

```

FOR each row t1 in T1 {
  BOOL f1:=FALSE;
  FOR each row t2 in T2 such that P1(t1,t2) {
    BOOL f2:=FALSE;
    FOR each row t3 in T3 such that P2(t2,t3) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
      f2=TRUE;
      f1=TRUE;
    }
    IF (!f2) {
      IF P(t1,t2,NULL) {
        t:=t1||t2||NULL; OUTPUT t;
      }
      f1=TRUE;
    }
  }
  IF (!f1) {
    IF P(t1,NULL,NULL) {
      t:=t1||NULL||NULL; OUTPUT t;
    }
  }
}

```

In general, for any nested loop for the first inner table in an outer join operation, a flag is introduced that is turned off before the loop and is checked after the loop. The flag is turned on when for the current row from the outer table a match from the table representing the inner operand is found. If at the end of the loop cycle the flag is still off, no match has been found for the current row of the outer table. In this case, the row is complemented by **NULL** values for the columns of the inner tables. The result row is passed to the final check for the output or into the next nested loop, but only if the row satisfies the join condition of all embedded outer joins.

In our example, the outer join table expressed by the following expression is embedded:

```
(T2 LEFT JOIN T3 ON P2(T2,T3))
```

Note that for the query with inner joins, the optimizer could choose a different order of nested loops, such as this one:

```

FOR each row t3 in T3 {
  FOR each row t2 in T2 such that P2(t2,t3) {
    FOR each row t1 in T1 such that P1(t1,t2) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
    }
  }
}

```

For the queries with outer joins, the optimizer can choose only such an order where loops for outer tables precede loops for inner tables. Thus, for our query with outer joins, only one nesting order is possible. For the following query, the optimizer will evaluate two different nestings:

```

SELECT * T1 LEFT JOIN (T2,T3) ON P1(T1,T2) AND P2(T1,T3)
WHERE P(T1,T2,T3)

```

The nestings are these:

```

FOR each row t1 in T1 {
  BOOL f1:=FALSE;
  FOR each row t2 in T2 such that P1(t1,t2) {
    FOR each row t3 in T3 such that P2(t1,t3) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
      f1:=TRUE
    }
  }
}

```

```

IF (!f1) {
  IF P(t1,NULL,NULL) {
    t:=t1||NULL||NULL; OUTPUT t;
  }
}
}

```

and:

```

FOR each row t1 in T1 {
  BOOL f1:=FALSE;
  FOR each row t3 in T3 such that P2(t1,t3) {
    FOR each row t2 in T2 such that P1(t1,t2) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
      f1:=TRUE
    }
  }
  IF (!f1) {
    IF P(t1,NULL,NULL) {
      t:=t1||NULL||NULL; OUTPUT t;
    }
  }
}

```

In both nestings, **T1** must be processed in the outer loop because it is used in an outer join. **T2** and **T3** are used in an inner join, so that join must be processed in the inner loop. However, because the join is an inner join, **T2** and **T3** can be processed in either order.

When discussing the nested-loop algorithm for inner joins, we omitted some details whose impact on the performance of query execution may be huge. We did not mention so-called “pushed-down” conditions. Suppose that our **WHERE** condition  $P(T1, T2, T3)$  can be represented by a conjunctive formula:

```
P(T1,T2,T2) = C1(T1) AND C2(T2) AND C3(T3).
```

In this case, MySQL actually uses the following nested-loop schema for the execution of the query with inner joins:

```

FOR each row t1 in T1 such that C1(t1) {
  FOR each row t2 in T2 such that P1(t1,t2) AND C2(t2) {
    FOR each row t3 in T3 such that P2(t2,t3) AND C3(t3) {
      IF P(t1,t2,t3) {
        t:=t1||t2||t3; OUTPUT t;
      }
    }
  }
}

```

You see that each of the conjuncts  $C1(T1)$ ,  $C2(T2)$ ,  $C3(T3)$  are pushed out of the most inner loop to the most outer loop where it can be evaluated. If  $C1(T1)$  is a very restrictive condition, this condition pushdown may greatly reduce the number of rows from table **T1** passed to the inner loops. As a result, the execution time for the query may improve immensely.

For a query with outer joins, the **WHERE** condition is to be checked only after it has been found that the current row from the outer table has a match in the inner tables. Thus, the optimization of pushing conditions out of the inner nested loops cannot be applied directly to queries with outer joins. Here we have to introduce conditional pushed-down predicates guarded by the flags that are turned on when a match has been encountered.

For our example with outer joins with:

```
P(T1,T2,T3)=C1(T1) AND C(T2) AND C3(T3)
```

the nested-loop schema using guarded pushed-down conditions looks like this:

```

FOR each row t1 in T1 such that C1(t1) {
  BOOL f1:=FALSE;
  FOR each row t2 in T2
    such that P1(t1,t2) AND (f1?C2(t2):TRUE) {
    BOOL f2:=FALSE;
    FOR each row t3 in T3
      such that P2(t2,t3) AND (f1&&f2?C3(t3):TRUE) {
        IF (f1&&f2?TRUE:(C2(t2) AND C3(t3))) {
          t:=t1||t2||t3; OUTPUT t;
        }
        f2=TRUE;
        f1=TRUE;
      }
    }
  IF (!f2) {
    IF (f1?TRUE:C2(t2) && P(t1,t2,NULL)) {
      t:=t1||t2||NULL; OUTPUT t;
    }
    f1=TRUE;
  }
}
IF (!f1 && P(t1,NULL,NULL)) {
  t:=t1||NULL||NULL; OUTPUT t;
}
}

```

In general, pushed-down predicates can be extracted from join conditions such as `P1(T1,T2)` and `P(T2,T3)`. In this case, a pushed-down predicate is guarded also by a flag that prevents checking the predicate for the `NULL`-complemented row generated by the corresponding outer join operation.

Note that access by key from one inner table to another in the same nested join is prohibited if it is induced by a predicate from the `WHERE` condition. (We could use conditional key access in this case, but this technique is not employed yet in MySQL 5.6.)

### 8.13.10. Outer Join Simplification

Table expressions in the `FROM` clause of a query are simplified in many cases.

At the parser stage, queries with right outer joins operations are converted to equivalent queries containing only left join operations. In the general case, the conversion is performed according to the following rule:

```

(T1, ...) RIGHT JOIN (T2,...) ON P(T1,...,T2,...) =
(T2, ...) LEFT JOIN (T1,...) ON P(T1,...,T2,...)

```

All inner join expressions of the form `T1 INNER JOIN T2 ON P(T1,T2)` are replaced by the list `T1,T2,P(T1,T2)` being joined as a conjunct to the `WHERE` condition (or to the join condition of the embedding join, if there is any).

When the optimizer evaluates plans for join queries with outer join operation, it takes into consideration only the plans where, for each such operation, the outer tables are accessed before the inner tables. The optimizer options are limited because only such plans enables us to execute queries with outer joins operations by the nested loop schema.

Suppose that we have a query of the form:

```

SELECT * T1 LEFT JOIN T2 ON P1(T1,T2)
WHERE P(T1,T2) AND R(T2)

```

with `R(T2)` narrowing greatly the number of matching rows from table `T2`. If we executed the query as it is, the optimizer would have no other choice besides to access table `T1` before table `T2` that may lead to a very inefficient execution plan.

Fortunately, MySQL converts such a query into a query without an outer join operation if the `WHERE` condition is null-rejected. A condition is called null-rejected for an outer join operation if it evaluates to `FALSE` or to `UNKNOWN` for any `NULL`-complemented row built for the operation.

Thus, for this outer join:

```
T1 LEFT JOIN T2 ON T1.A=T2.A
```

Conditions such as these are null-rejected:

```
T2.B IS NOT NULL,
T2.B > 3,
T2.C <= T1.C,
T2.B < 2 OR T2.C > 1
```

Conditions such as these are not null-rejected:

```
T2.B IS NULL,
T1.B < 3 OR T2.B IS NOT NULL,
T1.B < 3 OR T2.B > 3
```

The general rules for checking whether a condition is null-rejected for an outer join operation are simple. A condition is null-rejected in the following cases:

- If it is of the form `A IS NOT NULL`, where `A` is an attribute of any of the inner tables
- If it is a predicate containing a reference to an inner table that evaluates to `UNKNOWN` when one of its arguments is `NULL`
- If it is a conjunction containing a null-rejected condition as a conjunct
- If it is a disjunction of null-rejected conditions

A condition can be null-rejected for one outer join operation in a query and not null-rejected for another. In the query:

```
SELECT * FROM T1 LEFT JOIN T2 ON T2.A=T1.A
           LEFT JOIN T3 ON T3.B=T1.B
WHERE T3.C > 0
```

the `WHERE` condition is null-rejected for the second outer join operation but is not null-rejected for the first one.

If the `WHERE` condition is null-rejected for an outer join operation in a query, the outer join operation is replaced by an inner join operation.

For example, the preceding query is replaced with the query:

```
SELECT * FROM T1 LEFT JOIN T2 ON T2.A=T1.A
           INNER JOIN T3 ON T3.B=T1.B
WHERE T3.C > 0
```

For the original query, the optimizer would evaluate plans compatible with only one access order `T1, T2, T3`. For the replacing query, it additionally considers the access sequence `T3, T1, T2`.

A conversion of one outer join operation may trigger a conversion of another. Thus, the query:

```
SELECT * FROM T1 LEFT JOIN T2 ON T2.A=T1.A
           LEFT JOIN T3 ON T3.B=T2.B
WHERE T3.C > 0
```

will be first converted to the query:

```
SELECT * FROM T1 LEFT JOIN T2 ON T2.A=T1.A
           INNER JOIN T3 ON T3.B=T2.B
WHERE T3.C > 0
```

which is equivalent to the query:

```
SELECT * FROM (T1 LEFT JOIN T2 ON T2.A=T1.A), T3
```

```
WHERE T3.C > 0 AND T3.B=T2.B
```

Now the remaining outer join operation can be replaced by an inner join, too, because the condition `T3.B=T2.B` is null-rejected and we get a query without outer joins at all:

```
SELECT * FROM (T1 INNER JOIN T2 ON T2.A=T1.A), T3
WHERE T3.C > 0 AND T3.B=T2.B
```

Sometimes we succeed in replacing an embedded outer join operation, but cannot convert the embedding outer join. The following query:

```
SELECT * FROM T1 LEFT JOIN
      (T2 LEFT JOIN T3 ON T3.B=T2.B)
      ON T2.A=T1.A
WHERE T3.C > 0
```

is converted to:

```
SELECT * FROM T1 LEFT JOIN
      (T2 INNER JOIN T3 ON T3.B=T2.B)
      ON T2.A=T1.A
WHERE T3.C > 0,
```

That can be rewritten only to the form still containing the embedding outer join operation:

```
SELECT * FROM T1 LEFT JOIN
      (T2,T3)
      ON (T2.A=T1.A AND T3.B=T2.B)
WHERE T3.C > 0.
```

When trying to convert an embedded outer join operation in a query, we must take into account the join condition for the embedding outer join together with the `WHERE` condition. In the query:

```
SELECT * FROM T1 LEFT JOIN
      (T2 LEFT JOIN T3 ON T3.B=T2.B)
      ON T2.A=T1.A AND T3.C=T1.C
WHERE T3.D > 0 OR T1.D > 0
```

the `WHERE` condition is not null-rejected for the embedded outer join, but the join condition of the embedding outer join `T2.A=T1.A AND T3.C=T1.C` is null-rejected. So the query can be converted to:

```
SELECT * FROM T1 LEFT JOIN
      (T2, T3)
      ON T2.A=T1.A AND T3.C=T1.C AND T3.B=T2.B
WHERE T3.D > 0 OR T1.D > 0
```

### 8.13.11. Multi-Range Read Optimization

Reading rows using a range scan on a secondary index can result in many random disk accesses to the base table when the table is large and not stored in the storage engine's cache. With the Disk-Sweep Multi-Range Read (MRR) optimization, MySQL tries to reduce the number of random disk access for range scans by first scanning the index only and collecting the keys for the relevant rows. Then the keys are sorted and finally the rows are retrieved from the base table using the order of the primary key. The motivation for Disk-sweep MRR is to reduce the number of random disk accesses and instead achieve a more sequential scan of the base table data.

The Multi-Range Read optimization provides these benefits:

- MRR enables data rows to be accessed sequentially rather than in random order, based on index tuples. The server obtains a set of index tuples that satisfy the query conditions, sorts them according to data row ID order, and uses the sorted tuples to retrieve data rows in order. This makes data access more efficient and less expensive.
- MRR enables batch processing of requests for key access for operations that require access to data rows through index tuples, such as range index scans and equi-joins that use an index for the join

attribute. MRR iterates over a sequence of index ranges to obtain qualifying index tuples. As these results accumulate, they are used to access the corresponding data rows. It is not necessary to acquire all index tuples before starting to read data rows.

The following scenarios illustrate when MRR optimization can be advantageous:

Scenario A: MRR can be used for [InnoDB](#) and [MyISAM](#) tables for index range scans and equi-join operations.

1. A portion of the index tuples are accumulated in a buffer.
2. The tuples in the buffer are sorted by their data row ID.
3. Data rows are accessed according to the sorted index tuple sequence.

Scenario B: MRR can be used for [NDB](#) tables for multiple-range index scans or when performing an equi-join by an attribute.

1. A portion of ranges, possibly single-key ranges, is accumulated in a buffer on the central node where the query is submitted.
2. The ranges are sent to the execution nodes that access data rows.
3. The accessed rows are packed into packages and sent back to the central node.
4. The received packages with data rows are placed in a buffer.
5. Data rows are read from the buffer.

When MRR is used, the [Extra](#) column in [EXPLAIN](#) output shows [Using MRR](#).

[InnoDB](#) and [MyISAM](#) do not use MRR if full table rows need not be accessed to produce the query result. This is the case if results can be produced entirely on the basis on information in the index tuples (through a [covering index](#)); MRR provides no benefit.

Example query for which MRR can be used, assuming that there is an index on [\(key\\_part1, key\\_part2\)](#):

```
SELECT * FROM t
WHERE key_part1 >= 1000 AND key_part1 < 2000
AND key_part2 = 10000;
```

The index consists of tuples of [\(key\\_part1, key\\_part2\)](#) values, ordered first by [key\\_part1](#) and then by [key\\_part2](#).

Without MRR, an index scan covers all index tuples for the [key\\_part1](#) range from 1000 up to 2000, regardless of the [key\\_part2](#) value in these tuples. The scan does extra work to the extent that tuples in the range contain [key\\_part2](#) values other than 10000.

With MRR, the scan is broken up into multiple ranges, each for a single value of [key\\_part1](#) (1000, 1001, ..., 1999). Each of these scans need look only for tuples with [key\\_part2](#) = 10000. If the index contains many tuples for which [key\\_part2](#) is not 10000, MRR results in many fewer index tuples being read.

To express this using interval notation, the non-MRR scan must examine the index range [\[ {1000, 10000}, {2000, MIN\\_INT} \)](#), which may include many tuples other than those for which [key\\_part2](#) = 10000. The MRR scan examines multiple single-point intervals [\[ {1000, 10000} \]](#), ..., [\[ {1999, 10000} \]](#), which includes only tuples with [key\\_part2](#) = 10000.

Two [optimizer\\_switch \[520\]](#) system variable flags provide an interface to the use of MRR optimization. The [mrr](#) flag controls whether MRR is enabled. If [mrr](#) is enabled ([on](#)), the [mrr\\_cost\\_based](#) flag controls whether the optimizer attempts to make a cost-based choice between



using and not using MRR (`on`) or uses MRR whenever possible (`off`). By default, `mrr` is `on` and `mrr_cost_based` is `off` (use MRR whenever possible). See [Section 8.8.5.2, “Controlling Switchable Optimizations”](#).

For MRR, a storage engine uses the value of the `read_rnd_buffer_size` [533] system variable as a guideline for how much memory it can allocate for its buffer. The engine uses up to `read_rnd_buffer_size` [533] bytes and determines the number of ranges to process in a single pass.

## 8.13.12. Block Nested-Loop and Batched Key Access Joins

In MySQL 5.6, a Batched Key Access (BKA) Join algorithm is available that uses both index access to the joined table and a join buffer. The BKA algorithm supports inner join, outer join, and semi-join operations, including nested outer joins. Benefits of BKA include improved join performance due to more efficient table scanning. Also, the Block Nested-Loop (BNL) Join algorithm previously used only for inner joins is extended and can be employed for outer join and semi-join operations, including nested outer joins.

The following sections discuss the join buffer management that underlies the extension of the original BNL algorithm, the extended BNL algorithm, and the BKA algorithm. For information about semi-join strategies, see [Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”](#)

### 8.13.12.1. Join Buffer Management for Block Nested-Loop and Batched Key Access Algorithms

In MySQL 5.6, MySQL Server can employ join buffers to execute not only inner joins without index access to the inner table, but also outer joins and semi-joins that appear after subquery flattening. Moreover, a join buffer can be effectively used when there is an index access to the inner table.

The join buffer management code slightly more efficiently utilizes join buffer space when storing the values of the interesting row columns: No additional bytes are allocated in buffers for a row column if its value is `NULL`, and the minimum number of bytes is allocated for any value of the `VARCHAR` type.

The code supports two types of buffers, regular and incremental. Suppose that join buffer `B1` is employed to join tables `t1` and `t2` and the result of this operation is joined with table `t3` using join buffer `B2`:

- A regular join buffer contains columns from each join operand. If `B2` is a regular join buffer, each row `r` put into `B2` is composed of the columns of a row `r1` from `B1` and the interesting columns of a matching row `r2` from table `t2`.
- An incremental join buffer contains only columns from rows of the table produced by the second join operand. That is, it is incremental to a row from the first operand buffer. If `B2` is an incremental join buffer, it contains the interesting columns of the row `r2` together with a link to the row `r1` from `B1`.

Incremental join buffers are always incremental relative to a join buffer from an earlier join operation, so the buffer from the first join operation is always a regular buffer. In the example just given, the buffer `B1` used to join tables `t1` and `t2` must be a regular buffer.

Each row of the incremental buffer used for a join operation contains only the interesting columns of a row from the table to be joined. These columns are augmented with a reference to the interesting columns of the matched row from the table produced by the first join operand. Several rows in the incremental buffer can refer to the same row `r` whose columns are stored in the previous join buffers insofar as all these rows match row `r`.

Incremental buffers enable less frequent copying of columns from buffers used for previous join operations. This provides a savings in buffer space because in the general case a row produced by the first join operand can be matched by several rows produced by the second join operand. It is unnecessary to make several copies of a row from the first operand. Incremental buffers also provide a savings in processing time due to the reduction in copying time.

As of MySQL 5.6.3, the `block_nested_loop` and `batched_key_access` flags of the `optimizer_switch` [520] system variable control how the optimizer uses the Block Nested-Loop and Batched Key Access join algorithms. By default, `block_nested_loop` is `on` and `batched_key_access` is `off`. See [Section 8.8.5.2, “Controlling Switchable Optimizations”](#).

Before MySQL 5.6.3, the `optimizer_join_cache_level` [519] system variable controls join buffer management. For the possible values of this variable and their meanings, see the description in [Section 5.1.4, “Server System Variables”](#).

For information about semi-join strategies, see [Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”](#)

### 8.13.12.2. Block Nested-Loop Algorithm for Outer Joins and Semi-Joins

In MySQL 5.6, the original implementation of the BNL algorithm is extended to support outer join and semi-join operations.

When these operations are executed with a join buffer, each row put into the buffer is supplied with a match flag.

If an outer join operation is executed using a join buffer, each row of the table produced by the second operand is checked for a match against each row in the join buffer. When a match is found, a new extended row is formed (the original row plus columns from the second operand) and sent for further extensions by the remaining join operations. In addition, the match flag of the matched row in the buffer is enabled. After all rows of the table to be joined have been examined, the join buffer is scanned. Each row from the buffer that does not have its match flag enabled is extended by `NULL` complements (`NULL` values for each column in the second operand) and sent for further extensions by the remaining join operations.

As of MySQL 5.6.3, the `block_nested_loop` flag of the `optimizer_switch` [520] system variable controls how the optimizer uses the Block Nested-Loop algorithm. By default, `block_nested_loop` is `on`. See [Section 8.8.5.2, “Controlling Switchable Optimizations”](#).

Before MySQL 5.6.3, the `optimizer_join_cache_level` [519] system variable controls join buffer management. For the possible values of this variable and their meanings, see the description in [Section 5.1.4, “Server System Variables”](#).

In `EXPLAIN` output, use of BNL for a table is signified when the `Extra` value contains `Using join buffer (Block Nested Loop)` and the `type` value is `ALL` [900], `index` [899], or `range` [899].

For information about semi-join strategies, see [Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”](#)

### 8.13.12.3. Batched Key Access Joins

MySQL 5.6.3 implements a method of joining tables called the Batched Key Access (BKA) join algorithm. BKA can be applied when there is an index access to the table produced by the second join operand. Like the BNL join algorithm, the BKA join algorithm employs a join buffer to accumulate the interesting columns of the rows produced by the first operand of the join operation. Then the BKA algorithm builds keys to access the table to be joined for all rows in the buffer and submits these keys in a batch to the database engine for index lookups. The keys are submitted to the engine through the Multi-Range Read (MRR) interface (see [Section 8.13.11, “Multi-Range Read Optimization”](#)). After submission of the keys, the MRR engine functions perform lookups in the index in an optimal way, fetching the rows of the joined table found by these keys, and starts feeding the BKA join algorithm with matching rows. Each matching row is coupled with a reference to a row in the join buffer.

For BKA to be used, the `batched_key_access` flag of the `optimizer_switch` [520] system variable must be set to `on`. BKA uses MRR, so the `mrr` flag must also be `on`. Currently, the cost estimation for MRR is too pessimistic. Hence, it is also necessary for `mrr_cost_based` to be `off` for BKA to be used. The following setting enables BKA:

```
mysql> SET optimizer_switch='mrr=on,mrr_cost_based=off,batched_key_access=on';
```

There are two scenarios by which MRR functions execute:

- The first scenario is used for conventional disk-based storage engines such as [InnoDB](#) and [MyISAM](#). For these engines, usually the keys for all rows from the join buffer are submitted to the MRR interface at once. Engine-specific MRR functions perform index lookups for the submitted keys, get row IDs (or primary keys) from them, and then fetch rows for all these selected row IDs one by one by request from BKA algorithm. Every row is returned with an association reference that enables access to the matched row in the join buffer. The rows are fetched by the MRR functions in an optimal way: They are fetched in the row ID (primary key) order. This improves performance because reads are in disk order rather than random order.
- The second scenario is used for remote storage engines such as [NDB](#). A package of keys for a portion of rows from the join buffer, together with their associations, is sent by a MySQL Server (SQL node) to MySQL Cluster data nodes. In return, the SQL node receives a package (or several packages) of matching rows coupled with corresponding associations. The BKA join algorithm takes these rows and builds new joined rows. Then a new set of keys is sent to the data nodes and the rows from the returned packages are used to build new joined rows. The process continues until the last keys from the join buffer are sent to the data nodes, and the SQL node has received and joined all rows matching these keys. This improves performance because fewer key-bearing packages sent by the SQL node to the data nodes means fewer round trips between it and the data nodes to perform the join operation.

With the first scenario, a portion of the join buffer is reserved to store row IDs (primary keys) selected by index lookups and passed as a parameter to the MRR functions.

There is no special buffer to store keys built for rows from the join buffer. Instead, a function that builds the key for the next row in the buffer is passed as a parameter to the MRR functions.

In [EXPLAIN](#) output, use of BKA for a table is signified when the [Extra](#) value contains [Using join buffer \(Batched Key Access\)](#) and the [type](#) value is [ref \[898\]](#) or [eq\\_ref \[898\]](#).

### 8.13.13. ORDER BY Optimization

In some cases, MySQL can use an index to satisfy an [ORDER BY](#) clause without doing any extra sorting.

The index can also be used even if the [ORDER BY](#) does not match the index exactly, as long as all of the unused portions of the index and all the extra [ORDER BY](#) columns are constants in the [WHERE](#) clause. The following queries use the index to resolve the [ORDER BY](#) part:

```
SELECT * FROM t1
  ORDER BY key_part1,key_part2,... ;

SELECT * FROM t1
  WHERE key_part1 = constant
  ORDER BY key_part2;

SELECT * FROM t1
  ORDER BY key_part1 DESC, key_part2 DESC;

SELECT * FROM t1
  WHERE key_part1 = 1
  ORDER BY key_part1 DESC, key_part2 DESC;

SELECT * FROM t1
  WHERE key_part1 > constant
  ORDER BY key_part1 ASC;

SELECT * FROM t1
  WHERE key_part1 < constant
  ORDER BY key_part1 DESC;
```

```
SELECT * FROM t1
WHERE key_part1 = constant1 AND key_part2 > constant2
ORDER BY key_part2;
```

In some cases, MySQL *cannot* use indexes to resolve the `ORDER BY`, although it still uses indexes to find the rows that match the `WHERE` clause. These cases include the following:

- You use `ORDER BY` on different keys:

```
SELECT * FROM t1 ORDER BY key1, key2;
```

- You use `ORDER BY` on nonconsecutive parts of a key:

```
SELECT * FROM t1 WHERE key2=constant ORDER BY key_part2;
```

- You mix `ASC` and `DESC`:

```
SELECT * FROM t1 ORDER BY key_part1 DESC, key_part2 ASC;
```

- The key used to fetch the rows is not the same as the one used in the `ORDER BY`:

```
SELECT * FROM t1 WHERE key2=constant ORDER BY key1;
```

- You use `ORDER BY` with an expression that includes terms other than the key column name:

```
SELECT * FROM t1 ORDER BY ABS(key);
SELECT * FROM t1 ORDER BY -key;
```

- You are joining many tables, and the columns in the `ORDER BY` are not all from the first nonconstant table that is used to retrieve rows. (This is the first table in the `EXPLAIN` output that does not have a `const [898]` join type.)
- You have different `ORDER BY` and `GROUP BY` expressions.
- You index only a prefix of a column named in the `ORDER BY` clause. In this case, the index cannot be used to fully resolve the sort order. For example, if you have a `CHAR(20)` column, but index only the first 10 bytes, the index cannot distinguish values past the 10th byte and a `filesort` will be needed.
- The type of table index used does not store rows in order. For example, this is true for a `HASH` index in a `MEMORY` table.

Availability of an index for sorting may be affected by the use of column aliases. Suppose that the column `t1.a` is indexed. In this statement, the name of the column in the select list is `a`. It refers to `t1.a`, so for the reference to `a` in the `ORDER BY`, the index can be used:

```
SELECT a FROM t1 ORDER BY a;
```

In this statement, the name of the column in the select list is also `a`, but it is the alias name. It refers to `ABS(a)`, so for the reference to `a` in the `ORDER BY`, the index cannot be used:

```
SELECT ABS(a) AS a FROM t1 ORDER BY a;
```

In the following statement, the `ORDER BY` refers to a name that is not the name of a column in the select list. But there is a column in `t1` named `a`, so the `ORDER BY` uses that, and the index can be used. (The resulting sort order may be completely different from the order for `ABS(a)`, of course.)

```
SELECT ABS(a) AS b FROM t1 ORDER BY a;
```

By default, MySQL sorts all `GROUP BY col1, col2, ...` queries as if you specified `ORDER BY col1, col2, ...` in the query as well. If you include an `ORDER BY` clause explicitly that contains the same column list, MySQL optimizes it away without any speed penalty, although the sorting still occurs. If a query includes `GROUP BY` but you want to avoid the overhead of sorting the result, you can suppress sorting by specifying `ORDER BY NULL`. For example:

```
INSERT INTO foo
SELECT a, COUNT(*) FROM bar GROUP BY a ORDER BY NULL;
```

Relying on implicit `GROUP BY` sorting in MySQL 5.6 is deprecated. To achieve a specific sort order of grouped results, it is preferable to use an explicit `ORDER BY` clause. `GROUP BY` sorting is a MySQL extension that may change in a future release; for example, to make it possible for the optimizer to order groupings in whatever manner it deems most efficient and to avoid the sorting overhead.

With `EXPLAIN SELECT ... ORDER BY`, you can check whether MySQL can use indexes to resolve the query. It cannot if you see `Using filesort` in the `Extra` column. See [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#). `Filesort` uses a fixed-length row-storage format similar to that used by the `MEMORY` storage engine. Variable-length types such as `VARCHAR` are stored using a fixed length.

MySQL has two `filesort` algorithms for sorting and retrieving results. The original method uses only the `ORDER BY` columns. The modified method uses not just the `ORDER BY` columns, but all the columns used in the query.

The optimizer selects which `filesort` algorithm to use. It normally uses the modified algorithm except when `BLOB` or `TEXT` columns are involved, in which case it uses the original algorithm.

The original `filesort` algorithm works as follows:

1. Read all rows according to key or by table scanning. Rows that do not match the `WHERE` clause are skipped.
2. For each row, store a pair of values in a buffer (the sort key and the row pointer). The size of the buffer is the value of the `sort_buffer_size` [542] system variable.
3. When the buffer gets full, run a `qsort` (quicksort) on it and store the result in a temporary file. Save a pointer to the sorted block. (If all pairs fit into the sort buffer, no temporary file is created.)
4. Repeat the preceding steps until all rows have been read.
5. Do a multi-merge of up to `MERGEBUFF` (7) regions to one block in another temporary file. Repeat until all blocks from the first file are in the second file.
6. Repeat the following until there are fewer than `MERGEBUFF2` (15) blocks left.
7. On the last multi-merge, only the pointer to the row (the last part of the sort key) is written to a result file.
8. Read the rows in sorted order by using the row pointers in the result file. To optimize this, we read in a big block of row pointers, sort them, and use them to read the rows in sorted order into a row buffer. The size of the buffer is the value of the `read_rnd_buffer_size` [533] system variable. The code for this step is in the `sql/records.cc` source file.

One problem with this approach is that it reads rows twice: One time when evaluating the `WHERE` clause, and again after sorting the pair values. And even if the rows were accessed successively the first time (for example, if a table scan is done), the second time they are accessed randomly. (The sort keys are ordered, but the row positions are not.)

The modified `filesort` algorithm incorporates an optimization such that it records not only the sort key value and row position, but also the columns required for the query. This avoids reading the rows twice. The modified `filesort` algorithm works like this:

1. Read the rows that match the `WHERE` clause.
2. For each row, record a tuple of values consisting of the sort key value and row position, and also the columns required for the query.
3. Sort the tuples by sort key value

- Retrieve the rows in sorted order, but read the required columns directly from the sorted tuples rather than by accessing the table a second time.

Using the modified `filesort` algorithm, the tuples are longer than the pairs used in the original method, and fewer of them fit in the sort buffer (the size of which is given by `sort_buffer_size` [542]). As a result, it is possible for the extra I/O to make the modified approach slower, not faster. To avoid a slowdown, the optimization is used only if the total size of the extra columns in the sort tuple does not exceed the value of the `max_length_for_sort_data` [506] system variable. (A symptom of setting the value of this variable too high is a combination of high disk activity and low CPU activity.)

For slow queries for which `filesort` is not used, try lowering `max_length_for_sort_data` [506] to a value that is appropriate to trigger a `filesort`.

If you want to increase `ORDER BY` speed, check whether you can get MySQL to use indexes rather than an extra sorting phase. If this is not possible, you can try the following strategies:

- Increase the size of the `sort_buffer_size` [542] variable.
- Increase the size of the `read_rnd_buffer_size` [533] variable.
- Use less RAM per row by declaring columns only as large as they need to be to hold the values stored in them. For example, `CHAR(16)` is better than `CHAR(200)` if values never exceed 16 characters.
- Change `tmpdir` [559] to point to a dedicated file system with large amounts of free space. Also, this option accepts several paths that are used in round-robin fashion, so you can use this feature to spread the load across several directories. Paths should be separated by colon characters (":") on Unix and semicolon characters(";") on Windows. The paths should be for directories in file systems that are located on different *physical* disks, not different partitions on the same disk.

If an index is not used for `ORDER BY` but a `LIMIT` clause is also present, the optimizer may be able to avoid using a merge file and sort the rows in memory. For details, see [Section 8.2.1.3, "Optimizing LIMIT Queries"](#).

## 8.13.14. GROUP BY Optimization

The most general way to satisfy a `GROUP BY` clause is to scan the whole table and create a new temporary table where all rows from each group are consecutive, and then use this temporary table to discover groups and apply aggregate functions (if any). In some cases, MySQL is able to do much better than that and to avoid creation of temporary tables by using index access.

The most important preconditions for using indexes for `GROUP BY` are that all `GROUP BY` columns reference attributes from the same index, and that the index stores its keys in order (for example, this is a `BTREE` index and not a `HASH` index). Whether use of temporary tables can be replaced by index access also depends on which parts of an index are used in a query, the conditions specified for these parts, and the selected aggregate functions.

There are two ways to execute a `GROUP BY` query through index access, as detailed in the following sections. In the first method, the grouping operation is applied together with all range predicates (if any). The second method first performs a range scan, and then groups the resulting tuples.

In MySQL, `GROUP BY` is used for sorting, so the server may also apply `ORDER BY` optimizations to grouping. See [Section 8.13.13, "ORDER BY Optimization"](#).

### 8.13.14.1. Loose Index Scan

The most efficient way to process `GROUP BY` is when an index is used to directly retrieve the grouping columns. With this access method, MySQL uses the property of some index types that the keys are ordered (for example, `BTREE`). This property enables use of lookup groups in an index without having to consider all keys in the index that satisfy all `WHERE` conditions. This access method considers only



a fraction of the keys in an index, so it is called a *loose index scan*. When there is no `WHERE` clause, a loose index scan reads as many keys as the number of groups, which may be a much smaller number than that of all keys. If the `WHERE` clause contains range predicates (see the discussion of the [range \[899\]](#) join type in [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#)), a loose index scan looks up the first key of each group that satisfies the range conditions, and again reads the least possible number of keys. This is possible under the following conditions:

- The query is over a single table.
- The `GROUP BY` names only columns that form a leftmost prefix of the index and no other columns. (If, instead of `GROUP BY`, the query has a `DISTINCT` clause, all distinct attributes refer to columns that form a leftmost prefix of the index.) For example, if a table `t1` has an index on `(c1, c2, c3)`, loose index scan is applicable if the query has `GROUP BY c1, c2,`. It is not applicable if the query has `GROUP BY c2, c3` (the columns are not a leftmost prefix) or `GROUP BY c1, c2, c4` (`c4` is not in the index).
- The only aggregate functions used in the select list (if any) are `MIN()` [1284] and `MAX()` [1284], and all of them refer to the same column. The column must be in the index and must follow the columns in the `GROUP BY`.
- Any other parts of the index than those from the `GROUP BY` referenced in the query must be constants (that is, they must be referenced in equalities with constants), except for the argument of `MIN()` [1284] or `MAX()` [1284] functions.
- For columns in the index, full column values must be indexed, not just a prefix. For example, with `c1 VARCHAR(20)`, `INDEX (c1(10))`, the index cannot be used for loose index scan.

If loose index scan is applicable to a query, the `EXPLAIN` output shows `Using index for group-by` in the `Extra` column.

Assume that there is an index `idx(c1, c2, c3)` on table `t1(c1, c2, c3, c4)`. The loose index scan access method can be used for the following queries:

```
SELECT c1, c2 FROM t1 GROUP BY c1, c2;
SELECT DISTINCT c1, c2 FROM t1;
SELECT c1, MIN(c2) FROM t1 GROUP BY c1;
SELECT c1, c2 FROM t1 WHERE c1 < const GROUP BY c1, c2;
SELECT MAX(c3), MIN(c3), c1, c2 FROM t1 WHERE c2 > const GROUP BY c1, c2;
SELECT c2 FROM t1 WHERE c1 < const GROUP BY c1, c2;
SELECT c1, c2 FROM t1 WHERE c3 = const GROUP BY c1, c2;
```

The following queries cannot be executed with this quick select method, for the reasons given:

- There are aggregate functions other than `MIN()` [1284] or `MAX()` [1284]:

```
SELECT c1, SUM(c2) FROM t1 GROUP BY c1;
```

- The columns in the `GROUP BY` clause do not form a leftmost prefix of the index:

```
SELECT c1, c2 FROM t1 GROUP BY c2, c3;
```

- The query refers to a part of a key that comes after the `GROUP BY` part, and for which there is no equality with a constant:

```
SELECT c1, c3 FROM t1 GROUP BY c1, c2;
```

Were the query to include `WHERE c3 = const`, loose index scan could be used.

The loose index scan access method can be applied to other forms of aggregate function references in the select list, in addition to the `MIN()` [1284] and `MAX()` [1284] references already supported:

- `AVG(DISTINCT)` [1282], `SUM(DISTINCT)` [1285], and `COUNT(DISTINCT)` [1283] are supported. `AVG(DISTINCT)` [1282] and `SUM(DISTINCT)` [1285] take a single argument. `COUNT(DISTINCT)` [1283] can have more than one column argument.



- There must be no `GROUP BY` or `DISTINCT` clause in the query.
- The loose scan limitations described earlier still apply.

Assume that there is an index `idx(c1, c2, c3)` on table `t1(c1, c2, c3, c4)`. The loose index scan access method can be used for the following queries:

```
SELECT COUNT(DISTINCT c1), SUM(DISTINCT c1) FROM t1;
SELECT COUNT(DISTINCT c1, c2), COUNT(DISTINCT c2, c1) FROM t1;
```

Loose index scan is not applicable for the following queries:

```
SELECT DISTINCT COUNT(DISTINCT c1) FROM t1;
SELECT COUNT(DISTINCT c1) FROM t1 GROUP BY c1;
```

### 8.13.14.2. Tight Index Scan

A tight index scan may be either a full index scan or a range index scan, depending on the query conditions.

When the conditions for a loose index scan are not met, it still may be possible to avoid creation of temporary tables for `GROUP BY` queries. If there are range conditions in the `WHERE` clause, this method reads only the keys that satisfy these conditions. Otherwise, it performs an index scan. Because this method reads all keys in each range defined by the `WHERE` clause, or scans the whole index if there are no range conditions, we term it a *tight index scan*. With a tight index scan, the grouping operation is performed only after all keys that satisfy the range conditions have been found.

For this method to work, it is sufficient that there is a constant equality condition for all columns in a query referring to parts of the key coming before or in between parts of the `GROUP BY` key. The constants from the equality conditions fill in any “gaps” in the search keys so that it is possible to form complete prefixes of the index. These index prefixes then can be used for index lookups. If we require sorting of the `GROUP BY` result, and it is possible to form search keys that are prefixes of the index, MySQL also avoids extra sorting operations because searching with prefixes in an ordered index already retrieves all the keys in order.

Assume that there is an index `idx(c1, c2, c3)` on table `t1(c1, c2, c3, c4)`. The following queries do not work with the loose index scan access method described earlier, but still work with the tight index scan access method.

- There is a gap in the `GROUP BY`, but it is covered by the condition `c2 = 'a'`:

```
SELECT c1, c2, c3 FROM t1 WHERE c2 = 'a' GROUP BY c1, c3;
```

- The `GROUP BY` does not begin with the first part of the key, but there is a condition that provides a constant for that part:

```
SELECT c1, c2, c3 FROM t1 WHERE c1 = 'a' GROUP BY c2, c3;
```

### 8.13.15. DISTINCT Optimization

`DISTINCT` combined with `ORDER BY` needs a temporary table in many cases.

Because `DISTINCT` may use `GROUP BY`, learn how MySQL works with columns in `ORDER BY` or `HAVING` clauses that are not part of the selected columns. See [Section 12.17.3, “MySQL Extensions to GROUP BY”](#).

In most cases, a `DISTINCT` clause can be considered as a special case of `GROUP BY`. For example, the following two queries are equivalent:

```
SELECT DISTINCT c1, c2, c3 FROM t1
WHERE c1 > const;
```

```
SELECT c1, c2, c3 FROM t1
WHERE c1 > const GROUP BY c1, c2, c3;
```

Due to this equivalence, the optimizations applicable to `GROUP BY` queries can be also applied to queries with a `DISTINCT` clause. Thus, for more details on the optimization possibilities for `DISTINCT` queries, see [Section 8.13.14, “GROUP BY Optimization”](#).

When combining `LIMIT row_count` with `DISTINCT`, MySQL stops as soon as it finds `row_count` unique rows.

If you do not use columns from all tables named in a query, MySQL stops scanning any unused tables as soon as it finds the first match. In the following case, assuming that `t1` is used before `t2` (which you can check with `EXPLAIN`), MySQL stops reading from `t2` (for any particular row in `t1`) when it finds the first row in `t2`:

```
SELECT DISTINCT t1.a FROM t1, t2 where t1.a=t2.a;
```

## 8.13.16. Subquery Optimization

The MySQL query optimizer has different strategies available to evaluate subqueries. For `IN` (or `=ANY`) subqueries, the optimizer has these choices:

- Semi-join
- Materialization
- `EXISTS` strategy

For `NOT IN` (or `<>ANY`) subqueries, the optimizer has these choices:

- Materialization
- `EXISTS` strategy

The following sections provide more information about these optimization strategies.

### 8.13.16.1. Optimizing Subqueries with Semi-Join Transformations

As of MySQL 5.6.5, the optimizer uses semi-join strategies to improve subquery execution, as described in this section.

For an inner join between two tables, the join returns a row from one table as many times as there are matches in the other table. But for some questions, the only information that matters is whether there is a match, not the number of matches. Suppose that there are tables named `class` and `roster` that list classes in a course curriculum and class rosters (students enrolled in each class), respectively. To list the classes that actually have students enrolled, you could use this join:

```
SELECT class.class_num, class.class_name
FROM class INNER JOIN roster
WHERE class.class_num = roster.class_num;
```

However, the result lists each class once for each enrolled student. For the question being asked, this is unnecessary duplication of information.

Assuming that `class_num` is a primary key in the `class` table, duplicate suppression could be achieved by using `SELECT DISTINCT`, but it is inefficient to generate all matching rows first only to eliminate duplicates later.

The same duplicate-free result can be obtained by using a subquery:

```
SELECT class_num, class_name
FROM class
```

```
WHERE class_num IN (SELECT class_num FROM roster);
```

Here, the optimizer can recognize that the `IN` clause requires the subquery to return only one instance of each class number from the `roster` table. In this case, the query can be executed as a *semi-join*—that is, an operation that returns only one instance of each row in `class` that is matched by rows in `roster`.

Before MySQL 5.6.6, the outer query specification was limited to simple table scans or inner joins using comma syntax, and view references were not possible. As of 5.6.6, outer join and inner join syntax is permitted in the outer query specification, and the restriction that table references must be base tables has been lifted.

In MySQL, a subquery must satisfy these criteria to be handled as a semi-join:

- It must be an `IN` (or `=ANY`) subquery that appears at the top level of the `WHERE` or `ON` clause, possibly as a term in an `AND` expression. For example:

```
SELECT ...
FROM ot1, ...
WHERE (oe1, ...) IN (SELECT ie1, ... FROM it1, ... WHERE ...);
```

Here, `ot_i` and `it_i` represent tables in the outer and inner parts of the query, and `oe_i` and `ie_i` represent expressions that refer to columns in the outer and inner tables.

- It must be a single `SELECT` without `UNION` constructs.
- It must not contain a `GROUP BY` or `HAVING` clause or aggregate functions.
- It must not have `ORDER BY` with `LIMIT`.
- The number of outer and inner tables together must be less than the maximum number of tables permitted in a join.

The subquery may be correlated or uncorrelated. `DISTINCT` is permitted, as is `LIMIT` unless `ORDER BY` is also used.

If a subquery meets the preceding criteria, MySQL converts it to a semi-join and makes a cost-based choice from these strategies:

- Convert the subquery to a join, or use table pullout and run the query as an inner join between subquery tables and outer tables. Table pullout pulls a table out from the subquery to the outer query.
- Duplicate Weedout: Run the semi-join as if it was a join and remove duplicate records using a temporary table.
- FirstMatch: When scanning the inner tables for row combinations and there are multiple instances of a given value group, choose one rather than returning them all. This "shortcuts" scanning and eliminates production of unnecessary rows.
- LooseScan: Scan a subquery table using an index that enables a single value to be chosen from each subquery's value group.
- Materialize the subquery into a temporary table with an index and use the temporary table to perform a join. The index is used to remove duplicates. The index might also be used later for lookups when joining the temporary table with the outer tables; if not, the table is scanned.

Each of these strategies except Duplicate Weedout can be enabled or disabled using the `optimizer_switch` [520] system variable. The `semijoin` flag controls whether semi-joins are used. If it is set to `on`, the `firstmatch`, `loosescan`, and `materialization` flags enable finer control over the permitted semi-join strategies. These flags are `on` by default. See [Section 8.8.5.2, "Controlling Switchable Optimizations"](#).

The use of semi-join strategies is indicated in `EXPLAIN` output as follows:

- Semi-joined tables show up in the outer select. `EXPLAIN EXTENDED` plus `SHOW WARNINGS` shows the rewritten query, which displays the semi-join structure. From this you can get an idea about which tables were pulled out of the semi-join. If a subquery was converted to a semi-join, you will see that the subquery predicate is gone and its tables and `WHERE` clause were merged into the outer query join list and `WHERE` clause.
- Temporary table use for Duplicate Weedout is indicated by `Start temporary` and `End temporary` in the `Extra` column. Tables that were not pulled out and are in the range of `EXPLAIN` output rows covered by `Start temporary` and `End temporary` will have their `rowid` in the temporary table.
- `FirstMatch(tbl_name)` in the `Extra` column indicates join shortcutting.
- `LooseScan(m..n)` in the `Extra` column indicates use of the LooseScan strategy. `m` and `n` are key part numbers.
- As of MySQL 5.6.7, temporary table use for materialization is indicated by rows with a `select_type` value of `MATERIALIZED` and rows with a `table` value of `<subqueryN>`.

Before MySQL 5.6.7, temporary table use for materialization is indicated in the `Extra` column by `Materialize` if a single table is used, or by `Start materialize` and `End materialize` if multiple tables are used. If `Scan` is present, no temporary table index is used for table reads. Otherwise, an index lookup is used.

### 8.13.16.2. Optimizing Subqueries with Subquery Materialization

As of MySQL 5.6.5, the optimizer uses subquery materialization as a strategy that enables more efficient subquery processing.

If materialization is not used, the optimizer sometimes rewrites a noncorrelated subquery as a correlated subquery. For example, the following `IN` subquery is noncorrelated (`where_condition` involves only columns from `t2` and not `t1`):

```
SELECT * FROM t1
WHERE t1.a IN (SELECT t2.b FROM t2 WHERE where_condition);
```

The optimizer might rewrite this as an `EXISTS` correlated subquery:

```
SELECT * FROM t1
WHERE EXISTS (SELECT t2.b FROM t2 WHERE where_condition AND t1.a=t2.b);
```

Subquery materialization using a temporary table avoids such rewrites and makes it possible to execute the subquery only once rather than once per row of the outer query. Materialization speeds up query execution by generating a subquery result as a temporary table, normally in memory. The first time MySQL needs the subquery result, it materializes that result into a temporary table. Any subsequent time the result is needed, MySQL refers again to the temporary table. The table is indexed with a hash index to make lookups fast and inexpensive. The index is unique, which makes the table smaller because it has no duplicates.

Subquery materialization attempts to use an in-memory temporary table when possible, falling back to on-disk storage if the table becomes too large. See [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).

For subquery materialization to be used in MySQL, the `materialization` flag of the `optimizer_switch` [520] system variable must be `on`. Materialization then applies to subquery predicates that appear anywhere (in the select list, `WHERE`, `ON`, `GROUP BY`, `HAVING`, or `ORDER BY`), for predicates that fall into any of these use cases:

- The predicate has this form, when no outer expression `oe_i` or inner expression `ie_i` is nullable. `N` can be 1 or larger.

```
(oe_1, oe_2, ..., oe_N) [NOT] IN (SELECT ie_1, ie_2, ..., ie_N ...)
```

- The predicate has this form, when there is a single outer expression *oe* and inner expression *ie*. The expressions can be nullable.

```
oe [NOT] IN (SELECT ie ...)
```

- The predicate is `IN` or `NOT IN` and a result of `UNKNOWN (NULL)` has the same meaning as a result of `FALSE`.

The following examples illustrate how the requirement for equivalence of `UNKNOWN` and `FALSE` predicate evaluation affects whether subquery materialization can be used. Assume that *where\_condition* involves columns only from *t2* and not *t1* so that the subquery is noncorrelated.

This query is subject to materialization:

```
SELECT * FROM t1
WHERE t1.a IN (SELECT t2.b FROM t2 WHERE where_condition);
```

Here, it does not matter whether the `IN` predicate returns `UNKNOWN` or `FALSE`. Either way, the row from *t1* is not included in the query result.

An example where subquery materialization will not be used is the following query, where *t2.b* is a nullable column.

```
SELECT * FROM t1
WHERE (t1.a,t1.b) NOT IN (SELECT t2.a,t2.b FROM t2
                        WHERE where_condition);
```

Use of `EXPLAIN` with a query can give some indication of whether the optimizer uses subquery materialization. Compared to query execution that does not use materialization, *select\_type* may change from `DEPENDENT SUBQUERY` to `SUBQUERY`. This indicates that, for a subquery that would be executed once per outer row, materialization enables the subquery to be executed just once. In addition, for `EXPLAIN EXTENDED`, the text displayed by a following `SHOW WARNINGS` will include `materialize materialize` and `materialized-subquery (materialized subselect` before MySQL 5.6.6).

### 8.13.16.3. Optimizing Subqueries in the `FROM` Clause (Derived Tables)

As of MySQL 5.6.3, the optimizer more efficiently handles subqueries in the `FROM` clause (that is, derived tables):

- Materialization of subqueries in the `FROM` clause is postponed until their contents are needed during query execution, which improves performance:
  - Previously, subqueries in the `FROM` clause were materialized for `EXPLAIN SELECT` statements. This resulted in partial `SELECT` execution, even though the purpose of `EXPLAIN` is to obtain query plan information, not to execute the query. This materialization no longer occurs, so `EXPLAIN` is faster for such queries.
  - For non-`EXPLAIN` queries, delay of materialization may result in not having to do it at all. Consider a query that joins the result of a subquery in the `FROM` clause to another table: If the optimizer processes that other table first and finds that it returns no rows, the join need not be carried out further and the optimizer can completely skip materializing the subquery.
- During query execution, the optimizer may add an index to a derived table to speed up row retrieval from it.

Consider the following `EXPLAIN` statement, for which a subquery appears in the `FROM` clause of a `SELECT` query:

```
EXPLAIN SELECT * FROM (SELECT * FROM t1);
```

The optimizer avoids materializing the subquery by delaying it until the result is needed during `SELECT` execution. In this case, the query is not executed, so the result is never needed.

Even for queries that are executed, delay of subquery materialization may permit the optimizer to avoid materialization entirely. Consider the following query, which joins the result of a subquery in the `FROM` clause to another table:

```
SELECT * FROM t1
JOIN (SELECT t2.f1 FROM t2) AS derived_t2 ON t1.f2=derived_t2.f1
WHERE t1.f1 > 0;
```

If the optimization processes `t1` first and the `WHERE` clause produces an empty result, the join must necessarily be empty and the subquery need not be materialized.

In the worst case (derived tables are materialized), query execution will take the same time as before MySQL 5.6.3 because no additional work is done. In the best case (derived tables are not materialized), query execution will be quicker by the time needed to perform materialization.

For cases when materialization is required for a subquery in the `FROM` clause, the optimizer may speed up access to the result by adding an index to the materialized table. If such an index would permit [ref \[898\]](#) access to the table, it can greatly reduce amount of data that must be read during query execution. Consider the following query:

```
SELECT * FROM t1
JOIN (SELECT * FROM t2) AS derived_t2 ON t1.f1=derived_t2.f1;
```

The optimizer constructs an index over column `f1` from `derived_t2` if doing so would permit the use of [ref \[898\]](#) access for the lowest cost execution plan. After adding the index, the optimizer can treat the materialized derived table the same as a usual table with an index, and it benefits similarly from the generated index. The overhead of index creation is negligible compared to the cost of query execution without the index. If [ref \[898\]](#) access would result in higher cost than some other access method, no index is created and the optimizer loses nothing.

#### 8.13.16.4. Optimizing Subqueries with `EXISTS` Strategy

Certain optimizations are applicable to comparisons that use the `IN` operator to test subquery results (or that use `=ANY`, which is equivalent). This section discusses these optimizations, particularly with regard to the challenges that `NULL` values present. The last part of the discussion includes suggestions on what you can do to help the optimizer.

Consider the following subquery comparison:

```
outer_expr IN (SELECT inner_expr FROM ... WHERE subquery_where)
```

MySQL evaluates queries “from outside to inside.” That is, it first obtains the value of the outer expression `outer_expr`, and then runs the subquery and captures the rows that it produces.

A very useful optimization is to “inform” the subquery that the only rows of interest are those where the inner expression `inner_expr` is equal to `outer_expr`. This is done by pushing down an appropriate equality into the subquery's `WHERE` clause. That is, the comparison is converted to this:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where AND outer_expr=inner_expr)
```

After the conversion, MySQL can use the pushed-down equality to limit the number of rows that it must examine when evaluating the subquery.

More generally, a comparison of `N` values to a subquery that returns `N`-value rows is subject to the same conversion. If `oe_i` and `ie_i` represent corresponding outer and inner expression values, this subquery comparison:

```
(oe_1, ..., oe_N) IN
(SELECT ie_1, ..., ie_N FROM ... WHERE subquery_where)
```

Becomes:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where
      AND oe_1 = ie_1
      AND ...
      AND oe_N = ie_N)
```

The following discussion assumes a single pair of outer and inner expression values for simplicity.

The conversion just described has its limitations. It is valid only if we ignore possible `NULL` values. That is, the “pushdown” strategy works as long as both of these two conditions are true:

- `outer_expr` and `inner_expr` cannot be `NULL`.
- You do not need to distinguish `NULL` from `FALSE` subquery results. (If the subquery is a part of an `OR` [1172] or `AND` [1171] expression in the `WHERE` clause, MySQL assumes that you do not care.)

When either or both of those conditions do not hold, optimization is more complex.

Suppose that `outer_expr` is known to be a non-`NULL` value but the subquery does not produce a row such that `outer_expr = inner_expr`. Then `outer_expr IN (SELECT ...)` evaluates as follows:

- `NULL`, if the `SELECT` produces any row where `inner_expr` is `NULL`
- `FALSE`, if the `SELECT` produces only non-`NULL` values or produces nothing

In this situation, the approach of looking for rows with `outer_expr = inner_expr` is no longer valid. It is necessary to look for such rows, but if none are found, also look for rows where `inner_expr` is `NULL`. Roughly speaking, the subquery can be converted to:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where AND
      (outer_expr=inner_expr OR inner_expr IS NULL))
```

The need to evaluate the extra `IS NULL` [1168] condition is why MySQL has the `ref_or_null` [899] access method:

```
mysql> EXPLAIN
-> SELECT outer_expr IN (SELECT t2.maybe_null_key
->                        FROM t2, t3 WHERE ...)
-> FROM t1;
***** 1. row *****
      id: 1
      select_type: PRIMARY
      table: t1
...
***** 2. row *****
      id: 2
      select_type: DEPENDENT SUBQUERY
      table: t2
      type: ref_or_null
possible_keys: maybe_null_key
      key: maybe_null_key
      key_len: 5
      ref: func
      rows: 2
      Extra: Using where; Using index
...

```

The `unique_subquery` [899] and `index_subquery` [899] subquery-specific access methods also have “or `NULL`” variants. However, they are not visible in `EXPLAIN` output, so you must use `EXPLAIN EXTENDED` followed by `SHOW WARNINGS` (note the `checking NULL` in the warning message):

```
mysql> EXPLAIN EXTENDED
-> SELECT outer_expr IN (SELECT maybe_null_key FROM t2) FROM t1\G
***** 1. row *****
      id: 1

```



```

select_type: PRIMARY
table: t1
...
***** 2. row *****
      id: 2
select_type: DEPENDENT SUBQUERY
table: t2
type: index_subquery
possible_keys: maybe_null_key
key: maybe_null_key
key_len: 5
ref: func
rows: 2
Extra: Using index

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Note
Code: 1003
Message: select (`test`.`t1`.`outer_expr`,
              ((`test`.`t1`.`outer_expr`) in t2 on
              maybe_null_key checking NULL)) AS `outer_expr` IN (SELECT
              maybe_null_key FROM t2)` from `test`.`t1`

```

The additional `OR ... IS NULL` condition makes query execution slightly more complicated (and some optimizations within the subquery become inapplicable), but generally this is tolerable.

The situation is much worse when `outer_expr` can be `NULL`. According to the SQL interpretation of `NULL` as “unknown value,” `NULL IN (SELECT inner_expr ...)` should evaluate to:

- `NULL`, if the `SELECT` produces any rows
- `FALSE`, if the `SELECT` produces no rows

For proper evaluation, it is necessary to be able to check whether the `SELECT` has produced any rows at all, so `outer_expr = inner_expr` cannot be pushed down into the subquery. This is a problem, because many real world subqueries become very slow unless the equality can be pushed down.

Essentially, there must be different ways to execute the subquery depending on the value of `outer_expr`.

The optimizer chooses SQL compliance over speed, so it accounts for the possibility that `outer_expr` might be `NULL`.

If `outer_expr` is `NULL`, to evaluate the following expression, it is necessary to run the `SELECT` to determine whether it produces any rows:

```
NULL IN (SELECT inner_expr FROM ... WHERE subquery_where)
```

It is necessary to run the original `SELECT` here, without any pushed-down equalities of the kind mentioned earlier.

On the other hand, when `outer_expr` is not `NULL`, it is absolutely essential that this comparison:

```
outer_expr IN (SELECT inner_expr FROM ... WHERE subquery_where)
```

be converted to this expression that uses a pushed-down condition:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where AND outer_expr=inner_expr)
```

Without this conversion, subqueries will be slow. To solve the dilemma of whether to push down or not push down conditions into the subquery, the conditions are wrapped in “trigger” functions. Thus, an expression of the following form:

```
outer_expr IN (SELECT inner_expr FROM ... WHERE subquery_where)
```

is converted into:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where
      AND trigcond(outer_expr=inner_expr))
```

More generally, if the subquery comparison is based on several pairs of outer and inner expressions, the conversion takes this comparison:

```
(oe_1, ..., oe_N) IN (SELECT ie_1, ..., ie_N FROM ... WHERE subquery_where)
```

and converts it to this expression:

```
EXISTS (SELECT 1 FROM ... WHERE subquery_where
      AND trigcond(oe_1=ie_1)
      AND ...
      AND trigcond(oe_N=ie_N)
    )
```

Each `trigcond(X)` is a special function that evaluates to the following values:

- `X` when the “linked” outer expression `oe_i` is not `NULL`
- `TRUE` when the “linked” outer expression `oe_i` is `NULL`

Note that trigger functions are *not* triggers of the kind that you create with `CREATE TRIGGER`.

Equalities that are wrapped into `trigcond()` functions are not first class predicates for the query optimizer. Most optimizations cannot deal with predicates that may be turned on and off at query execution time, so they assume any `trigcond(X)` to be an unknown function and ignore it. At the moment, triggered equalities can be used by those optimizations:

- Reference optimizations: `trigcond(X=Y [OR Y IS NULL])` can be used to construct `ref` [898], `eq_ref` [898], or `ref_or_null` [899] table accesses.
- Index lookup-based subquery execution engines: `trigcond(X=Y)` can be used to construct `unique_subquery` [899] or `index_subquery` [899] accesses.
- Table-condition generator: If the subquery is a join of several tables, the triggered condition will be checked as soon as possible.

When the optimizer uses a triggered condition to create some kind of index lookup-based access (as for the first two items of the preceding list), it must have a fallback strategy for the case when the condition is turned off. This fallback strategy is always the same: Do a full table scan. In `EXPLAIN` output, the fallback shows up as `Full scan on NULL key` in the `Extra` column:

```
mysql> EXPLAIN SELECT t1.col1,
-> t1.col1 IN (SELECT t2.key1 FROM t2 WHERE t2.col2=t1.col2) FROM t1\G
***** 1. row *****
      id: 1
      select_type: PRIMARY
      table: t1
      ...
***** 2. row *****
      id: 2
      select_type: DEPENDENT SUBQUERY
      table: t2
      type: index_subquery
possible_keys: key1
      key: key1
      key_len: 5
      ref: func
      rows: 2
      Extra: Using where; Full scan on NULL key
```

If you run `EXPLAIN EXTENDED` followed by `SHOW WARNINGS`, you can see the triggered condition:

```
***** 1. row *****
Level: Note
Code: 1003
```

```

Message: select `test`.`t1`.`col1` AS `col1`,
<in_optimizer>(`test`.`t1`.`col1`,
<exists>(<index_lookup>(<cache>(`test`.`t1`.`col1`) in t2
on key1 checking NULL
where (`test`.`t2`.`col2` = `test`.`t1`.`col2`) having
trigcond(<is_not_null_test>(`test`.`t2`.`key1`)))) AS
`t1.col1 IN (select t2.key1 from t2 where t2.col2=t1.col2)`
from `test`.`t1`

```

The use of triggered conditions has some performance implications. A `NULL IN (SELECT ...)` expression now may cause a full table scan (which is slow) when it previously did not. This is the price paid for correct results (the goal of the trigger-condition strategy was to improve compliance and not speed).

For multiple-table subqueries, execution of `NULL IN (SELECT ...)` will be particularly slow because the join optimizer does not optimize for the case where the outer expression is `NULL`. It assumes that subquery evaluations with `NULL` on the left side are very rare, even if there are statistics that indicate otherwise. On the other hand, if the outer expression might be `NULL` but never actually is, there is no performance penalty.

To help the query optimizer better execute your queries, use these tips:

- Declare a column as `NOT NULL` if it really is. (This also helps other aspects of the optimizer by simplifying condition testing for the column.)
- If you do not need to distinguish a `NULL` from `FALSE` subquery result, you can easily avoid the slow execution path. Replace a comparison that looks like this:

```
outer_expr IN (SELECT inner_expr FROM ...)
```

with this expression:

```
(outer_expr IS NOT NULL) AND (outer_expr IN (SELECT inner_expr FROM ...))
```

Then `NULL IN (SELECT ...)` will never be evaluated because MySQL stops evaluating `AND [1171]` parts as soon as the expression result is clear.

The `subquery_materialization_cost_based` enables control over the choice between subquery materialization and `IN -> EXISTS` subquery transformation. See [Section 8.8.5.2, "Controlling Switchable Optimizations"](#).



---

# Chapter 9. Language Structure

## Table of Contents

9.1. Literal Values .....	1011
9.1.1. String Literals .....	1011
9.1.2. Number Literals .....	1013
9.1.3. Date and Time Literals .....	1014
9.1.4. Hexadecimal Literals .....	1016
9.1.5. Boolean Literals .....	1016
9.1.6. Bit-Field Literals .....	1016
9.1.7. <code>NULL</code> Values .....	1017
9.2. Schema Object Names .....	1017
9.2.1. Identifier Qualifiers .....	1019
9.2.2. Identifier Case Sensitivity .....	1020
9.2.3. Mapping of Identifiers to File Names .....	1021
9.2.4. Function Name Parsing and Resolution .....	1023
9.3. Reserved Words .....	1026
9.4. User-Defined Variables .....	1029
9.5. Expression Syntax .....	1032
9.6. Comment Syntax .....	1034

This chapter discusses the rules for writing the following elements of [SQL](#) statements when using MySQL:

- Literal values such as strings and numbers
- Identifiers such as database, table, and column names
- Reserved words
- User-defined and system variables
- Comments

## 9.1. Literal Values

This section describes how to write literal values in MySQL. These include strings, numbers, hexadecimal values, boolean values, and `NULL`. The section also covers the various nuances and “gotchas” that you may run into when dealing with these basic types in MySQL.

### 9.1.1. String Literals

A string is a sequence of bytes or characters, enclosed within either single quote (“’”) or double quote (“””) characters. Examples:

```
'a string'
"another string"
```

Quoted strings placed next to each other are concatenated to a single string. The following lines are equivalent:

```
'a string'
'a' ' ' 'string'
```

If the [ANSI\\_QUOTES \[603\]](#) SQL mode is enabled, string literals can be quoted only within single quotation marks because a string quoted within double quotation marks is interpreted as an identifier.

A *binary string* is a string of bytes that has no character set or collation. A *nonbinary string* is a string of characters that has a character set and collation. For both types of strings, comparisons are based on

the numeric values of the string unit. For binary strings, the unit is the byte. For nonbinary strings the unit is the character and some character sets support multi-byte characters. Character value ordering is a function of the string collation.

String literals may have an optional character set introducer and `COLLATE` clause:

```
[_charset_name]'string' [COLLATE collation_name]
```

Examples:

```
SELECT _latin1'string';
SELECT _latin1'string' COLLATE latin1_danish_ci;
```

You can use `N'literal'` (or `n'literal'`) to create a string in the national character set. These statements are equivalent:

```
SELECT N'some text';
SELECT n'some text';
SELECT _utf8'some text';
```

For more information about these forms of string syntax, see [Section 10.1.3.5, “Character String Literal Character Set and Collation”](#), and [Section 10.1.3.6, “National Character Set”](#).

Within a string, certain sequences have special meaning unless the `NO_BACKSLASH_ESCAPES` [604] SQL mode is enabled. Each of these sequences begins with a backslash (“\”), known as the *escape character*. MySQL recognizes the escape sequences shown in [Table 9.1, “Special Character Escape Sequences”](#). For all other escape sequences, backslash is ignored. That is, the escaped character is interpreted as if it was not escaped. For example, “\x” is just “x”. These sequences are case sensitive. For example, “\b” is interpreted as a backspace, but “\B” is interpreted as “B”. Escape processing is done according to the character set indicated by the `character_set_connection` [467] system variable. This is true even for strings that are preceded by an introducer that indicates a different character set, as discussed in [Section 10.1.3.5, “Character String Literal Character Set and Collation”](#).

**Table 9.1. Special Character Escape Sequences**

Escape Sequence	Character Represented by Sequence
\0	An ASCII NUL (0x00) character.
\'	A single quote (“'”) character.
\"	A double quote (“””) character.
\b	A backspace character.
\n	A newline (linefeed) character.
\r	A carriage return character.
\t	A tab character.
\z	ASCII 26 (Control+Z). See note following the table.
\\	A backslash (“\”) character.
\%	A “%” character. See note following the table.
\_	A “_” character. See note following the table.

The ASCII 26 character can be encoded as “\z” to enable you to work around the problem that ASCII 26 stands for END-OF-FILE on Windows. ASCII 26 within a file causes problems if you try to use `mysql db_name < file_name`.

The “\%” and “\\_” sequences are used to search for literal instances of “%” and “\_” in pattern-matching contexts where they would otherwise be interpreted as wildcard characters. See the description of the `LIKE` [1190] operator in [Section 12.5.1, “String Comparison Functions”](#). If you use “\%” or “\\_” outside of pattern-matching contexts, they evaluate to the strings “\%” and “\\_”, not to “%” and “\_”.

There are several ways to include quote characters within a string:

- A “'” inside a string quoted with “'” may be written as “''”.
- A “”” inside a string quoted with “”” may be written as “"""”.
- Precede the quote character by an escape character (“\”).
- A “'” inside a string quoted with “”” needs no special treatment and need not be doubled or escaped. In the same way, “”” inside a string quoted with “'” needs no special treatment.

The following `SELECT` statements demonstrate how quoting and escaping work:

```
mysql> SELECT 'hello', '"hello"', '"hello"', 'hel'lo', '\hello';
+-----+-----+-----+-----+-----+
| hello | "hello" | '"hello"' | hel'lo | \hello |
+-----+-----+-----+-----+-----+

mysql> SELECT "hello", "'hello'", "'hello'", "hel"lo", "\"hello";
+-----+-----+-----+-----+-----+
| hello | 'hello' | "'hello'" | hel"lo | "\"hello" |
+-----+-----+-----+-----+-----+

mysql> SELECT 'This\nIs\nFour\nLines';
+-----+
| This
| Is
| Four
| Lines |
+-----+

mysql> SELECT 'disappearing\ backslash';
+-----+
| disappearing backslash |
+-----+
```

If you want to insert binary data into a string column (such as a `BLOB` column), you should represent certain characters by escape sequences. Backslash (“\”) and the quote character used to quote the string must be escaped. In certain client environments, it may also be necessary to escape `NUL` or Control+Z. The `mysql` client truncates quoted strings containing `NUL` characters if they are not escaped, and Control+Z may be taken for END-OF-FILE on Windows if not escaped. For the escape sequences that represent each of these characters, see [Table 9.1, “Special Character Escape Sequences”](#).

When writing application programs, any string that might contain any of these special characters must be properly escaped before the string is used as a data value in an SQL statement that is sent to the MySQL server. You can do this in two ways:

- Process the string with a function that escapes the special characters. In a C program, you can use the `mysql_real_escape_string()` C API function to escape characters. See [Section 21.8.6.54, “mysql\\_real\\_escape\\_string\(\)”](#). Within SQL statements that construct other SQL statements, you can use the `QUOTE()` [1183] function. The Perl DBI interface provides a `quote` method to convert special characters to the proper escape sequences. See [Section 21.10, “MySQL Perl API”](#). Other language interfaces may provide a similar capability.
- As an alternative to explicitly escaping special characters, many MySQL APIs provide a placeholder capability that enables you to insert special markers into a statement string, and then bind data values to them when you issue the statement. In this case, the API takes care of escaping special characters in the values for you.

## 9.1.2. Number Literals

Integers are represented as a sequence of digits. Floats use “.” as a decimal separator. Either type of number may be preceded by “-” or “+” to indicate a negative or positive value, respectively

Examples of valid integers:

```
1221
```



```
0
-32
```

Examples of valid floating-point numbers:

```
294.42
-32032.6809e+10
148.00
```

An integer may be used in a floating-point context; it is interpreted as the equivalent floating-point number.

### 9.1.3. Date and Time Literals

Date and time values can be represented in several formats, such as quoted strings or as numbers, depending on the exact type of the value and other factors. For example, in contexts where MySQL expects a date, it interprets any of '2015-07-21', '20150721', and 20150721 as a date.

This section describes the acceptable formats for date and time literals. For more information about the temporal data types, such as the range of permitted values, consult these sections:

- [Section 11.1.2, “Date and Time Type Overview”](#)
- [Section 11.3, “Date and Time Types”](#)

**Standard SQL and ODBC Date and Time Literals.** Standard SQL permits temporal literals to be specified using a type keyword and a string. The space between the keyword and string is optional.

```
DATE 'str'
TIME 'str'
TIMESTAMP 'str'
```

MySQL recognizes those constructions and also the corresponding ODBC syntax:

```
{ d 'str' }
{ t 'str' }
{ ts 'str' }
```

Before MySQL 5.6.4, MySQL ignores the type keyword and each of the preceding constructions produces the string value 'str', with a type of `VARCHAR`.

As of 5.6.4, MySQL uses the type keyword and these constructions produce `DATE`, `TIME`, and `DATETIME` values, respectively, including a trailing fractional seconds part if specified. The `TIMESTAMP` syntax produces a `DATETIME` value in MySQL because `DATETIME` has a range that more closely corresponds to the standard SQL `TIMESTAMP` type, which has a year range from 0001 to 9999. (The MySQL `TIMESTAMP` year range is 1970 to 2038.)

**String and Numeric Literals in Date and Time Context.** MySQL recognizes `DATE` values in these formats:

- As a string in either 'YYYY-MM-DD' or 'YY-MM-DD' format. A “relaxed” syntax is permitted: Any punctuation character may be used as the delimiter between date parts. For example, '2012-12-31', '2012/12/31', '2012^12^31', and '2012@12@31' are equivalent.
- As a string with no delimiters in either 'YYYYMMDD' or 'YYMMDD' format, provided that the string makes sense as a date. For example, '20070523' and '070523' are interpreted as '2007-05-23', but '071332' is illegal (it has nonsensical month and day parts) and becomes '0000-00-00'.
- As a number in either YYYYMMDD or YYMMDD format, provided that the number makes sense as a date. For example, 19830905 and 830905 are interpreted as '1983-09-05'.

MySQL recognizes `DATETIME` and `TIMESTAMP` values in these formats:

- As a string in either 'YYYY-MM-DD HH:MM:SS' or 'YY-MM-DD HH:MM:SS' format. A “relaxed” syntax is permitted here, too: Any punctuation character may be used as the delimiter between

date parts or time parts. For example, `'2012-12-31 11:30:45'`, `'2012^12^31 11+30+45'`, `'2012/12/31 11*30*45'`, and `'2012@12@31 11^30^45'` are equivalent.

- As a string with no delimiters in either `'YYYYMMDDHHMMSS'` or `'YYMMDDHHMMSS'` format, provided that the string makes sense as a date. For example, `'20070523091528'` and `'070523091528'` are interpreted as `'2007-05-23 09:15:28'`, but `'071122129015'` is illegal (it has a nonsensical minute part) and becomes `'0000-00-00 00:00:00'`.
- As a number in either `YYYYMMDDHHMMSS` or `YYMMDDHHMMSS` format, provided that the number makes sense as a date. For example, `19830905132800` and `830905132800` are interpreted as `'1983-09-05 13:28:00'`.

A `DATETIME` or `TIMESTAMP` value can include a trailing fractional seconds part in up to microseconds (6 digits) precision. For information about fractional seconds support in MySQL, see [Section 11.3.6, “Fractional Seconds in Time Values”](#).

Dates containing two-digit year values are ambiguous because the century is unknown. MySQL interprets two-digit year values using these rules:

- Year values in the range `70-99` are converted to `1970-1999`.
- Year values in the range `00-69` are converted to `2000-2069`.

See also [Section 11.3.8, “Two-Digit Years in Dates”](#).

For values specified as strings that include date part delimiters, it is unnecessary to specify two digits for month or day values that are less than 10. `'2015-6-9'` is the same as `'2015-06-09'`. Similarly, for values specified as strings that include time part delimiters, it is unnecessary to specify two digits for hour, minute, or second values that are less than 10. `'2015-10-30 1:2:3'` is the same as `'2015-10-30 01:02:03'`.

Values specified as numbers should be 6, 8, 12, or 14 digits long. If a number is 8 or 14 digits long, it is assumed to be in `YYYYMMDD` or `YYYYMMDDHHMMSS` format and that the year is given by the first 4 digits. If the number is 6 or 12 digits long, it is assumed to be in `YYMMDD` or `YYMMDDHHMMSS` format and that the year is given by the first 2 digits. Numbers that are not one of these lengths are interpreted as though padded with leading zeros to the closest length.

Values specified as nondelimited strings are interpreted according their length. For a string 8 or 14 characters long, the year is assumed to be given by the first 4 characters. Otherwise, the year is assumed to be given by the first 2 characters. The string is interpreted from left to right to find year, month, day, hour, minute, and second values, for as many parts as are present in the string. This means you should not use strings that have fewer than 6 characters. For example, if you specify `'9903'`, thinking that represents March, 1999, MySQL converts it to the “zero” date value. This occurs because the year and month values are `99` and `03`, but the day part is completely missing. However, you can explicitly specify a value of zero to represent missing month or day parts. For example, to insert the value `'1999-03-00'`, use `'990300'`.

MySQL recognizes `TIME` values in these formats:

- As a string in `'D HH:MM:SS'` format. You can also use one of the following “relaxed” syntaxes: `'HH:MM:SS'`, `'HH:MM'`, `'D HH:MM'`, `'D HH'`, or `'SS'`. Here `D` represents days and can have a value from 0 to 34.
- As a string with no delimiters in `'HHMMSS'` format, provided that it makes sense as a time. For example, `'101112'` is understood as `'10:11:12'`, but `'109712'` is illegal (it has a nonsensical minute part) and becomes `'00:00:00'`.
- As a number in `HHMMSS` format, provided that it makes sense as a time. For example, `101112` is understood as `'10:11:12'`. The following alternative formats are also understood: `SS`, `MMSS`, or `HHMMSS`.

A trailing fractional seconds part is recognized in the `'D HH:MM:SS.fraction'`, `'HH:MM:SS.fraction'`, `'HHMMSS.fraction'`, and `HHMMSS.fraction` time formats, where

`fraction` is the fractional part in up to microseconds (6 digits) precision. For information about fractional seconds support in MySQL, see [Section 11.3.6, “Fractional Seconds in Time Values”](#).

For `TIME` values specified as strings that include a time part delimiter, it is unnecessary to specify two digits for hours, minutes, or seconds values that are less than 10. `'8:3:2'` is the same as `'08:03:02'`.

## 9.1.4. Hexadecimal Literals

MySQL supports hexadecimal values, written using `X'val'`, `x'val'`, or `0xval` format, where `val` contains hexadecimal digits (`0..9, A..F`). Lettercase of the digits does not matter. For values written using `X'val'` or `x'val'` format, `val` must contain an even number of digits. For values written using `0xval` syntax, values that contain an odd number of digits are treated as having an extra leading 0. For example, `0x0a` and `0xaaa` are interpreted as `0x0a` and `0x0aaa`.

In numeric contexts, hexadecimal values act like integers (64-bit precision). In string contexts, they act like binary strings, where each pair of hex digits is converted to a character:

```
mysql> SELECT X'4D7953514C';
      -> 'MySQL'
mysql> SELECT 0x0a+0;
      -> 10
mysql> SELECT 0x5061756c;
      -> 'Paul'
```

The default type of a hexadecimal value is a string. If you want to ensure that the value is treated as a number, you can use `CAST(... AS UNSIGNED)` [\[1245\]](#):

```
mysql> SELECT 0x41, CAST(0x41 AS UNSIGNED);
      -> 'A', 65
```

The `X'hexstring'` syntax is based on standard SQL. The `0x` syntax is based on ODBC. Hexadecimal strings are often used by ODBC to supply values for `BLOB` columns.

You can convert a string or a number to a string in hexadecimal format with the `HEX()` [\[1180\]](#) function:

```
mysql> SELECT HEX('cat');
      -> '636174'
mysql> SELECT 0x636174;
      -> 'cat'
```

## 9.1.5. Boolean Literals

The constants `TRUE` and `FALSE` evaluate to 1 and 0, respectively. The constant names can be written in any lettercase.

```
mysql> SELECT TRUE, true, FALSE, false;
      -> 1, 1, 0, 0
```

## 9.1.6. Bit-Field Literals

Bit-field values can be written using `b'value'` or `0bvalue` notation. `value` is a binary value written using zeros and ones.

Bit-field notation is convenient for specifying values to be assigned to `BIT` columns:

```
mysql> CREATE TABLE t (b BIT(8));
mysql> INSERT INTO t SET b = b'11111111';
mysql> INSERT INTO t SET b = b'1010';
mysql> INSERT INTO t SET b = b'0101';
```

Bit values are returned as binary values. To display them in printable form, add 0 or use a conversion function such as `BIN()` [\[1177\]](#). High-order 0 bits are not displayed in the converted value.

```
mysql> SELECT b+0, BIN(b+0), OCT(b+0), HEX(b+0) FROM t;
```

b+0	BIN(b+0)	OCT(b+0)	HEX(b+0)
255	11111111	377	FF
10	1010	12	A
5	101	5	5

Bit values assigned to user variables are treated as binary strings. To assign a bit value as a number to a user variable, use `CAST( ) [1245]` or `+0`:

```
mysql> SET @v1 = 0b1000001;
mysql> SET @v2 = CAST(0b1000001 AS UNSIGNED), @v3 = 0b1000001+0;
mysql> SELECT @v1, @v2, @v3;
+-----+-----+-----+
| @v1 | @v2 | @v3 |
+-----+-----+-----+
| A | 65 | 65 |
+-----+-----+-----+
```

### 9.1.7. NULL Values

The `NULL` value means “no data.” `NULL` can be written in any lettercase. A synonym is `\N` (case sensitive).

For text file import or export operations performed with `LOAD DATA INFILE` or `SELECT ... INTO OUTFILE`, `NULL` is represented by the `\N` sequence. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

Be aware that the `NULL` value is different from values such as `0` for numeric types or the empty string for string types. For more information, see [Section C.5.5.3, “Problems with NULL Values”](#).

## 9.2. Schema Object Names

Certain objects within MySQL, including database, table, index, column, alias, view, stored procedure, partition, tablespace, and other object names are known as identifiers. This section describes the permissible syntax for identifiers in MySQL. [Section 9.2.2, “Identifier Case Sensitivity”](#), describes which types of identifiers are case sensitive and under what conditions.

An identifier may be quoted or unquoted. If an identifier contains special characters or is a reserved word, you *must* quote it whenever you refer to it. (Exception: A reserved word that follows a period in a qualified name must be an identifier, so it need not be quoted.) Reserved words are listed at [Section 9.3, “Reserved Words”](#).

Identifiers are converted to Unicode internally. They may contain these characters:

- Permitted characters in unquoted identifiers:
  - ASCII: [0-9,a-z,A-Z\$\_] (basic Latin letters, digits 0-9, dollar, underscore)
  - Extended: U+0080 .. U+FFFF
- Permitted characters in quoted identifiers include the full Unicode Basic Multilingual Plane (BMP), except U+0000:
  - ASCII: U+0001 .. U+007F
  - Extended: U+0080 .. U+FFFF
- ASCII NUL (U+0000) and supplementary characters (U+10000 and higher) are not permitted in quoted or unquoted identifiers.
- Identifiers may begin with a digit but unless quoted may not consist solely of digits.
- Database, table, and column names cannot end with space characters.

The identifier quote character is the backtick (“`”):

```
mysql> SELECT * FROM `select` WHERE `select`.id > 100;
```

If the [ANSI\\_QUOTES \[603\]](#) SQL mode is enabled, it is also permissible to quote identifiers within double quotation marks:

```
mysql> CREATE TABLE "test" (col INT);
ERROR 1064: You have an error in your SQL syntax...
mysql> SET sql_mode='ANSI_QUOTES';
mysql> CREATE TABLE "test" (col INT);
Query OK, 0 rows affected (0.00 sec)
```

The [ANSI\\_QUOTES \[603\]](#) mode causes the server to interpret double-quoted strings as identifiers. Consequently, when this mode is enabled, string literals must be enclosed within single quotation marks. They cannot be enclosed within double quotation marks. The server SQL mode is controlled as described in [Section 5.1.7, “Server SQL Modes”](#).

Identifier quote characters can be included within an identifier if you quote the identifier. If the character to be included within the identifier is the same as that used to quote the identifier itself, then you need to double the character. The following statement creates a table named `a`b` that contains a column named `c"d`:

```
mysql> CREATE TABLE `a``b` (`c"d` INT);
```

In the select list of a query, a quoted column alias can be specified using identifier or string quoting characters:

```
mysql> SELECT 1 AS `one`, 2 AS 'two';
+-----+-----+
| one | two |
+-----+-----+
| 1 | 2 |
+-----+-----+
```

Elsewhere in the statement, quoted references to the alias must use identifier quoting or the reference is treated as a string literal.

It is recommended that you do not use names that begin with *Me* or *MeN*, where *M* and *N* are integers. For example, avoid using `1e` as an identifier, because an expression such as `1e+3` is ambiguous. Depending on context, it might be interpreted as the expression `1e + 3` or as the number `1e+3`.

Be careful when using `MD5()` [\[1263\]](#) to produce table names because it can produce names in illegal or ambiguous formats such as those just described.

A user variable cannot be used directly in an SQL statement as an identifier or as part of an identifier. See [Section 9.4, “User-Defined Variables”](#), for more information and examples of workarounds.

Special characters in database and table names are encoded in the corresponding file system names as described in [Section 9.2.3, “Mapping of Identifiers to File Names”](#). If you have databases or tables from an older version of MySQL that contain special characters and for which the underlying directory names or file names have not been updated to use the new encoding, the server displays their names with a prefix of `#mysql50#`. For information about referring to such names or converting them to the newer encoding, see that section.

The following table describes the maximum length for each type of identifier.

Identifier	Maximum Length (characters)
Database	64
Table	64
Column	64
Index	64
Constraint	64

Identifier	Maximum Length (characters)
Stored Procedure or Function	64
Trigger	64
View	64
Event	64
Tablespace	64
Server	64
Log File Group	64
Alias	256 (see exception following table)
Compound Statement Label	16

Aliases for column names in `CREATE VIEW` statements are checked against the maximum column length of 64 characters (not the maximum alias length of 256 characters).

Identifiers are stored using Unicode (UTF-8). This applies to identifiers in table definitions that are stored in `.frm` files and to identifiers stored in the grant tables in the `mysql` database. The sizes of the identifier string columns in the grant tables are measured in characters. You can use multi-byte characters without reducing the number of characters permitted for values stored in these columns, something not true prior to MySQL 4.1. As indicated earlier, the permissible Unicode characters are those in the Basic Multilingual Plane (BMP). Supplementary characters are not permitted.

### 9.2.1. Identifier Qualifiers

MySQL permits names that consist of a single identifier or multiple identifiers. The components of a multiple-part name must be separated by period (“.”) characters. The initial parts of a multiple-part name act as qualifiers that affect the context within which the final identifier is interpreted.

In MySQL, you can refer to a table column using any of the following forms.

Column Reference	Meaning
<code>col_name</code>	The column <code>col_name</code> from whichever table used in the statement contains a column of that name.
<code>tbl_name.col_name</code>	The column <code>col_name</code> from table <code>tbl_name</code> of the default database.
<code>db_name.tbl_name.col_name</code>	The column <code>col_name</code> from table <code>tbl_name</code> of the database <code>db_name</code> .

The qualifier character is a separate token and need not be contiguous with the associated identifiers. For example, `tbl_name.col_name` and `tbl_name . col_name` are equivalent.

If any components of a multiple-part name require quoting, quote them individually rather than quoting the name as a whole. For example, write ``my-table`.`my-column``, not ``my-table.my-column``.

A reserved word that follows a period in a qualified name must be an identifier, so in that context it need not be quoted.

You need not specify a `tbl_name` or `db_name.tbl_name` prefix for a column reference in a statement unless the reference would be ambiguous. Suppose that tables `t1` and `t2` each contain a column `c`, and you retrieve `c` in a `SELECT` statement that uses both `t1` and `t2`. In this case, `c` is ambiguous because it is not unique among the tables used in the statement. You must qualify it with a table name as `t1.c` or `t2.c` to indicate which table you mean. Similarly, to retrieve from a table `t` in database `db1` and from a table `t` in database `db2` in the same statement, you must refer to columns in those tables as `db1.t.col_name` and `db2.t.col_name`.

The syntax `.tbl_name` means the table `tbl_name` in the default database. This syntax is accepted for ODBC compatibility because some ODBC programs prefix table names with a “.” character.

## 9.2.2. Identifier Case Sensitivity

In MySQL, databases correspond to directories within the data directory. Each table within a database corresponds to at least one file within the database directory (and possibly more, depending on the storage engine). Triggers also correspond to files. Consequently, the case sensitivity of the underlying operating system plays a part in the case sensitivity of database, table, and trigger names. This means such names are not case sensitive in Windows, but are case sensitive in most varieties of Unix. One notable exception is Mac OS X, which is Unix-based but uses a default file system type (HFS+) that is not case sensitive. However, Mac OS X also supports UFS volumes, which are case sensitive just as on any Unix. See [Section 1.8.4, “MySQL Extensions to Standard SQL”](#). The `lower_case_table_names` [501] system variable also affects how the server handles identifier case sensitivity, as described later in this section.



### Note

Although database, table, and trigger names are not case sensitive on some platforms, you should not refer to one of these using different cases within the same statement. The following statement would not work because it refers to a table both as `my_table` and as `MY_TABLE`:

```
mysql> SELECT * FROM my_table WHERE MY_TABLE.col=1;
```

Column, index, stored routine, and event names are not case sensitive on any platform, nor are column aliases.

However, names of logfile groups are case sensitive. This differs from standard SQL.

By default, table aliases are case sensitive on Unix, but not so on Windows or Mac OS X. The following statement would not work on Unix, because it refers to the alias both as `a` and as `A`:

```
mysql> SELECT col_name FROM tbl_name AS a
-> WHERE a.col_name = 1 OR A.col_name = 2;
```

However, this same statement is permitted on Windows. To avoid problems caused by such differences, it is best to adopt a consistent convention, such as always creating and referring to databases and tables using lowercase names. This convention is recommended for maximum portability and ease of use.

How table and database names are stored on disk and used in MySQL is affected by the `lower_case_table_names` [501] system variable, which you can set when starting `mysqld`. `lower_case_table_names` [501] can take the values shown in the following table. This variable does *not* affect case sensitivity of trigger identifiers. On Unix, the default value of `lower_case_table_names` [501] is 0. On Windows the default value is 1. On Mac OS X, the default value is 2.

Value	Meaning
0	Table and database names are stored on disk using the lettercase specified in the <code>CREATE TABLE</code> or <code>CREATE DATABASE</code> statement. Name comparisons are case sensitive. You should <i>not</i> set this variable to 0 if you are running MySQL on a system that has case-insensitive file names (such as Windows or Mac OS X). If you force this variable to 0 with <code>--lower-case-table-names=0</code> [501] on a case-insensitive file system and access <code>MyISAM</code> tablename using different lettercases, index corruption may result.
1	Table names are stored in lowercase on disk and name comparisons are not case sensitive. MySQL converts all table names to lowercase on storage and lookup. This behavior also applies to database names and table aliases.
2	Table and database names are stored on disk using the lettercase specified in the <code>CREATE TABLE</code> or <code>CREATE DATABASE</code> statement, but MySQL converts them to lowercase on lookup. Name comparisons are not case sensitive. This works <i>only</i> on file systems that are not case sensitive! <code>InnoDB</code> table names are stored in lowercase, as for <code>lower_case_table_names=1</code> .



If you are using MySQL on only one platform, you do not normally have to change the `lower_case_table_names` [501] variable from its default value. However, you may encounter difficulties if you want to transfer tables between platforms that differ in file system case sensitivity. For example, on Unix, you can have two different tables named `my_table` and `MY_TABLE`, but on Windows these two names are considered identical. To avoid data transfer problems arising from lettercase of database or table names, you have two options:

- Use `lower_case_table_names=1` on all systems. The main disadvantage with this is that when you use `SHOW TABLES` or `SHOW DATABASES`, you do not see the names in their original lettercase.
- Use `lower_case_table_names=0` on Unix and `lower_case_table_names=2` on Windows. This preserves the lettercase of database and table names. The disadvantage of this is that you must ensure that your statements always refer to your database and table names with the correct lettercase on Windows. If you transfer your statements to Unix, where lettercase is significant, they do not work if the lettercase is incorrect.

**Exception:** If you are using InnoDB tables and you are trying to avoid these data transfer problems, you should set `lower_case_table_names` [501] to 1 on all platforms to force names to be converted to lowercase.

If you plan to set the `lower_case_table_names` [501] system variable to 1 on Unix, you must first convert your old database and table names to lowercase before stopping `mysqld` and restarting it with the new variable setting. To do this for an individual table, use `RENAME TABLE`:

```
RENAME TABLE T1 TO t1;
```

To convert one or more entire databases, dump them before setting `lower_case_table_names` [501], then drop the databases, and reload them after setting `lower_case_table_names` [501]:

1. Use `mysqldump` to dump each database:

```
mysqldump --databases db1 > db1.sql
mysqldump --databases db2 > db2.sql
...
```

Do this for each database that must be recreated.

2. Use `DROP DATABASE` to drop each database.
3. Stop the server, set `lower_case_table_names` [501], and restart the server.
4. Reload the dump file for each database. Because `lower_case_table_names` [501] is set, each database and table name will be converted to lowercase as it is recreated:

```
mysql < db1.sql
mysql < db2.sql
...
```

Object names may be considered duplicates if their uppercase forms are equal according to a binary collation. That is true for names of cursors, conditions, procedures, functions, savepoints, stored routine parameters, stored program local variables, and plugins. It is not true for names of columns, constraints, databases, partitions, statements prepared with `PREPARE`, tables, triggers, users, and user-defined variables.

File system case sensitivity can affect searches in string columns of `INFORMATION_SCHEMA` tables. For more information, see [Section 10.1.7.9, “Collation and INFORMATION\\_SCHEMA Searches”](#).

### 9.2.3. Mapping of Identifiers to File Names

There is a correspondence between database and table identifiers and names in the file system. For the basic structure, MySQL represents each database as a directory in the data directory, and each table by one or more files in the appropriate database directory. For the table format files (`.FRM`), the data is always stored in this structure and location.

For the data and index files, the exact representation on disk is storage engine specific. These files may be stored in the same location as the [FRM](#) files, or the information may be stored in a separate file. [InnoDB](#) data is stored in the InnoDB data files. If you are using tablespaces with [InnoDB](#), then the specific tablespace files you create are used instead.

Any character is legal in database or table identifiers except ASCII NUL (0x00). MySQL encodes any characters that are problematic in the corresponding file system objects when it creates database directories or table files:

- Basic Latin letters ([a..zA..Z](#)), digits ([0..9](#)) and underscore ([\\_](#)) are encoded as is. Consequently, their case sensitivity directly depends on file system features.
- All other national letters from alphabets that have uppercase/lowercase mapping are encoded as shown in the following table. Values in the Code Range column are UCS-2 values.

Code Range	Pattern	Number	Used	Unused	Blocks
00C0..017F	[@][0..4][g..z]	5*20= 100	97	3	Latin-1 Supplement + Latin Extended-A
0370..03FF	[@][5..9][g..z]	5*20= 100	88	12	Greek and Coptic
0400..052F	[@][g..z][0..6]	20*7= 140	137	3	Cyrillic + Cyrillic Supplement
0530..058F	[@][g..z][7..8]	20*2= 40	38	2	Armenian
2160..217F	[@][g..z][9]	20*1= 20	16	4	Number Forms
0180..02AF	[@][g..z][a..k]	20*11=220	203	17	Latin Extended-B + IPA Extensions
1E00..1EFF	[@][g..z][l..r]	20*7= 140	136	4	Latin Extended Additional
1F00..1FFF	[@][g..z][s..z]	20*8= 160	144	16	Greek Extended
.... ....	[@][a..f][g..z]	6*20= 120	0	120	RESERVED
24B6..24E9	[@][@][a..z]	26	26	0	Enclosed Alphanumerics
FF21..FF5A	[@][a..z][@]	26	26	0	Halfwidth and Fullwidth forms

One of the bytes in the sequence encodes lettercase. For example: [LATIN CAPITAL LETTER A WITH GRAVE](#) is encoded as @0G, whereas [LATIN SMALL LETTER A WITH GRAVE](#) is encoded as @0g. Here the third byte (G or g) indicates lettercase. (On a case-insensitive file system, both letters will be treated as the same.)

For some blocks, such as Cyrillic, the second byte determines lettercase. For other blocks, such as Latin1 Supplement, the third byte determines lettercase. If two bytes in the sequence are letters (as in Greek Extended), the leftmost letter character stands for lettercase. All other letter bytes must be in lowercase.

- All nonletter characters except underscore ([\\_](#)), as well as letters from alphabets that do not have uppercase/lowercase mapping (such as Hebrew) are encoded using hexadecimal representation using lowercase letters for hex digits [a..f](#):

```
0x003F -> @003f
0xFFFF -> @ffff
```

The hexadecimal values correspond to character values in the [ucs2](#) double-byte character set.

On Windows, some names such as [nul](#), [prn](#), and [aux](#) are encoded by appending @@@ to the name when the server creates the corresponding file or directory. This occurs on all platforms for portability of the corresponding database object between platforms.

If you have databases or tables from a version of MySQL older than 5.1.6 that contain special characters and for which the underlying directory names or file names have not been updated to use the new encoding, the server displays their names with a prefix of `#mysql50#` in the output from `INFORMATION_SCHEMA` tables or `SHOW` statements. For example, if you have a table named `a@b` and its name encoding has not been updated, `SHOW TABLES` displays it like this:

```
mysql> SHOW TABLES;
+-----+
| Tables_in_test |
+-----+
| #mysql50#a@b   |
+-----+
```

To refer to such a name for which the encoding has not been updated, you must supply the `#mysql50#` prefix:

```
mysql> SHOW COLUMNS FROM `a@b`;
ERROR 1146 (42S02): Table 'test.a@b' doesn't exist

mysql> SHOW COLUMNS FROM `#mysql50#a@b`;
+-----+-----+-----+-----+-----+-----+
| Field | Type   | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| i     | int(11)| YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
```

To update old names to eliminate the need to use the special prefix to refer to them, re-encode them with `mysqlcheck`. The following command updates all names to the new encoding:

```
shell> mysqlcheck --check-upgrade --fix-db-names --fix-table-names --all-databases
```

To check only specific databases or tables, omit `--all-databases` [276] and provide the appropriate database or table arguments. For information about `mysqlcheck` invocation syntax, see [Section 4.5.3](#), “`mysqlcheck` — A Table Maintenance Program”.



#### Note

The `#mysql50#` prefix is intended only to be used internally by the server. You should not create databases or tables with names that use this prefix.

Also, `mysqlcheck` cannot fix names that contain literal instances of the `@` character that is used for encoding special characters. If you have databases or tables that contain this character, use `mysqldump` to dump them before upgrading to MySQL 5.1.6 or later, and then reload the dump file after upgrading.

## 9.2.4. Function Name Parsing and Resolution

MySQL 5.6 supports built-in (native) functions, user-defined functions (UDFs), and stored functions. This section describes how the server recognizes whether the name of a built-in function is used as a function call or as an identifier, and how the server determines which function to use in cases when functions of different types exist with a given name.

### Built-In Function Name Parsing

The parser uses default rules for parsing names of built-in functions. These rules can be changed by enabling the `IGNORE_SPACE` [604] SQL mode.

When the parser encounters a word that is the name of a built-in function, it must determine whether the name signifies a function call or is instead a nonexpression reference to an identifier such as a table or column name. For example, in the following statements, the first reference to `count` is a function call, whereas the second reference is a table name:

```
SELECT COUNT(*) FROM mytable;
CREATE TABLE count (i INT);
```

The parser should recognize the name of a built-in function as indicating a function call only when parsing what is expected to be an expression. That is, in nonexpression context, function names are permitted as identifiers.

However, some built-in functions have special parsing or implementation considerations, so the parser uses the following rules by default to distinguish whether their names are being used as function calls or as identifiers in nonexpression context:

- To use the name as a function call in an expression, there must be no whitespace between the name and the following “(” parenthesis character.
- Conversely, to use the function name as an identifier, it must not be followed immediately by a parenthesis.

The requirement that function calls be written with no whitespace between the name and the parenthesis applies only to the built-in functions that have special considerations. `COUNT` is one such name. The exact list of function names for which following whitespace determines their interpretation are those listed in the `sql_functions[]` array of the `sql/lex.h` source file. Before MySQL 5.1, these are rather numerous (about 200), so you may find it easiest to treat the no-whitespace requirement as applying to all function calls. In MySQL 5.1 and later, parser improvements reduce to about 30 the number of affected function names.

For functions not listed in the `sql_functions[]` array, whitespace does not matter. They are interpreted as function calls only when used in expression context and may be used freely as identifiers otherwise. `ASCII` is one such name. However, for these nonaffected function names, interpretation may vary in expression context: `func_name ()` is interpreted as a built-in function if there is one with the given name; if not, `func_name ()` is interpreted as a user-defined function or stored function if one exists with that name.

The `IGNORE_SPACE` [604] SQL mode can be used to modify how the parser treats function names that are whitespace-sensitive:

- With `IGNORE_SPACE` [604] disabled, the parser interprets the name as a function call when there is no whitespace between the name and the following parenthesis. This occurs even when the function name is used in nonexpression context:

```
mysql> CREATE TABLE count(i INT);
ERROR 1064 (42000): You have an error in your SQL syntax ...
near 'count(i INT)'
```

To eliminate the error and cause the name to be treated as an identifier, either use whitespace following the name or write it as a quoted identifier (or both):

```
CREATE TABLE count (i INT);
CREATE TABLE `count`(i INT);
CREATE TABLE `count` (i INT);
```

- With `IGNORE_SPACE` [604] enabled, the parser loosens the requirement that there be no whitespace between the function name and the following parenthesis. This provides more flexibility in writing function calls. For example, either of the following function calls are legal:

```
SELECT COUNT(*) FROM mytable;
SELECT COUNT (*) FROM mytable;
```

However, enabling `IGNORE_SPACE` [604] also has the side effect that the parser treats the affected function names as reserved words (see Section 9.3, “Reserved Words”). This means that a space following the name no longer signifies its use as an identifier. The name can be used in function calls with or without following whitespace, but causes a syntax error in nonexpression context unless it is quoted. For example, with `IGNORE_SPACE` [604] enabled, both of the following statements fail with a syntax error because the parser interprets `count` as a reserved word:

```
CREATE TABLE count(i INT);
CREATE TABLE count (i INT);
```

To use the function name in nonexpression context, write it as a quoted identifier:

```
CREATE TABLE `count`(i INT);
CREATE TABLE `count` (i INT);
```

To enable the `IGNORE_SPACE` [604] SQL mode, use this statement:

```
SET sql_mode = 'IGNORE_SPACE';
```

`IGNORE_SPACE` [604] is also enabled by certain other composite modes such as `ANSI` [608] that include it in their value:

```
SET sql_mode = 'ANSI';
```

Check [Section 5.1.7, “Server SQL Modes”](#), to see which composite modes enable `IGNORE_SPACE` [604].

To minimize the dependency of SQL code on the `IGNORE_SPACE` [604] setting, use these guidelines:

- Avoid creating UDFs or stored functions that have the same name as a built-in function.
- Avoid using function names in nonexpression context. For example, these statements use `count` (one of the affected function names affected by `IGNORE_SPACE` [604]), so they fail with or without whitespace following the name if `IGNORE_SPACE` [604] is enabled:

```
CREATE TABLE count(i INT);
CREATE TABLE count (i INT);
```

If you must use a function name in nonexpression context, write it as a quoted identifier:

```
CREATE TABLE `count`(i INT);
CREATE TABLE `count` (i INT);
```

The number of function names affected by `IGNORE_SPACE` [604] was reduced significantly in MySQL 5.1.13, from about 200 to about 30. As of MySQL 5.1.13, only the following functions are still affected by the `IGNORE_SPACE` [604] setting.

ADDDATE	BIT_AND	BIT_OR	BIT_XOR
CAST	COUNT	CURDATE	CURTIME
DATE_ADD	DATE_SUB	EXTRACT	GROUP_CONCAT
MAX	MID	MIN	NOW
POSITION	SESSION_USER	STD	STDDEV
STDDEV_POP	STDDEV_SAMP	SUBDATE	SUBSTR
SUBSTRING	SUM	SYSDATE	SYSTEM_USER
TRIM	VARIANCE	VAR_POP	VAR_SAMP

For earlier versions of MySQL, check the contents of the `sql_functions[]` array in the `sql/lex.h` source file to see which functions are affected by `IGNORE_SPACE` [604].

**Incompatibility warning:** The change in MySQL 5.1.13 that reduces the number of function names affected by `IGNORE_SPACE` [604] improves the consistency of parser operation. However, it also introduces the possibility of incompatibility for old SQL code that relies on the following conditions:

- `IGNORE_SPACE` [604] is disabled.
- The presence or absence of whitespace following a function name is used to distinguish between a built-in function and stored function that have the same name, such as `PI()` [1205] versus `PI ( )`.

For functions that are no longer affected by `IGNORE_SPACE` [604] as of MySQL 5.1.13, that strategy no longer works. Either of the following approaches can be used if you have code that is subject to the preceding incompatibility:

- If a stored function has a name that conflicts with a built-in function, refer to the stored function with a schema name qualifier, regardless of whether whitespace is present. For example, write `schema_name.PI()` or `schema_name.PI ()`.
- Alternatively, rename the stored function to use a nonconflicting name and change invocations of the function to use the new name.

### Function Name Resolution

The following rules describe how the server resolves references to function names for function creation and invocation:

- Built-in functions and user-defined functions

An error occurs if you try to create a UDF with the same name as a built-in function.

- Built-in functions and stored functions

It is possible to create a stored function with the same name as a built-in function, but to invoke the stored function it is necessary to qualify it with a schema name. For example, if you create a stored function named `PI` in the `test` schema, you invoke it as `test.PI()` because the server resolves `PI()` [1205] as a reference to the built-in function. The server creates a warning if the stored function name collides with a built-in function name. The warning can be displayed with `SHOW WARNINGS`.

- User-defined functions and stored functions

User-defined functions and stored functions share the same namespace, so you cannot create a UDF and a stored function with the same name.

The preceding function name resolution rules have implications for upgrading to versions of MySQL that implement new built-in functions:

- If you have already created a user-defined function with a given name and upgrade MySQL to a version that implements a new built-in function with the same name, the UDF becomes inaccessible. To correct this, use `DROP FUNCTION` to drop the UDF, and then use `CREATE FUNCTION` to re-create the UDF with a different nonconflicting name.
- If a new version of MySQL implements a built-in function with the same name as an existing stored function, you have two choices: Rename the stored function to use a nonconflicting name, or change calls to the function so that they use a schema qualifier (that is, use `schema_name.func_name()` syntax).

## 9.3. Reserved Words

Certain words such as `SELECT`, `DELETE`, or `BIGINT` are reserved and require special treatment for use as identifiers such as table and column names. This may also be true for the names of built-in functions.

Reserved words are permitted as identifiers if you quote them as described in [Section 9.2, “Schema Object Names”](#):

```
mysql> CREATE TABLE interval (begin INT, end INT);
ERROR 1064 (42000): You have an error in your SQL syntax ...
near 'interval (begin INT, end INT)'
```

```
mysql> CREATE TABLE `interval` (begin INT, end INT);
Query OK, 0 rows affected (0.01 sec)
```

Exception: A word that follows a period in a qualified name must be an identifier, so it need not be quoted even if it is reserved:

```
mysql> CREATE TABLE mydb.interval (begin INT, end INT);
Query OK, 0 rows affected (0.01 sec)
```

Names of built-in functions are permitted as identifiers but may require care to be used as such. For example, `COUNT` is acceptable as a column name. However, by default, no whitespace is permitted in function invocations between the function name and the following “(” character. This requirement enables the parser to distinguish whether the name is used in a function call or in nonfunction context. For further detail on recognition of function names, see [Section 9.2.4, “Function Name Parsing and Resolution”](#).

The words in the following table are explicitly reserved in MySQL 5.6. In addition, `_FILENAME` is reserved. At some point, you might upgrade to a higher version, so it is a good idea to have a look at future reserved words, too. You can find these in the manuals that cover higher versions of MySQL. Most of the words in the table are forbidden by standard SQL as column or table names (for example, `GROUP`). A few are reserved because MySQL needs them and uses a `yacc` parser. A reserved word can be used as an identifier if you quote it.

For a more detailed list of reserved words, including differences between versions, see [Reserved Words in MySQL 5.6](#).

**Table 9.2. Reserved Words in MySQL 5.6.11**

ACCESSIBLE	ADD	ALL
ALTER	ANALYZE	AND
AS	ASC	ASENSITIVE
BEFORE	BETWEEN	BIGINT
BINARY	BLOB	BOTH
BY	CALL	CASCADE
CASE	CHANGE	CHAR
CHARACTER	CHECK	COLLATE
COLUMN	CONDITION	CONSTRAINT
CONTINUE	CONVERT	CREATE
CROSS	CURRENT_DATE	CURRENT_TIME
CURRENT_TIMESTAMP	CURRENT_USER	CURSOR
DATABASE	DATABASES	DAY_HOUR
DAY_MICROSECOND	DAY_MINUTE	DAY_SECOND
DEC	DECIMAL	DECLARE
DEFAULT	DELAYED	DELETE
DESC	DESCRIBE	DETERMINISTIC
DISTINCT	DISTINCTROW	DIV
DOUBLE	DROP	DUAL
EACH	ELSE	ELSEIF
ENCLOSED	ESCAPED	EXISTS
EXIT	EXPLAIN	FALSE
FETCH	FLOAT	FLOAT4
FLOAT8	FOR	FORCE
FOREIGN	FROM	FULLTEXT
GET	GRANT	GROUP
HAVING	HIGH_PRIORITY	HOURL_MICROSECOND
HOURL_MINUTE	HOURL_SECOND	IF
IGNORE	IN	INDEX



Reserved Words

INFILE	INNER	INOUT
INSENSITIVE	INSERT	INT
INT1	INT2	INT3
INT4	INT8	INTEGER
INTERVAL	INTO	IO_AFTER_GTIDS
IO_BEFORE_GTIDS	IS	ITERATE
JOIN	KEY	KEYS
KILL	LEADING	LEAVE
LEFT	LIKE	LIMIT
LINEAR	LINES	LOAD
LOCALTIME	LOCALTIMESTAMP	LOCK
LONG	LONGBLOB	LONGTEXT
LOOP	LOW_PRIORITY	MASTER_BIND
MASTER_SSL_VERIFY_SERVER_CERT	MATCH	MAXVALUE
MEDIUMBLOB	MEDIUMINT	MEDIUMTEXT
MIDDLEINT	MINUTE_MICROSECOND	MINUTE_SECOND
MOD	MODIFIES	NATURAL
NOT	NO_WRITE_TO_BINLOG	NULL
NUMERIC	ON	OPTIMIZE
OPTION	OPTIONALLY	OR
ORDER	OUT	OUTER
OUTFILE	PARTITION	PRECISION
PRIMARY	PROCEDURE	PURGE
RANGE	READ	READS
READ_WRITE	REAL	REFERENCES
REGEXP	RELEASE	RENAME
REPEAT	REPLACE	REQUIRE
RESIGNAL	RESTRICT	RETURN
REVOKE	RIGHT	RLIKE
SCHEMA	SCHEMAS	SECOND_MICROSECOND
SELECT	SENSITIVE	SEPARATOR
SET	SHOW	SIGNAL
SMALLINT	SPATIAL	SPECIFIC
SQL	SQLEXCEPTION	SQLSTATE
SQLWARNING	SQL_BIG_RESULT	SQL_CALC_FOUND_ROWS
SQL_SMALL_RESULT	SSL	STARTING
STRAIGHT_JOIN	TABLE	TERMINATED
THEN	TINYBLOB	TINYINT
TINYTEXT	TO	TRAILING
TRIGGER	TRUE	UNDO
UNION	UNIQUE	UNLOCK
UNSIGNED	UPDATE	USAGE

USE	USING	UTC_DATE
UTC_TIME	UTC_TIMESTAMP	VALUES
VARBINARY	VARCHAR	VARCHARACTER
VARYING	WHEN	WHERE
WHILE	WITH	WRITE
XOR	YEAR_MONTH	ZEROFILL

**Table 9.3. New Reserved Words in MySQL 5.6**

GET	IO_AFTER_GTIDS	IO_BEFORE_GTIDS
MASTER_BIND	ONE_SHOT	PARTITION
SQL_AFTER_GTIDS	SQL_BEFORE_GTIDS	

MySQL permits some keywords to be used as unquoted identifiers because many people previously used them. Examples are those in the following list:

- ACTION
- BIT
- DATE
- ENUM
- NO
- TEXT
- TIME
- TIMESTAMP

## 9.4. User-Defined Variables

You can store a value in a user-defined variable in one statement and then refer to it later in another statement. This enables you to pass values from one statement to another. *User-defined variables are session-specific.* That is, a user variable defined by one client cannot be seen or used by other clients. All variables for a given client session are automatically freed when that client exits.

User variables are written as `@var_name`, where the variable name `var_name` consists of alphanumeric characters, “.”, “\_”, and “\$”. A user variable name can contain other characters if you quote it as a string or identifier (for example, `@'my-var'`, `@"my-var"`, or `@`my-var``).

User variable names are not case sensitive in MySQL 5.0 and up.

One way to set a user-defined variable is by issuing a `SET` statement:

```
SET @var_name = expr [, @var_name = expr] ...
```

For `SET`, either `= [1173]` or `:= [1172]` can be used as the assignment operator.

You can also assign a value to a user variable in statements other than `SET`. In this case, the assignment operator must be `:= [1172]` and not `= [1173]` because the latter is treated as the comparison operator `= [1166]` in non-`SET` statements:

```
mysql> SET @t1=1, @t2=2, @t3:=4;
mysql> SELECT @t1, @t2, @t3, @t4 := @t1+@t2+@t3;
+-----+-----+-----+-----+

```

```

| @t1 | @t2 | @t3 | @t4 := @t1+@t2+@t3 |
+-----+-----+-----+-----+
| 1 | 2 | 4 | 7 |
+-----+-----+-----+

```

User variables can be assigned a value from a limited set of data types: integer, decimal, floating-point, binary or nonbinary string, or `NULL` value. Assignment of decimal and real values does not preserve the precision or scale of the value. A value of a type other than one of the permissible types is converted to a permissible type. For example, a value having a temporal or spatial data type is converted to a binary string.

If a user variable is assigned a nonbinary (character) string value, it has the same character set and collation as the string. The coercibility of user variables is implicit. (This is the same coercibility as for table column values.)

Bit values assigned to user variables are treated as binary strings. To assign a bit value as a number to a user variable, use `CAST() [1245]` or `+0`:

```

mysql> SET @v1 = b'1000001';
mysql> SET @v2 = CAST(b'1000001' AS UNSIGNED), @v3 = b'1000001'+0;
mysql> SELECT @v1, @v2, @v3;
+-----+-----+-----+
| @v1 | @v2 | @v3 |
+-----+-----+-----+
| A | 65 | 65 |
+-----+-----+-----+

```

If the value of a user variable is selected in a result set, it is returned to the client as a string.

If you refer to a variable that has not been initialized, it has a value of `NULL` and a type of string.

User variables may be used in most contexts where expressions are permitted. This does not currently include contexts that explicitly require a literal value, such as in the `LIMIT` clause of a `SELECT` statement, or the `IGNORE N LINES` clause of a `LOAD DATA` statement.

As a general rule, you should never assign a value to a user variable and read the value within the same statement. You might get the results you expect, but this is not guaranteed. The order of evaluation for expressions involving user variables is undefined and may change based on the elements contained within a given statement; in addition, this order is not guaranteed to be the same between releases of the MySQL Server. In `SELECT @a, @a:=@a+1, ...`, you might think that MySQL will evaluate `@a` first and then do an assignment second. However, changing the statement (for example, by adding a `GROUP BY`, `HAVING`, or `ORDER BY` clause) may cause MySQL to select an execution plan with a different order of evaluation.

Another issue with assigning a value to a variable and reading the value within the same statement is that the default result type of a variable is based on its type at the start of the statement. The following example illustrates this:

```

mysql> SET @a='test';
mysql> SELECT @a,(@a:=20) FROM tbl_name;

```

For this `SELECT` statement, MySQL reports to the client that column one is a string and converts all accesses of `@a` to strings, even though `@a` is set to a number for the second row. After the `SELECT` statement executes, `@a` is regarded as a number for the next statement.

To avoid problems with this behavior, either do not assign a value to and read the value of the same variable within a single statement, or else set the variable to `0`, `0.0`, or `' '` to define its type before you use it.

In a `SELECT` statement, each select expression is evaluated only when sent to the client. This means that in a `HAVING`, `GROUP BY`, or `ORDER BY` clause, referring to a variable that is assigned a value in the select expression list does *not* work as expected:

```
mysql> SELECT (@aa:=id) AS a, (@aa+3) AS b FROM tbl_name HAVING b=5;
```

The reference to `b` in the `HAVING` clause refers to an alias for an expression in the select list that uses `@aa`. This does not work as expected: `@aa` contains the value of `id` from the previous selected row, not from the current row.

User variables are intended to provide data values. They cannot be used directly in an SQL statement as an identifier or as part of an identifier, such as in contexts where a table or database name is expected, or as a reserved word such as `SELECT`. This is true even if the variable is quoted, as shown in the following example:

```
mysql> SELECT c1 FROM t;
+-----+
| c1 |
+-----+
| 0 |
+-----+
| 1 |
+-----+
2 rows in set (0.00 sec)

mysql> SET @col = "c1";
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @col FROM t;
+-----+
| @col |
+-----+
| c1 |
+-----+
1 row in set (0.00 sec)

mysql> SELECT `@col` FROM t;
ERROR 1054 (42S22): Unknown column '@col' in 'field list'

mysql> SET @col = "`c1`";
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @col FROM t;
+-----+
| @col |
+-----+
| `c1` |
+-----+
1 row in set (0.00 sec)
```

An exception to this principle that user variables cannot be used to provide identifiers is that if you are constructing a string for use as a prepared statement to be executed later. In this case, user variables can be used to provide any part of the statement. The following example illustrates how this can be done:

```
mysql> SET @c = "c1";
Query OK, 0 rows affected (0.00 sec)

mysql> SET @s = CONCAT("SELECT ", @c, " FROM t");
Query OK, 0 rows affected (0.00 sec)

mysql> PREPARE stmt FROM @s;
Query OK, 0 rows affected (0.04 sec)
Statement prepared

mysql> EXECUTE stmt;
+-----+
| c1 |
+-----+
| 0 |
+-----+
| 1 |
+-----+
2 rows in set (0.00 sec)
```

```
mysql> DEALLOCATE PREPARE stmt;
Query OK, 0 rows affected (0.00 sec)
```

See [Section 13.5, “SQL Syntax for Prepared Statements”](#), for more information.

A similar technique can be used in application programs to construct SQL statements using program variables, as shown here using PHP 5:

```
<?php
    $mysqli = new mysqli("localhost", "user", "pass", "test");

    if( mysqli_connect_errno() )
        die("Connection failed: %s\n", mysqli_connect_error());

    $col = "c1";

    $query = "SELECT $col FROM t";

    $result = $mysqli->query($query);

    while($row = $result->fetch_assoc())
    {
        echo "<p>" . $row["$col"] . "</p>\n";
    }

    $result->close();

    $mysqli->close();
?>
```

Assembling an SQL statement in this fashion is sometimes known as “Dynamic SQL”.

## 9.5. Expression Syntax

The following rules define expression syntax in MySQL. The grammar shown here is based on that given in the `sql/sql_yacc.yy` file of MySQL source distributions. See the notes after the grammar for additional information about some of the terms. Operator precedence is given in [Section 12.3.1, “Operator Precedence”](#).

```
expr:
    expr OR expr
  | expr || expr
  | expr XOR expr
  | expr AND expr
  | expr && expr
  | NOT expr
  | ! expr
  | boolean_primary IS [NOT] {TRUE | FALSE | UNKNOWN}
  | boolean_primary

boolean_primary:
    boolean_primary IS [NOT] NULL
  | boolean_primary <=> predicate
  | boolean_primary comparison_operator predicate
  | boolean_primary comparison_operator {ALL | ANY} (subquery)
  | predicate

comparison_operator: = | >= | > | <= | < | <> | !=

predicate:
    bit_expr [NOT] IN (subquery)
  | bit_expr [NOT] IN (expr [, expr] ...)
  | bit_expr [NOT] BETWEEN bit_expr AND predicate
  | bit_expr SOUNDS LIKE bit_expr
  | bit_expr [NOT] LIKE simple_expr [ESCAPE simple_expr]
  | bit_expr [NOT] REGEXP bit_expr
  | bit_expr
```

```

bit_expr:
  bit_expr | bit_expr
  | bit_expr & bit_expr
  | bit_expr << bit_expr
  | bit_expr >> bit_expr
  | bit_expr + bit_expr
  | bit_expr - bit_expr
  | bit_expr * bit_expr
  | bit_expr / bit_expr
  | bit_expr DIV bit_expr
  | bit_expr MOD bit_expr
  | bit_expr % bit_expr
  | bit_expr ^ bit_expr
  | bit_expr + interval_expr
  | bit_expr - interval_expr
  | simple_expr

simple_expr:
  literal
  | identifier
  | function_call
  | simple_expr COLLATE collation_name
  | param_marker
  | variable
  | simple_expr || simple_expr
  | + simple_expr
  | - simple_expr
  | ~ simple_expr
  | ! simple_expr
  | BINARY simple_expr
  | (expr [, expr] ...)
  | ROW (expr, expr [, expr] ...)
  | (subquery)
  | EXISTS (subquery)
  | {identifier expr}
  | match_expr
  | case_expr
  | interval_expr

```

**Notes:**

For literal value syntax, see [Section 9.1, “Literal Values”](#).

For identifier syntax, see [Section 9.2, “Schema Object Names”](#).

Variables can be user variables, system variables, or stored program local variables or parameters:

- User variables: [Section 9.4, “User-Defined Variables”](#)
- System variables: [Section 5.1.5, “Using System Variables”](#)
- Local variables: [Section 13.6.4.1, “Local Variable DECLARE Syntax”](#)
- Parameters: [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#)

*param\_marker* is '?' as used in prepared statements for placeholders. See [Section 13.5.1, “PREPARE Syntax”](#).

(*subquery*) indicates a subquery that returns a single value; that is, a scalar subquery. See [Section 13.2.10.1, “The Subquery as Scalar Operand”](#).

{*identifier expr*} is ODBC escape syntax and is accepted for ODBC compatibility. The value is *expr*. The curly braces in the syntax should be written literally; they are not metasyntax as used elsewhere in syntax descriptions.

*match\_expr* indicates a `MATCH [1229]` expression. See [Section 12.9, “Full-Text Search Functions”](#).

*case\_expr* indicates a `CASE [1174]` expression. See [Section 12.4, “Control Flow Functions”](#).

*interval\_expr* represents a time interval. The syntax is `INTERVAL expr unit`, where *unit* is a specifier such as `HOUR`, `DAY`, or `WEEK`. For the full list of *unit* specifiers, see the description of the `DATE_ADD()` [1213] function in Section 12.7, “Date and Time Functions”.

The meaning of some operators depends on the SQL mode:

- By default, `||` [1172] is a logical `OR` [1172] operator. With `PIPES_AS_CONCAT` [607] enabled, `||` [1172] is string concatenation, with a precedence between `^` [1258] and the unary operators.
- By default, `!` [1171] has a higher precedence than `NOT`. With `HIGH_NOT_PRECEDENCE` [603] enabled, `!` [1171] and `NOT` have the same precedence.

See Section 5.1.7, “Server SQL Modes”.

## 9.6. Comment Syntax

MySQL Server supports three comment styles:

- From a “#” character to the end of the line.
- From a “-- ” sequence to the end of the line. In MySQL, the “-- ” (double-dash) comment style requires the second dash to be followed by at least one whitespace or control character (such as a space, tab, newline, and so on). This syntax differs slightly from standard SQL comment syntax, as discussed in Section 1.8.5.5, “‘--’ as the Start of a Comment”.
- From a `/*` sequence to the following `*/` sequence, as in the C programming language. This syntax enables a comment to extend over multiple lines because the beginning and closing sequences need not be on the same line.

The following example demonstrates all three comment styles:

```
mysql> SELECT 1+1;      # This comment continues to the end of line
mysql> SELECT 1+1;      -- This comment continues to the end of line
mysql> SELECT 1 /* this is an in-line comment */ + 1;
mysql> SELECT 1+
/*
this is a
multiple-line comment
*/
1;
```

Nested comments are not supported. (Under some conditions, nested comments might be permitted, but usually are not, and users should avoid them.)

MySQL Server supports some variants of C-style comments. These enable you to write code that includes MySQL extensions, but is still portable, by using comments of the following form:

```
/*! MySQL-specific code */
```

In this case, MySQL Server parses and executes the code within the comment as it would any other SQL statement, but other SQL servers will ignore the extensions. For example, MySQL Server recognizes the `STRAIGHT_JOIN` keyword in the following statement, but other servers will not:

```
SELECT /*! STRAIGHT_JOIN */ col1 FROM table1,table2 WHERE ...
```

If you add a version number after the “!” character, the syntax within the comment is executed only if the MySQL version is greater than or equal to the specified version number. The `TEMPORARY` keyword in the following comment is executed only by servers from MySQL 3.23.02 or higher:

```
CREATE /*!32302 TEMPORARY */ TABLE t (a INT);
```

The comment syntax just described applies to how the `mysqld` server parses SQL statements. The `mysql` client program also performs some parsing of statements before sending them to the server. (It does this to determine statement boundaries within a multiple-statement input line.)



Comments in this format, `/*!12345 ... */`, are not stored on the server. If this format is used to comment stored routines, the comments will not be retained on the server.

The use of short-form `mysql` commands such as `\C` within multi-line `/* ... */` comments is not supported.



---

# Chapter 10. Globalization

## Table of Contents

10.1. Character Set Support .....	1037
10.1.1. Character Sets and Collations in General .....	1038
10.1.2. Character Sets and Collations in MySQL .....	1039
10.1.3. Specifying Character Sets and Collations .....	1040
10.1.4. Connection Character Sets and Collations .....	1046
10.1.5. Configuring the Character Set and Collation for Applications .....	1049
10.1.6. Character Set for Error Messages .....	1050
10.1.7. Collation Issues .....	1051
10.1.8. String Repertoire .....	1060
10.1.9. Operations Affected by Character Set Support .....	1061
10.1.10. Unicode Support .....	1064
10.1.11. Upgrading from Previous to Current Unicode Support .....	1068
10.1.12. UTF-8 for Metadata .....	1070
10.1.13. Column Character Set Conversion .....	1071
10.1.14. Character Sets and Collations That MySQL Supports .....	1072
10.2. Setting the Error Message Language .....	1085
10.3. Adding a Character Set .....	1086
10.3.1. Character Definition Arrays .....	1088
10.3.2. String Collating Support for Complex Character Sets .....	1089
10.3.3. Multi-Byte Character Support for Complex Character Sets .....	1089
10.4. Adding a Collation to a Character Set .....	1089
10.4.1. Collation Implementation Types .....	1090
10.4.2. Choosing a Collation ID .....	1093
10.4.3. Adding a Simple Collation to an 8-Bit Character Set .....	1094
10.4.4. Adding a UCA Collation to a Unicode Character Set .....	1095
10.5. Character Set Configuration .....	1101
10.6. MySQL Server Time Zone Support .....	1102
10.6.1. Staying Current with Time Zone Changes .....	1104
10.6.2. Time Zone Leap Second Support .....	1105
10.7. MySQL Server Locale Support .....	1107

This chapter covers issues of globalization, which includes internationalization (MySQL's capabilities for adapting to local use) and localization (selecting particular local conventions):

- MySQL support for character sets in SQL statements.
- How to configure the server to support different character sets.
- Selecting the language for error messages.
- How to set the server's time zone and enable per-connection time zone support.
- Selecting the locale for day and month names.

## 10.1. Character Set Support

MySQL includes character set support that enables you to store data using a variety of character sets and perform comparisons according to a variety of collations. You can specify character sets at the server, database, table, and column level. MySQL supports the use of character sets for the [MyISAM](#), [MEMORY](#), and [InnoDB](#) storage engines.

This chapter discusses the following topics:

- What are character sets and collations?
- The multiple-level default system for character set assignment.
- Syntax for specifying character sets and collations.
- Affected functions and operations.
- Unicode support.
- The character sets and collations that are available, with notes.

Character set issues affect not only data storage, but also communication between client programs and the MySQL server. If you want the client program to communicate with the server using a character set different from the default, you'll need to indicate which one. For example, to use the `utf8` Unicode character set, issue this statement after connecting to the server:

```
SET NAMES 'utf8';
```

For more information about configuring character sets for application use and character set-related issues in client/server communication, see [Section 10.1.5, “Configuring the Character Set and Collation for Applications”](#), and [Section 10.1.4, “Connection Character Sets and Collations”](#).

### 10.1.1. Character Sets and Collations in General

A *character set* is a set of symbols and encodings. A *collation* is a set of rules for comparing characters in a character set. Let's make the distinction clear with an example of an imaginary character set.

Suppose that we have an alphabet with four letters: “A”, “B”, “a”, “b”. We give each letter a number: “A” = 0, “B” = 1, “a” = 2, “b” = 3. The letter “A” is a symbol, the number 0 is the **encoding** for “A”, and the combination of all four letters and their encodings is a **character set**.

Suppose that we want to compare two string values, “A” and “B”. The simplest way to do this is to look at the encodings: 0 for “A” and 1 for “B”. Because 0 is less than 1, we say “A” is less than “B”. What we've just done is apply a collation to our character set. The collation is a set of rules (only one rule in this case): “compare the encodings.” We call this simplest of all possible collations a *binary* collation.

But what if we want to say that the lowercase and uppercase letters are equivalent? Then we would have at least two rules: (1) treat the lowercase letters “a” and “b” as equivalent to “A” and “B”; (2) then compare the encodings. We call this a *case-insensitive* collation. It is a little more complex than a binary collation.

In real life, most character sets have many characters: not just “A” and “B” but whole alphabets, sometimes multiple alphabets or eastern writing systems with thousands of characters, along with many special symbols and punctuation marks. Also in real life, most collations have many rules, not just for whether to distinguish lettercase, but also for whether to distinguish accents (an “accent” is a mark attached to a character as in German “ö”), and for multiple-character mappings (such as the rule that “ö” = “oe” in one of the two German collations).

MySQL can do these things for you:

- Store strings using a variety of character sets.
- Compare strings using a variety of collations.
- Mix strings with different character sets or collations in the same server, the same database, or even the same table.
- Enable specification of character set and collation at any level.

In these respects, MySQL is far ahead of most other database management systems. However, to use these features effectively, you need to know what character sets and collations are available, how to change the defaults, and how they affect the behavior of string operators and functions.

## 10.1.2. Character Sets and Collations in MySQL

The MySQL server can support multiple character sets. To list the available character sets, use the `SHOW CHARACTER SET` statement. A partial listing follows. For more complete information, see [Section 10.1.14, “Character Sets and Collations That MySQL Supports”](#).

```
mysql> SHOW CHARACTER SET;
+-----+-----+-----+-----+
| Charset | Description | Default collation | Maxlen |
+-----+-----+-----+-----+
| big5    | Big5 Traditional Chinese | big5_chinese_ci | 2 |
| dec8    | DEC West European | dec8_swedish_ci | 1 |
| cp850   | DOS West European | cp850_general_ci | 1 |
| hp8     | HP West European | hp8_english_ci | 1 |
| koi8r   | KOI8-R Relcom Russian | koi8r_general_ci | 1 |
| latin1  | cp1252 West European | latin1_swedish_ci | 1 |
| latin2  | ISO 8859-2 Central European | latin2_general_ci | 1 |
| swe7    | 7bit Swedish | swe7_swedish_ci | 1 |
| ascii   | US ASCII | ascii_general_ci | 1 |
| ujis    | EUC-JP Japanese | ujis_japanese_ci | 3 |
| sjis    | Shift-JIS Japanese | sjis_japanese_ci | 2 |
| hebrew  | ISO 8859-8 Hebrew | hebrew_general_ci | 1 |
| tis620  | TIS620 Thai | tis620_thai_ci | 1 |
| euckr   | EUC-KR Korean | euckr_korean_ci | 2 |
| koi8u   | KOI8-U Ukrainian | koi8u_general_ci | 1 |
| gb2312  | GB2312 Simplified Chinese | gb2312_chinese_ci | 2 |
| greek   | ISO 8859-7 Greek | greek_general_ci | 1 |
| cp1250  | Windows Central European | cp1250_general_ci | 1 |
| gbk     | GBK Simplified Chinese | gbk_chinese_ci | 2 |
| latin5  | ISO 8859-9 Turkish | latin5_turkish_ci | 1 |
...

```

Any given character set always has at least one collation. It may have several collations. To list the collations for a character set, use the `SHOW COLLATION` statement. For example, to see the collations for the `latin1` (cp1252 West European) character set, use this statement to find those collation names that begin with `latin1`:

```
mysql> SHOW COLLATION LIKE 'latin1%';
+-----+-----+-----+-----+-----+-----+
| Collation | Charset | Id | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| latin1_german1_ci | latin1 | 5 | | | 0 |
| latin1_swedish_ci | latin1 | 8 | Yes | Yes | 1 |
| latin1_danish_ci | latin1 | 15 | | | 0 |
| latin1_german2_ci | latin1 | 31 | | Yes | 2 |
| latin1_bin | latin1 | 47 | | Yes | 1 |
| latin1_general_ci | latin1 | 48 | | | 0 |
| latin1_general_cs | latin1 | 49 | | | 0 |
| latin1_spanish_ci | latin1 | 94 | | | 0 |
+-----+-----+-----+-----+-----+-----+

```

The `latin1` collations have the following meanings.

Collation	Meaning
<code>latin1_german1_ci</code>	German DIN-1
<code>latin1_swedish_ci</code>	Swedish/Finnish
<code>latin1_danish_ci</code>	Danish/Norwegian
<code>latin1_german2_ci</code>	German DIN-2
<code>latin1_bin</code>	Binary according to <code>latin1</code> encoding
<code>latin1_general_ci</code>	Multilingual (Western European)
<code>latin1_general_cs</code>	Multilingual (ISO Western European), case sensitive
<code>latin1_spanish_ci</code>	Modern Spanish

Collations have these general characteristics:

- Two different character sets cannot have the same collation.
- Each character set has one collation that is the *default collation*. For example, the default collation for `latin1` is `latin1_swedish_ci`. The output for `SHOW CHARACTER SET` indicates which collation is the default for each displayed character set.
- There is a convention for collation names: They start with the name of the character set with which they are associated, they usually include a language name, and they end with `_ci` (case insensitive), `_cs` (case sensitive), or `_bin` (binary).

In cases where a character set has multiple collations, it might not be clear which collation is most suitable for a given application. To avoid choosing the wrong collation, it can be helpful to perform some comparisons with representative data values to make sure that a given collation sorts values the way you expect.

[Collation-Charts.Org](http://Collation-Charts.Org) is a useful site for information that shows how one collation compares to another.

### 10.1.3. Specifying Character Sets and Collations

There are default settings for character sets and collations at four levels: server, database, table, and column. The description in the following sections may appear complex, but it has been found in practice that multiple-level defaulting leads to natural and obvious results.

`CHARACTER SET` is used in clauses that specify a character set. `CHARSET` can be used as a synonym for `CHARACTER SET`.

Character set issues affect not only data storage, but also communication between client programs and the MySQL server. If you want the client program to communicate with the server using a character set different from the default, you'll need to indicate which one. For example, to use the `utf8` Unicode character set, issue this statement after connecting to the server:

```
SET NAMES 'utf8';
```

For more information about character set-related issues in client/server communication, see [Section 10.1.4, “Connection Character Sets and Collations”](#).

#### 10.1.3.1. Server Character Set and Collation

MySQL Server has a server character set and a server collation. These can be set at server startup on the command line or in an option file and changed at runtime.

Initially, the server character set and collation depend on the options that you use when you start `mysqld`. You can use `--character-set-server` [417] for the character set. Along with it, you can add `--collation-server` [417] for the collation. If you don't specify a character set, that is the same as saying `--character-set-server=latin1` [417]. If you specify only a character set (for example, `latin1`) but not a collation, that is the same as saying `--character-set-server=latin1` [417] `--collation-server=latin1_swedish_ci` [417] because `latin1_swedish_ci` is the default collation for `latin1`. Therefore, the following three commands all have the same effect:

```
shell> mysqld
shell> mysqld --character-set-server=latin1
shell> mysqld --character-set-server=latin1 \
--collation-server=latin1_swedish_ci
```

One way to change the settings is by recompiling. To change the default server character set and collation when building from sources, use the `DEFAULT_CHARSET` [134] and `DEFAULT_COLLATION` [134] options for `CMake`. For example:

```
shell> cmake . -DDEFAULT_CHARSET=latin1
```

Or:

```
shell> cmake . -DDEFAULT_CHARSET=latin1 \
-DDEFAULT_COLLATION=latin1_german1_ci
```

Both `mysqld` and `CMake` verify that the character set/collation combination is valid. If not, each program displays an error message and terminates.

The server character set and collation are used as default values if the database character set and collation are not specified in `CREATE DATABASE` statements. They have no other purpose.

The current server character set and collation can be determined from the values of the `character_set_server` [468] and `collation_server` [470] system variables. These variables can be changed at runtime.

### 10.1.3.2. Database Character Set and Collation

Every database has a database character set and a database collation. The `CREATE DATABASE` and `ALTER DATABASE` statements have optional clauses for specifying the database character set and collation:

```
CREATE DATABASE db_name
  [[DEFAULT] CHARACTER SET charset_name]
  [[DEFAULT] COLLATE collation_name]

ALTER DATABASE db_name
  [[DEFAULT] CHARACTER SET charset_name]
  [[DEFAULT] COLLATE collation_name]
```

The keyword `SCHEMA` can be used instead of `DATABASE`.

All database options are stored in a text file named `db.opt` that can be found in the database directory.

The `CHARACTER SET` and `COLLATE` clauses make it possible to create databases with different character sets and collations on the same MySQL server.

Example:

```
CREATE DATABASE db_name CHARACTER SET latin1 COLLATE latin1_swedish_ci;
```

MySQL chooses the database character set and database collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` are specified, character set `X` and collation `Y` are used.
- If `CHARACTER SET X` is specified without `COLLATE`, character set `X` and its default collation are used. To see the default collation for each character set, use the `SHOW COLLATION` statement.
- If `COLLATE Y` is specified without `CHARACTER SET`, the character set associated with `Y` and collation `Y` are used.
- Otherwise, the server character set and server collation are used.

The database character set and collation are used as default values for table definitions if the table character set and collation are not specified in `CREATE TABLE` statements. The database character set also is used by `LOAD DATA INFILE`. The character set and collation have no other purposes.

The character set and collation for the default database can be determined from the values of the `character_set_database` [468] and `collation_database` [469] system variables. The server sets these variables whenever the default database changes. If there is no default



database, the variables have the same value as the corresponding server-level system variables, `character_set_server` [468] and `collation_server` [470].

### 10.1.3.3. Table Character Set and Collation

Every table has a table character set and a table collation. The `CREATE TABLE` and `ALTER TABLE` statements have optional clauses for specifying the table character set and collation:

```
CREATE TABLE tbl_name (column_list)
  [[DEFAULT] CHARACTER SET charset_name]
  [COLLATE collation_name]]

ALTER TABLE tbl_name
  [[DEFAULT] CHARACTER SET charset_name]
  [COLLATE collation_name]
```

Example:

```
CREATE TABLE t1 ( ... )
CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

MySQL chooses the table character set and collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` are specified, character set `X` and collation `Y` are used.
- If `CHARACTER SET X` is specified without `COLLATE`, character set `X` and its default collation are used. To see the default collation for each character set, use the `SHOW COLLATION` statement.
- If `COLLATE Y` is specified without `CHARACTER SET`, the character set associated with `Y` and collation `Y` are used.
- Otherwise, the database character set and collation are used.

The table character set and collation are used as default values for column definitions if the column character set and collation are not specified in individual column definitions. The table character set and collation are MySQL extensions; there are no such things in standard SQL.

### 10.1.3.4. Column Character Set and Collation

Every “character” column (that is, a column of type `CHAR`, `VARCHAR`, or `TEXT`) has a column character set and a column collation. Column definition syntax for `CREATE TABLE` and `ALTER TABLE` has optional clauses for specifying the column character set and collation:

```
col_name {CHAR | VARCHAR | TEXT} (col_length)
  [CHARACTER SET charset_name]
  [COLLATE collation_name]
```

These clauses can also be used for `ENUM` and `SET` columns:

```
col_name {ENUM | SET} (val_list)
  [CHARACTER SET charset_name]
  [COLLATE collation_name]
```

Examples:

```
CREATE TABLE t1
(
  col1 VARCHAR(5)
  CHARACTER SET latin1
  COLLATE latin1_german1_ci
);

ALTER TABLE t1 MODIFY
col1 VARCHAR(5)
CHARACTER SET latin1
```

```
COLLATE latin1_swedish_ci;
```

MySQL chooses the column character set and collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` are specified, character set `X` and collation `Y` are used.

```
CREATE TABLE t1
(
  coll CHAR(10) CHARACTER SET utf8 COLLATE utf8_unicode_ci
) CHARACTER SET latin1 COLLATE latin1_bin;
```

The character set and collation are specified for the column, so they are used. The column has character set `utf8` and collation `utf8_unicode_ci`.

- If `CHARACTER SET X` is specified without `COLLATE`, character set `X` and its default collation are used.

```
CREATE TABLE t1
(
  coll CHAR(10) CHARACTER SET utf8
) CHARACTER SET latin1 COLLATE latin1_bin;
```

The character set is specified for the column, but the collation is not. The column has character set `utf8` and the default collation for `utf8`, which is `utf8_general_ci`. To see the default collation for each character set, use the `SHOW COLLATION` statement.

- If `COLLATE Y` is specified without `CHARACTER SET`, the character set associated with `Y` and collation `Y` are used.

```
CREATE TABLE t1
(
  coll CHAR(10) COLLATE utf8_polish_ci
) CHARACTER SET latin1 COLLATE latin1_bin;
```

The collation is specified for the column, but the character set is not. The column has collation `utf8_polish_ci` and the character set is the one associated with the collation, which is `utf8`.

- Otherwise, the table character set and collation are used.

```
CREATE TABLE t1
(
  coll CHAR(10)
) CHARACTER SET latin1 COLLATE latin1_bin;
```

Neither the character set nor collation are specified for the column, so the table defaults are used. The column has character set `latin1` and collation `latin1_bin`.

The `CHARACTER SET` and `COLLATE` clauses are standard SQL.

If you use `ALTER TABLE` to convert a column from one character set to another, MySQL attempts to map the data values, but if the character sets are incompatible, there may be data loss.

### 10.1.3.5. Character String Literal Character Set and Collation

Every character string literal has a character set and a collation.

A character string literal may have an optional character set introducer and `COLLATE` clause:

```
[_charset_name]'string' [COLLATE collation_name]
```

Examples:

```
SELECT 'string';
SELECT _latin1'string';
SELECT _latin1'string' COLLATE latin1_danish_ci;
```

For the simple statement `SELECT 'string'`, the string has the character set and collation defined by the `character_set_connection` [467] and `collation_connection` [469] system variables.

The `_charset_name` expression is formally called an *introducer*. It tells the parser, “the string that is about to follow uses character set *X*.” Because this has confused people in the past, we emphasize that an introducer does not change the string to the introducer character set like `CONVERT()` [1245] would do. It does not change the string's value, although padding may occur. The introducer is just a signal. An introducer is also legal before standard hex literal and numeric hex literal notation (`x'literal'` and `0xnnnn`), or before bit-field literal notation (`b'literal'` and `0bnnnn`).

Examples:

```
SELECT _latin1 x'AABBCC';
SELECT _latin1 0xAABBCC;
SELECT _latin1 b'1100011';
SELECT _latin1 0b1100011;
```

MySQL determines a literal's character set and collation in the following manner:

- If both `_X` and `COLLATE Y` are specified, character set *X* and collation *Y* are used.
- If `_X` is specified but `COLLATE` is not specified, character set *X* and its default collation are used. To see the default collation for each character set, use the `SHOW COLLATION` statement.
- Otherwise, the character set and collation given by the `character_set_connection` [467] and `collation_connection` [469] system variables are used.

Examples:

- A string with `latin1` character set and `latin1_german1_ci` collation:

```
SELECT _latin1'Müller' COLLATE latin1_german1_ci;
```

- A string with `latin1` character set and its default collation (that is, `latin1_swedish_ci`):

```
SELECT _latin1'Müller';
```

- A string with the connection default character set and collation:

```
SELECT 'Müller';
```

Character set introducers and the `COLLATE` clause are implemented according to standard SQL specifications.

An introducer indicates the character set for the following string, but does not change now how the parser performs escape processing within the string. Escapes are always interpreted by the parser according to the character set given by `character_set_connection` [467].

The following examples show that escape processing occurs using `character_set_connection` [467] even in the presence of an introducer. The examples use `SET NAMES` (which changes `character_set_connection` [467], as discussed in Section 10.1.4, “Connection Character Sets and Collations”), and display the resulting strings using the `HEX()` [1180] function so that the exact string contents can be seen.

Example 1:

```
mysql> SET NAMES latin1;
Query OK, 0 rows affected (0.01 sec)

mysql> SELECT HEX('à\n'), HEX(_sjis'à\n');
+-----+-----+
| HEX('à\n') | HEX(_sjis'à\n') |
+-----+-----+
| E00A      | E00A              |
```

```
+-----+-----+
1 row in set (0.00 sec)
```

Here, “à” (hex value `E0`) is followed by “\n”, the escape sequence for newline. The escape sequence is interpreted using the `character_set_connection` [467] value of `latin1` to produce a literal newline (hex value `0A`). This happens even for the second string. That is, the introducer of `_sjis` does not affect the parser's escape processing.

Example 2:

```
mysql> SET NAMES sjis;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT HEX('à\n'), HEX(_latin1'à\n');
+-----+-----+
| HEX('à\n') | HEX(_latin1'à\n') |
+-----+-----+
| E05C6E     | E05C6E             |
+-----+-----+
1 row in set (0.04 sec)
```

Here, `character_set_connection` [467] is `sjis`, a character set in which the sequence of “à” followed by “\” (hex values `05` and `5C`) is a valid multi-byte character. Hence, the first two bytes of the string are interpreted as a single `sjis` character, and the “\” is not interpreted as an escape character. The following “n” (hex value `6E`) is not interpreted as part of an escape sequence. This is true even for the second string; the introducer of `_latin1` does not affect escape processing.

### 10.1.3.6. National Character Set

Standard SQL defines `NCHAR` or `NATIONAL CHAR` as a way to indicate that a `CHAR` column should use some predefined character set. MySQL 5.6 uses `utf8` as this predefined character set. For example, these data type declarations are equivalent:

```
CHAR(10) CHARACTER SET utf8
NATIONAL CHARACTER(10)
NCHAR(10)
```

As are these:

```
VARCHAR(10) CHARACTER SET utf8
NATIONAL VARCHAR(10)
NCHAR VARCHAR(10)
NATIONAL CHARACTER VARYING(10)
NATIONAL CHAR VARYING(10)
```

You can use `N'literal'` (or `n'literal'`) to create a string in the national character set. These statements are equivalent:

```
SELECT N'some text';
SELECT n'some text';
SELECT _utf8'some text';
```

For information on upgrading character sets to MySQL 5.6 from versions prior to 4.1, see the *MySQL 3.23, 4.0, 4.1 Reference Manual*.

### 10.1.3.7. Examples of Character Set and Collation Assignment

The following examples show how MySQL determines default character set and collation values.

#### Example 1: Table and Column Definition

```
CREATE TABLE t1
(
  c1 CHAR(10) CHARACTER SET latin1 COLLATE latin1_german1_ci
) DEFAULT CHARACTER SET latin2 COLLATE latin2_bin;
```

Here we have a column with a `latin1` character set and a `latin1_german1_ci` collation. The definition is explicit, so that is straightforward. Notice that there is no problem with storing a `latin1` column in a `latin2` table.

### Example 2: Table and Column Definition

```
CREATE TABLE t1
(
  c1 CHAR(10) CHARACTER SET latin1
) DEFAULT CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

This time we have a column with a `latin1` character set and a default collation. Although it might seem natural, the default collation is not taken from the table level. Instead, because the default collation for `latin1` is always `latin1_swedish_ci`, column `c1` has a collation of `latin1_swedish_ci` (not `latin1_danish_ci`).

### Example 3: Table and Column Definition

```
CREATE TABLE t1
(
  c1 CHAR(10)
) DEFAULT CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

We have a column with a default character set and a default collation. In this circumstance, MySQL checks the table level to determine the column character set and collation. Consequently, the character set for column `c1` is `latin1` and its collation is `latin1_danish_ci`.

### Example 4: Database, Table, and Column Definition

```
CREATE DATABASE d1
  DEFAULT CHARACTER SET latin2 COLLATE latin2_czech_ci;
USE d1;
CREATE TABLE t1
(
  c1 CHAR(10)
);
```

We create a column without specifying its character set and collation. We're also not specifying a character set and a collation at the table level. In this circumstance, MySQL checks the database level to determine the table settings, which thereafter become the column settings.) Consequently, the character set for column `c1` is `latin2` and its collation is `latin2_czech_ci`.

## 10.1.3.8. Compatibility with Other DBMSs

For MaxDB compatibility these two statements are the same:

```
CREATE TABLE t1 (f1 CHAR(N) UNICODE);
CREATE TABLE t1 (f1 CHAR(N) CHARACTER SET ucs2);
```

## 10.1.4. Connection Character Sets and Collations

Several character set and collation system variables relate to a client's interaction with the server. Some of these have been mentioned in earlier sections:

- The server character set and collation are the values of the `character_set_server` [468] and `collation_server` [470] system variables.
- The character set and collation of the default database are the values of the `character_set_database` [468] and `collation_database` [469] system variables.

Additional character set and collation system variables are involved in handling traffic for the connection between a client and the server. Every client has connection-related character set and collation system variables.

A “connection” is what you make when you connect to the server. The client sends SQL statements, such as queries, over the connection to the server. The server sends responses, such as result sets or error messages, over the connection back to the client. This leads to several questions about character set and collation handling for client connections, each of which can be answered in terms of system variables:

- What character set is the statement in when it leaves the client?

The server takes the `character_set_client` [467] system variable to be the character set in which statements are sent by the client.

- What character set should the server translate a statement to after receiving it?

For this, the server uses the `character_set_connection` [467] and `collation_connection` [469] system variables. It converts statements sent by the client from `character_set_client` [467] to `character_set_connection` [467] (except for string literals that have an introducer such as `_latin1` or `_utf8`). `collation_connection` [469] is important for comparisons of literal strings. For comparisons of strings with column values, `collation_connection` [469] does not matter because columns have their own collation, which has a higher collation precedence.

- What character set should the server translate to before shipping result sets or error messages back to the client?

The `character_set_results` [468] system variable indicates the character set in which the server returns query results to the client. This includes result data such as column values, and result metadata such as column names and error messages.

Clients can fine-tune the settings for these variables, or depend on the defaults (in which case, you can skip the rest of this section). If you do not use the defaults, you must change the character settings *for each connection to the server*.

Two statements affect the connection-related character set variables as a group:

- `SET NAMES 'charset_name' [COLLATE 'collation_name']`

`SET NAMES` indicates what character set the client will use to send SQL statements to the server. Thus, `SET NAMES 'cp1251'` tells the server, “future incoming messages from this client are in character set `cp1251`.” It also specifies the character set that the server should use for sending results back to the client. (For example, it indicates what character set to use for column values if you use a `SELECT` statement.)

A `SET NAMES 'charset_name'` statement is equivalent to these three statements:

```
SET character_set_client = charset_name;
SET character_set_results = charset_name;
SET character_set_connection = charset_name;
```

Setting `character_set_connection` [467] to `charset_name` also implicitly sets `collation_connection` [469] to the default collation for `charset_name`. It is unnecessary to set that collation explicitly. To specify a particular collation, use the optional `COLLATE` clause:

```
SET NAMES 'charset_name' COLLATE 'collation_name'
```

- `SET CHARACTER SET charset_name`

`SET CHARACTER SET` is similar to `SET NAMES` but sets `character_set_connection` [467] and `collation_connection` [469] to `character_set_database` [468] and `collation_database` [469]. A `SET CHARACTER SET charset_name` statement is equivalent to these three statements:

```
SET character_set_client = charset_name;
```

```
SET character_set_results = charset_name ;
SET collation_connection = @@collation_database ;
```

Setting `collation_connection` [469] also implicitly sets `character_set_connection` [467] to the character set associated with the collation (equivalent to executing `SET character_set_connection = @@character_set_database`). It is unnecessary to set `character_set_connection` [467] explicitly.



**Note**

`ucs2`, `utf16`, `utf16le`, and `utf32` cannot be used as a client character set, which means that they do not work for `SET NAMES` or `SET CHARACTER SET`.

The MySQL client programs `mysql`, `mysqladmin`, `mysqlcheck`, `mysqlimport`, and `mysqlshow` determine the default character set to use as follows:

- In the absence of other information, the programs use the compiled-in default character set, usually `latin1`.
- The programs can autodetect which character set to use based on the operating system setting, such as the value of the `LANG` or `LC_ALL` locale environment variable on Unix systems or the code page setting on Windows systems. For systems on which the locale is available from the OS, the client uses it to set the default character set rather than using the compiled-in default. For example, setting `LANG` to `ru_RU.KOI8-R` causes the `koi8r` character set to be used. Thus, users can configure the locale in their environment for use by MySQL clients.

The OS character set is mapped to the closest MySQL character set if there is no exact match. If the client does not support the matching character set, it uses the compiled-in default. For example, `ucs2` is not supported as a connection character set.

C applications can use character set autodetection based on the OS setting by invoking `mysql_options()` as follows before connecting to the server:

```
mysql_options(mysql,
              MYSQL_SET_CHARSET_NAME,
              MYSQL_AUTODETECT_CHARSET_NAME) ;
```

- The programs support a `--default-character-set` [249] option, which enables users to specify the character set explicitly to override whatever default the client otherwise determines.

When a client connects to the server, it sends the name of the character set that it wants to use. The server uses the name to set the `character_set_client` [467], `character_set_results` [468], and `character_set_connection` [467] system variables. In effect, the server performs a `SET NAMES` operation using the character set name.

With the `mysql` client, to use a character set different from the default, you could explicitly execute `SET NAMES` every time you start up. To accomplish the same result more easily, add the `--default-character-set` [249] option setting to your `mysql` command line or in your option file. For example, the following option file setting changes the three connection-related character set variables set to `koi8r` each time you invoke `mysql`:

```
[mysql]
default-character-set=koi8r
```

If you are using the `mysql` client with auto-reconnect enabled (which is not recommended), it is preferable to use the `charset` command rather than `SET NAMES`. For example:

```
mysql> charset utf8
Charset changed
```

The `charset` command issues a `SET NAMES` statement, and also changes the default character set that `mysql` uses when it reconnects after the connection has dropped.



Example: Suppose that `column1` is defined as `CHAR(5) CHARACTER SET latin2`. If you do not say `SET NAMES` or `SET CHARACTER SET`, then for `SELECT column1 FROM t`, the server sends back all the values for `column1` using the character set that the client specified when it connected. On the other hand, if you say `SET NAMES 'latin1'` or `SET CHARACTER SET latin1` before issuing the `SELECT` statement, the server converts the `latin2` values to `latin1` just before sending results back. Conversion may be lossy if there are characters that are not in both character sets.

If you want the server to perform no conversion of result sets or error messages, set `character_set_results` [468] to `NULL` or `binary`:

```
SET character_set_results = NULL;
```

To see the values of the character set and collation system variables that apply to your connection, use these statements:

```
SHOW VARIABLES LIKE 'character_set%';
SHOW VARIABLES LIKE 'collation%';
```

You must also consider the environment within which your MySQL applications execute. See [Section 10.1.5, “Configuring the Character Set and Collation for Applications”](#).

For more information about character sets and error messages, see [Section 10.1.6, “Character Set for Error Messages”](#).

## 10.1.5. Configuring the Character Set and Collation for Applications

For applications that store data using the default MySQL character set and collation (`latin1`, `latin1_swedish_ci`), no special configuration should be needed. If applications require data storage using a different character set or collation, you can configure character set information several ways:

- Specify character settings per database. For example, applications that use one database might require `utf8`, whereas applications that use another database might require `sjis`.
- Specify character settings at server startup. This causes the server to use the given settings for all applications that do not make other arrangements.
- Specify character settings at configuration time, if you build MySQL from source. This causes the server to use the given settings for all applications, without having to specify them at server startup.

When different applications require different character settings, the per-database technique provides a good deal of flexibility. If most or all applications use the same character set, specifying character settings at server startup or configuration time may be most convenient.

For the per-database or server-startup techniques, the settings control the character set for data storage. Applications must also tell the server which character set to use for client/server communications, as described in the following instructions.

The examples shown here assume use of the `utf8` character set and `utf8_general_ci` collation.

**Specify character settings per database.** To create a database such that its tables will use a given default character set and collation for data storage, use a `CREATE DATABASE` statement like this:

```
CREATE DATABASE mydb
  DEFAULT CHARACTER SET utf8
  DEFAULT COLLATE utf8_general_ci;
```

Tables created in the database will use `utf8` and `utf8_general_ci` by default for any character columns.

Applications that use the database should also configure their connection to the server each time they connect. This can be done by executing a `SET NAMES 'utf8'` statement after connecting. The statement can be used regardless of connection method: The `mysql` client, PHP scripts, and so forth.

In some cases, it may be possible to configure the connection to use the desired character set some other way. For example, for connections made using `mysql`, you can specify the `--default-character-set=utf8` [249] command-line option to achieve the same effect as `SET NAMES 'utf8'`.

For more information about configuring client connections, see [Section 10.1.4, “Connection Character Sets and Collations”](#).

If you change the default character set or collation for a database, stored routines that use the database defaults must be dropped and recreated so that they use the new defaults. (In a stored routine, variables with character data types use the database defaults if the character set or collation are not specified explicitly. See [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#).)

**Specify character settings at server startup.** To select a character set and collation at server startup, use the `--character-set-server` [417] and `--collation-server` [417] options. For example, to specify the options in an option file, include these lines:

```
[mysqld]
character-set-server=utf8
collation-server=utf8_general_ci
```

These settings apply server-wide and apply as the defaults for databases created by any application, and for tables created in those databases.

It is still necessary for applications to configure their connection using `SET NAMES` or equivalent after they connect, as described previously. You might be tempted to start the server with the `--init_connect="SET NAMES 'utf8'"` [489] option to cause `SET NAMES` to be executed automatically for each client that connects. However, this will yield inconsistent results because the `init_connect` [489] value is not executed for users who have the `SUPER` [756] privilege.

**Specify character settings at MySQL configuration time.** To select a character set and collation when you configure and build MySQL from source, use the `DEFAULT_CHARSET` [134] and `DEFAULT_COLLATION` [134] options for CMake:

```
shell> cmake . -DDEFAULT_CHARSET=utf8 \
-DDEFAULT_COLLATION=utf8_general_ci
```

The resulting server uses `utf8` and `utf8_general_ci` as the default for databases and tables and for client connections. It is unnecessary to use `--character-set-server` [417] and `--collation-server` [417] to specify those defaults at server startup. It is also unnecessary for applications to configure their connection using `SET NAMES` or equivalent after they connect to the server.

Regardless of how you configure the MySQL character set for application use, you must also consider the environment within which those applications execute. If you will send statements using UTF-8 text taken from a file that you create in an editor, you should edit the file with the locale of your environment set to UTF-8 so that the file encoding is correct and so that the operating system handles it correctly. If you use the `mysql` client from within a terminal window, the window must be configured to use UTF-8 or characters may not display properly. For a script that executes in a Web environment, the script must handle character encoding properly for its interaction with the MySQL server, and it must generate pages that correctly indicate the encoding so that browsers know how to display the content of the pages. For example, you can include this `<meta>` tag within your `<head>` element:

```
<meta http-equiv="Content-Type" content="text/html; charset=utf-8" />
```

## 10.1.6. Character Set for Error Messages

This section describes how the server uses character sets for constructing error messages and returning them to clients. For information about the language of error messages (rather than the character set), see [Section 10.2, “Setting the Error Message Language”](#).

In MySQL 5.6, the server constructs error messages using UTF-8 and returns them to clients in the character set specified by the `character_set_results` [468] system variable.

The server constructs error messages as follows:

- The message template uses UTF-8.
- Parameters in the message template are replaced with values that apply to a specific error occurrence:
  - Identifiers such as table or column names use UTF-8 internally so they are copied as is.
  - Character (nonbinary) string values are converted from their character set to UTF-8.
  - Binary string values are copied as is for bytes in the range `0x20` to `0x7E`, and using `\x` hex encoding for bytes outside that range. For example, if a duplicate-key error occurs for an attempt to insert `0x41CF9F` into a `VARBINARY` unique column, the resulting error message uses UTF-8 with some bytes hex encoded:

```
Duplicate entry 'A\xC3\x9F' for key 1
```

To return a message to the client after it has been constructed, the server converts it from UTF-8 to the character set specified by the `character_set_results` [468] system variable. If `character_set_results` [468] has a value of `NULL` or `binary`, no conversion occurs. No conversion occurs if the variable value is `utf8`, either, because that matches the original error message character set.

For characters that cannot be represented in `character_set_results` [468], some encoding may occur during the conversion. The encoding uses Unicode code point values:

- Characters in the Basic Multilingual Plane (BMP) range (`0x0000` to `0xFFFF`) are written using `\nnnn` notation.
- Characters outside the BMP range (`0x01000` to `0x10FFFF`) are written using `\+nnnnnnn` notation.

Clients can set `character_set_results` [468] to control the character set in which they receive error messages. The variable can be set directly, or indirectly by means such as `SET NAMES`. For more information about `character_set_results` [468], see [Section 10.1.4, "Connection Character Sets and Collations"](#).

The encoding that occurs during the conversion to `character_set_results` [468] before returning error messages to clients can result in different message content compared to earlier versions (before MySQL 5.5). For example, if an error occurs for an attempt to drop a table named `^` (KATAKANA LETTER PE) and `character_set_results` [468] is a character set such as `latin1` that does not contain that character, the resulting message sent to the client has an encoded table name:

```
ERROR 1051 (42S02): Unknown table '\30DA'
```

Before MySQL 5.5, the name is not encoded:

```
ERROR 1051 (42S02): Unknown table '^'
```

## 10.1.7. Collation Issues

The following sections discuss various aspects of character set collations.

### 10.1.7.1. Collation Names

MySQL collation names follow these rules:

- A name ending in `_ci` indicates a case-insensitive collation.

- A name ending in `_cs` indicates a case-sensitive collation.
- A name ending in `_bin` indicates a binary collation. Character comparisons are based on character binary code values.
- Unicode collation names may include a version number to indicate the version of the Unicode Collation Algorithm (UCA) on which the collation is based. UCA-based collations without a version number in the name use the version-4.0.0 UCA weight keys: <http://www.unicode.org/Public/UCA/4.0.0/allkeys-4.0.0.txt>. A collation name such as `utf8_unicode_520_ci` is based on UCA 5.2.0 weight keys: <http://www.unicode.org/Public/UCA/5.2.0/allkeys.txt>.

### 10.1.7.2. Using `COLLATE` in SQL Statements

With the `COLLATE` clause, you can override whatever the default collation is for a comparison. `COLLATE` may be used in various parts of SQL statements. Here are some examples:

- With `ORDER BY`:

```
SELECT k
FROM t1
ORDER BY k COLLATE latin1_german2_ci;
```

- With `AS`:

```
SELECT k COLLATE latin1_german2_ci AS k1
FROM t1
ORDER BY k1;
```

- With `GROUP BY`:

```
SELECT k
FROM t1
GROUP BY k COLLATE latin1_german2_ci;
```

- With aggregate functions:

```
SELECT MAX(k COLLATE latin1_german2_ci)
FROM t1;
```

- With `DISTINCT`:

```
SELECT DISTINCT k COLLATE latin1_german2_ci
FROM t1;
```

- With `WHERE`:

```
SELECT *
FROM t1
WHERE _latin1 'Müller' COLLATE latin1_german2_ci = k;
```

```
SELECT *
FROM t1
WHERE k LIKE _latin1 'Müller' COLLATE latin1_german2_ci;
```

- With `HAVING`:

```
SELECT k
FROM t1
GROUP BY k
HAVING k = _latin1 'Müller' COLLATE latin1_german2_ci;
```

### 10.1.7.3. `COLLATE` Clause Precedence

The `COLLATE` clause has high precedence (higher than `||` [1172]), so the following two expressions are equivalent:

```
x || y COLLATE z
```

```
x || (y COLLATE z)
```

#### 10.1.7.4. Collations Must Be for the Right Character Set

Each character set has one or more collations, but each collation is associated with one and only one character set. Therefore, the following statement causes an error message because the `latin2_bin` collation is not legal with the `latin1` character set:

```
mysql> SELECT _latin1 'x' COLLATE latin2_bin;
ERROR 1253 (42000): COLLATION 'latin2_bin' is not valid
for CHARACTER SET 'latin1'
```

#### 10.1.7.5. Collation of Expressions

In the great majority of statements, it is obvious what collation MySQL uses to resolve a comparison operation. For example, in the following cases, it should be clear that the collation is the collation of column `charset_name`:

```
SELECT x FROM T ORDER BY x;
SELECT x FROM T WHERE x = x;
SELECT DISTINCT x FROM T;
```

However, with multiple operands, there can be ambiguity. For example:

```
SELECT x FROM T WHERE x = 'Y';
```

Should the comparison use the collation of the column `x`, or of the string literal `'Y'`? Both `x` and `'Y'` have collations, so which collation takes precedence?

Standard SQL resolves such questions using what used to be called “coercibility” rules. MySQL assigns coercibility values as follows:

- An explicit `COLLATE` clause has a coercibility of 0. (Not coercible at all.)
- The concatenation of two strings with different collations has a coercibility of 1.
- The collation of a column or a stored routine parameter or local variable has a coercibility of 2.
- A “system constant” (the string returned by functions such as `USER()` [1273] or `VERSION()` [1273]) has a coercibility of 3.
- The collation of a literal has a coercibility of 4.
- `NULL` or an expression that is derived from `NULL` has a coercibility of 5.

MySQL uses coercibility values with the following rules to resolve ambiguities:

- Use the collation with the lowest coercibility value.
- If both sides have the same coercibility, then:
  - If both sides are Unicode, or both sides are not Unicode, it is an error.
  - If one of the sides has a Unicode character set, and another side has a non-Unicode character set, the side with Unicode character set wins, and automatic character set conversion is applied to the non-Unicode side. For example, the following statement does not return an error:

```
SELECT CONCAT(utf8_column, latin1_column) FROM t1;
```

It returns a result that has a character set of `utf8` and the same collation as `utf8_column`. Values of `latin1_column` are automatically converted to `utf8` before concatenating.

- For an operation with operands from the same character set but that mix a `_bin` collation and a `_ci` or `_cs` collation, the `_bin` collation is used. This is similar to how operations that mix

nonbinary and binary strings evaluate the operands as binary strings, except that it is for collations rather than data types.

Although automatic conversion is not in the SQL standard, the SQL standard document does say that every character set is (in terms of supported characters) a “subset” of Unicode. Because it is a well-known principle that “what applies to a superset can apply to a subset,” we believe that a collation for Unicode can apply for comparisons with non-Unicode strings.

Examples:

Comparison	Collation Used
<code>column1 = 'A'</code>	Use collation of <code>column1</code>
<code>column1 = 'A' COLLATE x</code>	Use collation of <code>'A' COLLATE x</code>
<code>column1 COLLATE x = 'A' COLLATE y</code>	Error

The `COERCIBILITY()` [1266] function can be used to determine the coercibility of a string expression:

```
mysql> SELECT COERCIBILITY('A' COLLATE latin1_swedish_ci);
-> 0
mysql> SELECT COERCIBILITY(VERSION());
-> 3
mysql> SELECT COERCIBILITY('A');
-> 4
```

See [Section 12.14, “Information Functions”](#).

For implicit conversion of a numeric or temporal value to a string, such as occurs for the argument `1` in the expression `CONCAT(1, 'abc')` [1178], the result is a character (nonbinary) string that has a character set and collation determined by the `character_set_connection` [467] and `collation_connection` [469] system variables. See [Section 12.2, “Type Conversion in Expression Evaluation”](#).

### 10.1.7.6. The `_bin` and `binary` Collations

This section describes how `_bin` collations for nonbinary strings differ from the `binary` “collation” for binary strings.

Nonbinary strings (as stored in the `CHAR`, `VARCHAR`, and `TEXT` data types) have a character set and collation. A given character set can have several collations, each of which defines a particular sorting and comparison order for the characters in the set. One of these is the binary collation for the character set, indicated by a `_bin` suffix in the collation name. For example, `latin1` and `utf8` have binary collations named `latin1_bin` and `utf8_bin`.

Binary strings (as stored in the `BINARY`, `VARBINARY`, and `BLOB` data types) have no character set or collation in the sense that nonbinary strings do. (Applied to a binary string, the `CHARSET()` and `COLLATION()` functions both return a value of `binary`.) Binary strings are sequences of bytes and the numeric values of those bytes determine sort order.

The `_bin` collations differ from the `binary` collation in several respects.

**The unit for sorting and comparison.** Binary strings are sequences of bytes. Sorting and comparison is always based on numeric byte values. Nonbinary strings are sequences of characters, which might be multi-byte. Collations for nonbinary strings define an ordering of the character values for sorting and comparison. For the `_bin` collation, this ordering is based solely on binary code values of the characters (which is similar to ordering for binary strings except that a `_bin` collation must take into account that a character might contain multiple bytes). For other collations, character ordering might take additional factors such as lettercase into account.

**Character set conversion.** A nonbinary string has a character set and is converted to another character set in many cases, even when the string has a `_bin` collation:

- When assigning column values from another column that has a different character set:

```
UPDATE t1 SET utf8_bin_column=latin1_column;
INSERT INTO t1 (latin1_column) SELECT utf8_bin_column FROM t2;
```

- When assigning column values for `INSERT` or `UPDATE` using a string literal:

```
SET NAMES latin1;
INSERT INTO t1 (utf8_bin_column) VALUES ('string-in-latin1');
```

- When sending results from the server to a client:

```
SET NAMES latin1;
SELECT utf8_bin_column FROM t2;
```

For binary string columns, no conversion occurs. For the preceding cases, the string value is copied byte-wise.

**Lettercase conversion.** Collations provide information about lettercase of characters, so characters in a nonbinary string can be converted from one lettercase to another, even for `_bin` collations that ignore lettercase for ordering:

```
mysql> SET NAMES latin1 COLLATE latin1_bin;
Query OK, 0 rows affected (0.02 sec)

mysql> SELECT LOWER('aA'), UPPER('zZ');
+-----+-----+
| LOWER('aA') | UPPER('zZ') |
+-----+-----+
| aa          | ZZ          |
+-----+-----+
1 row in set (0.13 sec)
```

The concept of lettercase does not apply to bytes in a binary string. To perform lettercase conversion, the string must be converted to a nonbinary string:

```
mysql> SET NAMES binary;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT LOWER('aA'), LOWER(CONVERT('aA' USING latin1));
+-----+-----+
| LOWER('aA') | LOWER(CONVERT('aA' USING latin1)) |
+-----+-----+
| aA          | aa          |
+-----+-----+
1 row in set (0.00 sec)
```

**Trailing space handling in comparisons.** Nonbinary strings have `PADSPACE` behavior for all collations, including `_bin` collations. Trailing spaces are insignificant in comparisons:

```
mysql> SET NAMES utf8 COLLATE utf8_bin;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT 'a ' = 'a';
+-----+
| 'a ' = 'a' |
+-----+
| 1          |
+-----+
1 row in set (0.00 sec)
```

For binary strings, all characters are significant in comparisons, including trailing spaces:

```
mysql> SET NAMES binary;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT 'a ' = 'a';
+-----+
| 'a ' = 'a' |
```



```

+-----+
|          0 |
+-----+
1 row in set (0.00 sec)

```

**Trailing space handling for inserts and retrievals.** `CHAR(N)` columns store nonbinary strings. Values shorter than `N` characters are extended with spaces on insertion. For retrieval, trailing spaces are removed.

`BINARY(N)` columns store binary strings. Values shorter than `N` bytes are extended with `0x00` bytes on insertion. For retrieval, nothing is removed; a value of the declared length is always returned.

```

mysql> CREATE TABLE t1 (
->   a CHAR(10) CHARACTER SET utf8 COLLATE utf8_bin,
->   b BINARY(10)
-> );
Query OK, 0 rows affected (0.09 sec)

mysql> INSERT INTO t1 VALUES ('a','a');
Query OK, 1 row affected (0.01 sec)

mysql> SELECT HEX(a), HEX(b) FROM t1;
+-----+-----+
| HEX(a) | HEX(b) |
+-----+-----+
| 61     | 61000000000000000000 |
+-----+-----+
1 row in set (0.04 sec)

```

### 10.1.7.7. The `BINARY` Operator

The `BINARY` [1244] operator casts the string following it to a binary string. This is an easy way to force a comparison to be done byte by byte rather than character by character. `BINARY` [1244] also causes trailing spaces to be significant.

```

mysql> SELECT 'a' = 'A';
-> 1
mysql> SELECT BINARY 'a' = 'A';
-> 0
mysql> SELECT 'a' = 'a ';
-> 1
mysql> SELECT BINARY 'a' = 'a ';
-> 0

```

`BINARY str` is shorthand for `CAST(str AS BINARY)` [1245].

The `BINARY` attribute in character column definitions has a different effect. A character column defined with the `BINARY` attribute is assigned the binary collation of the column character set. Every character set has a binary collation. For example, the binary collation for the `latin1` character set is `latin1_bin`, so if the table default character set is `latin1`, these two column definitions are equivalent:

```

CHAR(10) BINARY
CHAR(10) CHARACTER SET latin1 COLLATE latin1_bin

```

The effect of `BINARY` as a column attribute differs from its effect prior to MySQL 4.1. Formerly, `BINARY` resulted in a column that was treated as a binary string. A binary string is a string of bytes that has no character set or collation, which differs from a nonbinary character string that has a binary collation. For both types of strings, comparisons are based on the numeric values of the string unit, but for nonbinary strings the unit is the character and some character sets support multi-byte characters. Section 11.4.2, “The `BINARY` and `VARBINARY` Types”.

The use of `CHARACTER SET binary` in the definition of a `CHAR`, `VARCHAR`, or `TEXT` column causes the column to be treated as a binary data type. For example, the following pairs of definitions are equivalent:

```
CHAR(10) CHARACTER SET binary
BINARY(10)

VARCHAR(10) CHARACTER SET binary
VARBINARY(10)

TEXT CHARACTER SET binary
BLOB
```

### 10.1.7.8. Examples of the Effect of Collation

#### Example 1: Sorting German Umlauts

Suppose that column `X` in table `T` has these `latin1` column values:

```
Muffler
Müller
MX Systems
MySQL
```

Suppose also that the column values are retrieved using the following statement:

```
SELECT X FROM T ORDER BY X COLLATE collation_name;
```

The following table shows the resulting order of the values if we use `ORDER BY` with different collations.

<code>latin1_swedish_ci</code>	<code>latin1_german1_ci</code>	<code>latin1_german2_ci</code>
Muffler	Muffler	Müller
MX Systems	Müller	Muffler
Müller	MX Systems	MX Systems
MySQL	MySQL	MySQL

The character that causes the different sort orders in this example is the U with two dots over it (ü), which the Germans call “U-umlaut.”

- The first column shows the result of the `SELECT` using the Swedish/Finnish collating rule, which says that U-umlaut sorts with Y.
- The second column shows the result of the `SELECT` using the German DIN-1 rule, which says that U-umlaut sorts with U.
- The third column shows the result of the `SELECT` using the German DIN-2 rule, which says that U-umlaut sorts with UE.

#### Example 2: Searching for German Umlauts

Suppose that you have three tables that differ only by the character set and collation used:

```
mysql> SET NAMES utf8;
mysql> CREATE TABLE german1 (
  ->   c CHAR(10)
  -> ) CHARACTER SET latin1 COLLATE latin1_german1_ci;
mysql> CREATE TABLE german2 (
  ->   c CHAR(10)
  -> ) CHARACTER SET latin1 COLLATE latin1_german2_ci;
mysql> CREATE TABLE germanutf8 (
  ->   c CHAR(10)
  -> ) CHARACTER SET utf8 COLLATE utf8_unicode_ci;
```

Each table contains two records:

```
mysql> INSERT INTO german1 VALUES ('Bar'), ('Bär');
mysql> INSERT INTO german2 VALUES ('Bar'), ('Bär');
```

```
mysql> INSERT INTO germanutf8 VALUES ('Bar'), ('Bär');
```

Two of the above collations have an `A = Ä` equality, and one has no such equality (`latin1_german2_ci`). For that reason, you'll get these results in comparisons:

```
mysql> SELECT * FROM german1 WHERE c = 'Bär';
+-----+
| c     |
+-----+
| Bar   |
| Bär   |
+-----+
mysql> SELECT * FROM german2 WHERE c = 'Bär';
+-----+
| c     |
+-----+
| Bär   |
+-----+
mysql> SELECT * FROM germanutf8 WHERE c = 'Bär';
+-----+
| c     |
+-----+
| Bar   |
| Bär   |
+-----+
```

This is not a bug but rather a consequence of the sorting properties of `latin1_german1_ci` and `utf8_unicode_ci` (the sorting shown is done according to the German DIN 5007 standard).

### 10.1.7.9. Collation and `INFORMATION_SCHEMA` Searches

String columns in `INFORMATION_SCHEMA` tables have a collation of `utf8_general_ci`, which is case insensitive. However, searches in `INFORMATION_SCHEMA` string columns are also affected by file system case sensitivity. For values that correspond to objects that are represented in the file system, such as names of databases and tables, searches may be case sensitive if the file system is case sensitive. This section describes how to work around this issue if necessary; see also Bug #34921.

Suppose that a query searches the `SCHEMATA.SCHEMA_NAME` column for the `test` database. On Linux, file systems are case sensitive, so comparisons of `SCHEMATA.SCHEMA_NAME` with `'test'` match, but comparisons with `'TEST'` do not:

```
mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
-> WHERE SCHEMA_NAME = 'test';
+-----+
| SCHEMA_NAME |
+-----+
| test        |
+-----+
1 row in set (0.01 sec)

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
-> WHERE SCHEMA_NAME = 'TEST';
Empty set (0.00 sec)
```

On Windows or Mac OS X where file systems are not case sensitive, comparisons match both `'test'` and `'TEST'`:

```
mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
-> WHERE SCHEMA_NAME = 'test';
+-----+
| SCHEMA_NAME |
+-----+
| test        |
+-----+
1 row in set (0.00 sec)

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
-> WHERE SCHEMA_NAME = 'TEST';
```

```

+-----+
| SCHEMA_NAME |
+-----+
| TEST        |
+-----+
1 row in set (0.00 sec)

```

The value of the `lower_case_table_names` [501] system variable makes no difference in this context.

This behavior occurs because the `utf8_general_ci` collation is not used for `INFORMATION_SCHEMA` queries when searching the file system for database objects. It is a result of optimizations implemented for `INFORMATION_SCHEMA` searches in MySQL. For information about these optimizations, see [Section 8.2.4, “Optimizing INFORMATION\\_SCHEMA Queries”](#).

Searches in `INFORMATION_SCHEMA` string columns for values that refer to `INFORMATION_SCHEMA` itself do use the `utf8_general_ci` collation because `INFORMATION_SCHEMA` is a “virtual” database and is not represented in the file system. For example, comparisons with `SCHEMATA.SCHEMA_NAME` match `'information_schema'` or `'INFORMATION_SCHEMA'` regardless of platform:

```

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
      -> WHERE SCHEMA_NAME = 'information_schema';
+-----+
| SCHEMA_NAME |
+-----+
| information_schema |
+-----+
1 row in set (0.00 sec)

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
      -> WHERE SCHEMA_NAME = 'INFORMATION_SCHEMA';
+-----+
| SCHEMA_NAME |
+-----+
| information_schema |
+-----+
1 row in set (0.00 sec)

```

If the result of a string operation on an `INFORMATION_SCHEMA` column differs from expectations, a workaround is to use an explicit `COLLATE` clause to force a suitable collation ([Section 10.1.7.2, “Using COLLATE in SQL Statements”](#)). For example, to perform a case-insensitive search, use `COLLATE` with the `INFORMATION_SCHEMA` column name:

```

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
      -> WHERE SCHEMA_NAME COLLATE utf8_general_ci = 'test';
+-----+
| SCHEMA_NAME |
+-----+
| test        |
+-----+
1 row in set (0.00 sec)

mysql> SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA
      -> WHERE SCHEMA_NAME COLLATE utf8_general_ci = 'TEST';
| SCHEMA_NAME |
+-----+
| test        |
+-----+
1 row in set (0.00 sec)

```

You can also use the `UPPER()` [1187] or `LOWER()` [1182] function:

```

WHERE UPPER(SCHEMA_NAME) = 'TEST'
WHERE LOWER(SCHEMA_NAME) = 'test'

```

Although a case-insensitive comparison can be performed even on platforms with case-sensitive file systems, as just shown, it is not necessarily always the right thing to do. On such platforms, it is possible to have multiple objects with names that differ only in lettercase. For example, tables named

`city`, `CITY`, and `City` can all exist simultaneously. Consider whether a search should match all such names or just one and write queries accordingly:

```
WHERE TABLE_NAME COLLATE utf8_bin = 'City'
WHERE TABLE_NAME COLLATE utf8_general_ci = 'city'
WHERE UPPER(TABLE_NAME) = 'CITY'
WHERE LOWER(TABLE_NAME) = 'city'
```

The first of those comparisons (with `utf8_bin`) is case sensitive; the others are not.

## 10.1.8. String Repertoire

The *repertoire* of a character set is the collection of characters in the set.

String expressions have a repertoire attribute, which can have two values:

- **ASCII**: The expression can contain only characters in the Unicode range `U+0000` to `U+007F`.
- **UNICODE**: The expression can contain characters in the Unicode range `U+0000` to `U+FFFF`.

The **ASCII** range is a subset of **UNICODE** range, so a string with **ASCII** repertoire can be converted safely without loss of information to the character set of any string with **UNICODE** repertoire or to a character set that is a superset of **ASCII**. (All MySQL character sets are supersets of **ASCII** with the exception of `swe7`, which reuses some punctuation characters for Swedish accented characters.) The use of repertoire enables character set conversion in expressions for many cases where MySQL would otherwise return an “illegal mix of collations” error.

The following discussion provides examples of expressions and their repertoires, and describes how the use of repertoire changes string expression evaluation:

- The repertoire for string constants depends on string content:

```
SET NAMES utf8; SELECT 'abc';
SELECT _utf8'def';
SELECT N'MySQL';
```

Although the character set is `utf8` in each of the preceding cases, the strings do not actually contain any characters outside the ASCII range, so their repertoire is **ASCII** rather than **UNICODE**.

- Columns having the `ascii` character set have **ASCII** repertoire because of their character set. In the following table, `c1` has **ASCII** repertoire:

```
CREATE TABLE t1 (c1 CHAR(1) CHARACTER SET ascii);
```

The following example illustrates how repertoire enables a result to be determined in a case where an error occurs without repertoire:

```
CREATE TABLE t1 (
  c1 CHAR(1) CHARACTER SET latin1,
  c2 CHAR(1) CHARACTER SET ascii
);
INSERT INTO t1 VALUES ('a','b');
SELECT CONCAT(c1,c2) FROM t1;
```

Without repertoire, this error occurs:

```
ERROR 1267 (HY000): Illegal mix of collations (latin1_swedish_ci,IMPLICIT)
and (ascii_general_ci,IMPLICIT) for operation 'concat'
```

Using repertoire, subset to superset (`ascii` to `latin1`) conversion can occur and a result is returned:

```
+-----+
| CONCAT(c1,c2) |
+-----+
```

```
| ab |
+-----+
```

- Functions with one string argument inherit the repertoire of their argument. The result of `UPPER(_utf8'abc')` [1187] has `ASCII` repertoire because its argument has `ASCII` repertoire.
- For functions that return a string but do not have string arguments and use `character_set_connection` [467] as the result character set, the result repertoire is `ASCII` if `character_set_connection` [467] is `ascii`, and `UNICODE` otherwise:

```
FORMAT(numeric_column, 4);
```

Use of repertoire changes how MySQL evaluates the following example:

```
SET NAMES ascii;
CREATE TABLE t1 (a INT, b VARCHAR(10) CHARACTER SET latin1);
INSERT INTO t1 VALUES (1,'b');
SELECT CONCAT(FORMAT(a, 4), b) FROM t1;
```

Without repertoire, this error occurs:

```
ERROR 1267 (HY000): Illegal mix of collations (ascii_general_ci,COERCIBLE)
and (latin1_swedish_ci,IMPLICIT) for operation 'concat'
```

With repertoire, a result is returned:

```
+-----+
| CONCAT(FORMAT(a, 4), b) |
+-----+
| 1.0000b                |
+-----+
```

- Functions with two or more string arguments use the “widest” argument repertoire for the result repertoire (`UNICODE` is wider than `ASCII`). Consider the following `CONCAT()` [1178] calls:

```
CONCAT(_ucs2 0x0041, _ucs2 0x0042)
CONCAT(_ucs2 0x0041, _ucs2 0x00C2)
```

For the first call, the repertoire is `ASCII` because both arguments are within the range of the `ascii` character set. For the second call, the repertoire is `UNICODE` because the second argument is outside the `ascii` character set range.

- The repertoire for function return values is determined based only on the repertoire of the arguments that affect the result's character set and collation.

```
IF(column1 < column2, 'smaller', 'greater')
```

The result repertoire is `ASCII` because the two string arguments (the second argument and the third argument) both have `ASCII` repertoire. The first argument does not matter for the result repertoire, even if the expression uses string values.

## 10.1.9. Operations Affected by Character Set Support

This section describes operations that take character set information into account.

### 10.1.9.1. Result Strings

MySQL has many operators and functions that return a string. This section answers the question: What is the character set and collation of such a string?

For simple functions that take string input and return a string result as output, the output's character set and collation are the same as those of the principal input value. For example, `UPPER(X)` [1187] returns a string whose character string and collation are the same as that of `X`. The same applies for `INSTR()` [1181], `LCASE()` [1181], `LOWER()` [1182], `LTRIM()` [1182], `MID()` [1183],

`REPEAT()` [1183], `REPLACE()` [1183], `REVERSE()` [1184], `RIGHT()` [1184], `RPAD()` [1184], `RTRIM()` [1184], `SOUNDEX()` [1184], `SUBSTRING()` [1185], `TRIM()` [1186], `UCASE()` [1186], and `UPPER()` [1187].

Note: The `REPLACE()` [1183] function, unlike all other functions, always ignores the collation of the string input and performs a case-sensitive comparison.

If a string input or function result is a binary string, the string has no character set or collation. This can be checked by using the `CHARSET()` [1266] and `COLLATION()` [1267] functions, both of which return `binary` to indicate that their argument is a binary string:

```
mysql> SELECT CHARSET(BINARY 'a'), COLLATION(BINARY 'a');
+-----+-----+
| CHARSET(BINARY 'a') | COLLATION(BINARY 'a') |
+-----+-----+
| binary              | binary                 |
+-----+-----+
```

For operations that combine multiple string inputs and return a single string output, the “aggregation rules” of standard SQL apply for determining the collation of the result:

- If an explicit `COLLATE X` occurs, use `X`.
- If explicit `COLLATE X` and `COLLATE Y` occur, raise an error.
- Otherwise, if all collations are `X`, use `X`.
- Otherwise, the result has no collation.

For example, with `CASE ... WHEN a THEN b WHEN b THEN c COLLATE X END`, the resulting collation is `X`. The same applies for `UNION, ||` [1172], `CONCAT()` [1178], `ELT()` [1179], `GREATEST()` [1169], `IF()` [1174], and `LEAST()` [1170].

For operations that convert to character data, the character set and collation of the strings that result from the operations are defined by the `character_set_connection` [467] and `collation_connection` [469] system variables. This applies only to `CAST()` [1245], `CONV()` [1203], `FORMAT()` [1180], `HEX()` [1180], and `SPACE()` [1185].

If you are uncertain about the character set or collation of the result returned by a string function, you can use the `CHARSET()` [1266] or `COLLATION()` [1267] function to find out:

```
mysql> SELECT USER(), CHARSET(USER()), COLLATION(USER());
+-----+-----+-----+
| USER()          | CHARSET(USER()) | COLLATION(USER()) |
+-----+-----+-----+
| test@localhost | utf8            | utf8_general_ci  |
+-----+-----+-----+
```

### 10.1.9.2. `CONVERT()` and `CAST()`

`CONVERT()` [1245] provides a way to convert data between different character sets. The syntax is:

```
CONVERT(expr USING transcoding_name)
```

In MySQL, transcoding names are the same as the corresponding character set names.

Examples:

```
SELECT CONVERT(_latin1'Müller' USING utf8);
INSERT INTO utf8table (utf8column)
  SELECT CONVERT(latin1field USING utf8) FROM latin1table;
```

`CONVERT(... USING ...)` [1245] is implemented according to the standard SQL specification.

You may also use `CAST()` [1245] to convert a string to a different character set. The syntax is:



```
CAST(character_string AS character_data_type CHARACTER SET charset_name)
```

Example:

```
SELECT CAST(_latin1'test' AS CHAR CHARACTER SET utf8);
```

If you use `CAST()` [1245] without specifying `CHARACTER SET`, the resulting character set and collation are defined by the `character_set_connection` [467] and `collation_connection` [469] system variables. If you use `CAST()` [1245] with `CHARACTER SET X`, the resulting character set and collation are `X` and the default collation of `X`.

You may not use a `COLLATE` clause inside a `CONVERT()` [1245] or `CAST()` [1245] call, but you may use it outside. For example, `CAST(... COLLATE ...)` [1245] is illegal, but `CAST(...)  
COLLATE ...` [1245] is legal:

```
SELECT CAST(_latin1'test' AS CHAR CHARACTER SET utf8) COLLATE utf8_bin;
```

### 10.1.9.3. SHOW Statements and INFORMATION\_SCHEMA

Several `SHOW` statements provide additional character set information. These include `SHOW CHARACTER SET`, `SHOW COLLATION`, `SHOW CREATE DATABASE`, `SHOW CREATE TABLE` and `SHOW COLUMNS`. These statements are described here briefly. For more information, see [Section 13.7.5, “SHOW Syntax”](#).

`INFORMATION_SCHEMA` has several tables that contain information similar to that displayed by the `SHOW` statements. For example, the `CHARACTER_SETS` and `COLLATIONS` tables contain the information displayed by `SHOW CHARACTER SET` and `SHOW COLLATION`. See [Chapter 19, INFORMATION\\_SCHEMA Tables](#).

The `SHOW CHARACTER SET` statement shows all available character sets. It takes an optional `LIKE` [1190] clause that indicates which character set names to match. For example:

```
mysql> SHOW CHARACTER SET LIKE 'latin%';
```

Charset	Description	Default collation	Maxlen
latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
latin5	ISO 8859-9 Turkish	latin5_turkish_ci	1
latin7	ISO 8859-13 Baltic	latin7_general_ci	1

The output from `SHOW COLLATION` includes all available character sets. It takes an optional `LIKE` [1190] clause that indicates which collation names to match. For example:

```
mysql> SHOW COLLATION LIKE 'latin1%';
```

Collation	Charset	Id	Default	Compiled	Sortlen
latin1_german1_ci	latin1	5			0
latin1_swedish_ci	latin1	8	Yes	Yes	0
latin1_danish_ci	latin1	15			0
latin1_german2_ci	latin1	31		Yes	2
latin1_bin	latin1	47		Yes	0
latin1_general_ci	latin1	48			0
latin1_general_cs	latin1	49			0
latin1_spanish_ci	latin1	94			0

`SHOW CREATE DATABASE` displays the `CREATE DATABASE` statement that creates a given database:

```
mysql> SHOW CREATE DATABASE test;
```

Database	Create Database
test	CREATE DATABASE test



Before MySQL 5.5	MySQL 5.5 and up
All Unicode 3.0 characters	All Unicode 5.0 and 6.0 characters
No supplementary characters	With supplementary characters
<code>ucs2</code> character set, BMP only	No change
<code>utf8</code> character set for up to three bytes, BMP only	No change
	New <code>utf8mb4</code> character set for up to four bytes, BMP or supplemental
	New <code>utf16</code> character set, BMP or supplemental
	New <code>utf16le</code> character set, BMP or supplemental (5.6.1 and up)
	New <code>utf32</code> character set, BMP or supplemental

These changes are upward compatible. If you want to use the new character sets, there are potential incompatibility issues for your applications; see [Section 10.1.11, “Upgrading from Previous to Current Unicode Support”](#). That section also describes how to convert tables from `utf8` to the (4-byte) `utf8mb4` character set, and what constraints may apply in doing so.

MySQL 5.6 supports these Unicode character sets:

- `ucs2`, the UCS-2 encoding of the Unicode character set using 16 bits per character.
- `utf16`, the UTF-16 encoding for the Unicode character set; like `ucs2` but with an extension for supplementary characters.
- `utf16le`, the UTF-16LE encoding for the Unicode character set; like `utf16` but little-endian rather than big-endian.
- `utf32`, the UTF-32 encoding for the Unicode character set using 32 bits per character.
- `utf8`, a UTF-8 encoding of the Unicode character set using one to three bytes per character.
- `utf8mb4`, a UTF-8 encoding of the Unicode character set using one to four bytes per character.

`ucs2` and `utf8` support BMP characters. `utf8mb4`, `utf16`, `utf16le`, and `utf32` support BMP and supplementary characters.

A similar set of collations is available for most Unicode character sets. For example, each has a Danish collation, the names of which are `ucs2_danish_ci`, `utf16_danish_ci`, `utf32_danish_ci`, `utf8_danish_ci`, and `utf8mb4_danish_ci`. The exception is `utf16le`, which has only two collations. All Unicode collations are listed at [Section 10.1.14.1, “Unicode Character Sets”](#), which also describes collation properties for supplementary characters.

Note that although many of the supplementary characters come from East Asian languages, what MySQL 5.6 adds is support for more Japanese and Chinese characters in Unicode character sets, not support for new Japanese and Chinese character sets.

The MySQL implementation of UCS-2, UTF-16, and UTF-32 stores characters in big-endian byte order and does not use a byte order mark (BOM) at the beginning of values. Other database systems might use little-endian byte order or a BOM. In such cases, conversion of values will need to be performed when transferring data between those systems and MySQL. The implementation of UTF-16LE is little-endian.

MySQL uses no BOM for UTF-8 values.

Client applications that need to communicate with the server using Unicode should set the client character set accordingly; for example, by issuing a `SET NAMES 'utf8'` statement. `ucs2`, `utf16`,

`utf16le`, and `utf32` cannot be used as a client character set, which means that they do not work for `SET NAMES` or `SET CHARACTER SET`. (See [Section 10.1.4, “Connection Character Sets and Collations”](#).)

The following sections provide additional detail on the Unicode character sets in MySQL.

### 10.1.10.1. The `ucs2` Character Set (UCS-2 Unicode Encoding)

In UCS-2, every character is represented by a 2-byte Unicode code with the most significant byte first. For example: `LATIN CAPITAL LETTER A` has the code `0x0041` and it is stored as a 2-byte sequence: `0x00 0x41`. `CYRILLIC SMALL LETTER YERU` (Unicode `0x044B`) is stored as a 2-byte sequence: `0x04 0x4B`. For Unicode characters and their codes, please refer to the [Unicode Home Page](#).

In MySQL, the `ucs2` character set is a fixed-length 16-bit encoding for Unicode BMP characters.

### 10.1.10.2. The `utf16` Character Set (UTF-16 Unicode Encoding)

The `utf16` character set is the `ucs2` character set with an extension that enables encoding of supplementary characters:

- For a BMP character, `utf16` and `ucs2` have identical storage characteristics: same code values, same encoding, same length.
- For a supplementary character, `utf16` has a special sequence for representing the character using 32 bits. This is called the “surrogate” mechanism: For a number greater than `0xffff`, take 10 bits and add them to `0xd800` and put them in the first 16-bit word, take 10 more bits and add them to `0xdc00` and put them in the next 16-bit word. Consequently, all supplementary characters require 32 bits, where the first 16 bits are a number between `0xd800` and `0xdbff`, and the last 16 bits are a number between `0xdc00` and `0xdfff`. Examples are in [Section 15.5 Surrogates Area](#) of the Unicode 4.0 document.

Because `utf16` supports surrogates and `ucs2` does not, there is a validity check that applies only in `utf16`: You cannot insert a top surrogate without a bottom surrogate, or vice versa. For example:

```
INSERT INTO t (ucs2_column) VALUES (0xd800); /* legal */
INSERT INTO t (utf16_column)VALUES (0xd800); /* illegal */
```

There is no validity check for characters that are technically valid but are not true Unicode (that is, characters that Unicode considers to be “unassigned code points” or “private use” characters or even “illegals” like `0xffff`). For example, since `U+F8FF` is the Apple Logo, this is legal:

```
INSERT INTO t (utf16_column)VALUES (0xf8ff); /* legal */
```

Such characters cannot be expected to mean the same thing to everyone.

Because MySQL must allow for the worst case (that one character requires four bytes) the maximum length of a `utf16` column or index is only half of the maximum length for a `ucs2` column or index. For example, in MySQL 5.6, the maximum length of a `MEMORY` table index key is 3072 bytes, so these statements create tables with the longest permitted indexes for `ucs2` and `utf16` columns:

```
CREATE TABLE tf (s1 VARCHAR(1536) CHARACTER SET ucs2) ENGINE=MEMORY;
CREATE INDEX i ON tf (s1);
CREATE TABLE tg (s1 VARCHAR(768) CHARACTER SET utf16) ENGINE=MEMORY;
CREATE INDEX i ON tg (s1);
```

### 10.1.10.3. The `utf16le` Character Set (UTF-16LE Unicode Encoding)

This is the same as `utf16` but is little-endian rather than big-endian.

### 10.1.10.4. The `utf32` Character Set (UTF-32 Unicode Encoding)

The `utf32` character set is fixed length (like `ucs2` and unlike `utf16`). `utf32` uses 32 bits for every character, unlike `ucs2` (which uses 16 bits for every character), and unlike `utf16` (which uses 16 bits for some characters and 32 bits for others).

`utf32` takes twice as much space as `ucs2` and more space than `utf16`, but `utf32` has the same advantage as `ucs2` that it is predictable for storage: The required number of bytes for `utf32` equals the number of characters times 4. Also, unlike `utf16`, there are no tricks for encoding in `utf32`, so the stored value equals the code value.

To demonstrate how the latter advantage is useful, here is an example that shows how to determine a `utf8mb4` value given the `utf32` code value:

```
/* Assume code value = 100cc LINEAR B WHEELED CHARIOT */
CREATE TABLE tmp (utf32_col CHAR(1) CHARACTER SET utf32,
                  utf8mb4_col CHAR(1) CHARACTER SET utf8mb4);
INSERT INTO tmp VALUES (0x000100cc,NULL);
UPDATE tmp SET utf8mb4_col = utf32_col;
SELECT HEX(utf32_col),HEX(utf8mb4_col) FROM tmp;
```

MySQL is very forgiving about additions of unassigned Unicode characters or private-use-area characters. There is in fact only one validity check for `utf32`: No code value may be greater than `0x10ffff`. For example, this is illegal:

```
INSERT INTO t (utf32_column) VALUES (0x110000); /* illegal */
```

### 10.1.10.5. The `utf8` Character Set (3-Byte UTF-8 Unicode Encoding)

UTF-8 (Unicode Transformation Format with 8-bit units) is an alternative way to store Unicode data. It is implemented according to RFC 3629, which describes encoding sequences that take from one to four bytes. (An older standard for UTF-8 encoding, RFC 2279, describes UTF-8 sequences that take from one to six bytes. RFC 3629 renders RFC 2279 obsolete; for this reason, sequences with five and six bytes are no longer used.)

The idea of UTF-8 is that various Unicode characters are encoded using byte sequences of different lengths:

- Basic Latin letters, digits, and punctuation signs use one byte.
- Most European and Middle East script letters fit into a 2-byte sequence: extended Latin letters (with tilde, macron, acute, grave and other accents), Cyrillic, Greek, Armenian, Hebrew, Arabic, Syriac, and others.
- Korean, Chinese, and Japanese ideographs use 3-byte or 4-byte sequences.

The `utf8` character set is the same in MySQL 5.6 as before 5.6 and has exactly the same characteristics:

- No support for supplementary characters (BMP characters only).
- A maximum of three bytes per multi-byte character.

Exactly the same set of characters is available in `utf8` as in `ucs2`. That is, they have the same repertoire.

### 10.1.10.6. The `utf8mb3` “Character Set” (Alias for `utf8`)

In a future version of MySQL, it is possible that `utf8` will become the 4-byte `utf8`, and that users who want to indicate 3-byte `utf8` will have to say `utf8mb3`. To avoid some future problems which might occur with replication when master and slave servers have different MySQL versions, it is possible for users to specify `utf8mb3` in `CHARACTER SET` clauses, and `utf8mb3_collation_substring` in `COLLATE` clauses, where `collation_substring` is `bin`, `czech_ci`, `danish_ci`, `esperanto_ci`, `estonian_ci`, and so forth. For example:

```
CREATE TABLE t (s1 CHAR(1) CHARACTER SET utf8mb3;
SELECT * FROM t WHERE s1 COLLATE utf8mb3_general_ci = 'x';
DECLARE x VARCHAR(5) CHARACTER SET utf8mb3 COLLATE utf8mb3_danish_ci;
SELECT CAST('a' AS CHAR CHARACTER SET utf8) COLLATE utf8_czech_ci;
```

MySQL immediately converts instances of `utf8mb3` in an alias to `utf8`, so in statements such as `SHOW CREATE TABLE` or `SELECT CHARACTER_SET_NAME FROM INFORMATION_SCHEMA.COLUMNS` or `SELECT COLLATION_NAME FROM INFORMATION_SCHEMA.COLUMNS`, users will see the true name, `utf8` or `utf8_collation_substring`.

The `utf8mb3` alias is valid only in `CHARACTER SET` clauses, and in certain other places. For example, these are legal:

```
mysql> --character-set-server=utf8mb3
SET NAMES 'utf8mb3'; /* and other SET statements that have similar effect */
SELECT _utf8mb3 'a';
```

There is no `utf8mb3` alias to the corresponding `utf8` collation for collation names that include a version number (for example, `utf8_unicode_520_ci`) to indicate the Unicode Collation Algorithm version on which the collation is based.

### 10.1.10.7. The `utf8mb4` Character Set (4-Byte UTF-8 Unicode Encoding)

The character set named `utf8` uses a maximum of three bytes per character and contains only BMP characters. The `utf8mb4` character set uses a maximum of four bytes per character supports supplemental characters:

- For a BMP character, `utf8` and `utf8mb4` have identical storage characteristics: same code values, same encoding, same length.
- For a supplementary character, `utf8` cannot store the character at all, while `utf8mb4` requires four bytes to store it. Since `utf8` cannot store the character at all, you do not have any supplementary characters in `utf8` columns and you need not worry about converting characters or losing data when upgrading `utf8` data from older versions of MySQL.

`utf8mb4` is a superset of `utf8`, so for an operation such as the following concatenation, the result has character set `utf8mb4` and the collation of `utf8mb4_col`:

```
SELECT CONCAT(utf8_col, utf8mb4_col);
```

Similarly, the following comparison in the `WHERE` clause works according to the collation of `utf8mb_col`:

```
SELECT * FROM utf8_tbl, utf8mb4_tbl
WHERE utf8_tbl.utf8_col = utf8mb4_tbl.utf8mb4_col;
```



#### Tip

To save space with UTF-8, use `VARCHAR` instead of `CHAR`. Otherwise, MySQL must reserve three (or four) bytes for each character in a `CHAR CHARACTER SET utf8` (or `utf8mb4`) column because that is the maximum possible length. For example, MySQL must reserve 40 bytes for a `CHAR(10) CHARACTER SET utf8mb4` column.

### 10.1.11. Upgrading from Previous to Current Unicode Support

This section describes issues pertaining to Unicode support that you may face when upgrading to MySQL 5.6 from an older MySQL release. It also provides guidelines for downgrading from MySQL 5.6 to an older release.

In most respects, upgrading to MySQL 5.6 should present few problems with regard to Unicode usage, although there are some potential areas of incompatibility. These are the primary areas of concern:

- For the variable-length character data types ([VARCHAR](#) and the [TEXT](#) types), the maximum length in characters is less for [utf8mb4](#) columns than for [utf8](#) columns.
- For all character data types ([CHAR](#), [VARCHAR](#), and the [TEXT](#) types), the maximum number of characters that can be indexed is less for [utf8mb4](#) columns than for [utf8](#) columns.

Consequently, if you want to upgrade tables from [utf8](#) to [utf8mb4](#) to take advantage of supplementary-character support, it may be necessary to change some column or index definitions.

Tables can be converted from [utf8](#) to [utf8mb4](#) by using [ALTER TABLE](#). Suppose that a table was originally defined as follows:

```
CREATE TABLE t1 (
  col1 CHAR(10) CHARACTER SET utf8 COLLATE utf8_unicode_ci NOT NULL,
  col2 CHAR(10) CHARACTER SET utf8 COLLATE utf8_bin NOT NULL
) CHARACTER SET utf8;
```

The following statement converts [t1](#) to use [utf8mb4](#):

```
ALTER TABLE t1
  DEFAULT CHARACTER SET utf8mb4,
  MODIFY col1 CHAR(10)
  CHARACTER SET utf8mb4 COLLATE utf8mb4_unicode_ci NOT NULL,
  MODIFY col2 CHAR(10)
  CHARACTER SET utf8mb4 COLLATE utf8mb4_bin NOT NULL;
```

In terms of table content, conversion from [utf8](#) to [utf8mb4](#) presents no problems:

- For a BMP character, [utf8](#) and [utf8mb4](#) have identical storage characteristics: same code values, same encoding, same length.
- For a supplementary character, [utf8](#) cannot store the character at all, while [utf8mb4](#) requires four bytes to store it. Since [utf8](#) cannot store the character at all, you do not have any supplementary characters in [utf8](#) columns and you need not worry about converting characters or losing data when upgrading [utf8](#) data from older versions of MySQL.

In terms of table structure, the catch when converting from [utf8](#) to [utf8mb4](#) is that the maximum length of a column or index key is unchanged in terms of *bytes*. Therefore, it is smaller in terms of *characters* because the maximum length of a character is four bytes instead of three. For the [CHAR](#), [VARCHAR](#), and [TEXT](#) data types, watch for these things when converting your MySQL tables:

- Check all definitions of [utf8](#) columns and make sure they will not exceed the maximum length for the storage engine.
- Check all indexes on [utf8](#) columns and make sure they will not exceed the maximum length for the storage engine. Sometimes the maximum can change due to storage engine enhancements.

If the preceding conditions apply, you must either reduce the defined length of columns or indexes, or continue to use [utf8](#) rather than [utf8mb4](#).

Here are some examples where structural changes may be needed:

- A [TINYTEXT](#) column can hold up to 255 bytes, so it can hold up to 85 3-byte or 63 4-byte characters. Suppose that you have a [TINYTEXT](#) column that uses [utf8](#) but must be able to contain more than 63 characters. You cannot convert it to [utf8mb4](#) unless you also change the data type to a longer type such as [TEXT](#).

Similarly, a very long [VARCHAR](#) column may need to be changed to one of the longer [TEXT](#) types if you want to convert it from [utf8](#) to [utf8mb4](#).

- [InnoDB](#) has a maximum index length of 767 bytes, so for [utf8](#) or [utf8mb4](#) columns, you can index a maximum of 255 or 191 characters, respectively. If you currently have [utf8](#) columns with indexes longer than 191 characters, you will need to index a smaller number of characters. In an [InnoDB](#) table, these column and index definitions are legal:



```
coll VARCHAR(500) CHARACTER SET utf8, INDEX (coll(255))
```

To use `utf8mb4` instead, the index must be smaller:

```
coll VARCHAR(500) CHARACTER SET utf8mb4, INDEX (coll(191))
```

The preceding types of changes are most likely to be required only if you have very long columns or indexes. Otherwise, you should be able to convert your tables from `utf8` to `utf8mb4` without problems. You can do this by using `ALTER TABLE` as described earlier in this section after upgrading in place to 5.6.

The following items summarize other potential areas of incompatibility:

- Performance of 4-byte UTF-8 (`utf8mb4`) is slower than for 3-byte UTF-8 (`utf8`). If you do not want to incur this penalty, continue to use `utf8`.
- `SET NAMES 'utf8mb4'` causes use of the 4-byte character set for connection character sets. As long as no 4-byte characters are sent from the server, there should be no problems. Otherwise, applications that expect to receive a maximum of three bytes per character may have problems. Conversely, applications that expect to send 4-byte characters must ensure that the server understands them.
- Applications cannot send `utf16`, `utf16le`, or `utf32` character data to an older server that does not understand them.
- For replication, if the character sets that support supplementary characters are going to be used on the master, all slaves must understand them as well. If you attempt to replicate from a MySQL 5.6 master to an older slave, `utf8` data will be seen as `utf8` by the slave and should replicate correctly. But you cannot send `utf8mb4`, `utf16`, `utf16le`, or `utf32` data.

Also, keep in mind the general principle that if a table has different definitions on the master and slave, this can lead to unexpected results. For example, the differences in limitations on index key length makes it risky to use `utf8` on the master and `utf8mb4` on the slave.

If you have upgraded to MySQL 5.6, and then decide to downgrade back to an older release, these considerations apply:

- `ucs2` and `utf8` data should present no problems.
- Any definitions that refer to the `utf8mb4`, `utf16`, `utf16le`, or `utf32` character sets will not be recognized by the older server.
- For object definitions that refer to the `utf8mb4` character set, you can dump them with `mysqldump` in MySQL 5.6, edit the dump file to change instances of `utf8mb4` to `utf8`, and reload the file in the older server, as long as there are no 4-byte characters in the data. The older server will see `utf8` in the dump file object definitions and create new objects that use the (3-byte) `utf8` character set.

## 10.1.12. UTF-8 for Metadata

*Metadata* is “the data about the data.” Anything that *describes* the database—as opposed to being the *contents* of the database—is metadata. Thus column names, database names, user names, version names, and most of the string results from `SHOW` are metadata. This is also true of the contents of tables in `INFORMATION_SCHEMA` because those tables by definition contain information about database objects.

Representation of metadata must satisfy these requirements:

- All metadata must be in the same character set. Otherwise, neither the `SHOW` statements nor `SELECT` statements for tables in `INFORMATION_SCHEMA` would work properly because different rows in the same column of the results of these operations would be in different character sets.

- Metadata must include all characters in all languages. Otherwise, users would not be able to name columns and tables using their own languages.

To satisfy both requirements, MySQL stores metadata in a Unicode character set, namely UTF-8. This does not cause any disruption if you never use accented or non-Latin characters. But if you do, you should be aware that metadata is in UTF-8.

The metadata requirements mean that the return values of the `USER()` [1273], `CURRENT_USER()` [1267], `SESSION_USER()` [1273], `SYSTEM_USER()` [1273], `DATABASE()` [1268], and `VERSION()` [1273] functions have the UTF-8 character set by default.

The server sets the `character_set_system` [469] system variable to the name of the metadata character set:

```
mysql> SHOW VARIABLES LIKE 'character_set_system';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_system | utf8 |
+-----+-----+
```

Storage of metadata using Unicode does *not* mean that the server returns headers of columns and the results of `DESCRIBE` functions in the `character_set_system` [469] character set by default. When you use `SELECT column1 FROM t`, the name `column1` itself is returned from the server to the client in the character set determined by the value of the `character_set_results` [468] system variable, which has a default value of `latin1`. If you want the server to pass metadata results back in a different character set, use the `SET NAMES` statement to force the server to perform character set conversion. `SET NAMES` sets the `character_set_results` [468] and other related system variables. (See Section 10.1.4, “Connection Character Sets and Collations”.) Alternatively, a client program can perform the conversion after receiving the result from the server. It is more efficient for the client to perform the conversion, but this option is not always available for all clients.

If `character_set_results` [468] is set to `NULL`, no conversion is performed and the server returns metadata using its original character set (the set indicated by `character_set_system` [469]).

Error messages returned from the server to the client are converted to the client character set automatically, as with metadata.

If you are using (for example) the `USER()` [1273] function for comparison or assignment within a single statement, don't worry. MySQL performs some automatic conversion for you.

```
SELECT * FROM t1 WHERE USER() = latin1_column;
```

This works because the contents of `latin1_column` are automatically converted to UTF-8 before the comparison.

```
INSERT INTO t1 (latin1_column) SELECT USER();
```

This works because the contents of `USER()` [1273] are automatically converted to `latin1` before the assignment.

Although automatic conversion is not in the SQL standard, the SQL standard document does say that every character set is (in terms of supported characters) a “subset” of Unicode. Because it is a well-known principle that “what applies to a superset can apply to a subset,” we believe that a collation for Unicode can apply for comparisons with non-Unicode strings. For more information about coercion of strings, see Section 10.1.7.5, “Collation of Expressions”.

### 10.1.13. Column Character Set Conversion

To convert a binary or nonbinary string column to use a particular character set, use `ALTER TABLE`. For successful conversion to occur, one of the following conditions must apply:

- If the column has a binary data type (`BINARY`, `VARBINARY`, `BLOB`), all the values that it contains must be encoded using a single character set (the character set you're converting the column to). If you use a binary column to store information in multiple character sets, MySQL has no way to know which values use which character set and cannot convert the data properly.
- If the column has a nonbinary data type (`CHAR`, `VARCHAR`, `TEXT`), its contents should be encoded in the column character set, not some other character set. If the contents are encoded in a different character set, you can convert the column to use a binary data type first, and then to a nonbinary column with the desired character set.

Suppose that a table `t` has a binary column named `coll` defined as `VARBINARY(50)`. Assuming that the information in the column is encoded using a single character set, you can convert it to a nonbinary column that has that character set. For example, if `coll` contains binary data representing characters in the `greek` character set, you can convert it as follows:

```
ALTER TABLE t MODIFY coll VARCHAR(50) CHARACTER SET greek;
```

If your original column has a type of `BINARY(50)`, you could convert it to `CHAR(50)`, but the resulting values will be padded with `0x00` bytes at the end, which may be undesirable. To remove these bytes, use the `TRIM()` [1186] function:

```
UPDATE t SET coll = TRIM(TRAILING 0x00 FROM coll);
```

Suppose that table `t` has a nonbinary column named `coll` defined as `CHAR(50) CHARACTER SET latin1` but you want to convert it to use `utf8` so that you can store values from many languages. The following statement accomplishes this:

```
ALTER TABLE t MODIFY coll CHAR(50) CHARACTER SET utf8;
```

Conversion may be lossy if the column contains characters that are not in both character sets.

A special case occurs if you have old tables from before MySQL 4.1 where a nonbinary column contains values that actually are encoded in a character set different from the server's default character set. For example, an application might have stored `sjis` values in a column, even though MySQL's default character set was `latin1`. It is possible to convert the column to use the proper character set but an additional step is required. Suppose that the server's default character set was `latin1` and `coll` is defined as `CHAR(50)` but its contents are `sjis` values. The first step is to convert the column to a binary data type, which removes the existing character set information without performing any character conversion:

```
ALTER TABLE t MODIFY coll BLOB;
```

The next step is to convert the column to a nonbinary data type with the proper character set:

```
ALTER TABLE t MODIFY coll CHAR(50) CHARACTER SET sjis;
```

This procedure requires that the table not have been modified already with statements such as `INSERT` or `UPDATE` after an upgrade to MySQL 4.1 or later. In that case, MySQL would store new values in the column using `latin1`, and the column will contain a mix of `sjis` and `latin1` values and cannot be converted properly.

If you specified attributes when creating a column initially, you should also specify them when altering the table with `ALTER TABLE`. For example, if you specified `NOT NULL` and an explicit `DEFAULT` value, you should also provide them in the `ALTER TABLE` statement. Otherwise, the resulting column definition will not include those attributes.

## 10.1.14. Character Sets and Collations That MySQL Supports

MySQL supports 70+ collations for 30+ character sets. This section indicates which character sets MySQL supports. There is one subsection for each group of related character sets. For each character set, the permissible collations are listed.

You can always list the available character sets and their default collations with the `SHOW CHARACTER SET` statement:

```
mysql> SHOW CHARACTER SET;
```

Charset	Description	Default collation
big5	Big5 Traditional Chinese	big5_chinese_ci
dec8	DEC West European	dec8_swedish_ci
cp850	DOS West European	cp850_general_ci
hp8	HP West European	hp8_english_ci
koi8r	KOI8-R Relcom Russian	koi8r_general_ci
latin1	cp1252 West European	latin1_swedish_ci
latin2	ISO 8859-2 Central European	latin2_general_ci
swe7	7bit Swedish	swe7_swedish_ci
ascii	US ASCII	ascii_general_ci
ujis	EUC-JP Japanese	ujis_japanese_ci
sjis	Shift-JIS Japanese	sjis_japanese_ci
hebrew	ISO 8859-8 Hebrew	hebrew_general_ci
tis620	TIS620 Thai	tis620_thai_ci
euckr	EUC-KR Korean	euckr_korean_ci
koi8u	KOI8-U Ukrainian	koi8u_general_ci
gb2312	GB2312 Simplified Chinese	gb2312_chinese_ci
greek	ISO 8859-7 Greek	greek_general_ci
cp1250	Windows Central European	cp1250_general_ci
gbk	GBK Simplified Chinese	gbk_chinese_ci
latin5	ISO 8859-9 Turkish	latin5_turkish_ci
armscii8	ARMSCII-8 Armenian	armscii8_general_ci
utf8	UTF-8 Unicode	utf8_general_ci
ucs2	UCS-2 Unicode	ucs2_general_ci
cp866	DOS Russian	cp866_general_ci
keybcs2	DOS Kamenicky Czech-Slovak	keybcs2_general_ci
macce	Mac Central European	macce_general_ci
macroman	Mac West European	macroman_general_ci
cp852	DOS Central European	cp852_general_ci
latin7	ISO 8859-13 Baltic	latin7_general_ci
utf8mb4	UTF-8 Unicode	utf8mb4_general_ci
cp1251	Windows Cyrillic	cp1251_general_ci
utf16	UTF-16 Unicode	utf16_general_ci
utf16le	UTF-16LE Unicode	utf16le_general_ci
cp1256	Windows Arabic	cp1256_general_ci
cp1257	Windows Baltic	cp1257_general_ci
utf32	UTF-32 Unicode	utf32_general_ci
binary	Binary pseudo charset	binary
geostd8	GEOSTD8 Georgian	geostd8_general_ci
cp932	SJIS for Windows Japanese	cp932_japanese_ci
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci

In cases where a character set has multiple collations, it might not be clear which collation is most suitable for a given application. To avoid choosing the wrong collation, it can be helpful to perform some comparisons with representative data values to make sure that a given collation sorts values the way you expect.

[Collation-Charts.Org](http://Collation-Charts.Org) is a useful site for information that shows how one collation compares to another.

### 10.1.14.1. Unicode Character Sets

MySQL 5.6 supports these Unicode character sets:

- `ucs2`, the UCS-2 encoding of the Unicode character set using 16 bits per character.
- `utf16`, the UTF-16 encoding for the Unicode character set; like `ucs2` but with an extension for supplementary characters.
- `utf16le`, the UTF-16LE encoding for the Unicode character set; like `utf16` but little-endian rather than big-endian.
- `utf32`, the UTF-32 encoding for the Unicode character set using 32 bits per character.

- `utf8`, a UTF-8 encoding of the Unicode character set using one to three bytes per character.
- `utf8mb4`, a UTF-8 encoding of the Unicode character set using one to four bytes per character.

`ucs2` and `utf8` support Basic Multilingual Plane (BMP) characters. `utf8mb4`, `utf16`, `utf16le`, and `utf32` support BMP and supplementary characters. `utf16le` was added in MySQL 5.6.1.

You can store text in about 650 languages using these character sets. This section lists the collations available for each Unicode character set and describes their differentiating properties. For general information about the character sets, see [Section 10.1.10, “Unicode Support”](#).

A similar set of collations is available for most Unicode character sets. These are shown in the following list, where `xxx` represents the character set name. For example, `xxx_danish_ci` represents the Danish collations, the specific names of which are `ucs2_danish_ci`, `utf16_danish_ci`, `utf32_danish_ci`, `utf8_danish_ci`, and `utf8mb4_danish_ci`.

Collation support for `utf16le` is more limited. The only collations available are `utf16le_general_ci` and `utf16le_bin`. These are similar to `utf16_general_ci` and `utf16_bin`.

Unicode collation names may also include a version number (for example, `xxx_unicode_520_ci`) to indicate the Unicode Collation Algorithm version on which the collation is based, as described later in this section. For such collations, there is no `utf8mb3` alias to the corresponding `utf8` collation. See [Section 10.1.10.6, “The `utf8mb3` “Character Set” \(Alias for `utf8`\)”](#).

- `xxx_bin`
- `xxx_croatian_ci`
- `xxx_czech_ci`
- `xxx_danish_ci`
- `xxx_esperanto_ci`
- `xxx_estonian_ci`
- `xxx_general_ci` (default)
- `xxx_german2_ci`
- `xxx_general_mysql500_ci`
- `xxx_hungarian_ci`
- `xxx_icelandic_ci`
- `xxx_latvian_ci`
- `xxx_lithuanian_ci`
- `xxx_persian_ci`
- `xxx_polish_ci`
- `xxx_roman_ci`
- `xxx_romanian_ci`
- `xxx_sinhala_ci`
- `xxx_slovak_ci`
- `xxx_slovenian_ci`

- `xxx_spanish_ci`
- `xxx_spanish2_ci`
- `xxx_swedish_ci`
- `xxx_turkish_ci`
- `xxx_unicode_ci`
- `xxx_vietnamese_ci`

MySQL implements the `xxx_unicode_ci` collations according to the Unicode Collation Algorithm (UCA) described at <http://www.unicode.org/reports/tr10/>. The collation uses the version-4.0.0 UCA weight keys: <http://www.unicode.org/Public/UCA/4.0.0/allkeys-4.0.0.txt>. Currently, the `xxx_unicode_ci` collations have only partial support for the Unicode Collation Algorithm. Some characters are not supported yet. Also, combining marks are not fully supported. This affects primarily Vietnamese, Yoruba, and some smaller languages such as Navajo.

MySQL implements language-specific Unicode collations only if the ordering with `xxx_unicode_ci` does not work well for a language. Language-specific collations are UCA-based. They are derived from `xxx_unicode_ci` with additional language tailoring rules.

Collations based on UCA versions later than 4.0.0 include the version in the collation name. Thus, `xxx_unicode_520_ci` collations are based on UCA 5.2.0 weight keys: <http://www.unicode.org/Public/UCA/5.2.0/allkeys.txt>.

`LOWER()` [1182] and `UPPER()` [1187] perform case folding according to the collation of their argument. A character that has uppercase and lowercase versions only in a Unicode version more recent than 4.0.0 will be converted by these functions only if the argument has a collation that uses a recent enough UCA version.

For any Unicode character set, operations performed using the `xxx_general_ci` collation are faster than those for the `xxx_unicode_ci` collation. For example, comparisons for the `utf8_general_ci` collation are faster, but slightly less correct, than comparisons for `utf8_unicode_ci`. The reason for this is that `utf8_unicode_ci` supports mappings such as expansions; that is, when one character compares as equal to combinations of other characters. For example, in German and some other languages “ß” is equal to “ss”. `utf8_unicode_ci` also supports contractions and ignorable characters. `utf8_general_ci` is a legacy collation that does not support expansions, contractions, or ignorable characters. It can make only one-to-one comparisons between characters.

To further illustrate, the following equalities hold in both `utf8_general_ci` and `utf8_unicode_ci` (for the effect this has in comparisons or when doing searches, see [Section 10.1.7.8, “Examples of the Effect of Collation”](#)):

```
Ä = A
Ö = O
Ü = U
```

A difference between the collations is that this is true for `utf8_general_ci`:

```
ß = s
```

Whereas this is true for `utf8_unicode_ci`, which supports the German DIN-1 ordering (also known as dictionary order):

```
ß = ss
```

MySQL implements language-specific collations for the `utf8` character set only if the ordering with `utf8_unicode_ci` does not work well for a language. For example, `utf8_unicode_ci` works fine for German dictionary order and French, so there is no need to create special `utf8` collations.

`utf8_general_ci` also is satisfactory for both German and French, except that “ß” is equal to “s”, and not to “ss”. If this is acceptable for your application, you should use `utf8_general_ci` because it is faster. If this is not acceptable (for example, if you require German dictionary order), use `utf8_unicode_ci` because it is more accurate.

If you require German DIN-2 (phone book) ordering, use the `utf8_german2_ci` collation, which compares the following sets of characters equal:

```

Ä = Æ = AE
Ö = Œ = OE
Ü = UE
ß = ss

```

`utf8_german2_ci` is similar to `latin1_german2_ci`, but the latter does not compare “Æ” equal to “AE” or “Œ” equal to “OE”. There is no `utf8_german_ci` corresponding to `latin1_german_ci` for German dictionary order because `utf8_general_ci` suffices.

`xxx_swedish_ci` includes Swedish rules. For example, in Swedish, the following relationship holds, which is not something expected by a German or French speaker:

```

Û = Y < Ö

```

The `xxx_spanish_ci` and `xxx_spanish2_ci` collations correspond to modern Spanish and traditional Spanish, respectively. In both collations, “ñ” (n-tilde) is a separate letter between “n” and “o”. In addition, for traditional Spanish, “ch” is a separate letter between “c” and “d”, and “ll” is a separate letter between “l” and “m”.

In the `xxx_roman_ci` collations, I and J compare as equal, and U and V compare as equal.

The `xxx_croatian_ci` collations are tailored for these Croatian letters: Č, Ć, DŽ, Đ, Lj, Nj, Š, Ž.

For all Unicode collations except the “binary” (`xxx_bin`) collations, MySQL performs a table lookup to find a character’s collating weight. This weight can be displayed using the `WEIGHT_STRING()` [1187] function. (See Section 12.5, “String Functions”.) If a character is not in the table (for example, because it is a “new” character), collating weight determination becomes more complex:

- For BMP characters in general collations (`xxx_general_ci`), weight = code point.
- For BMP characters in UCA collations (for example, `xxx_unicode_ci` and language-specific collations), the following algorithm applies:

```

if (code >= 0x3400 && code <= 0x4DB5)
    base= 0xFB80; /* CJK Ideograph Extension */
else if (code >= 0x4E00 && code <= 0x9FA5)
    base= 0xFB40; /* CJK Ideograph */
else
    base= 0xFBC0; /* All other characters */
aaaa= base + (code >> 15);
bbbb= (code & 0x7FFF) | 0x8000;

```

The result is a sequence of two collating elements, `aaaa` followed by `bbbb`. For example:

```

mysql> SELECT HEX(WEIGHT_STRING(ucs2 0x04CF COLLATE ucs2_unicode_ci));
+-----+
| HEX(WEIGHT_STRING(ucs2 0x04CF COLLATE ucs2_unicode_ci)) |
+-----+
| FBC084CF |
+-----+

```

Thus, `U+04cf` **CYRILLIC SMALL LETTER PALOCHKA** is, with all UCA 4.0.0 collations, greater than `U+04c0` **CYRILLIC LETTER PALOCHKA**. With UCA 5.2.0 collations, all palochkas sort together.

- For supplementary characters in general collations, the weight is the weight for `0xffff` **REPLACEMENT CHARACTER**. For supplementary characters in UCA 4.0.0 collations, their collating



weight is `0xffffd`. That is, to MySQL, all supplementary characters are equal to each other, and greater than almost all BMP characters.

An example with Deseret characters and `COUNT(DISTINCT)`:

```
CREATE TABLE t (s1 VARCHAR(5) CHARACTER SET utf32 COLLATE utf32_unicode_ci);
INSERT INTO t VALUES (0xffffd); /* REPLACEMENT CHARACTER */
INSERT INTO t VALUES (0x010412); /* DESERET CAPITAL LETTER BEE */
INSERT INTO t VALUES (0x010413); /* DESERET CAPITAL LETTER TEE */
SELECT COUNT(DISTINCT s1) FROM t;
```

The result is 2 because in the MySQL `xxx_unicode_ci` collations, the replacement character has a weight of `0x0dc6`, whereas Deseret Bee and Deseret Tee both have a weight of `0xffffd`. (Were the `utf32_general_ci` collation used instead, the result would be 1 because all three characters have a weight of `0xffffd` in that collation.)

An example with cuneiform characters and `WEIGHT_STRING()` [1187]:

```
/*
The four characters in the INSERT string are
00000041 # LATIN CAPITAL LETTER A
0001218F # CUNEIFORM SIGN KAB
000121A7 # CUNEIFORM SIGN KISH
00000042 # LATIN CAPITAL LETTER B
*/
CREATE TABLE t (s1 CHAR(4) CHARACTER SET utf32 COLLATE utf32_unicode_ci);
INSERT INTO t VALUES (0x000000410001218f000121a700000042);
SELECT HEX(WEIGHT_STRING(s1)) FROM t;
```

The result is:

```
0E33 FFFD FFFD 0E4A
```

`0E33` and `0E4A` are primary weights as in [UCA 4.0.0](#). `FFFD` is the weight for KAB and also for KISH.

The rule that all supplementary characters are equal to each other is nonoptimal but is not expected to cause trouble. These characters are very rare, so it will be very rare that a multi-character string consists entirely of supplementary characters. In Japan, since the supplementary characters are obscure Kanji ideographs, the typical user does not care what order they are in, anyway. If you really want rows sorted by MySQL's rule and secondarily by code point value, it is easy:

```
ORDER BY s1 COLLATE utf32_unicode_ci, s1 COLLATE utf32_bin
```

- For supplementary characters based on UCA versions later than 4.0.0 (for example, `xxx_unicode_520_ci`), supplementary characters do not necessarily all have the same collation weight. Some have explicit weights from the UCA `allkeys.txt` file. Others have weights calculated from this algorithm:

```
aaaa= base + (code >> 15);
bbbb= (code & 0x7FFF) | 0x8000;
```

### The `utf16_bin` Collation

There is a difference between “ordering by the character's code value” and “ordering by the character's binary representation,” a difference that appears only with `utf16_bin`, because of surrogates.

Suppose that `utf16_bin` (the binary collation for `utf16`) was a binary comparison “byte by byte” rather than “character by character.” If that were so, the order of characters in `utf16_bin` would differ from the order in `utf8_bin`. For example, the following chart shows two rare characters. The first character is in the range `E000-FFFF`, so it is greater than a surrogate but less than a supplementary. The second character is a supplementary.

Code point	Character	utf8	utf16
0FF9D	HALFWIDTH KATAKANA LETTER N	EF BE 9D	FF 9D
10384	UGARITIC LETTER DELTA	F0 90 8E 84	D8 00 DF 84

The two characters in the chart are in order by code point value because `0xff9d < 0x10384`. And they are in order by `utf8` value because `0xef < 0xf0`. But they are not in order by `utf16` value, if we use byte-by-byte comparison, because `0xff > 0xd8`.

So MySQL's `utf16_bin` collation is not “byte by byte.” It is “by code point.” When MySQL sees a supplementary-character encoding in `utf16`, it converts to the character's code-point value, and then compares. Therefore, `utf8_bin` and `utf16_bin` are the same ordering. This is consistent with the SQL:2008 standard requirement for a UCS\_BASIC collation: “UCS\_BASIC is a collation in which the ordering is determined entirely by the Unicode scalar values of the characters in the strings being sorted. It is applicable to the UCS character repertoire. Since every character repertoire is a subset of the UCS repertoire, the UCS\_BASIC collation is potentially applicable to every character set. NOTE 11: The Unicode scalar value of a character is its code point treated as an unsigned integer.”

If the character set is `ucs2`, comparison is byte-by-byte, but `ucs2` strings should not contain surrogates, anyway.

The `xxx_general_mysql500_ci` collations were added in MySQL 5.6.5. They preserve the pre-5.1.24 ordering of the original `xxx_general_ci` collations and permit upgrades for tables created before MySQL 5.1.24. For more information, see [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#), and [Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”](#).

For additional information about Unicode collations in MySQL, see [Collation-Charts.Org](#) (`utf8`).

#### 10.1.14.2. West European Character Sets

Western European character sets cover most West European languages, such as French, Spanish, Catalan, Basque, Portuguese, Italian, Albanian, Dutch, German, Danish, Swedish, Norwegian, Finnish, Faroese, Icelandic, Irish, Scottish, and English.

- `ascii` (US ASCII) collations:
  - `ascii_bin`
  - `ascii_general_ci` (default)
- `cp850` (DOS West European) collations:
  - `cp850_bin`
  - `cp850_general_ci` (default)
- `dec8` (DEC Western European) collations:
  - `dec8_bin`
  - `dec8_swedish_ci` (default)
- `hp8` (HP Western European) collations:
  - `hp8_bin`
  - `hp8_english_ci` (default)
- `latin1` (cp1252 West European) collations:
  - `latin1_bin`
  - `latin1_danish_ci`
  - `latin1_general_ci`
  - `latin1_general_cs`

- `latin1_german1_ci`
- `latin1_german2_ci`
- `latin1_spanish_ci`
- `latin1_swedish_ci` (default)

`latin1` is the default character set. MySQL's `latin1` is the same as the Windows `cp1252` character set. This means it is the same as the official ISO 8859-1 or IANA (Internet Assigned Numbers Authority) `latin1`, except that IANA `latin1` treats the code points between `0x80` and `0x9f` as “undefined,” whereas `cp1252`, and therefore MySQL's `latin1`, assign characters for those positions. For example, `0x80` is the Euro sign. For the “undefined” entries in `cp1252`, MySQL translates `0x81` to Unicode `0x0081`, `0x8d` to `0x008d`, `0x8f` to `0x008f`, `0x90` to `0x0090`, and `0x9d` to `0x009d`.

The `latin1_swedish_ci` collation is the default that probably is used by the majority of MySQL customers. Although it is frequently said that it is based on the Swedish/Finnish collation rules, there are Swedes and Finns who disagree with this statement.

The `latin1_german1_ci` and `latin1_german2_ci` collations are based on the DIN-1 and DIN-2 standards, where DIN stands for *Deutsches Institut für Normung* (the German equivalent of ANSI). DIN-1 is called the “dictionary collation” and DIN-2 is called the “phone book collation.” For an example of the effect this has in comparisons or when doing searches, see [Section 10.1.7.8, “Examples of the Effect of Collation”](#).

- `latin1_german1_ci` (dictionary) rules:

```

Ä = A
Ö = O
Û = U
ß = s

```

- `latin1_german2_ci` (phone-book) rules:

```

Ä = AE
Ö = OE
Û = UE
ß = ss

```

In the `latin1_spanish_ci` collation, “ñ” (n-tilde) is a separate letter between “n” and “o”.

- `macroman` (Mac West European) collations:
  - `macroman_bin`
  - `macroman_general_ci` (default)
- `swe7` (7bit Swedish) collations:
  - `swe7_bin`
  - `swe7_swedish_ci` (default)

For additional information about Western European collations in MySQL, see [Collation-Charts.Org](#) ([ascii](#), [cp850](#), [dec8](#), [hp8](#), [latin1](#), [macroman](#), [swe7](#)).

### 10.1.14.3. Central European Character Sets

MySQL provides some support for character sets used in the Czech Republic, Slovakia, Hungary, Romania, Slovenia, Croatia, Poland, and Serbia (Latin).

- `cp1250` (Windows Central European) collations:

- `cp1250_bin`
- `cp1250_croatian_ci`
- `cp1250_czech_cs`
- `cp1250_general_ci` (default)
- `cp1250_polish_ci`
- `cp852` (DOS Central European) collations:
  - `cp852_bin`
  - `cp852_general_ci` (default)
- `keybcs2` (DOS Kamenicky Czech-Slovak) collations:
  - `keybcs2_bin`
  - `keybcs2_general_ci` (default)
- `latin2` (ISO 8859-2 Central European) collations:
  - `latin2_bin`
  - `latin2_croatian_ci`
  - `latin2_czech_cs`
  - `latin2_general_ci` (default)
  - `latin2_hungarian_ci`
- `macce` (Mac Central European) collations:
  - `macce_bin`
  - `macce_general_ci` (default)

For additional information about Central European collations in MySQL, see [Collation-Charts.Org](http://Collation-Charts.Org) ([cp1250](#), [cp852](#), [keybcs2](#), [latin2](#), [macce](#)).

#### 10.1.14.4. South European and Middle East Character Sets

South European and Middle Eastern character sets supported by MySQL include Armenian, Arabic, Georgian, Greek, Hebrew, and Turkish.

- `armscii8` (ARMScii-8 Armenian) collations:
  - `armscii8_bin`
  - `armscii8_general_ci` (default)
- `cp1256` (Windows Arabic) collations:
  - `cp1256_bin`
  - `cp1256_general_ci` (default)
- `geostd8` (GEOSTD8 Georgian) collations:
  - `geostd8_bin`

- `geostd8_general_ci` (default)
- `greek` (ISO 8859-7 Greek) collations:
  - `greek_bin`
  - `greek_general_ci` (default)
- `hebrew` (ISO 8859-8 Hebrew) collations:
  - `hebrew_bin`
  - `hebrew_general_ci` (default)
- `latin5` (ISO 8859-9 Turkish) collations:
  - `latin5_bin`
  - `latin5_turkish_ci` (default)

For additional information about South European and Middle Eastern collations in MySQL, see Collation-Charts.Org ([armscii8](#), [cp1256](#), [geostd8](#), [greek](#), [hebrew](#), [latin5](#)).

#### 10.1.14.5. Baltic Character Sets

The Baltic character sets cover Estonian, Latvian, and Lithuanian languages.

- `cp1257` (Windows Baltic) collations:
  - `cp1257_bin`
  - `cp1257_general_ci` (default)
  - `cp1257_lithuanian_ci`
- `latin7` (ISO 8859-13 Baltic) collations:
  - `latin7_bin`
  - `latin7_estonian_cs`
  - `latin7_general_ci` (default)
  - `latin7_general_cs`

For additional information about Baltic collations in MySQL, see Collation-Charts.Org ([cp1257](#), [latin7](#)).

#### 10.1.14.6. Cyrillic Character Sets

The Cyrillic character sets and collations are for use with Belarusian, Bulgarian, Russian, Ukrainian, and Serbian (Cyrillic) languages.

- `cp1251` (Windows Cyrillic) collations:
  - `cp1251_bin`
  - `cp1251_bulgarian_ci`
  - `cp1251_general_ci` (default)
  - `cp1251_general_cs`
  - `cp1251_ukrainian_ci`

- `cp866` (DOS Russian) collations:
  - `cp866_bin`
  - `cp866_general_ci` (default)
- `koi8r` (KOI8-R Relcom Russian) collations:
  - `koi8r_bin`
  - `koi8r_general_ci` (default)
- `koi8u` (KOI8-U Ukrainian) collations:
  - `koi8u_bin`
  - `koi8u_general_ci` (default)

For additional information about Cyrillic collations in MySQL, see [Collation-Charts.Org](http://Collation-Charts.Org) ([cp1251](#), [cp866](#), [koi8r](#), [koi8u](#)). ).

### 10.1.14.7. Asian Character Sets

The Asian character sets that we support include Chinese, Japanese, Korean, and Thai. These can be complicated. For example, the Chinese sets must allow for thousands of different characters. See [Section 10.1.14.7.1, “The `cp932` Character Set”](#), for additional information about the `cp932` and `sjis` character sets.

For answers to some common questions and problems relating support for Asian character sets in MySQL, see [Section B.11, “MySQL 5.6 FAQ: MySQL Chinese, Japanese, and Korean Character Sets”](#).

- `big5` (Big5 Traditional Chinese) collations:
  - `big5_bin`
  - `big5_chinese_ci` (default)
- `cp932` (SJIS for Windows Japanese) collations:
  - `cp932_bin`
  - `cp932_japanese_ci` (default)
- `ucjpm` (UJIS for Windows Japanese) collations:
  - `ucjpm_bin`
  - `ucjpm_japanese_ci` (default)
- `euckr` (EUC-KR Korean) collations:
  - `euckr_bin`
  - `euckr_korean_ci` (default)
- `gb2312` (GB2312 Simplified Chinese) collations:
  - `gb2312_bin`
  - `gb2312_chinese_ci` (default)
- `gbk` (GBK Simplified Chinese) collations:
  - `gbk_bin`

- `gbk_chinese_ci` (default)
- `sjis` (Shift-JIS Japanese) collations:
  - `sjis_bin`
  - `sjis_japanese_ci` (default)
- `tis620` (TIS620 Thai) collations:
  - `tis620_bin`
  - `tis620_thai_ci` (default)
- `ujis` (EUC-JP Japanese) collations:
  - `ujis_bin`
  - `ujis_japanese_ci` (default)

The `big5_chinese_ci` collation sorts on number of strokes.

For additional information about Asian collations in MySQL, see Collation-Charts.Org ([big5](#), [cp932](#), [eucjpm](#), [euckr](#), [gb2312](#), [gbk](#), [sjis](#), [tis620](#), [ujis](#)).

#### 10.1.14.7.1. The `cp932` Character Set

##### Why is `cp932` needed?

In MySQL, the `sjis` character set corresponds to the `Shift_JIS` character set defined by IANA, which supports JIS X0201 and JIS X0208 characters. (See <http://www.iana.org/assignments/character-sets>.)

However, the meaning of “SHIFT JIS” as a descriptive term has become very vague and it often includes the extensions to `Shift_JIS` that are defined by various vendors.

For example, “SHIFT JIS” used in Japanese Windows environments is a Microsoft extension of `Shift_JIS` and its exact name is `Microsoft Windows Codepage : 932` or `cp932`. In addition to the characters supported by `Shift_JIS`, `cp932` supports extension characters such as NEC special characters, NEC selected—IBM extended characters, and IBM selected characters.

Many Japanese users have experienced problems using these extension characters. These problems stem from the following factors:

- MySQL automatically converts character sets.
- Character sets are converted using Unicode (`ucs2`).
- The `sjis` character set does not support the conversion of these extension characters.
- There are several conversion rules from so-called “SHIFT JIS” to Unicode, and some characters are converted to Unicode differently depending on the conversion rule. MySQL supports only one of these rules (described later).

The MySQL `cp932` character set is designed to solve these problems.

Because MySQL supports character set conversion, it is important to separate IANA `Shift_JIS` and `cp932` into two different character sets because they provide different conversion rules.

##### How does `cp932` differ from `sjis`?

The `cp932` character set differs from `sjis` in the following ways:



- `cp932` supports NEC special characters, NEC selected—IBM extended characters, and IBM selected characters.
- Some `cp932` characters have two different code points, both of which convert to the same Unicode code point. When converting from Unicode back to `cp932`, one of the code points must be selected. For this “round trip conversion,” the rule recommended by Microsoft is used. (See <http://support.microsoft.com/kb/170559/EN-US/>.)

The conversion rule works like this:

- If the character is in both JIS X 0208 and NEC special characters, use the code point of JIS X 0208.
- If the character is in both NEC special characters and IBM selected characters, use the code point of NEC special characters.
- If the character is in both IBM selected characters and NEC selected—IBM extended characters, use the code point of IBM extended characters.

The table shown at <http://www.microsoft.com/globaldev/reference/dbcs/932.htm> provides information about the Unicode values of `cp932` characters. For `cp932` table entries with characters under which a four-digit number appears, the number represents the corresponding Unicode (`ucs2`) encoding. For table entries with an underlined two-digit value appears, there is a range of `cp932` character values that begin with those two digits. Clicking such a table entry takes you to a page that displays the Unicode value for each of the `cp932` characters that begin with those digits.

The following links are of special interest. They correspond to the encodings for the following sets of characters:

- NEC special characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_87.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_87.htm)

- NEC selected—IBM extended characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_ED.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_ED.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_EE.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_EE.htm)

- IBM selected characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FA.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FA.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FB.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FB.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FC.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FC.htm)

- `cp932` supports conversion of user-defined characters in combination with `eucajpms`, and solves the problems with `sjis/ujis` conversion. For details, please refer to <http://www.opengroup.org/jp/jvc/cde/sjis-euc-e.html>.

For some characters, conversion to and from `ucs2` is different for `sjis` and `cp932`. The following tables illustrate these differences.

Conversion to `ucs2`:

<code>sjis/cp932</code> Value	<code>sjis</code> -> <code>ucs2</code> Conversion	<code>cp932</code> -> <code>ucs2</code> Conversion
5C	005C	005C
7E	007E	007E
815C	2015	2015
815F	005C	FF3C
8160	301C	FF5E
8161	2016	2225

<b>sjis/cp932 Value</b>	<b>sjis -&gt; ucs2 Conversion</b>	<b>cp932 -&gt; ucs2 Conversion</b>
817C	2212	FF0D
8191	00A2	FFE0
8192	00A3	FFE1
81CA	00AC	FFE2

Conversion from `ucs2`:

<b>ucs2 value</b>	<b>ucs2 -&gt; sjis Conversion</b>	<b>ucs2 -&gt; cp932 Conversion</b>
005C	815F	5C
007E	7E	7E
00A2	8191	3F
00A3	8192	3F
00AC	81CA	3F
2015	815C	815C
2016	8161	3F
2212	817C	3F
2225	3F	8161
301C	8160	3F
FF0D	3F	817C
FF3C	3F	815F
FF5E	3F	8160
FFE0	3F	8191
FFE1	3F	8192
FFE2	3F	81CA

Users of any Japanese character sets should be aware that using `--character-set-client-handshake` [416] (or `--skip-character-set-client-handshake` [416]) has an important effect. See Section 5.1.3, “Server Command Options”.

## 10.2. Setting the Error Message Language

By default, `mysqld` produces error messages in English, but they can also be displayed in any of several other languages: Czech, Danish, Dutch, Estonian, French, German, Greek, Hungarian, Italian, Japanese, Korean, Norwegian, Norwegian-ny, Polish, Portuguese, Romanian, Russian, Slovak, Spanish, or Swedish.

You can select which language the server uses for error messages using the instructions in this section.

In MySQL 5.6, the server searches for the error message file in two locations:

- It tries to find the file in a directory constructed from two system variable values, `lc_messages_dir` [495] and `lc_messages` [495], with the latter converted to a language name. Suppose that you start the server using this command:

```
shell> mysqld --lc_messages_dir=/usr/share/mysql --lc_messages=fr_FR
```

In this case, `mysqld` maps the locale `fr_FR` to the language `french` and looks for the error file in the `/usr/share/mysql/french` directory.

- If the message file cannot be found in the directory constructed as just described, the server ignores the `lc_messages` [495] value and uses only the `lc_messages_dir` [495] value as the location in which to look.

The `lc_messages_dir` [495] system variable has only a global value and is read only. `lc_messages` [495] has global and session values and can be modified at runtime, so the error message language can be changed while the server is running, and individual clients each can have a different error message language by changing their session `lc_messages` [495] value to a different locale name. For example, if the server is using the `fr_FR` locale for error messages, a client can execute this statement to receive error messages in English:

```
mysql> SET lc_messages = 'en_US';
```

By default, the language files are located in the `share/mysql/LANGUAGE` directory under the MySQL base directory.

For information about changing the character set for error messages (rather than the language), see [Section 10.1.6, “Character Set for Error Messages”](#).

You can change the content of the error messages produced by the server using the instructions in the MySQL Internals manual, available at [MySQL Internals: Error Messages](#). If you do change the content of error messages, remember to repeat your changes after each upgrade to a newer version of MySQL.

## 10.3. Adding a Character Set

This section discusses the procedure for adding a character set to MySQL. The proper procedure depends on whether the character set is simple or complex:

- If the character set does not need special string collating routines for sorting and does not need multi-byte character support, it is simple.
- If the character set needs either of those features, it is complex.

For example, `greek` and `swe7` are simple character sets, whereas `big5` and `czech` are complex character sets.

To use the following instructions, you must have a MySQL source distribution. In the instructions, `MYSET` represents the name of the character set that you want to add.

1. Add a `<charset>` element for `MYSET` to the `sql/share/charsets/Index.xml` file. Use the existing contents in the file as a guide to adding new contents. A partial listing for the `latin1` `<charset>` element follows:

```
<charset name="latin1">
  <family>Western</family>
  <description>cp1252 West European</description>
  ...
  <collation name="latin1_swedish_ci" id="8" order="Finnish, Swedish">
    <flag>primary</flag>
    <flag>compiled</flag>
  </collation>
  <collation name="latin1_danish_ci" id="15" order="Danish"/>
  ...
  <collation name="latin1_bin" id="47" order="Binary">
    <flag>binary</flag>
    <flag>compiled</flag>
  </collation>
  ...
</charset>
```

The `<charset>` element must list all the collations for the character set. These must include at least a binary collation and a default (primary) collation. The default collation is often named using a suffix of `general_ci` (general, case insensitive). It is possible for the binary collation to be the

default collation, but usually they are different. The default collation should have a [primary](#) flag. The binary collation should have a [binary](#) flag.

You must assign a unique ID number to each collation. The range of IDs from 1024 to 2047 is reserved for user-defined collations. To find the maximum of the currently used collation IDs, use this query:

```
SELECT MAX(ID) FROM INFORMATION_SCHEMA.COLLATIONS;
```

2. This step depends on whether you are adding a simple or complex character set. A simple character set requires only a configuration file, whereas a complex character set requires C source file that defines collation functions, multi-byte functions, or both.

For a simple character set, create a configuration file, `MYSET.xml`, that describes the character set properties. Create this file in the `sql/share/charsets` directory. You can use a copy of `latin1.xml` as the basis for this file. The syntax for the file is very simple:

- Comments are written as ordinary XML comments (`<!-- text -->`).
- Words within `<map>` array elements are separated by arbitrary amounts of whitespace.
- Each word within `<map>` array elements must be a number in hexadecimal format.
- The `<map>` array element for the `<ctype>` element has 257 words. The other `<map>` array elements after that have 256 words. See [Section 10.3.1, “Character Definition Arrays”](#).
- For each collation listed in the `<charset>` element for the character set in `Index.xml`, `MYSET.xml` must contain a `<collation>` element that defines the character ordering.

For a complex character set, create a C source file that describes the character set properties and defines the support routines necessary to properly perform operations on the character set:

- Create the file `ctype-MYSET.c` in the `strings` directory. Look at one of the existing `ctype-*.c` files (such as `ctype-big5.c`) to see what needs to be defined. The arrays in your file must have names like `ctype_MYSET`, `to_lower_MYSET`, and so on. These correspond to the arrays for a simple character set. See [Section 10.3.1, “Character Definition Arrays”](#).
  - For each `<collation>` element listed in the `<charset>` element for the character set in `Index.xml`, the `ctype-MYSET.c` file must provide an implementation of the collation.
  - If the character set requires string collating functions, see [Section 10.3.2, “String Collating Support for Complex Character Sets”](#).
  - If the character set requires multi-byte character support, see [Section 10.3.3, “Multi-Byte Character Support for Complex Character Sets”](#).
3. Modify the configuration information. Use the existing configuration information as a guide to adding information for `MYSYS`. The example here assumes that the character set has default and binary collations, but more lines are needed if `MYSET` has additional collations.

- a. Edit `mysys/charset-def.c`, and “register” the collations for the new character set.

Add these lines to the “declaration” section:

```
#ifdef HAVE_CHARSET_MYSET
extern CHARSET_INFO my_charset_MYSET_general_ci;
extern CHARSET_INFO my_charset_MYSET_bin;
#endif
```

Add these lines to the “registration” section:

```
#ifdef HAVE_CHARSET_MYSET
add_compiled_collation(&my_charset_MYSET_general_ci);
```

```
add_compiled_collation(&my_charset_MYSET_bin);
#endif
```

- b. If the character set uses `ctype-MYSET.c`, edit `strings/CMakeLists.txt` and add `ctype-MYSET.c` to the definition of the `STRINGS_SOURCES` variable.
  - c. Edit `cmake/character_sets.cmake`:
    - i. Add `MYSET` to the value of with `CHARSETS_AVAILABLE` in alphabetic order.
    - ii. Add `MYSET` to the value of `CHARSETS_COMPLEX` in alphabetic order. This is needed even for simple character sets, or CMake will not recognize `-DDEFAULT_CHARSET=MYSET` [134].
4. Reconfigure, recompile, and test.

### 10.3.1. Character Definition Arrays

Each simple character set has a configuration file located in the `sql/share/charsets` directory. For a character set named `MYSYS`, the file is named `MYSET.xml`. It uses `<map>` array elements to list character set properties. `<map>` elements appear within these elements:

- `<ctype>` defines attributes for each character.
- `<lower>` and `<upper>` list the lowercase and uppercase characters.
- `<unicode>` maps 8-bit character values to Unicode values.
- `<collation>` elements indicate character ordering for comparisons and sorts, one element per collation. Binary collations need no `<map>` element because the character codes themselves provide the ordering.

For a complex character set as implemented in a `ctype-MYSET.c` file in the `strings` directory, there are corresponding arrays: `ctype_MYSET[ ]`, `to_lower_MYSET[ ]`, and so forth. Not every complex character set has all of the arrays. See also the existing `ctype-*.c` files for examples. See the `CHARSET_INFO.txt` file in the `strings` directory for additional information.

Most of the arrays are indexed by character value and have 256 elements. The `<ctype>` array is indexed by character value + 1 and has 257 elements. This is a legacy convention for handling EOF.

`<ctype>` array elements are bit values. Each element describes the attributes of a single character in the character set. Each attribute is associated with a bitmask, as defined in `include/m_ctype.h`:

```
#define _MY_U 01 /* Upper case */
#define _MY_L 02 /* Lower case */
#define _MY_NMR 04 /* Numeral (digit) */
#define _MY_SPC 010 /* Spacing character */
#define _MY_PNT 020 /* Punctuation */
#define _MY_CTR 040 /* Control character */
#define _MY_B 0100 /* Blank */
#define _MY_X 0200 /* hexadecimal digit */
```

The `<ctype>` value for a given character should be the union of the applicable bitmask values that describe the character. For example, 'A' is an uppercase character (`_MY_U`) as well as a hexadecimal digit (`_MY_X`), so its `ctype` value should be defined like this:

```
ctype['A'+1] = _MY_U | _MY_X = 01 | 0200 = 0201
```

The bitmask values in `m_ctype.h` are octal values, but the elements of the `<ctype>` array in `MYSET.xml` should be written as hexadecimal values.

The `<lower>` and `<upper>` arrays hold the lowercase and uppercase characters corresponding to each member of the character set. For example:

```
lower['A'] should contain 'a'
```

```
upper['a'] should contain 'A'
```

Each `<collation>` array indicates how characters should be ordered for comparison and sorting purposes. MySQL sorts characters based on the values of this information. In some cases, this is the same as the `<upper>` array, which means that sorting is case-insensitive. For more complicated sorting rules (for complex character sets), see the discussion of string collating in [Section 10.3.2, “String Collating Support for Complex Character Sets”](#).

### 10.3.2. String Collating Support for Complex Character Sets

For a simple character set named *MYSET*, sorting rules are specified in the *MYSET.xml* configuration file using `<map>` array elements within `<collation>` elements. If the sorting rules for your language are too complex to be handled with simple arrays, you must define string collating functions in the *ctype-MYSET.c* source file in the *strings* directory.

The existing character sets provide the best documentation and examples to show how these functions are implemented. Look at the *ctype-\*.c* files in the *strings* directory, such as the files for the *big5*, *czech*, *gbk*, *sjis*, and *tis160* character sets. Take a look at the *MY\_COLLATION\_HANDLER* structures to see how they are used. See also the *CHARSET\_INFO.txt* file in the *strings* directory for additional information.

### 10.3.3. Multi-Byte Character Support for Complex Character Sets

If you want to add support for a new character set named *MYSET* that includes multi-byte characters, you must use multi-byte character functions in the *ctype-MYSET.c* source file in the *strings* directory.

The existing character sets provide the best documentation and examples to show how these functions are implemented. Look at the *ctype-\*.c* files in the *strings* directory, such as the files for the *euc\_kr*, *gb2312*, *gbk*, *sjis*, and *ujis* character sets. Take a look at the *MY\_CHARSET\_HANDLER* structures to see how they are used. See also the *CHARSET\_INFO.txt* file in the *strings* directory for additional information.

## 10.4. Adding a Collation to a Character Set

A collation is a set of rules that defines how to compare and sort character strings. Each collation in MySQL belongs to a single character set. Every character set has at least one collation, and most have two or more collations.

A collation orders characters based on weights. Each character in a character set maps to a weight. Characters with equal weights compare as equal, and characters with unequal weights compare according to the relative magnitude of their weights.

The `WEIGHT_STRING()` [1187] function can be used to see the weights for the characters in a string. The value that it returns to indicate weights is a binary string, so it is convenient to use `HEX(WEIGHT_STRING(str))` to display the weights in printable form. The following example shows that weights do not differ for lettercase for the letters in `'AaBb'` if it is a nonbinary case-insensitive string, but do differ if it is a binary string:

```
mysql> SELECT HEX(WEIGHT_STRING('AaBb' COLLATE latin1_swedish_ci));
+-----+
| HEX(WEIGHT_STRING('AaBb' COLLATE latin1_swedish_ci)) |
+-----+
| 41414242                                             |
+-----+
mysql> SELECT HEX(WEIGHT_STRING(BINARY 'AaBb'));
+-----+
| HEX(WEIGHT_STRING(BINARY 'AaBb')) |
+-----+
| 41614262                                             |
+-----+
```

MySQL supports several collation implementations, as discussed in [Section 10.4.1, “Collation Implementation Types”](#). Some of these can be added to MySQL without recompiling:

- Simple collations for 8-bit character sets.
- UCA-based collations for Unicode character sets.
- Binary (`xxx_bin`) collations.

The following sections describe how to add collations of the first two types to existing character sets. All existing character sets already have a binary collation, so there is no need here to describe how to add one.

Summary of the procedure for adding a new collation:

1. Choose a collation ID.
2. Add configuration information that names the collation and describes the character-ordering rules.
3. Restart the server.
4. Verify that the collation is present.

The instructions here cover only collations that can be added without recompiling MySQL. To add a collation that does require recompiling (as implemented by means of functions in a C source file), use the instructions in [Section 10.3, “Adding a Character Set”](#). However, instead of adding all the information required for a complete character set, just modify the appropriate files for an existing character set. That is, based on what is already present for the character set's current collations, add data structures, functions, and configuration information for the new collation.



#### Note

If you modify an existing collation, that may affect the ordering of rows for indexes on columns that use the collation. In this case, rebuild any such indexes to avoid problems such as incorrect query results. For further information, see [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#).

## Additional Resources

- The Unicode Collation Algorithm (UCA) specification: <http://www.unicode.org/reports/tr10/>
- The Locale Data Markup Language (LDML) specification: <http://www.unicode.org/reports/tr35/>
- MySQL Blog article “Instructions for adding a new Unicode collation”: <http://blogs.mysql.com/peterg/2008/05/19/instructions-for-adding-a-new-unicode-collation/>

### 10.4.1. Collation Implementation Types

MySQL implements several types of collations:

#### Simple collations for 8-bit character sets

This kind of collation is implemented using an array of 256 weights that defines a one-to-one mapping from character codes to weights. `latin1_swedish_ci` is an example. It is a case-insensitive collation, so the uppercase and lowercase versions of a character have the same weights and they compare as equal.

```
mysql> SET NAMES 'latin1' COLLATE 'latin1_swedish_ci';
Query OK, 0 rows affected (0.01 sec)

mysql> SELECT HEX(WEIGHT_STRING('a')), HEX(WEIGHT_STRING('A'));
+-----+-----+

```



```

| HEX(WEIGHT_STRING('a')) | HEX(WEIGHT_STRING('A')) |
+-----+-----+
| 41 | 41 |
+-----+-----+
1 row in set (0.01 sec)

mysql> SELECT 'a' = 'A';
+-----+
| 'a' = 'A' |
+-----+
| 1 |
+-----+
1 row in set (0.12 sec)

```

For implementation instructions, see [Section 10.4.3, “Adding a Simple Collation to an 8-Bit Character Set”](#).

### Complex collations for 8-bit character sets

This kind of collation is implemented using functions in a C source file that define how to order characters, as described in [Section 10.3, “Adding a Character Set”](#).

### Collations for non-Unicode multi-byte character sets

For this type of collation, 8-bit (single-byte) and multi-byte characters are handled differently. For 8-bit characters, character codes map to weights in case-insensitive fashion. (For example, the single-byte characters 'a' and 'A' both have a weight of 0x41.) For multi-byte characters, there are two types of relationship between character codes and weights:

- Weights equal character codes. `sjis_japanese_ci` is an example of this kind of collation. The multi-byte character 'ぢ' has a character code of 0x82C0, and the weight is also 0x82C0.

```

mysql> CREATE TABLE t1
  -> (c1 VARCHAR(2) CHARACTER SET sjis COLLATE sjis_japanese_ci);
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO t1 VALUES ('a'),('A'),(0x82C0);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT c1, HEX(c1), HEX(WEIGHT_STRING(c1)) FROM t1;
+-----+-----+-----+
| c1 | HEX(c1) | HEX(WEIGHT_STRING(c1)) |
+-----+-----+-----+
| a | 61 | 41 |
| A | 41 | 41 |
| ぢ | 82C0 | 82C0 |
+-----+-----+-----+
3 rows in set (0.00 sec)

```

- Character codes map one-to-one to weights, but a code is not necessarily equal to the weight. `gbk_chinese_ci` is an example of this kind of collation. The multi-byte character '臙' has a character code of 0x81B0 but a weight of 0xC286.

```

mysql> CREATE TABLE t1
  -> (c1 VARCHAR(2) CHARACTER SET gbk COLLATE gbk_chinese_ci);
Query OK, 0 rows affected (0.33 sec)

mysql> INSERT INTO t1 VALUES ('a'),('A'),(0x81B0);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT c1, HEX(c1), HEX(WEIGHT_STRING(c1)) FROM t1;
+-----+-----+-----+
| c1 | HEX(c1) | HEX(WEIGHT_STRING(c1)) |
+-----+-----+-----+
| a | 61 | 41 |
| A | 41 | 41 |

```

```
| 膳      | 81B0      | C286      |
+-----+-----+-----+
3 rows in set (0.00 sec)
```

For implementation instructions, see [Section 10.3, “Adding a Character Set”](#).

### Collations for Unicode multi-byte character sets

Some of these collations are based on the Unicode Collation Algorithm (UCA), others are not.

Non-UCA collations have a one-to-one mapping from character code to weight. In MySQL, such collations are case insensitive and accent insensitive. `utf8_general_ci` is an example: `'a'`, `'A'`, `'À'`, and `'á'` each have different character codes but all have a weight of `0x0041` and compare as equal.

```
mysql> SET NAMES 'utf8' COLLATE 'utf8_general_ci';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t1
-> (c1 CHAR(1) CHARACTER SET UTF8 COLLATE utf8_general_ci);
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO t1 VALUES ('a'),('A'),('À'),('á');
Query OK, 4 rows affected (0.00 sec)
Records: 4  Duplicates: 0  Warnings: 0

mysql> SELECT c1, HEX(c1), HEX(WEIGHT_STRING(c1)) FROM t1;
+-----+-----+-----+
| c1    | HEX(c1) | HEX(WEIGHT_STRING(c1)) |
+-----+-----+-----+
| a     | 61      | 0041                    |
| A     | 41      | 0041                    |
| À     | C380    | 0041                    |
| á     | C3A1    | 0041                    |
+-----+-----+-----+
4 rows in set (0.00 sec)
```

UCA-based collations in MySQL have these properties:

- If a character has weights, each weight uses 2 bytes (16 bits).
- A character may have zero weights (or an empty weight). In this case, the character is ignorable. Example: "U+0000 NULL" does not have a weight and is ignorable.
- A character may have one weight. Example: `'a'` has a weight of `0x0E33`.

```
mysql> SET NAMES 'utf8' COLLATE 'utf8_unicode_ci';
Query OK, 0 rows affected (0.05 sec)

mysql> SELECT HEX('a'), HEX(WEIGHT_STRING('a'));
+-----+-----+
| HEX('a') | HEX(WEIGHT_STRING('a')) |
+-----+-----+
| 61      | 0E33                     |
+-----+-----+
1 row in set (0.02 sec)
```

- A character may have many weights. This is an expansion. Example: The German letter `'ß'` (SZ ligature, or SHARP S) has a weight of `0x0FEA0FEA`.

```
mysql> SET NAMES 'utf8' COLLATE 'utf8_unicode_ci';
Query OK, 0 rows affected (0.11 sec)

mysql> SELECT HEX('ß'), HEX(WEIGHT_STRING('ß'));
+-----+-----+
| HEX('ß') | HEX(WEIGHT_STRING('ß')) |
+-----+-----+
| C39F    | 0FEA0FEA                 |
+-----+-----+
```

```
1 row in set (0.00 sec)
```

- Many characters may have one weight. This is a contraction. Example: 'ch' is a single letter in Czech and has a weight of 0x0EE2.

```
mysql> SET NAMES 'utf8' COLLATE 'utf8_czech_ci';
Query OK, 0 rows affected (0.09 sec)

mysql> SELECT HEX('ch'), HEX(WEIGHT_STRING('ch'));
+-----+-----+
| HEX('ch') | HEX(WEIGHT_STRING('ch')) |
+-----+-----+
| 6368      | 0EE2                      |
+-----+-----+
1 row in set (0.00 sec)
```

A many-characters-to-many-weights mapping is also possible (this is contraction with expansion), but is not supported by MySQL.

For implementation instructions, for a non-UCA collation, see [Section 10.3, “Adding a Character Set”](#). For a UCA collation, see [Section 10.4.4, “Adding a UCA Collation to a Unicode Character Set”](#).

### Miscellaneous collations

There are also a few collations that do not fall into any of the previous categories.

## 10.4.2. Choosing a Collation ID

Each collation must have a unique ID. To add a collation, you must choose an ID value that is not currently used. The range of IDs from 1024 to 2047 is reserved for user-defined collations. The collation ID that you choose will appear in these contexts:

- The `ID` column of the `INFORMATION_SCHEMA.COLLATIONS` table.
- The `Id` column of `SHOW COLLATION` output.
- The `charsetnr` member of the `MYSQL_FIELD` C API data structure.
- The `number` member of the `MY_CHARSET_INFO` data structure returned by the `mysql_get_character_set_info()` C API function.

To determine the largest currently used ID, issue the following statement:

```
mysql> SELECT MAX(ID) FROM INFORMATION_SCHEMA.COLLATIONS;
+-----+
| MAX(ID) |
+-----+
|      210 |
+-----+
```

To display a list of all currently used IDs, issue this statement:

```
mysql> SELECT ID FROM INFORMATION_SCHEMA.COLLATIONS ORDER BY ID;
+-----+
| ID |
+-----+
|  1 |
|  2 |
| ... |
| 52 |
| 53 |
| 57 |
| 58 |
| ... |
| 98 |
| 99 |
+-----+
```

```
| 128 |
| 129 |
| ... |
| 210 |
+-----+
```



### Warning

Before MySQL 5.5, which provides for a range of user-defined collation IDs, you must choose an ID in the range from 1 to 254. In this case, if you upgrade MySQL, you may find that the collation ID you choose has been assigned to a collation included in the new MySQL distribution. In this case, you will need to choose a new value for your own collation.

In addition, before upgrading, you should save the configuration files that you change. If you upgrade in place, the process will replace the your modified files.

## 10.4.3. Adding a Simple Collation to an 8-Bit Character Set

This section describes how to add a simple collation for an 8-bit character set by writing the `<collation>` elements associated with a `<charset>` character set description in the MySQL `Index.xml` file. The procedure described here does not require recompiling MySQL. The example adds a collation named `latin1_test_ci` to the `latin1` character set.

1. Choose a collation ID, as shown in [Section 10.4.2, “Choosing a Collation ID”](#). The following steps use an ID of 1024.
2. Modify the `Index.xml` and `latin1.xml` configuration files. These files will be located in the directory named by the `character_sets_dir` [469] system variable. You can check the variable value as follows, although the path name might be different on your system:

```
mysql> SHOW VARIABLES LIKE 'character_sets_dir';
+-----+
| Variable_name | Value |
+-----+
| character_sets_dir | /user/local/mysql/share/mysql/charsets/ |
+-----+
```

3. Choose a name for the collation and list it in the `Index.xml` file. Find the `<charset>` element for the character set to which the collation is being added, and add a `<collation>` element that indicates the collation name and ID, to associate the name with the ID. For example:

```
<charset name="latin1">
...
  <collation name="latin1_test_ci" id="1024"/>
...
</charset>
```

4. In the `latin1.xml` configuration file, add a `<collation>` element that names the collation and that contains a `<map>` element that defines a character code-to-weight mapping table for character codes 0 to 255. Each value within the `<map>` element must be a number in hexadecimal format.

```
<collation name="latin1_test_ci">
<map>
00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
10 11 12 13 14 15 16 17 18 19 1A 1B 1C 1D 1E 1F
20 21 22 23 24 25 26 27 28 29 2A 2B 2C 2D 2E 2F
30 31 32 33 34 35 36 37 38 39 3A 3B 3C 3D 3E 3F
40 41 42 43 44 45 46 47 48 49 4A 4B 4C 4D 4E 4F
50 51 52 53 54 55 56 57 58 59 5A 5B 5C 5D 5E 5F
60 41 42 43 44 45 46 47 48 49 4A 4B 4C 4D 4E 4F
50 51 52 53 54 55 56 57 58 59 5A 7B 7C 7D 7E 7F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
90 91 92 93 94 95 96 97 98 99 9A 9B 9C 9D 9E 9F
A0 A1 A2 A3 A4 A5 A6 A7 A8 A9 AA AB AC AD AE AF
B0 B1 B2 B3 B4 B5 B6 B7 B8 B9 BA BB BC BD BE BF
```

```

41 41 41 41 5B 5D 5B 43 45 45 45 45 49 49 49 49
44 4E 4F 4F 4F 4F 5C D7 5C 55 55 55 59 59 DE DF
41 41 41 41 5B 5D 5B 43 45 45 45 45 49 49 49 49
44 4E 4F 4F 4F 4F 5C F7 5C 55 55 55 59 59 DE FF
</map>
</collation>

```

- Restart the server and use this statement to verify that the collation is present:

```

mysql> SHOW COLLATION LIKE 'latin1_test_ci';
+-----+-----+-----+-----+-----+-----+
| Collation      | Charset | Id    | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| latin1_test_ci | latin1  | 1024  |         |          | 1       |
+-----+-----+-----+-----+-----+-----+

```

### 10.4.4. Adding a UCA Collation to a Unicode Character Set

This section describes how to add a UCA collation for a Unicode character set by writing the `<collation>` element within a `<charset>` character set description in the MySQL `Index.xml` file. The procedure described here does not require recompiling MySQL. It uses a subset of the Locale Data Markup Language (LDML) specification, which is available at <http://www.unicode.org/reports/tr35/>. With this method, you need not define the entire collation. Instead, you begin with an existing “base” collation and describe the new collation in terms of how it differs from the base collation. The following table lists the base collations of the Unicode character sets for which UCA collations can be defined. It is not possible to create user-defined UCA collations for `utf16le`; there is no `utf16le_unicode_ci` collation that would serve as the basis for such collations.

**Table 10.1. MySQL Character Sets Available for User-Defined UCA Collations**

Character Set	Base Collation
<code>utf8</code>	<code>utf8_unicode_ci</code>
<code>ucs2</code>	<code>ucs2_unicode_ci</code>
<code>utf16</code>	<code>utf16_unicode_ci</code>
<code>utf32</code>	<code>utf32_unicode_ci</code>

The following sections show how to add a collation that is defined using LDML syntax, and provide a summary of LDML rules supported in MySQL.

#### 10.4.4.1. Defining a UCA Collation using LDML Syntax

To add a UCA collation for a Unicode character set without recompiling MySQL, use the following procedure. If you are unfamiliar with the LDML rules used to describe the collation's sort characteristics, see [Section 10.4.4.2, “LDML Syntax Supported in MySQL”](#).

The example adds a collation named `utf8_phone_ci` to the `utf8` character set. The collation is designed for a scenario involving a Web application for which users post their names and phone numbers. Phone numbers can be given in very different formats:

```

+7-12345-67
+7-12-345-67
+7 12 345 67
+7 (12) 345 67
+71234567

```

The problem raised by dealing with these kinds of values is that the varying permissible formats make searching for a specific phone number very difficult. The solution is to define a new collation that reorders punctuation characters, making them ignorable.

- Choose a collation ID, as shown in [Section 10.4.2, “Choosing a Collation ID”](#). The following steps use an ID of 1029.

- To modify the `Index.xml` configuration file. This file will be located in the directory named by the `character_sets_dir` [469] system variable. You can check the variable value as follows, although the path name might be different on your system:

```
mysql> SHOW VARIABLES LIKE 'character_sets_dir';
+-----+-----+
| Variable_name | Value                               |
+-----+-----+
| character_sets_dir | /user/local/mysql/share/mysql/charsets/ |
+-----+-----+
```

- Choose a name for the collation and list it in the `Index.xml` file. In addition, you'll need to provide the collation ordering rules. Find the `<charset>` element for the character set to which the collation is being added, and add a `<collation>` element that indicates the collation name and ID, to associate the name with the ID. Within the `<collation>` element, provide a `<rules>` element containing the ordering rules:

```
<charset name="utf8">
...
  <collation name="utf8_phone_ci" id="1029">
    <rules>
      <reset>\u0000</reset>
      <i>\u0020</i> <!-- space -->
      <i>\u0028</i> <!-- left parenthesis -->
      <i>\u0029</i> <!-- right parenthesis -->
      <i>\u002B</i> <!-- plus -->
      <i>\u002D</i> <!-- hyphen -->
    </rules>
  </collation>
...
</charset>
```

- If you want a similar collation for other Unicode character sets, add other `<collation>` elements. For example, to define `ucs2_phone_ci`, add a `<collation>` element to the `<charset name="ucs2">` element. Remember that each collation must have its own unique ID.
- Restart the server and use this statement to verify that the collation is present:

```
mysql> SHOW COLLATION LIKE 'utf8_phone_ci';
+-----+-----+-----+-----+-----+-----+
| Collation | Charset | Id | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| utf8_phone_ci | utf8 | 1029 | | | 8 |
+-----+-----+-----+-----+-----+-----+
```

Now test the collation to make sure that it has the desired properties.

Create a table containing some sample phone numbers using the new collation:

```
mysql> CREATE TABLE phonebook (
->   name VARCHAR(64),
->   phone VARCHAR(64) CHARACTER SET utf8 COLLATE utf8_phone_ci
-> );
Query OK, 0 rows affected (0.09 sec)

mysql> INSERT INTO phonebook VALUES ('Svoj','+7 912 800 80 02');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO phonebook VALUES ('Hf','+7 (912) 800 80 04');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO phonebook VALUES ('Bar','+7-912-800-80-01');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO phonebook VALUES ('Ramil','(7912) 800 80 03');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO phonebook VALUES ('Sanja','+380 (912) 8008005');
Query OK, 1 row affected (0.00 sec)
```

Run some queries to see whether the ignored punctuation characters are in fact ignored for sorting and comparisons:

```
mysql> SELECT * FROM phonebook ORDER BY phone;
+-----+-----+
| name | phone                |
+-----+-----+
| Sanja | +380 (912) 8008005 |
| Bar   | +7-912-800-80-01 |
| Svoj  | +7 912 800 80 02 |
| Ramil | (7912) 800 80 03 |
| Hf    | +7 (912) 800 80 04 |
+-----+-----+
5 rows in set (0.00 sec)

mysql> SELECT * FROM phonebook WHERE phone='+7(912)800-80-01';
+-----+-----+
| name | phone                |
+-----+-----+
| Bar   | +7-912-800-80-01 |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT * FROM phonebook WHERE phone='79128008001';
+-----+-----+
| name | phone                |
+-----+-----+
| Bar   | +7-912-800-80-01 |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT * FROM phonebook WHERE phone='7 9 1 2 8 0 0 8 0 0 1';
+-----+-----+
| name | phone                |
+-----+-----+
| Bar   | +7-912-800-80-01 |
+-----+-----+
1 row in set (0.00 sec)
```

#### 10.4.4.2. LDML Syntax Supported in MySQL

This section describes the LDML syntax that MySQL recognizes. This is a subset of the syntax described in the LDML specification available at <http://www.unicode.org/reports/tr35/>, which should be consulted for further information. MySQL recognizes a large enough subset of the syntax that, in many cases, it is possible to download a collation definition from the Unicode Common Locale Data Repository and paste the relevant part (that is, the part between the `<rules>` and `</rules>` tags) into the MySQL `Index.xml` file. The rules described here are all supported except that character sorting occurs only at the primary level. Rules that specify differences at secondary or higher sort levels are recognized (and thus can be included in collation definitions) but are treated as equality at the primary level.

The MySQL server generates diagnostics when it finds problems while parsing the `Index.xml` file. See [Section 10.4.4.3, “Diagnostics During `Index.xml` Parsing”](#).

##### Character Representation

Characters named in LDML rules can be written literally or in `\unnnn` format, where `nnnn` is the hexadecimal Unicode code point value. For example, `A` and `á` can be written literally or as `\u0041` and `\u00E1`. Within hexadecimal values, the digits `A` through `F` are not case sensitive; `\u00E1` and `\u00e1` are equivalent. For UCA 4.0.0 collations, hexadecimal notation can be used only for characters in the Basic Multilingual Plane, not for characters outside the BMP range of `0000` to `FFFF`. For UCA 5.2.0 collations, hexadecimal notation can be used for any character.

The `Index.xml` file itself should be written using UTF-8 encoding.

##### Syntax Rules



LDML has reset rules and shift rules to specify character ordering. Orderings are given as a set of rules that begin with a reset rule that establishes an anchor point, followed by shift rules that indicate how characters sort relative to the anchor point.

- A `<reset>` rule does not specify any ordering in and of itself. Instead, it “resets” the ordering for subsequent shift rules to cause them to be taken in relation to a given character. Either of the following rules resets subsequent shift rules to be taken in relation to the letter 'A':

```
<reset>A</reset>
<reset>\u0041</reset>
```

- The `<p>`, `<s>`, and `<t>` shift rules define primary, secondary, and tertiary differences of a character from another character:
  - Use primary differences to distinguish separate letters.
  - Use secondary differences to distinguish accent variations.
  - Use tertiary differences to distinguish lettercase variations.

Either of these rules specifies a primary shift rule for the 'G' character:

```
<p>G</p>
<p>\u0047</p>
```

- The `<i>` shift rule indicates that one character sorts identically to another. The following rules cause 'b' to sort the same as 'a':

```
<reset>a</reset>
<i>b</i>
```

- Abbreviated shift syntax specifies multiple shift rules using a single pair of tags. The following table shows the correspondence between abbreviated syntax rules and the equivalent nonabbreviated rules.

**Table 10.2. Abbreviated Shift Syntax**

Abbreviated Syntax	Nonabbreviated Syntax
<code>&lt;pc&gt;xyz&lt;/pc&gt;</code>	<code>&lt;p&gt;x&lt;/p&gt;&lt;p&gt;y&lt;/p&gt;&lt;p&gt;z&lt;/p&gt;</code>
<code>&lt;sc&gt;xyz&lt;/sc&gt;</code>	<code>&lt;s&gt;x&lt;/s&gt;&lt;s&gt;y&lt;/s&gt;&lt;s&gt;z&lt;/s&gt;</code>
<code>&lt;tc&gt;xyz&lt;/tc&gt;</code>	<code>&lt;t&gt;x&lt;/t&gt;&lt;t&gt;y&lt;/t&gt;&lt;t&gt;z&lt;/t&gt;</code>
<code>&lt;ic&gt;xyz&lt;/ic&gt;</code>	<code>&lt;i&gt;x&lt;/i&gt;&lt;i&gt;y&lt;/i&gt;&lt;i&gt;z&lt;/i&gt;</code>

- An expansion is a reset rule that establishes an anchor point for a multiple-character sequence. MySQL supports expansions 2 to 6 characters long. The following rules put 'z' greater at the primary level than the sequence of three characters 'abc':

```
<reset>abc</reset>
<p>z</p>
```

- A contraction is a shift rule that sorts a multiple-character sequence. MySQL supports contractions 2 to 6 characters long. The following rules put the sequence of three characters 'xyz' greater at the primary level than 'a':

```
<reset>a</reset>
<p>xyz</p>
```

- Long expansions and long contractions can be used together. These rules put the sequence of three characters 'xyz' greater at the primary level than the sequence of three characters 'abc':

```
<reset>abc</reset>
```

```
<p>xyz</p>
```

- Normal expansion syntax uses `<x>` plus `<extend>` elements to specify an expansion. The following rules put the character 'k' greater at the secondary level than the sequence 'ch'. That is, 'k' behaves as if it expands to a character after 'c' followed by 'h':

```
<reset>c</reset>
<x><s>k</s><extend>h</extend></x>
```

This syntax permits long sequences. These rules sort the sequence 'ccs' greater at the tertiary level than the sequence 'cscs':

```
<reset>cs</reset>
<x><t>ccs</t><extend>cs</extend></x>
```

The LDML specification describes normal expansion syntax as “tricky.” See that specification for details.

- Previous context syntax uses `<x>` plus `<context>` elements to specify that the context before a character affects how it sorts. The following rules put '-' greater at the secondary level than 'a', but only when '-' occurs after 'b':

```
<reset>a</reset>
<x><context>b</context><s>-</s></x>
```

- Previous context syntax can include the `<extend>` element. These rules put 'def' greater at the primary level than 'aghi', but only when 'def' comes after 'abc':

```
<reset>a</reset>
<x><context>abc</context><p>def</p><extend>ghi</extend></x>
```

- Reset rules permit a `before` attribute. Normally, shift rules after a reset rule indicate characters that sort after the reset character. Shift rules after a reset rule that has the `before` attribute indicate characters that sort before the reset character. The following rules put the character 'b' immediately before 'a' at the primary level:

```
<reset before="primary">a</reset>
<p>b</p>
```

Permissible `before` attribute values specify the sort level by name or the equivalent numeric value:

```
<reset before="primary">
<reset before="1">

<reset before="secondary">
<reset before="2">

<reset before="tertiary">
<reset before="3">
```

- A reset rule can name a logical reset position rather than a literal character:

```
<first_tertiary_ignorable/>
<last_tertiary_ignorable/>
<first_secondary_ignorable/>
<last_secondary_ignorable/>
<first_primary_ignorable/>
<last_primary_ignorable/>
<first_variable/>
<last_variable/>
<first_non_ignorable/>
<last_non_ignorable/>
<first_trailing/>
<last_trailing/>
```

These rules put 'z' greater at the primary level than nonignorable characters that have a Default Unicode Collation Element Table (DUCET) entry and that are not CJK:

```
<reset><last_non_ignorable/></reset>
<p>z</p>
```

Logical positions have the code points shown in the following table.

**Table 10.3. Logical Reset Position Code Points**

Logical Position	Unicode 4.0.0 Code Point	Unicode 5.2.0 Code Point
<code>&lt;first_non_ignorable/&gt;</code>	U+02D0	U+02D0
<code>&lt;last_non_ignorable/&gt;</code>	U+A48C	U+1342E
<code>&lt;first_primary_ignorable/&gt;</code>	U+0332	U+0332
<code>&lt;last_primary_ignorable/&gt;</code>	U+20EA	U+101FD
<code>&lt;first_secondary_ignorable/&gt;</code>	U+0000	U+0000
<code>&lt;last_secondary_ignorable/&gt;</code>	U+FE73	U+FE73
<code>&lt;first_tertiary_ignorable/&gt;</code>	U+0000	U+0000
<code>&lt;last_tertiary_ignorable/&gt;</code>	U+FE73	U+FE73
<code>&lt;first_trailing/&gt;</code>	U+0000	U+0000
<code>&lt;last_trailing/&gt;</code>	U+0000	U+0000
<code>&lt;first_variable/&gt;</code>	U+0009	U+0009
<code>&lt;last_variable/&gt;</code>	U+2183	U+1D371

- The `<collation>` element permits a `shift-after-method` attribute that affects character weight calculation for shift rules. The attribute has these permitted values:
  - `simple`: Calculate character weights as for reset rules that do not have a `before` attribute. This is the default if the attribute is not given.
  - `expand`: Use expansions for shifts after reset rules.

Suppose that '0' and '1' have weights of 0E29 and 0E2A and we want to put all basic Latin letters between '0' and '1':

```
<reset>0</reset>
<pc>abcdefghijklmnopqrstuvwxyz</pc>
```

For simple shift mode, weights are calculated as follows:

```
'a' has weight 0E29+1
'b' has weight 0E29+2
'c' has weight 0E29+3
...
```

However, there are not enough vacant positions to put 26 characters between '0' and '1'. The result is that digits and letters are intermixed.

To solve this, use `shift-after-method="expand"`. Then weights are calculated like this:

```
'a' has weight [0E29][233D+1]
'b' has weight [0E29][233D+2]
'c' has weight [0E29][233D+3]
...
```

233D is the UCA 4.0.0 weight for character 0xA48C, which is the last nonignorable character (a sort of the greatest character in the collation, excluding CJK). UCA 5.2.0 is similar but uses 3ACA, for character 0x1342E.

In MySQL 5.6, an extension to LDML rules permits the `<collation>` element to include an optional `version` attribute in `<collation>` tags to indicate the UCA version on which the collation is based. If the `version` attribute is omitted, its default value is 4.0.0. For example, this specification indicates a collation that is based on UCA 5.2.0:

```
<collation id="nnn" name="utf8_xxx_ci" version="5.2.0">
...
</collation>
```

#### 10.4.4.3. Diagnostics During `Index.xml` Parsing

The MySQL server generates diagnostics when it finds problems while parsing the `Index.xml` file:

- Unknown tags are written to the error log. For example, the following message results if a collation definition contains a `<aaa>` tag:

```
[Warning] Buffered warning: Unknown LDML tag:
'charsets/charset/collation/rules/aaa'
```

- If collation initialization is not possible, the server reports an “Unknown collation” error, and also generates warnings explaining the problems, such as in the previous example. In other cases, when a collation description is generally correct but contains some unknown tags, the collation is initialized and is available for use. The unknown parts are ignored, but a warning is generated in the error log.
- Problems with collations generate warnings that clients can display with `SHOW WARNINGS`. Suppose that a reset rule contains an expansion longer than the maximum supported length of 6 characters:

```
<reset>abcdefghi</reset>
<i>x</i>
```

An attempt to use the collation produces warnings:

```
mysql> SELECT _utf8'test' COLLATE utf8_test_ci;
ERROR 1273 (HY000): Unknown collation: 'utf8_test_ci'
mysql> SHOW WARNINGS;
```

Level	Code	Message
Error	1273	Unknown collation: 'utf8_test_ci'
Warning	1273	Expansion is too long at 'abcdefghi=x'

## 10.5. Character Set Configuration

You can change the default server character set and collation with the `--character-set-server` [417] and `--collation-server` [417] options when you start the server. The collation must be a legal collation for the default character set. (Use the `SHOW COLLATION` statement to determine which collations are available for each character set.) See [Section 5.1.3, “Server Command Options”](#).

If you try to use a character set that is not compiled into your binary, you might run into the following problems:

- Your program uses an incorrect path to determine where the character sets are stored (which is typically the `share/mysql/charsets` or `share/charsets` directory under the MySQL installation directory). This can be fixed by using the `--character-sets-dir` option when you run the program in question. For example, to specify a directory to be used by MySQL client programs, list it in the `[client]` group of your option file. The examples given here show what the setting might look like for Unix or Windows, respectively:

```
[client]
character-sets-dir=/usr/local/mysql/share/mysql/charsets

[client]
character-sets-dir="C:/Program Files/MySQL/MySQL Server 5.6/share/charsets"
```

- The character set is a complex character set that cannot be loaded dynamically. In this case, you must recompile the program with support for the character set.

For Unicode character sets, you can define collations without recompiling by using LDML notation. See [Section 10.4.4, “Adding a UCA Collation to a Unicode Character Set”](#).

- The character set is a dynamic character set, but you do not have a configuration file for it. In this case, you should install the configuration file for the character set from a new MySQL distribution.
- If your character set index file does not contain the name for the character set, your program displays an error message. The file is named `Index.xml` and the message is:

```
Character set 'charset_name' is not a compiled character set and is not
specified in the '/usr/share/mysql/charsets/Index.xml' file
```

To solve this problem, you should either get a new index file or manually add the name of any missing character sets to the current file.

You can force client programs to use specific character set as follows:

```
[client]
default-character-set=charset_name
```

This is normally unnecessary. However, when `character_set_system` [469] differs from `character_set_server` [468] or `character_set_client` [467], and you input characters manually (as database object identifiers, column values, or both), these may be displayed incorrectly in output from the client or the output itself may be formatted incorrectly. In such cases, starting the mysql client with `--default-character-set=system_character_set` [249]—that is, setting the client character set to match the system character set—should fix the problem.

For MyISAM tables, you can check the character set name and number for a table with `myisamchk -dvv tbl_name`.

## 10.6. MySQL Server Time Zone Support

The MySQL server maintains several time zone settings:

- The system time zone. When the server starts, it attempts to determine the time zone of the host machine and uses it to set the `system_time_zone` [550] system variable. The value does not change thereafter.

You can set the system time zone for MySQL Server at startup with the `--timezone=timezone_name` [229] option to `mysqld_safe`. You can also set it by setting the `TZ` environment variable before you start `mysqld`. The permissible values for `--timezone` [229] or `TZ` are system dependent. Consult your operating system documentation to see what values are acceptable.

- The server's current time zone. The global `time_zone` [557] system variable indicates the time zone the server currently is operating in. The initial value for `time_zone` [557] is 'SYSTEM', which indicates that the server time zone is the same as the system time zone.

The initial global server time zone value can be specified explicitly at startup with the `--default-time-zone=timezone` [420] option on the command line, or you can use the following line in an option file:

```
default-time-zone='timezone'
```

If you have the `SUPER` [756] privilege, you can set the global server time zone value at runtime with this statement:

```
mysql> SET GLOBAL time_zone = timezone;
```

- Per-connection time zones. Each client that connects has its own time zone setting, given by the session `time_zone` [557] variable. Initially, the session variable takes its value from the global `time_zone` [557] variable, but the client can change its own time zone with this statement:

```
mysql> SET time_zone = timezone;
```

The current session time zone setting affects display and storage of time values that are zone-sensitive. This includes the values displayed by functions such as `NOW()` [1220] or `CURTIME()` [1213], and values stored in and retrieved from `TIMESTAMP` columns. Values for `TIMESTAMP` columns are converted from the current time zone to UTC for storage, and from UTC to the current time zone for retrieval.

The current time zone setting does not affect values displayed by functions such as `UTC_TIMESTAMP()` [1227] or values in `DATE`, `TIME`, or `DATETIME` columns. Nor are values in those data types stored in UTC; the time zone applies for them only when converting from `TIMESTAMP` values. If you want locale-specific arithmetic for `DATE`, `TIME`, or `DATETIME` values, convert them to UTC, perform the arithmetic, and then convert back.

The current values of the global and client-specific time zones can be retrieved like this:

```
mysql> SELECT @@global.time_zone, @@session.time_zone;
```

`timezone` values can be given in several formats, none of which are case sensitive:

- The value `'SYSTEM'` indicates that the time zone should be the same as the system time zone.
- The value can be given as a string indicating an offset from UTC, such as `'+10:00'` or `'-6:00'`.
- The value can be given as a named time zone, such as `'Europe/Helsinki'`, `'US/Eastern'`, or `'MET'`. Named time zones can be used only if the time zone information tables in the `mysql` database have been created and populated.

The MySQL installation procedure creates the time zone tables in the `mysql` database, but does not load them. You must do so manually using the following instructions. (If you are upgrading to MySQL 4.1.3 or later from an earlier version, you can create the tables by upgrading your `mysql` database. Use the instructions in [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#). After creating the tables, you can load them.)



#### Note

Loading the time zone information is not necessarily a one-time operation because the information changes occasionally. For example, the rules for Daylight Saving Time in the United States, Mexico, and parts of Canada changed in 2007. When such changes occur, applications that use the old rules become out of date and you may find it necessary to reload the time zone tables to keep the information used by your MySQL server current. See the notes at the end of this section.

If your system has its own `zoneinfo` database (the set of files describing time zones), you should use the `mysql_tzinfo_to_sql` program for filling the time zone tables. Examples of such systems are Linux, FreeBSD, Solaris, and Mac OS X. One likely location for these files is the `/usr/share/zoneinfo` directory. If your system does not have a `zoneinfo` database, you can use the downloadable package described later in this section.

The `mysql_tzinfo_to_sql` program is used to load the time zone tables. On the command line, pass the `zoneinfo` directory path name to `mysql_tzinfo_to_sql` and send the output into the `mysql` program. For example:

```
shell> mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root mysql
```

`mysql_tzinfo_to_sql` reads your system's time zone files and generates SQL statements from them. `mysql` processes those statements to load the time zone tables.

`mysql_tzinfo_to_sql` also can be used to load a single time zone file or to generate leap second information:

- To load a single time zone file `tz_file` that corresponds to a time zone name `tz_name`, invoke `mysql_tzinfo_to_sql` like this:

```
shell> mysql_tzinfo_to_sql tz_file tz_name | mysql -u root mysql
```

With this approach, you must execute a separate command to load the time zone file for each named zone that the server needs to know about.

- If your time zone needs to account for leap seconds, initialize the leap second information like this, where `tz_file` is the name of your time zone file:

```
shell> mysql_tzinfo_to_sql --leap tz_file | mysql -u root mysql
```

- After running `mysql_tzinfo_to_sql`, it is best to restart the server so that it does not continue to use any previously cached time zone data.

If your system is one that has no zoneinfo database (for example, Windows or HP-UX), you can use the package of pre-built time zone tables that is available for download at the MySQL Developer Zone:

<http://dev.mysql.com/downloads/timezones.html>

This time zone package contains `.frm`, `.MYD`, and `.MYI` files for the `MyISAM` time zone tables. These tables should be part of the `mysql` database, so you should place the files in the `mysql` subdirectory of your MySQL server's data directory. The server should be stopped while you do this and restarted afterward.



#### Warning

Do not use the downloadable package if your system has a zoneinfo database. Use the `mysql_tzinfo_to_sql` utility instead. Otherwise, you may cause a difference in datetime handling between MySQL and other applications on your system.

For information about time zone settings in replication setup, please see [Section 16.4.1, “Replication Features and Issues”](#).

## 10.6.1. Staying Current with Time Zone Changes

As mentioned earlier, when the time zone rules change, applications that use the old rules become out of date. To stay current, it is necessary to make sure that your system uses current time zone information is used. For MySQL, there are two factors to consider in staying current:

- The operating system time affects the value that the MySQL server uses for times if its time zone is set to `SYSTEM`. Make sure that your operating system is using the latest time zone information. For most operating systems, the latest update or service pack prepares your system for the time changes. Check the Web site for your operating system vendor for an update that addresses the time changes.
- If you replace the system's `/etc/localtime` timezone file with a version that uses rules differing from those in effect at `mysqld` startup, you should restart `mysqld` so that it uses the updated rules. Otherwise, `mysqld` might not notice when the system changes its time.
- If you use named time zones with MySQL, make sure that the time zone tables in the `mysql` database are up to date. If your system has its own zoneinfo database, you should reload the MySQL time zone tables whenever the zoneinfo database is updated, using the instructions given earlier in this section. For systems that do not have their own zoneinfo database, check the MySQL Developer Zone for updates. When a new update is available, download it and use it to replace your current time zone tables. `mysqld` caches time zone information that it looks up, so after replacing



the time zone tables, you should restart `mysqld` to make sure that it does not continue to serve outdated time zone data.

If you are uncertain whether named time zones are available, for use either as the server's time zone setting or by clients that set their own time zone, check whether your time zone tables are empty. The following query determines whether the table that contains time zone names has any rows:

```
mysql> SELECT COUNT(*) FROM mysql.time_zone_name;
+-----+
| COUNT(*) |
+-----+
|         0 |
+-----+
```

A count of zero indicates that the table is empty. In this case, no one can be using named time zones, and you don't need to update the tables. A count greater than zero indicates that the table is not empty and that its contents are available to be used for named time zone support. In this case, you should be sure to reload your time zone tables so that anyone who uses named time zones will get correct query results.

To check whether your MySQL installation is updated properly for a change in Daylight Saving Time rules, use a test like the one following. The example uses values that are appropriate for the 2007 DST 1-hour change that occurs in the United States on March 11 at 2 a.m.

The test uses these two queries:

```
SELECT CONVERT_TZ('2007-03-11 2:00:00','US/Eastern','US/Central');
SELECT CONVERT_TZ('2007-03-11 3:00:00','US/Eastern','US/Central');
```

The two time values indicate the times at which the DST change occurs, and the use of named time zones requires that the time zone tables be used. The desired result is that both queries return the same result (the input time, converted to the equivalent value in the 'US/Central' time zone).

Before updating the time zone tables, you would see an incorrect result like this:

```
mysql> SELECT CONVERT_TZ('2007-03-11 2:00:00','US/Eastern','US/Central');
+-----+
| CONVERT_TZ('2007-03-11 2:00:00','US/Eastern','US/Central') |
+-----+
| 2007-03-11 01:00:00 |
+-----+

mysql> SELECT CONVERT_TZ('2007-03-11 3:00:00','US/Eastern','US/Central');
+-----+
| CONVERT_TZ('2007-03-11 3:00:00','US/Eastern','US/Central') |
+-----+
| 2007-03-11 02:00:00 |
+-----+
```

After updating the tables, you should see the correct result:

```
mysql> SELECT CONVERT_TZ('2007-03-11 2:00:00','US/Eastern','US/Central');
+-----+
| CONVERT_TZ('2007-03-11 2:00:00','US/Eastern','US/Central') |
+-----+
| 2007-03-11 01:00:00 |
+-----+

mysql> SELECT CONVERT_TZ('2007-03-11 3:00:00','US/Eastern','US/Central');
+-----+
| CONVERT_TZ('2007-03-11 3:00:00','US/Eastern','US/Central') |
+-----+
| 2007-03-11 01:00:00 |
+-----+
```

## 10.6.2. Time Zone Leap Second Support

Leap second values are returned with a time part that ends with `:59:59`. This means that a function such as `NOW()` [1220] can return the same value for two or three consecutive seconds during the leap second. It remains true that literal temporal values having a time part that ends with `:59:60` or `:59:61` are considered invalid.

If it is necessary to search for `TIMESTAMP` values one second before the leap second, anomalous results may be obtained if you use a comparison with `'YYYY-MM-DD hh:mm:ss'` values. The following example demonstrates this. It changes the local time zone to UTC so there is no difference between internal values (which are in UTC) and displayed values (which have time zone correction applied).

```
mysql> CREATE TABLE t1 (
->   a INT,
->   ts TIMESTAMP DEFAULT NOW(),
->   PRIMARY KEY (ts)
-> );
Query OK, 0 rows affected (0.01 sec)

mysql> -- change to UTC
mysql> SET time_zone = '+00:00';
Query OK, 0 rows affected (0.00 sec)

mysql> -- Simulate NOW() = '2008-12-31 23:59:59'
mysql> SET timestamp = 1230767999;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t1 (a) VALUES (1);
Query OK, 1 row affected (0.00 sec)

mysql> -- Simulate NOW() = '2008-12-31 23:59:60'
mysql> SET timestamp = 1230768000;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t1 (a) VALUES (2);
Query OK, 1 row affected (0.00 sec)

mysql> -- values differ internally but display the same
mysql> SELECT a, ts, UNIX_TIMESTAMP(ts) FROM t1;
+-----+-----+-----+
| a    | ts                | UNIX_TIMESTAMP(ts) |
+-----+-----+-----+
| 1    | 2008-12-31 23:59:59 | 1230767999         |
| 2    | 2008-12-31 23:59:59 | 1230768000         |
+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> -- only the non-leap value matches
mysql> SELECT * FROM t1 WHERE ts = '2008-12-31 23:59:59';
+-----+-----+
| a    | ts                |
+-----+-----+
| 1    | 2008-12-31 23:59:59 |
+-----+-----+
1 row in set (0.00 sec)

mysql> -- the leap value with seconds=60 is invalid
mysql> SELECT * FROM t1 WHERE ts = '2008-12-31 23:59:60';
Empty set, 2 warnings (0.00 sec)
```

To work around this, you can use a comparison based on the UTC value actually stored in column, which has the leap second correction applied:

```
mysql> -- selecting using UNIX_TIMESTAMP value return leap value
mysql> SELECT * FROM t1 WHERE UNIX_TIMESTAMP(ts) = 1230768000;
+-----+-----+
| a    | ts                |
+-----+-----+
| 2    | 2008-12-31 23:59:59 |
+-----+-----+
1 row in set (0.00 sec)
```

## 10.7. MySQL Server Locale Support

The locale indicated by the `lc_time_names` [495] system variable controls the language used to display day and month names and abbreviations. This variable affects the output from the `DATE_FORMAT()` [1215], `DAYNAME()` [1217], and `MONTHNAME()` [1220] functions.

`lc_time_names` [495] does not affect the `STR_TO_DATE()` [1221] or `GET_FORMAT()` [1218] function.

The `lc_time_names` [495] value does not affect the result from `FORMAT()` [1180], but this function takes an optional third parameter that enables a locale to be specified to be used for the result number's decimal point, thousands separator, and grouping between separators. Permissible locale values are the same as the legal values for the `lc_time_names` [495] system variable.

Locale names have language and region subtags listed by IANA (<http://www.iana.org/assignments/language-subtag-registry>) such as `'ja_JP'` or `'pt_BR'`. The default value is `'en_US'` regardless of your system's locale setting, but you can set the value at server startup or set the `GLOBAL` value if you have the `SUPER` [756] privilege. Any client can examine the value of `lc_time_names` [495] or set its `SESSION` value to affect the locale for its own connection.

```
mysql> SET NAMES 'utf8';
Query OK, 0 rows affected (0.09 sec)

mysql> SELECT @@lc_time_names;
+-----+
| @@lc_time_names |
+-----+
| en_US           |
+-----+
1 row in set (0.00 sec)

mysql> SELECT DAYNAME('2010-01-01'), MONTHNAME('2010-01-01');
+-----+-----+
| DAYNAME('2010-01-01') | MONTHNAME('2010-01-01') |
+-----+-----+
| Friday                | January                  |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT DATE_FORMAT('2010-01-01','%W %a %M %b');
+-----+
| DATE_FORMAT('2010-01-01','%W %a %M %b') |
+-----+
| Friday Fri January Jan                    |
+-----+
1 row in set (0.00 sec)

mysql> SET lc_time_names = 'es_MX';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @@lc_time_names;
+-----+
| @@lc_time_names |
+-----+
| es_MX           |
+-----+
1 row in set (0.00 sec)

mysql> SELECT DAYNAME('2010-01-01'), MONTHNAME('2010-01-01');
+-----+-----+
| DAYNAME('2010-01-01') | MONTHNAME('2010-01-01') |
+-----+-----+
| viernes                | enero                    |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT DATE_FORMAT('2010-01-01','%W %a %M %b');
+-----+
| DATE_FORMAT('2010-01-01','%W %a %M %b') |
+-----+
```

```
| DATE_FORMAT('2010-01-01', '%W %a %M %b') |
+-----+
| viernes vie enero ene |
+-----+
1 row in set (0.00 sec)
```

The day or month name for each of the affected functions is converted from `utf8` to the character set indicated by the `character_set_connection` [467] system variable.

`lc_time_names` [495] may be set to any of the following locale values. The set of locales supported by MySQL may differ from those supported by your operating system.

<code>ar_AE</code> : Arabic - United Arab Emirates	<code>ar_BH</code> : Arabic - Bahrain
<code>ar_DZ</code> : Arabic - Algeria	<code>ar_EG</code> : Arabic - Egypt
<code>ar_IN</code> : Arabic - India	<code>ar_IQ</code> : Arabic - Iraq
<code>ar_JO</code> : Arabic - Jordan	<code>ar_KW</code> : Arabic - Kuwait
<code>ar_LB</code> : Arabic - Lebanon	<code>ar_LY</code> : Arabic - Libya
<code>ar_MA</code> : Arabic - Morocco	<code>ar_OM</code> : Arabic - Oman
<code>ar_QA</code> : Arabic - Qatar	<code>ar_SA</code> : Arabic - Saudi Arabia
<code>ar_SD</code> : Arabic - Sudan	<code>ar_SY</code> : Arabic - Syria
<code>ar_TN</code> : Arabic - Tunisia	<code>ar_YE</code> : Arabic - Yemen
<code>be_BY</code> : Belarusian - Belarus	<code>bg_BG</code> : Bulgarian - Bulgaria
<code>ca_ES</code> : Catalan - Spain	<code>cs_CZ</code> : Czech - Czech Republic
<code>da_DK</code> : Danish - Denmark	<code>de_AT</code> : German - Austria
<code>de_BE</code> : German - Belgium	<code>de_CH</code> : German - Switzerland
<code>de_DE</code> : German - Germany	<code>de_LU</code> : German - Luxembourg
<code>el_GR</code> : Greek - Greece	<code>en_AU</code> : English - Australia
<code>en_CA</code> : English - Canada	<code>en_GB</code> : English - United Kingdom
<code>en_IN</code> : English - India	<code>en_NZ</code> : English - New Zealand
<code>en_PH</code> : English - Philippines	<code>en_US</code> : English - United States
<code>en_ZA</code> : English - South Africa	<code>en_ZW</code> : English - Zimbabwe
<code>es_AR</code> : Spanish - Argentina	<code>es_BO</code> : Spanish - Bolivia
<code>es_CL</code> : Spanish - Chile	<code>es_CO</code> : Spanish - Columbia
<code>es_CR</code> : Spanish - Costa Rica	<code>es_DO</code> : Spanish - Dominican Republic
<code>es_EC</code> : Spanish - Ecuador	<code>es_ES</code> : Spanish - Spain
<code>es_GT</code> : Spanish - Guatemala	<code>es_HN</code> : Spanish - Honduras
<code>es_MX</code> : Spanish - Mexico	<code>es_NI</code> : Spanish - Nicaragua
<code>es_PA</code> : Spanish - Panama	<code>es_PE</code> : Spanish - Peru
<code>es_PR</code> : Spanish - Puerto Rico	<code>es_PY</code> : Spanish - Paraguay
<code>es_SV</code> : Spanish - El Salvador	<code>es_US</code> : Spanish - United States
<code>es_UY</code> : Spanish - Uruguay	<code>es_VE</code> : Spanish - Venezuela
<code>et_EE</code> : Estonian - Estonia	<code>eu_ES</code> : Basque - Basque
<code>fi_FI</code> : Finnish - Finland	<code>fo_FO</code> : Faroese - Faroe Islands
<code>fr_BE</code> : French - Belgium	<code>fr_CA</code> : French - Canada
<code>fr_CH</code> : French - Switzerland	<code>fr_FR</code> : French - France
<code>fr_LU</code> : French - Luxembourg	<code>gl_ES</code> : Galician - Spain

<a href="#">gu_IN</a> : Gujarati - India	<a href="#">he_IL</a> : Hebrew - Israel
<a href="#">hi_IN</a> : Hindi - India	<a href="#">hr_HR</a> : Croatian - Croatia
<a href="#">hu_HU</a> : Hungarian - Hungary	<a href="#">id_ID</a> : Indonesian - Indonesia
<a href="#">is_IS</a> : Icelandic - Iceland	<a href="#">it_CH</a> : Italian - Switzerland
<a href="#">it_IT</a> : Italian - Italy	<a href="#">ja_JP</a> : Japanese - Japan
<a href="#">ko_KR</a> : Korean - Republic of Korea	<a href="#">lt_LT</a> : Lithuanian - Lithuania
<a href="#">lv_LV</a> : Latvian - Latvia	<a href="#">mk_MK</a> : Macedonian - FYROM
<a href="#">mn_MN</a> : Mongolia - Mongolian	<a href="#">ms_MY</a> : Malay - Malaysia
<a href="#">nb_NO</a> : Norwegian(Bokmål) - Norway	<a href="#">nl_BE</a> : Dutch - Belgium
<a href="#">nl_NL</a> : Dutch - The Netherlands	<a href="#">no_NO</a> : Norwegian - Norway
<a href="#">pl_PL</a> : Polish - Poland	<a href="#">pt_BR</a> : Portugese - Brazil
<a href="#">pt_PT</a> : Portugese - Portugal	<a href="#">rm_CH</a> : Romansh - Switzerland
<a href="#">ro_RO</a> : Romanian - Romania	<a href="#">ru_RU</a> : Russian - Russia
<a href="#">ru_UA</a> : Russian - Ukraine	<a href="#">sk_SK</a> : Slovak - Slovakia
<a href="#">sl_SI</a> : Slovenian - Slovenia	<a href="#">sq_AL</a> : Albanian - Albania
<a href="#">sr_RS</a> : Serbian - Yugoslavia	<a href="#">sv_FI</a> : Swedish - Finland
<a href="#">sv_SE</a> : Swedish - Sweden	<a href="#">ta_IN</a> : Tamil - India
<a href="#">te_IN</a> : Telugu - India	<a href="#">th_TH</a> : Thai - Thailand
<a href="#">tr_TR</a> : Turkish - Turkey	<a href="#">uk_UA</a> : Ukrainian - Ukraine
<a href="#">ur_PK</a> : Urdu - Pakistan	<a href="#">vi_VN</a> : Vietnamese - Viet Nam
<a href="#">zh_CN</a> : Chinese - China	<a href="#">zh_HK</a> : Chinese - Hong Kong
<a href="#">zh_TW</a> : Chinese - Taiwan Province of China	



---

# Chapter 11. Data Types

## Table of Contents

11.1. Data Type Overview .....	1112
11.1.1. Numeric Type Overview .....	1112
11.1.2. Date and Time Type Overview .....	1115
11.1.3. String Type Overview .....	1117
11.2. Numeric Types .....	1120
11.2.1. Integer Types (Exact Value) - <code>INTEGER</code> , <code>INT</code> , <code>SMALLINT</code> , <code>TINYINT</code> , <code>MEDIUMINT</code> , <code>BIGINT</code> .....	1121
11.2.2. Fixed-Point Types (Exact Value) - <code>DECIMAL</code> , <code>NUMERIC</code> .....	1121
11.2.3. Floating-Point Types (Approximate Value) - <code>FLOAT</code> , <code>DOUBLE</code> .....	1122
11.2.4. Bit-Value Type - <code>BIT</code> .....	1122
11.2.5. Numeric Type Attributes .....	1122
11.2.6. Out-of-Range and Overflow Handling .....	1123
11.3. Date and Time Types .....	1124
11.3.1. The <code>DATE</code> , <code>DATETIME</code> , and <code>TIMESTAMP</code> Types .....	1126
11.3.2. The <code>TIME</code> Type .....	1127
11.3.3. The <code>YEAR</code> Type .....	1128
11.3.4. <code>YEAR(2)</code> Limitations and Migrating to <code>YEAR(4)</code> .....	1128
11.3.5. Automatic Initialization and Updating for <code>TIMESTAMP</code> and <code>DATETIME</code> .....	1131
11.3.6. Fractional Seconds in Time Values .....	1134
11.3.7. Conversion Between Date and Time Types .....	1135
11.3.8. Two-Digit Years in Dates .....	1137
11.4. String Types .....	1137
11.4.1. The <code>CHAR</code> and <code>VARCHAR</code> Types .....	1137
11.4.2. The <code>BINARY</code> and <code>VARBINARY</code> Types .....	1139
11.4.3. The <code>BLOB</code> and <code>TEXT</code> Types .....	1140
11.4.4. The <code>ENUM</code> Type .....	1141
11.4.5. The <code>SET</code> Type .....	1144
11.5. Data Type Default Values .....	1146
11.6. Data Type Storage Requirements .....	1147
11.7. Choosing the Right Type for a Column .....	1151
11.8. Using Data Types from Other Database Engines .....	1151

MySQL supports a number of [SQL](#) data types in several categories: numeric types, date and time types, and string (character and byte) types. This chapter provides an overview of these data types, a more detailed description of the properties of the types in each category, and a summary of the data type storage requirements. The initial overview is intentionally brief. The more detailed descriptions later in the chapter should be consulted for additional information about particular data types, such as the permissible formats in which you can specify values.

MySQL also supports extensions for handling spatial data. For information about these data types, see [Section 12.18, “Spatial Extensions”](#).

Data type descriptions use these conventions:

- *M* indicates the maximum display width for integer types. For floating-point and fixed-point types, *M* is the total number of digits that can be stored (the precision). For string types, *M* is the maximum length. The maximum permissible value of *M* depends on the data type.
- *D* applies to floating-point and fixed-point types and indicates the number of digits following the decimal point (the scale). The maximum possible value is 30, but should be no greater than *M*–2.
- *fsp* applies to the `TIME`, `DATETIME`, and `TIMESTAMP` types and represents fractional seconds precision; that is, the number of digits following the decimal point for fractional parts of seconds. The *fsp* value, if given, must be in the range 0 to 6. A value of 0 signifies that there is no fractional part.



If omitted, the default precision is 0. (This differs from the standard SQL default of 6, for compatibility with previous MySQL versions.)

- Square brackets (“[” and “]”) indicate optional parts of type definitions.

## 11.1. Data Type Overview

### 11.1.1. Numeric Type Overview

A summary of the numeric data types follows. For additional information about properties and storage requirements of the numeric types, see [Section 11.2, “Numeric Types”](#), and [Section 11.6, “Data Type Storage Requirements”](#).

*M* indicates the maximum display width for integer types. The maximum display width is 255. Display width is unrelated to the range of values a type can contain, as described in [Section 11.2, “Numeric Types”](#). For floating-point and fixed-point types, *M* is the total number of digits that can be stored.

If you specify `ZEROFILL` for a numeric column, MySQL automatically adds the `UNSIGNED` attribute to the column.

Numeric data types that permit the `UNSIGNED` attribute also permit `SIGNED`. However, these data types are signed by default, so the `SIGNED` attribute has no effect.

`SERIAL` is an alias for `BIGINT UNSIGNED NOT NULL AUTO_INCREMENT UNIQUE`.

`SERIAL DEFAULT VALUE` in the definition of an integer column is an alias for `NOT NULL AUTO_INCREMENT UNIQUE`.



#### Warning

When you use subtraction between integer values where one is of type `UNSIGNED`, the result is unsigned unless the `NO_UNSIGNED_SUBTRACTION [605]` SQL mode is enabled. See [Section 12.10, “Cast Functions and Operators”](#).

- `BIT[ (M) ]`

A bit-field type. *M* indicates the number of bits per value, from 1 to 64. The default is 1 if *M* is omitted.

- `TINYINT[ (M) ] [UNSIGNED] [ZEROFILL]`

A very small integer. The signed range is `-128` to `127`. The unsigned range is `0` to `255`.

- `BOOL`, `BOOLEAN`

These types are synonyms for `TINYINT(1)`. A value of zero is considered false. Nonzero values are considered true:

```
mysql> SELECT IF(0, 'true', 'false');
+-----+
| IF(0, 'true', 'false') |
+-----+
| false                   |
+-----+

mysql> SELECT IF(1, 'true', 'false');
+-----+
| IF(1, 'true', 'false') |
+-----+
| true                   |
+-----+

mysql> SELECT IF(2, 'true', 'false');
+-----+
| IF(2, 'true', 'false') |
+-----+
```

```
+-----+
| true  |
+-----+
```

However, the values `TRUE` and `FALSE` are merely aliases for `1` and `0`, respectively, as shown here:

```
mysql> SELECT IF(0 = FALSE, 'true', 'false');
+-----+
| IF(0 = FALSE, 'true', 'false') |
+-----+
| true                            |
+-----+

mysql> SELECT IF(1 = TRUE, 'true', 'false');
+-----+
| IF(1 = TRUE, 'true', 'false')  |
+-----+
| true                            |
+-----+

mysql> SELECT IF(2 = TRUE, 'true', 'false');
+-----+
| IF(2 = TRUE, 'true', 'false')  |
+-----+
| false                           |
+-----+

mysql> SELECT IF(2 = FALSE, 'true', 'false');
+-----+
| IF(2 = FALSE, 'true', 'false') |
+-----+
| false                           |
+-----+
```

The last two statements display the results shown because `2` is equal to neither `1` nor `0`.

- `SMALLINT[(M)] [UNSIGNED] [ZEROFILL]`

A small integer. The signed range is `-32768` to `32767`. The unsigned range is `0` to `65535`.

- `MEDIUMINT[(M)] [UNSIGNED] [ZEROFILL]`

A medium-sized integer. The signed range is `-8388608` to `8388607`. The unsigned range is `0` to `16777215`.

- `INT[(M)] [UNSIGNED] [ZEROFILL]`

A normal-size integer. The signed range is `-2147483648` to `2147483647`. The unsigned range is `0` to `4294967295`.

- `INTEGER[(M)] [UNSIGNED] [ZEROFILL]`

This type is a synonym for `INT`.

- `BIGINT[(M)] [UNSIGNED] [ZEROFILL]`

A large integer. The signed range is `-9223372036854775808` to `9223372036854775807`. The unsigned range is `0` to `18446744073709551615`.

`SERIAL` is an alias for `BIGINT UNSIGNED NOT NULL AUTO_INCREMENT UNIQUE`.

Some things you should be aware of with respect to `BIGINT` columns:

- All arithmetic is done using signed `BIGINT` or `DOUBLE` values, so you should not use unsigned big integers larger than `9223372036854775807` (63 bits) except with bit functions! If you do that, some of the last digits in the result may be wrong because of rounding errors when converting a `BIGINT` value to a `DOUBLE`.

MySQL can handle `BIGINT` in the following cases:

- When using integers to store large unsigned values in a `BIGINT` column.
- In `MIN(col_name) [1284]` or `MAX(col_name) [1284]`, where `col_name` refers to a `BIGINT` column.
- When using operators (`+ [1200]`, `- [1200]`, `* [1200]`, and so on) where both operands are integers.
- You can always store an exact integer value in a `BIGINT` column by storing it using a string. In this case, MySQL performs a string-to-number conversion that involves no intermediate double-precision representation.
- The `- [1200]`, `+ [1200]`, and `* [1200]` operators use `BIGINT` arithmetic when both operands are integer values. This means that if you multiply two big integers (or results from functions that return integers), you may get unexpected results when the result is larger than `9223372036854775807`.
- `DECIMAL[(M[,D])] [UNSIGNED] [ZEROFILL]`

A packed “exact” fixed-point number. `M` is the total number of digits (the precision) and `D` is the number of digits after the decimal point (the scale). The decimal point and (for negative numbers) the “-” sign are not counted in `M`. If `D` is 0, values have no decimal point or fractional part. The maximum number of digits (`M`) for `DECIMAL` is 65. The maximum number of supported decimals (`D`) is 30. If `D` is omitted, the default is 0. If `M` is omitted, the default is 10.

`UNSIGNED`, if specified, disallows negative values.

All basic calculations (`+`, `-`, `*`, `/`) with `DECIMAL` columns are done with a precision of 65 digits.

- `DEC[(M[,D])] [UNSIGNED] [ZEROFILL]`, `NUMERIC[(M[,D])] [UNSIGNED] [ZEROFILL]`, `FIXED[(M[,D])] [UNSIGNED] [ZEROFILL]`

These types are synonyms for `DECIMAL`. The `FIXED` synonym is available for compatibility with other database systems.

- `FLOAT[(M,D)] [UNSIGNED] [ZEROFILL]`

A small (single-precision) floating-point number. Permissible values are `-3.402823466E+38` to `-1.175494351E-38`, 0, and `1.175494351E-38` to `3.402823466E+38`. These are the theoretical limits, based on the IEEE standard. The actual range might be slightly smaller depending on your hardware or operating system.

`M` is the total number of digits and `D` is the number of digits following the decimal point. If `M` and `D` are omitted, values are stored to the limits permitted by the hardware. A single-precision floating-point number is accurate to approximately 7 decimal places.

`UNSIGNED`, if specified, disallows negative values.

Using `FLOAT` might give you some unexpected problems because all calculations in MySQL are done with double precision. See [Section C.5.5.7, “Solving Problems with No Matching Rows”](#).

- `DOUBLE[(M,D)] [UNSIGNED] [ZEROFILL]`

A normal-size (double-precision) floating-point number. Permissible values are `-1.7976931348623157E+308` to `-2.2250738585072014E-308`, 0, and `2.2250738585072014E-308` to `1.7976931348623157E+308`. These are the theoretical limits, based on the IEEE standard. The actual range might be slightly smaller depending on your hardware or operating system.

$M$  is the total number of digits and  $D$  is the number of digits following the decimal point. If  $M$  and  $D$  are omitted, values are stored to the limits permitted by the hardware. A double-precision floating-point number is accurate to approximately 15 decimal places.

`UNSIGNED`, if specified, disallows negative values.

- `DOUBLE PRECISION[( $M$ , $D$ )] [UNSIGNED] [ZEROFILL], REAL[( $M$ , $D$ )] [UNSIGNED] [ZEROFILL]`

These types are synonyms for `DOUBLE`. Exception: If the `REAL_AS_FLOAT` [607] SQL mode is enabled, `REAL` is a synonym for `FLOAT` rather than `DOUBLE`.

- `FLOAT( $p$ ) [UNSIGNED] [ZEROFILL]`

A floating-point number.  $p$  represents the precision in bits, but MySQL uses this value only to determine whether to use `FLOAT` or `DOUBLE` for the resulting data type. If  $p$  is from 0 to 24, the data type becomes `FLOAT` with no  $M$  or  $D$  values. If  $p$  is from 25 to 53, the data type becomes `DOUBLE` with no  $M$  or  $D$  values. The range of the resulting column is the same as for the single-precision `FLOAT` or double-precision `DOUBLE` data types described earlier in this section.

`FLOAT( $p$ )` syntax is provided for ODBC compatibility.

## 11.1.2. Date and Time Type Overview

A summary of the temporal data types follows. For additional information about properties and storage requirements of the temporal types, see [Section 11.3, “Date and Time Types”](#), and [Section 11.6, “Data Type Storage Requirements”](#). For descriptions of functions that operate on temporal values, see [Section 12.7, “Date and Time Functions”](#).

For the `DATE` and `DATETIME` range descriptions, “supported” means that although earlier values might work, there is no guarantee.

MySQL 5.6.4 and up permits fractional seconds for `TIME`, `DATETIME`, and `TIMESTAMP` values, with up to microseconds (6 digits) precision. To define a column that includes a fractional seconds part, use the syntax `type_name( $fsp$ )`, where `type_name` is `TIME`, `DATETIME`, or `TIMESTAMP`, and  $fsp$  is the fractional seconds precision. For example:

```
CREATE TABLE t1 (t TIME(3), dt DATETIME(6));
```

The  $fsp$  value, if given, must be in the range 0 to 6. A value of 0 signifies that there is no fractional part. If omitted, the default precision is 0. (This differs from the standard SQL default of 6, for compatibility with previous MySQL versions.)

MySQL 5.6.5 introduces expanded automatic initialization and updating of temporal types. Any `TIMESTAMP` column in a table can have these properties, rather than at most one column per table. In addition, these properties are now available for `DATETIME` columns.

The `YEAR(2)` data type has certain issues that you should consider before choosing to use it. As of MySQL 5.6.6, `YEAR(2)` is deprecated. `YEAR(2)` columns in existing tables are treated as before, but `YEAR(2)` in new or altered tables are converted to `YEAR(4)`. For more information, see [Section 11.3.4, “YEAR\(2\) Limitations and Migrating to YEAR\(4\)”](#).

- `DATE`

A date. The supported range is '1000-01-01' to '9999-12-31'. MySQL displays `DATE` values in 'YYYY-MM-DD' format, but permits assignment of values to `DATE` columns using either strings or numbers.

- `DATETIME[( $fsp$ )]`

A date and time combination. The supported range is '1000-01-01 00:00:00.000000' to '9999-12-31 23:59:59.999999'. MySQL displays `DATETIME` values in 'YYYY-MM-DD

`HH:MM:SS[.fraction]` format, but permits assignment of values to `DATETIME` columns using either strings or numbers.

As of MySQL 5.6.4, an optional `fsp` value in the range from 0 to 6 may be given to specify fractional seconds precision. A value of 0 signifies that there is no fractional part. If omitted, the default precision is 0.

As of MySQL 5.6.5, automatic initialization and updating to the current date and time for `DATETIME` columns can be specified using `DEFAULT` and `ON UPDATE` column definition clauses, as described in [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#).

- `TIMESTAMP[ (fsp) ]`

A timestamp. The range is `'1970-01-01 00:00:01.000000'` UTC to `'2038-01-19 03:14:07.999999'` UTC. `TIMESTAMP` values are stored as the number of seconds since the epoch (`'1970-01-01 00:00:00'` UTC). A `TIMESTAMP` cannot represent the value `'1970-01-01 00:00:00'` because that is equivalent to 0 seconds from the epoch and the value 0 is reserved for representing `'0000-00-00 00:00:00'`, the “zero” `TIMESTAMP` value.

As of MySQL 5.6.4, an optional `fsp` value in the range from 0 to 6 may be given to specify fractional seconds precision. A value of 0 signifies that there is no fractional part. If omitted, the default precision is 0.

The way the server handles `TIMESTAMP` definitions depends on the value of the `explicit_defaults_for_timestamp` [\[481\]](#) system variable (see [Section 5.1.4, “Server System Variables”](#)). By default, `explicit_defaults_for_timestamp` [\[481\]](#) is disabled and the server handles `TIMESTAMP` as follows:

Unless specified otherwise, the first `TIMESTAMP` column in a table is defined to be automatically set to the date and time of the most recent modification if not explicitly assigned a value. This makes `TIMESTAMP` useful for recording the timestamp of an `INSERT` or `UPDATE` operation. You can also set any `TIMESTAMP` column to the current date and time by assigning it a `NULL` value, unless it has been defined with the `NULL` attribute to permit `NULL` values.

Automatic initialization and updating to the current date and time can be specified using `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` column definition clauses. By default, the first `TIMESTAMP` column has these properties, as previously noted. As of MySQL 5.6.5, any `TIMESTAMP` column in a table can be defined to have these properties. Before 5.6.5, at most one `TIMESTAMP` column per table can have them, but it is possible to suppress them for the first column and instead assign them to a different `TIMESTAMP` column. See [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#).

If `explicit_defaults_for_timestamp` [\[481\]](#) is enabled, there is no automatic assignment of the `DEFAULT CURRENT_TIMESTAMP` or `ON UPDATE CURRENT_TIMESTAMP` attributes to any `TIMESTAMP` column. They must be included explicitly in the column definition. Also, any `TIMESTAMP` not explicitly declared as `NOT NULL` permits `NULL` values.

`explicit_defaults_for_timestamp` [\[481\]](#) is available as of MySQL 5.6.6. Before 5.6.6, the server handles `TIMESTAMP` as discussed for `explicit_defaults_for_timestamp` [\[481\]](#) disabled. Those behaviors, while they remain the default, are nonstandard and are deprecated as of 5.6.6. For discussion regarding upgrading to an installation with `explicit_defaults_for_timestamp` [\[481\]](#) enabled, see [Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”](#).

- `TIME[ (fsp) ]`

A time. The range is `'-838:59:59.000000'` to `'838:59:59.000000'`. MySQL displays `TIME` values in `'HH:MM:SS[.fraction]'` format, but permits assignment of values to `TIME` columns using either strings or numbers.

As of MySQL 5.6.4, an optional *fsp* value in the range from 0 to 6 may be given to specify fractional seconds precision. A value of 0 signifies that there is no fractional part. If omitted, the default precision is 0.

- `YEAR[(2|4)]`

A year in two-digit or four-digit format. The default is four-digit format. `YEAR(2)` or `YEAR(4)` differ in display format, but have the same range of values. In four-digit format, values display as 1901 to 2155, and 0000. In two-digit format, values display as 70 to 69, representing years from 1970 to 2069. MySQL displays `YEAR` values in `YYYY` or `YY` format, but permits assignment of values to `YEAR` columns using either strings or numbers.



#### Note

The `YEAR(2)` data type has certain issues that you should consider before choosing to use it. As of MySQL 5.6.6, `YEAR(2)` is deprecated. `YEAR(2)` columns in existing tables are treated as before, but `YEAR(2)` in new or altered tables are converted to `YEAR(4)`. For more information, see [Section 11.3.4, “YEAR\(2\) Limitations and Migrating to YEAR\(4\)”](#).

For additional information about `YEAR` display format and interpretation of input values, see [Section 11.3.3, “The YEAR Type”](#).

The `SUM()` [1285] and `AVG()` [1282] aggregate functions do not work with temporal values. (They convert the values to numbers, losing everything after the first nonnumeric character.) To work around this problem, convert to numeric units, perform the aggregate operation, and convert back to a temporal value. Examples:

```
SELECT SEC_TO_TIME(SUM(TIME_TO_SEC(time_col))) FROM tbl_name;  
SELECT FROM_DAYS(SUM(TO_DAYS(date_col))) FROM tbl_name;
```



#### Note

The MySQL server can be run with the `MAXDB` [608] SQL mode enabled. In this case, `TIMESTAMP` is identical with `DATETIME`. If this mode is enabled at the time that a table is created, `TIMESTAMP` columns are created as `DATETIME` columns. As a result, such columns use `DATETIME` display format, have the same range of values, and there is no automatic initialization or updating to the current date and time. See [Section 5.1.7, “Server SQL Modes”](#).

## 11.1.3. String Type Overview

A summary of the string data types follows. For additional information about properties and storage requirements of the string types, see [Section 11.4, “String Types”](#), and [Section 11.6, “Data Type Storage Requirements”](#).

In some cases, MySQL may change a string column to a type different from that given in a `CREATE TABLE` or `ALTER TABLE` statement. See [Section 13.1.14.3, “Silent Column Specification Changes”](#).

MySQL interprets length specifications in character column definitions in character units. This applies to `CHAR`, `VARCHAR`, and the `TEXT` types.

Column definitions for many string data types can include attributes that specify the character set or collation of the column. These attributes apply to the `CHAR`, `VARCHAR`, the `TEXT` types, `ENUM`, and `SET` data types:

- The `CHARACTER SET` attribute specifies the character set, and the `COLLATE` attribute specifies a collation for the character set. For example:

```
CREATE TABLE t  
(  
  c1 VARCHAR(20) CHARACTER SET utf8,
```

```
c2 TEXT CHARACTER SET latin1 COLLATE latin1_general_cs
);
```

This table definition creates a column named `c1` that has a character set of `utf8` with the default collation for that character set, and a column named `c2` that has a character set of `latin1` and a case-sensitive collation.

The rules for assigning the character set and collation when either or both of the `CHARACTER SET` and `COLLATE` attributes are missing are described in [Section 10.1.3.4, “Column Character Set and Collation”](#).

`CHARSET` is a synonym for `CHARACTER SET`.

- Specifying the `CHARACTER SET binary` attribute for a character data type causes the column to be created as the corresponding binary data type: `CHAR` becomes `BINARY`, `VARCHAR` becomes `VARBINARY`, and `TEXT` becomes `BLOB`. For the `ENUM` and `SET` data types, this does not occur; they are created as declared. Suppose that you specify a table using this definition:

```
CREATE TABLE t
(
  c1 VARCHAR(10) CHARACTER SET binary,
  c2 TEXT CHARACTER SET binary,
  c3 ENUM('a','b','c') CHARACTER SET binary
);
```

The resulting table has this definition:

```
CREATE TABLE t
(
  c1 VARBINARY(10),
  c2 BLOB,
  c3 ENUM('a','b','c') CHARACTER SET binary
);
```

- The `ASCII` attribute is shorthand for `CHARACTER SET latin1`.
- The `UNICODE` attribute is shorthand for `CHARACTER SET ucs2`.
- The `BINARY` attribute is shorthand for specifying the binary collation of the column character set. In this case, sorting and comparison are based on numeric character values.

Character column sorting and comparison are based on the character set assigned to the column. For the `CHAR`, `VARCHAR`, `TEXT`, `ENUM`, and `SET` data types, you can declare a column with a binary collation or the `BINARY` attribute to cause sorting and comparison to use the underlying character code values rather than a lexical ordering.

[Section 10.1, “Character Set Support”](#), provides additional information about use of character sets in MySQL.

- `[NATIONAL] CHAR[(M)] [CHARACTER SET charset_name] [COLLATE collation_name]`

A fixed-length string that is always right-padded with spaces to the specified length when stored. *M* represents the column length in characters. The range of *M* is 0 to 255. If *M* is omitted, the length is 1.



#### Note

Trailing spaces are removed when `CHAR` values are retrieved unless the `PAD_CHAR_TO_FULL_LENGTH [606]` SQL mode is enabled.

`CHAR` is shorthand for `CHARACTER`. `NATIONAL CHAR` (or its equivalent short form, `NCHAR`) is the standard SQL way to define that a `CHAR` column should use some predefined character set. MySQL 4.1 and up uses `utf8` as this predefined character set. [Section 10.1.3.6, “National Character Set”](#).

The `CHAR BYTE` data type is an alias for the `BINARY` data type. This is a compatibility feature.



MySQL permits you to create a column of type `CHAR(0)`. This is useful primarily when you have to be compliant with old applications that depend on the existence of a column but that do not actually use its value. `CHAR(0)` is also quite nice when you need a column that can take only two values: A column that is defined as `CHAR(0) NULL` occupies only one bit and can take only the values `NULL` and `''` (the empty string).

- `[NATIONAL] VARCHAR(M) [CHARACTER SET charset_name] [COLLATE collation_name]`

A variable-length string. *M* represents the maximum column length in characters. The range of *M* is 0 to 65,535. The effective maximum length of a `VARCHAR` is subject to the maximum row size (65,535 bytes, which is shared among all columns) and the character set used. For example, `utf8` characters can require up to three bytes per character, so a `VARCHAR` column that uses the `utf8` character set can be declared to be a maximum of 21,844 characters. See [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

MySQL stores `VARCHAR` values as a 1-byte or 2-byte length prefix plus data. The length prefix indicates the number of bytes in the value. A `VARCHAR` column uses one length byte if values require no more than 255 bytes, two length bytes if values may require more than 255 bytes.



#### Note

MySQL 5.6 follows the standard SQL specification, and does *not* remove trailing spaces from `VARCHAR` values.

`VARCHAR` is shorthand for `CHARACTER VARYING`. `NATIONAL VARCHAR` is the standard SQL way to define that a `VARCHAR` column should use some predefined character set. MySQL 4.1 and up uses `utf8` as this predefined character set. [Section 10.1.3.6, “National Character Set”](#). `NVARCHAR` is shorthand for `NATIONAL VARCHAR`.

- `BINARY(M)`

The `BINARY` type is similar to the `CHAR` type, but stores binary byte strings rather than nonbinary character strings. *M* represents the column length in bytes.

- `VARBINARY(M)`

The `VARBINARY` type is similar to the `VARCHAR` type, but stores binary byte strings rather than nonbinary character strings. *M* represents the maximum column length in bytes.

- `TINYBLOB`

A `BLOB` column with a maximum length of 255 ( $2^8 - 1$ ) bytes. Each `TINYBLOB` value is stored using a 1-byte length prefix that indicates the number of bytes in the value.

- `TINYTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 255 ( $2^8 - 1$ ) characters. The effective maximum length is less if the value contains multi-byte characters. Each `TINYTEXT` value is stored using a 1-byte length prefix that indicates the number of bytes in the value.

- `BLOB[ (M) ]`

A `BLOB` column with a maximum length of 65,535 ( $2^{16} - 1$ ) bytes. Each `BLOB` value is stored using a 2-byte length prefix that indicates the number of bytes in the value.

An optional length *M* can be given for this type. If this is done, MySQL creates the column as the smallest `BLOB` type large enough to hold values *M* bytes long.

- `TEXT[(M)] [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 65,535 ( $2^{16} - 1$ ) characters. The effective maximum length is less if the value contains multi-byte characters. Each `TEXT` value is stored using a 2-byte length prefix that indicates the number of bytes in the value.

An optional length *M* can be given for this type. If this is done, MySQL creates the column as the smallest `TEXT` type large enough to hold values *M* characters long.

- `MEDIUMBLOB`

A `BLOB` column with a maximum length of 16,777,215 ( $2^{24} - 1$ ) bytes. Each `MEDIUMBLOB` value is stored using a 3-byte length prefix that indicates the number of bytes in the value.

- `MEDIUMTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 16,777,215 ( $2^{24} - 1$ ) characters. The effective maximum length is less if the value contains multi-byte characters. Each `MEDIUMTEXT` value is stored using a 3-byte length prefix that indicates the number of bytes in the value.

- `LOBLOB`

A `BLOB` column with a maximum length of 4,294,967,295 or 4GB ( $2^{32} - 1$ ) bytes. The effective maximum length of `LOBLOB` columns depends on the configured maximum packet size in the client/server protocol and available memory. Each `LOBLOB` value is stored using a 4-byte length prefix that indicates the number of bytes in the value.

- `LONGTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 4,294,967,295 or 4GB ( $2^{32} - 1$ ) characters. The effective maximum length is less if the value contains multi-byte characters. The effective maximum length of `LONGTEXT` columns also depends on the configured maximum packet size in the client/server protocol and available memory. Each `LONGTEXT` value is stored using a 4-byte length prefix that indicates the number of bytes in the value.

- `ENUM('value1', 'value2', ...) [CHARACTER SET charset_name] [COLLATE collation_name]`

An enumeration. A string object that can have only one value, chosen from the list of values `'value1', 'value2', ..., NULL` or the special `' '` error value. `ENUM` values are represented internally as integers.

An `ENUM` column can have a maximum of 65,535 distinct elements. (The practical limit is less than 3000.) A table can have no more than 255 unique element list definitions among its `ENUM` and `SET` columns considered as a group. For more information on these limits, see [Section E.10.5, "Limits Imposed by .frm File Structure"](#).

- `SET('value1', 'value2', ...) [CHARACTER SET charset_name] [COLLATE collation_name]`

A set. A string object that can have zero or more values, each of which must be chosen from the list of values `'value1', 'value2', ...`. `SET` values are represented internally as integers.

A `SET` column can have a maximum of 64 distinct members. A table can have no more than 255 unique element list definitions among its `ENUM` and `SET` columns considered as a group. For more information on this limit, see [Section E.10.5, "Limits Imposed by .frm File Structure"](#).

## 11.2. Numeric Types

MySQL supports all standard SQL numeric data types. These types include the exact numeric data types (`INTEGER`, `SMALLINT`, `DECIMAL`, and `NUMERIC`), as well as the approximate numeric data types (`FLOAT`, `REAL`, and `DOUBLE PRECISION`). The keyword `INT` is a synonym for `INTEGER`, and

the keywords [DEC](#) and [FIXED](#) are synonyms for [DECIMAL](#). MySQL treats [DOUBLE](#) as a synonym for [DOUBLE PRECISION](#) (a nonstandard extension). MySQL also treats [REAL](#) as a synonym for [DOUBLE PRECISION](#) (a nonstandard variation), unless the [REAL\\_AS\\_FLOAT \[607\]](#) SQL mode is enabled.

The [BIT](#) data type stores bit-field values and is supported for [MyISAM](#), [MEMORY](#), [InnoDB](#), and [NDB](#) tables.

For information about how MySQL handles assignment of out-of-range values to columns and overflow during expression evaluation, see [Section 11.2.6, “Out-of-Range and Overflow Handling”](#).

For information about numeric type storage requirements, see [Section 11.6, “Data Type Storage Requirements”](#).

The data type used for the result of a calculation on numeric operands depends on the types of the operands and the operations performed on them. For more information, see [Section 12.6.1, “Arithmetic Operators”](#).

### 11.2.1. Integer Types (Exact Value) - [INTEGER](#), [INT](#), [SMALLINT](#), [TINYINT](#), [MEDIUMINT](#), [BIGINT](#)

MySQL supports the SQL standard integer types [INTEGER](#) (or [INT](#)) and [SMALLINT](#). As an extension to the standard, MySQL also supports the integer types [TINYINT](#), [MEDIUMINT](#), and [BIGINT](#). The following table shows the required storage and range for each integer type.

Type	Storage (Bytes)	Minimum Value (Signed/Unsigned)	Maximum Value (Signed/Unsigned)
<a href="#">TINYINT</a>	1	-128 0	127 255
<a href="#">SMALLINT</a>	2	-32768 0	32767 65535
<a href="#">MEDIUMINT</a>	3	-8388608 0	8388607 16777215
<a href="#">INT</a>	4	-2147483648 0	2147483647 4294967295
<a href="#">BIGINT</a>	8	-9223372036854775808 0	9223372036854775807 18446744073709551615

### 11.2.2. Fixed-Point Types (Exact Value) - [DECIMAL](#), [NUMERIC](#)

The [DECIMAL](#) and [NUMERIC](#) types store exact numeric data values. These types are used when it is important to preserve exact precision, for example with monetary data. In MySQL, [NUMERIC](#) is implemented as [DECIMAL](#), so the following remarks about [DECIMAL](#) apply equally to [NUMERIC](#).

MySQL 5.6 stores [DECIMAL](#) values in binary format. See [Section 12.19, “Precision Math”](#).

In a [DECIMAL](#) column declaration, the precision and scale can be (and usually is) specified; for example:

```
salary DECIMAL(5,2)
```

In this example, 5 is the precision and 2 is the scale. The precision represents the number of significant digits that are stored for values, and the scale represents the number of digits that can be stored following the decimal point.

Standard SQL requires that `DECIMAL(5,2)` be able to store any value with five digits and two decimals, so values that can be stored in the `salary` column range from `-999.99` to `999.99`.

In standard SQL, the syntax `DECIMAL(M)` is equivalent to `DECIMAL(M,0)`. Similarly, the syntax `DECIMAL` is equivalent to `DECIMAL(M,0)`, where the implementation is permitted to decide the value of `M`. MySQL supports both of these variant forms of `DECIMAL` syntax. The default value of `M` is 10.

If the scale is 0, `DECIMAL` values contain no decimal point or fractional part.

The maximum number of digits for `DECIMAL` is 65, but the actual range for a given `DECIMAL` column can be constrained by the precision or scale for a given column. When such a column is assigned a value with more digits following the decimal point than are permitted by the specified scale, the value is converted to that scale. (The precise behavior is operating system-specific, but generally the effect is truncation to the permissible number of digits.)

### 11.2.3. Floating-Point Types (Approximate Value) - `FLOAT`, `DOUBLE`

The `FLOAT` and `DOUBLE` types represent approximate numeric data values. MySQL uses four bytes for single-precision values and eight bytes for double-precision values.

For `FLOAT`, the SQL standard permits an optional specification of the precision (but not the range of the exponent) in bits following the keyword `FLOAT` in parentheses. MySQL also supports this optional precision specification, but the precision value is used only to determine storage size. A precision from 0 to 23 results in a 4-byte single-precision `FLOAT` column. A precision from 24 to 53 results in an 8-byte double-precision `DOUBLE` column.

MySQL permits a nonstandard syntax: `FLOAT(M,D)` or `REAL(M,D)` or `DOUBLE PRECISION(M,D)`. Here, “`(M,D)`” means that values can be stored with up to `M` digits in total, of which `D` digits may be after the decimal point. For example, a column defined as `FLOAT(7,4)` will look like `-999.9999` when displayed. MySQL performs rounding when storing values, so if you insert `999.00009` into a `FLOAT(7,4)` column, the approximate result is `999.0001`.

Because floating-point values are approximate and not stored as exact values, attempts to treat them as exact in comparisons may lead to problems. They are also subject to platform or implementation dependencies. For more information, see [Section C.5.5.8, “Problems with Floating-Point Values”](#)

For maximum portability, code requiring storage of approximate numeric data values should use `FLOAT` or `DOUBLE PRECISION` with no specification of precision or number of digits.

### 11.2.4. Bit-Value Type - `BIT`

The `BIT` data type is used to store bit-field values. A type of `BIT(M)` enables storage of `M`-bit values. `M` can range from 1 to 64.

To specify bit values, `b'value'` notation can be used. `value` is a binary value written using zeros and ones. For example, `b'111'` and `b'1000000'` represent 7 and 128, respectively. See [Section 9.1.6, “Bit-Field Literals”](#).

If you assign a value to a `BIT(M)` column that is less than `M` bits long, the value is padded on the left with zeros. For example, assigning a value of `b'101'` to a `BIT(6)` column is, in effect, the same as assigning `b'000101'`.

### 11.2.5. Numeric Type Attributes

MySQL supports an extension for optionally specifying the display width of integer data types in parentheses following the base keyword for the type. For example, `INT(4)` specifies an `INT` with a display width of four digits. This optional display width may be used by applications to display integer values having a width less than the width specified for the column by left-padding them with spaces.

(That is, this width is present in the metadata returned with result sets. Whether it is used or not is up to the application.)

The display width does *not* constrain the range of values that can be stored in the column. Nor does it prevent values wider than the column display width from being displayed correctly. For example, a column specified as `SMALLINT(3)` has the usual `SMALLINT` range of `-32768` to `32767`, and values outside the range permitted by three digits are displayed in full using more than three digits.

When used in conjunction with the optional (nonstandard) attribute `ZEROFILL`, the default padding of spaces is replaced with zeros. For example, for a column declared as `INT(4) ZEROFILL`, a value of `5` is retrieved as `0005`.



#### Note

The `ZEROFILL` attribute is ignored when a column is involved in expressions or `UNION` queries.

If you store values larger than the display width in an integer column that has the `ZEROFILL` attribute, you may experience problems when MySQL generates temporary tables for some complicated joins. In these cases, MySQL assumes that the data values fit within the column display width.

All integer types can have an optional (nonstandard) attribute `UNSIGNED`. Unsigned type can be used to permit only nonnegative numbers in a column or when you need a larger upper numeric range for the column. For example, if an `INT` column is `UNSIGNED`, the size of the column's range is the same but its endpoints shift from `-2147483648` and `2147483647` up to `0` and `4294967295`.

Floating-point and fixed-point types also can be `UNSIGNED`. As with integer types, this attribute prevents negative values from being stored in the column. Unlike the integer types, the upper range of column values remains the same.

If you specify `ZEROFILL` for a numeric column, MySQL automatically adds the `UNSIGNED` attribute to the column.

Integer or floating-point data types can have the additional attribute `AUTO_INCREMENT`. When you insert a value of `NULL` (recommended) or `0` into an indexed `AUTO_INCREMENT` column, the column is set to the next sequence value. Typically this is `value+1`, where `value` is the largest value for the column currently in the table. `AUTO_INCREMENT` sequences begin with `1`.

## 11.2.6. Out-of-Range and Overflow Handling

When MySQL stores a value in a numeric column that is outside the permissible range of the column data type, the result depends on the SQL mode in effect at the time:

- If strict SQL mode is enabled, MySQL rejects the out-of-range value with an error, and the insert fails, in accordance with the SQL standard.
- If no restrictive modes are enabled, MySQL clips the value to the appropriate endpoint of the range and stores the resulting value instead.

When an out-of-range value is assigned to an integer column, MySQL stores the value representing the corresponding endpoint of the column data type range. If you store `256` into a `TINYINT` or `TINYINT UNSIGNED` column, MySQL stores `127` or `255`, respectively.

When a floating-point or fixed-point column is assigned a value that exceeds the range implied by the specified (or default) precision and scale, MySQL stores the value representing the corresponding endpoint of that range.

Column-assignment conversions that occur due to clipping when MySQL is not operating in strict mode are reported as warnings for `ALTER TABLE`, `LOAD DATA INFILE`, `UPDATE`, and multiple-row `INSERT`

statements. In strict mode, these statements fail, and some or all the values will not be inserted or changed, depending on whether the table is a transactional table and other factors. For details, see [Section 5.1.7, “Server SQL Modes”](#).

In MySQL 5.6, overflow during numeric expression evaluation results in an error. For example, the largest signed `BIGINT` value is 9223372036854775807, so the following expression produces an error:

```
mysql> SELECT 9223372036854775807 + 1;
ERROR 1690 (22003): BIGINT value is out of range in '(9223372036854775807 + 1)'
```

To enable the operation to succeed in this case, convert the value to unsigned;

```
mysql> SELECT CAST(9223372036854775807 AS UNSIGNED) + 1;
+-----+
| CAST(9223372036854775807 AS UNSIGNED) + 1 |
+-----+
|                9223372036854775808 |
+-----+
```

Whether overflow occurs depends on the range of the operands, so another way to handle the preceding expression is to use exact-value arithmetic because `DECIMAL` values have a larger range than integers:

```
mysql> SELECT 9223372036854775807.0 + 1;
+-----+
| 9223372036854775807.0 + 1 |
+-----+
|                9223372036854775808.0 |
+-----+
```

Subtraction between integer values, where one is of type `UNSIGNED`, produces an unsigned result by default. Prior to MySQL 5.5.5, if the result would otherwise have been negative, it becomes the maximum integer value:

```
mysql> SET sql_mode = '';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
| CAST(0 AS UNSIGNED) - 1 |
+-----+
| 18446744073709551615 |
+-----+
```

As of MySQL 5.5.5, if the result would otherwise have been negative, an error results:

```
mysql> SET sql_mode = '';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT CAST(0 AS UNSIGNED) - 1;
ERROR 1690 (22003): BIGINT UNSIGNED value is out of range in '(cast(0 as unsigned) - 1)'
```

If the `NO_UNSIGNED_SUBTRACTION` [\[605\]](#) SQL mode is enabled, the result is negative:

```
mysql> SET sql_mode = 'NO_UNSIGNED_SUBTRACTION';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
| CAST(0 AS UNSIGNED) - 1 |
+-----+
|                -1 |
+-----+
```

If the result of such an operation is used to update an `UNSIGNED` integer column, the result is clipped to the maximum value for the column type, or clipped to 0 if `NO_UNSIGNED_SUBTRACTION` [\[605\]](#) is enabled. If strict SQL mode is enabled, an error occurs and the column remains unchanged.

## 11.3. Date and Time Types



The date and time types for representing temporal values are [DATE](#), [TIME](#), [DATETIME](#), [TIMESTAMP](#), and [YEAR](#). Each temporal type has a range of valid values, as well as a “zero” value that may be used when you specify an invalid value that MySQL cannot represent. The [TIMESTAMP](#) type has special automatic updating behavior, described later. For temporal type storage requirements, see [Section 11.6, “Data Type Storage Requirements”](#).

Keep in mind these general considerations when working with date and time types:

- MySQL retrieves values for a given date or time type in a standard output format, but it attempts to interpret a variety of formats for input values that you supply (for example, when you specify a value to be assigned to or compared to a date or time type). For a description of the permitted formats for date and time types, see [Section 9.1.3, “Date and Time Literals”](#). It is expected that you supply valid values. Unpredictable results may occur if you use values in other formats.
- Although MySQL tries to interpret values in several formats, date parts must always be given in year-month-day order (for example, `'98-09-04'`), rather than in the month-day-year or day-month-year orders commonly used elsewhere (for example, `'09-04-98'`, `'04-09-98'`).
- Dates containing two-digit year values are ambiguous because the century is unknown. MySQL interprets two-digit year values using these rules:
  - Year values in the range `70-99` are converted to `1970-1999`.
  - Year values in the range `00-69` are converted to `2000-2069`.

See also [Section 11.3.8, “Two-Digit Years in Dates”](#).

- Conversion of values from one temporal type to another occurs according to the rules in [Section 11.3.7, “Conversion Between Date and Time Types”](#).
- MySQL automatically converts a date or time value to a number if the value is used in a numeric context and vice versa.
- By default, when MySQL encounters a value for a date or time type that is out of range or otherwise invalid for the type, it converts the value to the “zero” value for that type. The exception is that out-of-range [TIME](#) values are clipped to the appropriate endpoint of the [TIME](#) range.
- By setting the SQL mode to the appropriate value, you can specify more exactly what kind of dates you want MySQL to support. (See [Section 5.1.7, “Server SQL Modes”](#).) You can get MySQL to accept certain dates, such as `'2009-11-31'`, by enabling the [ALLOW\\_INVALID\\_DATES \[603\]](#) SQL mode. This is useful when you want to store a “possibly wrong” value which the user has specified (for example, in a web form) in the database for future processing. Under this mode, MySQL verifies only that the month is in the range from 1 to 12 and that the day is in the range from 1 to 31.
- MySQL permits you to store dates where the day or month and day are zero in a [DATE](#) or [DATETIME](#) column. This is extremely useful for applications that need to store birthdates for which you may not know the exact date. In this case, you simply store the date as `'2009-00-00'` or `'2009-01-00'`. If you store dates such as these, you should not expect to get correct results for functions such as [DATE\\_SUB\(\) \[1217\]](#) or [DATE\\_ADD\(\) \[1213\]](#) that require complete dates. To disallow zero month or day parts in dates, enable the [NO\\_ZERO\\_IN\\_DATE \[606\]](#) SQL mode.
- MySQL permits you to store a “zero” value of `'0000-00-00'` as a “dummy date.” This is in some cases more convenient than using [NULL](#) values, and uses less data and index space. To disallow `'0000-00-00'`, enable the [NO\\_ZERO\\_DATE \[606\]](#) SQL mode.
- “Zero” date or time values used through Connector/ODBC are converted automatically to [NULL](#) because ODBC cannot handle such values.

The following table shows the format of the “zero” value for each type. The “zero” values are special, but you can store or refer to them explicitly using the values shown in the table. You can also do this



using the values `'0'` or `0`, which are easier to write. For temporal types that include a date part (`DATE`, `DATETIME`, and `TIMESTAMP`), use of these values produces warnings if the `NO_ZERO_DATE` [606] SQL mode is enabled.

Data Type	“Zero” Value
<code>DATE</code>	<code>'0000-00-00'</code>
<code>TIME</code>	<code>'00:00:00'</code>
<code>DATETIME</code>	<code>'0000-00-00 00:00:00'</code>
<code>TIMESTAMP</code>	<code>'0000-00-00 00:00:00'</code>
<code>YEAR</code>	<code>0000</code>

### 11.3.1. The `DATE`, `DATETIME`, and `TIMESTAMP` Types

The `DATE`, `DATETIME`, and `TIMESTAMP` types are related. This section describes their characteristics, how they are similar, and how they differ. MySQL recognizes `DATE`, `DATETIME`, and `TIMESTAMP` values in several formats, described in [Section 9.1.3, “Date and Time Literals”](#). For the `DATE` and `DATETIME` range descriptions, “supported” means that although earlier values might work, there is no guarantee.

The `DATE` type is used for values with a date part but no time part. MySQL retrieves and displays `DATE` values in `'YYYY-MM-DD'` format. The supported range is `'1000-01-01'` to `'9999-12-31'`.

The `DATETIME` type is used for values that contain both date and time parts. MySQL retrieves and displays `DATETIME` values in `'YYYY-MM-DD HH:MM:SS'` format. The supported range is `'1000-01-01 00:00:00'` to `'9999-12-31 23:59:59'`.

The `TIMESTAMP` data type is used for values that contain both date and time parts. `TIMESTAMP` has a range of `'1970-01-01 00:00:01'` UTC to `'2038-01-19 03:14:07'` UTC.

A `DATETIME` or `TIMESTAMP` value can include a trailing fractional seconds part in up to microseconds (6 digits) precision. In particular, as of MySQL 5.6.4, any fractional part in a value inserted into a `DATETIME` or `TIMESTAMP` column is stored rather than discarded. With the fractional part included, the format for these values is `'YYYY-MM-DD HH:MM:SS[.fraction]'`, the range for `DATETIME` values is `'1000-01-01 00:00:00.000000'` to `'9999-12-31 23:59:59.999999'`, and the range for `TIMESTAMP` values is `'1970-01-01 00:00:01.000000'` to `'2038-01-19 03:14:07.999999'`. For information about fractional seconds support in MySQL, see [Section 11.3.6, “Fractional Seconds in Time Values”](#).

The `TIMESTAMP` and (as of MySQL 5.6.5) `DATETIME` data types offer automatic initialization and updating to the current date and time. For more information, see [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#).

MySQL converts `TIMESTAMP` values from the current time zone to UTC for storage, and back from UTC to the current time zone for retrieval. (This does not occur for other types such as `DATETIME`.) By default, the current time zone for each connection is the server's time. The time zone can be set on a per-connection basis. As long as the time zone setting remains constant, you get back the same value you store. If you store a `TIMESTAMP` value, and then change the time zone and retrieve the value, the retrieved value is different from the value you stored. This occurs because the same time zone was not used for conversion in both directions. The current time zone is available as the value of the `time_zone` [557] system variable. For more information, see [Section 10.6, “MySQL Server Time Zone Support”](#).

Invalid `DATE`, `DATETIME`, or `TIMESTAMP` values are converted to the “zero” value of the appropriate type (`'0000-00-00'` or `'0000-00-00 00:00:00'`).

Be aware of certain properties of date value interpretation in MySQL:

- MySQL permits a “relaxed” format for values specified as strings, in which any punctuation character may be used as the delimiter between date parts or time parts. In some cases, this syntax can be deceiving. For example, a value such as `'10:11:12'` might look like a time value because of the “:” delimiter, but is interpreted as the year `'2010-11-12'` if used in a date context. The value `'10:45:15'` is converted to `'0000-00-00'` because `'45'` is not a valid month.
- The server requires that month and day values be valid, and not merely in the range 1 to 12 and 1 to 31, respectively. With strict mode disabled, invalid dates such as `'2004-04-31'` are converted to `'0000-00-00'` and a warning is generated. With strict mode enabled, invalid dates generate an error. To permit such dates, enable `ALLOW_INVALID_DATES` [603]. See Section 5.1.7, “Server SQL Modes”, for more information.
- MySQL does not accept `TIMESTAMP` values that include a zero in the day or month column or values that are not a valid date. The sole exception to this rule is the special “zero” value `'0000-00-00 00:00:00'`.
- Dates containing two-digit year values are ambiguous because the century is unknown. MySQL interprets two-digit year values using these rules:
  - Year values in the range `00-69` are converted to `2000-2069`.
  - Year values in the range `70-99` are converted to `1970-1999`.

See also Section 11.3.8, “Two-Digit Years in Dates”.



#### Note

The MySQL server can be run with the `MAXDB` [608] SQL mode enabled. In this case, `TIMESTAMP` is identical with `DATETIME`. If this mode is enabled at the time that a table is created, `TIMESTAMP` columns are created as `DATETIME` columns. As a result, such columns use `DATETIME` display format, have the same range of values, and there is no automatic initialization or updating to the current date and time. See Section 5.1.7, “Server SQL Modes”.

### 11.3.2. The `TIME` Type

MySQL retrieves and displays `TIME` values in `'HH:MM:SS'` format (or `'HHH:MM:SS'` format for large hours values). `TIME` values may range from `'-838:59:59'` to `'838:59:59'`. The hours part may be so large because the `TIME` type can be used not only to represent a time of day (which must be less than 24 hours), but also elapsed time or a time interval between two events (which may be much greater than 24 hours, or even negative).

MySQL recognizes `TIME` values in several formats, some of which can include a trailing fractional seconds part in up to microseconds (6 digits) precision. See Section 9.1.3, “Date and Time Literals”. For information about fractional seconds support in MySQL, see Section 11.3.6, “Fractional Seconds in Time Values”. In particular, as of MySQL 5.6.4, any fractional part in a value inserted into a `TIME` column is stored rather than discarded. With the fractional part included, the range for `TIME` values is `'-838:59:59.000000'` to `'838:59:59.000000'`.

Be careful about assigning abbreviated values to a `TIME` column. MySQL interprets abbreviated `TIME` values with colons as time of the day. That is, `'11:12'` means `'11:12:00'`, not `'00:11:12'`. MySQL interprets abbreviated values without colons using the assumption that the two rightmost digits represent seconds (that is, as elapsed time rather than as time of day). For example, you might think of `'1112'` and `1112` as meaning `'11:12:00'` (12 minutes after 11 o'clock), but MySQL interprets them as `'00:11:12'` (11 minutes, 12 seconds). Similarly, `'12'` and `12` are interpreted as `'00:00:12'`.

By default, values that lie outside the `TIME` range but are otherwise valid are clipped to the closest endpoint of the range. For example, `'-850:00:00'` and `'850:00:00'` are converted to `'-838:59:59'` and `'838:59:59'`. Invalid `TIME` values are converted to `'00:00:00'`. Note that

because '00:00:00' is itself a valid `TIME` value, there is no way to tell, from a value of '00:00:00' stored in a table, whether the original value was specified as '00:00:00' or whether it was invalid.

For more restrictive treatment of invalid `TIME` values, enable strict SQL mode to cause errors to occur. See [Section 5.1.7, “Server SQL Modes”](#).

### 11.3.3. The YEAR Type

The `YEAR` type is a 1-byte type used to represent year values. It can be declared as `YEAR(4)` or `YEAR(2)` to specify a display width of four or two characters. The default is four characters if no width is given.



#### Note

The `YEAR(2)` data type has certain issues that you should consider before choosing to use it. Also, as of MySQL 5.6.6, `YEAR(2)` is deprecated. `YEAR(2)` columns in existing tables are treated as before, but `YEAR(2)` in new or altered tables are converted to `YEAR(4)`. For more information, see [Section 11.3.4, “YEAR\(2\) Limitations and Migrating to YEAR\(4\)”](#).

`YEAR(4)` and `YEAR(2)` differ in display format, but have the same range of values. For 4-digit format, MySQL displays `YEAR` values in `YYYY` format, with a range of 1901 to 2155, or 0000. For 2-digit format, MySQL displays only the last two (least significant) digits; for example, 70 (1970 or 2070) or 69 (2069).

You can specify input `YEAR` values in a variety of formats:

- As a 4-digit string in the range '1901' to '2155'.
- As a 4-digit number in the range 1901 to 2155.
- As a 1- or 2-digit string in the range '0' to '99'. Values in the ranges '0' to '69' and '70' to '99' are converted to `YEAR` values in the ranges 2000 to 2069 and 1970 to 1999.
- As a 1- or 2-digit number in the range 1 to 99. Values in the ranges 1 to 69 and 70 to 99 are converted to `YEAR` values in the ranges 2001 to 2069 and 1970 to 1999.

Inserting a numeric 0 has a different effect for `YEAR(2)` and `YEAR(4)`. For `YEAR(2)`, the result has a display value of 00 and an internal value of 2000. For `YEAR(4)`, the result has a display value of 0000 and an internal value of 0000. To specify zero for `YEAR(4)` and have it be interpreted as 2000, specify it as a string '0' or '00'.

- As the result of a function that returns a value that is acceptable in a `YEAR` context, such as `NOW()` [1220].

Invalid `YEAR` values are converted to 0000.

See also [Section 11.3.8, “Two-Digit Years in Dates”](#).

### 11.3.4. YEAR(2) Limitations and Migrating to YEAR(4)

Although the internal range of values for `YEAR(4)` and `YEAR(2)` is the same (1901 to 2155, and 0000), the display width for `YEAR(2)` makes that type inherently ambiguous because displayed values indicate only the last two digits of the internal values. The result can be a loss of information under certain circumstances. For this reason, consider avoiding `YEAR(2)` throughout your applications and using `YEAR(4)` wherever you need a `YEAR` data type. This section describes problems that can occur when using `YEAR(2)` and provides information about migrating existing `YEAR(2)` columns to `YEAR(4)`. Note that migration will become necessary at some point because support for `YEAR` data types with display values other than 4, most notably `YEAR(2)`, is reduced as of MySQL 5.6.6 and will be removed entirely in a future release.

## YEAR(2) Limitations

Issues with the YEAR(2) data type include ambiguity of displayed values, and possible loss of information when values are dumped and reloaded or converted to strings.

- Displayed YEAR(2) values can be ambiguous. It is possible for up to three YEAR(2) values that have different internal values to have the same displayed value, as the following example demonstrates:

```
mysql> CREATE TABLE t (y2 YEAR(2), y4 YEAR(4));
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO t (y2) VALUES(1912),(2012),(2112);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> UPDATE t SET y4 = y2;
Query OK, 3 rows affected (0.00 sec)
Rows matched: 3 Changed: 3 Warnings: 0

mysql> SELECT * FROM t;
+-----+-----+
| y2   | y4   |
+-----+-----+
| 12   | 1912 |
| 12   | 2012 |
| 12   | 2112 |
+-----+-----+
3 rows in set (0.00 sec)
```

- If you use `mysqldump` to dump the table created in the preceding item, the dump file represents all y2 values using the same 2-digit representation (12). If you reload the table from the dump file, all resulting rows have internal value 2012 and display value 12, thus losing the distinctions among them.
- Conversion of a YEAR(2) or YEAR(4) data value to string form uses the display width of the YEAR type. Suppose that YEAR(2) and YEAR(4) columns both contain the value 1970. Assigning each column to a string results in a value of '70' or '1970', respectively. That is, loss of information occurs for conversion from YEAR(2) to string.
- Values outside the range from 1970 to 2069 are stored incorrectly when inserted into a YEAR(2) column in a CSV table. For example, inserting 2111 results in a display value of 11 but an internal value of 2011.

To avoid these problems, use YEAR(4) rather than YEAR(2). Suggestions regarding migration strategies appear later in this section.

## Reduced YEAR(2) Support in MySQL 5.6

As of MySQL 5.6.6, support for YEAR(2) is diminished:

- YEAR(2) in column definitions for new tables is converted (with a warning) to YEAR(4):

```
mysql> CREATE TABLE t1 (y YEAR(2));
Query OK, 0 rows affected, 1 warning (0.03 sec)

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1818
Message: YEAR(2) column type is deprecated. Creating YEAR(4) column instead.
1 row in set (0.00 sec)

mysql> SHOW CREATE TABLE t1\G
***** 1. row *****
Table: t1
Create Table: CREATE TABLE `t1` (
```

```
`y` year(4) DEFAULT NULL
) ENGINE=InnoDB DEFAULT CHARSET=latin1
1 row in set (0.00 sec)
```

- YEAR(2) in existing tables remains as YEAR(2) and is processed in queries as in older versions of MySQL. However, several programs or statements convert YEAR(2) to YEAR(4) automatically:
  - ALTER TABLE statements that result in a table rebuild.
  - REPAIR TABLE (which CHECK TABLE recommends you use if it finds that a table contains YEAR(2) columns).
  - mysql\_upgrade (which uses REPAIR TABLE).
  - Dumping with mysqldump and reloading the dump file. Unlike the conversions performed by the preceding three items, a dump and reload has the potential to change values.

A MySQL upgrade usually involves at least one of the last two items. However, with respect to YEAR(2), you should avoid dumping and reloading; as noted, that can change values.

## Migrating from YEAR(2) to YEAR(4)

Should you decide to convert YEAR(2) columns to YEAR(4), you can do so manually at any time without upgrading. Alternatively, you can upgrade to a version of MySQL with reduced support for YEAR(2) (MySQL 5.6.6 or later), then have MySQL convert YEAR(2) columns automatically. In the latter case, avoid upgrading by dumping and reloading your data because that can change data values. In addition, if you use replication, there are upgrade considerations you must take into account.

To convert YEAR(2) columns to YEAR(4) manually, use ALTER TABLE. Suppose that a table t1 has this definition:

```
CREATE TABLE t1 (ycol YEAR(2) NOT NULL DEFAULT '70');
```

Modify the column using ALTER TABLE as follows. Remember to include any column attributes such as NOT NULL or DEFAULT:

```
ALTER TABLE t1 MODIFY ycol YEAR(4) NOT NULL DEFAULT '1970';
```

The ALTER TABLE statement converts the table without changing YEAR(2) values. If the server is a replication master, the ALTER TABLE statement replicates to slaves and makes the corresponding table change on each one.

Another migration method is to perform a binary upgrade: Install MySQL 5.6.6 or later without dumping and reloading your data. Then run mysql\_upgrade, which uses REPAIR TABLE to convert YEAR(2) columns to YEAR(4) without changing data values. If the server is a replication master, the REPAIR TABLE statements replicate to slaves and make the corresponding table changes on each one, unless you invoke mysql\_upgrade with the --skip-write-binlog [244] option.

Upgrades to replication servers usually involve upgrading slaves to a newer version of MySQL, then upgrading the master. For example, if a master and slave both run MySQL 5.5, a typical upgrade sequence involves upgrading the slave to 5.6, then upgrading the master to 5.6. With regard to the different treatment of YEAR(2) as of MySQL 5.6.6, that upgrade sequence results in a problem: Suppose that the slave has been upgraded but not yet the master. Then creating a table containing a YEAR(2) column on the master results in a table containing a YEAR(4) column on the slave. Consequently, these operations will have a different result on the master and slave, if you use statement-based replication:

- Inserting numeric 0. The resulting value has an internal value of 2000 on the master but 0000 on the slave.
- Converting YEAR(2) to string. This operation uses the display value of YEAR(2) on the master but YEAR(4) on the slave.

To avoid such problems, use one of these strategies:

- Use row-based replication instead of statement-based replication.
- Modify all `YEAR(2)` columns on the master to `YEAR(4)` before upgrading. (Use `ALTER TABLE`, as described previously.) Then you can upgrade normally (slave first, then master) without introducing any `YEAR(2)` to `YEAR(4)` differences between the master and slave).

One migration method should be avoided: Do not dump your data with `mysqldump` and reload the dump file after upgrading. This has the potential to change `YEAR(2)` values, as described previously.

A migration from `YEAR(2)` to `YEAR(4)` should also involve examining application code for the possibility of changed behavior under conditions such as these:

- Code that expects selecting a `YEAR` column to produce exactly two digits.
- Code that does not account for different handling for inserts of numeric 0: Inserting 0 into `YEAR(2)` or `YEAR(4)` results in an internal value of 2000 or 0000, respectively.

### 11.3.5. Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`



#### Note

The following discussion applies under the assumption that the `explicit_defaults_for_timestamp` [481] system variable is disabled. If this variable is enabled, the `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` clauses that specify automatic initialization and updating are available, but are not assigned to any `TIMESTAMP` column unless explicitly included in the column definition.

As of MySQL 5.6.5, `TIMESTAMP` and `DATETIME` columns can be automatically initialized and updated to the current date and time (that is, the current timestamp). Before 5.6.5, this is true only for `TIMESTAMP`, and for at most one `TIMESTAMP` column per table. The following notes first describe automatic initialization and updating for MySQL 5.6.5 and up, then the differences for versions preceding 5.6.5.

For any `TIMESTAMP` or `DATETIME` column in a table, you can assign the current timestamp as the default value, the auto-update value, or both:

- An auto-initialized column is set to the current timestamp for inserted rows that specify no value for the column.
- An auto-updated column is automatically updated to the current timestamp when the value of any other column in the row is changed from its current value. An auto-updated column remains unchanged if all other columns are set to their current values. To prevent an auto-updated column from updating when other columns change, explicitly set it to its current value. To update an auto-updated column even when other columns do not change, explicitly set it to the value it should have (for example, set it to `CURRENT_TIMESTAMP` [1213]).

In addition, you can initialize or update any `TIMESTAMP` column to the current date and time by assigning it a `NULL` value, unless it has been defined with the `NULL` attribute to permit `NULL` values.

To specify automatic properties, use the `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` clauses in column definitions. The order of the clauses does not matter. If both are present in a column definition, either can occur first. Any of the synonyms for `CURRENT_TIMESTAMP` [1213] have the same meaning as `CURRENT_TIMESTAMP` [1213]. These are `CURRENT_TIMESTAMP()` [1213], `NOW()` [1220], `LOCALTIME` [1219], `LOCALTIME()` [1219], `LOCALTIMESTAMP` [1219], and `LOCALTIMESTAMP()` [1219].

Use of `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` is specific to `TIMESTAMP` and `DATETIME`. The `DEFAULT` clause also can be used to specify a constant (nonautomatic) default value; for example, `DEFAULT 0` or `DEFAULT '2000-01-01 00:00:00'`.



**Note**

The following examples that use `DEFAULT 0` do not work if the `NO_ZERO_DATE` [606] SQL mode is enabled because that mode causes “zero” date values (specified, for example, as `0 '0000-00-00 00:00:00'`) to be rejected. Be aware that the `TRADITIONAL` [609] SQL mode includes `NO_ZERO_DATE` [606].

`TIMESTAMP` or `DATETIME` column definitions can specify the current timestamp for both the default and auto-update values, for one but not the other, or for neither. Different columns can have different combinations of automatic properties. The following rules describe the possibilities:

- With both `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP`, the column has the current timestamp for its default value and is automatically updated to the current timestamp.

```
CREATE TABLE t1 (
  ts TIMESTAMP DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP,
  dt DATETIME DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
);
```

- With a `DEFAULT` clause but no `ON UPDATE CURRENT_TIMESTAMP` clause, the column has the given default value and is not automatically updated to the current timestamp.

The default depends on whether the `DEFAULT` clause specifies `CURRENT_TIMESTAMP` or a constant value. With `CURRENT_TIMESTAMP`, the default is the current timestamp.

```
CREATE TABLE t1 (
  ts TIMESTAMP DEFAULT CURRENT_TIMESTAMP,
  dt DATETIME DEFAULT CURRENT_TIMESTAMP
);
```

With a constant, the default is the given value. In this case, the column has no automatic properties at all.

```
CREATE TABLE t1 (
  ts TIMESTAMP DEFAULT 0,
  dt DATETIME DEFAULT 0
);
```

- With an `ON UPDATE CURRENT_TIMESTAMP` clause and a constant `DEFAULT` clause, the column is automatically updated to the current timestamp and has the given constant default value.

```
CREATE TABLE t1 (
  ts TIMESTAMP DEFAULT 0 ON UPDATE CURRENT_TIMESTAMP,
  dt DATETIME DEFAULT 0 ON UPDATE CURRENT_TIMESTAMP
);
```

- With an `ON UPDATE CURRENT_TIMESTAMP` clause but no `DEFAULT` clause, the column is automatically updated to the current timestamp but does not have the current timestamp for its default value.

The default in this case is type dependent. `TIMESTAMP` has a default of 0 unless defined with the `NULL` attribute, in which case the default is `NULL`.

```
CREATE TABLE t1 (
  ts1 TIMESTAMP ON UPDATE CURRENT_TIMESTAMP,      -- default 0
  ts2 TIMESTAMP NULL ON UPDATE CURRENT_TIMESTAMP -- default NULL
);
```

`DATETIME` has a default of `NULL` unless defined with the `NOT NULL` attribute, in which case the default is 0.

```
CREATE TABLE t1 (
  dt1 DATETIME ON UPDATE CURRENT_TIMESTAMP,      -- default NULL
  dt2 DATETIME NOT NULL ON UPDATE CURRENT_TIMESTAMP -- default 0
);
```



`TIMESTAMP` and `DATETIME` columns have no automatic properties unless they are specified explicitly, with this exception: By default, the *first* `TIMESTAMP` column has both `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` if neither is specified explicitly. To suppress automatic properties for the first `TIMESTAMP` column, do either of the following:

- Define the column with a `DEFAULT` clause that specifies a constant default value.
- Specify the `NULL` attribute. This also causes the column to permit `NULL` values, which means that you cannot assign the current timestamp by setting the column to `NULL`. Assigning `NULL` sets the column to `NULL`.

Consider these table definitions:

```
CREATE TABLE t1 (
  ts1 TIMESTAMP DEFAULT 0,
  ts2 TIMESTAMP DEFAULT CURRENT_TIMESTAMP
  ON UPDATE CURRENT_TIMESTAMP);
CREATE TABLE t2 (
  ts1 TIMESTAMP NULL,
  ts2 TIMESTAMP DEFAULT CURRENT_TIMESTAMP
  ON UPDATE CURRENT_TIMESTAMP);
CREATE TABLE t3 (
  ts1 TIMESTAMP NULL DEFAULT 0,
  ts2 TIMESTAMP DEFAULT CURRENT_TIMESTAMP
  ON UPDATE CURRENT_TIMESTAMP);
```

The tables have these properties:

- In each table definition, the first `TIMESTAMP` column has no automatic initialization or updating.
- The tables differ in how the `ts1` column handles `NULL` values. For `t1`, `ts1` is `NOT NULL` and assigning it a value of `NULL` sets it to the current timestamp. For `t2` and `t3`, `ts1` permits `NULL` and assigning it a value of `NULL` sets it to `NULL`.
- `t2` and `t3` differ in the default value for `ts1`. For `t2`, `ts1` is defined to permit `NULL`, so the default is also `NULL` in the absence of an explicit `DEFAULT` clause. For `t3`, `ts1` permits `NULL` but has an explicit default of 0.

If a `TIMESTAMP` or `DATETIME` column definition includes an explicit fractional seconds precision value anywhere, the same value must be used throughout the column definition. This is permitted:

```
CREATE TABLE t1 (
  ts TIMESTAMP(6) DEFAULT CURRENT_TIMESTAMP(6) ON UPDATE CURRENT_TIMESTAMP(6)
);
```

This is not permitted:

```
CREATE TABLE t1 (
  ts TIMESTAMP(6) DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP(3)
);
```

## Automatic Timestamp Properties Before MySQL 5.6.5

Before MySQL 5.6.5, support for automatic initialization and updating is more limited:

- `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` cannot be used with `DATETIME` columns.
- `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` can be used with at most one `TIMESTAMP` column per table. It is not possible to have the current timestamp be the default value for one column and the auto-update value for another column.

You can choose whether to use these properties and which `TIMESTAMP` column should have them. It need not be the first one in a table that is automatically initialized or updated to the current timestamp. To specify automatic initialization or updating for a different `TIMESTAMP` column, you must suppress

the automatic properties for the first one, as previously described. Then, for the other `TIMESTAMP` column, the rules for the `DEFAULT` and `ON UPDATE` clauses are the same as for the first `TIMESTAMP` column, except that if you omit both clauses, no automatic initialization or updating occurs.

### TIMESTAMP Initialization and the NULL Attribute

By default, `TIMESTAMP` columns are `NOT NULL`, cannot contain `NULL` values, and assigning `NULL` assigns the current timestamp. To permit a `TIMESTAMP` column to contain `NULL`, explicitly declare it with the `NULL` attribute. In this case, the default value also becomes `NULL` unless overridden with a `DEFAULT` clause that specifies a different default value. `DEFAULT NULL` can be used to explicitly specify `NULL` as the default value. (For a `TIMESTAMP` column not declared with the `NULL` attribute, `DEFAULT NULL` is invalid.) If a `TIMESTAMP` column permits `NULL` values, assigning `NULL` sets it to `NULL`, not to the current timestamp.

The following table contains several `TIMESTAMP` columns that permit `NULL` values:

```
CREATE TABLE t
(
  ts1 TIMESTAMP NULL DEFAULT NULL,
  ts2 TIMESTAMP NULL DEFAULT 0,
  ts3 TIMESTAMP NULL DEFAULT CURRENT_TIMESTAMP
);
```

A `TIMESTAMP` column that permits `NULL` values does *not* take on the current timestamp at insert time except under one of the following conditions:

- Its default value is defined as `CURRENT_TIMESTAMP [1213]` and no value is specified for the column
- `CURRENT_TIMESTAMP [1213]` or any of its synonyms such as `NOW() [1220]` is explicitly inserted into the column

In other words, a `TIMESTAMP` column defined to permit `NULL` values auto-initializes only if its definition includes `DEFAULT CURRENT_TIMESTAMP`:

```
CREATE TABLE t (ts TIMESTAMP NULL DEFAULT CURRENT_TIMESTAMP);
```

If the `TIMESTAMP` column permits `NULL` values but its definition does not include `DEFAULT CURRENT_TIMESTAMP`, you must explicitly insert a value corresponding to the current date and time. Suppose that tables `t1` and `t2` have these definitions:

```
CREATE TABLE t1 (ts TIMESTAMP NULL DEFAULT '0000-00-00 00:00:00');
CREATE TABLE t2 (ts TIMESTAMP NULL DEFAULT NULL);
```

To set the `TIMESTAMP` column in either table to the current timestamp at insert time, explicitly assign it that value. For example:

```
INSERT INTO t1 VALUES (NOW());
INSERT INTO t2 VALUES (CURRENT_TIMESTAMP);
```

## 11.3.6. Fractional Seconds in Time Values

Before MySQL 5.6.4, the instances are limited in which a fractional seconds part is permitted in temporal values. A trailing fractional part is permissible in contexts such as literal values, and in the arguments to or return values from some temporal functions. Example:

```
mysql> SELECT MICROSECOND('2010-12-10 14:12:09.019473');
+-----+
| MICROSECOND('2010-12-10 14:12:09.019473') |
+-----+
|                                     19473 |
+-----+
```

However, when MySQL stores a value into a column of any temporal data type, it discards any fractional part and does not store it.

MySQL 5.6.4 and up expands fractional seconds support for `TIME`, `DATETIME`, and `TIMESTAMP` values, with up to microseconds (6 digits) precision:

- To define a column that includes a fractional seconds part, use the syntax `type_name(fsp)`, where `type_name` is `TIME`, `DATETIME`, or `TIMESTAMP`, and `fsp` is the fractional seconds precision. For example:

```
CREATE TABLE t1 (t TIME(3), dt DATETIME(6));
```

The `fsp` value, if given, must be in the range 0 to 6. A value of 0 signifies that there is no fractional part. If omitted, the default precision is 0. (This differs from the standard SQL default of 6, for compatibility with previous MySQL versions.)

- Functions that take temporal arguments accept values with fractional seconds. Return values from temporal functions include fractional seconds as appropriate. For example, `NOW()` [1220] with no argument returns the current date and time with no fractional part, but takes an optional argument from 0 to 6 to specify that the return value includes a fractional seconds part of that many digits.
- Syntax for temporal literals produces temporal values: `DATE 'str'`, `TIME 'str'`, and `TIMESTAMP 'str'`, and the ODBC-syntax equivalents. The resulting value includes a trailing fractional seconds part if specified. Previously, the temporal type keyword was ignored and these constructs produced the string value. See [Standard SQL and ODBC Date and Time Literals](#)

In some cases, previously accepted syntax may produce different results. The following items indicate where existing code may need to be changed to avoid problems:

- Some expressions produce results that differ from previous results. Examples: The `timestamp` system variable returns a value that includes a microseconds fractional part rather than an integer. Functions that return a result that includes the current time (such as `CURTIME()` [1213], `SYSDATE()` [1223], or `UTC_TIMESTAMP()` [1227]) interpret an argument as an `fsp` value and the return value includes a fractional seconds part of that many digits. Previously, these functions permitted an argument but ignored it.
- `TIME` values are converted to `DATETIME` by adding the time to the current date. (This means that the date part of the result differs from the current date if the time value is outside the range from `'00:00:00'` to `'23:59:59'`.) Previously, conversion of `TIME` values to `DATETIME` was unreliable. See [Section 11.3.7, "Conversion Between Date and Time Types"](#).
- `TIMESTAMP(N)` was permitted in old MySQL versions, but `N` was a display width rather than fractional seconds precision. Support for this behavior was removed in MySQL 5.5.3, so applications that are reasonably up to date should not be subject to this issue. Otherwise, code must be rewritten.

### 11.3.7. Conversion Between Date and Time Types

To some extent, you can convert a value from one temporal type to another. However, there may be some alteration of the value or loss of information. In all cases, conversion between temporal types is subject to the range of valid values for the resulting type. For example, although `DATE`, `DATETIME`, and `TIMESTAMP` values all can be specified using the same set of formats, the types do not all have the same range of values. `TIMESTAMP` values cannot be earlier than 1970 UTC or later than `'2038-01-19 03:14:07'` UTC. This means that a date such as `'1968-01-01'`, while valid as a `DATE` or `DATETIME` value, is not valid as a `TIMESTAMP` value and is converted to 0.

Conversion of `DATE` values:

- Conversion to a `DATETIME` or `TIMESTAMP` value adds a time part of `'00:00:00'` because the `DATE` value contains no time information.
- Conversion to a `TIME` value is not useful; the result is `'00:00:00'`.

Conversion of `DATETIME` and `TIMESTAMP` values:

- Conversion to a `DATE` value discards the time part because the `DATE` type contains no time information.
- Conversion to a `TIME` value discards the date part because the `TIME` type contains no date information.

Conversion of `TIME` values to other temporal types is version specific:

- As of MySQL 5.6.4, the value of `CURRENT_DATE()` [1212] is used for the date part. The `TIME` is interpreted as elapsed time (not time of day) and added to the date. This means that the date part of the result differs from the current date if the time value is outside the range from `'00:00:00'` to `'23:59:59'`.

Suppose that the current date is `'2012-01-01'`. `TIME` values of `'12:00:00'`, `'24:00:00'`, and `'-12:00:00'`, when converted to `DATETIME` or `TIMESTAMP` values, result in `'2012-01-01 12:00:00'`, `'2012-01-02 00:00:00'`, and `'2011-12-31 12:00:00'`, respectively.

Conversion of `TIME` to `DATE` is similar but discards the time part from the result: `'2012-01-01'`, `'2012-01-02'`, and `'2011-12-31'`, respectively.

- Before 5.6.4, MySQL converts a time value to a date or date-and-time value by parsing the string value of the time as a date or date-and-time. This is unlikely to be useful. For example, `'23:12:31'` interpreted as a date becomes `'2032-12-31'`. Time values not valid as dates become `'0000-00-00'` or `NULL`.

Explicit conversion can be used to override implicit conversion. For example, in comparison of `DATE` and `DATETIME` values, the `DATE` value is coerced to the `DATETIME` type by adding a time part of `'00:00:00'`. To perform the comparison by ignoring the time part of the `DATETIME` value instead, use the `CAST()` [1245] function in the following way:

```
date_col = CAST(datetime_col AS DATE)
```

Conversion of `TIME` and `DATETIME` values to numeric form (for example, by adding `+0`) occurs as follows:

- As of MySQL 5.6.4, `TIME(N)` or `DATETIME(N)` is converted to integer when `N` is 0 (or omitted) and to a `DECIMAL` value with `N` decimal digits when `N` is greater than 0:

```
mysql> SELECT CURTIME(), CURTIME()+0, CURTIME(3)+0;
+-----+-----+-----+
| CURTIME() | CURTIME()+0 | CURTIME(3)+0 |
+-----+-----+-----+
| 09:28:00 |          92800 |      92800.887 |
+-----+-----+-----+
mysql> SELECT NOW(), NOW()+0, NOW(3)+0;
+-----+-----+-----+
| NOW() | NOW()+0 | NOW(3)+0 |
+-----+-----+-----+
| 2012-08-15 09:28:00 | 20120815092800 | 20120815092800.889 |
+-----+-----+-----+
```

- Before MySQL 5.6.4, the conversion results in a double-precision value with a microseconds part of `.000000`:

```
mysql> SELECT CURTIME(), CURTIME()+0;
+-----+-----+
| CURTIME() | CURTIME()+0 |
+-----+-----+
| 09:28:00 | 92800.000000 |
+-----+-----+
mysql> SELECT NOW(), NOW()+0;
+-----+-----+
| NOW() | NOW()+0 |
+-----+-----+
| 2012-08-15 09:28:00 | 20120815092800.000000 |
+-----+-----+
```

---

### 11.3.8. Two-Digit Years in Dates

Date values with two-digit years are ambiguous because the century is unknown. Such values must be interpreted into four-digit form because MySQL stores years internally using four digits.

For `DATETIME`, `DATE`, and `TIMESTAMP` types, MySQL interprets dates specified with ambiguous year values using these rules:

- Year values in the range `00–69` are converted to `2000–2069`.
- Year values in the range `70–99` are converted to `1970–1999`.

For `YEAR`, the rules are the same, with this exception: A numeric `00` inserted into `YEAR(4)` results in `0000` rather than `2000`. To specify zero for `YEAR(4)` and have it be interpreted as `2000`, specify it as a string `'0'` or `'00'`.

Remember that these rules are only heuristics that provide reasonable guesses as to what your data values mean. If the rules used by MySQL do not produce the values you require, you must provide unambiguous input containing four-digit year values.

`ORDER BY` properly sorts `YEAR` values that have two-digit years.

Some functions like `MIN()` [1284] and `MAX()` [1284] convert a `YEAR` to a number. This means that a value with a two-digit year does not work properly with these functions. The fix in this case is to convert the `YEAR` to four-digit year format.

## 11.4. String Types

The string types are `CHAR`, `VARCHAR`, `BINARY`, `VARBINARY`, `BLOB`, `TEXT`, `ENUM`, and `SET`. This section describes how these types work and how to use them in your queries. For string type storage requirements, see [Section 11.6, “Data Type Storage Requirements”](#).

### 11.4.1. The `CHAR` and `VARCHAR` Types

The `CHAR` and `VARCHAR` types are similar, but differ in the way they are stored and retrieved. They also differ in maximum length and in whether trailing spaces are retained.

The `CHAR` and `VARCHAR` types are declared with a length that indicates the maximum number of characters you want to store. For example, `CHAR(30)` can hold up to 30 characters.

The length of a `CHAR` column is fixed to the length that you declare when you create the table. The length can be any value from 0 to 255. When `CHAR` values are stored, they are right-padded with spaces to the specified length. When `CHAR` values are retrieved, trailing spaces are removed unless the `PAD_CHAR_TO_FULL_LENGTH` [606] SQL mode is enabled.

Values in `VARCHAR` columns are variable-length strings. The length can be specified as a value from 0 to 65,535. The effective maximum length of a `VARCHAR` is subject to the maximum row size (65,535 bytes, which is shared among all columns) and the character set used. See [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

In contrast to `CHAR`, `VARCHAR` values are stored as a 1-byte or 2-byte length prefix plus data. The length prefix indicates the number of bytes in the value. A column uses one length byte if values require no more than 255 bytes, two length bytes if values may require more than 255 bytes.

If strict SQL mode is not enabled and you assign a value to a `CHAR` or `VARCHAR` column that exceeds the column's maximum length, the value is truncated to fit and a warning is generated. For truncation of non-space characters, you can cause an error to occur (rather than a warning) and suppress insertion of the value by using strict SQL mode. See [Section 5.1.7, “Server SQL Modes”](#).

For VARCHAR columns, trailing spaces in excess of the column length are truncated prior to insertion and a warning is generated, regardless of the SQL mode in use. For CHAR columns, truncation of excess trailing spaces from inserted values is performed silently regardless of the SQL mode.

VARCHAR values are not padded when they are stored. Trailing spaces are retained when values are stored and retrieved, in conformance with standard SQL.

The following table illustrates the differences between CHAR and VARCHAR by showing the result of storing various string values into CHAR(4) and VARCHAR(4) columns (assuming that the column uses a single-byte character set such as latin1).

Value	CHAR(4)	Storage Required	VARCHAR(4)	Storage Required
' '	'    '	4 bytes	' '	1 byte
'ab'	'ab  '	4 bytes	'ab'	3 bytes
'abcd'	'abcd'	4 bytes	'abcd'	5 bytes
'abcdefgh'	'abcd'	4 bytes	'abcd'	5 bytes

The values shown as stored in the last row of the table apply *only when not using strict mode*; if MySQL is running in strict mode, values that exceed the column length are *not stored*, and an error results.

If a given value is stored into the CHAR(4) and VARCHAR(4) columns, the values retrieved from the columns are not always the same because trailing spaces are removed from CHAR columns upon retrieval. The following example illustrates this difference:

```
mysql> CREATE TABLE vc (v VARCHAR(4), c CHAR(4));
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO vc VALUES ('ab ', 'ab ');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT CONCAT('( ', v, ')'), CONCAT('( ', c, ')) FROM vc;
+-----+-----+
| CONCAT('( ', v, ')) | CONCAT('( ', c, ')) |
+-----+-----+
| (ab )              | (ab)                |
+-----+-----+
1 row in set (0.06 sec)
```

Values in CHAR and VARCHAR columns are sorted and compared according to the character set collation assigned to the column.

All MySQL collations are of type PADSPACE. This means that all CHAR, VARCHAR, and TEXT values in MySQL are compared without regard to any trailing spaces. “Comparison” in this context does not include the LIKE [1190] pattern-matching operator, for which trailing spaces are significant. For example:

```
mysql> CREATE TABLE names (myname CHAR(10));
Query OK, 0 rows affected (0.03 sec)

mysql> INSERT INTO names VALUES ('Monty');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT myname = 'Monty', myname = 'Monty ' FROM names;
+-----+-----+
| myname = 'Monty' | myname = 'Monty ' |
+-----+-----+
| 1                | 1                  |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT myname LIKE 'Monty', myname LIKE 'Monty ' FROM names;
+-----+-----+
| myname LIKE 'Monty' | myname LIKE 'Monty ' |
+-----+-----+
```

```

+-----+-----+
|          1 |          0 |
+-----+-----+
1 row in set (0.00 sec)

```

This is true for all MySQL versions, and is not affected by the server SQL mode.



#### Note

For more information about MySQL character sets and collations, see [Section 10.1, “Character Set Support”](#).

For those cases where trailing pad characters are stripped or comparisons ignore them, if a column has an index that requires unique values, inserting into the column values that differ only in number of trailing pad characters will result in a duplicate-key error. For example, if a table contains 'a', an attempt to store 'a ' causes a duplicate-key error.

## 11.4.2. The `BINARY` and `VARBINARY` Types

The `BINARY` and `VARBINARY` types are similar to `CHAR` and `VARCHAR`, except that they contain binary strings rather than nonbinary strings. That is, they contain byte strings rather than character strings. This means that they have no character set, and sorting and comparison are based on the numeric values of the bytes in the values.

The permissible maximum length is the same for `BINARY` and `VARBINARY` as it is for `CHAR` and `VARCHAR`, except that the length for `BINARY` and `VARBINARY` is a length in bytes rather than in characters.

The `BINARY` and `VARBINARY` data types are distinct from the `CHAR BINARY` and `VARCHAR BINARY` data types. For the latter types, the `BINARY` attribute does not cause the column to be treated as a binary string column. Instead, it causes the binary collation for the column character set to be used, and the column itself contains nonbinary character strings rather than binary byte strings. For example, `CHAR(5) BINARY` is treated as `CHAR(5) CHARACTER SET latin1 COLLATE latin1_bin`, assuming that the default character set is `latin1`. This differs from `BINARY(5)`, which stores 5-byte binary strings that have no character set or collation. For information about differences between nonbinary string binary collations and binary strings, see [Section 10.1.7.6, “The `\_bin` and binary Collations”](#).

If strict SQL mode is not enabled and you assign a value to a `BINARY` or `VARBINARY` column that exceeds the column's maximum length, the value is truncated to fit and a warning is generated. For cases of truncation, you can cause an error to occur (rather than a warning) and suppress insertion of the value by using strict SQL mode. See [Section 5.1.7, “Server SQL Modes”](#).

When `BINARY` values are stored, they are right-padded with the pad value to the specified length. The pad value is `0x00` (the zero byte). Values are right-padded with `0x00` on insert, and no trailing bytes are removed on select. All bytes are significant in comparisons, including `ORDER BY` and `DISTINCT` operations. `0x00` bytes and spaces are different in comparisons, with `0x00 < space`.

Example: For a `BINARY(3)` column, 'a ' becomes 'a \0' when inserted. 'a\0' becomes 'a \0\0' when inserted. Both inserted values remain unchanged when selected.

For `VARBINARY`, there is no padding on insert and no bytes are stripped on select. All bytes are significant in comparisons, including `ORDER BY` and `DISTINCT` operations. `0x00` bytes and spaces are different in comparisons, with `0x00 < space`.

For those cases where trailing pad bytes are stripped or comparisons ignore them, if a column has an index that requires unique values, inserting into the column values that differ only in number of trailing pad bytes will result in a duplicate-key error. For example, if a table contains 'a', an attempt to store 'a\0' causes a duplicate-key error.

You should consider the preceding padding and stripping characteristics carefully if you plan to use the `BINARY` data type for storing binary data and you require that the value retrieved be exactly the same



as the value stored. The following example illustrates how 0x00-padding of BINARY values affects column value comparisons:

```
mysql> CREATE TABLE t (c BINARY(3));
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO t SET c = 'a';
Query OK, 1 row affected (0.01 sec)

mysql> SELECT HEX(c), c = 'a', c = 'a\0\0' from t;
+-----+-----+-----+
| HEX(c) | c = 'a' | c = 'a\0\0' |
+-----+-----+-----+
| 610000 | 0 | 1 |
+-----+-----+-----+
1 row in set (0.09 sec)
```

If the value retrieved must be the same as the value specified for storage with no padding, it might be preferable to use VARBINARY or one of the BLOB data types instead.

### 11.4.3. The BLOB and TEXT Types

A BLOB is a binary large object that can hold a variable amount of data. The four BLOB types are TINYBLOB, BLOB, MEDIUMBLOB, and LONGBLOB. These differ only in the maximum length of the values they can hold. The four TEXT types are TINYTEXT, TEXT, MEDIUMTEXT, and LONGTEXT. These correspond to the four BLOB types and have the same maximum lengths and storage requirements. See [Section 11.6, “Data Type Storage Requirements”](#).

BLOB values are treated as binary strings (byte strings). They have no character set, and sorting and comparison are based on the numeric values of the bytes in column values. TEXT values are treated as nonbinary strings (character strings). They have a character set, and values are sorted and compared based on the collation of the character set.

If strict SQL mode is not enabled and you assign a value to a BLOB or TEXT column that exceeds the column's maximum length, the value is truncated to fit and a warning is generated. For truncation of nonspace characters, you can cause an error to occur (rather than a warning) and suppress insertion of the value by using strict SQL mode. See [Section 5.1.7, “Server SQL Modes”](#).

Truncation of excess trailing spaces from values to be inserted into TEXT columns always generates a warning, regardless of the SQL mode.

For TEXT and BLOB columns, there is no padding on insert and no bytes are stripped on select.

If a TEXT column is indexed, index entry comparisons are space-padded at the end. This means that, if the index requires unique values, duplicate-key errors will occur for values that differ only in the number of trailing spaces. For example, if a table contains 'a', an attempt to store 'a ' causes a duplicate-key error. This is not true for BLOB columns.

In most respects, you can regard a BLOB column as a VARBINARY column that can be as large as you like. Similarly, you can regard a TEXT column as a VARCHAR column. BLOB and TEXT differ from VARBINARY and VARCHAR in the following ways:

- For indexes on BLOB and TEXT columns, you must specify an index prefix length. For CHAR and VARCHAR, a prefix length is optional. See [Section 8.3.4, “Column Indexes”](#).
- BLOB and TEXT columns cannot have DEFAULT values.

If you use the BINARY attribute with a TEXT data type, the column is assigned the binary collation of the column character set.

LONG and LONG VARCHAR map to the MEDIUMTEXT data type. This is a compatibility feature.

MySQL Connector/ODBC defines BLOB values as LONGVARBINARY and TEXT values as LONGVARCHAR.

Because `BLOB` and `TEXT` values can be extremely long, you might encounter some constraints in using them:

- Only the first `max_sort_length` [508] bytes of the column are used when sorting. The default value of `max_sort_length` [508] is 1024. You can make more bytes significant in sorting or grouping by increasing the value of `max_sort_length` [508] at server startup or runtime. Any client can change the value of its session `max_sort_length` [508] variable:

```
mysql> SET max_sort_length = 2000;
mysql> SELECT id, comment FROM t
-> ORDER BY comment;
```

- Instances of `BLOB` or `TEXT` columns in the result of a query that is processed using a temporary table causes the server to use a table on disk rather than in memory because the `MEMORY` storage engine does not support those data types (see [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#)). Use of disk incurs a performance penalty, so include `BLOB` or `TEXT` columns in the query result only if they are really needed. For example, avoid using `SELECT *`, which selects all columns.
- The maximum size of a `BLOB` or `TEXT` object is determined by its type, but the largest value you actually can transmit between the client and server is determined by the amount of available memory and the size of the communications buffers. You can change the message buffer size by changing the value of the `max_allowed_packet` [502] variable, but you must do so for both the server and your client program. For example, both `mysql` and `mysqldump` enable you to change the client-side `max_allowed_packet` [502] value. See [Section 8.11.2, “Tuning Server Parameters”](#), [Section 4.5.1, “mysql — The MySQL Command-Line Tool”](#), and [Section 4.5.4, “mysqldump — A Database Backup Program”](#). You may also want to compare the packet sizes and the size of the data objects you are storing with the storage requirements, see [Section 11.6, “Data Type Storage Requirements”](#)

Each `BLOB` or `TEXT` value is represented internally by a separately allocated object. This is in contrast to all other data types, for which storage is allocated once per column when the table is opened.

In some cases, it may be desirable to store binary data such as media files in `BLOB` or `TEXT` columns. You may find MySQL's string handling functions useful for working with such data. See [Section 12.5, “String Functions”](#). For security and other reasons, it is usually preferable to do so using application code rather than giving application users the `FILE` [754] privilege. You can discuss specifics for various languages and platforms in the MySQL Forums (<http://forums.mysql.com/>).

#### 11.4.4. The ENUM Type

An `ENUM` is a string object with a value chosen from a list of permitted values that are enumerated explicitly in the column specification at table creation time. It has these advantages:

- Compact data storage in situations where a column has a limited set of possible values. The strings you specify as input values are automatically encoded as numbers. See [Section 11.6, “Data Type Storage Requirements”](#) for the storage requirements for `ENUM` types.
- Readable queries and output. The numbers are translated back to the corresponding strings in query results.

and these potential issues to consider:

- If you make enumeration values that look like numbers, it is easy to mix up the literal values with their internal index numbers, as explained in [Section 11.4.4, “Limitations”](#).
- Using `ENUM` columns in `ORDER BY` clauses requires extra care, as explained in [Section 11.4.4, “Sorting”](#).

#### Creating and Using ENUM Columns

An enumeration value must be a quoted string literal. For example, you can create a table with an `ENUM` column like this:

```
CREATE TABLE shirts (
  name VARCHAR(40),
  size ENUM('x-small', 'small', 'medium', 'large', 'x-large')
);
INSERT INTO shirts (name, size) VALUES ('dress shirt','large'), ('t-shirt','medium'),
('polo shirt','small');
SELECT name, size FROM shirts WHERE size = 'medium';
+-----+-----+
| name   | size   |
+-----+-----+
| t-shirt | medium |
+-----+-----+
UPDATE shirts SET size = 'small' WHERE size = 'large';
COMMIT;
```

Inserting 1 million rows into this table with a value of 'medium' would require 1 million bytes of storage, as opposed to 6 million bytes if you stored the actual string 'medium' in a VARCHAR column.

### Index Values for Enumeration Literals

Each enumeration value has an index:

- The elements listed in the column specification are assigned index numbers, beginning with 1.
- The index value of the empty string error value is 0. This means that you can use the following SELECT statement to find rows into which invalid ENUM values were assigned:

```
mysql> SELECT * FROM tbl_name WHERE enum_col=0;
```

- The index of the NULL value is NULL.
- The term “index” here refers to a position within the list of enumeration values. It has nothing to do with table indexes.

For example, a column specified as ENUM('Mercury', 'Venus', 'Earth') can have any of the values shown here. The index of each value is also shown.

Value	Index
NULL	NULL
' '	0
'Mercury'	1
'Venus'	2
'Earth'	3

An ENUM column can have a maximum of 65,535 distinct elements. (The practical limit is less than 3000.) A table can have no more than 255 unique element list definitions among its ENUM and SET columns considered as a group. For more information on these limits, see Section E.10.5, “Limits Imposed by .frm File Structure”.

If you retrieve an ENUM value in a numeric context, the column value's index is returned. For example, you can retrieve numeric values from an ENUM column like this:

```
mysql> SELECT enum_col+0 FROM tbl_name;
```

Functions such as SUM() [1285] or AVG() [1282] that expect a numeric argument cast the argument to a number if necessary. For ENUM values, the index number is used in the calculation.

### Handling of Enumeration Literals

Trailing spaces are automatically deleted from ENUM member values in the table definition when a table is created.

When retrieved, values stored into an `ENUM` column are displayed using the lettercase that was used in the column definition. Note that `ENUM` columns can be assigned a character set and collation. For binary or case-sensitive collations, lettercase is taken into account when assigning values to the column.

If you store a number into an `ENUM` column, the number is treated as the index into the possible values, and the value stored is the enumeration member with that index. (However, this does *not* work with `LOAD DATA`, which treats all input as strings.) If the numeric value is quoted, it is still interpreted as an index if there is no matching string in the list of enumeration values. For these reasons, it is not advisable to define an `ENUM` column with enumeration values that look like numbers, because this can easily become confusing. For example, the following column has enumeration members with string values of `'0'`, `'1'`, and `'2'`, but numeric index values of 1, 2, and 3:

```
numbers ENUM('0','1','2')
```

If you store 2, it is interpreted as an index value, and becomes `'1'` (the value with index 2). If you store `'2'`, it matches an enumeration value, so it is stored as `'2'`. If you store `'3'`, it does not match any enumeration value, so it is treated as an index and becomes `'2'` (the value with index 3).

```
mysql> INSERT INTO t (numbers) VALUES(2),('2'),('3');
mysql> SELECT * FROM t;
+-----+
| numbers |
+-----+
| 1       |
| 2       |
| 2       |
+-----+
```

To determine all possible values for an `ENUM` column, use `SHOW COLUMNS FROM tbl_name LIKE 'enum_col'` and parse the `ENUM` definition in the `Type` column of the output.

In the C API, `ENUM` values are returned as strings. For information about using result set metadata to distinguish them from other strings, see [Section 21.8.4, “C API Data Structures”](#).

## Empty or NULL Enumeration Values

An enumeration value can also be the empty string (`' '`) or `NULL` under certain circumstances:

- If you insert an invalid value into an `ENUM` (that is, a string not present in the list of permitted values), the empty string is inserted instead as a special error value. This string can be distinguished from a “normal” empty string by the fact that this string has the numeric value 0. See [Section 11.4.4, “Index Values for Enumeration Literals”](#) for details about the numeric indexes for the enumeration values.

If strict SQL mode is enabled, attempts to insert invalid `ENUM` values result in an error.

- If an `ENUM` column is declared to permit `NULL`, the `NULL` value is a valid value for the column, and the default value is `NULL`. If an `ENUM` column is declared `NOT NULL`, its default value is the first element of the list of permitted values.

## Sorting

`ENUM` values are sorted based on their index numbers, which depend on the order in which the enumeration members were listed in the column specification. For example, `'b'` sorts before `'a'` for `ENUM('b', 'a')`. The empty string sorts before nonempty strings, and `NULL` values sort before all other enumeration values.

To prevent unexpected results when using the `ORDER BY` clause on an `ENUM` column, use one of these techniques:

- Specify the `ENUM` list in alphabetic order.
- Make sure that the column is sorted lexically rather than by index number by coding `ORDER BY CAST(col AS CHAR)` or `ORDER BY CONCAT(col)`.

## Limitations

An enumeration value cannot be an expression, even one that evaluates to a string value.

For example, this `CREATE TABLE` statement does *not* work because the `CONCAT` function cannot be used to construct an enumeration value:

```
CREATE TABLE sizes (
  size ENUM('small', CONCAT('med','ium'), 'large')
);
```

You also cannot employ a user variable as an enumeration value. This pair of statements do *not* work:

```
SET @mysize = 'medium';

CREATE TABLE sizes (
  size ENUM('small', @mysize, 'large')
);
```

We strongly recommend that you do *not* use numbers as enumeration values, because it does not save on storage over the appropriate `TINYINT` or `SMALLINT` type, and it is easy to mix up the strings and the underlying number values (which might not be the same) if you quote the `ENUM` values incorrectly. If you do use a number as an enumeration value, always enclose it in quotation marks. If the quotation marks are omitted, the number is regarded as an index. See [Section 11.4.4, “Handling of Enumeration Literals”](#) to see how even a quoted number could be mistakenly used as a numeric index value.

Duplicate values in the definition cause a warning, or an error if strict SQL mode is enabled.

### 11.4.5. The SET Type

A `SET` is a string object that can have zero or more values, each of which must be chosen from a list of permitted values specified when the table is created. `SET` column values that consist of multiple set members are specified with members separated by commas (`,`). A consequence of this is that `SET` member values should not themselves contain commas.

For example, a column specified as `SET('one', 'two') NOT NULL` can have any of these values:

```
' '
'one'
'two'
'one,two'
```

A `SET` column can have a maximum of 64 distinct members. A table can have no more than 255 unique element list definitions among its `ENUM` and `SET` columns considered as a group. For more information on this limit, see [Section E.10.5, “Limits Imposed by .frm File Structure”](#).

Duplicate values in the definition cause a warning, or an error if strict SQL mode is enabled.

Trailing spaces are automatically deleted from `SET` member values in the table definition when a table is created.

When retrieved, values stored in a `SET` column are displayed using the lettercase that was used in the column definition. Note that `SET` columns can be assigned a character set and collation. For binary or case-sensitive collations, lettercase is taken into account when assigning values to the column.

MySQL stores `SET` values numerically, with the low-order bit of the stored value corresponding to the first set member. If you retrieve a `SET` value in a numeric context, the value retrieved has bits set corresponding to the set members that make up the column value. For example, you can retrieve numeric values from a `SET` column like this:

```
mysql> SELECT set_col+0 FROM tbl_name;
```

If a number is stored into a SET column, the bits that are set in the binary representation of the number determine the set members in the column value. For a column specified as SET('a','b','c','d'), the members have the following decimal and binary values.

SET Member	Decimal Value	Binary Value
'a'	1	0001
'b'	2	0010
'c'	4	0100
'd'	8	1000

If you assign a value of 9 to this column, that is 1001 in binary, so the first and fourth SET value members 'a' and 'd' are selected and the resulting value is 'a,d'.

For a value containing more than one SET element, it does not matter what order the elements are listed in when you insert the value. It also does not matter how many times a given element is listed in the value. When the value is retrieved later, each element in the value appears once, with elements listed according to the order in which they were specified at table creation time. For example, suppose that a column is specified as SET('a','b','c','d'):

```
mysql> CREATE TABLE myset (col SET('a', 'b', 'c', 'd'));
```

If you insert the values 'a,d', 'd,a', 'a,d,d', 'a,d,a', and 'd,a,d':

```
mysql> INSERT INTO myset (col) VALUES
-> ('a,d'), ('d,a'), ('a,d,a'), ('a,d,d'), ('d,a,d');
Query OK, 5 rows affected (0.01 sec)
Records: 5 Duplicates: 0 Warnings: 0
```

Then all these values appear as 'a,d' when retrieved:

```
mysql> SELECT col FROM myset;
+-----+
| col   |
+-----+
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
+-----+
5 rows in set (0.04 sec)
```

If you set a SET column to an unsupported value, the value is ignored and a warning is issued:

```
mysql> INSERT INTO myset (col) VALUES ('a,d,d,s');
Query OK, 1 row affected, 1 warning (0.03 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                                     |
+-----+-----+-----+
| Warning | 1265 | Data truncated for column 'col' at row 1 |
+-----+-----+-----+
1 row in set (0.04 sec)

mysql> SELECT col FROM myset;
+-----+
| col   |
+-----+
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
+-----+
```

```
+-----+
6 rows in set (0.01 sec)
```

If strict SQL mode is enabled, attempts to insert invalid `SET` values result in an error.

`SET` values are sorted numerically. `NULL` values sort before non-`NULL` `SET` values.

Functions such as `SUM()` [1285] or `AVG()` [1282] that expect a numeric argument cast the argument to a number if necessary. For `SET` values, the cast operation causes the numeric value to be used.

Normally, you search for `SET` values using the `FIND_IN_SET()` [1179] function or the `LIKE` [1190] operator:

```
mysql> SELECT * FROM tbl_name WHERE FIND_IN_SET('value',set_col)>0;
mysql> SELECT * FROM tbl_name WHERE set_col LIKE '%value%';
```

The first statement finds rows where `set_col` contains the `value` set member. The second is similar, but not the same: It finds rows where `set_col` contains `value` anywhere, even as a substring of another set member.

The following statements also are permitted:

```
mysql> SELECT * FROM tbl_name WHERE set_col & 1;
mysql> SELECT * FROM tbl_name WHERE set_col = 'val1,val2';
```

The first of these statements looks for values containing the first set member. The second looks for an exact match. Be careful with comparisons of the second type. Comparing set values to `'val1,val2'` returns different results than comparing values to `'val2,val1'`. You should specify the values in the same order they are listed in the column definition.

To determine all possible values for a `SET` column, use `SHOW COLUMNS FROM tbl_name LIKE set_col` and parse the `SET` definition in the `Type` column of the output.

In the C API, `SET` values are returned as strings. For information about using result set metadata to distinguish them from other strings, see [Section 21.8.4, “C API Data Structures”](#).

## 11.5. Data Type Default Values

The `DEFAULT value` clause in a data type specification indicates a default value for a column. With one exception, the default value must be a constant; it cannot be a function or an expression. This means, for example, that you cannot set the default for a date column to be the value of a function such as `NOW()` [1220] or `CURRENT_DATE` [1212]. The exception is that you can specify `CURRENT_TIMESTAMP` [1213] as the default for `TIMESTAMP` and `DATETIME` columns. See [Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”](#).

`BLOB` and `TEXT` columns cannot be assigned a default value.

If a column definition includes no explicit `DEFAULT` value, MySQL determines the default value as follows:

If the column can take `NULL` as a value, the column is defined with an explicit `DEFAULT NULL` clause.

If the column cannot take `NULL` as the value, MySQL defines the column with no explicit `DEFAULT` clause. Exception: If the column is defined as part of a `PRIMARY KEY` but not explicitly as `NOT NULL`, MySQL creates it as a `NOT NULL` column (because `PRIMARY KEY` columns must be `NOT NULL`), but also assigns it a `DEFAULT` clause using the implicit default value. To prevent this, include an explicit `NOT NULL` in the definition of any `PRIMARY KEY` column.

For data entry into a `NOT NULL` column that has no explicit `DEFAULT` clause, if an `INSERT` or `REPLACE` statement includes no value for the column, or an `UPDATE` statement sets the column to `NULL`, MySQL handles the column according to the SQL mode in effect at the time:



- If strict SQL mode is enabled, an error occurs for transactional tables and the statement is rolled back. For nontransactional tables, an error occurs, but if this happens for the second or subsequent row of a multiple-row statement, the preceding rows will have been inserted.
- If strict mode is not enabled, MySQL sets the column to the implicit default value for the column data type.

Suppose that a table `t` is defined as follows:

```
CREATE TABLE t (i INT NOT NULL);
```

In this case, `i` has no explicit default, so in strict mode each of the following statements produce an error and no row is inserted. When not using strict mode, only the third statement produces an error; the implicit default is inserted for the first two statements, but the third fails because `DEFAULT(i)` [1275] cannot produce a value:

```
INSERT INTO t VALUES();
INSERT INTO t VALUES(DEFAULT);
INSERT INTO t VALUES(DEFAULT(i));
```

See [Section 5.1.7, “Server SQL Modes”](#).

For a given table, you can use the `SHOW CREATE TABLE` statement to see which columns have an explicit `DEFAULT` clause.

Implicit defaults are defined as follows:

- For numeric types, the default is `0`, with the exception that for integer or floating-point types declared with the `AUTO_INCREMENT` attribute, the default is the next value in the sequence.
- For date and time types other than `TIMESTAMP`, the default is the appropriate “zero” value for the type. This is also true for `TIMESTAMP` if the `explicit_defaults_for_timestamp` [481] system variable is enabled (see [Section 5.1.4, “Server System Variables”](#)). Otherwise, for the first `TIMESTAMP` column in a table, the default value is the current date and time. See [Section 11.3, “Date and Time Types”](#).
- For string types other than `ENUM`, the default value is the empty string. For `ENUM`, the default is the first enumeration value.

`SERIAL DEFAULT VALUE` in the definition of an integer column is an alias for `NOT NULL AUTO_INCREMENT UNIQUE`.

## 11.6. Data Type Storage Requirements

The storage requirements for table data on disk depend on several factors. Different storage engines represent data types and store raw data differently. Table data might be compressed, either for a column or an entire row, complicating the calculation of storage requirements for a table or column.

Despite differences in storage layout on disk, the internal MySQL APIs that communicate and exchange information about table rows use a consistent data structure that applies across all storage engines.

This section includes guidelines and information for the storage requirements for each data type supported by MySQL, including the internal format and size for storage engines that use a fixed-size representation for data types. Information is listed by category or storage engine.

The internal representation of a table has a maximum row size of 65,535 bytes, even if the storage engine is capable of supporting larger rows. This figure excludes `BLOB` or `TEXT` columns, which contribute only 9 to 12 bytes toward this size. For `BLOB` and `TEXT` data, the information is stored internally in a different area of memory than the row buffer. Different storage engines handle the allocation and storage of this data in different ways, according to the method they use for handling the

corresponding types. For more information, see [Chapter 14, Storage Engines](#), and [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

## InnoDB Tables

See [Section 14.2.3.12.7, “Physical Row Structure”](#) for information about storage requirements for InnoDB tables.

## NDBCLUSTER Tables



### Important

NDB tables use *4-byte alignment*; all NDB data storage is done in multiples of 4 bytes. Thus, a column value that would typically take 15 bytes requires 16 bytes in an NDB table. For example, in NDB tables, the `TINYINT`, `SMALLINT`, `MEDIUMINT`, and `INTEGER (INT)` column types each require 4 bytes storage per record due to the alignment factor.

Each `BIT(M)` column takes  $M$  bits of storage space. Although an individual `BIT` column is *not* 4-byte aligned, NDB reserves 4 bytes (32 bits) per row for the first 1-32 bits needed for `BIT` columns, then another 4 bytes for bits 33-64, and so on.

While a `NULL` itself does not require any storage space, NDB reserves 4 bytes per row if the table definition contains any columns defined as `NULL`, up to 32 `NULL` columns. (If a MySQL Cluster table is defined with more than 32 `NULL` columns up to 64 `NULL` columns, then 8 bytes per row are reserved.)

Every table using the NDB storage engine requires a primary key; if you do not define a primary key, a “hidden” primary key is created by NDB. This hidden primary key consumes 31-35 bytes per table record.

You can use the `ndb_size.pl` Perl script to estimate NDB storage requirements. It connects to a current MySQL (not MySQL Cluster) database and creates a report on how much space that database would require if it used the NDB storage engine. See [ndb\\_size.pl — NDBCLUSTER Size Requirement Estimator](#) for more information.

## Storage Requirements for Numeric Types

Data Type	Storage Required
<code>TINYINT</code>	1 byte
<code>SMALLINT</code>	2 bytes
<code>MEDIUMINT</code>	3 bytes
<code>INT, INTEGER</code>	4 bytes
<code>BIGINT</code>	8 bytes
<code>FLOAT(p)</code>	4 bytes if $0 \leq p \leq 24$ , 8 bytes if $25 \leq p \leq 53$
<code>FLOAT</code>	4 bytes
<code>DOUBLE [PRECISION], REAL</code>	8 bytes
<code>DECIMAL(M,D), NUMERIC(M,D)</code>	Varies; see following discussion
<code>BIT(M)</code>	approximately $(M+7)/8$ bytes

Values for `DECIMAL` (and `NUMERIC`) columns are represented using a binary format that packs nine decimal (base 10) digits into four bytes. Storage for the integer and fractional parts of each value are determined separately. Each multiple of nine digits requires four bytes, and the “leftover” digits require some fraction of four bytes. The storage required for excess digits is given by the following table.

Leftover Digits	Number of Bytes
0	0
1	1
2	1
3	2
4	2
5	3
6	3
7	4
8	4

## Storage Requirements for Date and Time Types

For `TIME`, `DATETIME`, and `TIMESTAMP` columns, the storage required for tables created before MySQL 5.6.4 differs from tables created from 5.6.4 on. This is due to a change in 5.6.4 that permits these types to have a fractional part, which requires from 0 to 3 bytes.

Data Type	Storage Required Before MySQL 5.6.4	Storage Required as of MySQL 5.6.4
<code>YEAR</code>	1 byte	1 byte
<code>DATE</code>	3 bytes	3 bytes
<code>TIME</code>	3 bytes	3 bytes + fractional seconds storage
<code>DATETIME</code>	8 bytes	5 bytes + fractional seconds storage
<code>TIMESTAMP</code>	4 bytes	4 bytes + fractional seconds storage

As of MySQL 5.6.4, storage for `YEAR` and `DATE` remains unchanged. However, `TIME`, `DATETIME`, and `TIMESTAMP` are represented differently. `DATETIME` is packed more efficiently, requiring 5 rather than 8 bytes for the nonfractional part, and all three parts have a fractional part that requires from 0 to 3 bytes, depending on the fractional seconds precision of stored values.

Fractional Seconds Precision	Storage Required
0	0 bytes
1, 2	1 byte
3, 4	2 bytes
5, 6	3 bytes

For example, `TIME(0)`, `TIME(2)`, `TIME(4)`, and `TIME(6)` use 3, 4, 5, and 6 bytes, respectively. `TIME` and `TIME(0)` are equivalent and require the same storage.

For details about internal representation of temporal values, see [MySQL Internals: Important Algorithms and Structures](#).

## Storage Requirements for String Types

In the following table, *M* represents the declared column length in characters for nonbinary string types and bytes for binary string types. *L* represents the actual length in bytes of a given string value.

Data Type	Storage Required
<code>CHAR(M)</code>	$M \times w$ bytes, $0 \leq M \leq 255$ , where <i>w</i> is the number of bytes required for the maximum-length character in the character set

Data Type	Storage Required
<code>BINARY(M)</code>	$M$ bytes, $0 \leq M \leq 255$
<code>VARCHAR(M)</code> , <code>VARBINARY(M)</code>	$L + 1$ bytes if column values require 0 – 255 bytes, $L + 2$ bytes if values may require more than 255 bytes
<code>TINYBLOB</code> , <code>TINYTEXT</code>	$L + 1$ bytes, where $L < 2^8$
<code>BLOB</code> , <code>TEXT</code>	$L + 2$ bytes, where $L < 2^{16}$
<code>MEDIUMBLOB</code> , <code>MEDIUMTEXT</code>	$L + 3$ bytes, where $L < 2^{24}$
<code>LONGBLOB</code> , <code>LONGTEXT</code>	$L + 4$ bytes, where $L < 2^{32}$
<code>ENUM('value1', 'value2', ...)</code>	1 or 2 bytes, depending on the number of enumeration values (65,535 values maximum)
<code>SET('value1', 'value2', ...)</code>	1, 2, 3, 4, or 8 bytes, depending on the number of set members (64 members maximum)

Variable-length string types are stored using a length prefix plus data. The length prefix requires from one to four bytes depending on the data type, and the value of the prefix is  $L$  (the byte length of the string). For example, storage for a `MEDIUMTEXT` value requires  $L$  bytes to store the value plus three bytes to store the length of the value.

To calculate the number of bytes used to store a particular `CHAR`, `VARCHAR`, or `TEXT` column value, you must take into account the character set used for that column and whether the value contains multi-byte characters. In particular, when using the `utf8` (or `utf8mb4`) Unicode character set, you must keep in mind that not all characters use the same number of bytes and can require up to three (four) bytes per character. For a breakdown of the storage used for different categories of `utf8` or `utf8mb4` characters, see [Section 10.1.10, “Unicode Support”](#).

`VARCHAR`, `VARBINARY`, and the `BLOB` and `TEXT` types are variable-length types. For each, the storage requirements depend on these factors:

- The actual length of the column value
- The column's maximum possible length
- The character set used for the column, because some character sets contain multi-byte characters

For example, a `VARCHAR(255)` column can hold a string with a maximum length of 255 characters. Assuming that the column uses the `latin1` character set (one byte per character), the actual storage required is the length of the string ( $L$ ), plus one byte to record the length of the string. For the string `'abcd'`,  $L$  is 4 and the storage requirement is five bytes. If the same column is instead declared to use the `ucs2` double-byte character set, the storage requirement is 10 bytes: The length of `'abcd'` is eight bytes and the column requires two bytes to store lengths because the maximum length is greater than 255 (up to 510 bytes).

The effective maximum number of *bytes* that can be stored in a `VARCHAR` or `VARBINARY` column is subject to the maximum row size of 65,535 bytes, which is shared among all columns. For a `VARCHAR` column that stores multi-byte characters, the effective maximum number of *characters* is less. For example, `utf8` characters can require up to three bytes per character, so a `VARCHAR` column that uses the `utf8` character set can be declared to be a maximum of 21,844 characters. See [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

The size of an `ENUM` object is determined by the number of different enumeration values. One byte is used for enumerations with up to 255 possible values. Two bytes are used for enumerations having between 256 and 65,535 possible values. See [Section 11.4.4, “The ENUM Type”](#).

The size of a `SET` object is determined by the number of different set members. If the set size is  $N$ , the object occupies  $(N+7) / 8$  bytes, rounded up to 1, 2, 3, 4, or 8 bytes. A `SET` can have a maximum of 64 members. See [Section 11.4.5, “The SET Type”](#).

## 11.7. Choosing the Right Type for a Column

For optimum storage, you should try to use the most precise type in all cases. For example, if an integer column is used for values in the range from 1 to 99999, `MEDIUMINT UNSIGNED` is the best type. Of the types that represent all the required values, this type uses the least amount of storage.

All basic calculations (+, -, \*, and /) with `DECIMAL` columns are done with precision of 65 decimal (base 10) digits. See [Section 11.1.1, “Numeric Type Overview”](#).

If accuracy is not too important or if speed is the highest priority, the `DOUBLE` type may be good enough. For high precision, you can always convert to a fixed-point type stored in a `BIGINT`. This enables you to do all calculations with 64-bit integers and then convert results back to floating-point values as necessary.

`PROCEDURE ANALYSE` can be used to obtain suggestions for optimal column data types. For more information, see [Section 8.4.2.4, “Using PROCEDURE ANALYSE”](#).

## 11.8. Using Data Types from Other Database Engines

To facilitate the use of code written for SQL implementations from other vendors, MySQL maps data types as shown in the following table. These mappings make it easier to import table definitions from other database systems into MySQL.

Other Vendor Type	MySQL Type
BOOL	TINYINT
BOOLEAN	TINYINT
CHARACTER VARYING(M)	VARCHAR(M)
FIXED	DECIMAL
FLOAT4	FLOAT
FLOAT8	DOUBLE
INT1	TINYINT
INT2	SMALLINT
INT3	MEDIUMINT
INT4	INT
INT8	BIGINT
LONG VARBINARY	MEDIUMBLOB
LONG VARCHAR	MEDIUMTEXT
LONG	MEDIUMTEXT
MIDDLEINT	MEDIUMINT
NUMERIC	DECIMAL

Data type mapping occurs at table creation time, after which the original type specifications are discarded. If you create a table with types used by other vendors and then issue a `DESCRIBE tbl_name` statement, MySQL reports the table structure using the equivalent MySQL types. For example:

```
mysql> CREATE TABLE t (a BOOL, b FLOAT8, c LONG VARCHAR, d NUMERIC);
Query OK, 0 rows affected (0.00 sec)

mysql> DESCRIBE t;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| a     | tinyint(1)   | YES  |     | NULL    |       |
```

```
| b      | double | YES | | NULL | |  
| c      | mediumtext | YES | | NULL | |  
| d      | decimal(10,0) | YES | | NULL | |  
+-----+-----+-----+-----+-----+  
4 rows in set (0.01 sec)
```

---

# Chapter 12. Functions and Operators

## Table of Contents

12.1. Function and Operator Reference .....	1154
12.2. Type Conversion in Expression Evaluation .....	1161
12.3. Operators .....	1164
12.3.1. Operator Precedence .....	1165
12.3.2. Comparison Functions and Operators .....	1165
12.3.3. Logical Operators .....	1171
12.3.4. Assignment Operators .....	1172
12.4. Control Flow Functions .....	1173
12.5. String Functions .....	1175
12.5.1. String Comparison Functions .....	1190
12.5.2. Regular Expressions .....	1193
12.6. Numeric Functions and Operators .....	1198
12.6.1. Arithmetic Operators .....	1199
12.6.2. Mathematical Functions .....	1201
12.7. Date and Time Functions .....	1209
12.8. What Calendar Is Used By MySQL? .....	1229
12.9. Full-Text Search Functions .....	1229
12.9.1. Natural Language Full-Text Searches .....	1231
12.9.2. Boolean Full-Text Searches .....	1234
12.9.3. Full-Text Searches with Query Expansion .....	1236
12.9.4. Full-Text Stopwords .....	1237
12.9.5. Full-Text Restrictions .....	1240
12.9.6. Fine-Tuning MySQL Full-Text Search .....	1241
12.9.7. Adding a Collation for Full-Text Indexing .....	1243
12.10. Cast Functions and Operators .....	1244
12.11. XML Functions .....	1247
12.12. Bit Functions .....	1257
12.13. Encryption and Compression Functions .....	1259
12.14. Information Functions .....	1265
12.15. GTID Functions .....	1273
12.16. Miscellaneous Functions .....	1274
12.17. Functions and Modifiers for Use with <code>GROUP BY</code> Clauses .....	1281
12.17.1. <code>GROUP BY</code> (Aggregate) Functions .....	1281
12.17.2. <code>GROUP BY</code> Modifiers .....	1285
12.17.3. MySQL Extensions to <code>GROUP BY</code> .....	1288
12.18. Spatial Extensions .....	1289
12.18.1. Introduction to MySQL Spatial Support .....	1289
12.18.2. The OpenGIS Geometry Model .....	1290
12.18.3. Supported Spatial Data Formats .....	1296
12.18.4. Creating a Spatially Enabled MySQL Database .....	1297
12.18.5. Spatial Analysis Functions .....	1302
12.18.6. Optimizing Spatial Analysis .....	1312
12.18.7. MySQL Conformance and Compatibility .....	1315
12.19. Precision Math .....	1315
12.19.1. Types of Numeric Values .....	1316
12.19.2. <code>DECIMAL</code> Data Type Changes .....	1316
12.19.3. Expression Handling .....	1317
12.19.4. Rounding Behavior .....	1319
12.19.5. Precision Math Examples .....	1319

Expressions can be used at several points in `SQL` statements, such as in the `ORDER BY` or `HAVING` clauses of `SELECT` statements, in the `WHERE` clause of a `SELECT`, `DELETE`, or `UPDATE` statement,



or in `SET` statements. Expressions can be written using literal values, column values, `NULL`, built-in functions, stored functions, user-defined functions, and operators. This chapter describes the functions and operators that are permitted for writing expressions in MySQL. Instructions for writing stored functions and user-defined functions are given in [Section 18.2, “Using Stored Routines \(Procedures and Functions\)”](#), and [Section 22.3, “Adding New Functions to MySQL”](#). See [Section 9.2.4, “Function Name Parsing and Resolution”](#), for the rules describing how the server interprets references to different kinds of functions.

An expression that contains `NULL` always produces a `NULL` value unless otherwise indicated in the documentation for a particular function or operator.



#### Note

By default, there must be no whitespace between a function name and the parenthesis following it. This helps the MySQL parser distinguish between function calls and references to tables or columns that happen to have the same name as a function. However, spaces around function arguments are permitted.

You can tell the MySQL server to accept spaces after function names by starting it with the `--sql-mode=IGNORE_SPACE` [442] option. (See [Section 5.1.7, “Server SQL Modes”](#).) Individual client programs can request this behavior by using the `CLIENT_IGNORE_SPACE` option for `mysql_real_connect()`. In either case, all function names become reserved words.

For the sake of brevity, most examples in this chapter display the output from the `mysql` program in abbreviated form. Rather than showing examples in this format:

```
mysql> SELECT MOD(29,9);
+-----+
| mod(29,9) |
+-----+
|         2 |
+-----+
1 rows in set (0.00 sec)
```

This format is used instead:

```
mysql> SELECT MOD(29,9);
-> 2
```

## 12.1. Function and Operator Reference

Table 12.1. Functions/Operators

Name	Description
<a href="#">ABS()</a> [1202]	Return the absolute value
<a href="#">ACOS()</a> [1202]	Return the arc cosine
<a href="#">ADDDATE()</a> [1211]	Add time values (intervals) to a date value
<a href="#">ADDTIME()</a> [1212]	Add time
<a href="#">AES_DECRYPT()</a> [1260]	Decrypt using AES
<a href="#">AES_ENCRYPT()</a> [1260]	Encrypt using AES
<a href="#">AND, &amp;&amp;</a> [1171]	Logical AND
<a href="#">ASCII()</a> [1177]	Return numeric value of left-most character
<a href="#">ASIN()</a> [1202]	Return the arc sine
<a href="#">=</a> [1173]	Assign a value (as part of a <code>SET</code> statement, or as part of the <code>SET</code> clause in an <code>UPDATE</code> statement)
<a href="#">:=</a> [1172]	Assign a value
<a href="#">ATAN2()</a> , <a href="#">ATAN()</a> [1203]	Return the arc tangent of the two arguments
<a href="#">ATAN()</a> [1202]	Return the arc tangent

<b>Name</b>	<b>Description</b>
<a href="#">AVG ( ) [1282]</a>	Return the average value of the argument
<a href="#">BENCHMARK ( ) [1266]</a>	Repeatedly execute an expression
<a href="#">BETWEEN ... AND ... [1168]</a>	Check whether a value is within a range of values
<a href="#">BIN ( ) [1177]</a>	Return a string containing binary representation of a number
<a href="#">BINARY [1244]</a>	Cast a string to a binary string
<a href="#">BIT_AND ( ) [1282]</a>	Return bitwise and
<a href="#">BIT_COUNT ( ) [1259]</a>	Return the number of bits that are set
<a href="#">BIT_LENGTH ( ) [1177]</a>	Return length of argument in bits
<a href="#">BIT_OR ( ) [1282]</a>	Return bitwise or
<a href="#">BIT_XOR ( ) [1283]</a>	Return bitwise xor
<a href="#">&amp; [1258]</a>	Bitwise AND
<a href="#">~ [1259]</a>	Invert bits
<a href="#">  [1258]</a>	Bitwise OR
<a href="#">^ [1258]</a>	Bitwise XOR
<a href="#">CASE [1174]</a>	Case operator
<a href="#">CAST ( ) [1245]</a>	Cast a value as a certain type
<a href="#">CEIL ( ) [1203]</a>	Return the smallest integer value not less than the argument
<a href="#">CEILING ( ) [1203]</a>	Return the smallest integer value not less than the argument
<a href="#">CHAR_LENGTH ( ) [1178]</a>	Return number of characters in argument
<a href="#">CHAR ( ) [1177]</a>	Return the character for each integer passed
<a href="#">CHARACTER_LENGTH ( ) [1178]</a>	A synonym for CHAR_LENGTH()
<a href="#">CHARSET ( ) [1266]</a>	Return the character set of the argument
<a href="#">COALESCE ( ) [1169]</a>	Return the first non-NULL argument
<a href="#">COERCIBILITY ( ) [1266]</a>	Return the collation coercibility value of the string argument
<a href="#">COLLATION ( ) [1267]</a>	Return the collation of the string argument
<a href="#">COMPRESS ( ) [1261]</a>	Return result as a binary string
<a href="#">CONCAT_WS ( ) [1179]</a>	Return concatenate with separator
<a href="#">CONCAT ( ) [1178]</a>	Return concatenated string
<a href="#">CONNECTION_ID ( ) [1267]</a>	Return the connection ID (thread ID) for the connection
<a href="#">CONV ( ) [1203]</a>	Convert numbers between different number bases
<a href="#">CONVERT_TZ ( ) [1212]</a>	Convert from one timezone to another
<a href="#">CONVERT ( ) [1245]</a>	Cast a value as a certain type
<a href="#">COS ( ) [1203]</a>	Return the cosine
<a href="#">COT ( ) [1203]</a>	Return the cotangent
<a href="#">COUNT(DISTINCT) [1283]</a>	Return the count of a number of different values
<a href="#">COUNT ( ) [1283]</a>	Return a count of the number of rows returned
<a href="#">CRC32 ( ) [1203]</a>	Compute a cyclic redundancy check value
<a href="#">CURDATE ( ) [1212]</a>	Return the current date
<a href="#">CURRENT_DATE ( ), CURRENT_DATE [1212]</a>	Synonyms for CURDATE()

Name	Description
<code>CURRENT_TIME()</code> , <code>CURRENT_TIME</code> [1212]	Synonyms for <code>CURTIME()</code>
<code>CURRENT_TIMESTAMP()</code> , <code>CURRENT_TIMESTAMP</code> [1213]	Synonyms for <code>NOW()</code>
<code>CURRENT_USER()</code> , <code>CURRENT_USER</code> [1267]	The authenticated user name and host name
<code>CURTIME()</code> [1213]	Return the current time
<code>DATABASE()</code> [1268]	Return the default (current) database name
<code>DATE_ADD()</code> [1213]	Add time values (intervals) to a date value
<code>DATE_FORMAT()</code> [1215]	Format date as specified
<code>DATE_SUB()</code> [1217]	Subtract a time value (interval) from a date
<code>DATE()</code> [1213]	Extract the date part of a date or datetime expression
<code>DATEDIFF()</code> [1213]	Subtract two dates
<code>DAY()</code> [1217]	Synonym for <code>DAYOFMONTH()</code>
<code>DAYNAME()</code> [1217]	Return the name of the weekday
<code>DAYOFMONTH()</code> [1217]	Return the day of the month (0-31)
<code>DAYOFWEEK()</code> [1217]	Return the weekday index of the argument
<code>DAYOFYEAR()</code> [1217]	Return the day of the year (1-366)
<code>DECODE()</code> [1261]	Decodes a string encrypted using <code>ENCODE()</code>
<code>DEFAULT()</code> [1275]	Return the default value for a table column
<code>DEGREES()</code> [1204]	Convert radians to degrees
<code>DES_DECRYPT()</code> [1261]	Decrypt a string
<code>DES_ENCRYPT()</code> [1262]	Encrypt a string
<code>DIV</code> [1201]	Integer division
<code>/</code> [1200]	Division operator
<code>ELT()</code> [1179]	Return string at index number
<code>ENCODE()</code> [1262]	Encode a string
<code>ENCRYPT()</code> [1262]	Encrypt a string
<code>&lt;=&gt;</code> [1167]	NULL-safe equal to operator
<code>=</code> [1166]	Equal operator
<code>EXP()</code> [1204]	Raise to the power of
<code>EXPORT_SET()</code> [1179]	Return a string such that for every bit set in the value bits, you get an on string and for every unset bit, you get an off string
<code>EXTRACT()</code> [1217]	Extract part of a date
<code>ExtractValue()</code> [1249]	Extracts a value from an XML string using XPath notation
<code>FIELD()</code> [1179]	Return the index (position) of the first argument in the subsequent arguments
<code>FIND_IN_SET()</code> [1179]	Return the index position of the first argument within the second argument
<code>FLOOR()</code> [1204]	Return the largest integer value not greater than the argument

Name	Description
<a href="#">FORMAT ( ) [1180]</a>	Return a number formatted to specified number of decimal places
<a href="#">FOUND_ROWS ( ) [1268]</a>	For a SELECT with a LIMIT clause, the number of rows that would be returned were there no LIMIT clause
<a href="#">FROM_BASE64 ( ) [1180]</a>	Decode to a base-64 string and return result
<a href="#">FROM_DAYS ( ) [1218]</a>	Convert a day number to a date
<a href="#">FROM_UNIXTIME ( ) [1218]</a>	Format UNIX timestamp as a date
<a href="#">GET_FORMAT ( ) [1218]</a>	Return a date format string
<a href="#">GET_LOCK ( ) [1275]</a>	Get a named lock
<a href="#">&gt;= [1167]</a>	Greater than or equal operator
<a href="#">&gt; [1167]</a>	Greater than operator
<a href="#">GREATEST ( ) [1169]</a>	Return the largest argument
<a href="#">GROUP_CONCAT ( ) [1283]</a>	Return a concatenated string
<a href="#">GTID_SUBSET ( ) [1274]</a>	Return true if all GTIDs in subset are also in set; otherwise false.
<a href="#">GTID_SUBTRACT ( ) [1274]</a>	Return all GTIDs in set that are not in subset.
<a href="#">HEX ( ) [1180]</a>	Return a hexadecimal representation of a decimal or string value
<a href="#">HOUR ( ) [1219]</a>	Extract the hour
<a href="#">IF ( ) [1174]</a>	If/else construct
<a href="#">IFNULL ( ) [1175]</a>	Null if/else construct
<a href="#">IN ( ) [1169]</a>	Check whether a value is within a set of values
<a href="#">INET_ATON ( ) [1276]</a>	Return the numeric value of an IP address
<a href="#">INET_NTOA ( ) [1276]</a>	Return the IP address from a numeric value
<a href="#">INET6_ATON ( ) [1276]</a>	Return the numeric value of an IPv6 address
<a href="#">INET6_NTOA ( ) [1277]</a>	Return the IPv6 address from a numeric value
<a href="#">INSERT ( ) [1180]</a>	Insert a substring at the specified position up to the specified number of characters
<a href="#">INSTR ( ) [1181]</a>	Return the index of the first occurrence of substring
<a href="#">INTERVAL ( ) [1170]</a>	Return the index of the argument that is less than the first argument
<a href="#">IS_FREE_LOCK ( ) [1278]</a>	Checks whether the named lock is free
<a href="#">IS_IPV4_COMPAT ( ) [1278]</a>	Return true if argument is an IPv4-compatible address
<a href="#">IS_IPV4_MAPPED ( ) [1278]</a>	Return true if argument is an IPv4-mapped address
<a href="#">IS_IPV4 ( ) [1278]</a>	Return true if argument is an IPv4 address
<a href="#">IS_IPV6 ( ) [1279]</a>	Return true if argument is an IPv6 address
<a href="#">IS NOT NULL [1168]</a>	NOT NULL value test
<a href="#">IS NOT [1168]</a>	Test a value against a boolean
<a href="#">IS NULL [1168]</a>	NULL value test
<a href="#">IS_USED_LOCK ( ) [1279]</a>	Checks whether the named lock is in use. Return connection identifier if true.
<a href="#">IS [1167]</a>	Test a value against a boolean
<a href="#">ISNULL ( ) [1170]</a>	Test whether the argument is NULL

Name	Description
LAST_DAY [1219]	Return the last day of the month for the argument
LAST_INSERT_ID() [1269]	Value of the AUTOINCREMENT column for the last INSERT
LCASE() [1181]	Synonym for LOWER()
LEAST() [1170]	Return the smallest argument
<< [1258]	Left shift
LEFT() [1181]	Return the leftmost number of characters as specified
LENGTH() [1181]	Return the length of a string in bytes
<= [1167]	Less than or equal operator
< [1167]	Less than operator
LIKE [1190]	Simple pattern matching
LN() [1204]	Return the natural logarithm of the argument
LOAD_FILE() [1181]	Load the named file
LOCALTIME(), LOCALTIME [1219]	Synonym for NOW()
LOCALTIMESTAMP, LOCALTIMESTAMP() [1219]	Synonym for NOW()
LOCATE() [1181]	Return the position of the first occurrence of substring
LOG10() [1205]	Return the base-10 logarithm of the argument
LOG2() [1205]	Return the base-2 logarithm of the argument
LOG() [1204]	Return the natural logarithm of the first argument
LOWER() [1182]	Return the argument in lowercase
LPAD() [1182]	Return the string argument, left-padded with the specified string
LTRIM() [1182]	Remove leading spaces
MAKE_SET() [1182]	Return a set of comma-separated strings that have the corresponding bit in bits set
MAKEDATE() [1219]	Create a date from the year and day of year
MAKETIME [1219]	MAKETIME()
MASTER_POS_WAIT() [1279]	Block until the slave has read and applied all updates up to the specified position
MATCH [1229]	Perform full-text search
MAX() [1284]	Return the maximum value
MD5() [1263]	Calculate MD5 checksum
MICROSECOND() [1220]	Return the microseconds from argument
MID() [1183]	Return a substring starting from the specified position
MIN() [1284]	Return the minimum value
- [1200]	Minus operator
MINUTE() [1220]	Return the minute from the argument
MOD() [1205]	Return the remainder
% or MOD [1201]	Modulo operator
MONTH() [1220]	Return the month from the date passed
MONTHNAME() [1220]	Return the name of the month
NAME_CONST() [1279]	Causes the column to have the given name

Name	Description
NOT BETWEEN ... AND ... [1169]	Check whether a value is not within a range of values
!=, <> [1167]	Not equal operator
NOT IN() [1170]	Check whether a value is not within a set of values
NOT LIKE [1192]	Negation of simple pattern matching
NOT REGEXP [1193]	Negation of REGEXP
NOT, ! [1171]	Negates value
NOW() [1220]	Return the current date and time
NULLIF() [1175]	Return NULL if expr1 = expr2
OCT() [1183]	Return a string containing octal representation of a number
OCTET_LENGTH() [1183]	A synonym for LENGTH()
OLD_PASSWORD() [1263]	Return the value of the pre-4.1 implementation of PASSWORD
, OR [1172]	Logical OR
ORD() [1183]	Return character code for leftmost character of the argument
PASSWORD() [1263]	Calculate and return a password string
PERIOD_ADD() [1221]	Add a period to a year-month
PERIOD_DIFF() [1221]	Return the number of months between periods
PI() [1205]	Return the value of pi
+ [1200]	Addition operator
POSITION() [1183]	A synonym for LOCATE()
POW() [1206]	Return the argument raised to the specified power
POWER() [1206]	Return the argument raised to the specified power
PROCEDURE ANALYSE()	Analyze the results of a query
QUARTER() [1221]	Return the quarter from a date argument
QUOTE() [1183]	Escape the argument for use in an SQL statement
RADIANS() [1206]	Return argument converted to radians
RAND() [1206]	Return a random floating-point value
REGEXP [1193]	Pattern matching using regular expressions
RELEASE_LOCK() [1280]	Releases the named lock
REPEAT() [1183]	Repeat a string the specified number of times
REPLACE() [1183]	Replace occurrences of a specified string
REVERSE() [1184]	Reverse the characters in a string
>> [1258]	Right shift
RIGHT() [1184]	Return the specified rightmost number of characters
RLIKE [1193]	Synonym for REGEXP
ROUND() [1207]	Round the argument
ROW_COUNT() [1272]	The number of rows updated
RPAD() [1184]	Append string the specified number of times
RTRIM() [1184]	Remove trailing spaces
SCHEMA() [1273]	A synonym for DATABASE()
SEC_TO_TIME() [1221]	Converts seconds to 'HH:MM:SS' format

Name	Description
SECOND ( ) [1221]	Return the second (0-59)
SESSION_USER ( ) [1273]	Synonym for USER()
SHA1 ( ), SHA ( ) [1264]	Calculate an SHA-1 160-bit checksum
SHA2 ( ) [1264]	Calculate an SHA-2 checksum
SIGN ( ) [1208]	Return the sign of the argument
SIN ( ) [1208]	Return the sine of the argument
SLEEP ( ) [1280]	Sleep for a number of seconds
SOUNDEX ( ) [1184]	Return a soundex string
SOUNDS LIKE [1185]	Compare sounds
SPACE ( ) [1185]	Return a string of the specified number of spaces
SQL_THREAD_WAIT_AFTER_GTIDS ( ) [1221] (ObsOLETE Replaced by 5.6.9)	WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS()
SQRT ( ) [1208]	Return the square root of the argument
STD ( ) [1284]	Return the population standard deviation
STDDEV_POP ( ) [1285]	Return the population standard deviation
STDDEV_SAMP ( ) [1285]	Return the sample standard deviation
STDDEV ( ) [1285]	Return the population standard deviation
STR_TO_DATE ( ) [1221]	Convert a string to a date
STRCMP ( ) [1192]	Compare two strings
SUBDATE ( ) [1222]	A synonym for DATE_SUB() when invoked with three arguments
SUBSTR ( ) [1185]	Return the substring as specified
SUBSTRING_INDEX ( ) [1185]	Return a substring from a string before the specified number of occurrences of the delimiter
SUBSTRING ( ) [1185]	Return the substring as specified
SUBTIME ( ) [1223]	Subtract times
SUM ( ) [1285]	Return the sum
SYSDATE ( ) [1223]	Return the time at which the function executes
SYSTEM_USER ( ) [1273]	Synonym for USER()
TAN ( ) [1208]	Return the tangent of the argument
TIME_FORMAT ( ) [1224]	Format as time
TIME_TO_SEC ( ) [1225]	Return the argument converted to seconds
TIME ( ) [1223]	Extract the time portion of the expression passed
TIMEDIFF ( ) [1223]	Subtract time
* [1200]	Multiplication operator
TIMESTAMP ( ) [1224]	With a single argument, this function returns the date or datetime expression; with two arguments, the sum of the arguments
TIMESTAMPADD ( ) [1224]	Add an interval to a datetime expression
TIMESTAMPDIFF ( ) [1224]	Subtract an interval from a datetime expression
TO_BASE64 ( ) [1186]	Return the argument converted to a base-64 string
TO_DAYS ( ) [1225]	Return the date argument converted to days



Name	Description
<a href="#">TO_SECONDS ( ) [1226]</a>	Return the date or datetime argument converted to seconds since Year 0
<a href="#">TRIM ( ) [1186]</a>	Remove leading and trailing spaces
<a href="#">TRUNCATE ( ) [1209]</a>	Truncate to specified number of decimal places
<a href="#">UCASE ( ) [1186]</a>	Synonym for UPPER()
<a href="#">- [1200]</a>	Change the sign of the argument
<a href="#">UNCOMPRESS ( ) [1265]</a>	Uncompress a string compressed
<a href="#">UNCOMPRESSED_LENGTH ( ) [1265]</a>	Return the length of a string before compression
<a href="#">UNHEX ( ) [1186]</a>	Return a string containing hex representation of a number
<a href="#">UNIX_TIMESTAMP ( ) [1226]</a>	Return a UNIX timestamp
<a href="#">UpdateXML ( ) [1251]</a>	Return replaced XML fragment
<a href="#">UPPER ( ) [1187]</a>	Convert to uppercase
<a href="#">USER ( ) [1273]</a>	The user name and host name provided by the client
<a href="#">UTC_DATE ( ) [1227]</a>	Return the current UTC date
<a href="#">UTC_TIME ( ) [1227]</a>	Return the current UTC time
<a href="#">UTC_TIMESTAMP ( ) [1227]</a>	Return the current UTC date and time
<a href="#">UUID_SHORT ( ) [1281]</a>	Return an integer-valued universal identifier
<a href="#">UUID ( ) [1280]</a>	Return a Universal Unique Identifier (UUID)
<a href="#">VALIDATE_PASSWORD_STRENGTH ( ) [1205]</a>	Determine strength of password
<a href="#">VALUES ( ) [1281]</a>	Defines the values to be used during an INSERT
<a href="#">VAR_POP ( ) [1285]</a>	Return the population standard variance
<a href="#">VAR_SAMP ( ) [1285]</a>	Return the sample variance
<a href="#">VARIANCE ( ) [1285]</a>	Return the population standard variance
<a href="#">VERSION ( ) [1273]</a>	Returns a string that indicates the MySQL server version
<a href="#">WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS ( ) [1274]</a>	Wait until the slave SQL thread has executed all the given GTIDs. Returns: the number of events that were executed (or NULL, if GTID mode is not enabled).
<a href="#">WEEK ( ) [1228]</a>	Return the week number
<a href="#">WEEKDAY ( ) [1228]</a>	Return the weekday index
<a href="#">WEEKOFYEAR ( ) [1229]</a>	Return the calendar week of the date (0-53)
<a href="#">WEIGHT_STRING ( ) [1187]</a>	Return the weight string for a string
<a href="#">XOR [1172]</a>	Logical XOR
<a href="#">YEAR ( ) [1229]</a>	Return the year
<a href="#">YEARWEEK ( ) [1229]</a>	Return the year and week

## 12.2. Type Conversion in Expression Evaluation

When an operator is used with operands of different types, type conversion occurs to make the operands compatible. Some conversions occur implicitly. For example, MySQL automatically converts numbers to strings as necessary, and vice versa.

```
mysql> SELECT 1+'1';
-> 2
mysql> SELECT CONCAT(2,' test');
-> '2 test'
```

It is also possible to convert a number to a string explicitly using the `CAST()` [1245] function. Conversion occurs implicitly with the `CONCAT()` [1178] function because it expects string arguments.

```
mysql> SELECT 38.8, CAST(38.8 AS CHAR);
-> 38.8, '38.8'
mysql> SELECT 38.8, CONCAT(38.8);
-> 38.8, '38.8'
```

See later in this section for information about the character set of implicit number-to-string conversions.

The following rules describe how conversion occurs for comparison operations:

- If one or both arguments are `NULL`, the result of the comparison is `NULL`, except for the `NULL`-safe `<=>` [1167] equality comparison operator. For `NULL <=> NULL`, the result is true. No conversion is needed.
- If both arguments in a comparison operation are strings, they are compared as strings.
- If both arguments are integers, they are compared as integers.
- Hexadecimal values are treated as binary strings if not compared to a number.
- If one of the arguments is a `TIMESTAMP` or `DATETIME` column and the other argument is a constant, the constant is converted to a timestamp before the comparison is performed. This is done to be more ODBC-friendly. Note that this is not done for the arguments to `IN()` [1169]! To be safe, always use complete datetime, date, or time strings when doing comparisons. For example, to achieve best results when using `BETWEEN` [1168] with date or time values, use `CAST()` [1245] to explicitly convert the values to the desired data type.

A single-row subquery from a table or tables is not considered a constant. For example, if a subquery returns an integer to be compared to a `DATETIME` value, the comparison is done as two integers. The integer is not converted to a temporal value. To compare the operands as `DATETIME` values, use `CAST()` [1245] to explicitly convert the subquery value to `DATETIME`.

- If one of the arguments is a decimal value, comparison depends on the other argument. The arguments are compared as decimal values if the other argument is a decimal or integer value, or as floating-point values if the other argument is a floating-point value.
- In all other cases, the arguments are compared as floating-point (real) numbers.

For information about conversion of values from one temporal type to another, see [Section 11.3.7, “Conversion Between Date and Time Types”](#).

The following examples illustrate conversion of strings to numbers for comparison operations:

```
mysql> SELECT 1 > '6x';
-> 0
mysql> SELECT 7 > '6x';
-> 1
mysql> SELECT 0 > 'x6';
-> 0
mysql> SELECT 0 = 'x6';
-> 1
```

For comparisons of a string column with a number, MySQL cannot use an index on the column to look up the value quickly. If `str_col` is an indexed string column, the index cannot be used when performing the lookup in the following statement:

```
SELECT * FROM tbl_name WHERE str_col=1;
```

The reason for this is that there are many different strings that may convert to the value `1`, such as `'1'`, `' 1'`, or `'1a'`.

Comparisons that use floating-point numbers (or values that are converted to floating-point numbers) are approximate because such numbers are inexact. This might lead to results that appear inconsistent:

```
mysql> SELECT '18015376320243458' = 18015376320243458;
-> 1
mysql> SELECT '18015376320243459' = 18015376320243459;
-> 0
```

Such results can occur because the values are converted to floating-point numbers, which have only 53 bits of precision and are subject to rounding:

```
mysql> SELECT '18015376320243459'+0.0;
-> 1.8015376320243e+16
```

Furthermore, the conversion from string to floating-point and from integer to floating-point do not necessarily occur the same way. The integer may be converted to floating-point by the CPU, whereas the string is converted digit by digit in an operation that involves floating-point multiplications.

The results shown will vary on different systems, and can be affected by factors such as computer architecture or the compiler version or optimization level. One way to avoid such problems is to use `CAST()` [1245] so that a value will not be converted implicitly to a float-point number:

```
mysql> SELECT CAST('18015376320243459' AS UNSIGNED) = 18015376320243459;
-> 1
```

For more information about floating-point comparisons, see [Section C.5.5.8, “Problems with Floating-Point Values”](#).

In MySQL 5.6, the server includes `dtoa`, a conversion library that provides the basis for improved conversion between string or `DECIMAL` values and approximate-value (`FLOAT/DOUBLE`) numbers:

- Consistent conversion results across platforms, which eliminates, for example, Unix versus Windows conversion differences.
- Accurate representation of values in cases where results previously did not provide sufficient precision, such as for values close to IEEE limits.
- Conversion of numbers to string format with the best possible precision. The precision of `dtoa` is always the same or better than that of the standard C library functions.

Because the conversions produced by this library differ in some cases from previous results, the potential exists for incompatibilities in applications that rely on previous results. For example, applications that depend on a specific exact result from previous conversions might need adjustment to accommodate additional precision.

The `dtoa` library provides conversions with the following properties. `D` represents a value with a `DECIMAL` or string representation, and `F` represents a floating-point number in native binary (IEEE) format.

- `F -> D` conversion is done with the best possible precision, returning `D` as the shortest string that yields `F` when read back in and rounded to the nearest value in native binary format as specified by IEEE.
- `D -> F` conversion is done such that `F` is the nearest native binary number to the input decimal string `D`.

These properties imply that `F -> D -> F` conversions are lossless unless `F` is `-inf`, `+inf`, or `NaN`. The latter values are not supported because the SQL standard defines them as invalid values for `FLOAT` or `DOUBLE`.

For `D -> F -> D` conversions, a sufficient condition for losslessness is that `D` uses 15 or fewer digits of precision, is not a denormal value, `-inf`, `+inf`, or `NaN`. In some cases, the conversion is lossless even if `D` has more than 15 digits of precision, but this is not always the case.

In MySQL 5.6, implicit conversion of a numeric or temporal value to string produces a value that has a character set and collation determined by the `character_set_connection` [467] and

`collation_connection` [469] system variables. (These variables commonly are set with `SET NAMES`. For information about connection character sets, see [Section 10.1.4, “Connection Character Sets and Collations”](#).)

This means that such a conversion results in a character (nonbinary) string (a `CHAR`, `VARCHAR`, or `LONGTEXT` value), except in the case that the connection character set is set to `binary`. In that case, the conversion result is a binary string (a `BINARY`, `VARBINARY`, or `LONGBLOB` value).

## 12.3. Operators

**Table 12.2. Operators**

Name	Description
<code>AND</code> , <code>&amp;&amp;</code> [1171]	Logical AND
<code>=</code> [1173]	Assign a value (as part of a <code>SET</code> statement, or as part of the <code>SET</code> clause in an <code>UPDATE</code> statement)
<code>:=</code> [1172]	Assign a value
<code>BETWEEN ... AND ...</code> [1168]	Check whether a value is within a range of values
<code>BINARY</code> [1244]	Cast a string to a binary string
<code>&amp;</code> [1258]	Bitwise AND
<code>~</code> [1259]	Invert bits
<code> </code> [1258]	Bitwise OR
<code>^</code> [1258]	Bitwise XOR
<code>CASE</code> [1174]	Case operator
<code>DIV</code> [1201]	Integer division
<code>/</code> [1200]	Division operator
<code>&lt;=&gt;</code> [1167]	NULL-safe equal to operator
<code>=</code> [1166]	Equal operator
<code>&gt;=</code> [1167]	Greater than or equal operator
<code>&gt;</code> [1167]	Greater than operator
<code>IS NOT NULL</code> [1168]	NOT NULL value test
<code>IS NOT</code> [1168]	Test a value against a boolean
<code>IS NULL</code> [1168]	NULL value test
<code>IS</code> [1167]	Test a value against a boolean
<code>&lt;&lt;</code> [1258]	Left shift
<code>&lt;=</code> [1167]	Less than or equal operator
<code>&lt;</code> [1167]	Less than operator
<code>LIKE</code> [1190]	Simple pattern matching
<code>-</code> [1200]	Minus operator
<code>%</code> or <code>MOD</code> [1201]	Modulo operator
<code>NOT BETWEEN ... AND ...</code> [1169]	Check whether a value is not within a range of values
<code>!=</code> , <code>&lt;&gt;</code> [1167]	Not equal operator
<code>NOT LIKE</code> [1192]	Negation of simple pattern matching
<code>NOT REGEXP</code> [1193]	Negation of REGEXP
<code>NOT</code> , <code>!</code> [1171]	Negates value
<code>  </code> , <code>OR</code> [1172]	Logical OR

Name	Description
+ [1200]	Addition operator
REGEXP [1193]	Pattern matching using regular expressions
>> [1258]	Right shift
RLIKE [1193]	Synonym for REGEXP
SOUNDS LIKE [1185]	Compare sounds
* [1200]	Multiplication operator
- [1200]	Change the sign of the argument
XOR [1172]	Logical XOR

### 12.3.1. Operator Precedence

Operator precedences are shown in the following list, from highest precedence to the lowest. Operators that are shown together on a line have the same precedence.

```

INTERVAL
BINARY, COLLATE
!
- (unary minus), ~ (unary bit inversion)
^
*, /, DIV, %, MOD
-, +
<<, >>
&
|
= (comparison), <=>, >=, >, <=, <, <>, !=, IS, LIKE, REGEXP, IN
BETWEEN, CASE, WHEN, THEN, ELSE
NOT
&&, AND
XOR
||, OR
= (assignment), :=

```

The precedence of = depends on whether it is used as a comparison operator (= [1166]) or as an assignment operator (= [1173]). When used as a comparison operator, it has the same precedence as <=> [1167], >= [1167], > [1167], <= [1167], < [1167], <> [1167], != [1167], IS [1167], LIKE [1190], REGEXP [1193], and IN [1169]. When used as an assignment operator, it has the same precedence as := [1172]. Section 13.7.4, “SET Syntax”, and Section 9.4, “User-Defined Variables”, explain how MySQL determines which interpretation of = should apply.

The meaning of some operators depends on the SQL mode:

- By default, || [1172] is a logical OR [1172] operator. With PIPES\_AS\_CONCAT [607] enabled, || [1172] is string concatenation, with a precedence between ^ [1258] and the unary operators.
- By default, ! [1171] has a higher precedence than NOT. With HIGH\_NOT\_PRECEDENCE [603] enabled, ! [1171] and NOT have the same precedence.

See Section 5.1.7, “Server SQL Modes”.

The precedence of operators determines the order of evaluation of terms in an expression. To override this order and group terms explicitly, use parentheses. For example:

```

mysql> SELECT 1+2*3;
-> 7
mysql> SELECT (1+2)*3;
-> 9

```

### 12.3.2. Comparison Functions and Operators

Table 12.3. Comparison Operators

Name	Description
<code>BETWEEN ... AND ...</code> [1168]	Check whether a value is within a range of values
<code>COALESCE()</code> [1169]	Return the first non-NULL argument
<code>&lt;=&gt;</code> [1167]	NULL-safe equal to operator
<code>=</code> [1166]	Equal operator
<code>&gt;=</code> [1167]	Greater than or equal operator
<code>&gt;</code> [1167]	Greater than operator
<code>GREATEST()</code> [1169]	Return the largest argument
<code>IN()</code> [1169]	Check whether a value is within a set of values
<code>INTERVAL()</code> [1170]	Return the index of the argument that is less than the first argument
<code>IS NOT NULL</code> [1168]	NOT NULL value test
<code>IS NOT</code> [1168]	Test a value against a boolean
<code>IS NULL</code> [1168]	NULL value test
<code>IS</code> [1167]	Test a value against a boolean
<code>ISNULL()</code> [1170]	Test whether the argument is NULL
<code>LEAST()</code> [1170]	Return the smallest argument
<code>&lt;=</code> [1167]	Less than or equal operator
<code>&lt;</code> [1167]	Less than operator
<code>LIKE</code> [1190]	Simple pattern matching
<code>NOT BETWEEN ... AND ...</code> [1169]	Check whether a value is not within a range of values
<code>!=, &lt;&gt;</code> [1167]	Not equal operator
<code>NOT IN()</code> [1170]	Check whether a value is not within a set of values
<code>NOT LIKE</code> [1192]	Negation of simple pattern matching
<code>STRCMP()</code> [1192]	Compare two strings

Comparison operations result in a value of 1 (`TRUE`), 0 (`FALSE`), or `NULL`. These operations work for both numbers and strings. Strings are automatically converted to numbers and numbers to strings as necessary.

The following relational comparison operators can be used to compare not only scalar operands, but row operands:

```
= > < >= <= <> !=
```

For examples of row comparisons, see [Section 13.2.10.5, “Row Subqueries”](#).

Some of the functions in this section return values other than 1 (`TRUE`), 0 (`FALSE`), or `NULL`. For example, `LEAST()` [1170] and `GREATEST()` [1169]. However, the value they return is based on comparison operations performed according to the rules described in [Section 12.2, “Type Conversion in Expression Evaluation”](#).

To convert a value to a specific type for comparison purposes, you can use the `CAST()` [1245] function. String values can be converted to a different character set using `CONVERT()` [1245]. See [Section 12.10, “Cast Functions and Operators”](#).

By default, string comparisons are not case sensitive and use the current character set. The default is `latin1` (cp1252 West European), which also works well for English.

- `=` [1166]

Equal:

```
mysql> SELECT 1 = 0;
-> 0
mysql> SELECT '0' = 0;
-> 1
mysql> SELECT '0.0' = 0;
-> 1
mysql> SELECT '0.01' = 0;
-> 0
mysql> SELECT '.01' = 0.01;
-> 1
```

- [<=> \[1167\]](#)

**NULL-safe equal.** This operator performs an equality comparison like the = [\[1166\]](#) operator, but returns 1 rather than NULL if both operands are NULL, and 0 rather than NULL if one operand is NULL.

```
mysql> SELECT 1 <=> 1, NULL <=> NULL, 1 <=> NULL;
-> 1, 1, 0
mysql> SELECT 1 = 1, NULL = NULL, 1 = NULL;
-> 1, NULL, NULL
```

- [<> \[1167\]](#), [!= \[1167\]](#)

Not equal:

```
mysql> SELECT '.01' <> '0.01';
-> 1
mysql> SELECT .01 <> '0.01';
-> 0
mysql> SELECT 'zapp' <> 'zapp';
-> 1
```

- [<= \[1167\]](#)

Less than or equal:

```
mysql> SELECT 0.1 <= 2;
-> 1
```

- [< \[1167\]](#)

Less than:

```
mysql> SELECT 2 < 2;
-> 0
```

- [>= \[1167\]](#)

Greater than or equal:

```
mysql> SELECT 2 >= 2;
-> 1
```

- [> \[1167\]](#)

Greater than:

```
mysql> SELECT 2 > 2;
-> 0
```

- [IS boolean\\_value \[1167\]](#)

Tests a value against a boolean value, where *boolean\_value* can be TRUE, FALSE, or UNKNOWN.

```
mysql> SELECT 1 IS TRUE, 0 IS FALSE, NULL IS UNKNOWN;
```



```
-> 1, 1, 1
```

- `IS NOT boolean_value` [1168]

Tests a value against a boolean value, where *boolean\_value* can be `TRUE`, `FALSE`, or `UNKNOWN`.

```
mysql> SELECT 1 IS NOT UNKNOWN, 0 IS NOT UNKNOWN, NULL IS NOT UNKNOWN;
-> 1, 1, 0
```

- `IS NULL` [1168]

Tests whether a value is `NULL`.

```
mysql> SELECT 1 IS NULL, 0 IS NULL, NULL IS NULL;
-> 0, 0, 1
```

To work well with ODBC programs, MySQL supports the following extra features when using `IS NULL` [1168]:

- If `sql_auto_is_null` [543] variable is set to 1, then after a statement that successfully inserts an automatically generated `AUTO_INCREMENT` value, you can find that value by issuing a statement of the following form:

```
SELECT * FROM tbl_name WHERE auto_col IS NULL
```

If the statement returns a row, the value returned is the same as if you invoked the `LAST_INSERT_ID()` [1269] function. For details, including the return value after a multiple-row insert, see [Section 12.14, “Information Functions”](#). If no `AUTO_INCREMENT` value was successfully inserted, the `SELECT` statement returns no row.

The behavior of retrieving an `AUTO_INCREMENT` value by using an `IS NULL` [1168] comparison can be disabled by setting `sql_auto_is_null = 0` [543]. See [Section 5.1.4, “Server System Variables”](#).

The default value of `sql_auto_is_null` [543] is 0 in MySQL 5.6.

- For `DATE` and `DATETIME` columns that are declared as `NOT NULL`, you can find the special date `'0000-00-00'` by using a statement like this:

```
SELECT * FROM tbl_name WHERE date_column IS NULL
```

This is needed to get some ODBC applications to work because ODBC does not support a `'0000-00-00'` date value.

See [Section 21.1.7.1.1, “Obtaining Auto-Increment Values”](#), and the description for the `FLAG_AUTO_IS_NULL` option at [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#).

- `IS NOT NULL` [1168]

Tests whether a value is not `NULL`.

```
mysql> SELECT 1 IS NOT NULL, 0 IS NOT NULL, NULL IS NOT NULL;
-> 1, 1, 0
```

- `expr BETWEEN min AND max` [1168]

If *expr* is greater than or equal to *min* and *expr* is less than or equal to *max*, `BETWEEN` [1168] returns 1, otherwise it returns 0. This is equivalent to the expression `(min <= expr AND expr <= max)` if all the arguments are of the same type. Otherwise type conversion takes place according to the rules described in [Section 12.2, “Type Conversion in Expression Evaluation”](#), but applied to all the three arguments.

```
mysql> SELECT 2 BETWEEN 1 AND 3, 2 BETWEEN 3 AND 1;
-> 1, 0
```

```
mysql> SELECT 1 BETWEEN 2 AND 3;
-> 0
mysql> SELECT 'b' BETWEEN 'a' AND 'c';
-> 1
mysql> SELECT 2 BETWEEN 2 AND '3';
-> 1
mysql> SELECT 2 BETWEEN 2 AND 'x-3';
-> 0
```

For best results when using `BETWEEN` [1168] with date or time values, use `CAST()` [1245] to explicitly convert the values to the desired data type. Examples: If you compare a `DATETIME` to two `DATE` values, convert the `DATE` values to `DATETIME` values. If you use a string constant such as `'2001-1-1'` in a comparison to a `DATE`, cast the string to a `DATE`.

- `expr NOT BETWEEN min AND max` [1169]

This is the same as `NOT (expr BETWEEN min AND max)`.

- `COALESCE(value,...)` [1169]

Returns the first non-`NULL` value in the list, or `NULL` if there are no non-`NULL` values.

```
mysql> SELECT COALESCE(NULL,1);
-> 1
mysql> SELECT COALESCE(NULL,NULL,NULL);
-> NULL
```

- `GREATEST(value1,value2,...)` [1169]

With two or more arguments, returns the largest (maximum-valued) argument. The arguments are compared using the same rules as for `LEAST()` [1170].

```
mysql> SELECT GREATEST(2,0);
-> 2
mysql> SELECT GREATEST(34.0,3.0,5.0,767.0);
-> 767.0
mysql> SELECT GREATEST('B','A','C');
-> 'C'
```

`GREATEST()` [1169] returns `NULL` if any argument is `NULL`.

- `expr IN (value,...)` [1169]

Returns `1` if `expr` is equal to any of the values in the `IN` list, else returns `0`. If all values are constants, they are evaluated according to the type of `expr` and sorted. The search for the item then is done using a binary search. This means `IN` is very quick if the `IN` value list consists entirely of constants. Otherwise, type conversion takes place according to the rules described in [Section 12.2, "Type Conversion in Expression Evaluation"](#), but applied to all the arguments.

```
mysql> SELECT 2 IN (0,3,5,7);
-> 0
mysql> SELECT 'wefwf' IN ('wee','wefwf','weg');
-> 1
```

You should never mix quoted and unquoted values in an `IN` list because the comparison rules for quoted values (such as strings) and unquoted values (such as numbers) differ. Mixing types may therefore lead to inconsistent results. For example, do not write an `IN` expression like this:

```
SELECT val1 FROM tbl1 WHERE val1 IN (1,2,'a');
```

Instead, write it like this:

```
SELECT val1 FROM tbl1 WHERE val1 IN ('1','2','a');
```

The number of values in the `IN` list is only limited by the `max_allowed_packet` [502] value.

To comply with the SQL standard, `IN` returns `NULL` not only if the expression on the left hand side is `NULL`, but also if no match is found in the list and one of the expressions in the list is `NULL`.

`IN()` syntax can also be used to write certain types of subqueries. See [Section 13.2.10.3, “Subqueries with ANY, IN, or SOME”](#).

- `expr NOT IN (value,...)` [1170]

This is the same as `NOT (expr IN (value,...))`.

- `ISNULL(expr)` [1170]

If `expr` is `NULL`, `ISNULL()` [1170] returns 1, otherwise it returns 0.

```
mysql> SELECT ISNULL(1+1);
-> 0
mysql> SELECT ISNULL(1/0);
-> 1
```

`ISNULL()` [1170] can be used instead of `=` [1166] to test whether a value is `NULL`. (Comparing a value to `NULL` using `=` [1166] always yields false.)

The `ISNULL()` [1170] function shares some special behaviors with the `IS NULL` [1168] comparison operator. See the description of `IS NULL` [1168].

- `INTERVAL(N,N1,N2,N3,...)` [1170]

Returns 0 if  $N < N1$ , 1 if  $N < N2$  and so on or -1 if  $N$  is `NULL`. All arguments are treated as integers. It is required that  $N1 < N2 < N3 < \dots < Nn$  for this function to work correctly. This is because a binary search is used (very fast).

```
mysql> SELECT INTERVAL(23, 1, 15, 17, 30, 44, 200);
-> 3
mysql> SELECT INTERVAL(10, 1, 10, 100, 1000);
-> 2
mysql> SELECT INTERVAL(22, 23, 30, 44, 200);
-> 0
```

- `LEAST(value1,value2,...)` [1170]

With two or more arguments, returns the smallest (minimum-valued) argument. The arguments are compared using the following rules:

- If any argument is `NULL`, the result is `NULL`. No comparison is needed.
- If the return value is used in an `INTEGER` context or all arguments are integer-valued, they are compared as integers.
- If the return value is used in a `REAL` context or all arguments are real-valued, they are compared as reals.
- If the arguments comprise a mix of numbers and strings, they are compared as numbers.
- If any argument is a nonbinary (character) string, the arguments are compared as nonbinary strings.
- In all other cases, the arguments are compared as binary strings.

```
mysql> SELECT LEAST(2,0);
-> 0
mysql> SELECT LEAST(34.0,3.0,5.0,767.0);
-> 3.0
mysql> SELECT LEAST('B','A','C');
-> 'A'
```

Note that the preceding conversion rules can produce strange results in some borderline cases:

```
mysql> SELECT CAST(LEAST(3600, 9223372036854775808.0) as SIGNED);
-> -9223372036854775808
```

This happens because MySQL reads 9223372036854775808.0 in an integer context. The integer representation is not good enough to hold the value, so it wraps to a signed integer.

### 12.3.3. Logical Operators

Table 12.4. Logical Operators

Name	Description
AND, && [1171]	Logical AND
NOT, ! [1171]	Negates value
, OR [1172]	Logical OR
XOR [1172]	Logical XOR

In SQL, all logical operators evaluate to `TRUE`, `FALSE`, or `NULL` (`UNKNOWN`). In MySQL, these are implemented as 1 (`TRUE`), 0 (`FALSE`), and `NULL`. Most of this is common to different SQL database servers, although some servers may return any nonzero value for `TRUE`.

MySQL evaluates any nonzero, non-`NULL` value to `TRUE`. For example, the following statements all assess to `TRUE`:

```
mysql> SELECT 10 IS TRUE;
-> 1
mysql> SELECT -10 IS TRUE;
-> 1
mysql> SELECT 'string' IS NOT NULL;
-> 1
```

- NOT [1171], ! [1171]

Logical NOT. Evaluates to 1 if the operand is 0, to 0 if the operand is nonzero, and `NOT NULL` returns `NULL`.

```
mysql> SELECT NOT 10;
-> 0
mysql> SELECT NOT 0;
-> 1
mysql> SELECT NOT NULL;
-> NULL
mysql> SELECT ! (1+1);
-> 0
mysql> SELECT ! 1+1;
-> 1
```

The last example produces 1 because the expression evaluates the same way as `(!1)+1`.

- AND [1171], && [1171]

Logical AND. Evaluates to 1 if all operands are nonzero and not `NULL`, to 0 if one or more operands are 0, otherwise `NULL` is returned.

```
mysql> SELECT 1 && 1;
-> 1
mysql> SELECT 1 && 0;
-> 0
mysql> SELECT 1 && NULL;
-> NULL
mysql> SELECT 0 && NULL;
-> 0
mysql> SELECT NULL && 0;
-> 0
```

- [OR \[1172\]](#), [|| \[1172\]](#)

Logical OR. When both operands are non-NULL, the result is 1 if any operand is nonzero, and 0 otherwise. With a NULL operand, the result is 1 if the other operand is nonzero, and NULL otherwise. If both operands are NULL, the result is NULL.

```
mysql> SELECT 1 || 1;
-> 1
mysql> SELECT 1 || 0;
-> 1
mysql> SELECT 0 || 0;
-> 0
mysql> SELECT 0 || NULL;
-> NULL
mysql> SELECT 1 || NULL;
-> 1
```

- [XOR \[1172\]](#)

Logical XOR. Returns NULL if either operand is NULL. For non-NULL operands, evaluates to 1 if an odd number of operands is nonzero, otherwise 0 is returned.

```
mysql> SELECT 1 XOR 1;
-> 0
mysql> SELECT 1 XOR 0;
-> 1
mysql> SELECT 1 XOR NULL;
-> NULL
mysql> SELECT 1 XOR 1 XOR 1;
-> 1
```

$a \text{ XOR } b$  is mathematically equal to  $(a \text{ AND } (\text{NOT } b)) \text{ OR } ((\text{NOT } a) \text{ and } b)$ .

## 12.3.4. Assignment Operators

Table 12.5. Assignment Operators

Name	Description
<a href="#">= [1173]</a>	Assign a value (as part of a SET statement, or as part of the SET clause in an UPDATE statement)
<a href="#">:= [1172]</a>	Assign a value

- [:= \[1172\]](#)

Assignment operator. Causes the user variable on the left hand side of the operator to take on the value to its right. The value on the right hand side may be a literal value, another variable storing a value, or any legal expression that yields a scalar value, including the result of a query (provided that this value is a scalar value). You can perform multiple assignments in the same SET statement. You can perform multiple assignments in the same statement-

Unlike [= \[1173\]](#), the [:= \[1172\]](#) operator is never interpreted as a comparison operator. This means you can use [:= \[1172\]](#) in any valid SQL statement (not just in SET statements) to assign a value to a variable.

```
mysql> SELECT @var1, @var2;
-> NULL, NULL
mysql> SELECT @var1 := 1, @var2;
-> 1, NULL
mysql> SELECT @var1, @var2;
-> 1, NULL
mysql> SELECT @var1, @var2 := @var1;
-> 1, 1
mysql> SELECT @var1, @var2;
-> 1, 1
```

```
mysql> SELECT @var1:=COUNT(*) FROM t1;
-> 4
mysql> SELECT @var1;
-> 4
```

You can make value assignments using `:=` [1172] in other statements besides `SELECT`, such as `UPDATE`, as shown here:

```
mysql> SELECT @var1;
-> 4
mysql> SELECT * FROM t1;
-> 1, 3, 5, 7

mysql> UPDATE t1 SET c1 = 2 WHERE c1 = @var1:= 1;
Query OK, 1 row affected (0.00 sec)
Rows matched: 1 Changed: 1 Warnings: 0

mysql> SELECT @var1;
-> 1
mysql> SELECT * FROM t1;
-> 2, 3, 5, 7
```

While it is also possible both to set and to read the value of the same variable in a single SQL statement using the `:=` [1172] operator, this is not recommended. Section 9.4, “User-Defined Variables”, explains why you should avoid doing this.

- `=` [1173]

This operator is used to perform value assignments in two cases, described in the next two paragraphs.

Within a `SET` statement, `=` is treated as an assignment operator that causes the user variable on the left hand side of the operator to take on the value to its right. (In other words, when used in a `SET` statement, `=` is treated identically to `:=` [1172].) The value on the right hand side may be a literal value, another variable storing a value, or any legal expression that yields a scalar value, including the result of a query (provided that this value is a scalar value). You can perform multiple assignments in the same `SET` statement.

In the `SET` clause of an `UPDATE` statement, `=` also acts as an assignment operator; in this case, however, it causes the column named on the left hand side of the operator to assume the value given to the right, provided any `WHERE` conditions that are part of the `UPDATE` are met. You can make multiple assignments in the same `SET` clause of an `UPDATE` statement.

In any other context, `=` is treated as a [comparison operator](#) [1166].

```
mysql> SELECT @var1, @var2;
-> NULL, NULL
mysql> SELECT @var1 := 1, @var2;
-> 1, NULL
mysql> SELECT @var1, @var2;
-> 1, NULL
mysql> SELECT @var1, @var2 := @var1;
-> 1, 1
mysql> SELECT @var1, @var2;
-> 1, 1
```

For more information, see [Section 13.7.4, “SET Syntax”](#), [Section 13.2.11, “UPDATE Syntax”](#), and [Section 13.2.10, “Subquery Syntax”](#).

## 12.4. Control Flow Functions

Table 12.6. Flow Control Operators

Name	Description
<code>CASE</code> [1174]	Case operator

Name	Description
<code>IF()</code> [1174]	If/else construct
<code>IFNULL()</code> [1175]	Null if/else construct
<code>NULLIF()</code> [1175]	Return NULL if <code>expr1 = expr2</code>

- `CASE value WHEN [compare_value] THEN result [WHEN [compare_value] THEN result ...] [ELSE result] END` [1174]

`CASE WHEN [condition] THEN result [WHEN [condition] THEN result ...] [ELSE result] END` [1174]

The first version returns the *result* where *value=compare\_value*. The second version returns the result for the first condition that is true. If there was no matching result value, the result after `ELSE` is returned, or `NULL` if there is no `ELSE` part.

```
mysql> SELECT CASE 1 WHEN 1 THEN 'one'
->      WHEN 2 THEN 'two' ELSE 'more' END;
-> 'one'
mysql> SELECT CASE WHEN 1>0 THEN 'true' ELSE 'false' END;
-> 'true'
mysql> SELECT CASE BINARY 'B'
->      WHEN 'a' THEN 1 WHEN 'b' THEN 2 END;
-> NULL
```

The return type of a `CASE` [1174] expression is the compatible aggregated type of all return values, but also depends on the context in which it is used. If used in a string context, the result is returned as a string. If used in a numeric context, the result is returned as a decimal, real, or integer value.



#### Note

The syntax of the `CASE` [1174] expression shown here differs slightly from that of the SQL `CASE statement` described in Section 13.6.5.1, “`CASE Syntax`”, for use inside stored programs. The `CASE` statement cannot have an `ELSE NULL` clause, and it is terminated with `END CASE` instead of `END`.

- `IF(expr1,expr2,expr3)` [1174]

If *expr1* is `TRUE` (`expr1 <> 0` and `expr1 <> NULL`) then `IF()` [1174] returns *expr2*; otherwise it returns *expr3*. `IF()` [1174] returns a numeric or string value, depending on the context in which it is used.

```
mysql> SELECT IF(1>2,2,3);
-> 3
mysql> SELECT IF(1<2,'yes','no');
-> 'yes'
mysql> SELECT IF(STRCMP('test','test1'),'no','yes');
-> 'no'
```

If only one of *expr2* or *expr3* is explicitly `NULL`, the result type of the `IF()` [1174] function is the type of the non-`NULL` expression.

The default return type of `IF()` [1174] (which may matter when it is stored into a temporary table) is calculated as follows.

Expression	Return Value
<i>expr2</i> or <i>expr3</i> returns a string	string
<i>expr2</i> or <i>expr3</i> returns a floating-point value	floating-point
<i>expr2</i> or <i>expr3</i> returns an integer	integer

If *expr2* and *expr3* are both strings, the result is case sensitive if either string is case sensitive.



**Note**

There is also an *IF statement*, which differs from the `IF()` [1174] function described here. See [Section 13.6.5.2, “IF Syntax”](#).

- `IFNULL(expr1,expr2)` [1175]

If `expr1` is not `NULL`, `IFNULL()` [1175] returns `expr1`; otherwise it returns `expr2`. `IFNULL()` [1175] returns a numeric or string value, depending on the context in which it is used.

```
mysql> SELECT IFNULL(1,0);
-> 1
mysql> SELECT IFNULL(NULL,10);
-> 10
mysql> SELECT IFNULL(1/0,10);
-> 10
mysql> SELECT IFNULL(1/0,'yes');
-> 'yes'
```

The default result value of `IFNULL(expr1,expr2)` [1175] is the more “general” of the two expressions, in the order `STRING`, `REAL`, or `INTEGER`. Consider the case of a table based on expressions or where MySQL must internally store a value returned by `IFNULL()` [1175] in a temporary table:

```
mysql> CREATE TABLE tmp SELECT IFNULL(1,'test') AS test;
mysql> DESCRIBE tmp;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| test  | varbinary(4) | NO   |     |         |       |
+-----+-----+-----+-----+-----+-----+
```

In this example, the type of the `test` column is `VARBINARY(4)`.

- `NULLIF(expr1,expr2)` [1175]

Returns `NULL` if `expr1 = expr2` is true, otherwise returns `expr1`. This is the same as `CASE WHEN expr1 = expr2 THEN NULL ELSE expr1 END` [1174].

```
mysql> SELECT NULLIF(1,1);
-> NULL
mysql> SELECT NULLIF(1,2);
-> 1
```

Note that MySQL evaluates `expr1` twice if the arguments are not equal.

## 12.5. String Functions

Table 12.7. String Operators

Name	Description
<code>ASCII()</code> [1177]	Return numeric value of left-most character
<code>BIN()</code> [1177]	Return a string containing binary representation of a number
<code>BIT_LENGTH()</code> [1177]	Return length of argument in bits
<code>CHAR_LENGTH()</code> [1178]	Return number of characters in argument
<code>CHAR()</code> [1177]	Return the character for each integer passed
<code>CHARACTER_LENGTH()</code> [1178]	A synonym for <code>CHAR_LENGTH()</code>
<code>CONCAT_WS()</code> [1179]	Return concatenate with separator
<code>CONCAT()</code> [1178]	Return concatenated string

Name	Description
<a href="#">ELT() [1179]</a>	Return string at index number
<a href="#">EXPORT_SET() [1179]</a>	Return a string such that for every bit set in the value bits, you get an on string and for every unset bit, you get an off string
<a href="#">FIELD() [1179]</a>	Return the index (position) of the first argument in the subsequent arguments
<a href="#">FIND_IN_SET() [1179]</a>	Return the index position of the first argument within the second argument
<a href="#">FORMAT() [1180]</a>	Return a number formatted to specified number of decimal places
<a href="#">FROM_BASE64() [1180]</a>	Decode to a base-64 string and return result
<a href="#">HEX() [1180]</a>	Return a hexadecimal representation of a decimal or string value
<a href="#">INSERT() [1180]</a>	Insert a substring at the specified position up to the specified number of characters
<a href="#">INSTR() [1181]</a>	Return the index of the first occurrence of substring
<a href="#">LCASE() [1181]</a>	Synonym for LOWER()
<a href="#">LEFT() [1181]</a>	Return the leftmost number of characters as specified
<a href="#">LENGTH() [1181]</a>	Return the length of a string in bytes
<a href="#">LIKE [1190]</a>	Simple pattern matching
<a href="#">LOAD_FILE() [1181]</a>	Load the named file
<a href="#">LOCATE() [1181]</a>	Return the position of the first occurrence of substring
<a href="#">LOWER() [1182]</a>	Return the argument in lowercase
<a href="#">LPAD() [1182]</a>	Return the string argument, left-padded with the specified string
<a href="#">LTRIM() [1182]</a>	Remove leading spaces
<a href="#">MAKE_SET() [1182]</a>	Return a set of comma-separated strings that have the corresponding bit in bits set
<a href="#">MATCH [1229]</a>	Perform full-text search
<a href="#">MID() [1183]</a>	Return a substring starting from the specified position
<a href="#">NOT LIKE [1192]</a>	Negation of simple pattern matching
<a href="#">NOT REGEXP [1193]</a>	Negation of REGEXP
<a href="#">OCT() [1183]</a>	Return a string containing octal representation of a number
<a href="#">OCTET_LENGTH() [1183]</a>	A synonym for LENGTH()
<a href="#">ORD() [1183]</a>	Return character code for leftmost character of the argument
<a href="#">POSITION() [1183]</a>	A synonym for LOCATE()
<a href="#">QUOTE() [1183]</a>	Escape the argument for use in an SQL statement
<a href="#">REGEXP [1193]</a>	Pattern matching using regular expressions
<a href="#">REPEAT() [1183]</a>	Repeat a string the specified number of times
<a href="#">REPLACE() [1183]</a>	Replace occurrences of a specified string
<a href="#">REVERSE() [1184]</a>	Reverse the characters in a string
<a href="#">RIGHT() [1184]</a>	Return the specified rightmost number of characters
<a href="#">RLIKE [1193]</a>	Synonym for REGEXP

Name	Description
<a href="#">RPAD() [1184]</a>	Append string the specified number of times
<a href="#">RTRIM() [1184]</a>	Remove trailing spaces
<a href="#">SOUNDEX() [1184]</a>	Return a soundex string
<a href="#">SOUNDS LIKE [1185]</a>	Compare sounds
<a href="#">SPACE() [1185]</a>	Return a string of the specified number of spaces
<a href="#">STRCMP() [1192]</a>	Compare two strings
<a href="#">SUBSTR() [1185]</a>	Return the substring as specified
<a href="#">SUBSTRING_INDEX() [1185]</a>	Return a substring from a string before the specified number of occurrences of the delimiter
<a href="#">SUBSTRING() [1185]</a>	Return the substring as specified
<a href="#">TO_BASE64() [1186]</a>	Return the argument converted to a base-64 string
<a href="#">TRIM() [1186]</a>	Remove leading and trailing spaces
<a href="#">UCASE() [1186]</a>	Synonym for UPPER()
<a href="#">UNHEX() [1186]</a>	Return a string containing hex representation of a number
<a href="#">UPPER() [1187]</a>	Convert to uppercase
<a href="#">WEIGHT_STRING() [1187]</a>	Return the weight string for a string

String-valued functions return `NULL` if the length of the result would be greater than the value of the `max_allowed_packet` [502] system variable. See [Section 8.11.2, “Tuning Server Parameters”](#).

For functions that operate on string positions, the first position is numbered 1.

For functions that take length arguments, noninteger arguments are rounded to the nearest integer.

- [ASCII\(\*str\*\) \[1177\]](#)

Returns the numeric value of the leftmost character of the string *str*. Returns 0 if *str* is the empty string. Returns `NULL` if *str* is `NULL`. [ASCII\(\)](#) [1177] works for 8-bit characters.

```
mysql> SELECT ASCII('2');
-> 50
mysql> SELECT ASCII(2);
-> 50
mysql> SELECT ASCII('dx');
-> 100
```

See also the [ORD\(\)](#) [1183] function.

- [BIN\(\*N\*\) \[1177\]](#)

Returns a string representation of the binary value of *N*, where *N* is a longlong (`BIGINT`) number. This is equivalent to [CONV\(\*N\*, 10, 2\)](#) [1203]. Returns `NULL` if *N* is `NULL`.

```
mysql> SELECT BIN(12);
-> '1100'
```

- [BIT\\_LENGTH\(\*str\*\) \[1177\]](#)

Returns the length of the string *str* in bits.

```
mysql> SELECT BIT_LENGTH('text');
-> 32
```

- [CHAR\(\*N\*, ... \[USING \*charset\\_name\*\]\) \[1177\]](#)

`CHAR()` [1177] interprets each argument *N* as an integer and returns a string consisting of the characters given by the code values of those integers. `NULL` values are skipped.

```
mysql> SELECT CHAR(77,121,83,81,'76');
-> 'MySQL'
mysql> SELECT CHAR(77,77.3,'77.3');
-> 'MMM'
```

`CHAR()` [1177] arguments larger than 255 are converted into multiple result bytes. For example, `CHAR(256)` [1177] is equivalent to `CHAR(1,0)` [1177], and `CHAR(256*256)` [1177] is equivalent to `CHAR(1,0,0)` [1177]:

```
mysql> SELECT HEX(CHAR(1,0)), HEX(CHAR(256));
+-----+-----+
| HEX(CHAR(1,0)) | HEX(CHAR(256)) |
+-----+-----+
| 0100          | 0100          |
+-----+-----+
mysql> SELECT HEX(CHAR(1,0,0)), HEX(CHAR(256*256));
+-----+-----+
| HEX(CHAR(1,0,0)) | HEX(CHAR(256*256)) |
+-----+-----+
| 010000         | 010000         |
+-----+-----+
```

By default, `CHAR()` [1177] returns a binary string. To produce a string in a given character set, use the optional `USING` clause:

```
mysql> SELECT CHARSET(CHAR(0x65)), CHARSET(CHAR(0x65 USING utf8));
+-----+-----+
| CHARSET(CHAR(0x65)) | CHARSET(CHAR(0x65 USING utf8)) |
+-----+-----+
| binary              | utf8                             |
+-----+-----+
```

If `USING` is given and the result string is illegal for the given character set, a warning is issued. Also, if strict SQL mode is enabled, the result from `CHAR()` [1177] becomes `NULL`.

- `CHAR_LENGTH(str)` [1178]

Returns the length of the string *str*, measured in characters. A multi-byte character counts as a single character. This means that for a string containing five 2-byte characters, `LENGTH()` [1181] returns 10, whereas `CHAR_LENGTH()` [1178] returns 5.

- `CHARACTER_LENGTH(str)` [1178]

`CHARACTER_LENGTH()` [1178] is a synonym for `CHAR_LENGTH()` [1178].

- `CONCAT(str1,str2,...)` [1178]

Returns the string that results from concatenating the arguments. May have one or more arguments. If all arguments are nonbinary strings, the result is a nonbinary string. If the arguments include any binary strings, the result is a binary string. A numeric argument is converted to its equivalent nonbinary string form.

`CONCAT()` [1178] returns `NULL` if any argument is `NULL`.

```
mysql> SELECT CONCAT('My', 'S', 'QL');
-> 'MySQL'
mysql> SELECT CONCAT('My', NULL, 'QL');
-> NULL
mysql> SELECT CONCAT(14.3);
-> '14.3'
```

For quoted strings, concatenation can be performed by placing the strings next to each other:

```
mysql> SELECT 'My' 'S' 'QL';
-> 'MySQL'
```

- `CONCAT_WS(separator, str1, str2, ...)` [1179]

`CONCAT_WS()` [1179] stands for Concatenate With Separator and is a special form of `CONCAT()` [1178]. The first argument is the separator for the rest of the arguments. The separator is added between the strings to be concatenated. The separator can be a string, as can the rest of the arguments. If the separator is `NULL`, the result is `NULL`.

```
mysql> SELECT CONCAT_WS(',', 'First name', 'Second name', 'Last Name');
-> 'First name,Second name,Last Name'
mysql> SELECT CONCAT_WS(',', 'First name', NULL, 'Last Name');
-> 'First name,Last Name'
```

`CONCAT_WS()` [1179] does not skip empty strings. However, it does skip any `NULL` values after the separator argument.

- `ELT(N, str1, str2, str3, ...)` [1179]

`ELT()` [1179] returns the  $N$ th element of the list of strings: `str1` if  $N = 1$ , `str2` if  $N = 2$ , and so on. Returns `NULL` if  $N$  is less than 1 or greater than the number of arguments. `ELT()` [1179] is the complement of `FIELD()` [1179].

```
mysql> SELECT ELT(1, 'ej', 'Heja', 'hej', 'foo');
-> 'ej'
mysql> SELECT ELT(4, 'ej', 'Heja', 'hej', 'foo');
-> 'foo'
```

- `EXPORT_SET(bits, on, off[, separator[, number_of_bits]])` [1179]

Returns a string such that for every bit set in the value `bits`, you get an `on` string and for every bit not set in the value, you get an `off` string. Bits in `bits` are examined from right to left (from low-order to high-order bits). Strings are added to the result from left to right, separated by the `separator` string (the default being the comma character “,”). The number of bits examined is given by `number_of_bits`, which has a default of 64 if not specified. `number_of_bits` is silently clipped to 64 if larger than 64. It is treated as an unsigned integer, so a value of  $-1$  is effectively the same as 64.

```
mysql> SELECT EXPORT_SET(5, 'Y', 'N', ',', 4);
-> 'Y,N,Y,N'
mysql> SELECT EXPORT_SET(6, '1', '0', ',', 10);
-> '0,1,1,0,0,0,0,0,0,0'
```

- `FIELD(str, str1, str2, str3, ...)` [1179]

Returns the index (position) of `str` in the `str1, str2, str3, ...` list. Returns 0 if `str` is not found.

If all arguments to `FIELD()` [1179] are strings, all arguments are compared as strings. If all arguments are numbers, they are compared as numbers. Otherwise, the arguments are compared as double.

If `str` is `NULL`, the return value is 0 because `NULL` fails equality comparison with any value. `FIELD()` [1179] is the complement of `ELT()` [1179].

```
mysql> SELECT FIELD('ej', 'Hej', 'ej', 'Heja', 'hej', 'foo');
-> 2
mysql> SELECT FIELD('fo', 'Hej', 'ej', 'Heja', 'hej', 'foo');
-> 0
```

- `FIND_IN_SET(str, strlist)` [1179]

Returns a value in the range of 1 to  $N$  if the string `str` is in the string list `strlist` consisting of  $N$  substrings. A string list is a string composed of substrings separated by “,” characters. If the first

argument is a constant string and the second is a column of type `SET`, the `FIND_IN_SET()` [1179] function is optimized to use bit arithmetic. Returns `0` if `str` is not in `strlist` or if `strlist` is the empty string. Returns `NULL` if either argument is `NULL`. This function does not work properly if the first argument contains a comma (",") character.

```
mysql> SELECT FIND_IN_SET('b','a,b,c,d');
-> 2
```

- `FORMAT(X,D[,locale])` [1180]

Formats the number `x` to a format like '#,###,###.##', rounded to `D` decimal places, and returns the result as a string. If `D` is `0`, the result has no decimal point or fractional part.

The optional third parameter enables a locale to be specified to be used for the result number's decimal point, thousands separator, and grouping between separators. Permissible locale values are the same as the legal values for the `lc_time_names` [495] system variable (see [Section 10.7, "MySQL Server Locale Support"](#)). If no locale is specified, the default is `'en_US'`.

```
mysql> SELECT FORMAT(12332.123456, 4);
-> '12,332.1235'
mysql> SELECT FORMAT(12332.1,4);
-> '12,332.1000'
mysql> SELECT FORMAT(12332.2,0);
-> '12,332'
mysql> SELECT FORMAT(12332.2,2,'de_DE');
-> '12.332,20'
```

- `FROM_BASE64(str)` [1180]

Takes a string encoded with the base-64 encoded rules used by `TO_BASE64()` [1186] and returns the decoded result as a binary string. The result is `NULL` if the argument is `NULL` or not a valid base-64 string. See the description of `TO_BASE64()` [1186] for details about the encoding and decoding rules.

This function was added in MySQL 5.6.1.

```
mysql> SELECT TO_BASE64('abc'), FROM_BASE64(TO_BASE64('abc'));
-> 'JWJj', 'abc'
```

- `HEX(str)` [1180], `HEX(N)` [1180]

For a string argument `str`, `HEX()` [1180] returns a hexadecimal string representation of `str` where each character in `str` is converted to two hexadecimal digits. The inverse of this operation is performed by the `UNHEX()` [1186] function.

For a numeric argument `N`, `HEX()` [1180] returns a hexadecimal string representation of the value of `N` treated as a longlong (`BIGINT`) number. This is equivalent to `CONV(N,10,16)` [1203]. The inverse of this operation is performed by `CONV(HEX(N),16,10)` [1203].

```
mysql> SELECT 0x616263, HEX('abc'), UNHEX(HEX('abc'));
-> 'abc', 616263, 'abc'
mysql> SELECT HEX(255), CONV(HEX(255),16,10);
-> 'FF', 255
```

- `INSERT(str,pos,len,newstr)` [1180]

Returns the string `str`, with the substring beginning at position `pos` and `len` characters long replaced by the string `newstr`. Returns the original string if `pos` is not within the length of the string. Replaces the rest of the string from position `pos` if `len` is not within the length of the rest of the string. Returns `NULL` if any argument is `NULL`.

```
mysql> SELECT INSERT('Quadratic', 3, 4, 'What');
-> 'QuWhattic'
mysql> SELECT INSERT('Quadratic', -1, 4, 'What');
-> 'Quadratic'
```

```
mysql> SELECT INSERT('Quadratic', 3, 100, 'What');
-> 'QuWhat'
```

This function is multi-byte safe.

- [INSTR\(\*str\*,\*substr\*\) \[1181\]](#)

Returns the position of the first occurrence of substring *substr* in string *str*. This is the same as the two-argument form of [LOCATE\(\) \[1181\]](#), except that the order of the arguments is reversed.

```
mysql> SELECT INSTR('foobarbar', 'bar');
-> 4
mysql> SELECT INSTR('xbar', 'foobar');
-> 0
```

This function is multi-byte safe, and is case sensitive only if at least one argument is a binary string.

- [LCASE\(\*str\*\) \[1181\]](#)

[LCASE\(\) \[1181\]](#) is a synonym for [LOWER\(\) \[1182\]](#).

- [LEFT\(\*str\*,\*len\*\) \[1181\]](#)

Returns the leftmost *len* characters from the string *str*, or [NULL](#) if any argument is [NULL](#).

```
mysql> SELECT LEFT('foobarbar', 5);
-> 'fooba'
```

This function is multi-byte safe.

- [LENGTH\(\*str\*\) \[1181\]](#)

Returns the length of the string *str*, measured in bytes. A multi-byte character counts as multiple bytes. This means that for a string containing five 2-byte characters, [LENGTH\(\) \[1181\]](#) returns 10, whereas [CHAR\\_LENGTH\(\) \[1178\]](#) returns 5.

```
mysql> SELECT LENGTH('text');
-> 4
```

- [LOAD\\_FILE\(\*file\\_name\*\) \[1181\]](#)

Reads the file and returns the file contents as a string. To use this function, the file must be located on the server host, you must specify the full path name to the file, and you must have the [FILE \[754\]](#) privilege. The file must be readable by all and its size less than [max\\_allowed\\_packet \[502\]](#) bytes. If the [secure\\_file\\_priv \[538\]](#) system variable is set to a nonempty directory name, the file to be loaded must be located in that directory.

If the file does not exist or cannot be read because one of the preceding conditions is not satisfied, the function returns [NULL](#).

The [character\\_set\\_filesystem \[468\]](#) system variable controls interpretation of file names that are given as literal strings.

```
mysql> UPDATE t
      SET blob_col=LOAD_FILE('/tmp/picture')
      WHERE id=1;
```

- [LOCATE\(\*substr\*,\*str\*\) \[1181\]](#), [LOCATE\(\*substr\*,\*str\*,\*pos\*\) \[1181\]](#)

The first syntax returns the position of the first occurrence of substring *substr* in string *str*. The second syntax returns the position of the first occurrence of substring *substr* in string *str*, starting at position *pos*. Returns 0 if *substr* is not in *str*.

```
mysql> SELECT LOCATE('bar', 'foobarbar');
-> 4
```



```
mysql> SELECT LOCATE('xbar', 'foobar');
-> 0
mysql> SELECT LOCATE('bar', 'foobarbar', 5);
-> 7
```

This function is multi-byte safe, and is case-sensitive only if at least one argument is a binary string.

- [LOWER\(\*str\*\)](#) [1182]

Returns the string *str* with all characters changed to lowercase according to the current character set mapping. The default is `latin1` (cp1252 West European).

```
mysql> SELECT LOWER('QUADRATICALLY');
-> 'quadratically'
```

[LOWER\(\)](#) [1182] (and [UPPER\(\)](#) [1187]) are ineffective when applied to binary strings (`BINARY`, `VARBINARY`, `BLOB`). To perform lettercase conversion, convert the string to a nonbinary string:

```
mysql> SET @str = BINARY 'New York';
mysql> SELECT LOWER(@str), LOWER(CONVERT(@str USING latin1));
+-----+-----+
| LOWER(@str) | LOWER(CONVERT(@str USING latin1)) |
+-----+-----+
| New York    | new york                          |
+-----+-----+
```

For Unicode character sets, [LOWER\(\)](#) [1182] and [UPPER\(\)](#) [1187] work accounting to Unicode Collation Algorithm (UCA) 5.2.0 for `xxx_unicode_520_ci` collations and for language-specific collations that are derived from them. For other Unicode collations, [LOWER\(\)](#) [1182] and [UPPER\(\)](#) [1187] work accounting to Unicode Collation Algorithm (UCA) 4.0.0. See [Section 10.1.14.1, “Unicode Character Sets”](#).

This function is multi-byte safe.

- [LPAD\(\*str\*, \*len\*, \*padstr\*\)](#) [1182]

Returns the string *str*, left-padded with the string *padstr* to a length of *len* characters. If *str* is longer than *len*, the return value is shortened to *len* characters.

```
mysql> SELECT LPAD('hi',4,'??');
-> '??hi'
mysql> SELECT LPAD('hi',1,'??');
-> 'h'
```

- [LTRIM\(\*str\*\)](#) [1182]

Returns the string *str* with leading space characters removed.

```
mysql> SELECT LTRIM('  barbar');
-> 'barbar'
```

This function is multi-byte safe.

- [MAKE\\_SET\(\*bits\*, \*str1\*, \*str2\*, ...\)](#) [1182]

Returns a set value (a string containing substrings separated by “,” characters) consisting of the strings that have the corresponding bit in *bits* set. *str1* corresponds to bit 0, *str2* to bit 1, and so on. `NULL` values in *str1*, *str2*, ... are not appended to the result.

```
mysql> SELECT MAKE_SET(1,'a','b','c');
-> 'a'
mysql> SELECT MAKE_SET(1 | 4,'hello','nice','world');
-> 'hello,world'
mysql> SELECT MAKE_SET(1 | 4,'hello','nice',NULL,'world');
-> 'hello'
mysql> SELECT MAKE_SET(0,'a','b','c');
-> ''
```

- `MID(str,pos,len)` [1183]

`MID(str,pos,len)` [1183] is a synonym for `SUBSTRING(str,pos,len)` [1185].

- `OCT(N)` [1183]

Returns a string representation of the octal value of *N*, where *N* is a longlong (`BIGINT`) number. This is equivalent to `CONV(N,10,8)` [1203]. Returns `NULL` if *N* is `NULL`.

```
mysql> SELECT OCT(12);
-> '14'
```

- `OCTET_LENGTH(str)` [1183]

`OCTET_LENGTH()` [1183] is a synonym for `LENGTH()` [1181].

- `ORD(str)` [1183]

If the leftmost character of the string *str* is a multi-byte character, returns the code for that character, calculated from the numeric values of its constituent bytes using this formula:

```
(1st byte code)
+ (2nd byte code * 256)
+ (3rd byte code * 2562) ...
```

If the leftmost character is not a multi-byte character, `ORD()` [1183] returns the same value as the `ASCII()` [1177] function.

```
mysql> SELECT ORD('2');
-> 50
```

- `POSITION(substr IN str)` [1183]

`POSITION(substr IN str)` [1183] is a synonym for `LOCATE(substr,str)` [1181].

- `QUOTE(str)` [1183]

Quotes a string to produce a result that can be used as a properly escaped data value in an SQL statement. The string is returned enclosed by single quotation marks and with each instance of backslash (“\”), single quote (“'”), ASCII `NULL`, and Control+Z preceded by a backslash. If the argument is `NULL`, the return value is the word “NULL” without enclosing single quotation marks.

```
mysql> SELECT QUOTE('Don\'t!');
-> 'Don\'t!'
mysql> SELECT QUOTE(NULL);
-> NULL
```

For comparison, see the quoting rules for literal strings and within the C API in [Section 9.1.1, “String Literals”](#), and [Section 21.8.6.54, “mysql\\_real\\_escape\\_string\(\)”](#).

- `REPEAT(str,count)` [1183]

Returns a string consisting of the string *str* repeated *count* times. If *count* is less than 1, returns an empty string. Returns `NULL` if *str* or *count* are `NULL`.

```
mysql> SELECT REPEAT('MySQL', 3);
-> 'MySQLMySQLMySQL'
```

- `REPLACE(str,from_str,to_str)` [1183]

Returns the string *str* with all occurrences of the string *from\_str* replaced by the string *to\_str*. `REPLACE()` [1183] performs a case-sensitive match when searching for *from\_str*.

```
mysql> SELECT REPLACE('www.mysql.com', 'w', 'Ww');
-> 'WwWwWw.mysql.com'
```

This function is multi-byte safe.

- [REVERSE\(\*str\*\)](#) [1184]

Returns the string *str* with the order of the characters reversed.

```
mysql> SELECT REVERSE('abc');
-> 'cba'
```

This function is multi-byte safe.

- [RIGHT\(\*str\*, \*len\*\)](#) [1184]

Returns the rightmost *len* characters from the string *str*, or `NULL` if any argument is `NULL`.

```
mysql> SELECT RIGHT('foobarbar', 4);
-> 'rbar'
```

This function is multi-byte safe.

- [RPAD\(\*str\*, \*len\*, \*padstr\*\)](#) [1184]

Returns the string *str*, right-padded with the string *padstr* to a length of *len* characters. If *str* is longer than *len*, the return value is shortened to *len* characters.

```
mysql> SELECT RPAD('hi',5,'?');
-> 'hi???'
mysql> SELECT RPAD('hi',1,'?');
-> 'h'
```

This function is multi-byte safe.

- [RTRIM\(\*str\*\)](#) [1184]

Returns the string *str* with trailing space characters removed.

```
mysql> SELECT RTRIM('barbar ');
-> 'barbar'
```

This function is multi-byte safe.

- [SOUNDEX\(\*str\*\)](#) [1184]

Returns a soundex string from *str*. Two strings that sound almost the same should have identical soundex strings. A standard soundex string is four characters long, but the [SOUNDEX\(\)](#) [1184] function returns an arbitrarily long string. You can use [SUBSTRING\(\)](#) [1185] on the result to get a standard soundex string. All nonalphabetic characters in *str* are ignored. All international alphabetic characters outside the A-Z range are treated as vowels.



#### Important

When using [SOUNDEX\(\)](#) [1184], you should be aware of the following limitations:

- This function, as currently implemented, is intended to work well with strings that are in the English language only. Strings in other languages may not produce reliable results.
- This function is not guaranteed to provide consistent results with strings that use multi-byte character sets, including `utf-8`.

We hope to remove these limitations in a future release. See Bug #22638 for more information.

```
mysql> SELECT SOUNDEX('Hello');
-> 'H400'
```

```
mysql> SELECT SOUNDEX('Quadratically');
-> 'Q36324'
```

**Note**

This function implements the original Soundex algorithm, not the more popular enhanced version (also described by D. Knuth). The difference is that original version discards vowels first and duplicates second, whereas the enhanced version discards duplicates first and vowels second.

- `expr1 SOUNDS LIKE expr2` [1185]

This is the same as `SOUNDEX(expr1) = SOUNDEX(expr2)` [1184].

- `SPACE(N)` [1185]

Returns a string consisting of *N* space characters.

```
mysql> SELECT SPACE(6);
-> '      '
```

- `SUBSTR(str, pos)` [1185], `SUBSTR(str FROM pos)` [1185], `SUBSTR(str, pos, len)` [1185], `SUBSTR(str FROM pos FOR len)` [1185]

`SUBSTR()` [1185] is a synonym for `SUBSTRING()` [1185].

- `SUBSTRING(str, pos)` [1185], `SUBSTRING(str FROM pos)` [1185], `SUBSTRING(str, pos, len)` [1185], `SUBSTRING(str FROM pos FOR len)` [1185]

The forms without a *len* argument return a substring from string *str* starting at position *pos*. The forms with a *len* argument return a substring *len* characters long from string *str*, starting at position *pos*. The forms that use `FROM` are standard SQL syntax. It is also possible to use a negative value for *pos*. In this case, the beginning of the substring is *pos* characters from the end of the string, rather than the beginning. A negative value may be used for *pos* in any of the forms of this function.

For all forms of `SUBSTRING()` [1185], the position of the first character in the string from which the substring is to be extracted is reckoned as 1.

```
mysql> SELECT SUBSTRING('Quadratically',5);
-> 'ratically'
mysql> SELECT SUBSTRING('foobarbar' FROM 4);
-> 'barbar'
mysql> SELECT SUBSTRING('Quadratically',5,6);
-> 'ratica'
mysql> SELECT SUBSTRING('Sakila', -3);
-> 'ila'
mysql> SELECT SUBSTRING('Sakila', -5, 3);
-> 'aki'
mysql> SELECT SUBSTRING('Sakila' FROM -4 FOR 2);
-> 'ki'
```

This function is multi-byte safe.

If *len* is less than 1, the result is the empty string.

- `SUBSTRING_INDEX(str, delim, count)` [1185]

Returns the substring from string *str* before *count* occurrences of the delimiter *delim*. If *count* is positive, everything to the left of the final delimiter (counting from the left) is returned. If *count* is negative, everything to the right of the final delimiter (counting from the right) is returned. `SUBSTRING_INDEX()` [1185] performs a case-sensitive match when searching for *delim*.

```
mysql> SELECT SUBSTRING_INDEX('www.mysql.com', '.', 2);
-> 'www.mysql'
```

```
mysql> SELECT SUBSTRING_INDEX('www.mysql.com', '.', -2);
-> 'mysql.com'
```

This function is multi-byte safe.

- [TO\\_BASE64\(\*str\*\)](#) [1186]

Converts the string argument to base-64 encoded form and returns the result as a character string with the connection character set and collation. If the argument is not a string, it is converted to a string before conversion takes place. The result is `NULL` if the argument is `NULL`. Base-64 encoded strings can be decoded using the [FROM\\_BASE64\(\)](#) [1180] function.

This function was added in MySQL 5.6.1.

```
mysql> SELECT TO_BASE64('abc'), FROM_BASE64(TO_BASE64('abc'));
-> 'JWJj', 'abc'
```

Different base-64 encoding schemes exist. These are the encoding and decoding rules used by [TO\\_BASE64\(\)](#) [1186] and [FROM\\_BASE64\(\)](#) [1180]:

- The encoding for alphabet value 62 is '+'.
- The encoding for alphabet value 63 is '/'.
- Encoded output consists of groups of 4 printable characters. Each 3 bytes of the input data are encoded using 4 characters. If the last group is incomplete, it is padded with '=' characters to a length of 4.
- A newline is added after each 76 characters of encoded output to divide long output into multiple lines.
- Decoding recognizes and ignores newline, carriage return, tab, and space.
- [TRIM\(\[{BOTH | LEADING | TRAILING} \[\*remstr\*\] FROM\] \*str\*\)](#) [1186], [TRIM\(\[\*remstr\*\] FROM\] \*str\*\)](#) [1186]

Returns the string *str* with all *remstr* prefixes or suffixes removed. If none of the specifiers `BOTH`, `LEADING`, or `TRAILING` is given, `BOTH` is assumed. *remstr* is optional and, if not specified, spaces are removed.

```
mysql> SELECT TRIM(' bar ');
-> 'bar'
mysql> SELECT TRIM(LEADING 'x' FROM 'xxxbarxxx');
-> 'barxxx'
mysql> SELECT TRIM(BOTH 'x' FROM 'xxxbarxxx');
-> 'bar'
mysql> SELECT TRIM(TRAILING 'xyz' FROM 'barxyz');
-> 'barx'
```

This function is multi-byte safe.

- [UCASE\(\*str\*\)](#) [1186]

[UCASE\(\)](#) [1186] is a synonym for [UPPER\(\)](#) [1187].

- [UNHEX\(\*str\*\)](#) [1186]

For a string argument *str*, [UNHEX\(\*str\*\)](#) [1186] performs the inverse operation of [HEX\(\*str\*\)](#) [1180]. That is, it interprets each pair of characters in the argument as a hexadecimal number and converts it to the character represented by the number. The return value is a binary string.

```
mysql> SELECT UNHEX('4D7953514C');
-> 'MySQL'
mysql> SELECT 0x4D7953514C;
```

```

-> 'MySQL'
mysql> SELECT UNHEX(HEX('string'));
-> 'string'
mysql> SELECT HEX(UNHEX('1267'));
-> '1267'

```

The characters in the argument string must be legal hexadecimal digits: '0' .. '9', 'A' .. 'F', 'a' .. 'f'. If the argument contains any nonhexadecimal digits, the result is `NULL`:

```

mysql> SELECT UNHEX('GG');
+-----+
| UNHEX('GG') |
+-----+
| NULL        |
+-----+

```

A `NULL` result can occur if the argument to `UNHEX()` [1186] is a `BINARY` column, because values are padded with `0x00` bytes when stored but those bytes are not stripped on retrieval. For example, '41' is stored into a `CHAR(3)` column as '41 ' and retrieved as '41' (with the trailing pad space stripped), so `UNHEX()` [1186] for the column value returns 'A'. By contrast '41' is stored into a `BINARY(3)` column as '41\0' and retrieved as '41\0' (with the trailing pad `0x00` byte not stripped). '\0' is not a legal hexadecimal digit, so `UNHEX()` [1186] for the column value returns `NULL`.

For a numeric argument *N*, the inverse of `HEX(N)` [1180] is not performed by `UNHEX()` [1186]. Use `CONV(HEX(N), 16, 10)` [1203] instead. See the description of `HEX()` [1180].

- `UPPER(str)` [1187]

Returns the string *str* with all characters changed to uppercase according to the current character set mapping. The default is `latin1` (cp1252 West European).

```

mysql> SELECT UPPER('Hej');
-> 'HEJ'

```

See the description of `LOWER()` [1182] for information that also applies to `UPPER()` [1187]. This included information about how to perform lettercase conversion of binary strings (`BINARY`, `VARBINARY`, `BLOB`) for which these functions are ineffective, and information about case folding for Unicode character sets.

This function is multi-byte safe.

- `WEIGHT_STRING(str [AS {CHAR|BINARY}(N)] [LEVEL levels] [flags])` [1187]

*levels*: *N* [ASC|DESC|REVERSE] [, *N* [ASC|DESC|REVERSE]] ...

This function returns the weight string for the input string. The return value is a binary string that represents the sorting and comparison value of the string. It has these properties:

- If `WEIGHT_STRING(str1)` [1187] = `WEIGHT_STRING(str2)` [1187], then *str1* = *str2* (*str1* and *str2* are considered equal)
- If `WEIGHT_STRING(str1)` [1187] < `WEIGHT_STRING(str2)` [1187], then *str1* < *str2* (*str1* sorts before *str2*)

`WEIGHT_STRING()` [1187] can be used for testing and debugging of collations, especially if you are adding a new collation. See [Section 10.4, “Adding a Collation to a Character Set”](#).

The input string, *str*, is a string expression. If the input is a nonbinary (character) string such as a `CHAR`, `VARCHAR`, or `TEXT` value, the return value contains the collation weights for the string. If the input is a binary (byte) string such as a `BINARY`, `VARBINARY`, or `BLOB` value, the return value is the same as the input (the weight for each byte in a binary string is the byte value). If the input is `NULL`, `WEIGHT_STRING()` [1187] returns `NULL`.

## Examples:

```
mysql> SET @s = _latin1 'AB' COLLATE latin1_swedish_ci;
mysql> SELECT @s, HEX(@s), HEX(WEIGHT_STRING(@s));
+-----+-----+-----+
| @s    | HEX(@s) | HEX(WEIGHT_STRING(@s)) |
+-----+-----+-----+
| AB    | 4142    | 4142                    |
+-----+-----+-----+
```

```
mysql> SET @s = _latin1 'ab' COLLATE latin1_swedish_ci;
mysql> SELECT @s, HEX(@s), HEX(WEIGHT_STRING(@s));
+-----+-----+-----+
| @s    | HEX(@s) | HEX(WEIGHT_STRING(@s)) |
+-----+-----+-----+
| ab    | 6162    | 4142                    |
+-----+-----+-----+
```

```
mysql> SET @s = CAST('AB' AS BINARY);
mysql> SELECT @s, HEX(@s), HEX(WEIGHT_STRING(@s));
+-----+-----+-----+
| @s    | HEX(@s) | HEX(WEIGHT_STRING(@s)) |
+-----+-----+-----+
| AB    | 4142    | 4142                    |
+-----+-----+-----+
```

```
mysql> SET @s = CAST('ab' AS BINARY);
mysql> SELECT @s, HEX(@s), HEX(WEIGHT_STRING(@s));
+-----+-----+-----+
| @s    | HEX(@s) | HEX(WEIGHT_STRING(@s)) |
+-----+-----+-----+
| ab    | 6162    | 6162                    |
+-----+-----+-----+
```

The preceding examples use `HEX()` [1180] to display the `WEIGHT_STRING()` [1187] result. Because the result is a binary value, `HEX()` [1180] can be especially useful when the result contains nonprinting values, to display it in printable form:

```
mysql> SET @s = CONVERT(0xC39F USING utf8) COLLATE utf8_czech_ci;
mysql> SELECT HEX(WEIGHT_STRING(@s));
+-----+
| HEX(WEIGHT_STRING(@s)) |
+-----+
| 0FEA0FEA                |
+-----+
```

For non-NULL return values, the data type of the value is `VARBINARY` if its length is within the maximum length for `VARBINARY`, otherwise the data type is `BLOB`.

The `AS` clause may be given to cast the input string to a nonbinary or binary string and to force it to a given length:

- `AS CHAR(N)` casts the string to a nonbinary string and pads it on the right with spaces to a length of `N` characters. `N` must be at least 1. If `N` is less than the length of the input string, the string is truncated to `N` characters. No warning occurs for truncation.
- `AS BINARY(N)` is similar but casts the string to a binary string, `N` is measured in bytes (not characters), and padding uses `0x00` bytes (not spaces).

```
mysql> SELECT HEX(WEIGHT_STRING('ab' AS CHAR(4)));
+-----+
| HEX(WEIGHT_STRING('ab' AS CHAR(4))) |
+-----+
| 41422020                            |
+-----+
```

```
mysql> SELECT HEX(WEIGHT_STRING('ab' AS BINARY(4)));
+-----+
```



```
| HEX(WEIGHT_STRING('ab' AS BINARY(4))) |
+-----+
| 61620000 |
+-----+
```

The `LEVEL` clause may be given to specify that the return value should contain weights for specific collation levels.

The `levels` specifier following the `LEVEL` keyword may be given either as a list of one or more integers separated by commas, or as a range of two integers separated by a dash. Whitespace around the punctuation characters does not matter.

Examples:

```
LEVEL 1
LEVEL 2, 3, 5
LEVEL 1-3
```

Any level less than 1 is treated as 1. Any level greater than the maximum for the input string collation is treated as maximum for the collation. The maximum varies per collation, but is never greater than 6.

In a list of levels, levels must be given in increasing order. In a range of levels, if the second number is less than the first, it is treated as the first number (for example, 4-2 is the same as 4-4).

If the `LEVEL` clause is omitted, MySQL assumes `LEVEL 1 - max`, where `max` is the maximum level for the collation.

If `LEVEL` is specified using list syntax (not range syntax), any level number can be followed by these modifiers:

- `ASC`: Return the weights without modification. This is the default.
- `DESC`: Return bitwise-inverted weights (for example, `0x78f0 DESC = 0x870f`).
- `REVERSE`: Return the weights in reverse order (that is, the weights for the reversed string, with the first character last and the last first).

Examples:

```
mysql> SELECT HEX(WEIGHT_STRING(0x007fff LEVEL 1));
+-----+
| HEX(WEIGHT_STRING(0x007fff LEVEL 1)) |
+-----+
| 007FFF |
+-----+
```

```
mysql> SELECT HEX(WEIGHT_STRING(0x007fff LEVEL 1 DESC));
+-----+
| HEX(WEIGHT_STRING(0x007fff LEVEL 1 DESC)) |
+-----+
| FF8000 |
+-----+
```

```
mysql> SELECT HEX(WEIGHT_STRING(0x007fff LEVEL 1 REVERSE));
+-----+
| HEX(WEIGHT_STRING(0x007fff LEVEL 1 REVERSE)) |
+-----+
| FF7F00 |
+-----+
```

```
mysql> SELECT HEX(WEIGHT_STRING(0x007fff LEVEL 1 DESC REVERSE));
+-----+
| HEX(WEIGHT_STRING(0x007fff LEVEL 1 DESC REVERSE)) |
+-----+
| 0080FF |
+-----+
```

The *flags* clause currently is unused.

## 12.5.1. String Comparison Functions

**Table 12.8. String Comparison Operators**

Name	Description
<a href="#">LIKE [1190]</a>	Simple pattern matching
<a href="#">NOT LIKE [1192]</a>	Negation of simple pattern matching
<a href="#">STRCMP() [1192]</a>	Compare two strings

If a string function is given a binary string as an argument, the resulting string is also a binary string. A number converted to a string is treated as a binary string. This affects only comparisons.

Normally, if any expression in a string comparison is case sensitive, the comparison is performed in case-sensitive fashion.

- `expr LIKE pat [ESCAPE 'escape_char'] [1190]`

Pattern matching using SQL simple regular expression comparison. Returns `1` (`TRUE`) or `0` (`FALSE`). If either *expr* or *pat* is `NULL`, the result is `NULL`.

The pattern need not be a literal string. For example, it can be specified as a string expression or table column.

Per the SQL standard, [LIKE \[1190\]](#) performs matching on a per-character basis, thus it can produce results different from the `=` [\[1166\]](#) comparison operator:

```
mysql> SELECT 'ä' LIKE 'ae' COLLATE latin1_german2_ci;
+-----+
| 'ä' LIKE 'ae' COLLATE latin1_german2_ci |
+-----+
|                                          0 |
+-----+
mysql> SELECT 'ä' = 'ae' COLLATE latin1_german2_ci;
+-----+
| 'ä' = 'ae' COLLATE latin1_german2_ci |
+-----+
|                                          1 |
+-----+
```

In particular, trailing spaces are significant, which is not true for `CHAR` or `VARCHAR` comparisons performed with the `=` [\[1166\]](#) operator:

```
mysql> SELECT 'a' = 'a ', 'a' LIKE 'a ';
+-----+-----+
| 'a' = 'a ' | 'a' LIKE 'a ' |
+-----+-----+
|          1 |              0 |
+-----+-----+
1 row in set (0.00 sec)
```

With [LIKE \[1190\]](#) you can use the following two wildcard characters in the pattern.

Character	Description
<code>%</code>	Matches any number of characters, even zero characters
<code>_</code>	Matches exactly one character

```
mysql> SELECT 'David!' LIKE 'David_';
-> 1
mysql> SELECT 'David!' LIKE '%D%v%';
```

```
-> 1
```

To test for literal instances of a wildcard character, precede it by the escape character. If you do not specify the `ESCAPE` character, “\” is assumed.

String	Description
<code>\%</code>	Matches one “%” character
<code>\_</code>	Matches one “_” character

```
mysql> SELECT 'David!' LIKE 'David\_';
-> 0
mysql> SELECT 'David_' LIKE 'David\_';
-> 1
```

To specify a different escape character, use the `ESCAPE` clause:

```
mysql> SELECT 'David_' LIKE 'David|_' ESCAPE '|';
-> 1
```

The escape sequence should be empty or one character long. The expression must evaluate as a constant at execution time. If the `NO_BACKSLASH_ESCAPES` [604] SQL mode is enabled, the sequence cannot be empty.

The following two statements illustrate that string comparisons are not case sensitive unless one of the operands is a binary string:

```
mysql> SELECT 'abc' LIKE 'ABC';
-> 1
mysql> SELECT 'abc' LIKE BINARY 'ABC';
-> 0
```

In MySQL, `LIKE` [1190] is permitted on numeric expressions. (This is an extension to the standard SQL `LIKE` [1190].)

```
mysql> SELECT 10 LIKE '1%';
-> 1
```



### Note

Because MySQL uses C escape syntax in strings (for example, “\n” to represent a newline character), you must double any “\” that you use in `LIKE` [1190] strings. For example, to search for “\n”, specify it as “\\n”. To search for “\”, specify it as “\\”;

Exception: At the end of the pattern string, backslash can be specified as “\\”. At the end of the string, backslash stands for itself because there is nothing following to escape. Suppose that a table contains the following values:

```
mysql> SELECT filename FROM t1;
+-----+
| filename |
+-----+
| C:      |
| C:\     |
| C:\Programs |
| C:\Programs\ |
+-----+
```

To test for values that end with backslash, you can match the values using either of the following patterns:

```
mysql> SELECT filename, filename LIKE '%\\' FROM t1;
+-----+-----+
| filename | filename LIKE '%\\' |
+-----+-----+
| C:       | 0 |
| C:\     | 1 |
| C:\Programs | 0 |
| C:\Programs\ | 1 |
+-----+-----+

mysql> SELECT filename, filename LIKE '%\\\' FROM t1;
+-----+-----+
| filename | filename LIKE '%\\\' |
+-----+-----+
| C:       | 0 |
| C:\     | 1 |
| C:\Programs | 0 |
| C:\Programs\ | 1 |
+-----+-----+
```

- `expr NOT LIKE pat [ESCAPE 'escape_char']` [1192]

This is the same as `NOT (expr LIKE pat [ESCAPE 'escape_char'])`.



### Note

Aggregate queries involving `NOT LIKE` [1192] comparisons with columns containing `NULL` may yield unexpected results. For example, consider the following table and data:

```
CREATE TABLE foo (bar VARCHAR(10));
INSERT INTO foo VALUES (NULL), (NULL);
```

The query `SELECT COUNT(*) FROM foo WHERE bar LIKE '%baz%';` returns 0. You might assume that `SELECT COUNT(*) FROM foo WHERE bar NOT LIKE '%baz%';` would return 2. However, this is not the case: The second query returns 0. This is because `NULL NOT LIKE expr` always returns `NULL`, regardless of the value of `expr`. The same is true for aggregate queries involving `NULL` and comparisons using `NOT RLIKE` [1193] or `NOT REGEXP` [1193]. In such cases, you must test explicitly for `NOT NULL` using `OR` [1172] (and not `AND` [1171]), as shown here:

```
SELECT COUNT(*) FROM foo WHERE bar NOT LIKE '%baz%' OR bar IS NULL;
```

- `STRCMP(expr1, expr2)` [1192]

`STRCMP()` [1192] returns 0 if the strings are the same, -1 if the first argument is smaller than the second according to the current sort order, and 1 otherwise.

```
mysql> SELECT STRCMP('text', 'text2');
-> -1
mysql> SELECT STRCMP('text2', 'text');
-> 1
mysql> SELECT STRCMP('text', 'text');
-> 0
```

`STRCMP()` [1192] performs the comparison using the collation of the arguments.

```
mysql> SET @s1 = _latin1 'x' COLLATE latin1_general_ci;
mysql> SET @s2 = _latin1 'X' COLLATE latin1_general_ci;
mysql> SET @s3 = _latin1 'x' COLLATE latin1_general_cs;
mysql> SET @s4 = _latin1 'X' COLLATE latin1_general_cs;
mysql> SELECT STRCMP(@s1, @s2), STRCMP(@s3, @s4);
+-----+-----+
| STRCMP(@s1, @s2) | STRCMP(@s3, @s4) |
+-----+-----+
```

0	1
+-----+	

If the collations are incompatible, one of the arguments must be converted to be compatible with the other. See [Section 10.1.7.5, “Collation of Expressions”](#).

```
mysql> SELECT STRCMP(@s1, @s3);
ERROR 1267 (HY000) at line 10: Illegal mix of collations (latin1_general_ci,IMPLICIT) and (latin1_ge
mysql> SELECT STRCMP(@s1, @s3 COLLATE latin1_general_ci);
+-----+
| STRCMP(@s1, @s3 COLLATE latin1_general_ci) |
+-----+
|                                           0 |
+-----+
```

## 12.5.2. Regular Expressions

**Table 12.9. String Regular Expression Operators**

Name	Description
<code>NOT REGEXP</code> [1193]	Negation of REGEXP
<code>REGEXP</code> [1193]	Pattern matching using regular expressions
<code>RLIKE</code> [1193]	Synonym for REGEXP

A regular expression is a powerful way of specifying a pattern for a complex search.

MySQL uses Henry Spencer's implementation of regular expressions, which is aimed at conformance with POSIX 1003.2. MySQL uses the extended version to support pattern-matching operations performed with the `REGEXP` [1193] operator in SQL statements.

This section summarizes, with examples, the special characters and constructs that can be used in MySQL for `REGEXP` [1193] operations. It does not contain all the details that can be found in Henry Spencer's `regex(7)` manual page. That manual page is included in MySQL source distributions, in the `regex.7` file under the `regex` directory. See also [Section 3.3.4.7, “Pattern Matching”](#).

- `expr NOT REGEXP pat` [1193], `expr NOT RLIKE pat` [1193]

This is the same as `NOT (expr REGEXP pat)`.

- `expr REGEXP pat` [1193], `expr RLIKE pat` [1193]

Performs a pattern match of a string expression `expr` against a pattern `pat`. The pattern can be an extended regular expression. The syntax for regular expressions is discussed in [Section 12.5.2, “Regular Expressions”](#). Returns 1 if `expr` matches `pat`; otherwise it returns 0. If either `expr` or `pat` is `NULL`, the result is `NULL`. `RLIKE` [1193] is a synonym for `REGEXP` [1193], provided for `mSQL` compatibility.

The pattern need not be a literal string. For example, it can be specified as a string expression or table column.



**Note**

Because MySQL uses the C escape syntax in strings (for example, “\n” to represent the newline character), you must double any “\” that you use in your `REGEXP` [1193] strings.

`REGEXP` [1193] is not case sensitive, except when used with binary strings.

```
mysql> SELECT 'Monty!' REGEXP '.*';
-> 1
mysql> SELECT 'new*\n*line' REGEXP 'new\\*.*\\*line';
-> 1
mysql> SELECT 'a' REGEXP 'A', 'a' REGEXP BINARY 'A';
```

```

-> 1 0
mysql> SELECT 'a' REGEXP '^[a-d]';
-> 1

```

`REGEXP` [1193] and `RLIKE` [1193] use the character set and collations of the arguments when deciding the type of a character and performing the comparison. If the arguments have different character sets or collations, coercibility rules apply as described in [Section 10.1.7.5, “Collation of Expressions”](#).



### Warning

The `REGEXP` [1193] and `RLIKE` [1193] operators work in byte-wise fashion, so they are not multi-byte safe and may produce unexpected results with multi-byte character sets. In addition, these operators compare characters by their byte values and accented characters may not compare as equal even if a given collation treats them as equal.

A regular expression describes a set of strings. The simplest regular expression is one that has no special characters in it. For example, the regular expression `hello` matches `hello` and nothing else.

Nontrivial regular expressions use certain special constructs so that they can match more than one string. For example, the regular expression `hello|word` matches either the string `hello` or the string `word`.

As a more complex example, the regular expression `B[an]*s` matches any of the strings `Bananas`, `Baaaaas`, `Bs`, and any other string starting with a `B`, ending with an `s`, and containing any number of `a` or `n` characters in between.

A regular expression for the `REGEXP` [1193] operator may use any of the following special characters and constructs:

- `^`

Match the beginning of a string.

```

mysql> SELECT 'fo\ngo' REGEXP '^fo$';          -> 0
mysql> SELECT 'fofo' REGEXP '^fo';           -> 1

```

- `$`

Match the end of a string.

```

mysql> SELECT 'fo\ngo' REGEXP 'fo\ngo$';      -> 1
mysql> SELECT 'fo\ngo' REGEXP 'fo$';         -> 0

```

- `.`

Match any character (including carriage return and newline).

```

mysql> SELECT 'fofo' REGEXP '^f.*$';          -> 1
mysql> SELECT 'fo\r\nfo' REGEXP '^f.*$';     -> 1

```

- `a*`

Match any sequence of zero or more `a` characters.

```

mysql> SELECT 'Ban' REGEXP '^Ba*n';          -> 1
mysql> SELECT 'Baaan' REGEXP '^Ba*n';       -> 1
mysql> SELECT 'Bn' REGEXP '^Ba*n';          -> 1

```

- `a+`

Match any sequence of one or more `a` characters.

```

mysql> SELECT 'Ban' REGEXP '^Ba+n';          -> 1
mysql> SELECT 'Bn' REGEXP '^Ba+n';          -> 0

```

- `a?`

Match either zero or one `a` character.

```
mysql> SELECT 'Bn' REGEXP '^Ba?n';      -> 1
mysql> SELECT 'Ban' REGEXP '^Ba?n';    -> 1
mysql> SELECT 'Baan' REGEXP '^Ba?n';   -> 0
```

- `de|abc`

Match either of the sequences `de` or `abc`.

```
mysql> SELECT 'pi' REGEXP 'pi|apa';     -> 1
mysql> SELECT 'axe' REGEXP 'pi|apa';    -> 0
mysql> SELECT 'apa' REGEXP 'pi|apa';    -> 1
mysql> SELECT 'apa' REGEXP '^pi|apa$';  -> 1
mysql> SELECT 'pi' REGEXP '^pi|apa$';   -> 1
mysql> SELECT 'pix' REGEXP '^pi|apa$';  -> 0
```

- `(abc)*`

Match zero or more instances of the sequence `abc`.

```
mysql> SELECT 'pi' REGEXP '^pi*$';      -> 1
mysql> SELECT 'pip' REGEXP '^pi*$';     -> 0
mysql> SELECT 'pipi' REGEXP '^pi*$';    -> 1
```

- `{1}, {2,3}`

`{n}` or `{m,n}` notation provides a more general way of writing regular expressions that match many occurrences of the previous atom (or “piece”) of the pattern. `m` and `n` are integers.

- `a*`

Can be written as `a{0,}`.

- `a+`

Can be written as `a{1,}`.

- `a?`

Can be written as `a{0,1}`.

To be more precise, `a{n}` matches exactly `n` instances of `a`. `a{n,}` matches `n` or more instances of `a`. `a{m,n}` matches `m` through `n` instances of `a`, inclusive.

`m` and `n` must be in the range from 0 to `RE_DUP_MAX` (default 255), inclusive. If both `m` and `n` are given, `m` must be less than or equal to `n`.

```
mysql> SELECT 'abcde' REGEXP 'a[bcd]{2}e'; -> 0
mysql> SELECT 'abcde' REGEXP 'a[bcd]{3}e'; -> 1
mysql> SELECT 'abcde' REGEXP 'a[bcd]{1,10}e'; -> 1
```

- `[a-dX], [^a-dX]`

Matches any character that is (or is not, if `^` is used) either `a`, `b`, `c`, `d` or `X`. A `-` character between two other characters forms a range that matches all characters from the first character to the second. For example, `[0-9]` matches any decimal digit. To include a literal `]` character, it must immediately follow the opening bracket `[`. To include a literal `-` character, it must be written first or last. Any character that does not have a defined special meaning inside a `[ ]` pair matches only itself.

```
mysql> SELECT 'aXbc' REGEXP '[a-dXYZ]'; -> 1
mysql> SELECT 'aXbc' REGEXP '^[a-dXYZ]$'; -> 0
mysql> SELECT 'aXbc' REGEXP '^[a-dXYZ]+$'; -> 1
mysql> SELECT 'aXbc' REGEXP '^[^a-dXYZ]+$'; -> 0
```



```
mysql> SELECT 'gheis' REGEXP '^[^a-dXYZ]+$';      -> 1
mysql> SELECT 'gheisa' REGEXP '^[^a-dXYZ]+$';     -> 0
```

- [.characters.]

Within a bracket expression (written using [ and ]), matches the sequence of characters of that collating element. `characters` is either a single character or a character name like `newline`. The following table lists the permissible character names.

The following table shows the permissible character names and the characters that they match. For characters given as numeric values, the values are represented in octal.

Name	Character	Name	Character
NUL	0	SOH	001
STX	002	ETX	003
EOT	004	ENQ	005
ACK	006	BEL	007
alert	007	BS	010
backspace	'\b'	HT	011
tab	'\t'	LF	012
newline	'\n'	VT	013
vertical-tab	'\v'	FF	014
form-feed	'\f'	CR	015
carriage-return	'\r'	SO	016
SI	017	DLE	020
DC1	021	DC2	022
DC3	023	DC4	024
NAK	025	SYN	026
ETB	027	CAN	030
EM	031	SUB	032
ESC	033	IS4	034
FS	034	IS3	035
GS	035	IS2	036
RS	036	IS1	037
US	037	space	' '
exclamation-mark	'!'	quotation-mark	'\"'
number-sign	'#'	dollar-sign	'\$'
percent-sign	'%'	ampersand	'&'
apostrophe	'\''	left-parenthesis	'('
right-parenthesis	')'	asterisk	'*'
plus-sign	'+'	comma	','
hyphen	'-'	hyphen-minus	'-'
period	'.'	full-stop	'.'
slash	'/'	solidus	'/'
zero	'0'	one	'1'
two	'2'	three	'3'

Name	Character	Name	Character
four	'4'	five	'5'
six	'6'	seven	'7'
eight	'8'	nine	'9'
colon	':'	semicolon	';'
less-than-sign	'<'	equals-sign	'='
greater-than-sign	'>'	question-mark	'?'
commercial-at	'@'	left-square-bracket	'['
backslash	'\'	reverse-solidus	'\/'
right-square-bracket	']'	circumflex	'^'
circumflex-accent	'^'	underscore	'_'
low-line	'_'	grave-accent	'`'
left-brace	'{'	left-curly-bracket	'{'
vertical-line	' '	right-brace	'}'
right-curly-bracket	'}'	tilde	'~'
DEL	177		

```
mysql> SELECT '~' REGEXP '[[.~.]]';           -> 1
mysql> SELECT '~' REGEXP '[[.tilde.]]';      -> 1
```

- [=character\_class=]

Within a bracket expression (written using [ and ]), [=character\_class=] represents an equivalence class. It matches all characters with the same collation value, including itself. For example, if o and (+) are the members of an equivalence class, [[=o=]], [[=(+)=]], and [o(+)] are all synonymous. An equivalence class may not be used as an endpoint of a range.

- [:character\_class:]

Within a bracket expression (written using [ and ]), [:character\_class:] represents a character class that matches all characters belonging to that class. The following table lists the standard class names. These names stand for the character classes defined in the [ctype\(3\)](#) manual page. A particular locale may provide other class names. A character class may not be used as an endpoint of a range.

Character Class Name	Meaning
alnum	Alphanumeric characters
alpha	Alphabetic characters
blank	Whitespace characters
cntrl	Control characters
digit	Digit characters
graph	Graphic characters
lower	Lowercase alphabetic characters
print	Graphic or space characters

Character Class Name	Meaning
<code>punct</code>	Punctuation characters
<code>space</code>	Space, tab, newline, and carriage return
<code>upper</code>	Uppercase alphabetic characters
<code>xdigit</code>	Hexadecimal digit characters

```
mysql> SELECT 'justalnums' REGEXP '[:alnum:]+';      -> 1
mysql> SELECT '!' REGEXP '[:alnum:]+';             -> 0
```

- `[[:<:]]`, `[[:>:]]`

These markers stand for word boundaries. They match the beginning and end of words, respectively. A word is a sequence of word characters that is not preceded by or followed by word characters. A word character is an alphanumeric character in the `alnum` class or an underscore (`_`).

```
mysql> SELECT 'a word a' REGEXP '[[:<:]]word[[:>:]]'; -> 1
mysql> SELECT 'a xword a' REGEXP '[[:<:]]word[[:>:]]'; -> 0
```

To use a literal instance of a special character in a regular expression, precede it by two backslash (`\`) characters. The MySQL parser interprets one of the backslashes, and the regular expression library interprets the other. For example, to match the string `1+2` that contains the special `+` character, only the last of the following regular expressions is the correct one:

```
mysql> SELECT '1+2' REGEXP '1+2';                  -> 0
mysql> SELECT '1+2' REGEXP '1\+2';                 -> 0
mysql> SELECT '1+2' REGEXP '1\\+2';                 -> 1
```

## 12.6. Numeric Functions and Operators

Table 12.10. Numeric Functions and Operators

Name	Description
<code>ABS()</code> [1202]	Return the absolute value
<code>ACOS()</code> [1202]	Return the arc cosine
<code>ASIN()</code> [1202]	Return the arc sine
<code>ATAN2()</code> , <code>ATAN()</code> [1203]	Return the arc tangent of the two arguments
<code>ATAN()</code> [1202]	Return the arc tangent
<code>CEIL()</code> [1203]	Return the smallest integer value not less than the argument
<code>CEILING()</code> [1203]	Return the smallest integer value not less than the argument
<code>CONV()</code> [1203]	Convert numbers between different number bases
<code>COS()</code> [1203]	Return the cosine
<code>COT()</code> [1203]	Return the cotangent
<code>CRC32()</code> [1203]	Compute a cyclic redundancy check value
<code>DEGREES()</code> [1204]	Convert radians to degrees
<code>DIV</code> [1201]	Integer division
<code>/</code> [1200]	Division operator
<code>EXP()</code> [1204]	Raise to the power of
<code>FLOOR()</code> [1204]	Return the largest integer value not greater than the argument
<code>LN()</code> [1204]	Return the natural logarithm of the argument
<code>LOG10()</code> [1205]	Return the base-10 logarithm of the argument

Name	Description
<a href="#">LOG2 ( ) [1205]</a>	Return the base-2 logarithm of the argument
<a href="#">LOG ( ) [1204]</a>	Return the natural logarithm of the first argument
<a href="#">- [1200]</a>	Minus operator
<a href="#">MOD ( ) [1205]</a>	Return the remainder
<a href="#">% or MOD [1201]</a>	Modulo operator
<a href="#">PI ( ) [1205]</a>	Return the value of pi
<a href="#">+ [1200]</a>	Addition operator
<a href="#">POW ( ) [1206]</a>	Return the argument raised to the specified power
<a href="#">POWER ( ) [1206]</a>	Return the argument raised to the specified power
<a href="#">RADIANS ( ) [1206]</a>	Return argument converted to radians
<a href="#">RAND ( ) [1206]</a>	Return a random floating-point value
<a href="#">ROUND ( ) [1207]</a>	Round the argument
<a href="#">SIGN ( ) [1208]</a>	Return the sign of the argument
<a href="#">SIN ( ) [1208]</a>	Return the sine of the argument
<a href="#">SQRT ( ) [1208]</a>	Return the square root of the argument
<a href="#">TAN ( ) [1208]</a>	Return the tangent of the argument
<a href="#">* [1200]</a>	Multiplication operator
<a href="#">TRUNCATE ( ) [1209]</a>	Truncate to specified number of decimal places
<a href="#">- [1200]</a>	Change the sign of the argument

## 12.6.1. Arithmetic Operators

**Table 12.11. Arithmetic Operators**

Name	Description
<a href="#">DIV [1201]</a>	Integer division
<a href="#">/ [1200]</a>	Division operator
<a href="#">- [1200]</a>	Minus operator
<a href="#">% or MOD [1201]</a>	Modulo operator
<a href="#">+ [1200]</a>	Addition operator
<a href="#">* [1200]</a>	Multiplication operator
<a href="#">- [1200]</a>	Change the sign of the argument

The usual arithmetic operators are available. The result is determined according to the following rules:

- In the case of [- \[1200\]](#), [+ \[1200\]](#), and [\\* \[1200\]](#), the result is calculated with [BIGINT \[64-bit\]](#) precision if both operands are integers.
- If both operands are integers and any of them are unsigned, the result is an unsigned integer. For subtraction, if the [NO\\_UNSIGNED\\_SUBTRACTION \[605\]](#) SQL mode is enabled, the result is signed even if any operand is unsigned.
- If any of the operands of a [+ \[1200\]](#), [- \[1200\]](#), [/ \[1200\]](#), [\\* \[1200\]](#), [% \[1201\]](#) is a real or string value, the precision of the result is the precision of the operand with the maximum precision.
- In division performed with [/ \[1200\]](#), the scale of the result when using two exact-value operands is the scale of the first operand plus the value of the [div\\_precision\\_increment \[478\]](#) system variable (which is 4 by default). For example, the result of the expression `5.05 / 0.014` has a scale of six decimal places (`360.714286`).

These rules are applied for each operation, such that nested calculations imply the precision of each component. Hence,  $(14620 / 9432456) / (24250 / 9432456)$ , resolves first to  $(0.0014) / (0.0026)$ , with the final result having 8 decimal places ( $0.60288653$ ).

Because of these rules and the way they are applied, care should be taken to ensure that components and subcomponents of a calculation use the appropriate level of precision. See [Section 12.10, “Cast Functions and Operators”](#).

For information about handling of overflow in numeric expression evaluation, see [Section 11.2.6, “Out-of-Range and Overflow Handling”](#).

Arithmetic operators apply to numbers. For other types of values, alternative operations may be available. For example, to add date values, use `DATE_ADD()` [1213]; see [Section 12.7, “Date and Time Functions”](#).

- [+ \[1200\]](#)

Addition:

```
mysql> SELECT 3+5;
-> 8
```

- [- \[1200\]](#)

Subtraction:

```
mysql> SELECT 3-5;
-> -2
```

- [- \[1200\]](#)

Unary minus. This operator changes the sign of the operand.

```
mysql> SELECT - 2;
-> -2
```



#### Note

If this operator is used with a `BIGINT`, the return value is also a `BIGINT`. This means that you should avoid using `-` on integers that may have the value of  $-2^{63}$ .

- [\\* \[1200\]](#)

Multiplication:

```
mysql> SELECT 3*5;
-> 15
mysql> SELECT 18014398509481984*18014398509481984.0;
-> 324518553658426726783156020576256.0
mysql> SELECT 18014398509481984*18014398509481984;
-> out-of-range error
```

The last expression produces an error because the result of the integer multiplication exceeds the 64-bit range of `BIGINT` calculations. (See [Section 11.2, “Numeric Types”](#).)

- [/ \[1200\]](#)

Division:

```
mysql> SELECT 3/5;
-> 0.60
```

Division by zero produces a `NULL` result:

```
mysql> SELECT 102/(1-1);
```

```
-> NULL
```

A division is calculated with `BIGINT` arithmetic only if performed in a context where its result is converted to an integer.

- `DIV` [1201]

Integer division. Similar to `FLOOR()` [1204], but is safe with `BIGINT` values.

In MySQL 5.6, if either operand has a noninteger type, the operands are converted to `DECIMAL` and divided using `DECIMAL` arithmetic before converting the result to `BIGINT`. If the result exceeds `BIGINT` range, an error occurs.

```
mysql> SELECT 5 DIV 2;
-> 2
```

- `N % M` [1201], `N MOD M` [1201]

Modulo operation. Returns the remainder of `N` divided by `M`. For more information, see the description for the `MOD()` [1205] function in Section 12.6.2, “Mathematical Functions”.

## 12.6.2. Mathematical Functions

**Table 12.12. Mathematical Functions**

Name	Description
<code>ABS()</code> [1202]	Return the absolute value
<code>ACOS()</code> [1202]	Return the arc cosine
<code>ASIN()</code> [1202]	Return the arc sine
<code>ATAN2()</code> , <code>ATAN()</code> [1203]	Return the arc tangent of the two arguments
<code>ATAN()</code> [1202]	Return the arc tangent
<code>CEIL()</code> [1203]	Return the smallest integer value not less than the argument
<code>CEILING()</code> [1203]	Return the smallest integer value not less than the argument
<code>CONV()</code> [1203]	Convert numbers between different number bases
<code>COS()</code> [1203]	Return the cosine
<code>COT()</code> [1203]	Return the cotangent
<code>CRC32()</code> [1203]	Compute a cyclic redundancy check value
<code>DEGREES()</code> [1204]	Convert radians to degrees
<code>EXP()</code> [1204]	Raise to the power of
<code>FLOOR()</code> [1204]	Return the largest integer value not greater than the argument
<code>LN()</code> [1204]	Return the natural logarithm of the argument
<code>LOG10()</code> [1205]	Return the base-10 logarithm of the argument
<code>LOG2()</code> [1205]	Return the base-2 logarithm of the argument
<code>LOG()</code> [1204]	Return the natural logarithm of the first argument
<code>MOD()</code> [1205]	Return the remainder
<code>PI()</code> [1205]	Return the value of pi
<code>POW()</code> [1206]	Return the argument raised to the specified power
<code>POWER()</code> [1206]	Return the argument raised to the specified power
<code>RADIANS()</code> [1206]	Return argument converted to radians
<code>RAND()</code> [1206]	Return a random floating-point value

Name	Description
<a href="#">ROUND() [1207]</a>	Round the argument
<a href="#">SIGN() [1208]</a>	Return the sign of the argument
<a href="#">SIN() [1208]</a>	Return the sine of the argument
<a href="#">SQRT() [1208]</a>	Return the square root of the argument
<a href="#">TAN() [1208]</a>	Return the tangent of the argument
<a href="#">TRUNCATE() [1209]</a>	Truncate to specified number of decimal places

All mathematical functions return `NULL` in the event of an error.

- [ABS\(X\) \[1202\]](#)

Returns the absolute value of *X*.

```
mysql> SELECT ABS(2);
      -> 2
mysql> SELECT ABS(-32);
      -> 32
```

This function is safe to use with `BIGINT` values.

- [ACOS\(X\) \[1202\]](#)

Returns the arc cosine of *X*, that is, the value whose cosine is *X*. Returns `NULL` if *X* is not in the range `-1` to `1`.

```
mysql> SELECT ACOS(1);
      -> 0
mysql> SELECT ACOS(1.0001);
      -> NULL
mysql> SELECT ACOS(0);
      -> 1.5707963267949
```

- [ASIN\(X\) \[1202\]](#)

Returns the arc sine of *X*, that is, the value whose sine is *X*. Returns `NULL` if *X* is not in the range `-1` to `1`.

```
mysql> SELECT ASIN(0.2);
      -> 0.20135792079033
mysql> SELECT ASIN('foo');

+-----+
| ASIN('foo') |
+-----+
|          0 |
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;

+-----+-----+-----+
| Level | Code | Message
+-----+-----+-----+
| Warning | 1292 | Truncated incorrect DOUBLE value: 'foo'
+-----+-----+-----+
```

- [ATAN\(X\) \[1202\]](#)

Returns the arc tangent of *X*, that is, the value whose tangent is *X*.

```
mysql> SELECT ATAN(2);
      -> 1.1071487177941
mysql> SELECT ATAN(-2);
      -> -1.1071487177941
```



- [ATAN\(\*Y\*,\*X\*\) \[1203\]](#), [ATAN2\(\*Y\*,\*X\*\) \[1203\]](#)

Returns the arc tangent of the two variables *X* and *Y*. It is similar to calculating the arc tangent of *Y* / *X*, except that the signs of both arguments are used to determine the quadrant of the result.

```
mysql> SELECT ATAN(-2,2);
      -> -0.78539816339745
mysql> SELECT ATAN2(PI(),0);
      -> 1.5707963267949
```

- [CEIL\(\*X\*\) \[1203\]](#)

[CEIL\(\)](#) [1203] is a synonym for [CEILING\(\)](#) [1203].

- [CEILING\(\*X\*\) \[1203\]](#)

Returns the smallest integer value not less than *X*.

```
mysql> SELECT CEILING(1.23);
      -> 2
mysql> SELECT CEILING(-1.23);
      -> -1
```

For exact-value numeric arguments, the return value has an exact-value numeric type. For string or floating-point arguments, the return value has a floating-point type.

- [CONV\(\*N\*,\*from\\_base\*,\*to\\_base\*\) \[1203\]](#)

Converts numbers between different number bases. Returns a string representation of the number *N*, converted from base *from\_base* to base *to\_base*. Returns *NULL* if any argument is *NULL*. The argument *N* is interpreted as an integer, but may be specified as an integer or a string. The minimum base is 2 and the maximum base is 36. If *to\_base* is a negative number, *N* is regarded as a signed number. Otherwise, *N* is treated as unsigned. [CONV\(\)](#) [1203] works with 64-bit precision.

```
mysql> SELECT CONV('a',16,2);
      -> '1010'
mysql> SELECT CONV('6E',18,8);
      -> '172'
mysql> SELECT CONV(-17,10,-18);
      -> '-H'
mysql> SELECT CONV(10+'10'+ '10'+0xa,10,10);
      -> '40'
```

- [COS\(\*X\*\) \[1203\]](#)

Returns the cosine of *X*, where *X* is given in radians.

```
mysql> SELECT COS(PI());
      -> -1
```

- [COT\(\*X\*\) \[1203\]](#)

Returns the cotangent of *X*.

```
mysql> SELECT COT(12);
      -> -1.5726734063977
mysql> SELECT COT(0);
      -> NULL
```

- [CRC32\(\*expr\*\) \[1203\]](#)

Computes a cyclic redundancy check value and returns a 32-bit unsigned value. The result is *NULL* if the argument is *NULL*. The argument is expected to be a string and (if possible) is treated as one if it is not.

```
mysql> SELECT CRC32('MySQL');
      -> 3259397556
```

```
mysql> SELECT CRC32('mysql');
-> 2501908538
```

- [DEGREES\(\*X\*\) \[1204\]](#)

Returns the argument *X*, converted from radians to degrees.

```
mysql> SELECT DEGREES(PI());
-> 180
mysql> SELECT DEGREES(PI() / 2);
-> 90
```

- [EXP\(\*X\*\) \[1204\]](#)

Returns the value of *e* (the base of natural logarithms) raised to the power of *X*. The inverse of this function is [LOG\(\*Y\*\) \[1204\]](#) (using a single argument only) or [LN\(\*Y\*\) \[1204\]](#).

```
mysql> SELECT EXP(2);
-> 7.3890560989307
mysql> SELECT EXP(-2);
-> 0.13533528323661
mysql> SELECT EXP(0);
-> 1
```

- [FLOOR\(\*X\*\) \[1204\]](#)

Returns the largest integer value not greater than *X*.

```
mysql> SELECT FLOOR(1.23);
-> 1
mysql> SELECT FLOOR(-1.23);
-> -2
```

For exact-value numeric arguments, the return value has an exact-value numeric type. For string or floating-point arguments, the return value has a floating-point type.

- [FORMAT\(\*X\*,\*D\*\) \[1180\]](#)

Formats the number *X* to a format like '#,###,###.##', rounded to *D* decimal places, and returns the result as a string. For details, see [Section 12.5, "String Functions"](#).

- [HEX\(\*N\\_or\\_S\*\) \[1180\]](#)

This function can be used to obtain a hexadecimal representation of a decimal number or a string; the manner in which it does so varies according to the argument's type. See this function's description in [Section 12.5, "String Functions"](#), for details.

- [LN\(\*X\*\) \[1204\]](#)

Returns the natural logarithm of *X*; that is, the base-*e* logarithm of *X*. If *X* is less than or equal to 0, then `NULL` is returned.

```
mysql> SELECT LN(2);
-> 0.69314718055995
mysql> SELECT LN(-2);
-> NULL
```

This function is synonymous with [LOG\(\*X\*\) \[1204\]](#). The inverse of this function is the [EXP\(\*Y\*\) \[1204\]](#) function.

- [LOG\(\*X\*\) \[1204\]](#), [LOG\(\*B\*,\*X\*\) \[1204\]](#)

If called with one parameter, this function returns the natural logarithm of *X*. If *X* is less than or equal to 0, then `NULL` is returned.

The inverse of this function (when called with a single argument) is the [EXP\(\*Y\*\) \[1204\]](#) function.

```
mysql> SELECT LOG(2);
      -> 0.69314718055995
mysql> SELECT LOG(-2);
      -> NULL
```

If called with two parameters, this function returns the logarithm of  $X$  to the base  $B$ . If  $X$  is less than or equal to 0, or if  $B$  is less than or equal to 1, then `NULL` is returned.

```
mysql> SELECT LOG(2,65536);
      -> 16
mysql> SELECT LOG(10,100);
      -> 2
mysql> SELECT LOG(1,100);
      -> NULL
```

`LOG(B, X)` [1204] is equivalent to `LOG(X) / LOG(B)` [1204].

- `LOG2(X)` [1205]

Returns the base-2 logarithm of  $X$ .

```
mysql> SELECT LOG2(65536);
      -> 16
mysql> SELECT LOG2(-100);
      -> NULL
```

`LOG2()` [1205] is useful for finding out how many bits a number requires for storage. This function is equivalent to the expression `LOG(X) / LOG(2)` [1204].

- `LOG10(X)` [1205]

Returns the base-10 logarithm of  $X$ .

```
mysql> SELECT LOG10(2);
      -> 0.30102999566398
mysql> SELECT LOG10(100);
      -> 2
mysql> SELECT LOG10(-100);
      -> NULL
```

`LOG10(X)` [1205] is equivalent to `LOG(10, X)` [1204].

- `MOD(N, M)` [1205], `N % M` [1201], `N MOD M` [1201]

Modulo operation. Returns the remainder of  $N$  divided by  $M$ .

```
mysql> SELECT MOD(234, 10);
      -> 4
mysql> SELECT 253 % 7;
      -> 1
mysql> SELECT MOD(29, 9);
      -> 2
mysql> SELECT 29 MOD 9;
      -> 2
```

This function is safe to use with `BIGINT` values.

`MOD()` [1205] also works on values that have a fractional part and returns the exact remainder after division:

```
mysql> SELECT MOD(34.5, 3);
      -> 1.5
```

`MOD(N, 0)` [1205] returns `NULL`.

- `PI()` [1205]

Returns the value of  $\pi$  (pi). The default number of decimal places displayed is seven, but MySQL uses the full double-precision value internally.

```
mysql> SELECT PI();
-> 3.141593
mysql> SELECT PI()+0.00000000000000000000;
-> 3.141592653589793116
```

- [POW\(X,Y\) \[1206\]](#)

Returns the value of  $X$  raised to the power of  $Y$ .

```
mysql> SELECT POW(2,2);
-> 4
mysql> SELECT POW(2,-2);
-> 0.25
```

- [POWER\(X,Y\) \[1206\]](#)

This is a synonym for [POW\(\)](#) [1206].

- [RADIANS\(X\) \[1206\]](#)

Returns the argument  $X$ , converted from degrees to radians. (Note that  $\pi$  radians equals 180 degrees.)

```
mysql> SELECT RADIANS(90);
-> 1.5707963267949
```

- [RAND\(\)](#) [1206], [RAND\(N\)](#) [1206]

Returns a random floating-point value  $v$  in the range  $0 \leq v < 1.0$ . If a constant integer argument  $N$  is specified, it is used as the seed value, which produces a repeatable sequence of column values. In the following example, note that the sequences of values produced by [RAND\(3\)](#) is the same both places where it occurs.

```
mysql> CREATE TABLE t (i INT);
Query OK, 0 rows affected (0.42 sec)

mysql> INSERT INTO t VALUES(1),(2),(3);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT i, RAND() FROM t;
+-----+-----+
| i    | RAND()          |
+-----+-----+
| 1    | 0.61914388706828 |
| 2    | 0.93845168309142 |
| 3    | 0.83482678498591 |
+-----+-----+
3 rows in set (0.00 sec)

mysql> SELECT i, RAND(3) FROM t;
+-----+-----+
| i    | RAND(3)         |
+-----+-----+
| 1    | 0.90576975597606 |
| 2    | 0.37307905813035 |
| 3    | 0.14808605345719 |
+-----+-----+
3 rows in set (0.00 sec)

mysql> SELECT i, RAND() FROM t;
+-----+-----+
| i    | RAND()          |
+-----+-----+
```

```

+-----+-----+
| 1 | 0.35877890638893 |
| 2 | 0.28941420772058 |
| 3 | 0.37073435016976 |
+-----+-----+
3 rows in set (0.00 sec)

mysql> SELECT i, RAND(3) FROM t;
+-----+-----+
| i | RAND(3) |
+-----+-----+
| 1 | 0.90576975597606 |
| 2 | 0.37307905813035 |
| 3 | 0.14808605345719 |
+-----+-----+
3 rows in set (0.01 sec)
    
```

With a constant initializer, the seed is initialized once when the statement is compiled, prior to execution. If a nonconstant initializer (such as a column name) is used as the argument, the seed is initialized with the value for each invocation of `RAND()` [1206]. (One implication of this is that for equal argument values, `RAND()` [1206] will return the same value each time.)

To obtain a random integer  $R$  in the range  $i \leq R < j$ , use the expression `FLOOR(i + RAND() * (j [1204] - i))`. For example, to obtain a random integer in the range the range  $7 \leq R < 12$ , you could use the following statement:

```
SELECT FLOOR(7 + (RAND() * 5));
```

`RAND()` [1206] in a `WHERE` clause is re-evaluated every time the `WHERE` is executed.

You cannot use a column with `RAND()` [1206] values in an `ORDER BY` clause, because `ORDER BY` would evaluate the column multiple times. However, you can retrieve rows in random order like this:

```
mysql> SELECT * FROM tbl_name ORDER BY RAND();
```

`ORDER BY RAND()` combined with `LIMIT` is useful for selecting a random sample from a set of rows:

```
mysql> SELECT * FROM table1, table2 WHERE a=b AND c<d -> ORDER BY RAND() LIMIT 1000;
```

`RAND()` [1206] is not meant to be a perfect random generator. It is a fast way to generate random numbers on demand that is portable between platforms for the same MySQL version.

This function is unsafe for statement-based replication. A warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #49222)

- `ROUND(X)` [1207], `ROUND(X,D)` [1207]

Rounds the argument  $X$  to  $D$  decimal places. The rounding algorithm depends on the data type of  $X$ .  $D$  defaults to 0 if not specified.  $D$  can be negative to cause  $D$  digits left of the decimal point of the value  $X$  to become zero.

```

mysql> SELECT ROUND(-1.23);
-> -1
mysql> SELECT ROUND(-1.58);
-> -2
mysql> SELECT ROUND(1.58);
-> 2
mysql> SELECT ROUND(1.298, 1);
-> 1.3
mysql> SELECT ROUND(1.298, 0);
-> 1
mysql> SELECT ROUND(23.298, -1);
-> 20
    
```

The return type is the same type as that of the first argument (assuming that it is integer, double, or decimal). This means that for an integer argument, the result is an integer (no decimal places):

```
mysql> SELECT ROUND(150.000,2), ROUND(150,2);
+-----+-----+
| ROUND(150.000,2) | ROUND(150,2) |
+-----+-----+
|          150.00 |          150 |
+-----+-----+
```

`ROUND()` [1207] uses the following rules depending on the type of the first argument:

- For exact-value numbers, `ROUND()` [1207] uses the “round half away from zero” or “round toward nearest” rule: A value with a fractional part of .5 or greater is rounded up to the next integer if positive or down to the next integer if negative. (In other words, it is rounded away from zero.) A value with a fractional part less than .5 is rounded down to the next integer if positive or up to the next integer if negative.
- For approximate-value numbers, the result depends on the C library. On many systems, this means that `ROUND()` [1207] uses the “round to nearest even” rule: A value with any fractional part is rounded to the nearest even integer.

The following example shows how rounding differs for exact and approximate values:

```
mysql> SELECT ROUND(2.5), ROUND(25E-1);
+-----+-----+
| ROUND(2.5) | ROUND(25E-1) |
+-----+-----+
| 3          | 2            |
+-----+-----+
```

For more information, see [Section 12.19, “Precision Math”](#).

- `SIGN(X)` [1208]

Returns the sign of the argument as `-1`, `0`, or `1`, depending on whether `X` is negative, zero, or positive.

```
mysql> SELECT SIGN(-32);
-> -1
mysql> SELECT SIGN(0);
-> 0
mysql> SELECT SIGN(234);
-> 1
```

- `SIN(X)` [1208]

Returns the sine of `X`, where `X` is given in radians.

```
mysql> SELECT SIN(PI());
-> 1.2246063538224e-16
mysql> SELECT ROUND(SIN(PI()));
-> 0
```

- `SQRT(X)` [1208]

Returns the square root of a nonnegative number `X`.

```
mysql> SELECT SQRT(4);
-> 2
mysql> SELECT SQRT(20);
-> 4.4721359549996
mysql> SELECT SQRT(-16);
-> NULL
```

- `TAN(X)` [1208]

Returns the tangent of `X`, where `X` is given in radians.

```
mysql> SELECT TAN(PI());
-> -1.2246063538224e-16
mysql> SELECT TAN(PI()+1);
-> 1.5574077246549
```

- `TRUNCATE(X,D)` [1209]

Returns the number  $X$ , truncated to  $D$  decimal places. If  $D$  is 0, the result has no decimal point or fractional part.  $D$  can be negative to cause  $D$  digits left of the decimal point of the value  $X$  to become zero.

```
mysql> SELECT TRUNCATE(1.223,1);
-> 1.2
mysql> SELECT TRUNCATE(1.999,1);
-> 1.9
mysql> SELECT TRUNCATE(1.999,0);
-> 1
mysql> SELECT TRUNCATE(-1.999,1);
-> -1.9
mysql> SELECT TRUNCATE(122,-2);
-> 100
mysql> SELECT TRUNCATE(10.28*100,0);
-> 1028
```

All numbers are rounded toward zero.

## 12.7. Date and Time Functions

This section describes the functions that can be used to manipulate temporal values. See [Section 11.3, “Date and Time Types”](#), for a description of the range of values each date and time type has and the valid formats in which values may be specified.

**Table 12.13. Date/Time Functions**

Name	Description
<code>ADDDATE()</code> [1211]	Add time values (intervals) to a date value
<code>ADDTIME()</code> [1212]	Add time
<code>CONVERT_TZ()</code> [1212]	Convert from one timezone to another
<code>CURDATE()</code> [1212]	Return the current date
<code>CURRENT_DATE()</code> , <code>CURRENT_DATE</code> [1212]	Synonyms for <code>CURDATE()</code>
<code>CURRENT_TIME()</code> , <code>CURRENT_TIME</code> [1212]	Synonyms for <code>CURTIME()</code>
<code>CURRENT_TIMESTAMP()</code> , <code>CURRENT_TIMESTAMP</code> [1213]	Synonyms for <code>NOW()</code>
<code>CURTIME()</code> [1213]	Return the current time
<code>DATE_ADD()</code> [1213]	Add time values (intervals) to a date value
<code>DATE_FORMAT()</code> [1215]	Format date as specified
<code>DATE_SUB()</code> [1217]	Subtract a time value (interval) from a date
<code>DATE()</code> [1213]	Extract the date part of a date or datetime expression
<code>DATEDIFF()</code> [1213]	Subtract two dates
<code>DAY()</code> [1217]	Synonym for <code>DAYOFMONTH()</code>
<code>DAYNAME()</code> [1217]	Return the name of the weekday
<code>DAYOFMONTH()</code> [1217]	Return the day of the month (0-31)
<code>DAYOFWEEK()</code> [1217]	Return the weekday index of the argument

Name	Description
DAYOFYEAR ( ) [1217]	Return the day of the year (1-366)
EXTRACT ( ) [1217]	Extract part of a date
FROM_DAYS ( ) [1218]	Convert a day number to a date
FROM_UNIXTIME ( ) [1218]	Format UNIX timestamp as a date
GET_FORMAT ( ) [1218]	Return a date format string
HOUR ( ) [1219]	Extract the hour
LAST_DAY [1219]	Return the last day of the month for the argument
LOCALTIME ( ), LOCALTIME [1219]	Synonym for NOW()
LOCALTIMESTAMP, LOCALTIMESTAMP ( ) [1219]	Synonym for NOW()
MAKEDATE ( ) [1219]	Create a date from the year and day of year
MAKETIME [1219]	MAKETIME()
MICROSECOND ( ) [1220]	Return the microseconds from argument
MINUTE ( ) [1220]	Return the minute from the argument
MONTH ( ) [1220]	Return the month from the date passed
MONTHNAME ( ) [1220]	Return the name of the month
NOW ( ) [1220]	Return the current date and time
PERIOD_ADD ( ) [1221]	Add a period to a year-month
PERIOD_DIFF ( ) [1221]	Return the number of months between periods
QUARTER ( ) [1221]	Return the quarter from a date argument
SEC_TO_TIME ( ) [1221]	Converts seconds to 'HH:MM:SS' format
SECOND ( ) [1221]	Return the second (0-59)
STR_TO_DATE ( ) [1221]	Convert a string to a date
SUBDATE ( ) [1222]	A synonym for DATE_SUB() when invoked with three arguments
SUBTIME ( ) [1223]	Subtract times
SYSDATE ( ) [1223]	Return the time at which the function executes
TIME_FORMAT ( ) [1224]	Format as time
TIME_TO_SEC ( ) [1225]	Return the argument converted to seconds
TIME ( ) [1223]	Extract the time portion of the expression passed
TIMEDIFF ( ) [1223]	Subtract time
TIMESTAMP ( ) [1224]	With a single argument, this function returns the date or datetime expression; with two arguments, the sum of the arguments
TIMESTAMPADD ( ) [1224]	Add an interval to a datetime expression
TIMESTAMPDIFF ( ) [1224]	Subtract an interval from a datetime expression
TO_DAYS ( ) [1225]	Return the date argument converted to days
TO_SECONDS ( ) [1226]	Return the date or datetime argument converted to seconds since Year 0
UNIX_TIMESTAMP ( ) [1226]	Return a UNIX timestamp
UTC_DATE ( ) [1227]	Return the current UTC date
UTC_TIME ( ) [1227]	Return the current UTC time



Name	Description
<a href="#">UTC_TIMESTAMP() [1227]</a>	Return the current UTC date and time
<a href="#">WEEK() [1228]</a>	Return the week number
<a href="#">WEEKDAY() [1228]</a>	Return the weekday index
<a href="#">WEEKOFYEAR() [1229]</a>	Return the calendar week of the date (0-53)
<a href="#">YEAR() [1229]</a>	Return the year
<a href="#">YEARWEEK() [1229]</a>	Return the year and week

Here is an example that uses date functions. The following query selects all rows with a `date_col` value from within the last 30 days:

```
mysql> SELECT something FROM tbl_name
  -> WHERE DATE_SUB(CURDATE(),INTERVAL 30 DAY) <= date_col;
```

The query also selects rows with dates that lie in the future.

Functions that expect date values usually accept datetime values and ignore the time part. Functions that expect time values usually accept datetime values and ignore the date part.

Functions that return the current date or time each are evaluated only once per query at the start of query execution. This means that multiple references to a function such as [NOW\(\) \[1220\]](#) within a single query always produce the same result. (For our purposes, a single query also includes a call to a stored program (stored routine, trigger, or event) and all subprograms called by that program.) This principle also applies to [CURDATE\(\) \[1212\]](#), [CURTIME\(\) \[1213\]](#), [UTC\\_DATE\(\) \[1227\]](#), [UTC\\_TIME\(\) \[1227\]](#), [UTC\\_TIMESTAMP\(\) \[1227\]](#), and to any of their synonyms.

The [CURRENT\\_TIMESTAMP\(\) \[1213\]](#), [CURRENT\\_TIME\(\) \[1212\]](#), [CURRENT\\_DATE\(\) \[1212\]](#), and [FROM\\_UNIXTIME\(\) \[1218\]](#) functions return values in the connection's current time zone, which is available as the value of the [time\\_zone \[557\]](#) system variable. In addition, [UNIX\\_TIMESTAMP\(\) \[1226\]](#) assumes that its argument is a datetime value in the current time zone. See [Section 10.6, "MySQL Server Time Zone Support"](#).

Some date functions can be used with "zero" dates or incomplete dates such as `'2001-11-00'`, whereas others cannot. Functions that extract parts of dates typically work with incomplete dates and thus can return 0 when you might otherwise expect a nonzero value. For example:

```
mysql> SELECT DAYOFMONTH('2001-11-00'), MONTH('2005-00-00');
  -> 0, 0
```

Other functions expect complete dates and return `NULL` for incomplete dates. These include functions that perform date arithmetic or that map parts of dates to names. For example:

```
mysql> SELECT DATE_ADD('2006-05-00',INTERVAL 1 DAY);
  -> NULL
mysql> SELECT DAYNAME('2006-05-00');
  -> NULL
```

As of MySQL 5.6.4, several functions are more strict when passed a [DATE\(\) \[1213\]](#) function value as their argument and reject incomplete dates with a day part of zero. These functions are affected: [CONVERT\\_TZ\(\) \[1212\]](#), [DATE\\_ADD\(\) \[1213\]](#), [DATE\\_SUB\(\) \[1217\]](#), [DAYOFYEAR\(\) \[1217\]](#), [LAST\\_DAY\(\) \[1219\]](#), [TIMESTAMPDIFF\(\) \[1224\]](#), [TO\\_DAYS\(\) \[1225\]](#), [TO\\_SECONDS\(\) \[1226\]](#), [WEEK\(\) \[1228\]](#), [WEEKDAY\(\) \[1228\]](#), [WEEKOFYEAR\(\) \[1229\]](#), [YEARWEEK\(\) \[1229\]](#). This restriction was relaxed for [LAST\\_DAY\(\) \[1219\]](#) in 5.6.5 to permit a day part of zero.

MySQL 5.6.4 and up supports fractional seconds for [TIME](#), [DATETIME](#), and [TIMESTAMP](#) values, with up to microsecond precision. Functions that take temporal arguments accept values with fractional seconds. Return values from temporal functions include fractional seconds as appropriate.

- [ADDDATE\(date,INTERVAL expr unit\) \[1211\]](#), [ADDDATE\(expr,days\) \[1211\]](#)

When invoked with the `INTERVAL` form of the second argument, `ADDDATE()` [1211] is a synonym for `DATE_ADD()` [1213]. The related function `SUBDATE()` [1222] is a synonym for `DATE_SUB()` [1217]. For information on the `INTERVAL unit` argument, see the discussion for `DATE_ADD()` [1213].

```
mysql> SELECT DATE_ADD('2008-01-02', INTERVAL 31 DAY);
-> '2008-02-02'
mysql> SELECT ADDDATE('2008-01-02', INTERVAL 31 DAY);
-> '2008-02-02'
```

When invoked with the `days` form of the second argument, MySQL treats it as an integer number of days to be added to `expr`.

```
mysql> SELECT ADDDATE('2008-01-02', 31);
-> '2008-02-02'
```

- `ADDTIME(expr1, expr2)` [1212]

`ADDTIME()` [1212] adds `expr2` to `expr1` and returns the result. `expr1` is a time or datetime expression, and `expr2` is a time expression.

```
mysql> SELECT ADDTIME('2007-12-31 23:59:59.999999', '1 1:1:1.000002');
-> '2008-01-02 01:01:01.000001'
mysql> SELECT ADDTIME('01:00:00.999999', '02:00:00.999998');
-> '03:00:01.999997'
```

- `CONVERT_TZ(dt, from_tz, to_tz)` [1212]

`CONVERT_TZ()` [1212] converts a datetime value `dt` from the time zone given by `from_tz` to the time zone given by `to_tz` and returns the resulting value. Time zones are specified as described in Section 10.6, “MySQL Server Time Zone Support”. This function returns `NULL` if the arguments are invalid.

If the value falls out of the supported range of the `TIMESTAMP` type when converted from `from_tz` to UTC, no conversion occurs. The `TIMESTAMP` range is described in Section 11.1.2, “Date and Time Type Overview”.

```
mysql> SELECT CONVERT_TZ('2004-01-01 12:00:00', 'GMT', 'MET');
-> '2004-01-01 13:00:00'
mysql> SELECT CONVERT_TZ('2004-01-01 12:00:00', '+00:00', '+10:00');
-> '2004-01-01 22:00:00'
```



#### Note

To use named time zones such as `'MET'` or `'Europe/Moscow'`, the time zone tables must be properly set up. See Section 10.6, “MySQL Server Time Zone Support”, for instructions.

- `CURDATE()` [1212]

Returns the current date as a value in `'YYYY-MM-DD'` or `YYYYMMDD` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT CURDATE();
-> '2008-06-13'
mysql> SELECT CURDATE() + 0;
-> 20080613
```

- `CURRENT_DATE` [1212], `CURRENT_DATE()` [1212]

`CURRENT_DATE` [1212] and `CURRENT_DATE()` [1212] are synonyms for `CURDATE()` [1212].

- `CURRENT_TIME` [1212], `CURRENT_TIME([fsp])` [1212]

`CURRENT_TIME` [1212] and `CURRENT_TIME()` [1212] are synonyms for `CURTIME()` [1213].

- `CURRENT_TIMESTAMP` [1213], `CURRENT_TIMESTAMP([fsp])` [1213]

`CURRENT_TIMESTAMP` [1213] and `CURRENT_TIMESTAMP()` [1213] are synonyms for `NOW()` [1220].

- `CURTIME([fsp])` [1213]

Returns the current time as a value in 'HH:MM:SS' or HHMMSS format, depending on whether the function is used in a string or numeric context. The value is expressed in the current time zone.

As of MySQL 5.6.4, if the *fsp* argument is given to specify a fractional seconds precision from 0 to 6, the return value includes a fractional seconds part of that many digits. Before 5.6.4, any argument is ignored.

```
mysql> SELECT CURTIME();
-> '23:50:26'
mysql> SELECT CURTIME() + 0;
-> 235026.000000
```

- `DATE(expr)` [1213]

Extracts the date part of the date or datetime expression *expr*.

```
mysql> SELECT DATE('2003-12-31 01:02:03');
-> '2003-12-31'
```

- `DATEDIFF(expr1,expr2)` [1213]

`DATEDIFF()` [1213] returns *expr1* – *expr2* expressed as a value in days from one date to the other. *expr1* and *expr2* are date or date-and-time expressions. Only the date parts of the values are used in the calculation.

```
mysql> SELECT DATEDIFF('2007-12-31 23:59:59','2007-12-30');
-> 1
mysql> SELECT DATEDIFF('2010-11-30 23:59:59','2010-12-31');
-> -31
```

- `DATE_ADD(date,INTERVAL expr unit)` [1213], `DATE_SUB(date,INTERVAL expr unit)` [1217]

These functions perform date arithmetic. The *date* argument specifies the starting date or datetime value. *expr* is an expression specifying the interval value to be added or subtracted from the starting date. *expr* is a string; it may start with a “-” for negative intervals. *unit* is a keyword indicating the units in which the expression should be interpreted.

The `INTERVAL` keyword and the *unit* specifier are not case sensitive.

The following table shows the expected form of the *expr* argument for each *unit* value.

<i>unit</i> Value	Expected <i>expr</i> Format
MICROSECOND	MICROSECONDS
SECOND	SECONDS
MINUTE	MINUTES
HOUR	HOURS
DAY	DAYS
WEEK	WEEKS
MONTH	MONTHS

<i>unit</i> Value	Expected <i>expr</i> Format
QUARTER	QUARTERS
YEAR	YEARS
SECOND_MICROSECOND	'SECONDS.MICROSECONDS'
MINUTE_MICROSECOND	'MINUTES:SECONDS.MICROSECONDS'
MINUTE_SECOND	'MINUTES:SECONDS'
HOURL_MICROSECOND	'HOURS:MINUTES:SECONDS.MICROSECONDS'
HOURL_SECOND	'HOURS:MINUTES:SECONDS'
HOURL_MINUTE	'HOURS:MINUTES'
DAY_MICROSECOND	'DAYS HOURS:MINUTES:SECONDS.MICROSECONDS'
DAY_SECOND	'DAYS HOURS:MINUTES:SECONDS'
DAY_MINUTE	'DAYS HOURS:MINUTES'
DAY_HOUR	'DAYS HOURS'
YEAR_MONTH	'YEARS-MONTHS'

The return value depends on the arguments:

- `DATETIME` if the first argument is a `DATETIME` (or `TIMESTAMP`) value, or if the first argument is a `DATE` and the *unit* value uses `HOURS`, `MINUTES`, or `SECONDS`.
- String otherwise.

To ensure that the result is `DATETIME`, you can use `CAST( )` [1245] to convert the first argument to `DATETIME`.

MySQL permits any punctuation delimiter in the *expr* format. Those shown in the table are the suggested delimiters. If the *date* argument is a `DATE` value and your calculations involve only `YEAR`, `MONTH`, and `DAY` parts (that is, no time parts), the result is a `DATE` value. Otherwise, the result is a `DATETIME` value.

Date arithmetic also can be performed using `INTERVAL` together with the `+` [1200] or `-` [1200] operator:

```
date + INTERVAL expr unit
date - INTERVAL expr unit
```

`INTERVAL expr unit` is permitted on either side of the `+` [1200] operator if the expression on the other side is a date or datetime value. For the `-` [1200] operator, `INTERVAL expr unit` is permitted only on the right side, because it makes no sense to subtract a date or datetime value from an interval.

```
mysql> SELECT '2008-12-31 23:59:59' + INTERVAL 1 SECOND;
-> '2009-01-01 00:00:00'
mysql> SELECT INTERVAL 1 DAY + '2008-12-31';
-> '2009-01-01'
mysql> SELECT '2005-01-01' - INTERVAL 1 SECOND;
-> '2004-12-31 23:59:59'
mysql> SELECT DATE_ADD('2000-12-31 23:59:59',
-> INTERVAL 1 SECOND);
-> '2001-01-01 00:00:00'
mysql> SELECT DATE_ADD('2010-12-31 23:59:59',
-> INTERVAL 1 DAY);
-> '2011-01-01 23:59:59'
mysql> SELECT DATE_ADD('2100-12-31 23:59:59',
-> INTERVAL '1:1' MINUTE_SECOND);
-> '2101-01-01 00:01:00'
mysql> SELECT DATE_SUB('2005-01-01 00:00:00',
```

```

-> INTERVAL '1 1:1:1' DAY_SECOND);
-> '2004-12-30 22:58:59'
mysql> SELECT DATE_ADD('1900-01-01 00:00:00',
-> INTERVAL '-1 10' DAY_HOUR);
-> '1899-12-30 14:00:00'
mysql> SELECT DATE_SUB('1998-01-02', INTERVAL 31 DAY);
-> '1997-12-02'
mysql> SELECT DATE_ADD('1992-12-31 23:59:59.000002',
-> INTERVAL '1.999999' SECOND_MICROSECOND);
-> '1993-01-01 00:00:01.000001'

```

If you specify an interval value that is too short (does not include all the interval parts that would be expected from the *unit* keyword), MySQL assumes that you have left out the leftmost parts of the interval value. For example, if you specify a *unit* of `DAY_SECOND`, the value of *expr* is expected to have days, hours, minutes, and seconds parts. If you specify a value like `'1:10'`, MySQL assumes that the days and hours parts are missing and the value represents minutes and seconds. In other words, `'1:10' DAY_SECOND` is interpreted in such a way that it is equivalent to `'1:10' MINUTE_SECOND`. This is analogous to the way that MySQL interprets `TIME` values as representing elapsed time rather than as a time of day.

Because *expr* is treated as a string, be careful if you specify a nonstring value with `INTERVAL`. For example, with an interval specifier of `HOURL_MINUTE`, `6/4` evaluates to `1.5000` and is treated as 1 hour, 5000 minutes:

```

mysql> SELECT 6/4;
-> 1.5000
mysql> SELECT DATE_ADD('2009-01-01', INTERVAL 6/4 HOUR_MINUTE);
-> '2009-01-04 12:20:00'

```

To ensure interpretation of the interval value as you expect, a `CAST()` [1245] operation may be used. To treat `6/4` as 1 hour, 5 minutes, cast it to a `DECIMAL` value with a single fractional digit:

```

mysql> SELECT CAST(6/4 AS DECIMAL(3,1));
-> 1.5
mysql> SELECT DATE_ADD('1970-01-01 12:00:00',
-> INTERVAL CAST(6/4 AS DECIMAL(3,1)) HOUR_MINUTE);
-> '1970-01-01 13:05:00'

```

If you add to or subtract from a date value something that contains a time part, the result is automatically converted to a datetime value:

```

mysql> SELECT DATE_ADD('2013-01-01', INTERVAL 1 DAY);
-> '2013-01-02'
mysql> SELECT DATE_ADD('2013-01-01', INTERVAL 1 HOUR);
-> '2013-01-01 01:00:00'

```

If you add `MONTH`, `YEAR_MONTH`, or `YEAR` and the resulting date has a day that is larger than the maximum day for the new month, the day is adjusted to the maximum days in the new month:

```

mysql> SELECT DATE_ADD('2009-01-30', INTERVAL 1 MONTH);
-> '2009-02-28'

```

Date arithmetic operations require complete dates and do not work with incomplete dates such as `'2006-07-00'` or badly malformed dates:

```

mysql> SELECT DATE_ADD('2006-07-00', INTERVAL 1 DAY);
-> NULL
mysql> SELECT '2005-03-32' + INTERVAL 1 MONTH;
-> NULL

```

- `DATE_FORMAT(date, format)` [1215]

Formats the *date* value according to the *format* string.

The following specifiers may be used in the *format* string. The “%” character is required before format specifier characters.

Specifier	Description
%a	Abbreviated weekday name (Sun..Sat)
%b	Abbreviated month name (Jan..Dec)
%c	Month, numeric (0..12)
%D	Day of the month with English suffix (0th, 1st, 2nd, 3rd, ...)
%d	Day of the month, numeric (00..31)
%e	Day of the month, numeric (0..31)
%f	Microseconds (000000..999999)
%H	Hour (00..23)
%h	Hour (01..12)
%I	Hour (01..12)
%i	Minutes, numeric (00..59)
%j	Day of year (001..366)
%k	Hour (0..23)
%l	Hour (1..12)
%M	Month name (January..December)
%m	Month, numeric (00..12)
%p	AM or PM
%r	Time, 12-hour (hh:mm:ss followed by AM or PM)
%S	Seconds (00..59)
%s	Seconds (00..59)
%T	Time, 24-hour (hh:mm:ss)
%U	Week (00..53), where Sunday is the first day of the week
%u	Week (00..53), where Monday is the first day of the week
%V	Week (01..53), where Sunday is the first day of the week; used with %X
%v	Week (01..53), where Monday is the first day of the week; used with %x
%W	Weekday name (Sunday..Saturday)
%w	Day of the week (0=Sunday..6=Saturday)
%X	Year for the week where Sunday is the first day of the week, numeric, four digits; used with %V
%x	Year for the week, where Monday is the first day of the week, numeric, four digits; used with %v
%Y	Year, numeric, four digits
%y	Year, numeric (two digits)
%%	A literal “%” character
%x	x, for any “x” not listed above

Ranges for the month and day specifiers begin with zero due to the fact that MySQL permits the storing of incomplete dates such as '2014-00-00'.

The language used for day and month names and abbreviations is controlled by the value of the `lc_time_names` [495] system variable (Section 10.7, “MySQL Server Locale Support”).

`DATE_FORMAT()` [1215] returns a string with a character set and collation given by `character_set_connection` [467] and `collation_connection` [469] so that it can return month and weekday names containing non-ASCII characters.

```
mysql> SELECT DATE_FORMAT('2009-10-04 22:23:00', '%W %M %Y');
-> 'Sunday October 2009'
mysql> SELECT DATE_FORMAT('2007-10-04 22:23:00', '%H:%i:%s');
-> '22:23:00'
mysql> SELECT DATE_FORMAT('1900-10-04 22:23:00',
->      '%D %y %a %d %m %b %j');
-> '4th 00 Thu 04 10 Oct 277'
mysql> SELECT DATE_FORMAT('1997-10-04 22:23:00',
->      '%H %k %I %r %T %S %w');
-> '22 22 10 10:23:00 PM 22:23:00 00 6'
mysql> SELECT DATE_FORMAT('1999-01-01', '%X %V');
-> '1998 52'
mysql> SELECT DATE_FORMAT('2006-06-00', '%d');
-> '00'
```

- `DATE_SUB(date, INTERVAL expr unit)` [1217]

See the description for `DATE_ADD()` [1213].

- `DAY(date)` [1217]

`DAY()` [1217] is a synonym for `DAYOFMONTH()` [1217].

- `DAYNAME(date)` [1217]

Returns the name of the weekday for *date*. The language used for the name is controlled by the value of the `lc_time_names` [495] system variable (Section 10.7, “MySQL Server Locale Support”).

```
mysql> SELECT DAYNAME('2007-02-03');
-> 'Saturday'
```

- `DAYOFMONTH(date)` [1217]

Returns the day of the month for *date*, in the range 1 to 31, or 0 for dates such as '0000-00-00' or '2008-00-00' that have a zero day part.

```
mysql> SELECT DAYOFMONTH('2007-02-03');
-> 3
```

- `DAYOFWEEK(date)` [1217]

Returns the weekday index for *date* (1 = Sunday, 2 = Monday, ..., 7 = Saturday). These index values correspond to the ODBC standard.

```
mysql> SELECT DAYOFWEEK('2007-02-03');
-> 7
```

- `DAYOFYEAR(date)` [1217]

Returns the day of the year for *date*, in the range 1 to 366.

```
mysql> SELECT DAYOFYEAR('2007-02-03');
-> 34
```

- `EXTRACT(unit FROM date)` [1217]

The `EXTRACT()` [1217] function uses the same kinds of unit specifiers as `DATE_ADD()` [1213] or `DATE_SUB()` [1217], but extracts parts from the date rather than performing date arithmetic.

```
mysql> SELECT EXTRACT(YEAR FROM '2009-07-02');
-> 2009
mysql> SELECT EXTRACT(YEAR_MONTH FROM '2009-07-02 01:02:03');
-> 200907
```

```
mysql> SELECT EXTRACT(DAY_MINUTE FROM '2009-07-02 01:02:03');
-> 20102
mysql> SELECT EXTRACT(MICROSECOND
-> FROM '2003-01-02 10:30:00.000123');
-> 123
```

- [FROM\\_DAYS\(\*N\*\)](#) [1218]

Given a day number *N*, returns a [DATE](#) value.

```
mysql> SELECT FROM_DAYS(730669);
-> '2007-07-03'
```

Use [FROM\\_DAYS\(\)](#) [1218] with caution on old dates. It is not intended for use with values that precede the advent of the Gregorian calendar (1582). See [Section 12.8, “What Calendar Is Used By MySQL?”](#).

- [FROM\\_UNIXTIME\(\*unix\\_timestamp\*\)](#) [1218],  
[FROM\\_UNIXTIME\(\*unix\\_timestamp\*,\*format\*\)](#) [1218]

Returns a representation of the *unix\_timestamp* argument as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context. The value is expressed in the current time zone. *unix\_timestamp* is an internal timestamp value such as is produced by the [UNIX\\_TIMESTAMP\(\)](#) [1226] function.

If *format* is given, the result is formatted according to the *format* string, which is used the same way as listed in the entry for the [DATE\\_FORMAT\(\)](#) [1215] function.

```
mysql> SELECT FROM_UNIXTIME(1196440219);
-> '2007-11-30 10:30:19'
mysql> SELECT FROM_UNIXTIME(1196440219) + 0;
-> 20071130103019.000000
mysql> SELECT FROM_UNIXTIME(UNIX_TIMESTAMP(),
-> '%Y %D %M %h:%i:%s %x');
-> '2007 30th November 10:30:59 2007'
```

Note: If you use [UNIX\\_TIMESTAMP\(\)](#) [1226] and [FROM\\_UNIXTIME\(\)](#) [1218] to convert between [TIMESTAMP](#) values and Unix timestamp values, the conversion is lossy because the mapping is not one-to-one in both directions. For details, see the description of the [UNIX\\_TIMESTAMP\(\)](#) [1226] function.

- [GET\\_FORMAT\({DATE|TIME|DATETIME}, { 'EUR' | 'USA' | 'JIS' | 'ISO' | 'INTERNAL' }\)](#) [1218]

Returns a format string. This function is useful in combination with the [DATE\\_FORMAT\(\)](#) [1215] and the [STR\\_TO\\_DATE\(\)](#) [1221] functions.

The possible values for the first and second arguments result in several possible format strings (for the specifiers used, see the table in the [DATE\\_FORMAT\(\)](#) [1215] function description). ISO format refers to ISO 9075, not ISO 8601.

Function Call	Result
<a href="#">GET_FORMAT(DATE, 'USA')</a> [1218]	'%m.%d.%Y'
<a href="#">GET_FORMAT(DATE, 'JIS')</a> [1218]	'%Y-%m-%d'
<a href="#">GET_FORMAT(DATE, 'ISO')</a> [1218]	'%Y-%m-%d'
<a href="#">GET_FORMAT(DATE, 'EUR')</a> [1218]	'%d.%m.%Y'
<a href="#">GET_FORMAT(DATE, 'INTERNAL')</a> [1218]	'%Y%m%d'
<a href="#">GET_FORMAT(DATETIME, 'USA')</a> [1218]	'%Y-%m-%d %H.%i.%s'
<a href="#">GET_FORMAT(DATETIME, 'JIS')</a> [1218]	'%Y-%m-%d %H:%i:%s'
<a href="#">GET_FORMAT(DATETIME, 'ISO')</a> [1218]	'%Y-%m-%d %H:%i:%s'



Function Call	Result
<code>GET_FORMAT(DATETIME, 'EUR')</code> [1218]	'%Y-%m-%d %H.%i.%s'
<code>GET_FORMAT(DATETIME, 'INTERNAL')</code> [1218]	'%Y%m%d%H%i%s'
<code>GET_FORMAT(TIME, 'USA')</code> [1218]	'%h:%i:%s %p'
<code>GET_FORMAT(TIME, 'JIS')</code> [1218]	'%H:%i:%s'
<code>GET_FORMAT(TIME, 'ISO')</code> [1218]	'%H:%i:%s'
<code>GET_FORMAT(TIME, 'EUR')</code> [1218]	'%H.%i.%s'
<code>GET_FORMAT(TIME, 'INTERNAL')</code> [1218]	'%H%i%s'

`TIMESTAMP` can also be used as the first argument to `GET_FORMAT()` [1218], in which case the function returns the same values as for `DATETIME`.

```
mysql> SELECT DATE_FORMAT('2003-10-03',GET_FORMAT(DATE,'EUR'));
-> '03.10.2003'
mysql> SELECT STR_TO_DATE('10.31.2003',GET_FORMAT(DATE,'USA'));
-> '2003-10-31'
```

- `HOOR(time)` [1219]

Returns the hour for *time*. The range of the return value is 0 to 23 for time-of-day values. However, the range of `TIME` values actually is much larger, so `HOOR` can return values greater than 23.

```
mysql> SELECT HOOR('10:05:03');
-> 10
mysql> SELECT HOOR('272:59:59');
-> 272
```

- `LAST_DAY(date)` [1219]

Takes a date or datetime value and returns the corresponding value for the last day of the month. Returns `NULL` if the argument is invalid.

```
mysql> SELECT LAST_DAY('2003-02-05');
-> '2003-02-28'
mysql> SELECT LAST_DAY('2004-02-05');
-> '2004-02-29'
mysql> SELECT LAST_DAY('2004-01-01 01:01:01');
-> '2004-01-31'
mysql> SELECT LAST_DAY('2003-03-32');
-> NULL
```

- `LOCALTIME` [1219], `LOCALTIME([fsp])` [1219]

`LOCALTIME` [1219] and `LOCALTIME()` [1219] are synonyms for `NOW()` [1220].

- `LOCALTIMESTAMP` [1219], `LOCALTIMESTAMP([fsp])` [1219]

`LOCALTIMESTAMP` [1219] and `LOCALTIMESTAMP()` [1219] are synonyms for `NOW()` [1220].

- `MAKEDATE(year,dayofyear)` [1219]

Returns a date, given year and day-of-year values. *dayofyear* must be greater than 0 or the result is `NULL`.

```
mysql> SELECT MAKEDATE(2011,31), MAKEDATE(2011,32);
-> '2011-01-31', '2011-02-01'
mysql> SELECT MAKEDATE(2011,365), MAKEDATE(2014,365);
-> '2011-12-31', '2014-12-31'
mysql> SELECT MAKEDATE(2011,0);
-> NULL
```

- `MAKETIME(hour,minute,second)` [1219]

Returns a time value calculated from the *hour*, *minute*, and *second* arguments.

As of MySQL 5.6.4, the *second* argument can have a fractional part.

```
mysql> SELECT MAKETIME(12,15,30);
-> '12:15:30'
```

- `MICROSECOND(expr)` [1220]

Returns the microseconds from the time or datetime expression *expr* as a number in the range from 0 to 999999.

```
mysql> SELECT MICROSECOND('12:00:00.123456');
-> 123456
mysql> SELECT MICROSECOND('2009-12-31 23:59:59.000010');
-> 10
```

- `MINUTE(time)` [1220]

Returns the minute for *time*, in the range 0 to 59.

```
mysql> SELECT MINUTE('2008-02-03 10:05:03');
-> 5
```

- `MONTH(date)` [1220]

Returns the month for *date*, in the range 1 to 12 for January to December, or 0 for dates such as '0000-00-00' or '2008-00-00' that have a zero month part.

```
mysql> SELECT MONTH('2008-02-03');
-> 2
```

- `MONTHNAME(date)` [1220]

Returns the full name of the month for *date*. The language used for the name is controlled by the value of the `lc_time_names` [495] system variable (Section 10.7, "MySQL Server Locale Support").

```
mysql> SELECT MONTHNAME('2008-02-03');
-> 'February'
```

- `NOW([fsp])` [1220]

Returns the current date and time as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context. The value is expressed in the current time zone.

As of MySQL 5.6.4, if the *fsp* argument is given to specify a fractional seconds precision from 0 to 6, the return value includes a fractional seconds part of that many digits. Before 5.6.4, any argument is ignored.

```
mysql> SELECT NOW();
-> '2007-12-15 23:50:26'
mysql> SELECT NOW() + 0;
-> 20071215235026.000000
```

`NOW()` [1220] returns a constant time that indicates the time at which the statement began to execute. (Within a stored function or trigger, `NOW()` [1220] returns the time at which the function or triggering statement began to execute.) This differs from the behavior for `SYSDATE()` [1223], which returns the exact time at which it executes.

```
mysql> SELECT NOW(), SLEEP(2), NOW();
+-----+-----+-----+
| NOW() | SLEEP(2) | NOW() |
+-----+-----+-----+
| 2006-04-12 13:47:36 | 0 | 2006-04-12 13:47:36 |
```

```

+-----+-----+-----+
mysql> SELECT SYSDATE(), SLEEP(2), SYSDATE();
+-----+-----+-----+
| SYSDATE()          | SLEEP(2) | SYSDATE()          |
+-----+-----+-----+
| 2006-04-12 13:47:44 |          0 | 2006-04-12 13:47:46 |
+-----+-----+-----+

```

In addition, the `SET TIMESTAMP` statement affects the value returned by `NOW()` [1220] but not by `SYSDATE()` [1223]. This means that timestamp settings in the binary log have no effect on invocations of `SYSDATE()` [1223]. Setting the timestamp to a nonzero value causes each subsequent invocation of `NOW()` [1220] to return that value. Setting the timestamp to zero cancels this effect so that `NOW()` [1220] once again returns the current date and time.

See the description for `SYSDATE()` [1223] for additional information about the differences between the two functions.

- `PERIOD_ADD(P,N)` [1221]

Adds *N* months to period *P* (in the format `YYMM` or `YYYYMM`). Returns a value in the format `YYYYMM`. Note that the period argument *P* is *not* a date value.

```
mysql> SELECT PERIOD_ADD(200801,2);
-> 200803
```

- `PERIOD_DIFF(P1,P2)` [1221]

Returns the number of months between periods *P1* and *P2*. *P1* and *P2* should be in the format `YYMM` or `YYYYMM`. Note that the period arguments *P1* and *P2* are *not* date values.

```
mysql> SELECT PERIOD_DIFF(200802,200703);
-> 11
```

- `QUARTER(date)` [1221]

Returns the quarter of the year for *date*, in the range 1 to 4.

```
mysql> SELECT QUARTER('2008-04-01');
-> 2
```

- `SECOND(time)` [1221]

Returns the second for *time*, in the range 0 to 59.

```
mysql> SELECT SECOND('10:05:03');
-> 3
```

- `SEC_TO_TIME(seconds)` [1221]

Returns the *seconds* argument, converted to hours, minutes, and seconds, as a `TIME` value. The range of the result is constrained to that of the `TIME` data type. A warning occurs if the argument corresponds to a value outside that range.

```
mysql> SELECT SEC_TO_TIME(2378);
-> '00:39:38'
mysql> SELECT SEC_TO_TIME(2378) + 0;
-> 3938
```

- `STR_TO_DATE(str,format)` [1221]

This is the inverse of the `DATE_FORMAT()` [1215] function. It takes a string *str* and a format string *format*. `STR_TO_DATE()` [1221] returns a `DATETIME` value if the format string contains both date and time parts, or a `DATE` or `TIME` value if the string contains only date or time parts. If the date, time, or datetime value extracted from *str* is illegal, `STR_TO_DATE()` [1221] returns `NULL` and produces a warning.

The server scans *str* attempting to match *format* to it. The format string can contain literal characters and format specifiers beginning with %. Literal characters in *format* must match literally in *str*. Format specifiers in *format* must match a date or time part in *str*. For the specifiers that can be used in *format*, see the `DATE_FORMAT()` [1215] function description.

```
mysql> SELECT STR_TO_DATE('01,5,2013','%d,%m,%Y');
-> '2013-05-01'
mysql> SELECT STR_TO_DATE('May 1, 2013','%M %d,%Y');
-> '2013-05-01'
```

Scanning starts at the beginning of *str* and fails if *format* is found not to match. Extra characters at the end of *str* are ignored.

```
mysql> SELECT STR_TO_DATE('a09:30:17','a%h:%i:%s');
-> '09:30:17'
mysql> SELECT STR_TO_DATE('a09:30:17','%h:%i:%s');
-> NULL
mysql> SELECT STR_TO_DATE('09:30:17a','%h:%i:%s');
-> '09:30:17'
```

Unspecified date or time parts have a value of 0, so incompletely specified values in *str* produce a result with some or all parts set to 0:

```
mysql> SELECT STR_TO_DATE('abc','abc');
-> '0000-00-00'
mysql> SELECT STR_TO_DATE('9','%m');
-> '0000-09-00'
mysql> SELECT STR_TO_DATE('9','%s');
-> '00:00:09'
```

Range checking on the parts of date values is as described in [Section 11.3.1, “The DATE, DATETIME, and TIMESTAMP Types”](#). This means, for example, that “zero” dates or dates with part values of 0 are permitted unless the SQL mode is set to disallow such values.

```
mysql> SELECT STR_TO_DATE('00/00/0000','%m/%d/%Y');
-> '0000-00-00'
mysql> SELECT STR_TO_DATE('04/31/2004','%m/%d/%Y');
-> '2004-04-31'
```



### Note

You cannot use format “%X%V” to convert a year-week string to a date because the combination of a year and week does not uniquely identify a year and month if the week crosses a month boundary. To convert a year-week to a date, you should also specify the weekday:

```
mysql> SELECT STR_TO_DATE('200442 Monday','%X%V %W');
-> '2004-10-18'
```

- `SUBDATE(date, INTERVAL expr unit)` [1222], `SUBDATE(expr, days)` [1222]

When invoked with the `INTERVAL` form of the second argument, `SUBDATE()` [1222] is a synonym for `DATE_SUB()` [1217]. For information on the `INTERVAL unit` argument, see the discussion for `DATE_ADD()` [1213].

```
mysql> SELECT DATE_SUB('2008-01-02', INTERVAL 31 DAY);
-> '2007-12-02'
mysql> SELECT SUBDATE('2008-01-02', INTERVAL 31 DAY);
-> '2007-12-02'
```

The second form enables the use of an integer value for *days*. In such cases, it is interpreted as the number of days to be subtracted from the date or datetime expression *expr*.

```
mysql> SELECT SUBDATE('2008-01-02 12:00:00', 31);
-> '2007-12-02 12:00:00'
```

- [SUBTIME\(\*expr1\*,\*expr2\*\)](#) [1223]

[SUBTIME\(\)](#) [1223] returns  $expr1 - expr2$  expressed as a value in the same format as *expr1*. *expr1* is a time or datetime expression, and *expr2* is a time expression.

```
mysql> SELECT SUBTIME('2007-12-31 23:59:59.999999','1 1:1:1.000002');
-> '2007-12-30 22:58:58.999997'
mysql> SELECT SUBTIME('01:00:00.999999', '02:00:00.999998');
-> '-00:59:59.999999'
```

- [SYSDATE\(\)](#) [1223]

Returns the current date and time as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context.

[SYSDATE\(\)](#) [1223] returns the time at which it executes. This differs from the behavior for [NOW\(\)](#) [1220], which returns a constant time that indicates the time at which the statement began to execute. (Within a stored function or trigger, [NOW\(\)](#) [1220] returns the time at which the function or triggering statement began to execute.)

```
mysql> SELECT NOW(), SLEEP(2), NOW();
+-----+-----+-----+
| NOW()          | SLEEP(2) | NOW()          |
+-----+-----+-----+
| 2006-04-12 13:47:36 |          0 | 2006-04-12 13:47:36 |
+-----+-----+-----+

mysql> SELECT SYSDATE(), SLEEP(2), SYSDATE();
+-----+-----+-----+
| SYSDATE()      | SLEEP(2) | SYSDATE()      |
+-----+-----+-----+
| 2006-04-12 13:47:44 |          0 | 2006-04-12 13:47:46 |
+-----+-----+-----+
```

In addition, the [SET TIMESTAMP](#) statement affects the value returned by [NOW\(\)](#) [1220] but not by [SYSDATE\(\)](#) [1223]. This means that timestamp settings in the binary log have no effect on invocations of [SYSDATE\(\)](#) [1223].

Because [SYSDATE\(\)](#) [1223] can return different values even within the same statement, and is not affected by [SET TIMESTAMP](#), it is nondeterministic and therefore unsafe for replication if statement-based binary logging is used. If that is a problem, you can use row-based logging.

Alternatively, you can use the [--sysdate-is-now](#) [444] option to cause [SYSDATE\(\)](#) [1223] to be an alias for [NOW\(\)](#) [1220]. This works if the option is used on both the master and the slave.

The nondeterministic nature of [SYSDATE\(\)](#) [1223] also means that indexes cannot be used for evaluating expressions that refer to it.

- [TIME\(\*expr\*\)](#) [1223]

Extracts the time part of the time or datetime expression *expr* and returns it as a string.

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when [binlog\\_format](#) [2010] is set to [STATEMENT](#). (Bug #47995)

```
mysql> SELECT TIME('2003-12-31 01:02:03');
-> '01:02:03'
mysql> SELECT TIME('2003-12-31 01:02:03.000123');
-> '01:02:03.000123'
```

- [TIMEDIFF\(\*expr1\*,\*expr2\*\)](#) [1223]

[TIMEDIFF\(\)](#) [1223] returns  $expr1 - expr2$  expressed as a time value. *expr1* and *expr2* are time or date-and-time expressions, but both must be of the same type.

The result returned by `TIMEDIFF()` is limited to the range allowed for `TIME` values. Alternatively, you can use either of the functions `TIMESTAMPDIFF()` [1224] and `UNIX_TIMESTAMP()` [1226], both of which return integers.

```
mysql> SELECT TIMEDIFF('2000:01:01 00:00:00',
-> '2000:01:01 00:00:00.000001');
-> '-00:00:00.000001'
mysql> SELECT TIMEDIFF('2008-12-31 23:59:59.000001',
-> '2008-12-30 01:01:01.000002');
-> '46:58:57.999999'
```

- `TIMESTAMP(expr)` [1224], `TIMESTAMP(expr1,expr2)` [1224]

With a single argument, this function returns the date or datetime expression `expr` as a datetime value. With two arguments, it adds the time expression `expr2` to the date or datetime expression `expr1` and returns the result as a datetime value.

```
mysql> SELECT TIMESTAMP('2003-12-31');
-> '2003-12-31 00:00:00'
mysql> SELECT TIMESTAMP('2003-12-31 12:00:00','12:00:00');
-> '2004-01-01 00:00:00'
```

- `TIMESTAMPADD(unit,interval,datetime_expr)` [1224]

Adds the integer expression `interval` to the date or datetime expression `datetime_expr`. The unit for `interval` is given by the `unit` argument, which should be one of the following values: `MICROSECOND` (microseconds), `SECOND`, `MINUTE`, `HOURL`, `DAY`, `WEEK`, `MONTH`, `QUARTER`, or `YEAR`.

The `unit` value may be specified using one of keywords as shown, or with a prefix of `SQL_TSI_`. For example, `DAY` and `SQL_TSI_DAY` both are legal.

```
mysql> SELECT TIMESTAMPADD(MINUTE,1,'2003-01-02');
-> '2003-01-02 00:01:00'
mysql> SELECT TIMESTAMPADD(WEEK,1,'2003-01-02');
-> '2003-01-09'
```

- `TIMESTAMPDIFF(unit,datetime_expr1,datetime_expr2)` [1224]

Returns `datetime_expr2 - datetime_expr1`, where `datetime_expr1` and `datetime_expr2` are date or datetime expressions. One expression may be a date and the other a datetime; a date value is treated as a datetime having the time part `'00:00:00'` where necessary. The unit for the result (an integer) is given by the `unit` argument. The legal values for `unit` are the same as those listed in the description of the `TIMESTAMPADD()` [1224] function.

```
mysql> SELECT TIMESTAMPDIFF(MONTH,'2003-02-01','2003-05-01');
-> 3
mysql> SELECT TIMESTAMPDIFF(YEAR,'2002-05-01','2001-01-01');
-> -1
mysql> SELECT TIMESTAMPDIFF(MINUTE,'2003-02-01','2003-05-01 12:05:55');
-> 128885
```



#### Note

The order of the date or datetime arguments for this function is the opposite of that used with the `TIMESTAMP()` [1224] function when invoked with 2 arguments.

- `TIME_FORMAT(time,format)` [1224]

This is used like the `DATE_FORMAT()` [1215] function, but the `format` string may contain format specifiers only for hours, minutes, seconds, and microseconds. Other specifiers produce a `NULL` value or `0`.

If the *time* value contains an hour part that is greater than 23, the %H and %k hour format specifiers produce a value larger than the usual range of 0..23. The other hour format specifiers produce the hour value modulo 12.

```
mysql> SELECT TIME_FORMAT('100:00:00', '%H %k %h %I %l');
-> '100 100 04 04 4'
```

- [TIME\\_TO\\_SEC\(\*time\*\) \[1225\]](#)

Returns the *time* argument, converted to seconds.

```
mysql> SELECT TIME_TO_SEC('22:23:00');
-> 80580
mysql> SELECT TIME_TO_SEC('00:39:38');
-> 2378
```

- [TO\\_DAYS\(\*date\*\) \[1225\]](#)

Given a date *date*, returns a day number (the number of days since year 0).

```
mysql> SELECT TO_DAYS(950501);
-> 728779
mysql> SELECT TO_DAYS('2007-10-07');
-> 733321
```

[TO\\_DAYS\(\)](#) [1225] is not intended for use with values that precede the advent of the Gregorian calendar (1582), because it does not take into account the days that were lost when the calendar was changed. For dates before 1582 (and possibly a later year in other locales), results from this function are not reliable. See [Section 12.8, "What Calendar Is Used By MySQL?"](#), for details.

Remember that MySQL converts two-digit year values in dates to four-digit form using the rules in [Section 11.3, "Date and Time Types"](#). For example, '2008-10-07' and '08-10-07' are seen as identical dates:

```
mysql> SELECT TO_DAYS('2008-10-07'), TO_DAYS('08-10-07');
-> 733687, 733687
```

In MySQL, the zero date is defined as '0000-00-00', even though this date is itself considered invalid. This means that, for '0000-00-00' and '0000-01-01', [TO\\_DAYS\(\)](#) [1225] returns the values shown here:

```
mysql> SELECT TO_DAYS('0000-00-00');
+-----+
| to_days('0000-00-00') |
+-----+
| NULL |
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Warning | 1292 | Incorrect datetime value: '0000-00-00' |
+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT TO_DAYS('0000-01-01');
+-----+
| to_days('0000-01-01') |
+-----+
| 1 |
+-----+
1 row in set (0.00 sec)
```

This is true whether or not the [ALLOW\\_INVALID\\_DATES \[603\]](#) SQL server mode is enabled.

- `TO_SECONDS(expr)` [1226]

Given a date or datetime *expr*, returns a the number of seconds since the year 0. If *expr* is not a valid date or datetime value, returns `NULL`.

```
mysql> SELECT TO_SECONDS(950501);
      -> 62966505600
mysql> SELECT TO_SECONDS('2009-11-29');
      -> 63426672000
mysql> SELECT TO_SECONDS('2009-11-29 13:43:32');
      -> 63426721412
mysql> SELECT TO_SECONDS(NOW());
      -> 63426721458
```

Like `TO_DAYS()` [1225], `TO_SECONDS()` is not intended for use with values that precede the advent of the Gregorian calendar (1582), because it does not take into account the days that were lost when the calendar was changed. For dates before 1582 (and possibly a later year in other locales), results from this function are not reliable. See [Section 12.8, “What Calendar Is Used By MySQL?”](#), for details.

Like `TO_DAYS()` [1225], `TO_SECONDS()`, converts two-digit year values in dates to four-digit form using the rules in [Section 11.3, “Date and Time Types”](#).

In MySQL, the zero date is defined as `'0000-00-00'`, even though this date is itself considered invalid. This means that, for `'0000-00-00'` and `'0000-01-01'`, `TO_SECONDS()` [1226] returns the values shown here:

```
mysql> SELECT TO_SECONDS('0000-00-00');
+-----+
| TO_SECONDS('0000-00-00') |
+-----+
| NULL |
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Warning | 1292 | Incorrect datetime value: '0000-00-00' |
+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT TO_SECONDS('0000-01-01');
+-----+
| TO_SECONDS('0000-01-01') |
+-----+
| 86400 |
+-----+
1 row in set (0.00 sec)
```

This is true whether or not the `ALLOW_INVALID_DATES` [603] SQL server mode is enabled.

- `UNIX_TIMESTAMP()` [1226], `UNIX_TIMESTAMP(date)` [1226]

If called with no argument, returns a Unix timestamp (seconds since `'1970-01-01 00:00:00'` UTC) as an unsigned integer. If `UNIX_TIMESTAMP()` [1226] is called with a *date* argument, it returns the value of the argument as seconds since `'1970-01-01 00:00:00'` UTC. *date* may be a `DATE` string, a `DATETIME` string, a `TIMESTAMP`, or a number in the format `YYMMDD` or `YYYYMMDD`. The server interprets *date* as a value in the current time zone and converts it to an internal value in UTC. Clients can set their time zone as described in [Section 10.6, “MySQL Server Time Zone Support”](#).

```
mysql> SELECT UNIX_TIMESTAMP();
      -> 1196440210
mysql> SELECT UNIX_TIMESTAMP('2007-11-30 10:30:19');
```



```
-> 1196440219
```

When `UNIX_TIMESTAMP()` [1226] is used on a `TIMESTAMP` column, the function returns the internal timestamp value directly, with no implicit “string-to-Unix-timestamp” conversion. If you pass an out-of-range date to `UNIX_TIMESTAMP()` [1226], it returns 0.

Note: If you use `UNIX_TIMESTAMP()` [1226] and `FROM_UNIXTIME()` [1218] to convert between `TIMESTAMP` values and Unix timestamp values, the conversion is lossy because the mapping is not one-to-one in both directions. For example, due to conventions for local time zone changes, it is possible for two `UNIX_TIMESTAMP()` [1226] to map two `TIMESTAMP` values to the same Unix timestamp value. `FROM_UNIXTIME()` [1218] will map that value back to only one of the original `TIMESTAMP` values. Here is an example, using `TIMESTAMP` values in the `CET` time zone:

```
mysql> SELECT UNIX_TIMESTAMP('2005-03-27 03:00:00');
+-----+
| UNIX_TIMESTAMP('2005-03-27 03:00:00') |
+-----+
| 1111885200 |
+-----+
mysql> SELECT UNIX_TIMESTAMP('2005-03-27 02:00:00');
+-----+
| UNIX_TIMESTAMP('2005-03-27 02:00:00') |
+-----+
| 1111885200 |
+-----+
mysql> SELECT FROM_UNIXTIME(1111885200);
+-----+
| FROM_UNIXTIME(1111885200) |
+-----+
| 2005-03-27 03:00:00 |
+-----+
```

If you want to subtract `UNIX_TIMESTAMP()` [1226] columns, you might want to cast the result to signed integers. See Section 12.10, “Cast Functions and Operators”.

- `UTC_DATE` [1227], `UTC_DATE()` [1227]

Returns the current UTC date as a value in 'YYYY-MM-DD' or YYYYMMDD format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT UTC_DATE(), UTC_DATE() + 0;
-> '2003-08-14', 20030814
```

- `UTC_TIME` [1227], `UTC_TIME([fsp])` [1227]

Returns the current UTC time as a value in 'HH:MM:SS' or HHMMSS format, depending on whether the function is used in a string or numeric context.

As of MySQL 5.6.4, if the `fsp` argument is given to specify a fractional seconds precision from 0 to 6, the return value includes a fractional seconds part of that many digits. Before 5.6.4, any argument is ignored.

```
mysql> SELECT UTC_TIME(), UTC_TIME() + 0;
-> '18:07:53', 180753.000000
```

- `UTC_TIMESTAMP` [1227], `UTC_TIMESTAMP([fsp])` [1227]

Returns the current UTC date and time as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context.

As of MySQL 5.6.4, if the `fsp` argument is given to specify a fractional seconds precision from 0 to 6, the return value includes a fractional seconds part of that many digits. Before 5.6.4, any argument is ignored.

```
mysql> SELECT UTC_TIMESTAMP(), UTC_TIMESTAMP() + 0;
```

```
-> '2003-08-14 18:08:04', 20030814180804.000000
```

- [WEEK\(\*date\* \[, \*mode\*\]\) \[1228\]](#)

This function returns the week number for *date*. The two-argument form of [WEEK\(\)](#) [1228] enables you to specify whether the week starts on Sunday or Monday and whether the return value should be in the range from 0 to 53 or from 1 to 53. If the *mode* argument is omitted, the value of the [default\\_week\\_format](#) [475] system variable is used. See [Section 5.1.4, “Server System Variables”](#).

The following table describes how the *mode* argument works.

Mode	First day of week	Range	Week 1 is the first week ...
0	Sunday	0-53	with a Sunday in this year
1	Monday	0-53	with more than 3 days this year
2	Sunday	1-53	with a Sunday in this year
3	Monday	1-53	with more than 3 days this year
4	Sunday	0-53	with more than 3 days this year
5	Monday	0-53	with a Monday in this year
6	Sunday	1-53	with more than 3 days this year
7	Monday	1-53	with a Monday in this year

```
mysql> SELECT WEEK('2008-02-20');
-> 7
mysql> SELECT WEEK('2008-02-20',0);
-> 7
mysql> SELECT WEEK('2008-02-20',1);
-> 8
mysql> SELECT WEEK('2008-12-31',1);
-> 53
```

Note that if a date falls in the last week of the previous year, MySQL returns 0 if you do not use 2, 3, 6, or 7 as the optional *mode* argument:

```
mysql> SELECT YEAR('2000-01-01'), WEEK('2000-01-01',0);
-> 2000, 0
```

One might argue that MySQL should return 52 for the [WEEK\(\)](#) [1228] function, because the given date actually occurs in the 52nd week of 1999. We decided to return 0 instead because we want the function to return “the week number in the given year.” This makes use of the [WEEK\(\)](#) [1228] function reliable when combined with other functions that extract a date part from a date.

If you would prefer the result to be evaluated with respect to the year that contains the first day of the week for the given date, use 0, 2, 5, or 7 as the optional *mode* argument.

```
mysql> SELECT WEEK('2000-01-01',2);
-> 52
```

Alternatively, use the [YEARWEEK\(\)](#) [1229] function:

```
mysql> SELECT YEARWEEK('2000-01-01');
-> 199952
mysql> SELECT MID(YEARWEEK('2000-01-01'),5,2);
-> '52'
```

- [WEEKDAY\(\*date\*\) \[1228\]](#)

Returns the weekday index for *date* (0 = Monday, 1 = Tuesday, ... 6 = Sunday).

```
mysql> SELECT WEEKDAY('2008-02-03 22:23:00');
-> 6
```

```
mysql> SELECT WEEKDAY('2007-11-06');
-> 1
```

- `WEEKOFYEAR(date)` [1229]

Returns the calendar week of the date as a number in the range from 1 to 53.

`WEEKOFYEAR()` [1229] is a compatibility function that is equivalent to `WEEK(date, 3)` [1228].

```
mysql> SELECT WEEKOFYEAR('2008-02-20');
-> 8
```

- `YEAR(date)` [1229]

Returns the year for *date*, in the range 1000 to 9999, or 0 for the “zero” date.

```
mysql> SELECT YEAR('1987-01-01');
-> 1987
```

- `YEARWEEK(date)` [1229], `YEARWEEK(date, mode)` [1229]

Returns year and week for a date. The *mode* argument works exactly like the *mode* argument to `WEEK()` [1228]. The year in the result may be different from the year in the date argument for the first and the last week of the year.

```
mysql> SELECT YEARWEEK('1987-01-01');
-> 198653
```

Note that the week number is different from what the `WEEK()` [1228] function would return (0) for optional arguments 0 or 1, as `WEEK()` [1228] then returns the week in the context of the given year.

## 12.8. What Calendar Is Used By MySQL?

MySQL uses what is known as a *proleptic Gregorian calendar*.

Every country that has switched from the Julian to the Gregorian calendar has had to discard at least ten days during the switch. To see how this works, consider the month of October 1582, when the first Julian-to-Gregorian switch occurred.

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
1	2	3	4	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

There are no dates between October 4 and October 15. This discontinuity is called the *cutover*. Any dates before the cutover are Julian, and any dates following the cutover are Gregorian. Dates during a cutover are nonexistent.

A calendar applied to dates when it was not actually in use is called *proleptic*. Thus, if we assume there was never a cutover and Gregorian rules always rule, we have a proleptic Gregorian calendar. This is what is used by MySQL, as is required by standard SQL. For this reason, dates prior to the cutover stored as MySQL `DATE` or `DATETIME` values must be adjusted to compensate for the difference. It is important to realize that the cutover did not occur at the same time in all countries, and that the later it happened, the more days were lost. For example, in Great Britain, it took place in 1752, when Wednesday September 2 was followed by Thursday September 14. Russia remained on the Julian calendar until 1918, losing 13 days in the process, and what is popularly referred to as its “October Revolution” occurred in November according to the Gregorian calendar.

## 12.9. Full-Text Search Functions

`MATCH(col1,col2,...) AGAINST(expr [search_modifier])` [1229]

```

search_modifier:
{
  | IN NATURAL LANGUAGE MODE
  | IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION
  | IN BOOLEAN MODE
  | WITH QUERY EXPANSION
}

```

MySQL has support for full-text indexing and searching:

- A full-text index in MySQL is an index of type `FULLTEXT`.
- Full-text indexes can be used only with `InnoDB` or `MyISAM` tables, and can be created only for `CHAR`, `VARCHAR`, or `TEXT` columns.
- A `FULLTEXT` index definition can be given in the `CREATE TABLE` statement when a table is created, or added later using `ALTER TABLE` or `CREATE INDEX`.
- For large data sets, it is much faster to load your data into a table that has no `FULLTEXT` index and then create the index after that, than to load data into a table that has an existing `FULLTEXT` index.

Full-text searching is performed using `MATCH() ... AGAINST [1229]` syntax. `MATCH() [1229]` takes a comma-separated list that names the columns to be searched. `AGAINST` takes a string to search for, and an optional modifier that indicates what type of search to perform. The search string must be a string value that is constant during query evaluation. This rules out, for example, a table column because that can differ for each row.

There are three types of full-text searches:

- A natural language search interprets the search string as a phrase in natural human language (a phrase in free text). There are no special operators. The stopword list applies, controlled by `innodb_ft_enable_stopword [1724]`, `innodb_ft_server_stopword_table [1725]`, and `innodb_ft_user_stopword_table [1726]` for `InnoDB` search indexes, and `ft_stopword_file [485]` for `MyISAM` ones.
- Full-text searches are natural language searches if the `IN NATURAL LANGUAGE MODE` modifier is given or if no modifier is given. For more information, see [Section 12.9.1, “Natural Language Full-Text Searches”](#).
- A boolean search interprets the search string using the rules of a special query language. The string contains the words to search for. It can also contain operators that specify requirements such that a word must be present or absent in matching rows, or that it should be weighted higher or lower than usual. Certain common words (stopwords) are omitted from the search index and do not match if present in the search string. The `IN BOOLEAN MODE` modifier specifies a boolean search. For more information, see [Section 12.9.2, “Boolean Full-Text Searches”](#).
- A query expansion search is a modification of a natural language search. The search string is used to perform a natural language search. Then words from the most relevant rows returned by the search are added to the search string and the search is done again. The query returns the rows from the second search. The `IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION` or `WITH QUERY EXPANSION` modifier specifies a query expansion search. For more information, see [Section 12.9.3, “Full-Text Searches with Query Expansion”](#).

For information about `FULLTEXT` query performance, see [Section 8.3.4, “Column Indexes”](#).

For more technical details about processing for `InnoDB FULLTEXT` indexes, see [Section 14.2.3.12.3, “FULLTEXT Indexes”](#).

Constraints on full-text searching are listed in [Section 12.9.5, “Full-Text Restrictions”](#).

The `myisam_ftdump` utility dumps the contents of a `MyISAM` full-text index. This may be helpful for debugging full-text queries. See [Section 4.6.2, “myisam\\_ftdump — Display Full-Text Index information”](#).

## 12.9.1. Natural Language Full-Text Searches

By default or with the `IN NATURAL LANGUAGE MODE` modifier, the `MATCH()` [1229] function performs a natural language search for a string against a *text collection*. A collection is a set of one or more columns included in a `FULLTEXT` index. The search string is given as the argument to `AGAINST()`. For each row in the table, `MATCH()` [1229] returns a relevance value; that is, a similarity measure between the search string and the text in that row in the columns named in the `MATCH()` [1229] list.

```
mysql> CREATE TABLE articles (
    id INT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,
    title VARCHAR(200),
    body TEXT,
    FULLTEXT (title,body)
) ENGINE=InnoDB;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO articles (title,body) VALUES
    ('MySQL Tutorial','DBMS stands for DataBase ...'),
    ('How To Use MySQL Well','After you went through a ...'),
    ('Optimizing MySQL','In this tutorial we will show ...'),
    ('1001 MySQL Tricks','1. Never run mysqld as root. 2. ...'),
    ('MySQL vs. YourSQL','In the following database comparison ...'),
    ('MySQL Security','When configured properly, MySQL ...');
Query OK, 6 rows affected (0.00 sec)
Records: 6 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM articles
    WHERE MATCH (title,body)
    AGAINST ('database' IN NATURAL LANGUAGE MODE);
+-----+-----+-----+
| id | title | body |
+-----+-----+-----+
| 1 | MySQL Tutorial | DBMS stands for DataBase ... |
| 5 | MySQL vs. YourSQL | In the following database comparison ... |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

By default, the search is performed in case-insensitive fashion. To perform a case-sensitive full-text search, use a binary collation for the indexed columns. For example, a column that uses the `latin1` character set of can be assigned a collation of `latin1_bin` to make it case sensitive for full-text searches.

When `MATCH()` [1229] is used in a `WHERE` clause, as in the example shown earlier, the rows returned are automatically sorted with the highest relevance first. Relevance values are nonnegative floating-point numbers. Zero relevance means no similarity. Relevance is computed based on the number of words in the row, the number of unique words in that row, the total number of words in the collection, and the number of documents (rows) that contain a particular word.

To simply count matches, you could use a query like this:

```
mysql> SELECT COUNT(*) FROM articles
    WHERE MATCH (title,body)
    AGAINST ('database' IN NATURAL LANGUAGE MODE);
+-----+
| COUNT(*) |
+-----+
| 2 |
+-----+
1 row in set (0.00 sec)
```

You might find it quicker to rewrite the query as follows:

```
mysql> SELECT
    COUNT(IF(MATCH (title,body) AGAINST ('database' IN NATURAL LANGUAGE MODE), 1, NULL))
    AS count
    FROM articles;
+-----+
| count |
+-----+
```

```

+-----+
|      2 |
+-----+
1 row in set (0.03 sec)

```

The first query does some extra work (sorting the results by relevance) but also can use an index lookup based on the `WHERE` clause. The index lookup might make the first query faster if the search matches few rows. The second query performs a full table scan, which might be faster than the index lookup if the search term was present in most rows.

For natural-language full-text searches, the columns named in the `MATCH()` [1229] function must be the same columns included in some `FULLTEXT` index in your table. For the preceding query, note that the columns named in the `MATCH()` [1229] function (`title` and `body`) are the same as those named in the definition of the `article` table's `FULLTEXT` index. To search the `title` or `body` separately, you would create separate `FULLTEXT` indexes for each column.

You can also perform a boolean search or a search with query expansion. These search types are described in [Section 12.9.2, "Boolean Full-Text Searches"](#), and [Section 12.9.3, "Full-Text Searches with Query Expansion"](#).

A full-text search that uses an index can name columns only from a single table in the `MATCH()` [1229] clause because an index cannot span multiple tables. A boolean search can be done in the absence of an index (albeit more slowly), in which case it is possible to name columns from multiple tables.

The preceding example is a basic illustration that shows how to use the `MATCH()` [1229] function where rows are returned in order of decreasing relevance. The next example shows how to retrieve the relevance values explicitly. Returned rows are not ordered because the `SELECT` statement includes neither `WHERE` nor `ORDER BY` clauses:

```

mysql> SELECT id, MATCH (title,body)
        AGAINST ('Tutorial' IN NATURAL LANGUAGE MODE) AS score
        FROM articles;

```

id	score
1	0.22764469683170319
2	0
3	0.22764469683170319
4	0
5	0
6	0

```

6 rows in set (0.00 sec)

```

The following example is more complex. The query returns the relevance values and it also sorts the rows in order of decreasing relevance. To achieve this result, specify `MATCH()` [1229] twice: once in the `SELECT` list and once in the `WHERE` clause. This causes no additional overhead, because the MySQL optimizer notices that the two `MATCH()` [1229] calls are identical and invokes the full-text search code only once.

```

mysql> SELECT id, body, MATCH (title,body) AGAINST
        ('Security implications of running MySQL as root'
         IN NATURAL LANGUAGE MODE) AS score
        FROM articles WHERE MATCH (title,body) AGAINST
        ('Security implications of running MySQL as root'
         IN NATURAL LANGUAGE MODE);

```

id	title	body
5	MySQL vs. YourSQL	In the following database comparison ...
1	MySQL Tutorial	DBMS stands for DataBase ...
3	Optimizing MySQL	In this tutorial we will show ...
6	MySQL Security	When configured properly, MySQL ...
2	How To Use MySQL Well	After you went through a ...
4	1001 MySQL Tricks	1. Never run mysqld as root. 2. ...

```

6 rows in set (0.00 sec)

```

The MySQL `FULLTEXT` implementation regards any sequence of true word characters (letters, digits, and underscores) as a word. That sequence may also contain apostrophes (“'”), but not more than one in a row. This means that `aaa 'bbb` is regarded as one word, but `aaa ' 'bbb` is regarded as two words. Apostrophes at the beginning or the end of a word are stripped by the `FULLTEXT` parser; `'aaa 'bbb` would be parsed as `aaa 'bbb`.

The `FULLTEXT` parser determines where words start and end by looking for certain delimiter characters; for example, “ ” (space), “,” (comma), and “.” (period). If words are not separated by delimiters (as in, for example, Chinese), the `FULLTEXT` parser cannot determine where a word begins or ends. To be able to add words or other indexed terms in such languages to a `FULLTEXT` index, you must preprocess them so that they are separated by some arbitrary delimiter such as “ ”.

In MySQL 5.6, it is possible to write a plugin that replaces the built-in full-text parser. For details, see [Section 22.2, “The MySQL Plugin API”](#). For example parser plugin source code, see the `plugin/fulltext` directory of a MySQL source distribution.

Some words are ignored in full-text searches:

- Any word that is too short is ignored. The default minimum length of words that are found by full-text searches is three characters for `InnoDB` search indexes, or four characters for `MyISAM`. You can control the cutoff by setting a configuration option before creating the index: `innodb_ft_min_token_size` [1725] configuration option for `InnoDB` search indexes, or `ft_min_word_len` [484] for `MyISAM`.
- Words in the stopword list are ignored. A stopword is a word such as “the” or “some” that is so common that it is considered to have zero semantic value. There is a built-in stopword list, but it can be overridden by a user-defined list. The stopword lists and related configuration options are different for `InnoDB` search indexes and `MyISAM` ones. Stopword processing is controlled by the configuration options `innodb_ft_enable_stopword` [1724], `innodb_ft_server_stopword_table` [1725], and `innodb_ft_user_stopword_table` [1726] for `InnoDB` search indexes, and `ft_stopword_file` [485] for `MyISAM` ones.

The default stopword lists are shown in [Section 12.9.4, “Full-Text Stopwords”](#). The default minimum word length and stopword list can be changed as described in [Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”](#).

Every correct word in the collection and in the query is weighted according to its significance in the collection or query. Thus, a word that is present in many documents has a lower weight, because it has lower semantic value in this particular collection. Conversely, if the word is rare, it receives a higher weight. The weights of the words are combined to compute the relevance of the row. This technique works best with large collections.



### MyISAM Limitation

For very small tables, word distribution does not adequately reflect their semantic value, and this model may sometimes produce bizarre results for search indexes on `MyISAM` tables. For example, although the word “MySQL” is present in every row of the `articles` table shown earlier, a search for the word in a `MyISAM` search index produces no results:

```
mysql> SELECT * FROM articles
      WHERE MATCH (title,body)
      AGAINST ('MySQL' IN NATURAL LANGUAGE MODE);
Empty set (0.00 sec)
```

The search result is empty because the word “MySQL” is present in at least 50% of the rows, and so is effectively treated as a stopword. This filtering technique is more suitable for large data sets, where you might not want the result set to return every second row from a 1GB table, than for small data sets where it might cause poor results for popular terms.



The 50% threshold can surprise you when you first try full-text searching to see how it works, and makes [InnoDB](#) tables more suited to experimentation with full-text searches. If you create a [MyISAM](#) table and insert only one or two rows of text into it, every word in the text occurs in at least 50% of the rows. As a result, no search returns any results until the table contains more rows. Users who need to bypass the 50% limitation can build search indexes on [InnoDB](#) tables, or the boolean search mode explained in [Section 12.9.2, “Boolean Full-Text Searches”](#).

## 12.9.2. Boolean Full-Text Searches

MySQL can perform boolean full-text searches using the `IN BOOLEAN MODE` modifier. With this modifier, certain characters have special meaning at the beginning or end of words in the search string. In the following query, the `+` and `-` operators indicate that a word must be present or absent, respectively, for a match to occur. Thus, the query retrieves all the rows that contain the word “MySQL” but that do *not* contain the word “YourSQL”:

```
mysql> SELECT * FROM articles WHERE MATCH (title,body)
      AGAINST ('+MySQL -YourSQL' IN BOOLEAN MODE);
```

id	title	body
1	MySQL Tutorial	DBMS stands for DataBase ...
2	How To Use MySQL Well	After you went through a ...
3	Optimizing MySQL	In this tutorial we will show ...
4	1001 MySQL Tricks	1. Never run mysqld as root. 2. ...
6	MySQL Security	When configured properly, MySQL ...



**Note**

In implementing this feature, MySQL uses what is sometimes referred to as *implied Boolean logic*, in which

- `+` stands for `AND`
- `-` stands for `NOT`
- `[no operator]` implies `OR`

Boolean full-text searches have these characteristics:

- They do not use the 50% threshold that applies to [MyISAM](#) search indexes.
- They do not automatically sort rows in order of decreasing relevance.
- Boolean queries against a [MyISAM](#) search index can work even without a `FULLTEXT` index, although a search executed in this fashion would be quite slow. [InnoDB](#) tables require a `FULLTEXT` index to perform boolean queries.
- The minimum and maximum word length full-text parameters apply: `innodb_ft_min_token_size` [1725] and `innodb_ft_max_token_size` [1724] for [InnoDB](#) search indexes, and `ft_min_word_len` [484] and `ft_max_word_len` [484] for [MyISAM](#) ones.
- The stopword list applies, controlled by `innodb_ft_enable_stopword` [1724], `innodb_ft_server_stopword_table` [1725], and `innodb_ft_user_stopword_table` [1726] for [InnoDB](#) search indexes, and `ft_stopword_file` [485] for [MyISAM](#) ones.

The boolean full-text search capability supports the following operators:

- `+`



A leading plus sign indicates that this word *must* be present in each row that is returned.

- -

A leading minus sign indicates that this word must *not* be present in any of the rows that are returned.

Note: The - operator acts only to exclude rows that are otherwise matched by other search terms. Thus, a boolean-mode search that contains only terms preceded by - returns an empty result. It does not return “all rows except those containing any of the excluded terms.”

- (no operator)

By default (when neither + nor - is specified), the word is optional, but the rows that contain it are rated higher. This mimics the behavior of `MATCH() ... AGAINST() [1229]` without the `IN BOOLEAN MODE` modifier.

- `@distance`

This operator works on `InnoDB` tables only. It tests whether two or more words all start within a specified distance from each other, measured in words. Specify the search words within a double-quoted string immediately before the `@distance` operator, for example, `MATCH(coll) AGAINST('"word1 word2 word3" @8' IN BOOLEAN MODE)`

- > <

These two operators are used to change a word's contribution to the relevance value that is assigned to a row. The > operator increases the contribution and the < operator decreases it. See the example following this list.

- ( )

Parentheses group words into subexpressions. Parenthesized groups can be nested.

- ~

A leading tilde acts as a negation operator, causing the word's contribution to the row's relevance to be negative. This is useful for marking “noise” words. A row containing such a word is rated lower than others, but is not excluded altogether, as it would be with the - operator.

- \*

The asterisk serves as the truncation (or wildcard) operator. Unlike the other operators, it is *appended* to the word to be affected. Words match if they begin with the word preceding the \* operator.

If a word is specified with the truncation operator, it is not stripped from a boolean query, even if it is too short or a stopword. Whether a word is too short is determined from the `innodb_ft_min_token_size [1725]` setting for `InnoDB` tables, or `ft_min_word_len [484]` for `MyISAM` tables. The wildcarded word is considered as a prefix that must be present at the start of one or more words. If the minimum word length is 4, a search for `'+word +the*'` could return fewer rows than a search for `'+word +the'`, because the second query ignores the too-short search term `the`.

- "

A phrase that is enclosed within double quote (“”) characters matches only rows that contain the phrase *literally, as it was typed*. The full-text engine splits the phrase into words and performs a search in the `FULLTEXT` index for the words. Nonword characters need not be matched exactly: Phrase searching requires only that matches contain exactly the same words as the phrase and in the same order. For example, `"test phrase"` matches `"test, phrase"`.

If the phrase contains no words that are in the index, the result is empty. The words might not be in the index because of a combination of factors: if they do not exist in the text, are stopwords, or are shorter than the minimum length of indexed words.

The following examples demonstrate some search strings that use boolean full-text operators:

- `'apple banana'`

Find rows that contain at least one of the two words.

- `'+apple +juice'`

Find rows that contain both words.

- `'+apple macintosh'`

Find rows that contain the word “apple”, but rank rows higher if they also contain “macintosh”.

- `'+apple -macintosh'`

Find rows that contain the word “apple” but not “macintosh”.

- `'+apple ~macintosh'`

Find rows that contain the word “apple”, but if the row also contains the word “macintosh”, rate it lower than if row does not. This is “softer” than a search for `'+apple -macintosh'`, for which the presence of “macintosh” causes the row not to be returned at all.

- `'+apple +(>turnover <strudel)'`

Find rows that contain the words “apple” and “turnover”, or “apple” and “strudel” (in any order), but rank “apple turnover” higher than “apple strudel”.

- `'apple*'`

Find rows that contain words such as “apple”, “apples”, “applesauce”, or “applet”.

- `'"some words"'`

Find rows that contain the exact phrase “some words” (for example, rows that contain “some words of wisdom” but not “some noise words”). Note that the “” characters that enclose the phrase are operator characters that delimit the phrase. They are not the quotation marks that enclose the search string itself.

### 12.9.3. Full-Text Searches with Query Expansion

Full-text search supports query expansion (and in particular, its variant “blind query expansion”). This is generally useful when a search phrase is too short, which often means that the user is relying on implied knowledge that the full-text search engine lacks. For example, a user searching for “database” may really mean that “MySQL”, “Oracle”, “DB2”, and “RDBMS” all are phrases that should match “databases” and should be returned, too. This is implied knowledge.

Blind query expansion (also known as automatic relevance feedback) is enabled by adding `WITH QUERY EXPANSION` or `IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION` following the search phrase. It works by performing the search twice, where the search phrase for the second search is the original search phrase concatenated with the few most highly relevant documents from the first search. Thus, if one of these documents contains the word “databases” and the word “MySQL”, the second search finds the documents that contain the word “MySQL” even if they do not contain the word “database”. The following example shows this difference:

```
mysql> SELECT * FROM articles
WHERE MATCH (title,body)
AGAINST ('database' IN NATURAL LANGUAGE MODE);
```

```

+-----+-----+-----+
| id | title                | body                |
+-----+-----+-----+
| 1  | MySQL Tutorial       | DBMS stands for DataBase ... |
| 5  | MySQL vs. YourSQL   | In the following database comparison ... |
+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT * FROM articles
      WHERE MATCH (title,body)
      AGAINST ('database' WITH QUERY EXPANSION);
+-----+-----+-----+
| id | title                | body                |
+-----+-----+-----+
| 5  | MySQL vs. YourSQL   | In the following database comparison ... |
| 1  | MySQL Tutorial       | DBMS stands for DataBase ... |
| 3  | Optimizing MySQL    | In this tutorial we will show ... |
| 6  | MySQL Security      | When configured properly, MySQL ... |
| 2  | How To Use MySQL Well | After you went through a ... |
| 4  | 1001 MySQL Tricks   | 1. Never run mysqld as root. 2. ... |
+-----+-----+-----+
6 rows in set (0.00 sec)

```

Another example could be searching for books by Georges Simenon about Maigret, when a user is not sure how to spell “Maigret”. A search for “Megre and the reluctant witnesses” finds only “Maigret and the Reluctant Witnesses” without query expansion. A search with query expansion finds all books with the word “Maigret” on the second pass.



#### Note

Because blind query expansion tends to increase noise significantly by returning nonrelevant documents, use it only when a search phrase is short.

## 12.9.4. Full-Text Stopwords

The stopwords list is loaded and searched for full-text queries using the server character set and collation (the values of the `character_set_server` and `collation_server` system variables). False hits or misses might occur for stopwords lookups if the stopwords file or columns used for full-text indexing or searches have a character set or collation different from `character_set_server` or `collation_server`.

Case sensitivity of stopwords lookups depends on the server collation. For example, lookups are case insensitive if the collation is `latin1_swedish_ci`, whereas lookups are case sensitive if the collation is `latin1_general_cs` or `latin1_bin`.

### Stopwords for InnoDB Search Indexes

InnoDB has a relatively short list of default stopwords, because documents from technical, literary, and so on sources often use short words as keywords or in significant phrases. For example, you might search for “to be or not to be” and expect to get a sensible result, rather than having all those words ignored.

To see the list, query the table `information_schema.innodb_ft_default_stopword`. To define your own stopwords list used for all InnoDB tables, define a table with the same structure as `innodb_ft_default_stopword`, fill it with the desired stopwords, and set the value of the `innodb_ft_server_stopword_table` [1725] option to a value of the form `db_name/table_name` before creating the search index. To create special stopwords lists on a table-by-table basis, define other tables to hold these lists and specify the appropriate one in the `innodb_ft_user_stopword_table` [1726] option before creating the search index.

### Stopwords for MyISAM Search Indexes

In MySQL 5.6, the stopwords file is loaded and searched using `latin1` if `character_set_server` is `ucs2`, `utf16`, `utf16le`, or `utf32`. If any table was created with `FULLTEXT` indexes while the server character set was `ucs2`, `utf16`, `utf16le`, or `utf32`, it should be repaired using this statement:

## Full-Text Stopwords

```
REPAIR TABLE tbl_name QUICK;
```

The following table shows the default list of stopwords for [MyISAM](#) search indexes. In a MySQL source distribution, you can find this list in the `storage/myisam/ft_static.c` file.

a's	able	about	above	according
accordingly	across	actually	after	afterwards
again	against	ain't	all	allow
allows	almost	alone	along	already
also	although	always	am	among
amongst	an	and	another	any
anybody	anyhow	anyone	anything	anyway
anyways	anywhere	apart	appear	appreciate
appropriate	are	aren't	around	as
aside	ask	asking	associated	at
available	away	awfully	be	became
because	become	becomes	becoming	been
before	beforehand	behind	being	believe
below	beside	besides	best	better
between	beyond	both	brief	but
by	c'mon	c's	came	can
can't	cannot	cant	cause	causes
certain	certainly	changes	clearly	co
com	come	comes	concerning	consequently
consider	considering	contain	containing	contains
corresponding	could	couldn't	course	currently
definitely	described	despite	did	didn't
different	do	does	doesn't	doing
don't	done	down	downwards	during
each	edu	eg	eight	either
else	elsewhere	enough	entirely	especially
et	etc	even	ever	every
everybody	everyone	everything	everywhere	ex
exactly	example	except	far	few
fifth	first	five	followed	following
follows	for	former	formerly	forth
four	from	further	furthermore	get
gets	getting	given	gives	go
goes	going	gone	got	gotten
greetings	had	hadn't	happens	hardly
has	hasn't	have	haven't	having
he	he's	hello	help	hence
her	here	here's	hereafter	hereby
herein	hereupon	hers	herself	hi

Full-Text Stopwords

him	himself	his	hither	hopefully
how	howbeit	however	i'd	i'll
i'm	i've	ie	if	ignored
immediate	in	inasmuch	inc	indeed
indicate	indicated	indicates	inner	insofar
instead	into	inward	is	isn't
it	it'd	it'll	it's	its
itself	just	keep	keeps	kept
know	known	knows	last	lately
later	latter	latterly	least	less
lest	let	let's	like	liked
likely	little	look	looking	looks
ltd	mainly	many	may	maybe
me	mean	meanwhile	merely	might
more	moreover	most	mostly	much
must	my	myself	name	namely
nd	near	nearly	necessary	need
needs	neither	never	nevertheless	new
next	nine	no	nobody	non
none	noone	nor	normally	not
nothing	novel	now	nowhere	obviously
of	off	often	oh	ok
okay	old	on	once	one
ones	only	onto	or	other
others	otherwise	ought	our	ours
ourselves	out	outside	over	overall
own	particular	particularly	per	perhaps
placed	please	plus	possible	presumably
probably	provides	que	quite	qv
rather	rd	re	really	reasonably
regarding	regardless	regards	relatively	respectively
right	said	same	saw	say
saying	says	second	secondly	see
seeing	seem	seemed	seeming	seems
seen	self	selves	sensible	sent
serious	seriously	seven	several	shall
she	should	shouldn't	since	six
so	some	somebody	somehow	someone
something	sometime	sometimes	somewhat	somewhere
soon	sorry	specified	specify	specifying
still	sub	such	sup	sure
t's	take	taken	tell	tends

th	than	thank	thanks	thanx
that	that's	thats	the	their
theirs	them	themselves	then	thence
there	there's	thereafter	thereby	therefore
therein	theres	thereupon	these	they
they'd	they'll	they're	they've	think
third	this	thorough	thoroughly	those
though	three	through	throughout	thru
thus	to	together	too	took
toward	towards	tried	tries	truly
try	trying	twice	two	un
under	unfortunately	unless	unlikely	until
unto	up	upon	us	use
used	useful	uses	using	usually
value	various	very	via	viz
vs	want	wants	was	wasn't
way	we	we'd	we'll	we're
we've	welcome	well	went	were
weren't	what	what's	whatever	when
whence	whenever	where	where's	whereafter
whereas	whereby	wherein	whereupon	wherever
whether	which	while	whither	who
who's	whoever	whole	whom	whose
why	will	willing	wish	with
within	without	won't	wonder	would
wouldn't	yes	yet	you	you'd
you'll	you're	you've	your	yours
yourself	yourselves	zero		

## 12.9.5. Full-Text Restrictions

- Full-text searches are supported for [InnoDB](#) and [MyISAM](#) tables only. `FULLTEXT` index support for [InnoDB](#) tables requires MySQL 5.6.4 or higher.
- Full-text searches are not supported for partitioned tables. See [Section 17.6, "Restrictions and Limitations on Partitioning"](#).
- Full-text searches can be used with most multi-byte character sets. The exception is that for Unicode, the `utf8` character set can be used, but not the `ucs2` character set. Although `FULLTEXT` indexes on `ucs2` columns cannot be used, you can perform `IN BOOLEAN MODE` searches on a `ucs2` column that has no such index.

The remarks for `utf8` also apply to `utf8mb4`, and the remarks for `ucs2` also apply to `utf16`, `utf16le`, and `utf32`.

- Ideographic languages such as Chinese and Japanese do not have word delimiters. Therefore, the `FULLTEXT` parser *cannot determine where words begin and end in these and other such languages*. The implications of this and some workarounds for the problem are described in [Section 12.9, "Full-Text Search Functions"](#).

- Although the use of multiple character sets within a single table is supported, all columns in a `FULLTEXT` index must use the same character set and collation.
- The `MATCH( )` [1229] column list must match exactly the column list in some `FULLTEXT` index definition for the table, unless this `MATCH( )` [1229] is `IN BOOLEAN MODE` on a `MyISAM` table. For `MyISAM` tables only, boolean-mode searches can be done on nonindexed columns, although they are likely to be slow.
- The argument to `AGAINST( )` must be a string value that is constant during query evaluation. This rules out, for example, a table column because that can differ for each row.
- Index hints are more limited for `FULLTEXT` searches than for non-`FULLTEXT` searches. See [Section 13.2.9.3, “Index Hint Syntax”](#).

## 12.9.6. Fine-Tuning MySQL Full-Text Search

MySQL's full-text search capability has few user-tunable parameters. You can exert more control over full-text searching behavior if you have a MySQL source distribution because some changes require source code modifications. See [Section 2.9, “Installing MySQL from Source”](#).

Note that full-text search is carefully tuned for effectiveness. Modifying the default behavior in most cases can actually decrease effectiveness. *Do not alter the MySQL sources unless you know what you are doing.*

Most full-text variables described in this section must be set at server startup time. A server restart is required to change them; they cannot be modified while the server is running.

Some variable changes require that you rebuild the `FULLTEXT` indexes in your tables. Instructions for doing so are given later in this section.

- The minimum and maximum lengths of words to be indexed are defined by the `innodb_ft_min_token_size` [1725] and `innodb_ft_max_token_size` [1724] for InnoDB search indexes, and `ft_min_word_len` [484] and `ft_max_word_len` [484] for `MyISAM` ones. After changing any of these options, rebuild your `FULLTEXT` indexes for the change to take effect. For example, to make two-character words searchable, you could put the following lines in an option file:

```
[mysqld]
innodb_ft_min_token_size=2
ft_min_word_len=2
```

Then restart the server and rebuild your `FULLTEXT` indexes. For `MyISAM` tables, note particularly the remarks regarding `myisamchk` in the instructions following this list.

- To override the default stopword list, set the `ft_stopword_file` [485] system variable. (See [Section 5.1.4, “Server System Variables”](#).) The variable value should be the path name of the file containing the stopword list, or the empty string to disable stopword filtering. The server looks for the file in the data directory unless an absolute path name is given to specify a different directory. After changing the value of this variable or the contents of the stopword file, restart the server and rebuild your `FULLTEXT` indexes.

The stopword list is free-form, separating stopwords with any nonalphanumeric character such as newline, space, or comma. Exceptions are the underscore character (“\_”) and a single apostrophe (“'”) which are treated as part of a word. The character set of the stopword list is the server's default character set; see [Section 10.1.3.1, “Server Character Set and Collation”](#).

- The 50% threshold for natural language searches is determined by the particular weighting scheme chosen. To disable it, look for the following line in `storage/myisam/ftdefs.h`:

```
#define GWS_IN_USE GWS_PROB
```

Change that line to this:

```
#define GWS_IN_USE GWS_FREQ
```

Then recompile MySQL. There is no need to rebuild the indexes in this case.



### Note

By making this change, you *severely* decrease MySQL's ability to provide adequate relevance values for the `MATCH( )` [1229] function. If you really need to search for such common words, it would be better to search using `IN BOOLEAN MODE` instead, which does not observe the 50% threshold.

- To change the operators used for boolean full-text searches on `MyISAM` tables, set the `ft_boolean_syntax` [483] system variable. (`InnoDB` does not have an equivalent setting.) This variable can be changed while the server is running, but you must have the `SUPER` [756] privilege to do so. No rebuilding of indexes is necessary in this case. See [Section 5.1.4, “Server System Variables”](#), which describes the rules governing how to set this variable.
- You can change the set of characters that are considered word characters in several ways, as described in the following list. After making the modification, rebuild the indexes for each table that contains any `FULLTEXT` indexes. Suppose that you want to treat the hyphen character ('-') as a word character. Use one of these methods:
  - Modify the MySQL source: In `storage/myisam/ftdefs.h`, see the `true_word_char( )` and `misc_word_char( )` macros. Add '-' to one of those macros and recompile MySQL.
  - Modify a character set file: This requires no recompilation. The `true_word_char( )` macro uses a “character type” table to distinguish letters and numbers from other characters. . You can edit the contents of the `<ctype><map>` array in one of the character set XML files to specify that '-' is a “letter.” Then use the given character set for your `FULLTEXT` indexes. For information about the `<ctype><map>` array format, see [Section 10.3.1, “Character Definition Arrays”](#).
  - Add a new collation for the character set used by the indexed columns, and alter the columns to use that collation. For general information about adding collations, see [Section 10.4, “Adding a Collation to a Character Set”](#). For an example specific to full-text indexing, see [Section 12.9.7, “Adding a Collation for Full-Text Indexing”](#).

If you modify full-text variables that affect indexing (`innodb_ft_min_token_size` [1725], `innodb_ft_max_token_size` [1724], `innodb_ft_server_stopword_table` [1725], `innodb_ft_user_stopword_table` [1726], `innodb_ft_enable_stopword` [1724], `ft_min_word_len` [484], `ft_max_word_len` [484], or `ft_stopword_file` [485]), or if you change the stopword file itself, you must rebuild your `FULLTEXT` indexes after making the changes and restarting the server. To rebuild the indexes in this case, it is sufficient to do a `QUICK` repair operation:

```
mysql> REPAIR TABLE tbl_name QUICK;
```

Alternatively, use `ALTER TABLE` with the `DROP INDEX` and `ADD INDEX` options to drop and re-create each `FULLTEXT` index. In some cases, this may be faster than a repair operation.

Each table that contains any `FULLTEXT` index must be repaired as just shown. Otherwise, queries for the table may yield incorrect results, and modifications to the table will cause the server to see the table as corrupt and in need of repair.

Note that if you use `myisamchk` to perform an operation that modifies table indexes (such as repair or analyze), the `FULLTEXT` indexes are rebuilt using the *default* full-text parameter values for minimum word length, maximum word length, and stopword file unless you specify otherwise. This can result in queries failing.

The problem occurs because these parameters are known only by the server. They are not stored in `MyISAM` index files. To avoid the problem if you have modified the minimum or maximum word length or stopword file values used by the server, specify the same `ft_min_word_len` [484],



`ft_max_word_len` [484], and `ft_stopword_file` [485] values for `myisamchk` that you use for `mysqld`. For example, if you have set the minimum word length to 3, you can repair a table with `myisamchk` like this:

```
shell> myisamchk --recover --ft_min_word_len=3 tbl_name.MYI
```

To ensure that `myisamchk` and the server use the same values for full-text parameters, place each one in both the `[mysqld]` and `[myisamchk]` sections of an option file:

```
[mysqld]
ft_min_word_len=3

[myisamchk]
ft_min_word_len=3
```

An alternative to using `myisamchk` for index modification is to use the `REPAIR TABLE`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `ALTER TABLE` statements. These statements are performed by the server, which knows the proper full-text parameter values to use.

## 12.9.7. Adding a Collation for Full-Text Indexing

This section describes how to add a new collation for full-text searches. The sample collation is like `latin1_swedish_ci` but treats the '-' character as a letter rather than as a punctuation character so that it can be indexed as a word character. General information about adding collations is given in [Section 10.4, "Adding a Collation to a Character Set"](#); it is assumed that you have read it and are familiar with the files involved.

To add a collation for full-text indexing, use this procedure:

1. Add a collation to the `Index.xml` file. The collation ID must be unused, so choose a value different from 1000 if that ID is already taken on your system.

```
<charset name="latin1">
...
<collation name="latin1_fulltext_ci" id="1000"/>
</charset>
```

2. Declare the sort order for the collation in the `latin1.xml` file. In this case, the order can be copied from `latin1_swedish_ci`:

```
<collation name="latin1_fulltext_ci">
<map>
00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
10 11 12 13 14 15 16 17 18 19 1A 1B 1C 1D 1E 1F
20 21 22 23 24 25 26 27 28 29 2A 2B 2C 2D 2E 2F
30 31 32 33 34 35 36 37 38 39 3A 3B 3C 3D 3E 3F
40 41 42 43 44 45 46 47 48 49 4A 4B 4C 4D 4E 4F
50 51 52 53 54 55 56 57 58 59 5A 5B 5C 5D 5E 5F
60 41 42 43 44 45 46 47 48 49 4A 4B 4C 4D 4E 4F
50 51 52 53 54 55 56 57 58 59 5A 7B 7C 7D 7E 7F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
90 91 92 93 94 95 96 97 98 99 9A 9B 9C 9D 9E 9F
A0 A1 A2 A3 A4 A5 A6 A7 A8 A9 AA AB AC AD AE AF
B0 B1 B2 B3 B4 B5 B6 B7 B8 B9 BA BB BC BD BE BF
41 41 41 41 5C 5B 5C 43 45 45 45 45 49 49 49 49
44 4E 4F 4F 4F 4F 5D D7 D8 55 55 55 59 59 DE DF
41 41 41 41 5C 5B 5C 43 45 45 45 45 49 49 49 49
44 4E 4F 4F 4F 4F 5D F7 D8 55 55 55 59 59 DE FF
</map>
</collation>
```

3. Modify the `ctype` array in `latin1.xml`. Change the value corresponding to 0x2D (which is the code for the '-' character) from 10 (punctuation) to 01 (small letter). In the following array, this is the element in the fourth row down, third value from the end.

```
<ctype>
```

```

<map>
00
20 20 20 20 20 20 20 20 20 28 28 28 28 28 20 20
20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20
48 10 10 10 10 10 10 10 10 10 10 10 10 01 10 10
84 84 84 84 84 84 84 84 84 84 10 10 10 10 10 10
10 81 81 81 81 81 81 01 01 01 01 01 01 01 01 01
01 01 01 01 01 01 01 01 01 01 01 01 10 10 10 10
10 82 82 82 82 82 82 02 02 02 02 02 02 02 02 02
02 02 02 02 02 02 02 02 02 02 02 10 10 10 10 20
10 00 10 02 10 10 10 10 10 10 01 10 01 00 01 00
00 10 10 10 10 10 10 10 10 10 02 10 02 00 02 01
48 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01
01 01 01 01 01 01 01 10 01 01 01 01 01 01 01 02
02 02 02 02 02 02 02 02 02 02 02 02 02 02 02 02
02 02 02 02 02 02 02 10 02 02 02 02 02 02 02 02
</map>
</ctype>

```

4. Restart the server.
5. To employ the new collation, include it in the definition of columns that are to use it:

```

mysql> DROP TABLE IF EXISTS t1;
Query OK, 0 rows affected (0.13 sec)

mysql> CREATE TABLE t1 (
  a TEXT CHARACTER SET latin1 COLLATE latin1_fulltext_ci,
  FULLTEXT INDEX(a)
) ENGINE=InnoDB;
Query OK, 0 rows affected (0.47 sec)

```

6. Test the collation to verify that hyphen is considered as a word character:

```

mysql> INSERT INTO t1 VALUES ('----'),('....'),('abcd');
Query OK, 3 rows affected (0.22 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM t1 WHERE MATCH a AGAINST ('----' IN BOOLEAN MODE);
+-----+
| a      |
+-----+
| ----  |
+-----+
1 row in set (0.00 sec)

```

## 12.10. Cast Functions and Operators

Table 12.14. Cast Functions

Name	Description
<a href="#">BINARY [1244]</a>	Cast a string to a binary string
<a href="#">CAST() [1245]</a>	Cast a value as a certain type
<a href="#">CONVERT() [1245]</a>	Cast a value as a certain type

- [BINARY \[1244\]](#)

The [BINARY \[1244\]](#) operator casts the string following it to a binary string. This is an easy way to force a column comparison to be done byte by byte rather than character by character. This causes the comparison to be case sensitive even if the column is not defined as [BINARY](#) or [BLOB](#). [BINARY \[1244\]](#) also causes trailing spaces to be significant.

```

mysql> SELECT 'a' = 'A';
-> 1

```

```
mysql> SELECT BINARY 'a' = 'A';
-> 0
mysql> SELECT 'a' = 'a ';
-> 1
mysql> SELECT BINARY 'a' = 'a ';
-> 0
```

In a comparison, `BINARY` [1244] affects the entire operation; it can be given before either operand with the same result.

`BINARY str` [1244] is shorthand for `CAST(str AS BINARY)` [1245].

Note that in some contexts, if you cast an indexed column to `BINARY`, MySQL is not able to use the index efficiently.

- `CAST(expr AS type)` [1245]

The `CAST()` [1245] function takes an expression of any type and produces a result value of a specified type, similar to `CONVERT()` [1245]. See the description of `CONVERT()` [1245] for more information.

- `CONVERT(expr, type)` [1245], `CONVERT(expr USING transcoding_name)` [1245]

The `CONVERT()` [1245] and `CAST()` [1245] functions take an expression of any type and produce a result value of a specified type.

The `type` for the result can be one of the following values:

- `BINARY[ (N) ]`
- `CHAR[ (N) ]`
- `DATE`
- `DATETIME`
- `DECIMAL[ (M[ , D] ) ]`
- `SIGNED [ INTEGER ]`
- `TIME`
- `UNSIGNED [ INTEGER ]`

`BINARY` produces a string with the `BINARY` data type. See [Section 11.4.2, “The `BINARY` and `VARBINARY` Types”](#) for a description of how this affects comparisons. If the optional length `N` is given, `BINARY(N)` causes the cast to use no more than `N` bytes of the argument. Values shorter than `N` bytes are padded with `0x00` bytes to a length of `N`.

`CHAR(N)` [1177] causes the cast to use no more than `N` characters of the argument.

`CAST()` [1245] and `CONVERT(... USING ...)` [1245] are standard SQL syntax. The non-`USING` form of `CONVERT()` [1245] is ODBC syntax.

`CONVERT()` [1245] with `USING` is used to convert data between different character sets. In MySQL, transcoding names are the same as the corresponding character set names. For example, this statement converts the string `'abc'` in the default character set to the corresponding string in the `utf8` character set:

```
SELECT CONVERT('abc' USING utf8);
```

Normally, you cannot compare a `BLOB` value or other binary string in case-insensitive fashion because binary strings have no character set, and thus no concept of lettercase. To perform a case-

insensitive comparison, use the `CONVERT()` [1245] function to convert the value to a nonbinary string. Comparisons of the result use the string collation. For example, if the character set of the result has a case-insensitive collation, a `LIKE` [1190] operation is not case sensitive:

```
SELECT 'A' LIKE CONVERT(blob_col USING latin1) FROM tbl_name;
```

To use a different character set, substitute its name for `latin1` in the preceding statement. To specify a particular collation for the converted string, use a `COLLATE` clause following the `CONVERT()` [1245] call, as described in Section 10.1.9.2, “`CONVERT()` and `CAST()`”. For example, to use `latin1_german1_ci`:

```
SELECT 'A' LIKE CONVERT(blob_col USING latin1) COLLATE latin1_german1_ci
FROM tbl_name;
```

`CONVERT()` [1245] can be used more generally for comparing strings that are represented in different character sets.

`LOWER()` [1182] (and `UPPER()` [1187]) are ineffective when applied to binary strings (`BINARY`, `VARBINARY`, `BLOB`). To perform lettercase conversion, convert the string to a nonbinary string:

```
mysql> SET @str = BINARY 'New York';
mysql> SELECT LOWER(@str), LOWER(CONVERT(@str USING latin1));
+-----+-----+
| LOWER(@str) | LOWER(CONVERT(@str USING latin1)) |
+-----+-----+
| New York   | new york                           |
+-----+-----+
```

The cast functions are useful when you want to create a column with a specific type in a `CREATE TABLE ... SELECT` statement:

```
CREATE TABLE new_table SELECT CAST('2000-01-01' AS DATE);
```

The functions also can be useful for sorting `ENUM` columns in lexical order. Normally, sorting of `ENUM` columns occurs using the internal numeric values. Casting the values to `CHAR` results in a lexical sort:

```
SELECT enum_col FROM tbl_name ORDER BY CAST(enum_col AS CHAR);
```

`CAST(str AS BINARY)` [1245] is the same thing as `BINARY str` [1244]. `CAST(expr AS CHAR)` [1245] treats the expression as a string with the default character set.

`CAST()` [1245] also changes the result if you use it as part of a more complex expression such as `CONCAT('Date: ',CAST(NOW() AS DATE))` [1178].

You should not use `CAST()` [1245] to extract data in different formats but instead use string functions like `LEFT()` [1181] or `EXTRACT()` [1217]. See Section 12.7, “Date and Time Functions”.

To cast a string to a numeric value in numeric context, you normally do not have to do anything other than to use the string value as though it were a number:

```
mysql> SELECT 1+'1';
-> 2
```

If you use a string in an arithmetic operation, it is converted to a floating-point number during expression evaluation.

If you use a number in string context, the number automatically is converted to a string:

```
mysql> SELECT CONCAT('hello you ',2);
-> 'hello you 2'
```

For information about implicit conversion of numbers to strings, see Section 12.2, “Type Conversion in Expression Evaluation”.

MySQL supports arithmetic with both signed and unsigned 64-bit values. If you are using numeric operators (such as + [1200] or - [1200]) and one of the operands is an unsigned integer, the result is unsigned by default (see [Section 12.6.1, “Arithmetic Operators”](#)). You can override this by using the `SIGNED` or `UNSIGNED` cast operator to cast a value to a signed or unsigned 64-bit integer, respectively.

```
mysql> SELECT CAST(1-2 AS UNSIGNED)
      -> 18446744073709551615
mysql> SELECT CAST(CAST(1-2 AS UNSIGNED) AS SIGNED);
      -> -1
```

If either operand is a floating-point value, the result is a floating-point value and is not affected by the preceding rule. (In this context, `DECIMAL` column values are regarded as floating-point values.)

```
mysql> SELECT CAST(1 AS UNSIGNED) - 2.0;
      -> -1.0
```

The SQL mode affects the result of conversion operations. Examples:

- If you convert a “zero” date string to a date, `CONVERT()` [1245] and `CAST()` [1245] return `NULL` and produce a warning when the `NO_ZERO_DATE` [606] SQL mode is enabled.
- For integer subtraction, if the `NO_UNSIGNED_SUBTRACTION` [605] SQL mode is enabled, the subtraction result is signed even if any operand is unsigned.

For more information, see [Section 5.1.7, “Server SQL Modes”](#).

## 12.11. XML Functions

**Table 12.15. XML Functions**

Name	Description
<code>ExtractValue()</code> [1249]	Extracts a value from an XML string using XPath notation
<code>UpdateXML()</code> [1251]	Return replaced XML fragment

This section discusses XML and related functionality in MySQL.



### Note

It is possible to obtain XML-formatted output from MySQL in the `mysql` and `mysqldump` clients by invoking them with the `--xml` [254] option. See [Section 4.5.1, “mysql — The MySQL Command-Line Tool”](#), and [Section 4.5.4, “mysqldump — A Database Backup Program”](#).

Two functions providing basic XPath 1.0 (XML Path Language, version 1.0) capabilities are available. Some basic information about XPath syntax and usage is provided later in this section; however, an in-depth discussion of these topics is beyond the scope of this Manual, and you should refer to the [XML Path Language \(XPath\) 1.0 standard](#) for definitive information. A useful resource for those new to XPath or who desire a refresher in the basics is the [Zvon.org XPath Tutorial](#), which is available in several languages.



### Note

These functions remain under development. We continue to improve these and other aspects of XML and XPath functionality in MySQL 5.6 and onwards. You may discuss these, ask questions about them, and obtain help from other users with them in the [MySQL XML User Forum](#).

XPath expressions used with these functions support user variables and local stored program variables. User variables are weakly checked; variables local to stored programs are strongly checked (see also Bug #26518):

- **User variables (weak checking).** Variables using the syntax `$_variable_name` (that is, user variables) are not checked. No warnings or errors are issued by the server if a variable has the wrong type or has previously not been assigned a value. This also means the user is fully responsible for any typographical errors, since no warnings will be given if (for example) `$_myvairable` is used where `$_myvariable` was intended.

Example:

```
mysql> SET @xml = '<a><b>X</b><b>Y</b></a>';
Query OK, 0 rows affected (0.00 sec)

mysql> SET @i =1, @j = 2;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @i, ExtractValue(@xml, '//b[$@i]');
+-----+-----+
| @i | ExtractValue(@xml, '//b[$@i]') |
+-----+-----+
| 1 | X |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT @j, ExtractValue(@xml, '//b[$@j]');
+-----+-----+
| @j | ExtractValue(@xml, '//b[$@j]') |
+-----+-----+
| 2 | Y |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT @k, ExtractValue(@xml, '//b[$@k]');
+-----+-----+
| @k | ExtractValue(@xml, '//b[$@k]') |
+-----+-----+
| NULL | |
+-----+-----+
1 row in set (0.00 sec)
```

- **Variables in stored programs (strong checking).** Variables using the syntax `$_variable_name` can be declared and used with these functions when they are called inside stored programs. Such variables are local to the stored program in which they are defined, and are strongly checked for type and value.

Example:

```
mysql> DELIMITER |

mysql> CREATE PROCEDURE myproc ()
-> BEGIN
-> DECLARE i INT DEFAULT 1;
-> DECLARE xml VARCHAR(25) DEFAULT '<a>X</a><a>Y</a><a>Z</a>';
->
-> WHILE i < 4 DO
-> SELECT xml, i, ExtractValue(xml, '//a[$i]');
-> SET i = i+1;
-> END WHILE;
-> END |
Query OK, 0 rows affected (0.01 sec)

mysql> DELIMITER ;

mysql> CALL myproc;
+-----+-----+-----+
| xml | i | ExtractValue(xml, '//a[$i]') |
+-----+-----+-----+
| <a>X</a><a>Y</a><a>Z</a> | 1 | X |
+-----+-----+-----+
1 row in set (0.00 sec)
```

```

+-----+-----+-----+
| xml          | i | ExtractValue(xml, '//a[$i]') |
+-----+-----+-----+
| <a>X</a><a>Y</a><a>Z</a> | 2 | Y
+-----+-----+-----+
1 row in set (0.01 sec)

+-----+-----+-----+
| xml          | i | ExtractValue(xml, '//a[$i]') |
+-----+-----+-----+
| <a>X</a><a>Y</a><a>Z</a> | 3 | Z
+-----+-----+-----+
1 row in set (0.01 sec)

```

**Parameters.** Variables used in XPath expressions inside stored routines that are passed in as parameters are also subject to strong checking.

Expressions containing user variables or variables local to stored programs must otherwise (except for notation) conform to the rules for XPath expressions containing variables as given in the XPath 1.0 specification.



### Note

Currently, a user variable used to store an XPath expression is treated as an empty string. Because of this, it is not possible to store an XPath expression as a user variable. (Bug #32911)

- [ExtractValue\(xml\\_frag, xpath\\_expr\) \[1249\]](#)

[ExtractValue\(\)](#) [1249] takes two string arguments, a fragment of XML markup *xml\_frag* and an XPath expression *xpath\_expr* (also known as a *locator*); it returns the text (CDATA) of the first text node which is a child of the elements or elements matched by the XPath expression. In MySQL 5.6.6 and earlier, the XPath expression could contain at most 127 characters. This limitation was lifted in MySQL 5.6.7. (Bug #13007062, Bug#62429)

Using this function is the equivalent of performing a match using the *xpath\_expr* after appending `/text()`. In other words, [ExtractValue\('<a><b>Sakila</b></a>', '/a/b'\)](#) [1249] and [ExtractValue\('<a><b>Sakila</b></a>', '/a/b/text\(\)'\)](#) [1249] produce the same result.

If multiple matches are found, the content of the first child text node of each matching element is returned (in the order matched) as a single, space-delimited string.

If no matching text node is found for the expression (including the implicit `/text()`)—for whatever reason, as long as *xpath\_expr* is valid, and *xml\_frag* consists of elements which are properly nested and closed—an empty string is returned. No distinction is made between a match on an empty element and no match at all. This is by design.

If you need to determine whether no matching element was found in *xml\_frag* or such an element was found but contained no child text nodes, you should test the result of an expression that uses the XPath `count()` function. For example, both of these statements return an empty string, as shown here:

```

mysql> SELECT ExtractValue('<a><b></b></a>', '/a/b');
+-----+
| ExtractValue('<a><b></b></a>', '/a/b') |
+-----+
|                                     |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue('<a><c></c></a>', '/a/b');
+-----+
| ExtractValue('<a><c></c></a>', '/a/b') |
+-----+
|                                     |
+-----+

```

```
+-----+
1 row in set (0.00 sec)
```

However, you can determine whether there was actually a matching element using the following:

```
mysql> SELECT ExtractValue('<a><b/></a>', 'count(/a/b)');
+-----+
| ExtractValue('<a><b/></a>', 'count(/a/b)') |
+-----+
| 1 |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue('<a><c/></a>', 'count(/a/b)');
+-----+
| ExtractValue('<a><c/></a>', 'count(/a/b)') |
+-----+
| 0 |
+-----+
1 row in set (0.01 sec)
```



### Important

`ExtractValue()` [1249] returns only `CDATA`, and does not return any tags that might be contained within a matching tag, nor any of their content (see the result returned as `val1` in the following example).

```
mysql> SELECT
-> ExtractValue('<a>ccc<b>ddd</b></a>', '/a') AS val1,
-> ExtractValue('<a>ccc<b>ddd</b></a>', '/a/b') AS val2,
-> ExtractValue('<a>ccc<b>ddd</b></a>', '//b') AS val3,
-> ExtractValue('<a>ccc<b>ddd</b></a>', '/b') AS val4,
-> ExtractValue('<a>ccc<b>ddd</b><b>eee</b></a>', '//b') AS val5;
+-----+-----+-----+-----+-----+
| val1 | val2 | val3 | val4 | val5 |
+-----+-----+-----+-----+-----+
| ccc | ddd | ddd |      | ddd eee |
+-----+-----+-----+-----+-----+
```

This function uses the current SQL collation for making comparisons with `contains()`, performing the same collation aggregation as other string functions (such as `CONCAT()` [1178]), in taking into account the collation coercibility of their arguments; see [Section 10.1.7.5, “Collation of Expressions”](#), for an explanation of the rules governing this behavior.

(Previously, binary—that is, case-sensitive—comparison was always used.)

`NULL` is returned if `xml_frag` contains elements which are not properly nested or closed, and a warning is generated, as shown in this example:

```
mysql> SELECT ExtractValue('<a>c</a><b', '//a');
+-----+
| ExtractValue('<a>c</a><b', '//a') |
+-----+
| NULL |
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Warning | 1523 | Incorrect XML value: 'parse error at line 1 pos 11: END-OF-INPUT unexpected ('>' wanted |
+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue('<a>c</a><b/>', '//a');
+-----+
| ExtractValue('<a>c</a><b/>', '//a') |
+-----+
```



```
+-----+
| c                                           |
+-----+
1 row in set (0.00 sec)
```

- `UpdateXML(xml_target, xpath_expr, new_xml)` [1251]

This function replaces a single portion of a given fragment of XML markup *xml\_target* with a new XML fragment *new\_xml*, and then returns the changed XML. The portion of *xml\_target* that is replaced matches an XPath expression *xpath\_expr* supplied by the user. In MySQL 5.6.6 and earlier, the XPath expression could contain at most 127 characters. This limitation is lifted in MySQL 5.6.7. (Bug #13007062, Bug #62429)

If no expression matching *xpath\_expr* is found, or if multiple matches are found, the function returns the original *xml\_target* XML fragment. All three arguments should be strings.

```
mysql> SELECT
-> UpdateXML('<a><b>ccc</b><d></d></a>', '/a', '<e>fff</e>') AS val1,
-> UpdateXML('<a><b>ccc</b><d></d></a>', '/b', '<e>fff</e>') AS val2,
-> UpdateXML('<a><b>ccc</b><d></d></a>', '//b', '<e>fff</e>') AS val3,
-> UpdateXML('<a><b>ccc</b><d></d></a>', '/a/d', '<e>fff</e>') AS val4,
-> UpdateXML('<a><d></d><b>ccc</b><d></d></a>', '/a/d', '<e>fff</e>') AS val5
-> \G

***** 1. row *****
val1: <e>fff</e>
val2: <a><b>ccc</b><d></d></a>
val3: <a><e>fff</e><d></d></a>
val4: <a><b>ccc</b><e>fff</e></a>
val5: <a><d></d><b>ccc</b><d></d></a>
```



### Note

A discussion in depth of XPath syntax and usage are beyond the scope of this Manual. Please see the [XML Path Language \(XPath\) 1.0 specification](#) for definitive information. A useful resource for those new to XPath or who are wishing a refresher in the basics is the [Zvon.org XPath Tutorial](#), which is available in several languages.

Descriptions and examples of some basic XPath expressions follow:

- `/tag`

Matches `<tag/>` if and only if `<tag/>` is the root element.

Example: `/a` has a match in `<a><b></b></a>` because it matches the outermost (root) tag. It does not match the inner `a` element in `<b><a></a></b>` because in this instance it is the child of another element.

- `/tag1/tag2`

Matches `<tag2/>` if and only if it is a child of `<tag1/>`, and `<tag1/>` is the root element.

Example: `/a/b` matches the `b` element in the XML fragment `<a><b></b></a>` because it is a child of the root element `a`. It does not have a match in `<b><a></a></b>` because in this case, `b` is the root element (and hence the child of no other element). Nor does the XPath expression have a match in `<a><c><b></b></c></a>`; here, `b` is a descendant of `a`, but not actually a child of `a`.

This construct is extendable to three or more elements. For example, the XPath expression `/a/b/c` matches the `c` element in the fragment `<a><b><c></c></b></a>`.

- `//tag`

Matches any instance of `<tag>`.

Example: `//a` matches the `a` element in any of the following: `<a><b><c/></b></a>`; `<c><a><b/></a></b>`; `<c><b><a/></b></c>`.

`//` can be combined with `/`. For example, `//a/b` matches the `b` element in either of the fragments `<a><b/></a>` or `<a><b><c/></b></a>`



### Note

`//tag` is the equivalent of `/descendant-or-self::*/tag`. A common error is to confuse this with `/descendant-or-self::tag`, although the latter expression can actually lead to very different results, as can be seen here:

```
mysql> SET @xml = '<a><b><c>w</c><b>x</b><d>y</d>z</b></a>';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @xml;
+-----+
| @xml |
+-----+
| <a><b><c>w</c><b>x</b><d>y</d>z</b></a> |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue(@xml, '//b[1]');
+-----+
| ExtractValue(@xml, '//b[1]') |
+-----+
| x z |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue(@xml, '//b[2]');
+-----+
| ExtractValue(@xml, '//b[2]') |
+-----+
| |
+-----+
1 row in set (0.01 sec)

mysql> SELECT ExtractValue(@xml, '/descendant-or-self::*b[1]');
+-----+
| ExtractValue(@xml, '/descendant-or-self::*b[1]') |
+-----+
| x z |
+-----+
1 row in set (0.06 sec)

mysql> SELECT ExtractValue(@xml, '/descendant-or-self::*b[2]');
+-----+
| ExtractValue(@xml, '/descendant-or-self::*b[2]') |
+-----+
| |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue(@xml, '/descendant-or-self::b[1]');
+-----+
| ExtractValue(@xml, '/descendant-or-self::b[1]') |
+-----+
| z |
+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue(@xml, '/descendant-or-self::b[2]');
+-----+
| ExtractValue(@xml, '/descendant-or-self::b[2]') |
+-----+
```

```
| x |
+-----+
1 row in set (0.00 sec)
```

- The `*` operator acts as a “wildcard” that matches any element. For example, the expression `/*/b` matches the `b` element in either of the XML fragments `<a><b/></a>` or `<c><b/></c>`. However, the expression does not produce a match in the fragment `<b><a/></b>` because `b` must be a child of some other element. The wildcard may be used in any position: The expression `/*/b/*` will match any child of a `b` element that is itself not the root element.
- You can match any of several locators using the `|` (UNION) operator. For example, the expression `//b|//c` matches all `b` and `c` elements in the XML target.
- It is also possible to match an element based on the value of one or more of its attributes. This done using the syntax `tag[@attribute="value"]`. For example, the expression `//b[@id="idB"]` matches the second `b` element in the fragment `<a><b id="idA"/><c/><b id="idB"/></a>`. To match against *any* element having `attribute="value"`, use the XPath expression `//*[attribute="value"]`.

To filter multiple attribute values, simply use multiple attribute-comparison clauses in succession. For example, the expression `//b[@c="x"][@d="y"]` matches the element `<b c="x" d="y"/>` occurring anywhere in a given XML fragment.

To find elements for which the same attribute matches any of several values, you can use multiple locators joined by the `|` operator. For example, to match all `b` elements whose `c` attributes have either of the values 23 or 17, use the expression `//b[@c="23"]|//b[@c="17"]`. You can also use the logical `or` operator for this purpose: `//b[@c="23" or @c="17"]`.



#### Note

The difference between `or` and `|` is that `or` joins conditions, while `|` joins result sets.

**XPath Limitations.** The XPath syntax supported by these functions is currently subject to the following limitations:

- Nodeset-to-nodeset comparison (such as `'/a/b[@c=@d]'`) is not supported.
- All of the standard XPath comparison operators are supported. (Bug #22823)
- Relative locator expressions are resolved in the context of the root node. For example, consider the following query and result:

```
mysql> SELECT ExtractValue(
->   '<a><b c="1">X</b><b c="2">Y</b></a>',
->   'a/b'
-> ) AS result;
+-----+
| result |
+-----+
| X Y    |
+-----+
1 row in set (0.03 sec)
```

In this case, the locator `a/b` resolves to `/a/b`.

Relative locators are also supported within predicates. In the following example, `d[../@c="1"]` is resolved as `/a/b[@c="1"]/d`:

```
mysql> SELECT ExtractValue(
->   '<a>
->     <b c="1"><d>X</d></b>
->     <b c="2"><d>X</d></b>
->   </a>',
->   'a/b/d[../@c="1"]')
```

```

-> AS result;
+-----+
| result |
+-----+
| x      |
+-----+
1 row in set (0.00 sec)

```

- Locators prefixed with expressions that evaluate as scalar values—including variable references, literals, numbers, and scalar function calls—are not permitted, and their use results in an error.
- The `::` operator is not supported in combination with node types such as the following:
  - `axis::comment()`
  - `axis::text()`
  - `axis::processing-instructions()`
  - `axis::node()`

However, name tests (such as `axis::name` and `axis::*`) are supported, as shown in these examples:

```

mysql> SELECT ExtractValue('<a><b>x</b><c>y</c></a>', '/a/child::b');
+-----+
| ExtractValue('<a><b>x</b><c>y</c></a>', '/a/child::b') |
+-----+
| x                                             |
+-----+
1 row in set (0.02 sec)

mysql> SELECT ExtractValue('<a><b>x</b><c>y</c></a>', '/a/child::*');
+-----+
| ExtractValue('<a><b>x</b><c>y</c></a>', '/a/child::*') |
+-----+
| x y                                           |
+-----+
1 row in set (0.01 sec)

```

- “Up-and-down” navigation is not supported in cases where the path would lead “above” the root element. That is, you cannot use expressions which match on descendants of ancestors of a given element, where one or more of the ancestors of the current element is also an ancestor of the root element (see Bug #16321).
- The following XPath functions are not supported, or have known issues as indicated:
  - `id()`
  - `lang()`
  - `local-name()`
  - `name()`
  - `namespace-uri()`
  - `normalize-space()`
  - `starts-with()`
  - `string()`
  - `substring-after()`
  - `substring-before()`

- `translate()`
- The following axes are not supported:
  - `following-sibling`
  - `following`
  - `preceding-sibling`
  - `preceding`

XPath expressions passed as arguments to `ExtractValue()` [1249] and `UpdateXML()` [1251] may contain the colon character (":") in element selectors, which enables their use with markup employing XML namespaces notation. For example:

```
mysql> SET @xml = '<a>111<b:c>222<d>333</d><e:f>444</e:f></b:c></a>';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT ExtractValue(@xml, '//e:f');
+-----+
| ExtractValue(@xml, '//e:f') |
+-----+
| 444                          |
+-----+
1 row in set (0.00 sec)

mysql> SELECT UpdateXML(@xml, '//b:c', '<g:h>555</g:h>');
+-----+
| UpdateXML(@xml, '//b:c', '<g:h>555</g:h>') |
+-----+
| <a>111<g:h>555</g:h></a>                  |
+-----+
1 row in set (0.00 sec)
```

This is similar in some respects to what is permitted by [Apache Xalan](#) and some other parsers, and is much simpler than requiring namespace declarations or the use of the `namespace-uri()` and `local-name()` functions.

**Error handling.** For both `ExtractValue()` [1249] and `UpdateXML()` [1251], the XPath locator used must be valid and the XML to be searched must consist of elements which are properly nested and closed. If the locator is invalid, an error is generated:

```
mysql> SELECT ExtractValue('<a>c</a><b/>', '//*[@a]');
ERROR 1105 (HY000): XPATH syntax error: '//*[@a]'
```

If `xml_frag` does not consist of elements which are properly nested and closed, `NULL` is returned and a warning is generated, as shown in this example:

```
mysql> SELECT ExtractValue('<a>c</a><b', '//*[@a]');
+-----+
| ExtractValue('<a>c</a><b', '//*[@a]') |
+-----+
| NULL                                |
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message
+-----+-----+-----+
| Warning | 1523 | Incorrect XML value: 'parse error at line 1 pos 11: END-OF-INPUT unexpected ('>' wan
+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT ExtractValue('<a>c</a><b/>', '//*[@a]');
```

```

+-----+
| ExtractValue('<a>c</a><b/>', '//a') |
+-----+
| c |
+-----+
1 row in set (0.00 sec)

```



### Important

The replacement XML used as the third argument to `UpdateXML()` [1251] is *not* checked to determine whether it consists solely of elements which are properly nested and closed.

**XPath Injection.** *code injection* occurs when malicious code is introduced into the system to gain unauthorized access to privileges and data. It is based on exploiting assumptions made by developers about the type and content of data input from users. XPath is no exception in this regard.

A common scenario in which this can happen is the case of application which handles authorization by matching the combination of a login name and password with those found in an XML file, using an XPath expression like this one:

```
//user[login/text()='neapolitan' and password/text()='1c3cr34m']/attribute::id
```

This is the XPath equivalent of an SQL statement like this one:

```
SELECT id FROM users WHERE login='neapolitan' AND password='1c3cr34m';
```

A PHP application employing XPath might handle the login process like this:

```

<?php

$file      = "users.xml";

$login     = $_POST["login"];
$password  = $_POST["password"];

$xmlpath = "//user[login/text()=$login and password/text()=$password]/attribute::id";

if( file_exists($file) )
{
    $xml = simplexml_load_file($file);

    if($result = $xml->xpath($xmlpath))
        echo "You are now logged in as user $result[0].";
    else
        echo "Invalid login name or password.";
}
else
    exit("Failed to open $file.");

?>

```

No checks are performed on the input. This means that a malevolent user can “short-circuit” the test by entering ' or 1=1 for both the login name and password, resulting in `$xmlpath` being evaluated as shown here:

```
//user[login/text()=' ' or 1=1 and password/text()=' ' or 1=1]/attribute::id
```

Since the expression inside the square brackets always evaluates as `true`, it is effectively the same as this one, which matches the `id` attribute of every `user` element in the XML document:

```
//user/attribute::id
```

One way in which this particular attack can be circumvented is simply by quoting the variable names to be interpolated in the definition of `$xmlpath`, forcing the values passed from a Web form to be converted to strings:

```
$xpath = "//user[login/text()=' $login' and password/text()=' $password']/attribute::id";
```

This is the same strategy that is often recommended for preventing SQL injection attacks. In general, the practices you should follow for preventing XPath injection attacks are the same as for preventing SQL injection:

- Never accepted untested data from users in your application.
- Check all user-submitted data for type; reject or convert data that is of the wrong type
- Test numeric data for out of range values; truncate, round, or reject values that are out of range. Test strings for illegal characters and either strip them out or reject input containing them.
- Do not output explicit error messages that might provide an unauthorized user with clues that could be used to compromise the system; log these to a file or database table instead.

Just as SQL injection attacks can be used to obtain information about database schemas, so can XPath injection be used to traverse XML files to uncover their structure, as discussed in Amit Klein's paper [Blind XPath Injection](#) (PDF file, 46KB).

It is also important to check the output being sent back to the client. Consider what can happen when we use the MySQL `ExtractValue()` [1249] function:

```
mysql> SELECT ExtractValue(
->     LOAD_FILE('users.xml'),
->     '//user[login/text()="" or 1=1 and password/text()="" or 1=1]/attribute::id'
-> ) AS id;
+-----+
| id      |
+-----+
| 00327 13579 02403 42354 28570 |
+-----+
1 row in set (0.01 sec)
```

Because `ExtractValue()` [1249] returns multiple matches as a single space-delimited string, this injection attack provides every valid ID contained within `users.xml` to the user as a single row of output. As an extra safeguard, you should also test output before returning it to the user. Here is a simple example:

```
mysql> SELECT @id = ExtractValue(
->     LOAD_FILE('users.xml'),
->     '//user[login/text()="" or 1=1 and password/text()="" or 1=1]/attribute::id'
-> );
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT IF(
->     INSTR(@id, ' ') = 0,
->     @id,
->     'Unable to retrieve user ID')
-> AS singleID;
+-----+
| singleID      |
+-----+
| Unable to retrieve user ID |
+-----+
1 row in set (0.00 sec)
```

In general, the guidelines for returning data to users securely are the same as for accepting user input. These can be summed up as:

- Always test outgoing data for type and permissible values.
- Never permit unauthorized users to view error messages that might provide information about the application that could be used to exploit it.

## 12.12. Bit Functions

Table 12.16. Bitwise Functions

Name	Description
<code>BIT_COUNT()</code> [1259]	Return the number of bits that are set
<code>&amp;</code> [1258]	Bitwise AND
<code>~</code> [1259]	Invert bits
<code> </code> [1258]	Bitwise OR
<code>^</code> [1258]	Bitwise XOR
<code>&lt;&lt;</code> [1258]	Left shift
<code>&gt;&gt;</code> [1258]	Right shift

MySQL uses `BIGINT` (64-bit) arithmetic for bit operations, so these operators have a maximum range of 64 bits.

- `|` [1258]

Bitwise OR:

```
mysql> SELECT 29 | 15;
-> 31
```

The result is an unsigned 64-bit integer.

- `&` [1258]

Bitwise AND:

```
mysql> SELECT 29 & 15;
-> 13
```

The result is an unsigned 64-bit integer.

- `^` [1258]

Bitwise XOR:

```
mysql> SELECT 1 ^ 1;
-> 0
mysql> SELECT 1 ^ 0;
-> 1
mysql> SELECT 11 ^ 3;
-> 8
```

The result is an unsigned 64-bit integer.

- `<<` [1258]

Shifts a longlong (`BIGINT`) number to the left.

```
mysql> SELECT 1 << 2;
-> 4
```

The result is an unsigned 64-bit integer. The value is truncated to 64 bits. In particular, if the shift count is greater or equal to the width of an unsigned 64-bit number, the result is zero.

- `>>` [1258]

Shifts a longlong (`BIGINT`) number to the right.

```
mysql> SELECT 4 >> 2;
```



```
-> 1
```

The result is an unsigned 64-bit integer. The value is truncated to 64 bits. In particular, if the shift count is greater or equal to the width of an unsigned 64-bit number, the result is zero.

- [~ \[1259\]](#)

Invert all bits.

```
mysql> SELECT 5 & ~1;
-> 4
```

The result is an unsigned 64-bit integer.

- [BIT\\_COUNT\(N\) \[1259\]](#)

Returns the number of bits that are set in the argument *N*.

```
mysql> SELECT BIT_COUNT(29), BIT_COUNT(b'101010');
-> 4, 3
```

## 12.13. Encryption and Compression Functions

**Table 12.17. Encryption Functions**

Name	Description
<a href="#">AES_DECRYPT() [1260]</a>	Decrypt using AES
<a href="#">AES_ENCRYPT() [1260]</a>	Encrypt using AES
<a href="#">COMPRESS() [1261]</a>	Return result as a binary string
<a href="#">DECODE() [1261]</a>	Decodes a string encrypted using ENCODE()
<a href="#">DES_DECRYPT() [1261]</a>	Decrypt a string
<a href="#">DES_ENCRYPT() [1262]</a>	Encrypt a string
<a href="#">ENCODE() [1262]</a>	Encode a string
<a href="#">ENCRYPT() [1262]</a>	Encrypt a string
<a href="#">MD5() [1263]</a>	Calculate MD5 checksum
<a href="#">OLD_PASSWORD() [1263]</a>	Return the value of the pre-4.1 implementation of PASSWORD
<a href="#">PASSWORD() [1263]</a>	Calculate and return a password string
<a href="#">SHA1(), SHA() [1264]</a>	Calculate an SHA-1 160-bit checksum
<a href="#">SHA2() [1264]</a>	Calculate an SHA-2 checksum
<a href="#">UNCOMPRESS() [1265]</a>	Uncompress a string compressed
<a href="#">UNCOMPRESSED_LENGTH() [1265]</a>	Return the length of a string before compression
<a href="#">VALIDATE_PASSWORD_STRENGTH() [1265]</a>	Determine strength of password

Many encryption and compression functions return strings for which the result might contain arbitrary byte values. If you want to store these results, use a column with a [VARBINARY](#) or [BLOB](#) binary string data type. This will avoid potential problems with trailing space removal or character set conversion that would change data values, such as may occur if you use a nonbinary string data type ([CHAR](#), [VARCHAR](#), [TEXT](#)).

Some encryption functions return strings of ASCII characters: [MD5\(\) \[1263\]](#), [OLD\\_PASSWORD\(\) \[1263\]](#), [PASSWORD\(\) \[1263\]](#), [SHA\(\) \[1264\]](#), [SHA1\(\) \[1264\]](#), [SHA2\(\) \[1264\]](#). In MySQL 5.6, their return value is a nonbinary string that has a character set and collation determined by the [character\\_set\\_connection \[467\]](#) and [collation\\_connection \[469\]](#) system variables.

For versions in which functions such as `MD5()` or `SHA1()` return a string of hex digits as a binary string, the return value cannot be converted to uppercase or compared in case-insensitive fashion as is. You must convert the value to a nonbinary string. See the discussion of binary string conversion in [Section 12.10, “Cast Functions and Operators”](#).

If an application stores values from a function such as `MD5()` [1263] or `SHA1()` [1264] that returns a string of hex digits, more efficient storage and comparisons can be obtained by converting the hex representation to binary using `UNHEX()` [1186] and storing the result in a `BINARY(N)` column. Each pair of hex digits requires one byte in binary form, so the value of *N* depends on the length of the hex string. *N* is 16 for an `MD5()` [1263] value and 20 for a `SHA1()` [1264] value. For `SHA2()` [1264], *N* ranges from 28 to 32 depending on the argument specifying the desired bit length of the result.

The size penalty for storing the hex string in a `CHAR` column is at least two times, up to eight times if the value is stored in a column that uses the `utf8` character set (where each character uses 4 bytes). Storing the string also results in slower comparisons because of the larger values and the need to take character set collation rules into account.

Suppose that an application stores `MD5()` [1263] string values in a `CHAR(32)` column:

```
CREATE TABLE md5_tbl (md5_val CHAR(32), ...);
INSERT INTO md5_tbl (md5_val, ...) VALUES(MD5('abcdef'), ...);
```

To convert hex strings to more compact form, modify the application to use `UNHEX()` [1186] and `BINARY(16)` instead as follows:

```
CREATE TABLE md5_tbl (md5_val BINARY(16), ...);
INSERT INTO md5_tbl (md5_val, ...) VALUES(UNHEX(MD5('abcdef')), ...);
```

Applications should be prepared to handle the very rare case that a hashing function produces the same value for two different input values. One way to make collisions detectable is to make the hash column a primary key.



#### Note

Exploits for the MD5 and SHA-1 algorithms have become known. You may wish to consider using one of the other encryption functions described in this section instead, such as `SHA2()` [1264].



#### Caution

Passwords or other sensitive values supplied as arguments to encryption functions are sent in plaintext to the MySQL server unless an SSL connection is used. Also, such values will appear in any MySQL logs to which they are written. To avoid these types of exposure, applications can encrypt sensitive values on the client side before sending them to the server. The same considerations apply to encryption keys. To avoid exposing these, applications can use stored procedures to encrypt and decrypt values on the server side.

- `AES_DECRYPT(crypt_str,key_str)` [1260]

This function decrypts data using the official AES (Advanced Encryption Standard) algorithm. For more information, see the description of `AES_ENCRYPT()` [1260].

- `AES_ENCRYPT(str,key_str)` [1260]

`AES_ENCRYPT()` [1260] and `AES_DECRYPT()` [1260] enable encryption and decryption of data using the official AES (Advanced Encryption Standard) algorithm, previously known as “Rijndael.” Encoding with a 128-bit key length is used, but you can extend it up to 256 bits by modifying the source. We chose 128 bits because it is much faster and it is secure enough for most purposes.

`AES_ENCRYPT()` [1260] encrypts a string and returns a binary string. `AES_DECRYPT()` [1260] decrypts the encrypted string and returns the original string. The input arguments may be any length. If either argument is `NULL`, the result of this function is also `NULL`.

Because AES is a block-level algorithm, padding is used to encode uneven length strings and so the result string length may be calculated using this formula:

```
16 * (trunc(string_length / 16) + 1)
```

If `AES_DECRYPT()` [1260] detects invalid data or incorrect padding, it returns `NULL`. However, it is possible for `AES_DECRYPT()` [1260] to return a non-`NULL` value (possibly garbage) if the input data or the key is invalid.

You can use the AES functions to store data in an encrypted form by modifying your queries:

```
INSERT INTO t VALUES (1,AES_ENCRYPT('text','password'));
```

`AES_ENCRYPT()` [1260] and `AES_DECRYPT()` [1260] can be considered the most cryptographically secure encryption functions currently available in MySQL.

- `COMPRESS(string_to_compress)` [1261]

Compresses a string and returns the result as a binary string. This function requires MySQL to have been compiled with a compression library such as `zlib`. Otherwise, the return value is always `NULL`. The compressed string can be uncompressed with `UNCOMPRESS()` [1265].

```
mysql> SELECT LENGTH(COMPRESS(REPEAT('a',1000)));
-> 21
mysql> SELECT LENGTH(COMPRESS(''));
-> 0
mysql> SELECT LENGTH(COMPRESS('a'));
-> 13
mysql> SELECT LENGTH(COMPRESS(REPEAT('a',16)));
-> 15
```

The compressed string contents are stored the following way:

- Empty strings are stored as empty strings.
- Nonempty strings are stored as a 4-byte length of the uncompressed string (low byte first), followed by the compressed string. If the string ends with space, an extra “.” character is added to avoid problems with endspace trimming should the result be stored in a `CHAR` or `VARCHAR` column. (However, use of nonbinary string data types such as `CHAR` or `VARCHAR` to store compressed strings is not recommended anyway because character set conversion may occur. Use a `VARBINARY` or `BLOB` binary string column instead.)
- `DECODE(crypt_str,pass_str)` [1261]

Decrypts the encrypted string `crypt_str` using `pass_str` as the password. `crypt_str` should be a string returned from `ENCODE()` [1262].

- `DES_DECRYPT(encrypt_str[,key_str])` [1261]

Decrypts a string encrypted with `DES_ENCRYPT()` [1262]. If an error occurs, this function returns `NULL`.

This function works only if MySQL has been configured with SSL support. See [Section 6.3.9, “Using SSL for Secure Connections”](#).

If no `key_str` argument is given, `DES_DECRYPT()` [1261] examines the first byte of the encrypted string to determine the DES key number that was used to encrypt the original string, and then reads the key from the DES key file to decrypt the message. For this to work, the user must have the `SUPER` [756] privilege. The key file can be specified with the `--des-key-file` [421] server option.

If you pass this function a `key_str` argument, that string is used as the key for decrypting the message.

If the `crypt_str` argument does not appear to be an encrypted string, MySQL returns the given `crypt_str`.

- `DES_ENCRYPT(str[, {key_num|key_str}])` [1262]

Encrypts the string with the given key using the Triple-DES algorithm.

This function works only if MySQL has been configured with SSL support. See [Section 6.3.9, “Using SSL for Secure Connections”](#).

The encryption key to use is chosen based on the second argument to `DES_ENCRYPT()` [1262], if one was given. With no argument, the first key from the DES key file is used. With a `key_num` argument, the given key number (0 to 9) from the DES key file is used. With a `key_str` argument, the given key string is used to encrypt `str`.

The key file can be specified with the `--des-key-file` [421] server option.

The return string is a binary string where the first character is `CHAR(128 | key_num)` [1177]. If an error occurs, `DES_ENCRYPT()` [1262] returns `NULL`.

The 128 is added to make it easier to recognize an encrypted key. If you use a string key, `key_num` is 127.

The string length for the result is given by this formula:

```
new_len = orig_len + (8 - (orig_len % 8)) + 1
```

Each line in the DES key file has the following format:

```
key_num des_key_str
```

Each `key_num` value must be a number in the range from 0 to 9. Lines in the file may be in any order. `des_key_str` is the string that is used to encrypt the message. There should be at least one space between the number and the key. The first key is the default key that is used if you do not specify any key argument to `DES_ENCRYPT()` [1262].

You can tell MySQL to read new key values from the key file with the `FLUSH DES_KEY_FILE` statement. This requires the `RELOAD` [755] privilege.

One benefit of having a set of default keys is that it gives applications a way to check for the existence of encrypted column values, without giving the end user the right to decrypt those values.

```
mysql> SELECT customer_address FROM customer_table
> WHERE crypt_credit_card = DES_ENCRYPT('credit_card_number');
```

- `ENCODE(str, pass_str)` [1262]

Encrypt `str` using `pass_str` as the password. To decrypt the result, use `DECODE()` [1261].

The result is a binary string of the same length as `str`.

The strength of the encryption is based on how good the random generator is. It should suffice for short strings.

- `ENCRYPT(str[, salt])` [1262]

Encrypts `str` using the Unix `crypt()` system call and returns a binary string. The `salt` argument must be a string with at least two characters or the result will be `NULL`. If no `salt` argument is given, a random value is used.

```
mysql> SELECT ENCRYPT('hello');
-> 'VxuFAJXVARROc'
```

`ENCRYPT()` [1262] ignores all but the first eight characters of `str`, at least on some systems. This behavior is determined by the implementation of the underlying `crypt()` system call.

The use of `ENCRYPT()` [1262] with the `ucs2`, `utf16`, `utf16le`, or `utf32` multi-byte character sets is not recommended because the system call expects a string terminated by a zero byte.

If `crypt()` is not available on your system (as is the case with Windows), `ENCRYPT()` [1262] always returns `NULL`.

- `MD5(str)` [1263]

Calculates an MD5 128-bit checksum for the string. The value is returned as a string of 32 hex digits, or `NULL` if the argument was `NULL`. The return value can, for example, be used as a hash key. See the notes at the beginning of this section about storing hash values efficiently.

The return value is a nonbinary string in the connection character set.

```
mysql> SELECT MD5('testing');
-> 'ae2b1fca515949e5d54fb22b8ed95575'
```

This is the “RSA Data Security, Inc. MD5 Message-Digest Algorithm.”

See the note regarding the MD5 algorithm at the beginning this section.

- `OLD_PASSWORD(str)` [1263]

`OLD_PASSWORD()` [1263] was added when the implementation of `PASSWORD()` [1263] was changed in MySQL 4.1 to improve security. `OLD_PASSWORD()` [1263] returns the value of the pre-4.1 implementation of `PASSWORD()` [1263] as a string, and is intended to permit you to reset passwords for any pre-4.1 clients that need to connect to your version 5.6 MySQL server without locking them out. See [Section 6.1.2.4, “Password Hashing in MySQL”](#).

The return value is a nonbinary string in the connection character set.

- `PASSWORD(str)` [1263]

Calculates and returns a hashed password string from the cleartext password `str` and returns a nonbinary string in the connection character set, or `NULL` if the argument is `NULL`. This function is the SQL interface to the algorithm used by the server to encrypt MySQL passwords for storage in the `mysql.user` grant table.

The password hashing method used by `PASSWORD()` [1263] depends on the value of the `old_passwords` [517] system variable:

```
mysql> SET old_passwords = 0;
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| *6C8989366EAF75BB670AD8EA7A7FC1176A95CEF4 |
+-----+

mysql> SET old_passwords = 1;
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| 6f8c114b58f2ce9e   |
+-----+
```

If `old_passwords=1`, `PASSWORD('str')` [1263] returns the same value as `OLD_PASSWORD('str')` [1263].

SHA-256 password hashing (`old_passwords=2` [517]) uses a random salt value, which makes the result from `PASSWORD()` [1263] nondeterministic. Consequently, statements that use this function are not safe for statement-based replication and cannot be stored in the query cache.

For descriptions of the permitted values of `old_passwords` [517], see Section 5.1.4, “Server System Variables”.

Encryption performed by `PASSWORD()` [1263] is one-way (not reversible). It is not the same type of encryption as used for Unix passwords; for that, use `ENCRYPT()` [1262].



#### Note

The `PASSWORD()` [1263] function is used by the authentication system in MySQL Server; you should *not* use it in your own applications. For that purpose, consider `MD5()` [1263] or `SHA2()` [1264] instead. Also see RFC 2195, section 2 (Challenge-Response Authentication Mechanism (CRAM)), for more information about handling passwords and authentication securely in your applications.



#### Important

Statements that invoke `PASSWORD()` [1263] may be recorded in server logs or in a history file such as `~/.mysql_history`, which means that cleartext passwords may be read by anyone having read access to that information. See Section 6.1.2, “Keeping Passwords Secure”.

- `SHA1(str)` [1264], `SHA(str)` [1264]

Calculates an SHA-1 160-bit checksum for the string, as described in RFC 3174 (Secure Hash Algorithm). The value is returned as a string of 40 hex digits, or `NULL` if the argument was `NULL`. One of the possible uses for this function is as a hash key. See the notes at the beginning of this section about storing hash values efficiently. You can also use `SHA1()` [1264] as a cryptographic function for storing passwords. `SHA()` [1264] is synonymous with `SHA1()` [1264].

The return value is a nonbinary string in the connection character set.

```
mysql> SELECT SHA1('abc');
-> 'a9993e364706816aba3e25717850c26c9cd0d89d'
```

`SHA1()` [1264] can be considered a cryptographically more secure equivalent of `MD5()` [1263]. However, see the note regarding the MD5 and SHA-1 algorithms at the beginning this section.

- `SHA2(str, hash_length)` [1264]

Calculates the SHA-2 family of hash functions (SHA-224, SHA-256, SHA-384, and SHA-512). The first argument is the cleartext string to be hashed. The second argument indicates the desired bit length of the result, which must have a value of 224, 256, 384, 512, or 0 (which is equivalent to 256). If either argument is `NULL` or the hash length is not one of the permitted values, the return value is `NULL`. Otherwise, the function result is a hash value containing the desired number of bits. See the notes at the beginning of this section about storing hash values efficiently.

The return value is a nonbinary string in the connection character set.

```
mysql> SELECT SHA2('abc', 224);
-> '23097d223405d8228642a477bda255b32aadbce4bda0b3f7e36c9da7'
```

This function works only if MySQL has been configured with SSL support. See Section 6.3.9, “Using SSL for Secure Connections”.

`SHA2()` [1264] can be considered cryptographically more secure than `MD5()` [1263] or `SHA1()` [1264].

- `UNCOMPRESS(string_to_uncompress)` [1265]

Uncompresses a string compressed by the `COMPRESS()` [1261] function. If the argument is not a compressed value, the result is `NULL`. This function requires MySQL to have been compiled with a compression library such as `zlib`. Otherwise, the return value is always `NULL`.

```
mysql> SELECT UNCOMPRESS(COMPRESS('any string'));
-> 'any string'
mysql> SELECT UNCOMPRESS('any string');
-> NULL
```

- `UNCOMPRESSED_LENGTH(compressed_string)` [1265]

Returns the length that the compressed string had before being compressed.

```
mysql> SELECT UNCOMPRESSED_LENGTH(COMPRESS(REPEAT('a',30)));
-> 30
```

- `VALIDATE_PASSWORD_STRENGTH(str)` [1265]

Given an argument representing a cleartext password, this function returns an integer to indicate how strong the password is. The return value ranges from 0 (weak) to 100 (strong).

The password is subjected to increasingly strict tests and the return value reflects which tests were satisfied, as shown in the following table.

Password Test	Return Value
Length < 4	0
Length ≥ 4 and < <code>validate_password_length</code> [744]	25
Satisfies policy 1 ( <code>LOW</code> )	50
Satisfies policy 2 ( <code>MEDIUM</code> )	75
Satisfies policy 3 ( <code>STRONG</code> )	100

Password assessment by `VALIDATE_PASSWORD_STRENGTH()` [1265] is done by the `validate_password` plugin. If that plugin is not installed, the function always returns 0. For information about installing the `validate_password` plugin, see [Section 6.1.2.6, “The Password Validation Plugin”](#). To examine or configure the parameters that affect password testing, check or set the system variables implemented by `validate_password` plugin. See [Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”](#).

This function was added in MySQL 5.6.6.

## 12.14. Information Functions

**Table 12.18. Information Functions**

Name	Description
<code>BENCHMARK()</code> [1266]	Repeatedly execute an expression
<code>CHARSET()</code> [1266]	Return the character set of the argument
<code>COERCIBILITY()</code> [1266]	Return the collation coercibility value of the string argument
<code>COLLATION()</code> [1267]	Return the collation of the string argument
<code>CONNECTION_ID()</code> [1267]	Return the connection ID (thread ID) for the connection
<code>CURRENT_USER()</code> , <code>CURRENT_USER</code> [1267]	The authenticated user name and host name
<code>DATABASE()</code> [1268]	Return the default (current) database name



Name	Description
<a href="#">FOUND_ROWS() [1268]</a>	For a SELECT with a LIMIT clause, the number of rows that would be returned were there no LIMIT clause
<a href="#">LAST_INSERT_ID() [1269]</a>	Value of the AUTOINCREMENT column for the last INSERT
<a href="#">ROW_COUNT() [1272]</a>	The number of rows updated
<a href="#">SCHEMA() [1273]</a>	A synonym for DATABASE()
<a href="#">SESSION_USER() [1273]</a>	Synonym for USER()
<a href="#">SYSTEM_USER() [1273]</a>	Synonym for USER()
<a href="#">USER() [1273]</a>	The user name and host name provided by the client
<a href="#">VERSION() [1273]</a>	Returns a string that indicates the MySQL server version

- [BENCHMARK\(count,expr\) \[1266\]](#)

The [BENCHMARK\(\) \[1266\]](#) function executes the expression *expr* repeatedly *count* times. It may be used to time how quickly MySQL processes the expression. The result value is always 0. The intended use is from within the `mysql` client, which reports query execution times:

```
mysql> SELECT BENCHMARK(1000000,ENCODE('hello','goodbye'));
+-----+
| BENCHMARK(1000000,ENCODE('hello','goodbye')) |
+-----+
|                                             0 |
+-----+
1 row in set (4.74 sec)
```

The time reported is elapsed time on the client end, not CPU time on the server end. It is advisable to execute [BENCHMARK\(\) \[1266\]](#) several times, and to interpret the result with regard to how heavily loaded the server machine is.

[BENCHMARK\(\) \[1266\]](#) is intended for measuring the runtime performance of scalar expressions, which has some significant implications for the way that you use it and interpret the results:

- Only scalar expressions can be used. Although the expression can be a subquery, it must return a single column and at most a single row. For example, [BENCHMARK\(10, \(SELECT \\* FROM t\)\) \[1266\]](#) will fail if the table *t* has more than one column or more than one row.
- Executing a `SELECT expr` statement *N* times differs from executing `SELECT BENCHMARK(N, expr)` in terms of the amount of overhead involved. The two have very different execution profiles and you should not expect them to take the same amount of time. The former involves the parser, optimizer, table locking, and runtime evaluation *N* times each. The latter involves only runtime evaluation *N* times, and all the other components just once. Memory structures already allocated are reused, and runtime optimizations such as local caching of results already evaluated for aggregate functions can alter the results. Use of [BENCHMARK\(\) \[1266\]](#) thus measures performance of the runtime component by giving more weight to that component and removing the “noise” introduced by the network, parser, optimizer, and so forth.
- [CHARSET\(str\) \[1266\]](#)

Returns the character set of the string argument.

```
mysql> SELECT CHARSET('abc');
-> 'latin1'
mysql> SELECT CHARSET(CONVERT('abc' USING utf8));
-> 'utf8'
mysql> SELECT CHARSET(USER());
-> 'utf8'
```

- [COERCIBILITY\(str\) \[1266\]](#)



Returns the collation coercibility value of the string argument.

```
mysql> SELECT COERCIBILITY('abc' COLLATE latin1_swedish_ci);
-> 0
mysql> SELECT COERCIBILITY(USER());
-> 3
mysql> SELECT COERCIBILITY('abc');
-> 4
```

The return values have the meanings shown in the following table. Lower values have higher precedence.

Coercibility	Meaning	Example
0	Explicit collation	Value with <code>COLLATE</code> clause
1	No collation	Concatenation of strings with different collations
2	Implicit collation	Column value, stored routine parameter or local variable
3	System constant	<code>USER()</code> [1273] return value
4	Coercible	Literal string
5	Ignorable	<code>NULL</code> or an expression derived from <code>NULL</code>

- [COLLATION\(\*str\*\)](#) [1267]

Returns the collation of the string argument.

```
mysql> SELECT COLLATION('abc');
-> 'latin1_swedish_ci'
mysql> SELECT COLLATION(_utf8'abc');
-> 'utf8_general_ci'
```

- [CONNECTION\\_ID\(\)](#) [1267]

Returns the connection ID (thread ID) for the connection. Every connection has an ID that is unique among the set of currently connected clients.

```
mysql> SELECT CONNECTION_ID();
-> 23786
```

- [CURRENT\\_USER](#) [1267], [CURRENT\\_USER\(\)](#) [1267]

Returns the user name and host name combination for the MySQL account that the server used to authenticate the current client. This account determines your access privileges. The return value is a string in the `utf8` character set.

The value of [CURRENT\\_USER\(\)](#) [1267] can differ from the value of [USER\(\)](#) [1273].

```
mysql> SELECT USER();
-> 'davida@localhost'
mysql> SELECT * FROM mysql.user;
ERROR 1044: Access denied for user ''@'localhost' to
database 'mysql'
mysql> SELECT CURRENT_USER();
-> '@localhost'
```

The example illustrates that although the client specified a user name of `davida` (as indicated by the value of the [USER\(\)](#) [1273] function), the server authenticated the client using an anonymous user account (as seen by the empty user name part of the [CURRENT\\_USER\(\)](#) [1267] value). One way this might occur is that there is no account listed in the grant tables for `davida`.

Within a stored program or view, `CURRENT_USER()` [1267] returns the account for the user who defined the object (as given by its `DEFINER` value) unless defined with the `SQL SECURITY INVOKER` characteristic. In the latter case, `CURRENT_USER()` [1267] returns the object's invoker.

Triggers and events have no option to define the `SQL SECURITY` characteristic, so for these objects, `CURRENT_USER()` [1267] returns the account for the user who defined the object. To return the invoker, use `USER()` [1273] or `SESSION_USER()` [1273].

The following statements support use of the `CURRENT_USER()` [1267] function to take the place of the name of (and, possibly, a host for) an affected user or a definer; in such cases, `CURRENT_USER()` [1267] is expanded where and as needed:

- `DROP USER`
- `RENAME USER`
- `GRANT`
- `REVOKE`
- `CREATE FUNCTION`
- `CREATE PROCEDURE`
- `CREATE TRIGGER`
- `CREATE EVENT`
- `CREATE VIEW`
- `ALTER EVENT`
- `ALTER VIEW`
- `SET PASSWORD`

For information about the implications that this expansion of `CURRENT_USER()` [1267] has for replication in different releases of MySQL 5.6, see [Section 16.4.1.6, “Replication of `CURRENT\_USER\(\)`”](#).

- `DATABASE()` [1268]

Returns the default (current) database name as a string in the `utf8` character set. If there is no default database, `DATABASE()` [1268] returns `NULL`. Within a stored routine, the default database is the database that the routine is associated with, which is not necessarily the same as the database that is the default in the calling context.

```
mysql> SELECT DATABASE();
-> 'test'
```

If there is no default database, `DATABASE()` [1268] returns `NULL`.

- `FOUND_ROWS()` [1268]

A `SELECT` statement may include a `LIMIT` clause to restrict the number of rows the server returns to the client. In some cases, it is desirable to know how many rows the statement would have returned without the `LIMIT`, but without running the statement again. To obtain this row count, include a `SQL_CALC_FOUND_ROWS` option in the `SELECT` statement, and then invoke `FOUND_ROWS()` [1268] afterward:

```
mysql> SELECT SQL_CALC_FOUND_ROWS * FROM tbl_name
-> WHERE id > 100 LIMIT 10;
```

```
mysql> SELECT FOUND_ROWS();
```

The second `SELECT` returns a number indicating how many rows the first `SELECT` would have returned had it been written without the `LIMIT` clause.

In the absence of the `SQL_CALC_FOUND_ROWS` option in the most recent successful `SELECT` statement, `FOUND_ROWS() [1268]` returns the number of rows in the result set returned by that statement. If the statement includes a `LIMIT` clause, `FOUND_ROWS() [1268]` returns the number of rows up to the limit. For example, `FOUND_ROWS() [1268]` returns 10 or 60, respectively, if the statement includes `LIMIT 10` or `LIMIT 50, 10`.

The row count available through `FOUND_ROWS() [1268]` is transient and not intended to be available past the statement following the `SELECT SQL_CALC_FOUND_ROWS` statement. If you need to refer to the value later, save it:

```
mysql> SELECT SQL_CALC_FOUND_ROWS * FROM ... ;
mysql> SET @rows = FOUND_ROWS();
```

If you are using `SELECT SQL_CALC_FOUND_ROWS`, MySQL must calculate how many rows are in the full result set. However, this is faster than running the query again without `LIMIT`, because the result set need not be sent to the client.

`SQL_CALC_FOUND_ROWS` and `FOUND_ROWS() [1268]` can be useful in situations when you want to restrict the number of rows that a query returns, but also determine the number of rows in the full result set without running the query again. An example is a Web script that presents a paged display containing links to the pages that show other sections of a search result. Using `FOUND_ROWS() [1268]` enables you to determine how many other pages are needed for the rest of the result.

The use of `SQL_CALC_FOUND_ROWS` and `FOUND_ROWS() [1268]` is more complex for `UNION` statements than for simple `SELECT` statements, because `LIMIT` may occur at multiple places in a `UNION`. It may be applied to individual `SELECT` statements in the `UNION`, or global to the `UNION` result as a whole.

The intent of `SQL_CALC_FOUND_ROWS` for `UNION` is that it should return the row count that would be returned without a global `LIMIT`. The conditions for use of `SQL_CALC_FOUND_ROWS` with `UNION` are:

- The `SQL_CALC_FOUND_ROWS` keyword must appear in the first `SELECT` of the `UNION`.
- The value of `FOUND_ROWS() [1268]` is exact only if `UNION ALL` is used. If `UNION` without `ALL` is used, duplicate removal occurs and the value of `FOUND_ROWS() [1268]` is only approximate.
- If no `LIMIT` is present in the `UNION`, `SQL_CALC_FOUND_ROWS` is ignored and returns the number of rows in the temporary table that is created to process the `UNION`.

Beyond the cases described here, the behavior of `FOUND_ROWS() [1268]` is undefined (for example, its value following a `SELECT` statement that fails with an error).



### Important

`FOUND_ROWS() [1268]` is not replicated reliably using statement-based replication. This function is automatically replicated using row-based replication.

- `LAST_INSERT_ID() [1269]`, `LAST_INSERT_ID(expr) [1269]`

`LAST_INSERT_ID() [1269]` (with no argument) returns a `BIGINT` (64-bit) value representing the first automatically generated value successfully inserted for an `AUTO_INCREMENT` column as a result of the most recently executed `INSERT` statement. The value of `LAST_INSERT_ID() [1269]` remains unchanged if no rows are successfully inserted.

For example, after inserting a row that generates an `AUTO_INCREMENT` value, you can get the value like this:

```
mysql> SELECT LAST_INSERT_ID();
-> 195
```

The currently executing statement does not affect the value of `LAST_INSERT_ID()` [1269]. Suppose that you generate an `AUTO_INCREMENT` value with one statement, and then refer to `LAST_INSERT_ID()` [1269] in a multiple-row `INSERT` statement that inserts rows into a table with its own `AUTO_INCREMENT` column. The value of `LAST_INSERT_ID()` [1269] will remain stable in the second statement; its value for the second and later rows is not affected by the earlier row insertions. (However, if you mix references to `LAST_INSERT_ID()` [1269] and `LAST_INSERT_ID(expr)` [1269], the effect is undefined.)

If the previous statement returned an error, the value of `LAST_INSERT_ID()` [1269] is undefined. For transactional tables, if the statement is rolled back due to an error, the value of `LAST_INSERT_ID()` [1269] is left undefined. For manual `ROLLBACK`, the value of `LAST_INSERT_ID()` [1269] is not restored to that before the transaction; it remains as it was at the point of the `ROLLBACK`.

Within the body of a stored routine (procedure or function) or a trigger, the value of `LAST_INSERT_ID()` [1269] changes the same way as for statements executed outside the body of these kinds of objects. The effect of a stored routine or trigger upon the value of `LAST_INSERT_ID()` [1269] that is seen by following statements depends on the kind of routine:

- If a stored procedure executes statements that change the value of `LAST_INSERT_ID()` [1269], the changed value is seen by statements that follow the procedure call.
- For stored functions and triggers that change the value, the value is restored when the function or trigger ends, so following statements will not see a changed value.

The ID that was generated is maintained in the server on a *per-connection* basis. This means that the value returned by the function to a given client is the first `AUTO_INCREMENT` value generated for most recent statement affecting an `AUTO_INCREMENT` column *by that client*. This value cannot be affected by other clients, even if they generate `AUTO_INCREMENT` values of their own. This behavior ensures that each client can retrieve its own ID without concern for the activity of other clients, and without the need for locks or transactions.

The value of `LAST_INSERT_ID()` [1269] is not changed if you set the `AUTO_INCREMENT` column of a row to a non-“magic” value (that is, a value that is not `NULL` and not `0`).



### Important

If you insert multiple rows using a single `INSERT` statement, `LAST_INSERT_ID()` [1269] returns the value generated for the *first* inserted row *only*. The reason for this is to make it possible to reproduce easily the same `INSERT` statement against some other server.

For example:

```
mysql> USE test;
Database changed
mysql> CREATE TABLE t (
->   id INT AUTO_INCREMENT NOT NULL PRIMARY KEY,
->   name VARCHAR(10) NOT NULL
-> );
Query OK, 0 rows affected (0.09 sec)

mysql> INSERT INTO t VALUES (NULL, 'Bob');
Query OK, 1 row affected (0.01 sec)

mysql> SELECT * FROM t;
```

```

+----+-----+
| id | name |
+----+-----+
|  1 | Bob  |
+----+-----+
1 row in set (0.01 sec)

mysql> SELECT LAST_INSERT_ID();
+-----+
| LAST_INSERT_ID() |
+-----+
|                1 |
+-----+
1 row in set (0.00 sec)

mysql> INSERT INTO t VALUES
-> (NULL, 'Mary'), (NULL, 'Jane'), (NULL, 'Lisa');
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM t;
+----+-----+
| id | name |
+----+-----+
|  1 | Bob  |
|  2 | Mary |
|  3 | Jane |
|  4 | Lisa |
+----+-----+
4 rows in set (0.01 sec)

mysql> SELECT LAST_INSERT_ID();
+-----+
| LAST_INSERT_ID() |
+-----+
|                2 |
+-----+
1 row in set (0.00 sec)

```

Although the second `INSERT` statement inserted three new rows into `t`, the ID generated for the first of these rows was 2, and it is this value that is returned by `LAST_INSERT_ID()` [1269] for the following `SELECT` statement.

If you use `INSERT IGNORE` and the row is ignored, the `AUTO_INCREMENT` counter is not incremented and `LAST_INSERT_ID()` [1269] returns 0, which reflects that no row was inserted.

If *expr* is given as an argument to `LAST_INSERT_ID()` [1269], the value of the argument is returned by the function and is remembered as the next value to be returned by `LAST_INSERT_ID()` [1269]. This can be used to simulate sequences:

1. Create a table to hold the sequence counter and initialize it:

```

mysql> CREATE TABLE sequence (id INT NOT NULL);
mysql> INSERT INTO sequence VALUES (0);

```

2. Use the table to generate sequence numbers like this:

```

mysql> UPDATE sequence SET id=LAST_INSERT_ID(id+1);
mysql> SELECT LAST_INSERT_ID();

```

The `UPDATE` statement increments the sequence counter and causes the next call to `LAST_INSERT_ID()` [1269] to return the updated value. The `SELECT` statement retrieves that value. The `mysql_insert_id()` C API function can also be used to get the value. See [Section 21.8.6.37](#), “`mysql_insert_id()`”.

You can generate sequences without calling `LAST_INSERT_ID()` [1269], but the utility of using the function this way is that the ID value is maintained in the server as the last automatically generated

value. It is multi-user safe because multiple clients can issue the `UPDATE` statement and get their own sequence value with the `SELECT` statement (or `mysql_insert_id()`), without affecting or being affected by other clients that generate their own sequence values.

Note that `mysql_insert_id()` is only updated after `INSERT` and `UPDATE` statements, so you cannot use the C API function to retrieve the value for `LAST_INSERT_ID(expr)` [1269] after executing other SQL statements like `SELECT` or `SET`.

- `ROW_COUNT()` [1272]

In MySQL 5.6, `ROW_COUNT()` returns a value as follows:

- DDL statements: 0. This applies to statements such as `CREATE TABLE` or `DROP TABLE`.
- DML statements other than `SELECT`: The number of affected rows. This applies to statements such as `UPDATE`, `INSERT`, or `DELETE` (as before), but now also to statements such as `ALTER TABLE` and `LOAD DATA INFILE`.
- `SELECT`: -1 if the statement returns a result set, or the number of rows “affected” if it does not. For example, for `SELECT * FROM t1`, `ROW_COUNT()` [1272] returns -1. For `SELECT * FROM t1 INTO OUTFILE 'file_name'`, `ROW_COUNT()` [1272] returns the number of rows written to the file.
- `SIGNAL` statements: 0.

For `UPDATE` statements, the affected-rows value by default is the number of rows actually changed. If you specify the `CLIENT_FOUND_ROWS` flag to `mysql_real_connect()` when connecting to `mysqld`, the affected-rows value is the number of rows “found”; that is, matched by the `WHERE` clause.

For `REPLACE` statements, the affected-rows value is 2 if the new row replaced an old row, because in this case, one row was inserted after the duplicate was deleted.

For `INSERT ... ON DUPLICATE KEY UPDATE` statements, the affected-rows value is 1 if the row is inserted as a new row and 2 if an existing row is updated.

The `ROW_COUNT()` [1272] value is similar to the value from the `mysql_affected_rows()` C API function and the row count that the `mysql` client displays following statement execution.

```
mysql> INSERT INTO t VALUES(1),(2),(3);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT ROW_COUNT();
+-----+
| ROW_COUNT() |
+-----+
|          3 |
+-----+
1 row in set (0.00 sec)

mysql> DELETE FROM t WHERE i IN(1,2);
Query OK, 2 rows affected (0.00 sec)

mysql> SELECT ROW_COUNT();
+-----+
| ROW_COUNT() |
+-----+
|          2 |
+-----+
1 row in set (0.00 sec)
```

**Important**

`ROW_COUNT()` [1272] is not replicated reliably using statement-based replication. This function is automatically replicated using row-based replication.

- `SCHEMA()` [1273]

This function is a synonym for `DATABASE()` [1268].

- `SESSION_USER()` [1273]

`SESSION_USER()` [1273] is a synonym for `USER()` [1273].

- `SYSTEM_USER()` [1273]

`SYSTEM_USER()` [1273] is a synonym for `USER()` [1273].

- `USER()` [1273]

Returns the current MySQL user name and host name as a string in the `utf8` character set.

```
mysql> SELECT USER();
-> 'david@localhost'
```

The value indicates the user name you specified when connecting to the server, and the client host from which you connected. The value can be different from that of `CURRENT_USER()` [1267].

You can extract only the user name part like this:

```
mysql> SELECT SUBSTRING_INDEX(USER(), '@', 1);
-> 'david'
```

- `VERSION()` [1273]

Returns a string that indicates the MySQL server version. The string uses the `utf8` character set. The value might have a suffix in addition to the version number. See the description of the `version` [563] system variable in [Section 5.1.4, “Server System Variables”](#).

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

```
mysql> SELECT VERSION();
-> '5.6.12-standard'
```

## 12.15. GTID Functions

The functions described in this section are used with GTID-based replication (available in MySQL 5.6.5 and later). For more information about GTIDs and how these functions are used, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

**Note**

The union of two GTID sets is simply their string representations joined together with an interposed comma. In other words, you can define a very simple function for obtaining the union of two GTID sets like this:

```
CREATE FUNCTION GTID_UNION(g1 TEXT, g2 TEXT)
RETURNS TEXT DETERMINISTIC
RETURN CONCAT(g1, ', ', g2);
```

Table 12.19. GTID Functions

Name	Description
<code>GTID_SUBSET()</code> [1274]	Return true if all GTIDs in subset are also in set; otherwise false.
<code>GTID_SUBTRACT()</code> [1274]	Return all GTIDs in set that are not in subset.
<code>SQL_THREAD_WAIT_AFTER_GTIDS()</code> [1274] (Deprecated since MySQL 5.6.9)	<b>OBSOLETE</b> Replaced by <code>WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS()</code>
<code>WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS()</code> [1274]	Wait until the slave SQL thread has executed all the given GTIDs. Returns: the number of events that were executed (or NULL, if GTID mode is not enabled).

- `GTID_SUBSET(subset, set)` [1274]

Given two sets of global transaction IDs *subset* and *set*, returns true if all GTIDs in *subset* are also in *set*. Returns false otherwise.

This function was added in MySQL 5.6.5.

- `GTID_SUBTRACT(set, subset)` [1274]

Given two sets of global transaction IDs *subset* and *set*, returns only those GTIDs from *set* that are not in *subset*.

This function was added in MySQL 5.6.5.

- `SQL_THREAD_WAIT_AFTER_GTIDS(gtid_set[, timeout])` [1274]

`SQL_THREAD_WAIT_AFTER_GTIDS()` was added in MySQL 5.6.5, and replaced by `WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS()` [1274] in MySQL 5.6.9. (Bug #14775984)

For more information, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

- `WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS(gtid_set[, timeout])` [1274]

Wait until the slave SQL thread has executed all of the transactions whose global transaction identifiers are contained in *gtid\_set* (see [Section 16.1.3.1, “GTID Concepts”](#), for a definition of “GTID sets”), or until *timeout* seconds have elapsed, whichever occurs first. *timeout* is optional; the default timeout is 0 seconds, in which case the master simply waits until all of the transactions in the GTID set have been executed.

Prior to MySQL 5.6.9, `WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS()` was named `SQL_THREAD_WAIT_AFTER_GTIDS()` [1274]. (Bug #14775984)

For more information, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

```
mysql> SELECT WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS('3E11FA47-71CA-11E1-9E33-C80AA9429562:1-5');
-> 5
```

The return value is the number of transactional events that were executed. Prior to MySQL 5.6.8, this function behaved unpredictably if no timeout was set and it was invoked while GTID-based replication was not active; in MySQL 5.6.8 and later, the function returns `NULL` whenever `gtid_mode` [2022] is `OFF`. (Bug #14640065)

## 12.16. Miscellaneous Functions

Table 12.20. Miscellaneous Functions

Name	Description
<code>DEFAULT()</code> [1275]	Return the default value for a table column



Name	Description
<a href="#">GET_LOCK() [1275]</a>	Get a named lock
<a href="#">INET_ATON() [1276]</a>	Return the numeric value of an IP address
<a href="#">INET_NTOA() [1276]</a>	Return the IP address from a numeric value
<a href="#">INET6_ATON() [1276]</a>	Return the numeric value of an IPv6 address
<a href="#">INET6_NTOA() [1277]</a>	Return the IPv6 address from a numeric value
<a href="#">IS_FREE_LOCK() [1278]</a>	Checks whether the named lock is free
<a href="#">IS_IPV4_COMPAT() [1278]</a>	Return true if argument is an IPv4-compatible address
<a href="#">IS_IPV4_MAPPED() [1278]</a>	Return true if argument is an IPv4-mapped address
<a href="#">IS_IPV4() [1278]</a>	Return true if argument is an IPv4 address
<a href="#">IS_IPV6() [1279]</a>	Return true if argument is an IPv6 address
<a href="#">IS_USED_LOCK() [1279]</a>	Checks whether the named lock is in use. Return connection identifier if true.
<a href="#">MASTER_POS_WAIT() [1279]</a>	Block until the slave has read and applied all updates up to the specified position
<a href="#">NAME_CONST() [1279]</a>	Causes the column to have the given name
<a href="#">RAND() [1206]</a>	Return a random floating-point value
<a href="#">RELEASE_LOCK() [1280]</a>	Releases the named lock
<a href="#">SLEEP() [1280]</a>	Sleep for a number of seconds
<a href="#">UUID_SHORT() [1281]</a>	Return an integer-valued universal identifier
<a href="#">UUID() [1280]</a>	Return a Universal Unique Identifier (UUID)
<a href="#">VALUES() [1281]</a>	Defines the values to be used during an INSERT

- [DEFAULT\(\*col\\_name\*\) \[1275\]](#)

Returns the default value for a table column. An error results if the column has no default value.

```
mysql> UPDATE t SET i = DEFAULT(i)+1 WHERE id < 100;
```

- [FORMAT\(\*X\*,\*D\*\) \[1180\]](#)

Formats the number *X* to a format like '#,###,###.##', rounded to *D* decimal places, and returns the result as a string. For details, see [Section 12.5, "String Functions"](#).

- [GET\\_LOCK\(\*str\*,\*timeout\*\) \[1275\]](#)

Tries to obtain a lock with a name given by the string *str*, using a timeout of *timeout* seconds. Returns 1 if the lock was obtained successfully, 0 if the attempt timed out (for example, because another client has previously locked the name), or NULL if an error occurred (such as running out of memory or the thread was killed with `mysqladmin kill`). If you have a lock obtained with [GET\\_LOCK\(\) \[1275\]](#), it is released when you execute [RELEASE\\_LOCK\(\) \[1280\]](#), execute a new [GET\\_LOCK\(\) \[1275\]](#), or your connection terminates (either normally or abnormally). Locks obtained with [GET\\_LOCK\(\) \[1275\]](#) do not interact with transactions. That is, committing a transaction does not release any such locks obtained during the transaction.

This function can be used to implement application locks or to simulate record locks. Names are locked on a server-wide basis. If a name has been locked by one client, [GET\\_LOCK\(\) \[1275\]](#) blocks any request by another client for a lock with the same name. This enables clients that agree on a given lock name to use the name to perform cooperative advisory locking. But be aware that it also enables a client that is not among the set of cooperating clients to lock a name, either inadvertently or deliberately, and thus prevent any of the cooperating clients from locking that name. One way to

reduce the likelihood of this is to use lock names that are database-specific or application-specific. For example, use lock names of the form `db_name.str` or `app_name.str`.

```
mysql> SELECT GET_LOCK('lock1',10);
-> 1
mysql> SELECT IS_FREE_LOCK('lock2');
-> 1
mysql> SELECT GET_LOCK('lock2',10);
-> 1
mysql> SELECT RELEASE_LOCK('lock2');
-> 1
mysql> SELECT RELEASE_LOCK('lock1');
-> NULL
```

The second `RELEASE_LOCK()` [1280] call returns `NULL` because the lock 'lock1' was automatically released by the second `GET_LOCK()` [1275] call.

If multiple clients are waiting for a lock, the order in which they will acquire it is undefined and depends on factors such as the thread library in use. In particular, applications should not assume that clients will acquire the lock in the same order that they issued the lock requests.

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `INET_ATON(expr)` [1276]

Given the dotted-quad representation of an IPv4 network address as a string, returns an integer that represents the numeric value of the address in network byte order (big endian). `INET_ATON()` [1276] returns `NULL` if it does not understand its argument.

```
mysql> SELECT INET_ATON('10.0.5.9');
-> 167773449
```

For this example, the return value is calculated as  $10 \times 256^3 + 0 \times 256^2 + 5 \times 256 + 9$ .

`INET_ATON()` [1276] may or may not return a non-`NULL` result for short-form IP addresses (such as '127.1' as a representation of '127.0.0.1'). Because of this, `INET_ATON()` [1276] should not be used for such addresses.



#### Note

To store values generated by `INET_ATON()` [1276], use an `INT UNSIGNED` column rather than `INT`, which is signed. If you use a signed column, values corresponding to IP addresses for which the first octet is greater than 127 cannot be stored correctly. See [Section 11.2.6, “Out-of-Range and Overflow Handling”](#).

- `INET_NTOA(expr)` [1276]

Given a numeric IPv4 network address in network byte order, returns the dotted-quad string representation of the address as a nonbinary string in the connection character set. `INET_NTOA()` [1276] returns `NULL` if it does not understand its argument.

```
mysql> SELECT INET_NTOA(167773449);
-> '10.0.5.9'
```

- `INET6_ATON(expr)` [1276]

Given an IPv6 or IPv4 network address as a string, returns a binary string that represents the numeric value of the address in network byte order (big endian). Because numeric-format IPv6 addresses require more bytes than the largest integer type, the representation returned by this function has the `VARBINARY` data type: `VARBINARY(16)` for IPv6 addresses and `VARBINARY(4)` for IPv4 addresses. If the argument is not a valid address, `INET6_ATON()` [1276] returns `NULL`.

The following examples use `HEX()` to display the `INET6_ATON()` result in printable form:

```
mysql> SELECT HEX(INET6_ATON('fdfe::5a55:caff:fefa:9089'));
-> 'FDFE0000000000005A55CAFFFEFA9089'
mysql> SELECT HEX(INET6_ATON('10.0.5.9'));
-> '0A000509'
```

`INET6_ATON()` observes several constraints on valid arguments. These are given in the following list along with examples.

- A trailing zone ID is not permitted, as in `fe80::3%1` or `fe80::3%eth0`.
- A trailing network mask is not permitted, as in `2001:45f:3:ba::/64` or `192.168.1.0/24`.
- For values representing IPv4 addresses, only classless addresses are supported. Classful addresses such as `192.168.1` are rejected. A trailing port number is not permitted, as in `192.168.1.2:8080`. Hexadecimal numbers in address components are not permitted, as in `192.0xa0.1.2`. Octal numbers are not supported: `192.168.010.1` is treated as `192.168.10.1`, not `192.168.8.1`. These IPv4 constraints also apply to IPv6 addresses that have IPv4 address parts, such as IPv4-compatible or IPv4-mapped addresses.

To convert an IPv4 address `expr` represented in numeric form as an `INT` value to an IPv6 address represented in numeric form as a `VARBINARY` value, use this expression:

```
INET6_ATON(INET_NTOA(expr))
```

For example:

```
mysql> SELECT HEX(INET6_ATON(INET_NTOA(167773449)));
-> '0A000509'
```

This function was added in MySQL 5.6.3.

- `INET6_NTOA(expr)` [1277]

Given an IPv6 or IPv4 network address represented in numeric form as a binary string, returns the string representation of the address as a nonbinary string in the connection character set. If the argument is not a valid address, `INET6_NTOA()` [1277] returns `NULL`.

`INET6_NTOA()` [1277] has these properties:

- It does not use operating system functions to perform conversions, thus the output string is platform independent.
- The return string has a maximum length of 39 (4 x 8 + 7). Given this statement:

```
CREATE TABLE t AS SELECT INET6_NTOA(expr) AS c1;
```

The resulting table would have this definition:

```
CREATE TABLE t (c1 VARCHAR(39) CHARACTER SET utf8 DEFAULT NULL);
```

- The return string uses lowercase letters for IPv6 addresses.

```
mysql> SELECT INET6_NTOA(INET6_ATON('fdfe::5a55:caff:fefa:9089'));
-> 'fdfe::5a55:caff:fefa:9089'
mysql> SELECT INET6_NTOA(INET6_ATON('10.0.5.9'));
-> '10.0.5.9'

mysql> SELECT INET6_NTOA(UNHEX('FDFE0000000000005A55CAFFFEFA9089'));
-> 'fdfe::5a55:caff:fefa:9089'
mysql> SELECT INET6_NTOA(UNHEX('0A000509'));
-> '10.0.5.9'
```

This function was added in MySQL 5.6.3.

- `IS_FREE_LOCK(str)` [1278]

Checks whether the lock named *str* is free to use (that is, not locked). Returns 1 if the lock is free (no one is using the lock), 0 if the lock is in use, and `NULL` if an error occurs (such as an incorrect argument).

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `IS_IPV4(expr)` [1278]

Returns 1 if the argument is a valid IPv4 address specified as a string, 0 otherwise.

```
mysql> SELECT IS_IPV4('10.0.5.9'), IS_IPV4('10.0.5.256');
-> 1, 0
```

For a given argument, if `IS_IPV4()` [1278] returns 1, `INET_ATON()` [1276] (and `INET6_ATON()` [1276]) will return non-`NULL`. The converse statement is not true: In some cases, `INET_ATON()` [1276] returns non-`NULL` when `IS_IPV4()` [1278] returns 0.

As implied by the preceding remarks, `IS_IPV4()` [1278] is more strict than `INET_ATON()` [1276] about what constitutes a valid IPv4 address, so it may be useful for applications that need to perform strong checks against invalid values. Alternatively, use `INET6_ATON()` [1276] to convert IPv4 addresses to internal form and check for a `NULL` result (which indicates an invalid address). `INET6_ATON()` [1276] is equally strong as `IS_IPV4()` [1278] about checking IPv4 addresses.

This function was added in MySQL 5.6.3.

- `IS_IPV4_COMPAT(expr)` [1278]

This function takes an IPv6 address represented in numeric form as a binary string, as returned by `INET6_ATON()` [1276]. It returns 1 if the argument is a valid IPv4-compatible IPv6 address, 0 otherwise. IPv4-compatible addresses have the form `::ipv4_address`.

```
mysql> SELECT IS_IPV4_COMPAT(INET6_ATON('::10.0.5.9'));
-> 1
mysql> SELECT IS_IPV4_COMPAT(INET6_ATON('::ffff:10.0.5.9'));
-> 0
```

The IPv4 part of an IPv4-compatible address can also be represented using hexadecimal notation. For example, `192.168.0.1` has this raw hexadecimal value:

```
mysql> SELECT HEX(INET6_ATON('192.168.0.1'));
-> 'C0A80001'
```

Expressed in IPv4-compatible form, `::192.168.0.1` is equivalent to `::c0a8:0001` or (without leading zeros) `::c0a8:1`

```
mysql> SELECT
-> IS_IPV4_COMPAT(INET6_ATON('::192.168.0.1')),
-> IS_IPV4_COMPAT(INET6_ATON('::c0a8:0001')),
-> IS_IPV4_COMPAT(INET6_ATON('::c0a8:1'));
-> 1, 1, 1
```

This function was added in MySQL 5.6.3.

- `IS_IPV4_MAPPED(expr)` [1278]

This function takes an IPv6 address represented in numeric form as a binary string, as returned by `INET6_ATON()` [1276]. It returns 1 if the argument is a valid IPv4-mapped IPv6 address, 0 otherwise. IPv4-mapped addresses have the form `::ffff:ipv4_address`.

```
mysql> SELECT IS_IPV4_MAPPED(INET6_ATON('::10.0.5.9'));
-> 0
mysql> SELECT IS_IPV4_MAPPED(INET6_ATON('::ffff:10.0.5.9'));
-> 1
```

As with `IS_IPV4_COMPAT()` the IPv4 part of an IPv4-mapped address can also be represented using hexadecimal notation:

```
mysql> SELECT
-> IS_IPV4_MAPPED(INET6_ATON('::ffff:192.168.0.1')),
-> IS_IPV4_MAPPED(INET6_ATON('::ffff:c0a8:0001')),
-> IS_IPV4_MAPPED(INET6_ATON('::ffff:c0a8:1'));
-> 1, 1, 1
```

This function was added in MySQL 5.6.3.

- `IS_IPV6(expr)` [1279]

Returns 1 if the argument is a valid IPv6 address specified as a string, 0 otherwise. This function does not consider IPv4 addresses to be valid IPv6 addresses.

```
mysql> SELECT IS_IPV6('10.0.5.9'), IS_IPV6('::1');
-> 0, 1
```

For a given argument, if `IS_IPV6()` [1279] returns 1, `INET6_ATON()` [1276] will return non-NULL.

This function was added in MySQL 5.6.3.

- `IS_USED_LOCK(str)` [1279]

Checks whether the lock named *str* is in use (that is, locked). If so, it returns the connection identifier of the client that holds the lock. Otherwise, it returns NULL.

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `MASTER_POS_WAIT(log_name, log_pos [, timeout])` [1279]

This function is useful for control of master/slave synchronization. It blocks until the slave has read and applied all updates up to the specified position in the master log. The return value is the number of log events the slave had to wait for to advance to the specified position. The function returns NULL if the slave SQL thread is not started, the slave's master information is not initialized, the arguments are incorrect, or an error occurs. It returns -1 if the timeout has been exceeded. If the slave SQL thread stops while `MASTER_POS_WAIT()` [1279] is waiting, the function returns NULL. If the slave is past the specified position, the function returns immediately.

If a *timeout* value is specified, `MASTER_POS_WAIT()` [1279] stops waiting when *timeout* seconds have elapsed. *timeout* must be greater than 0; a zero or negative *timeout* means no timeout.

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `NAME_CONST(name, value)` [1279]

Returns the given value. When used to produce a result set column, `NAME_CONST()` [1279] causes the column to have the given name. The arguments should be constants.

```
mysql> SELECT NAME_CONST('myname', 14);
+-----+
| myname |
+-----+
|      14 |
+-----+
```

This function is for internal use only. The server uses it when writing statements from stored programs that contain references to local program variables, as described in [Section 18.7, “Binary Logging of Stored Programs”](#). You might see this function in the output from `mysqlbinlog`.

- `RELEASE_LOCK(str)` [1280]

Releases the lock named by the string *str* that was obtained with `GET_LOCK()` [1275]. Returns 1 if the lock was released, 0 if the lock was not established by this thread (in which case the lock is not released), and `NULL` if the named lock did not exist. The lock does not exist if it was never obtained by a call to `GET_LOCK()` [1275] or if it has previously been released.

The `DO` statement is convenient to use with `RELEASE_LOCK()` [1280]. See [Section 13.2.3, “DO Syntax”](#).

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `SLEEP(duration)` [1280]

Sleeps (pauses) for the number of seconds given by the *duration* argument, then returns 0. If `SLEEP()` [1280] is interrupted, it returns 1. The duration may have a fractional part.

This function is unsafe for statement-based replication. In MySQL 5.6, a warning is logged if you use this function when `binlog_format` [2010] is set to `STATEMENT`. (Bug #47995)

- `UUID()` [1280]

Returns a Universal Unique Identifier (UUID) generated according to “DCE 1.1: Remote Procedure Call” (Appendix A) CAE (Common Applications Environment) Specifications published by The Open Group in October 1997 (Document Number C706, <http://www.opengroup.org/public/pubs/catalog/c706.htm>).

A UUID is designed as a number that is globally unique in space and time. Two calls to `UUID()` [1280] are expected to generate two different values, even if these calls are performed on two separate computers that are not connected to each other.

A UUID is a 128-bit number represented by a `utf8` string of five hexadecimal numbers in `aaaaaaaa-bbbb-cccc-dddd-eeeeeeeeeeee` format:

- The first three numbers are generated from a timestamp.
- The fourth number preserves temporal uniqueness in case the timestamp value loses monotonicity (for example, due to daylight saving time).
- The fifth number is an IEEE 802 node number that provides spatial uniqueness. A random number is substituted if the latter is not available (for example, because the host computer has no Ethernet card, or we do not know how to find the hardware address of an interface on your operating system). In this case, spatial uniqueness cannot be guaranteed. Nevertheless, a collision should have very low probability.

Currently, the MAC address of an interface is taken into account only on FreeBSD and Linux. On other operating systems, MySQL uses a randomly generated 48-bit number.

```
mysql> SELECT UUID();
-> '6ccd780c-baba-1026-9564-0040f4311e29'
```



### Warning

Although `UUID()` [1280] values are intended to be unique, they are not necessarily unguessable or unpredictable. If unpredictability is required, UUID values should be generated some other way.


**Note**

`UUID()` [1280] does not work with statement-based replication.

- `UUID_SHORT()` [1281]

Returns a “short” universal identifier as a 64-bit unsigned integer (rather than a string-form 128-bit identifier as returned by the `UUID()` [1280] function).

The value of `UUID_SHORT()` [1281] is guaranteed to be unique if the following conditions hold:

- The `server_id` [539] of the current host is unique among your set of master and slave servers
- `server_id` [539] is between 0 and 255
- You do not set back your system time for your server between `mysqld` restarts
- You do not invoke `UUID_SHORT()` [1281] on average more than 16 million times per second between `mysqld` restarts

The `UUID_SHORT()` [1281] return value is constructed this way:

```
(server_id & 255) << 56
+ (server_startup_time_in_seconds << 24)
+ incremented_variable++;
```

```
mysql> SELECT UUID_SHORT();
-> 92395783831158784
```

Note that `UUID_SHORT()` [1281] does not work with statement-based replication.

- `VALUES(col_name)` [1281]

In an `INSERT ... ON DUPLICATE KEY UPDATE` statement, you can use the `VALUES(col_name)` function in the `UPDATE` clause to refer to column values from the `INSERT` portion of the statement. In other words, `VALUES(col_name)` in the `UPDATE` clause refers to the value of `col_name` that would be inserted, had no duplicate-key conflict occurred. This function is especially useful in multiple-row inserts. The `VALUES()` [1281] function is meaningful only in the `ON DUPLICATE KEY UPDATE` clause of `INSERT` statements and returns `NULL` otherwise. See [Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”](#).

```
mysql> INSERT INTO table (a,b,c) VALUES (1,2,3),(4,5,6)
-> ON DUPLICATE KEY UPDATE c=VALUES(a)+VALUES(b);
```

## 12.17. Functions and Modifiers for Use with `GROUP BY` Clauses

### 12.17.1. `GROUP BY` (Aggregate) Functions

Table 12.21. Aggregate (`GROUP BY`) Functions

Name	Description
<code>AVG()</code> [1282]	Return the average value of the argument
<code>BIT_AND()</code> [1282]	Return bitwise and
<code>BIT_OR()</code> [1282]	Return bitwise or
<code>BIT_XOR()</code> [1283]	Return bitwise xor
<code>COUNT(DISTINCT)</code> [1283]	Return the count of a number of different values
<code>COUNT()</code> [1283]	Return a count of the number of rows returned
<code>GROUP_CONCAT()</code> [1283]	Return a concatenated string



Name	Description
<a href="#">MAX() [1284]</a>	Return the maximum value
<a href="#">MIN() [1284]</a>	Return the minimum value
<a href="#">STD() [1284]</a>	Return the population standard deviation
<a href="#">STDDEV_POP() [1285]</a>	Return the population standard deviation
<a href="#">STDDEV_SAMP() [1285]</a>	Return the sample standard deviation
<a href="#">STDDEV() [1285]</a>	Return the population standard deviation
<a href="#">SUM() [1285]</a>	Return the sum
<a href="#">VAR_POP() [1285]</a>	Return the population standard variance
<a href="#">VAR_SAMP() [1285]</a>	Return the sample variance
<a href="#">VARIANCE() [1285]</a>	Return the population standard variance

This section describes group (aggregate) functions that operate on sets of values. Unless otherwise stated, group functions ignore `NULL` values.

If you use a group function in a statement containing no `GROUP BY` clause, it is equivalent to grouping on all rows. For more information, see [Section 12.17.3, “MySQL Extensions to GROUP BY”](#).

For numeric arguments, the variance and standard deviation functions return a `DOUBLE` value. The [SUM\(\) \[1285\]](#) and [AVG\(\) \[1282\]](#) functions return a `DECIMAL` value for exact-value arguments (integer or `DECIMAL`), and a `DOUBLE` value for approximate-value arguments (`FLOAT` or `DOUBLE`).

The [SUM\(\) \[1285\]](#) and [AVG\(\) \[1282\]](#) aggregate functions do not work with temporal values. (They convert the values to numbers, losing everything after the first nonnumeric character.) To work around this problem, convert to numeric units, perform the aggregate operation, and convert back to a temporal value. Examples:

```
SELECT SEC_TO_TIME(SUM(TIME_TO_SEC(time_col))) FROM tbl_name;  
SELECT FROM_DAYS(SUM(TO_DAYS(date_col))) FROM tbl_name;
```

Functions such as [SUM\(\) \[1285\]](#) or [AVG\(\) \[1282\]](#) that expect a numeric argument cast the argument to a number if necessary. For `SET` or `ENUM` values, the cast operation causes the underlying numeric value to be used.

- [AVG\(\[DISTINCT\] \*expr\*\) \[1282\]](#)

Returns the average value of *expr*. The `DISTINCT` option can be used to return the average of the distinct values of *expr*.

[AVG\(\) \[1282\]](#) returns `NULL` if there were no matching rows.

```
mysql> SELECT student_name, AVG(test_score)  
-> FROM student  
-> GROUP BY student_name;
```

- [BIT\\_AND\(\*expr\*\) \[1282\]](#)

Returns the bitwise `AND` of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

This function returns `18446744073709551615` if there were no matching rows. (This is the value of an unsigned `BIGINT` value with all bits set to 1.)

- [BIT\\_OR\(\*expr\*\) \[1282\]](#)

Returns the bitwise `OR` of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

This function returns `0` if there were no matching rows.



- `BIT_XOR(expr)` [1283]

Returns the bitwise `XOR` [1172] of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

This function returns 0 if there were no matching rows.

- `COUNT(expr)` [1283]

Returns a count of the number of non-NULL values of *expr* in the rows retrieved by a `SELECT` statement. The result is a `BIGINT` value.

`COUNT()` [1283] returns 0 if there were no matching rows.

```
mysql> SELECT student.student_name,COUNT(*)
->      FROM student,course
->      WHERE student.student_id=course.student_id
->      GROUP BY student_name;
```

`COUNT(*)` [1283] is somewhat different in that it returns a count of the number of rows retrieved, whether or not they contain `NULL` values.

`COUNT(*)` [1283] is optimized to return very quickly if the `SELECT` retrieves from one table, no other columns are retrieved, and there is no `WHERE` clause. For example:

```
mysql> SELECT COUNT(*) FROM student;
```

This optimization applies only to `MyISAM` tables only, because an exact row count is stored for this storage engine and can be accessed very quickly. For transactional storage engines such as `InnoDB`, storing an exact row count is more problematic because multiple transactions may be occurring, each of which may affect the count.

- `COUNT(DISTINCT expr, [expr...])` [1283]

Returns a count of the number of rows with different non-NULL *expr* values.

`COUNT(DISTINCT)` [1283] returns 0 if there were no matching rows.

```
mysql> SELECT COUNT(DISTINCT results) FROM student;
```

In MySQL, you can obtain the number of distinct expression combinations that do not contain `NULL` by giving a list of expressions. In standard SQL, you would have to do a concatenation of all expressions inside `COUNT(DISTINCT ...)` [1283].

- `GROUP_CONCAT(expr)` [1283]

This function returns a string result with the concatenated non-NULL values from a group. It returns `NULL` if there are no non-NULL values. The full syntax is as follows:

```
GROUP_CONCAT([DISTINCT] expr [,expr ...]
             [ORDER BY {unsigned_integer | col_name | expr}
              [ASC | DESC] [,col_name ...]]
             [SEPARATOR str_val])
```

```
mysql> SELECT student_name,
->      GROUP_CONCAT(test_score)
->      FROM student
->      GROUP BY student_name;
```

Or:

```
mysql> SELECT student_name,
->      GROUP_CONCAT(DISTINCT test_score
->                    ORDER BY test_score DESC SEPARATOR ' ')
->      FROM student
```

```
-> GROUP BY student_name;
```

In MySQL, you can get the concatenated values of expression combinations. To eliminate duplicate values, use the `DISTINCT` clause. To sort values in the result, use the `ORDER BY` clause. To sort in reverse order, add the `DESC` (descending) keyword to the name of the column you are sorting by in the `ORDER BY` clause. The default is ascending order; this may be specified explicitly using the `ASC` keyword. The default separator between values in a group is comma (","). To specify a separator explicitly, use `SEPARATOR` followed by the string literal value that should be inserted between group values. To eliminate the separator altogether, specify `SEPARATOR ''`.

The result is truncated to the maximum length that is given by the `group_concat_max_len` [486] system variable, which has a default value of 1024. The value can be set higher, although the effective maximum length of the return value is constrained by the value of `max_allowed_packet` [502]. The syntax to change the value of `group_concat_max_len` [486] at runtime is as follows, where *val* is an unsigned integer:

```
SET [GLOBAL | SESSION] group_concat_max_len = val;
```

The return value is a nonbinary or binary string, depending on whether the arguments are nonbinary or binary strings. The result type is `TEXT` or `BLOB` unless `group_concat_max_len` [486] is less than or equal to 512, in which case the result type is `VARCHAR` or `VARBINARY`.

See also `CONCAT()` [1178] and `CONCAT_WS()` [1179]: Section 12.5, "String Functions".

- `MAX([DISTINCT] expr)` [1284]

Returns the maximum value of *expr*. `MAX()` [1284] may take a string argument; in such cases, it returns the maximum string value. See Section 8.3.1, "How MySQL Uses Indexes". The `DISTINCT` keyword can be used to find the maximum of the distinct values of *expr*, however, this produces the same result as omitting `DISTINCT`.

`MAX()` [1284] returns `NULL` if there were no matching rows.

```
mysql> SELECT student_name, MIN(test_score), MAX(test_score)
-> FROM student
-> GROUP BY student_name;
```

For `MAX()` [1284], MySQL currently compares `ENUM` and `SET` columns by their string value rather than by the string's relative position in the set. This differs from how `ORDER BY` compares them. This is expected to be rectified in a future MySQL release.

- `MIN([DISTINCT] expr)` [1284]

Returns the minimum value of *expr*. `MIN()` [1284] may take a string argument; in such cases, it returns the minimum string value. See Section 8.3.1, "How MySQL Uses Indexes". The `DISTINCT` keyword can be used to find the minimum of the distinct values of *expr*, however, this produces the same result as omitting `DISTINCT`.

`MIN()` [1284] returns `NULL` if there were no matching rows.

```
mysql> SELECT student_name, MIN(test_score), MAX(test_score)
-> FROM student
-> GROUP BY student_name;
```

For `MIN()` [1284], MySQL currently compares `ENUM` and `SET` columns by their string value rather than by the string's relative position in the set. This differs from how `ORDER BY` compares them. This is expected to be rectified in a future MySQL release.

- `STD(expr)` [1284]

Returns the population standard deviation of *expr*. This is an extension to standard SQL. The standard SQL function `STDDEV_POP()` [1285] can be used instead.

This function returns `NULL` if there were no matching rows.

- `STDDEV(expr)` [1285]

Returns the population standard deviation of *expr*. This function is provided for compatibility with Oracle. The standard SQL function `STDDEV_POP()` [1285] can be used instead.

This function returns `NULL` if there were no matching rows.

- `STDDEV_POP(expr)` [1285]

Returns the population standard deviation of *expr* (the square root of `VAR_POP()` [1285]). You can also use `STD()` [1284] or `STDDEV()` [1285], which are equivalent but not standard SQL.

`STDDEV_POP()` [1285] returns `NULL` if there were no matching rows.

- `STDDEV_SAMP(expr)` [1285]

Returns the sample standard deviation of *expr* (the square root of `VAR_SAMP()` [1285]).

`STDDEV_SAMP()` [1285] returns `NULL` if there were no matching rows.

- `SUM([DISTINCT] expr)` [1285]

Returns the sum of *expr*. If the return set has no rows, `SUM()` [1285] returns `NULL`. The `DISTINCT` keyword can be used to sum only the distinct values of *expr*.

`SUM()` [1285] returns `NULL` if there were no matching rows.

- `VAR_POP(expr)` [1285]

Returns the population standard variance of *expr*. It considers rows as the whole population, not as a sample, so it has the number of rows as the denominator. You can also use `VARIANCE()` [1285], which is equivalent but is not standard SQL.

`VAR_POP()` [1285] returns `NULL` if there were no matching rows.

- `VAR_SAMP(expr)` [1285]

Returns the sample variance of *expr*. That is, the denominator is the number of rows minus one.

`VAR_SAMP()` [1285] returns `NULL` if there were no matching rows.

- `VARIANCE(expr)` [1285]

Returns the population standard variance of *expr*. This is an extension to standard SQL. The standard SQL function `VAR_POP()` [1285] can be used instead.

`VARIANCE()` [1285] returns `NULL` if there were no matching rows.

## 12.17.2. GROUP BY Modifiers

The `GROUP BY` clause permits a `WITH ROLLUP` modifier that causes extra rows to be added to the summary output. These rows represent higher-level (or super-aggregate) summary operations. `ROLLUP` thus enables you to answer questions at multiple levels of analysis with a single query. It can be used, for example, to provide support for OLAP (Online Analytical Processing) operations.

Suppose that a table named `sales` has `year`, `country`, `product`, and `profit` columns for recording sales profitability:

```
CREATE TABLE sales
(
```

```

year      INT NOT NULL,
country  VARCHAR(20) NOT NULL,
product  VARCHAR(32) NOT NULL,
profit   INT
);

```

The table's contents can be summarized per year with a simple `GROUP BY` like this:

```

mysql> SELECT year, SUM(profit) FROM sales GROUP BY year;
+-----+-----+
| year | SUM(profit) |
+-----+-----+
| 2000 |          4525 |
| 2001 |          3010 |
+-----+-----+

```

This output shows the total profit for each year, but if you also want to determine the total profit summed over all years, you must add up the individual values yourself or run an additional query.

Or you can use `ROLLUP`, which provides both levels of analysis with a single query. Adding a `WITH ROLLUP` modifier to the `GROUP BY` clause causes the query to produce another row that shows the grand total over all year values:

```

mysql> SELECT year, SUM(profit) FROM sales GROUP BY year WITH ROLLUP;
+-----+-----+
| year | SUM(profit) |
+-----+-----+
| 2000 |          4525 |
| 2001 |          3010 |
| NULL |          7535 |
+-----+-----+

```

The grand total super-aggregate line is identified by the value `NULL` in the `year` column.

`ROLLUP` has a more complex effect when there are multiple `GROUP BY` columns. In this case, each time there is a “break” (change in value) in any but the last grouping column, the query produces an extra super-aggregate summary row.

For example, without `ROLLUP`, a summary on the `sales` table based on `year`, `country`, and `product` might look like this:

```

mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product;
+-----+-----+-----+-----+
| year | country | product | SUM(profit) |
+-----+-----+-----+-----+
| 2000 | Finland | Computer |          1500 |
| 2000 | Finland | Phone    |           100 |
| 2000 | India   | Calculator |           150 |
| 2000 | India   | Computer |          1200 |
| 2000 | USA     | Calculator |            75 |
| 2000 | USA     | Computer |          1500 |
| 2001 | Finland | Phone    |            10 |
| 2001 | USA     | Calculator |            50 |
| 2001 | USA     | Computer |          2700 |
| 2001 | USA     | TV       |            250 |
+-----+-----+-----+-----+

```

The output indicates summary values only at the year/country/product level of analysis. When `ROLLUP` is added, the query produces several extra rows:

```

mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product WITH ROLLUP;
+-----+-----+-----+-----+
| year | country | product | SUM(profit) |
+-----+-----+-----+-----+

```

2000	Finland	Computer	1500
2000	Finland	Phone	100
2000	Finland	NULL	1600
2000	India	Calculator	150
2000	India	Computer	1200
2000	India	NULL	1350
2000	USA	Calculator	75
2000	USA	Computer	1500
2000	USA	NULL	1575
2000	NULL	NULL	4525
2001	Finland	Phone	10
2001	Finland	NULL	10
2001	USA	Calculator	50
2001	USA	Computer	2700
2001	USA	TV	250
2001	USA	NULL	3000
2001	NULL	NULL	3010
NULL	NULL	NULL	7535

For this query, adding `ROLLUP` causes the output to include summary information at four levels of analysis, not just one. Here is how to interpret the `ROLLUP` output:

- Following each set of product rows for a given year and country, an extra summary row is produced showing the total for all products. These rows have the `product` column set to `NULL`.
- Following each set of rows for a given year, an extra summary row is produced showing the total for all countries and products. These rows have the `country` and `products` columns set to `NULL`.
- Finally, following all other rows, an extra summary row is produced showing the grand total for all years, countries, and products. This row has the `year`, `country`, and `products` columns set to `NULL`.

### Other Considerations When using `ROLLUP`

The following items list some behaviors specific to the MySQL implementation of `ROLLUP`:

When you use `ROLLUP`, you cannot also use an `ORDER BY` clause to sort the results. In other words, `ROLLUP` and `ORDER BY` are mutually exclusive. However, you still have some control over sort order. `GROUP BY` in MySQL sorts results, and you can use explicit `ASC` and `DESC` keywords with columns named in the `GROUP BY` list to specify sort order for individual columns. (The higher-level summary rows added by `ROLLUP` still appear after the rows from which they are calculated, regardless of the sort order.)

`LIMIT` can be used to restrict the number of rows returned to the client. `LIMIT` is applied after `ROLLUP`, so the limit applies against the extra rows added by `ROLLUP`. For example:

```
mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product WITH ROLLUP
-> LIMIT 5;
```

year	country	product	SUM(profit)
2000	Finland	Computer	1500
2000	Finland	Phone	100
2000	Finland	NULL	1600
2000	India	Calculator	150
2000	India	Computer	1200

Using `LIMIT` with `ROLLUP` may produce results that are more difficult to interpret, because you have less context for understanding the super-aggregate rows.

The `NULL` indicators in each super-aggregate row are produced when the row is sent to the client. The server looks at the columns named in the `GROUP BY` clause following the leftmost one that has

changed value. For any column in the result set with a name that is a lexical match to any of those names, its value is set to `NULL`. (If you specify grouping columns by column number, the server identifies which columns to set to `NULL` by number.)

Because the `NULL` values in the super-aggregate rows are placed into the result set at such a late stage in query processing, you cannot test them as `NULL` values within the query itself. For example, you cannot add `HAVING product IS NULL` to the query to eliminate from the output all but the super-aggregate rows.

On the other hand, the `NULL` values do appear as `NULL` on the client side and can be tested as such using any MySQL client programming interface.

### 12.17.3. MySQL Extensions to `GROUP BY`

In standard SQL, a query that includes a `GROUP BY` clause cannot refer to nonaggregated columns in the select list that are not named in the `GROUP BY` clause. For example, this query is illegal in standard SQL because the `name` column in the select list does not appear in the `GROUP BY`:

```
SELECT o.custid, c.name, MAX(o.payment)
   FROM orders AS o, customers AS c
  WHERE o.custid = c.custid
  GROUP BY o.custid;
```

For the query to be legal, the `name` column must be omitted from the select list or named in the `GROUP BY` clause.

MySQL extends the use of `GROUP BY` so that the select list can refer to nonaggregated columns not named in the `GROUP BY` clause. This means that the preceding query is legal in MySQL. You can use this feature to get better performance by avoiding unnecessary column sorting and grouping. However, this is useful primarily when all values in each nonaggregated column not named in the `GROUP BY` are the same for each group. The server is free to choose any value from each group, so unless they are the same, the values chosen are indeterminate. Furthermore, the selection of values from each group cannot be influenced by adding an `ORDER BY` clause. Sorting of the result set occurs after values have been chosen, and `ORDER BY` does not affect which values within each group the server chooses.

A similar MySQL extension applies to the `HAVING` clause. In standard SQL, a query that includes a `GROUP BY` clause cannot refer to nonaggregated columns in the `HAVING` clause that are not named in the `GROUP BY` clause. A MySQL extension permits references to such columns to simplify calculations. This extension assumes that the nongrouped columns will have the same group-wise values. Otherwise, the result is indeterminate.

To disable the MySQL `GROUP BY` extension, enable the `ONLY_FULL_GROUP_BY [606]` SQL mode. This enables standard SQL behavior: Columns not named in the `GROUP BY` clause cannot be used in the select list or `HAVING` clause unless enclosed in an aggregate function.

`ONLY_FULL_GROUP_BY [606]` also affects use of aliases in the `HAVING` clauses. For example, the following query returns `name` values that occur only once in table `orders`:

```
SELECT name, COUNT(name) FROM orders
   GROUP BY name
  HAVING COUNT(name) = 1;
```

MySQL extends this behavior to permit the use of an alias in the `HAVING` clause for the aggregated column:

```
SELECT name, COUNT(name) AS c FROM orders
   GROUP BY name
  HAVING c = 1;
```

Enabling `ONLY_FULL_GROUP_BY [606]` disables this MySQL extension and a `non-grouping field 'c' is used in HAVING clause` error occurs because the column `c` in the `HAVING` clause is not enclosed in an aggregate function (instead, it is an aggregate function).

The select list extension also applies to `ORDER BY`. That is, you can refer to nonaggregated columns in the `ORDER BY` clause that do not appear in the `GROUP BY` clause. (However, as mentioned previously, `ORDER BY` does not affect which values are chosen from nonaggregated columns; it only sorts them after they have been chosen.) This extension does not apply if the `ONLY_FULL_GROUP_BY` [606] SQL mode is enabled.

In some cases, you can use `MIN()` [1284] and `MAX()` [1284] to obtain a specific column value even if it is not unique. If the `sort` column contains integers no larger than 6 digits, the following query gives the value of `column` from the row containing the smallest `sort` value:

```
SUBSTR(MIN(CONCAT(LPAD(sort,6,'0'),column)),7)
```

See Section 3.6.4, “The Rows Holding the Group-wise Maximum of a Certain Column”.

If you are trying to follow standard SQL, you cannot use expressions in `GROUP BY` clauses. As a workaround, use an alias for the expression:

```
SELECT id, FLOOR(value/100) AS val
FROM tbl_name
GROUP BY id, val;
```

MySQL permits expressions in `GROUP BY` clauses, so the alias is unnecessary:

```
SELECT id, FLOOR(value/100)
FROM tbl_name
GROUP BY id, FLOOR(value/100);
```

## 12.18. Spatial Extensions

MySQL supports spatial extensions to enable the generation, storage, and analysis of geographic features. These features are available for `MyISAM`, `InnoDB`, `NDB`, and `ARCHIVE` tables.

For spatial columns, `MyISAM` supports both `SPATIAL` and non-`SPATIAL` indexes. Other storage engines support non-`SPATIAL` indexes, as described in Section 13.1.11, “`CREATE INDEX` Syntax”.

This chapter covers the following topics:

- The basis of these spatial extensions in the OpenGIS geometry model
- Data formats for representing spatial data
- How to use spatial data in MySQL
- Use of indexing for spatial data
- MySQL differences from the OpenGIS specification

### Additional Resources

- The Open Geospatial Consortium publishes the *OpenGIS® Simple Features Specifications For SQL*, a document that proposes several conceptual ways for extending an SQL RDBMS to support spatial data. This specification is available from the OGC Web site at <http://www.opengis.org/docs/99-049.pdf>.
- If you have questions or concerns about the use of the spatial extensions to MySQL, you can discuss them in the GIS forum: <http://forums.mysql.com/list.php?23>.

### 12.18.1. Introduction to MySQL Spatial Support

MySQL implements spatial extensions following the specification of the Open Geospatial Consortium (OGC). This is an international consortium of more than 250 companies, agencies, and universities participating in the development of publicly available conceptual solutions that can be useful



with all kinds of applications that manage spatial data. The OGC maintains a Web site at <http://www.opengis.org/>.

In 1997, the Open Geospatial Consortium published the *OpenGIS® Simple Features Specifications For SQL*, a document that proposes several conceptual ways for extending an SQL RDBMS to support spatial data. This specification is available from the OGC Web site at <http://www.opengis.org/docs/99-049.pdf>. It contains additional information relevant to this chapter.

MySQL implements a subset of the **SQL with Geometry Types** environment proposed by OGC. This term refers to an SQL environment that has been extended with a set of geometry types. A geometry-valued SQL column is implemented as a column that has a geometry type. The specification describe a set of SQL geometry types, as well as functions on those types to create and analyze geometry values.

A **geographic feature** is anything in the world that has a location. A feature can be:

- An entity. For example, a mountain, a pond, a city.
- A space. For example, town district, the tropics.
- A definable location. For example, a crossroad, as a particular place where two streets intersect.

Some documents use the term **geospatial feature** to refer to geographic features.

**Geometry** is another word that denotes a geographic feature. Originally the word **geometry** meant measurement of the earth. Another meaning comes from cartography, referring to the geometric features that cartographers use to map the world.

This chapter uses all of these terms synonymously: **geographic feature**, **geospatial feature**, **feature**, or **geometry**. Here, the term most commonly used is **geometry**, defined as *a point or an aggregate of points representing anything in the world that has a location*.

## 12.18.2. The OpenGIS Geometry Model

The set of geometry types proposed by OGC's **SQL with Geometry Types** environment is based on the **OpenGIS Geometry Model**. In this model, each geometric object has the following general properties:

- It is associated with a Spatial Reference System, which describes the coordinate space in which the object is defined.
- It belongs to some geometry class.

### 12.18.2.1. The Geometry Class Hierarchy

The geometry classes define a hierarchy as follows:

- **Geometry** (noninstantiable)
  - **Point** (instantiable)
  - **Curve** (noninstantiable)
    - **LineString** (instantiable)
      - **Line**
      - **LinearRing**
  - **Surface** (noninstantiable)
  - **Polygon** (instantiable)



- `GeometryCollection` (instantiable)
  - `MultiPoint` (instantiable)
  - `MultiCurve` (noninstantiable)
    - `MultiLineString` (instantiable)
  - `MultiSurface` (noninstantiable)
    - `MultiPolygon` (instantiable)

It is not possible to create objects in noninstantiable classes. It is possible to create objects in instantiable classes. All classes have properties, and instantiable classes may also have assertions (rules that define valid class instances).

`Geometry` is the base class. It is an abstract class. The instantiable subclasses of `Geometry` are restricted to zero-, one-, and two-dimensional geometric objects that exist in two-dimensional coordinate space. All instantiable geometry classes are defined so that valid instances of a geometry class are topologically closed (that is, all defined geometries include their boundary).

The base `Geometry` class has subclasses for `Point`, `Curve`, `Surface`, and `GeometryCollection`:

- `Point` represents zero-dimensional objects.
- `Curve` represents one-dimensional objects, and has subclass `LineString`, with sub-subclasses `Line` and `LinearRing`.
- `Surface` is designed for two-dimensional objects and has subclass `Polygon`.
- `GeometryCollection` has specialized zero-, one-, and two-dimensional collection classes named `MultiPoint`, `MultiLineString`, and `MultiPolygon` for modeling geometries corresponding to collections of `Points`, `LineStrings`, and `Polygons`, respectively. `MultiCurve` and `MultiSurface` are introduced as abstract superclasses that generalize the collection interfaces to handle `Curves` and `Surfaces`.

`Geometry`, `Curve`, `Surface`, `MultiCurve`, and `MultiSurface` are defined as noninstantiable classes. They define a common set of methods for their subclasses and are included for extensibility.

`Point`, `LineString`, `Polygon`, `GeometryCollection`, `MultiPoint`, `MultiLineString`, and `MultiPolygon` are instantiable classes.

### 12.18.2.2. Class `Geometry`

`Geometry` is the root class of the hierarchy. It is a noninstantiable class but has a number of properties that are common to all geometry values created from any of the `Geometry` subclasses. These properties are described in the following list. Particular subclasses have their own specific properties, described later.

#### Geometry Properties

A geometry value has the following properties:

- Its **type**. Each geometry belongs to one of the instantiable classes in the hierarchy.
- Its **SRID**, or Spatial Reference Identifier. This value identifies the geometry's associated Spatial Reference System that describes the coordinate space in which the geometry object is defined.

In MySQL, the SRID value is just an integer associated with the geometry value. All calculations are done assuming Euclidean (planar) geometry.

- Its **coordinates** in its Spatial Reference System, represented as double-precision (8-byte) numbers. All nonempty geometries include at least one pair of (X,Y) coordinates. Empty geometries contain no coordinates.

Coordinates are related to the SRID. For example, in different coordinate systems, the distance between two objects may differ even when objects have the same coordinates, because the distance on the **planar** coordinate system and the distance on the **geocentric** system (coordinates on the Earth's surface) are different things.

- Its **interior**, **boundary**, and **exterior**.

Every geometry occupies some position in space. The exterior of a geometry is all space not occupied by the geometry. The interior is the space occupied by the geometry. The boundary is the interface between the geometry's interior and exterior.

- Its **MBR** (Minimum Bounding Rectangle), or Envelope. This is the bounding geometry, formed by the minimum and maximum (X,Y) coordinates:

```
((MINX MINY, MAXX MINY, MAXX MAXY, MINX MAXY, MINX MINY))
```

- Whether the value is **simple** or **nonsimple**. Geometry values of types ([LineString](#), [MultiPoint](#), [MultiLineString](#)) are either simple or nonsimple. Each type determines its own assertions for being simple or nonsimple.
- Whether the value is **closed** or **not closed**. Geometry values of types ([LineString](#), [MultiString](#)) are either closed or not closed. Each type determines its own assertions for being closed or not closed.
- Whether the value is **empty** or **nonempty** A geometry is empty if it does not have any points. Exterior, interior, and boundary of an empty geometry are not defined (that is, they are represented by a [NULL](#) value). An empty geometry is defined to be always simple and has an area of 0.
- Its **dimension**. A geometry can have a dimension of -1, 0, 1, or 2:
  - -1 for an empty geometry.
  - 0 for a geometry with no length and no area.
  - 1 for a geometry with nonzero length and zero area.
  - 2 for a geometry with nonzero area.

[Point](#) objects have a dimension of zero. [LineString](#) objects have a dimension of 1. [Polygon](#) objects have a dimension of 2. The dimensions of [MultiPoint](#), [MultiLineString](#), and [MultiPolygon](#) objects are the same as the dimensions of the elements they consist of.

### 12.18.2.3. Class [Point](#)

A [Point](#) is a geometry that represents a single location in coordinate space.

#### [Point](#) Examples

- Imagine a large-scale map of the world with many cities. A [Point](#) object could represent each city.
- On a city map, a [Point](#) object could represent a bus stop.

#### [Point](#) Properties

- X-coordinate value.
- Y-coordinate value.
- [Point](#) is defined as a zero-dimensional geometry.

- The boundary of a [Point](#) is the empty set.

#### 12.18.2.4. Class [Curve](#)

A [Curve](#) is a one-dimensional geometry, usually represented by a sequence of points. Particular subclasses of [Curve](#) define the type of interpolation between points. [Curve](#) is a noninstantiable class.

##### [Curve](#) Properties

- A [Curve](#) has the coordinates of its points.
- A [Curve](#) is defined as a one-dimensional geometry.
- A [Curve](#) is simple if it does not pass through the same point twice.
- A [Curve](#) is closed if its start point is equal to its endpoint.
- The boundary of a closed [Curve](#) is empty.
- The boundary of a nonclosed [Curve](#) consists of its two endpoints.
- A [Curve](#) that is simple and closed is a [LinearRing](#).

#### 12.18.2.5. Class [LineString](#)

A [LineString](#) is a [Curve](#) with linear interpolation between points.

##### [LineString](#) Examples

- On a world map, [LineString](#) objects could represent rivers.
- In a city map, [LineString](#) objects could represent streets.

##### [LineString](#) Properties

- A [LineString](#) has coordinates of segments, defined by each consecutive pair of points.
- A [LineString](#) is a [Line](#) if it consists of exactly two points.
- A [LineString](#) is a [LinearRing](#) if it is both closed and simple.

#### 12.18.2.6. Class [Surface](#)

A [Surface](#) is a two-dimensional geometry. It is a noninstantiable class. Its only instantiable subclass is [Polygon](#).

##### [Surface](#) Properties

- A [Surface](#) is defined as a two-dimensional geometry.
- The OpenGIS specification defines a simple [Surface](#) as a geometry that consists of a single “patch” that is associated with a single exterior boundary and zero or more interior boundaries.
- The boundary of a simple [Surface](#) is the set of closed curves corresponding to its exterior and interior boundaries.

#### 12.18.2.7. Class [Polygon](#)

A [Polygon](#) is a planar [Surface](#) representing a multisided geometry. It is defined by a single exterior boundary and zero or more interior boundaries, where each interior boundary defines a hole in the [Polygon](#).

### Polygon Examples

- On a region map, `Polygon` objects could represent forests, districts, and so on.

### Polygon Assertions

- The boundary of a `Polygon` consists of a set of `LinearRing` objects (that is, `LineString` objects that are both simple and closed) that make up its exterior and interior boundaries.
- A `Polygon` has no rings that cross. The rings in the boundary of a `Polygon` may intersect at a `Point`, but only as a tangent.
- A `Polygon` has no lines, spikes, or punctures.
- A `Polygon` has an interior that is a connected point set.
- A `Polygon` may have holes. The exterior of a `Polygon` with holes is not connected. Each hole defines a connected component of the exterior.

The preceding assertions make a `Polygon` a simple geometry.

## 12.18.2.8. Class `GeometryCollection`

A `GeometryCollection` is a geometry that is a collection of one or more geometries of any class.

All the elements in a `GeometryCollection` must be in the same Spatial Reference System (that is, in the same coordinate system). There are no other constraints on the elements of a `GeometryCollection`, although the subclasses of `GeometryCollection` described in the following sections may restrict membership. Restrictions may be based on:

- Element type (for example, a `MultiPoint` may contain only `Point` elements)
- Dimension
- Constraints on the degree of spatial overlap between elements

## 12.18.2.9. Class `MultiPoint`

A `MultiPoint` is a geometry collection composed of `Point` elements. The points are not connected or ordered in any way.

### `MultiPoint` Examples

- On a world map, a `MultiPoint` could represent a chain of small islands.
- On a city map, a `MultiPoint` could represent the outlets for a ticket office.

### `MultiPoint` Properties

- A `MultiPoint` is a zero-dimensional geometry.
- A `MultiPoint` is simple if no two of its `Point` values are equal (have identical coordinate values).
- The boundary of a `MultiPoint` is the empty set.

## 12.18.2.10. Class `MultiCurve`

A `MultiCurve` is a geometry collection composed of `Curve` elements. `MultiCurve` is a noninstantiable class.

### `MultiCurve` Properties

- A `MultiCurve` is a one-dimensional geometry.

- A `MultiCurve` is simple if and only if all of its elements are simple; the only intersections between any two elements occur at points that are on the boundaries of both elements.
- A `MultiCurve` boundary is obtained by applying the “mod 2 union rule” (also known as the “odd-even rule”): A point is in the boundary of a `MultiCurve` if it is in the boundaries of an odd number of `MultiCurve` elements.
- A `MultiCurve` is closed if all of its elements are closed.
- The boundary of a closed `MultiCurve` is always empty.

#### 12.18.2.11. Class `MultiLineString`

A `MultiLineString` is a `MultiCurve` geometry collection composed of `LineString` elements.

##### `MultiLineString` Examples

- On a region map, a `MultiLineString` could represent a river system or a highway system.

#### 12.18.2.12. Class `MultiSurface`

A `MultiSurface` is a geometry collection composed of surface elements. `MultiSurface` is a noninstantiable class. Its only instantiable subclass is `MultiPolygon`.

##### `MultiSurface` Assertions

- Two `MultiSurface` surfaces have no interiors that intersect.
- Two `MultiSurface` elements have boundaries that intersect at most at a finite number of points.

#### 12.18.2.13. Class `MultiPolygon`

A `MultiPolygon` is a `MultiSurface` object composed of `Polygon` elements.

##### `MultiPolygon` Examples

- On a region map, a `MultiPolygon` could represent a system of lakes.

##### `MultiPolygon` Assertions

- A `MultiPolygon` has no two `Polygon` elements with interiors that intersect.
- A `MultiPolygon` has no two `Polygon` elements that cross (crossing is also forbidden by the previous assertion), or that touch at an infinite number of points.
- A `MultiPolygon` may not have cut lines, spikes, or punctures. A `MultiPolygon` is a regular, closed point set.
- A `MultiPolygon` that has more than one `Polygon` has an interior that is not connected. The number of connected components of the interior of a `MultiPolygon` is equal to the number of `Polygon` values in the `MultiPolygon`.

##### `MultiPolygon` Properties

- A `MultiPolygon` is a two-dimensional geometry.
- A `MultiPolygon` boundary is a set of closed curves (`LineString` values) corresponding to the boundaries of its `Polygon` elements.
- Each `Curve` in the boundary of the `MultiPolygon` is in the boundary of exactly one `Polygon` element.
- Every `Curve` in the boundary of an `Polygon` element is in the boundary of the `MultiPolygon`.

### 12.18.3. Supported Spatial Data Formats

This section describes the standard spatial data formats that are used to represent geometry objects in queries. They are:

- Well-Known Text (WKT) format
- Well-Known Binary (WKB) format

Internally, MySQL stores geometry values in a format that is not identical to either WKT or WKB format.

#### 12.18.3.1. Well-Known Text (WKT) Format

The Well-Known Text (WKT) representation of Geometry is designed to exchange geometry data in ASCII form. For a Backus-Naur grammar that specifies the formal production rules for writing WKT values, see the OpenGIS specification document referenced in [Section 12.18, “Spatial Extensions”](#).

Examples of WKT representations of geometry objects:

- A [Point](#):

```
POINT(15 20)
```

Note that point coordinates are specified with no separating comma. This differs from the syntax for the SQL [POINT\(\)](#) [1300] function, which requires a comma between the coordinates. Take care to use the syntax appropriate to the context of a given spatial operation. For example, the following statements both extract the X-coordinate from a [Point](#) object. The first produces the object directly using the [POINT\(\)](#) [1300] function. The second uses a WKT representation converted to a [Point](#) with [GeomFromText\(\)](#).

```
mysql> SELECT X(POINT(15, 20));
+-----+
| X(POINT(15, 20)) |
+-----+
|                15 |
+-----+

mysql> SELECT X(GeomFromText('POINT(15 20)'));
+-----+
| X(GeomFromText('POINT(15 20)')) |
+-----+
|                15 |
+-----+
```

- A [LineString](#) with four points:

```
LINESTRING(0 0, 10 10, 20 25, 50 60)
```

Note that point coordinate pairs are separated by commas.

- A [Polygon](#) with one exterior ring and one interior ring:

```
POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))
```

- A [MultiPoint](#) with three [Point](#) values:

```
MULTIPOINT(0 0, 20 20, 60 60)
```

- A [MultiLineString](#) with two [LineString](#) values:

```
MULTILINESTRING((10 10, 20 20), (15 15, 30 15))
```

- A [MultiPolygon](#) with two [Polygon](#) values:

```
MULTIPOLYGON(((0 0,10 0,10 10,0 10,0 0)),((5 5,7 5,7 7,5 7, 5 5)))
```

- A `GeometryCollection` consisting of two `Point` values and one `LineString`:

```
GEOMETRYCOLLECTION(POINT(10 10), POINT(30 30), LINESTRING(15 15, 20 20))
```

### 12.18.3.2. Well-Known Binary (WKB) Format

The Well-Known Binary (WKB) representation for geometric values is defined by the OpenGIS specification. It is also defined in the ISO *SQL/MM Part 3: Spatial* standard.

WKB is used to exchange geometry data as binary streams represented by `BLOB` values containing geometric WKB information.

WKB uses 1-byte unsigned integers, 4-byte unsigned integers, and 8-byte double-precision numbers (IEEE 754 format). A byte is eight bits.

For example, a WKB value that corresponds to `POINT(1 1)` consists of this sequence of 21 bytes (each represented here by two hex digits):

```
010100000000000000000000F03F000000000000F03F
```

The sequence may be broken down into these components:

```
Byte order : 01
WKB type   : 01000000
X          : 00000000000000F03F
Y          : 00000000000000F03F
```

Component representation is as follows:

- The byte order may be either 1 or 0 to indicate little-endian or big-endian storage. The little-endian and big-endian byte orders are also known as Network Data Representation (NDR) and External Data Representation (XDR), respectively.
- The WKB type is a code that indicates the geometry type. Values from 1 through 7 indicate `Point`, `LineString`, `Polygon`, `MultiPoint`, `MultiLineString`, `MultiPolygon`, and `GeometryCollection`.
- A `Point` value has X and Y coordinates, each represented as a double-precision value.

WKB values for more complex geometry values are represented by more complex data structures, as detailed in the OpenGIS specification.

## 12.18.4. Creating a Spatially Enabled MySQL Database

This section describes the data types you can use for representing spatial data in MySQL, and the functions available for creating and retrieving spatial values.

### 12.18.4.1. MySQL Spatial Data Types

MySQL has data types that correspond to OpenGIS classes. Some of these types hold single geometry values:

- `GEOMETRY`
- `POINT`
- `LINestring`
- `POLYGON`

`GEOMETRY` can store geometry values of any type. The other single-value types (`POINT`, `LINestring`, and `POLYGON`) restrict their values to a particular geometry type.

The other data types hold collections of values:

- `MULTIPOINT`
- `MULTILINESTRING`
- `MULTIPOLYGON`
- `GEOMETRYCOLLECTION`

`GEOMETRYCOLLECTION` can store a collection of objects of any type. The other collection types (`MULTIPOINT`, `MULTILINESTRING`, `MULTIPOLYGON`, and `GEOMETRYCOLLECTION`) restrict collection members to those having a particular geometry type.

## 12.18.4.2. Creating Spatial Values

This section describes how to create spatial values using Well-Known Text and Well-Known Binary functions that are defined in the OpenGIS standard, and using MySQL-specific functions.

### 12.18.4.2.1. Creating Geometry Values Using WKT Functions

MySQL provides a number of functions that take as arguments a Well-Known Text representation and, optionally, a spatial reference system identifier (SRID). They return the corresponding geometry.

`GeomFromText()` [1298] accepts a WKT of any geometry type as its first argument. An implementation also provides type-specific construction functions for construction of geometry values of each geometry type.

- `GeomCollFromText(wkt[,srid])` [1298],  
`GeometryCollectionFromText(wkt[,srid])` [1298]

Constructs a `GEOMETRYCOLLECTION` value using its WKT representation and SRID.

- `GeomFromText(wkt[,srid])` [1298], `GeometryFromText(wkt[,srid])` [1298]

Constructs a geometry value of any type using its WKT representation and SRID.

- `LineFromText(wkt[,srid])` [1298], `LineStringFromText(wkt[,srid])` [1298]

Constructs a `LINestring` value using its WKT representation and SRID.

- `MLineFromText(wkt[,srid])` [1298], `MultiLineStringFromText(wkt[,srid])` [1298]

Constructs a `MULTILINESTRING` value using its WKT representation and SRID.

- `MPointFromText(wkt[,srid])` [1298], `MultiPointFromText(wkt[,srid])` [1298]

Constructs a `MULTIPOINT` value using its WKT representation and SRID.

- `MPolyFromText(wkt[,srid])` [1298], `MultiPolygonFromText(wkt[,srid])` [1298]

Constructs a `MULTIPOLYGON` value using its WKT representation and SRID.

- `PointFromText(wkt[,srid])` [1298]

Constructs a `POINT` value using its WKT representation and SRID.

- `PolyFromText(wkt[,srid])` [1298], `PolygonFromText(wkt[,srid])` [1298]

Constructs a `POLYGON` value using its WKT representation and SRID.

The OpenGIS specification also defines the following optional functions, which MySQL does not implement. These functions construct `Polygon` or `MultiPolygon` values based on the WKT representation of a collection of rings or closed `LineString` values. These values may intersect.



- `BdMPolyFromText(wkt, srid)` [1299]

Constructs a `MultiPolygon` value from a `MultiLineString` value in WKT format containing an arbitrary collection of closed `LineString` values.

- `BdPolyFromText(wkt, srid)` [1299]

Constructs a `Polygon` value from a `MultiLineString` value in WKT format containing an arbitrary collection of closed `LineString` values.

#### 12.18.4.2.2. Creating Geometry Values Using WKB Functions

MySQL provides a number of functions that take as arguments a `BLOB` containing a Well-Known Binary representation and, optionally, a spatial reference system identifier (SRID). They return the corresponding geometry.

These functions also accept geometry objects for compatibility with the return value of the functions in Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”. Thus, those functions may be used to provide the first argument to the functions in this section.

- `GeomCollFromWKB(wkb[, srid])` [1299],  
`GeometryCollectionFromWKB(wkb[, srid])` [1299]

Constructs a `GEOMETRYCOLLECTION` value using its WKB representation and SRID.

- `GeomFromWKB(wkb[, srid])` [1299], `GeometryFromWKB(wkb[, srid])` [1299]

Constructs a geometry value of any type using its WKB representation and SRID.

- `LineFromWKB(wkb[, srid])` [1299], `LineStringFromWKB(wkb[, srid])` [1299]

Constructs a `LINESTRING` value using its WKB representation and SRID.

- `MLineFromWKB(wkb[, srid])` [1299], `MultiLineStringFromWKB(wkb[, srid])` [1299]

Constructs a `MULTILINESTRING` value using its WKB representation and SRID.

- `MPointFromWKB(wkb[, srid])` [1299], `MultiPointFromWKB(wkb[, srid])` [1299]

Constructs a `MULTIPOINT` value using its WKB representation and SRID.

- `MPolyFromWKB(wkb[, srid])` [1299], `MultiPolygonFromWKB(wkb[, srid])` [1299]

Constructs a `MULTIPOLYGON` value using its WKB representation and SRID.

- `PointFromWKB(wkb[, srid])` [1299]

Constructs a `POINT` value using its WKB representation and SRID.

- `PolyFromWKB(wkb[, srid])` [1299], `PolygonFromWKB(wkb[, srid])` [1299]

Constructs a `POLYGON` value using its WKB representation and SRID.

The OpenGIS specification also describes optional functions for constructing `Polygon` or `MultiPolygon` values based on the WKB representation of a collection of rings or closed `LineString` values. These values may intersect. MySQL does not implement these functions:

- `BdMPolyFromWKB(wkb, srid)` [1299]

Constructs a `MultiPolygon` value from a `MultiLineString` value in WKB format containing an arbitrary collection of closed `LineString` values.

- `BdPolyFromWKB(wkb, srid)` [1299]

Constructs a [Polygon](#) value from a [MultiLineString](#) value in WKB format containing an arbitrary collection of closed [LineString](#) values.

### 12.18.4.2.3. Creating Geometry Values Using MySQL-Specific Functions

MySQL provides a set of useful nonstandard functions for creating geometry values. The functions described in this section are MySQL extensions to the OpenGIS specification.

These functions produce geometry objects from either WKB values or geometry objects as arguments. If any argument is not a proper WKB or geometry representation of the proper object type, the return value is `NULL`.

For example, you can insert the geometry return value from `Point()` [1300] directly into a `Point` column:

```
INSERT INTO t1 (pt_col) VALUES(Point(1,2));
```

- [GeometryCollection\(\*g1,g2,...\*\)](#) [1300]

Constructs a [GeometryCollection](#).

- [LineString\(\*pt1,pt2,...\*\)](#) [1300]

Constructs a [LineString](#) value from a number of [Point](#) or WKB [Point](#) arguments. If the number of arguments is less than two, the return value is `NULL`.

- [MultiLineString\(\*ls1,ls2,...\*\)](#) [1300]

Constructs a [MultiLineString](#) value using [LineString](#) or WKB [LineString](#) arguments.

- [MultiPoint\(\*pt1,pt2,...\*\)](#) [1300]

Constructs a [MultiPoint](#) value using [Point](#) or WKB [Point](#) arguments.

- [MultiPolygon\(\*poly1,poly2,...\*\)](#) [1300]

Constructs a [MultiPolygon](#) value from a set of [Polygon](#) or WKB [Polygon](#) arguments.

- [Point\(\*x,y\*\)](#) [1300]

Constructs a [Point](#) using its coordinates.

- [Polygon\(\*ls1,ls2,...\*\)](#) [1300]

Constructs a [Polygon](#) value from a number of [LineString](#) or WKB [LineString](#) arguments. If any argument does not represent a [LinearRing](#) (that is, not a closed and simple [LineString](#)), the return value is `NULL`.

### 12.18.4.3. Creating Spatial Columns

MySQL provides a standard way of creating spatial columns for geometry types, for example, with `CREATE TABLE` or `ALTER TABLE`. Currently, spatial columns are supported for [MyISAM](#), [InnoDB](#), [NDB](#), and [ARCHIVE](#) tables. See also the annotations about spatial indexes under [Section 12.18.6.1, “Creating Spatial Indexes”](#).

- Use the `CREATE TABLE` statement to create a table with a spatial column:

```
CREATE TABLE geom (g GEOMETRY);
```

- Use the `ALTER TABLE` statement to add or drop a spatial column to or from an existing table:

```
ALTER TABLE geom ADD pt POINT;
ALTER TABLE geom DROP pt;
```

### 12.18.4.4. Populating Spatial Columns

After you have created spatial columns, you can populate them with spatial data.

Values should be stored in internal geometry format, but you can convert them to that format from either Well-Known Text (WKT) or Well-Known Binary (WKB) format. The following examples demonstrate how to insert geometry values into a table by converting WKT values into internal geometry format:

- Perform the conversion directly in the `INSERT` statement:

```
INSERT INTO geom VALUES (GeomFromText('POINT(1 1)'));

SET @g = 'POINT(1 1)';
INSERT INTO geom VALUES (GeomFromText(@g));
```

- Perform the conversion prior to the `INSERT`:

```
SET @g = GeomFromText('POINT(1 1)');
INSERT INTO geom VALUES (@g);
```

The following examples insert more complex geometries into the table:

```
SET @g = 'LINESTRING(0 0,1 1,2 2)';
INSERT INTO geom VALUES (GeomFromText(@g));

SET @g = 'POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))';
INSERT INTO geom VALUES (GeomFromText(@g));

SET @g =
'GEOMETRYCOLLECTION(POINT(1 1),LINESTRING(0 0,1 1,2 2,3 3,4 4))';
INSERT INTO geom VALUES (GeomFromText(@g));
```

The preceding examples all use `GeomFromText()` [1298] to create geometry values. You can also use type-specific functions:

```
SET @g = 'POINT(1 1)';
INSERT INTO geom VALUES (PointFromText(@g));

SET @g = 'LINESTRING(0 0,1 1,2 2)';
INSERT INTO geom VALUES (LineStringFromText(@g));

SET @g = 'POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))';
INSERT INTO geom VALUES (PolygonFromText(@g));

SET @g =
'GEOMETRYCOLLECTION(POINT(1 1),LINESTRING(0 0,1 1,2 2,3 3,4 4))';
INSERT INTO geom VALUES (GeomCollFromText(@g));
```

Note that if a client application program wants to use WKB representations of geometry values, it is responsible for sending correctly formed WKB in queries to the server. However, there are several ways of satisfying this requirement. For example:

- Inserting a `POINT(1 1)` value with hex literal syntax:

```
mysql> INSERT INTO geom VALUES
-> (GeomFromWKB(0x01010000000000000000000000F03F000000000000F03F));
```

- An ODBC application can send a WKB representation, binding it to a placeholder using an argument of `BLOB` type:

```
INSERT INTO geom VALUES (GeomFromWKB(?))
```

Other programming interfaces may support a similar placeholder mechanism.

- In a C program, you can escape a binary value using `mysql_real_escape_string()` and include the result in a query string that is sent to the server. See [Section 21.8.6.54](#), “`mysql_real_escape_string()`”.

### 12.18.4.5. Fetching Spatial Data

Geometry values stored in a table can be fetched in internal format. You can also convert them into WKT or WKB format.

- Fetching spatial data in internal format:

Fetching geometry values using internal format can be useful in table-to-table transfers:

```
CREATE TABLE geom2 (g GEOMETRY) SELECT g FROM geom;
```

- Fetching spatial data in WKT format:

The `AsText()` [1302] function converts a geometry from internal format into a WKT string.

```
SELECT AsText(g) FROM geom;
```

- Fetching spatial data in WKB format:

The `AsBinary()` [1302] function converts a geometry from internal format into a `BLOB` containing the WKB value.

```
SELECT AsBinary(g) FROM geom;
```

## 12.18.5. Spatial Analysis Functions

After populating spatial columns with values, you are ready to query and analyze them. MySQL provides a set of functions to perform various operations on spatial data. These functions can be grouped into four major categories according to the type of operation they perform:

- Functions that convert geometries between various formats
- Functions that provide access to qualitative or quantitative properties of a geometry
- Functions that describe relations between two geometries
- Functions that create new geometries from existing ones

Spatial analysis functions can be used in many contexts, such as:

- Any interactive SQL program, such as `mysql`.
- Application programs written in any language that supports a MySQL client API

### 12.18.5.1. Geometry Format Conversion Functions

MySQL supports the following functions for converting geometry values between internal format and either WKT or WKB format:

- `AsBinary(g)` [1302], `AsWKB(g)` [1302]

Converts a value in internal geometry format to its WKB representation and returns the binary result.

```
SELECT AsBinary(g) FROM geom;
```

- `AsText(g)` [1302], `AsWKT(g)` [1302]

Converts a value in internal geometry format to its WKT representation and returns the string result.

```
mysql> SET @g = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(GeomFromText(@g));
+-----+
| AsText(GeomFromText(@g)) |
+-----+
| LINESTRING(1 1,2 2,3 3) |
```

- `GeomFromText(wkt[, srid])` [1298]

Converts a string value from its WKT representation into internal geometry format and returns the result. A number of type-specific functions are also supported, such as `PointFromText()` [1298] and `LineFromText()` [1298]. See Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”.

- `GeomFromWKB(wkb[, srid])` [1299]

Converts a binary value from its WKB representation into internal geometry format and returns the result. A number of type-specific functions are also supported, such as `PointFromWKB()` [1299] and `LineFromWKB()` [1299]. See Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”.

## 12.18.5.2. Geometry Property Functions

Each function that belongs to this group takes a geometry value as its argument and returns some quantitative or qualitative property of the geometry. Some functions restrict their argument type. Such functions return `NULL` if the argument is of an incorrect geometry type. For example, `Area()` [1306] returns `NULL` if the object type is neither `Polygon` nor `MultiPolygon`.

### 12.18.5.2.1. General Geometry Functions

The functions listed in this section do not restrict their argument and accept a geometry value of any type.

- `Dimension(g)` [1303]

Returns the inherent dimension of the geometry value *g*. The result can be `-1`, `0`, `1`, or `2`. The meaning of these values is given in Section 12.18.2.2, “Class Geometry”.

```
mysql> SELECT Dimension(GeomFromText('LineString(1 1,2 2)'));
+-----+
| Dimension(GeomFromText('LineString(1 1,2 2)')) |
+-----+
| 1 |
+-----+
```

- `Envelope(g)` [1303]

Returns the Minimum Bounding Rectangle (MBR) for the geometry value *g*. The result is returned as a `Polygon` value.

The polygon is defined by the corner points of the bounding box:

```
POLYGON((MINX MINY, MAXX MINY, MAXX MAXY, MINX MAXY, MINX MINY))
```

```
mysql> SELECT AsText(Envelope(GeomFromText('LineString(1 1,2 2)')));
+-----+
| AsText(Envelope(GeomFromText('LineString(1 1,2 2)'))) |
+-----+
| POLYGON((1 1,2 1,2 2,1 2,1 1)) |
+-----+
```

- `GeometryType(g)` [1303]

Returns as a binary string the name of the geometry type of which the geometry instance *g* is a member. The name corresponds to one of the instantiable `Geometry` subclasses.

```
mysql> SELECT GeometryType(GeomFromText('POINT(1 1)'));
+-----+
| GeometryType(GeomFromText('POINT(1 1)')) |
+-----+
```

POINT
-------

- [IsSimple\(\*g\*\) \[1304\]](#)

Returns 1 if the geometry value *g* has no anomalous geometric points, such as self-intersection or self-tangency. [IsSimple\(\)](#) [1304] returns 0 if the argument is not simple, and `NULL` if it is `NULL`.

The description of each instantiable geometric class given earlier in the chapter includes the specific conditions that cause an instance of that class to be classified as not simple. (See [Section 12.18.2.1, “The Geometry Class Hierarchy”](#).)

Prior to MySQL 5.6.1, this function always returns 0.

- [SRID\(\*g\*\) \[1304\]](#)

Returns an integer indicating the Spatial Reference System ID for the geometry value *g*.

In MySQL, the SRID value is just an integer associated with the geometry value. All calculations are done assuming Euclidean (planar) geometry.

```
mysql> SELECT SRID(GeomFromText('LineString(1 1,2 2)',101));
+-----+
| SRID(GeomFromText('LineString(1 1,2 2)',101)) |
+-----+
| 101 |
+-----+
```

The OpenGIS specification also defines the following functions, which MySQL does not implement:

- [Boundary\(\*g\*\) \[1304\]](#)

Returns a geometry that is the closure of the combinatorial boundary of the geometry value *g*.

- [IsEmpty\(\*g\*\) \[1304\]](#)

This function is a placeholder that returns 0 for any valid geometry value, 1 for any invalid geometry value or `NULL`.

MySQL does not support GIS `EMPTY` values such as `POINT EMPTY`.

### 12.18.5.2.2. Point Functions

A `Point` consists of X and Y coordinates, which may be obtained using the following functions:

- [X\(\*p\*\) \[1304\]](#)

Returns the X-coordinate value for the `Point` object *p* as a double-precision number.

```
mysql> SELECT X(POINT(56.7, 53.34));
+-----+
| X(POINT(56.7, 53.34)) |
+-----+
| 56.7 |
+-----+
```

- [Y\(\*p\*\) \[1304\]](#)

Returns the Y-coordinate value for the `Point` object *p* as a double-precision number.

```
mysql> SELECT Y(POINT(56.7, 53.34));
+-----+
| Y(POINT(56.7, 53.34)) |
+-----+
| 53.34 |
+-----+
```

### 12.18.5.2.3. LineString Functions

A `LineString` consists of `Point` values. You can extract particular points of a `LineString`, count the number of points that it contains, or obtain its length.

- `EndPoint(ls)` [1305]

Returns the `Point` that is the endpoint of the `LineString` value `ls`.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(EndPoint(GeomFromText(@ls)));
+-----+
| AsText(EndPoint(GeomFromText(@ls))) |
+-----+
| POINT(3 3) |
+-----+
```

- `GLength(ls)` [1305]

Returns as a double-precision number the length of the `LineString` value `ls` in its associated spatial reference.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT GLength(GeomFromText(@ls));
+-----+
| GLength(GeomFromText(@ls)) |
+-----+
| 2.8284271247462 |
+-----+
```

`GLength()` [1305] is a nonstandard name. It corresponds to the OpenGIS `Length()` [1181] function.

- `NumPoints(ls)` [1305]

Returns the number of `Point` objects in the `LineString` value `ls`.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT NumPoints(GeomFromText(@ls));
+-----+
| NumPoints(GeomFromText(@ls)) |
+-----+
| 3 |
+-----+
```

- `PointN(ls,N)` [1305]

Returns the `N`-th `Point` in the `LineString` value `ls`. Points are numbered beginning with 1.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(PointN(GeomFromText(@ls),2));
+-----+
| AsText(PointN(GeomFromText(@ls),2)) |
+-----+
| POINT(2 2) |
+-----+
```

- `StartPoint(ls)` [1305]

Returns the `Point` that is the start point of the `LineString` value `ls`.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(StartPoint(GeomFromText(@ls)));
+-----+
| AsText(StartPoint(GeomFromText(@ls))) |
+-----+
| POINT(1 1) |
+-----+
```

The OpenGIS specification also defines the following function, which MySQL does not implement:

- `IsRing(ls)` [1306]

Returns 1 if the `LineString` value *ls* is closed (that is, its `StartPoint()` [1305] and `EndPoint()` [1305] values are the same) and is simple (does not pass through the same point more than once). Returns 0 if *ls* is not a ring, and `-1` if it is `NULL`.

#### 12.18.5.2.4. MultiLineString Functions

These functions return properties of `MultiLineString` values.

- `GLength(mIs)` [1305]

Returns as a double-precision number the length of the `MultiLineString` value *mIs*. The length of *mIs* is equal to the sum of the lengths of its elements.

```
mysql> SET @mIs = 'MultiLineString((1 1,2 2,3 3),(4 4,5 5))';
mysql> SELECT GLength(GeomFromText(@mIs));
+-----+
| GLength(GeomFromText(@mIs)) |
+-----+
|                4.2426406871193 |
+-----+
```

`GLength()` [1305] is a nonstandard name. It corresponds to the OpenGIS `Length()` [1181] function.

- `IsClosed(mIs)` [1306]

Returns 1 if the `MultiLineString` value *mIs* is closed (that is, the `StartPoint()` [1305] and `EndPoint()` [1305] values are the same for each `LineString` in *mIs*). Returns 0 if *mIs* is not closed, and `-1` if it is `NULL`.

```
mysql> SET @mIs = 'MultiLineString((1 1,2 2,3 3),(4 4,5 5))';
mysql> SELECT IsClosed(GeomFromText(@mIs));
+-----+
| IsClosed(GeomFromText(@mIs)) |
+-----+
|                                0 |
+-----+
```

#### 12.18.5.2.5. Polygon Functions

These functions return properties of `Polygon` values.

- `Area(poly)` [1306]

Returns as a double-precision number the area of the `Polygon` value *poly*, as measured in its spatial reference system.

```
mysql> SET @poly = 'Polygon((0 0,0 3,3 0,0 0),(1 1,1 2,2 1,1 1))';
mysql> SELECT Area(GeomFromText(@poly));
+-----+
| Area(GeomFromText(@poly)) |
+-----+
|                4 |
+-----+
```

- `ExteriorRing(poly)` [1306]

Returns the exterior ring of the `Polygon` value *poly* as a `LineString`.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT AsText(ExteriorRing(GeomFromText(@poly)));
+-----+
```



```
| AsText(ExteriorRing(GeomFromText(@poly))) |
+-----+
| LINESTRING(0 0,0 3,3 3,3 0,0 0) |
+-----+
```

- [InteriorRingN\(\*poly\*,\*N\*\)](#) [1307]

Returns the *N*-th interior ring for the [Polygon](#) value *poly* as a [LineString](#). Rings are numbered beginning with 1.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT AsText(InteriorRingN(GeomFromText(@poly),1));
+-----+
| AsText(InteriorRingN(GeomFromText(@poly),1)) |
+-----+
| LINESTRING(1 1,1 2,2 2,2 1,1 1) |
+-----+
```

- [NumInteriorRings\(\*poly\*\)](#) [1307]

Returns the number of interior rings in the [Polygon](#) value *poly*.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT NumInteriorRings(GeomFromText(@poly));
+-----+
| NumInteriorRings(GeomFromText(@poly)) |
+-----+
| 1 |
+-----+
```

### 12.18.5.2.6. MultiPolygon Functions

These functions return properties of [MultiPolygon](#) values.

- [Area\(\*mpoly\*\)](#) [1306]

Returns as a double-precision number the area of the [MultiPolygon](#) value *mpoly*, as measured in its spatial reference system.

```
mysql> SET @mpoly =
-> 'MultiPolygon(((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1)))';
mysql> SELECT Area(GeomFromText(@mpoly));
+-----+
| Area(GeomFromText(@mpoly)) |
+-----+
| 8 |
+-----+
```

- [Centroid\(\*mpoly\*\)](#) [1307]

Returns the mathematical centroid for the [MultiPolygon](#) value *mpoly* as a [Point](#). The result is not guaranteed to be on the [MultiPolygon](#).

```
mysql> SET @poly =
-> GeomFromText('POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7,5 5))');
mysql> SELECT GeometryType(@poly),AsText(Centroid(@poly));
+-----+
| GeometryType(@poly) | AsText(Centroid(@poly)) |
+-----+
| POLYGON | POINT(4.958333333333333 4.958333333333333) |
+-----+
```

The OpenGIS specification also defines the following function, which MySQL does not implement:

- [PointOnSurface\(\*mpoly\*\)](#)

Returns a [Point](#) value that is guaranteed to be on the [MultiPolygon](#) value *mpoly*.

### 12.18.5.2.7. GeometryCollection Functions

These functions return properties of `GeometryCollection` values.

- `GeometryN(gc,N)` [1308]

Returns the *N*-th geometry in the `GeometryCollection` value *gc*. Geometries are numbered beginning with 1.

```
mysql> SET @gc = 'GeometryCollection(Point(1 1),LineString(2 2, 3 3))';
mysql> SELECT AsText(GeometryN(GeomFromText(@gc),1));
+-----+
| AsText(GeometryN(GeomFromText(@gc),1)) |
+-----+
| POINT(1 1) |
+-----+
```

- `NumGeometries(gc)` [1308]

Returns the number of geometries in the `GeometryCollection` value *gc*.

```
mysql> SET @gc = 'GeometryCollection(Point(1 1),LineString(2 2, 3 3))';
mysql> SELECT NumGeometries(GeomFromText(@gc));
+-----+
| NumGeometries(GeomFromText(@gc)) |
+-----+
| 2 |
+-----+
```

### 12.18.5.3. Functions That Create New Geometries from Existing Ones

The following sections describe functions that take geometry values as arguments and return new geometry values.

#### 12.18.5.3.1. Geometry Functions That Produce New Geometries

Section 12.18.5.2, “[Geometry Property Functions](#)”, discusses several functions that construct new geometries from existing ones. See that section for descriptions of these functions:

- `Envelope(g)` [1303]
- `StartPoint(ls)` [1305]
- `EndPoint(ls)` [1305]
- `PointN(ls,N)` [1305]
- `ExteriorRing(poly)` [1306]
- `InteriorRingN(poly,N)` [1307]
- `GeometryN(gc,N)` [1308]

#### 12.18.5.3.2. Spatial Operators

OpenGIS proposes a number of other functions that can produce geometries. They are designed to implement spatial operators.

- `Buffer(g,d)` [1308]

Returns a geometry that represents all points whose distance from the geometry value *g* is less than or equal to a distance of *d*.

`Buffer()` [1308] supports negative distances for polygons, multipolygons, and geometry collections containing polygons or multipolygons. For point, multipoint, linestring, multilinestring,

and geometry collections not containing any polygons or multipolygons, `Buffer()` [1308] with a negative distance returns `NULL`.

Prior to MySQL 5.6.1, this function is unimplemented.

The OpenGIS specification also defines the following functions, which MySQL does not implement:

- `ConvexHull(g)` [1309]

Returns a geometry that represents the convex hull of the geometry value *g*.

- `Difference(g1,g2)` [1309]

Returns a geometry that represents the point set difference of the geometry value *g1* with *g2*.

- `Intersection(g1,g2)` [1309]

Returns a geometry that represents the point set intersection of the geometry values *g1* with *g2*.

- `SymDifference(g1,g2)` [1309]

Returns a geometry that represents the point set symmetric difference of the geometry value *g1* with *g2*.

- `Union(g1,g2)` [1309]

Returns a geometry that represents the point set union of the geometry values *g1* and *g2*.

#### 12.18.5.4. Functions for Testing Spatial Relations Between Geometric Objects

The functions described in these sections take two geometries as input parameters and return a qualitative or quantitative relation between them.

##### 12.18.5.4.1. Relations on Geometry Minimal Bounding Rectangles (MBRs)

MySQL provides several functions that test relations between minimal bounding rectangles of two geometries *g1* and *g2*. The return values 1 and 0 indicate true and false, respectively.

- `MBRContains(g1,g2)` [1309]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangle of *g1* contains the Minimum Bounding Rectangle of *g2*. This tests the opposite relationship as `MBRWithin()` [1310].

```
mysql> SET @g1 = GeomFromText('Polygon((0 0,0 3,3 3,3 0,0 0))');
mysql> SET @g2 = GeomFromText('Point(1 1)');
mysql> SELECT MBRContains(@g1,@g2), MBRContains(@g2,@g1);
+-----+-----+
| MBRContains(@g1,@g2) | MBRContains(@g2,@g1) |
+-----+-----+
| 1 | 0 |
+-----+-----+
```

- `MBRDisjoint(g1,g2)` [1309]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* are disjoint (do not intersect).

- `MBREqual(g1,g2)` [1309]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* are the same.

- `MBRIntersects(g1,g2)` [1309]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* intersect.

- `MBROverlaps(g1,g2)` [1310]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* overlap. The term *spatially overlaps* is used if two geometries intersect and their intersection results in a geometry of the same dimension but not equal to either of the given geometries.

- `MBRTouches(g1,g2)` [1310]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* touch. Two geometries *spatially touch* if the interiors of the geometries do not intersect, but the boundary of one of the geometries intersects either the boundary or the interior of the other.

- `MBRWithin(g1,g2)` [1310]

Returns 1 or 0 to indicate whether the Minimum Bounding Rectangle of *g1* is within the Minimum Bounding Rectangle of *g2*. This tests the opposite relationship as `MBRContains()` [1309].

```
mysql> SET @g1 = GeomFromText('Polygon((0 0,0 3,3 3,3 0,0 0))');
mysql> SET @g2 = GeomFromText('Polygon((0 0,0 5,5 5,5 0,0 0))');
mysql> SELECT MBRWithin(@g1,@g2), MBRWithin(@g2,@g1);
+-----+-----+
| MBRWithin(@g1,@g2) | MBRWithin(@g2,@g1) |
+-----+-----+
| 1 | 0 |
+-----+-----+
```

#### 12.18.5.4.2. Functions That Test Spatial Relationships Between Geometries

The OpenGIS specification defines the following functions. They test the relationship between two geometry values *g1* and *g2*.

The return values 1 and 0 indicate true and false, respectively.



#### Note

MySQL originally implemented these functions such that they used object bounding rectangles and returned the same result as the corresponding MBR-based functions. As of MySQL 5.6.1, corresponding versions are available that use precise object shapes. These versions are named with an `ST_` prefix. For example, `Contains()` [1311] uses object bounding rectangles, whereas `ST_Contains()` [1310] uses object shapes.

As of MySQL 5.6.1, there are also `ST_` aliases for existing spatial functions that were already exact. For example, `ST_IsEmpty()` is an alias for `IsEmpty()` [1304].

#### Functions That Use Object Shapes

- `ST_Contains(g1,g2)` [1310]

Returns 1 or 0 to indicate whether *g1* completely contains *g2*. This tests the opposite relationship as `ST_Within()` [1311].

- `ST_Crosses(g1,g2)` [1310]

Returns 1 if *g1* spatially crosses *g2*. Returns `NULL` if *g1* is a `Polygon` or a `MultiPolygon`, or if *g2* is a `Point` or a `MultiPoint`. Otherwise, returns 0.

The term *spatially crosses* denotes a spatial relation between two given geometries that has the following properties:

- The two geometries intersect
- Their intersection results in a geometry that has a dimension that is one less than the maximum dimension of the two given geometries
- Their intersection is not equal to either of the two given geometries
- `ST_Disjoint(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* is spatially disjoint from (does not intersect) *g2*.
- `ST_Equals(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* is spatially equal to *g2*.
- `ST_Intersects(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* spatially intersects *g2*.
- `ST_Overlaps(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* spatially overlaps *g2*. The term *spatially overlaps* is used if two geometries intersect and their intersection results in a geometry of the same dimension but not equal to either of the given geometries.
- `ST_Touches(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* spatially touches *g2*. Two geometries *spatially touch* if the interiors of the geometries do not intersect, but the boundary of one of the geometries intersects either the boundary or the interior of the other.
- `ST_Within(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* is spatially within *g2*. This tests the opposite relationship as `ST_Contains()` [1310].

### Functions That Use Bounding Rectangles

- `Contains(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* completely contains *g2*. This tests the opposite relationship as `Within()` [1312].
- `Crosses(g1,g2)` [1311]  
Returns 1 if *g1* spatially crosses *g2*. Returns `NULL` if *g1* is a `Polygon` or a `MultiPolygon`, or if *g2* is a `Point` or a `MultiPoint`. Otherwise, returns 0.  
  
The term *spatially crosses* denotes a spatial relation between two given geometries that has the following properties:
  - The two geometries intersect
  - Their intersection results in a geometry that has a dimension that is one less than the maximum dimension of the two given geometries
  - Their intersection is not equal to either of the two given geometries
- `Disjoint(g1,g2)` [1311]  
Returns 1 or 0 to indicate whether *g1* is spatially disjoint from (does not intersect) *g2*.
- `Equals(g1,g2)` [1311]

Returns 1 or 0 to indicate whether *g1* is spatially equal to *g2*.

- `Intersects(g1,g2)` [1312]

Returns 1 or 0 to indicate whether *g1* spatially intersects *g2*.

- `Overlaps(g1,g2)` [1312]

Returns 1 or 0 to indicate whether *g1* spatially overlaps *g2*. The term *spatially overlaps* is used if two geometries intersect and their intersection results in a geometry of the same dimension but not equal to either of the given geometries.

- `Touches(g1,g2)` [1312]

Returns 1 or 0 to indicate whether *g1* spatially touches *g2*. Two geometries *spatially touch* if the interiors of the geometries do not intersect, but the boundary of one of the geometries intersects either the boundary or the interior of the other.

- `Within(g1,g2)` [1312]

Returns 1 or 0 to indicate whether *g1* is spatially within *g2*. This tests the opposite relationship as `Contains()` [1311].

## 12.18.6. Optimizing Spatial Analysis

For `MyISAM` tables, Search operations in nonspatial databases can be optimized using `SPATIAL` indexes. This is true for spatial databases as well. With the help of a great variety of multi-dimensional indexing methods that have previously been designed, it is possible to optimize spatial searches. The most typical of these are:

- Point queries that search for all objects that contain a given point
- Region queries that search for all objects that overlap a given region

MySQL uses **R-Trees with quadratic splitting** for `SPATIAL` indexes on spatial columns. A `SPATIAL` index is built using the MBR of a geometry. For most geometries, the MBR is a minimum rectangle that surrounds the geometries. For a horizontal or a vertical linestring, the MBR is a rectangle degenerated into the linestring. For a point, the MBR is a rectangle degenerated into the point.

It is also possible to create normal indexes on spatial columns. In a non-`SPATIAL` index, you must declare a prefix for any spatial column except for `POINT` columns.

`MyISAM` supports both `SPATIAL` and non-`SPATIAL` indexes. Other storage engines support non-`SPATIAL` indexes, as described in [Section 13.1.11, “CREATE INDEX Syntax”](#).

### 12.18.6.1. Creating Spatial Indexes

For `MyISAM` tables, MySQL can create spatial indexes using syntax similar to that for creating regular indexes, but extended with the `SPATIAL` keyword. Currently, columns in spatial indexes must be declared `NOT NULL`. The following examples demonstrate how to create spatial indexes:

- With `CREATE TABLE`:

```
CREATE TABLE geom (g GEOMETRY NOT NULL, SPATIAL INDEX(g)) ENGINE=MyISAM;
```

- With `ALTER TABLE`:

```
ALTER TABLE geom ADD SPATIAL INDEX(g);
```

- With `CREATE INDEX`:

```
CREATE SPATIAL INDEX sp_index ON geom (g);
```

For `MyISAM` tables, `SPATIAL INDEX` creates an R-tree index. For storage engines that support nonspatial indexing of spatial columns, the engine creates a B-tree index. A B-tree index on spatial values will be useful for exact-value lookups, but not for range scans.

For more information on indexing spatial columns, see [Section 13.1.11](#), “`CREATE INDEX Syntax`”.

To drop spatial indexes, use `ALTER TABLE` or `DROP INDEX`:

- With `ALTER TABLE`:

```
ALTER TABLE geom DROP INDEX g;
```

- With `DROP INDEX`:

```
DROP INDEX sp_index ON geom;
```

Example: Suppose that a table `geom` contains more than 32,000 geometries, which are stored in the column `g` of type `GEOMETRY`. The table also has an `AUTO_INCREMENT` column `fid` for storing object ID values.

```
mysql> DESCRIBE geom;
+-----+-----+-----+-----+-----+-----+
| Field | Type      | Null | Key | Default | Extra          |
+-----+-----+-----+-----+-----+-----+
| fid   | int(11)   |      | PRI | NULL    | auto_increment |
| g     | geometry  |      |     |         |                |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT COUNT(*) FROM geom;
+-----+
| count(*) |
+-----+
|      32376 |
+-----+
1 row in set (0.00 sec)
```

To add a spatial index on the column `g`, use this statement:

```
mysql> ALTER TABLE geom ADD SPATIAL INDEX(g);
Query OK, 32376 rows affected (4.05 sec)
Records: 32376 Duplicates: 0 Warnings: 0
```

### 12.18.6.2. Using a Spatial Index

The optimizer investigates whether available spatial indexes can be involved in the search for queries that use a function such as `MBRContains()` [1309] or `MBRWithin()` [1310] in the `WHERE` clause. The following query finds all objects that are in the given rectangle:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,
           31000 15000,
           31000 16000,
           30000 16000,
           30000 15000));'
mysql> SELECT fid,AsText(g) FROM geom WHERE
-> MBRContains(GeomFromText(@poly),g);
+-----+-----+-----+-----+-----+-----+
| fid | AsText(g)
+-----+-----+-----+-----+-----+-----+
| 21 | LINESTRING(30350.4 15828.8,30350.6 15845,30333.8 15845,30
| 22 | LINESTRING(30350.6 15871.4,30350.6 15887.8,30334 15887.8, ...
| 23 | LINESTRING(30350.6 15914.2,30350.6 15930.4,30334 15930.4, ...
| 24 | LINESTRING(30290.2 15823,30290.2 15839.4,30273.4 15839.4, ...
| 25 | LINESTRING(30291.4 15866.2,30291.6 15882.4,30274.8 15882. ...
| 26 | LINESTRING(30291.6 15918.2,30291.6 15934.4,30275 15934.4, ...
| 249 | LINESTRING(30337.8 15938.6,30337.8 15946.8,30320.4 15946. ...
| 1 | LINESTRING(30250.4 15129.2,30248.8 15138.4,30238.2 15136. ...
```

```
| 2 | LINESTRING(30220.2 15122.8,30217.2 15137.8,30207.6 15136, ... |
| 3 | LINESTRING(30179 15114.4,30176.6 15129.4,30167 15128,3016 ... |
| 4 | LINESTRING(30155.2 15121.4,30140.4 15118.6,30142 15109,30 ... |
| 5 | LINESTRING(30192.4 15085,30177.6 15082.2,30179.2 15072.4, ... |
| 6 | LINESTRING(30244 15087,30229 15086.2,30229.4 15076.4,3024 ... |
| 7 | LINESTRING(30200.6 15059.4,30185.6 15058.6,30186 15048.8, ... |
| 10 | LINESTRING(30179.6 15017.8,30181 15002.8,30190.8 15003.6, ... |
| 11 | LINESTRING(30154.2 15000.4,30168.6 15004.8,30166 15014.2, ... |
| 13 | LINESTRING(30105 15065.8,30108.4 15050.8,30118 15053,3011 ... |
| 154 | LINESTRING(30276.2 15143.8,30261.4 15141,30263 15131.4,30 ... |
| 155 | LINESTRING(30269.8 15084,30269.4 15093.4,30258.6 15093,30 ... |
| 157 | LINESTRING(30128.2 15011,30113.2 15010.2,30113.6 15000.4, ... |
+-----+
20 rows in set (0.00 sec)
```

Use `EXPLAIN` to check the way this query is executed:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,
           31000 15000,
           31000 16000,
           30000 16000,
           30000 15000))';
mysql> EXPLAIN SELECT fid,AsText(g) FROM geom WHERE
-> MBRContains(GeomFromText(@poly),g)\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: geom
         type: range
possible_keys: g
          key: g
        key_len: 32
          ref: NULL
         rows: 50
      Extra: Using where
1 row in set (0.00 sec)
```

Check what would happen without a spatial index:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,
           31000 15000,
           31000 16000,
           30000 16000,
           30000 15000))';
mysql> EXPLAIN SELECT fid,AsText(g) FROM g IGNORE INDEX (g) WHERE
-> MBRContains(GeomFromText(@poly),g)\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: geom
         type: ALL
possible_keys: NULL
          key: NULL
        key_len: NULL
          ref: NULL
         rows: 32376
      Extra: Using where
1 row in set (0.00 sec)
```

Executing the `SELECT` statement without the spatial index yields the same result but causes the execution time to rise from 0.00 seconds to 0.46 seconds:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,
           31000 15000,
           31000 16000,
           30000 16000,
           30000 15000))';
mysql> SELECT fid,AsText(g) FROM geom IGNORE INDEX (g) WHERE
```



```

-> MBRContains(GeomFromText(@poly),g);
+-----+
| fid | AsText(g)
+-----+
| 1 | LINESTRING(30250.4 15129.2,30248.8 15138.4,30238.2 15136. ...
| 2 | LINESTRING(30220.2 15122.8,30217.2 15137.8,30207.6 15136, ...
| 3 | LINESTRING(30179 15114.4,30176.6 15129.4,30167 15128,3016 ...
| 4 | LINESTRING(30155.2 15121.4,30140.4 15118.6,30142 15109,30 ...
| 5 | LINESTRING(30192.4 15085,30177.6 15082.2,30179.2 15072.4, ...
| 6 | LINESTRING(30244 15087,30229 15086.2,30229.4 15076.4,3024 ...
| 7 | LINESTRING(30200.6 15059.4,30185.6 15058.6,30186 15048.8, ...
| 10 | LINESTRING(30179.6 15017.8,30181 15002.8,30190.8 15003.6, ...
| 11 | LINESTRING(30154.2 15000.4,30168.6 15004.8,30166 15014.2, ...
| 13 | LINESTRING(30105 15065.8,30108.4 15050.8,30118 15053,3011 ...
| 21 | LINESTRING(30350.4 15828.8,30350.6 15845,30333.8 15845,30 ...
| 22 | LINESTRING(30350.6 15871.4,30350.6 15887.8,30334 15887.8, ...
| 23 | LINESTRING(30350.6 15914.2,30350.6 15930.4,30334 15930.4, ...
| 24 | LINESTRING(30290.2 15823,30290.2 15839.4,30273.4 15839.4, ...
| 25 | LINESTRING(30291.4 15866.2,30291.6 15882.4,30274.8 15882. ...
| 26 | LINESTRING(30291.6 15918.2,30291.6 15934.4,30275 15934.4, ...
| 154 | LINESTRING(30276.2 15143.8,30261.4 15141,30263 15131.4,30 ...
| 155 | LINESTRING(30269.8 15084,30269.4 15093.4,30258.6 15093,30 ...
| 157 | LINESTRING(30128.2 15011,30113.2 15010.2,30113.6 15000.4, ...
| 249 | LINESTRING(30337.8 15938.6,30337.8 15946.8,30320.4 15946. ...
+-----+
20 rows in set (0.46 sec)

```

## 12.18.7. MySQL Conformance and Compatibility

MySQL does not yet implement the following GIS features:

- Additional Metadata Views

OpenGIS specifications propose several additional metadata views. For example, a system view named `GEOMETRY_COLUMNS` contains a description of geometry columns, one row for each geometry column in the database.

- The OpenGIS function `Length()` [1181] on `LineString` and `MultiLineString` currently should be called in MySQL as `GLength()` [1305]

The problem is that there is an existing SQL function `Length()` [1181] that calculates the length of string values, and sometimes it is not possible to distinguish whether the function is called in a textual or spatial context. We need either to solve this somehow, or decide on another function name.

## 12.19. Precision Math

MySQL 5.6 provides support for precision math: numeric value handling that results in extremely accurate results and a high degree control over invalid values. Precision math is based on these two features:

- SQL modes that control how strict the server is about accepting or rejecting invalid data.
- The MySQL library for fixed-point arithmetic.

These features have several implications for numeric operations and provide a high degree of compliance with standard SQL:

- **Precise calculations:** For exact-value numbers, calculations do not introduce floating-point errors. Instead, exact precision is used. For example, MySQL treats a number such as `.0001` as an exact value rather than as an approximation, and summing it 10,000 times produces a result of exactly `1`, not a value that is merely “close” to 1.
- **Well-defined rounding behavior:** For exact-value numbers, the result of `ROUND()` [1207] depends on its argument, not on environmental factors such as how the underlying C library works.

- **Platform independence:** Operations on exact numeric values are the same across different platforms such as Windows and Unix.
- **Control over handling of invalid values:** Overflow and division by zero are detectable and can be treated as errors. For example, you can treat a value that is too large for a column as an error rather than having the value truncated to lie within the range of the column's data type. Similarly, you can treat division by zero as an error rather than as an operation that produces a result of `NULL`. The choice of which approach to take is determined by the setting of the server SQL mode.

The following discussion covers several aspects of how precision math works, including possible incompatibilities with older applications. At the end, some examples are given that demonstrate how MySQL 5.6 handles numeric operations precisely. For information about controlling the SQL mode, see [Section 5.1.7, "Server SQL Modes"](#).

### 12.19.1. Types of Numeric Values

The scope of precision math for exact-value operations includes the exact-value data types (`DECIMAL` and integer types) and exact-value numeric literals. Approximate-value data types and numeric literals are handled as floating-point numbers.

Exact-value numeric literals have an integer part or fractional part, or both. They may be signed. Examples: `1`, `.2`, `3.4`, `-5`, `-6.78`, `+9.10`.

Approximate-value numeric literals are represented in scientific notation with a mantissa and exponent. Either or both parts may be signed. Examples: `1.2E3`, `1.2E-3`, `-1.2E3`, `-1.2E-3`.

Two numbers that look similar may be treated differently. For example, `2.34` is an exact-value (fixed-point) number, whereas `2.34E0` is an approximate-value (floating-point) number.

The `DECIMAL` data type is a fixed-point type and calculations are exact. In MySQL, the `DECIMAL` type has several synonyms: `NUMERIC`, `DEC`, `FIXED`. The integer types also are exact-value types.

The `FLOAT` and `DOUBLE` data types are floating-point types and calculations are approximate. In MySQL, types that are synonymous with `FLOAT` or `DOUBLE` are `DOUBLE PRECISION` and `REAL`.

### 12.19.2. `DECIMAL` Data Type Changes

This section discusses the characteristics of the `DECIMAL` data type (and its synonyms) in MySQL 5.6, with particular regard to the following topics:

- Maximum number of digits
- Storage format
- Storage requirements
- The nonstandard MySQL extension to the upper range of `DECIMAL` columns

Possible incompatibilities with applications that are written for older versions of MySQL (prior to 5.0.3) are noted throughout this section.

The declaration syntax for a `DECIMAL` column is `DECIMAL(M,D)`. The ranges of values for the arguments in MySQL 5.6 are as follows:

- `M` is the maximum number of digits (the precision). It has a range of 1 to 65. (Older versions of MySQL permitted a range of 1 to 254.)
- `D` is the number of digits to the right of the decimal point (the scale). It has a range of 0 to 30 and must be no larger than `M`.

The maximum value of 65 for `M` means that calculations on `DECIMAL` values are accurate up to 65 digits. This limit of 65 digits of precision also applies to exact-value numeric literals, so the maximum

range of such literals differs from before. (In older versions of MySQL, decimal values could have up to 254 digits. However, calculations were done using floating-point and thus were approximate, not exact.)

Values for `DECIMAL` columns in MySQL 5.6 are stored using a binary format that packs nine decimal digits into 4 bytes. The storage requirements for the integer and fractional parts of each value are determined separately. Each multiple of nine digits requires 4 bytes, and any remaining digits left over require some fraction of 4 bytes. The storage required for remaining digits is given by the following table.

Leftover Digits	Number of Bytes
0	0
1–2	1
3–4	2
5–6	3
7–9	4

For example, a `DECIMAL(18,9)` column has nine digits on either side of the decimal point, so the integer part and the fractional part each require 4 bytes. A `DECIMAL(20,6)` column has fourteen integer digits and six fractional digits. The integer digits require four bytes for nine of the digits and 3 bytes for the remaining five digits. The six fractional digits require 3 bytes.

Unlike some older versions of MySQL, `DECIMAL` columns in MySQL 5.6 do not store a leading `+` character or `-` character or leading `0` digits. If you insert `+0003.1` into a `DECIMAL(5,1)` column, it is stored as `3.1`. For negative numbers, a literal `-` character is not stored. Applications that rely on the older behavior must be modified to account for this change.

`DECIMAL` columns in MySQL 5.6 do not permit values larger than the range implied by the column definition. For example, a `DECIMAL(3,0)` column supports a range of `-999` to `999`. A `DECIMAL(M,D)` column permits at most  $M - D$  digits to the left of the decimal point. This is not compatible with applications relying on older versions of MySQL that permitted storing an extra digit in lieu of a `+` sign.

The SQL standard requires that the precision of `NUMERIC(M,D)` be *exactly*  $M$  digits. For `DECIMAL(M,D)`, the standard requires a precision of at least  $M$  digits but permits more. In MySQL, `DECIMAL(M,D)` and `NUMERIC(M,D)` are the same, and both have a precision of exactly  $M$  digits.

For more detailed information about porting applications that rely on the old treatment of the `DECIMAL` data type, see the *MySQL 5.0 Reference Manual*.

### 12.19.3. Expression Handling

With precision math, exact-value numbers are used as given whenever possible. For example, numbers in comparisons are used exactly as given without a change in value. In strict SQL mode, for `INSERT` into a column with an exact data type (`DECIMAL` or integer), a number is inserted with its exact value if it is within the column range. When retrieved, the value should be the same as what was inserted. (Without strict mode, truncation for `INSERT` is permissible.)

Handling of a numeric expression depends on what kind of values the expression contains:

- If any approximate values are present, the expression is approximate and is evaluated using floating-point arithmetic.
- If no approximate values are present, the expression contains only exact values. If any exact value contains a fractional part (a value following the decimal point), the expression is evaluated using `DECIMAL` exact arithmetic and has a precision of 65 digits. The term “exact” is subject to the limits of what can be represented in binary. For example, `1.0/3.0` can be approximated in decimal notation as `.333...`, but not written as an exact number, so `(1.0/3.0)*3.0` does not evaluate to exactly `1.0`.

- Otherwise, the expression contains only integer values. The expression is exact and is evaluated using integer arithmetic and has a precision the same as `BIGINT` (64 bits).

If a numeric expression contains any strings, they are converted to double-precision floating-point values and the expression is approximate.

Inserts into numeric columns are affected by the SQL mode, which is controlled by the `sql_mode` [545] system variable. (See Section 5.1.7, “Server SQL Modes”.) The following discussion mentions strict mode (selected by the `STRICT_ALL_TABLES` [607] or `STRICT_TRANS_TABLES` [607] mode values) and `ERROR_FOR_DIVISION_BY_ZERO` [603]. To turn on all restrictions, you can simply use `TRADITIONAL` [609] mode, which includes both strict mode values and `ERROR_FOR_DIVISION_BY_ZERO` [603]:

```
mysql> SET sql_mode='TRADITIONAL';
```

If a number is inserted into an exact type column (`DECIMAL` or integer), it is inserted with its exact value if it is within the column range.

If the value has too many digits in the fractional part, rounding occurs and a warning is generated. Rounding is done as described in Section 12.19.4, “Rounding Behavior”.

If the value has too many digits in the integer part, it is too large and is handled as follows:

- If strict mode is not enabled, the value is truncated to the nearest legal value and a warning is generated.
- If strict mode is enabled, an overflow error occurs.

Underflow is not detected, so underflow handling is undefined.

By default, division by zero produces a result of `NULL` and no warning. With the `ERROR_FOR_DIVISION_BY_ZERO` [603] SQL mode enabled, MySQL handles division by zero differently:

- If strict mode is not enabled, a warning occurs.
- If strict mode is enabled, inserts and updates involving division by zero are prohibited, and an error occurs.

In other words, inserts and updates involving expressions that perform division by zero can be treated as errors, but this requires `ERROR_FOR_DIVISION_BY_ZERO` [603] in addition to strict mode.

Suppose that we have this statement:

```
INSERT INTO t SET i = 1/0;
```

This is what happens for combinations of strict and `ERROR_FOR_DIVISION_BY_ZERO` [603] modes.

<code>sql_mode</code> [545] Value	Result
' ' (Default)	No warning, no error; <code>i</code> is set to <code>NULL</code> .
strict	No warning, no error; <code>i</code> is set to <code>NULL</code> .
<code>ERROR_FOR_DIVISION_BY_ZERO</code> [603]	Warning, no error; <code>i</code> is set to <code>NULL</code> .
strict, <code>ERROR_FOR_DIVISION_BY_ZERO</code> [603]	Error condition; no row is inserted.

For inserts of strings into numeric columns, conversion from string to number is handled as follows if the string has nonnumeric contents:

- A string that does not begin with a number cannot be used as a number and produces an error in strict mode, or a warning otherwise. This includes the empty string.

- A string that begins with a number can be converted, but the trailing nonnumeric portion is truncated. If the truncated portion contains anything other than spaces, this produces an error in strict mode, or a warning otherwise.

## 12.19.4. Rounding Behavior

This section discusses precision math rounding for the `ROUND()` [1207] function and for inserts into columns with exact-value types (`DECIMAL` and integer).

The `ROUND()` [1207] function rounds differently depending on whether its argument is exact or approximate:

- For exact-value numbers, `ROUND()` [1207] uses the “round half up” rule: A value with a fractional part of .5 or greater is rounded up to the next integer if positive or down to the next integer if negative. (In other words, it is rounded away from zero.) A value with a fractional part less than .5 is rounded down to the next integer if positive or up to the next integer if negative.
- For approximate-value numbers, the result depends on the C library. On many systems, this means that `ROUND()` [1207] uses the “round to nearest even” rule: A value with any fractional part is rounded to the nearest even integer.

The following example shows how rounding differs for exact and approximate values:

```
mysql> SELECT ROUND(2.5), ROUND(25E-1);
+-----+-----+
| ROUND(2.5) | ROUND(25E-1) |
+-----+-----+
| 3          | 2            |
+-----+-----+
```

For inserts into a `DECIMAL` or integer column, the target is an exact data type, so rounding uses “round half away from zero,” regardless of whether the value to be inserted is exact or approximate:

```
mysql> CREATE TABLE t (d DECIMAL(10,0));
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t VALUES(2.5),(2.5E0);
Query OK, 2 rows affected, 2 warnings (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 2

mysql> SELECT d FROM t;
+-----+
| d    |
+-----+
| 3    |
| 3    |
+-----+
```

## 12.19.5. Precision Math Examples

This section provides some examples that show precision math query results in MySQL 5.6. These examples demonstrate the principles described in [Section 12.19.3, “Expression Handling”](#), and [Section 12.19.4, “Rounding Behavior”](#).

**Example 1.** Numbers are used with their exact value as given when possible:

```
mysql> SELECT (.1 + .2) = .3;
+-----+
| (.1 + .2) = .3 |
+-----+
| 1              |
+-----+
```

For floating-point values, results are inexact:

```
mysql> SELECT (.1E0 + .2E0) = .3E0;
```

```

+-----+
| (.1E0 + .2E0) = .3E0 |
+-----+
|                          0 |
+-----+

```

Another way to see the difference in exact and approximate value handling is to add a small number to a sum many times. Consider the following stored procedure, which adds `.0001` to a variable 1,000 times.

```

CREATE PROCEDURE p ()
BEGIN
  DECLARE i INT DEFAULT 0;
  DECLARE d DECIMAL(10,4) DEFAULT 0;
  DECLARE f FLOAT DEFAULT 0;
  WHILE i < 10000 DO
    SET d = d + .0001;
    SET f = f + .0001E0;
    SET i = i + 1;
  END WHILE;
  SELECT d, f;
END;

```

The sum for both `d` and `f` logically should be 1, but that is true only for the decimal calculation. The floating-point calculation introduces small errors:

```

+-----+
| d      | f      |
+-----+
| 1.0000 | 0.999999999999991 |
+-----+

```

**Example 2.** Multiplication is performed with the scale required by standard SQL. That is, for two numbers `X1` and `X2` that have scale `S1` and `S2`, the scale of the result is `S1 + S2`:

```

mysql> SELECT .01 * .01;
+-----+
| .01 * .01 |
+-----+
| 0.0001    |
+-----+

```

**Example 3.** Rounding behavior for exact-value numbers is well-defined:

Rounding behavior (for example, with the `ROUND()` [1207] function) is independent of the implementation of the underlying C library, which means that results are consistent from platform to platform.

- Rounding for exact-value columns (`DECIMAL` and integer) and exact-valued numbers uses the “round half away from zero” rule. Values with a fractional part of `.5` or greater are rounded away from zero to the nearest integer, as shown here:

```

mysql> SELECT ROUND(2.5), ROUND(-2.5);
+-----+
| ROUND(2.5) | ROUND(-2.5) |
+-----+
| 3          | -3          |
+-----+

```

- Rounding for floating-point values uses the C library, which on many systems uses the “round to nearest even” rule. Values with any fractional part on such systems are rounded to the nearest even integer:

```

mysql> SELECT ROUND(2.5E0), ROUND(-2.5E0);
+-----+
| ROUND(2.5E0) | ROUND(-2.5E0) |
+-----+
| 2            | -2            |
+-----+

```

**Example 4.** In strict mode, inserting a value that is out of range for a column causes an error, rather than truncation to a legal value.

When MySQL is not running in strict mode, truncation to a legal value occurs:

```
mysql> SET sql_mode='';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t (i TINYINT);
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO t SET i = 128;
Query OK, 1 row affected, 1 warning (0.00 sec)

mysql> SELECT i FROM t;
+-----+
| i     |
+-----+
| 127  |
+-----+
1 row in set (0.00 sec)
```

However, an error occurs if strict mode is in effect:

```
mysql> SET sql_mode='STRICT_ALL_TABLES';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t (i TINYINT);
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t SET i = 128;
ERROR 1264 (22003): Out of range value adjusted for column 'i' at row 1

mysql> SELECT i FROM t;
Empty set (0.00 sec)
```

**Example 5:** In strict mode and with [ERROR\\_FOR\\_DIVISION\\_BY\\_ZERO \[603\]](#) set, division by zero causes an error, not a result of `NULL`.

In nonstrict mode, division by zero has a result of `NULL`:

```
mysql> SET sql_mode='';
Query OK, 0 rows affected (0.01 sec)

mysql> CREATE TABLE t (i TINYINT);
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t SET i = 1 / 0;
Query OK, 1 row affected (0.00 sec)

mysql> SELECT i FROM t;
+-----+
| i     |
+-----+
| NULL  |
+-----+
1 row in set (0.03 sec)
```

However, division by zero is an error if the proper SQL modes are in effect:

```
mysql> SET sql_mode='STRICT_ALL_TABLES,ERROR_FOR_DIVISION_BY_ZERO';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t (i TINYINT);
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t SET i = 1 / 0;
ERROR 1365 (22012): Division by 0
```

```
mysql> SELECT i FROM t;
Empty set (0.01 sec)
```

**Example 6.** Exact-value literals are evaluated as exact values.

Prior to MySQL 5.0.3, exact-value and approximate-value literals both are evaluated as double-precision floating-point values:

```
mysql> SELECT VERSION();
+-----+
| VERSION() |
+-----+
| 4.1.18-log |
+-----+
1 row in set (0.01 sec)

mysql> CREATE TABLE t SELECT 2.5 AS a, 25E-1 AS b;
Query OK, 1 row affected (0.07 sec)
Records: 1 Duplicates: 0 Warnings: 0

mysql> DESCRIBE t;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| a     | double(3,1)  |      |     | 0.0     |      |
| b     | double       |      |     | 0       |      |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.04 sec)
```

As of MySQL 5.0.3, the approximate-value literal is evaluated using floating point, but the exact-value literal is handled as **DECIMAL**:

```
mysql> SELECT VERSION();
+-----+
| VERSION() |
+-----+
| 5.1.6-alpha-log |
+-----+
1 row in set (0.11 sec)

mysql> CREATE TABLE t SELECT 2.5 AS a, 25E-1 AS b;
Query OK, 1 row affected (0.01 sec)
Records: 1 Duplicates: 0 Warnings: 0

mysql> DESCRIBE t;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| a     | decimal(2,1) unsigned | NO   |     | 0.0     |      |
| b     | double              | NO   |     | 0       |      |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.01 sec)
```

**Example 7.** If the argument to an aggregate function is an exact numeric type, the result is also an exact numeric type, with a scale at least that of the argument.

Consider these statements:

```
mysql> CREATE TABLE t (i INT, d DECIMAL, f FLOAT);
mysql> INSERT INTO t VALUES(1,1,1);
mysql> CREATE TABLE y SELECT AVG(i), AVG(d), AVG(f) FROM t;
```

Before MySQL 5.0.3, the result is a double no matter the argument type:

```
mysql> DESCRIBE y;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| AVG(i) | double(17,4) | YES  |     | NULL    |      |
+-----+-----+-----+-----+-----+-----+
```



AVG(d)	double(17,4)	YES		NULL	
AVG(f)	double	YES		NULL	

As of MySQL 5.0.3, the result is a double only for the floating-point argument. For exact type arguments, the result is also an exact type:

```
mysql> DESCRIBE y;
```

Field	Type	Null	Key	Default	Extra
AVG(i)	decimal(14,4)	YES		NULL	
AVG(d)	decimal(14,4)	YES		NULL	
AVG(f)	double	YES		NULL	

The result is a double only for the floating-point argument. For exact type arguments, the result is also an exact type.



---

# Chapter 13. SQL Statement Syntax

## Table of Contents

13.1. Data Definition Statements .....	1326
13.1.1. ALTER DATABASE Syntax .....	1326
13.1.2. ALTER EVENT Syntax .....	1327
13.1.3. ALTER FUNCTION Syntax .....	1329
13.1.4. ALTER PROCEDURE Syntax .....	1329
13.1.5. ALTER SERVER Syntax .....	1329
13.1.6. ALTER TABLE Syntax .....	1330
13.1.7. ALTER VIEW Syntax .....	1343
13.1.8. CREATE DATABASE Syntax .....	1344
13.1.9. CREATE EVENT Syntax .....	1344
13.1.10. CREATE FUNCTION Syntax .....	1349
13.1.11. CREATE INDEX Syntax .....	1349
13.1.12. CREATE PROCEDURE and CREATE FUNCTION Syntax .....	1352
13.1.13. CREATE SERVER Syntax .....	1357
13.1.14. CREATE TABLE Syntax .....	1358
13.1.15. CREATE TRIGGER Syntax .....	1384
13.1.16. CREATE VIEW Syntax .....	1388
13.1.17. DROP DATABASE Syntax .....	1391
13.1.18. DROP EVENT Syntax .....	1392
13.1.19. DROP FUNCTION Syntax .....	1392
13.1.20. DROP INDEX Syntax .....	1392
13.1.21. DROP PROCEDURE and DROP FUNCTION Syntax .....	1393
13.1.22. DROP SERVER Syntax .....	1393
13.1.23. DROP TABLE Syntax .....	1393
13.1.24. DROP TRIGGER Syntax .....	1394
13.1.25. DROP VIEW Syntax .....	1394
13.1.26. RENAME TABLE Syntax .....	1394
13.1.27. TRUNCATE TABLE Syntax .....	1395
13.2. Data Manipulation Statements .....	1396
13.2.1. CALL Syntax .....	1396
13.2.2. DELETE Syntax .....	1398
13.2.3. DO Syntax .....	1402
13.2.4. HANDLER Syntax .....	1402
13.2.5. INSERT Syntax .....	1403
13.2.6. LOAD DATA INFILE Syntax .....	1412
13.2.7. LOAD XML Syntax .....	1420
13.2.8. REPLACE Syntax .....	1426
13.2.9. SELECT Syntax .....	1427
13.2.10. Subquery Syntax .....	1445
13.2.11. UPDATE Syntax .....	1455
13.3. MySQL Transactional and Locking Statements .....	1458
13.3.1. START TRANSACTION, COMMIT, and ROLLBACK Syntax .....	1458
13.3.2. Statements That Cannot Be Rolled Back .....	1460
13.3.3. Statements That Cause an Implicit Commit .....	1461
13.3.4. SAVEPOINT, ROLLBACK TO SAVEPOINT, and RELEASE SAVEPOINT Syntax .....	1462
13.3.5. LOCK TABLES and UNLOCK TABLES Syntax .....	1462
13.3.6. SET TRANSACTION Syntax .....	1468
13.3.7. XA Transactions .....	1470
13.4. Replication Statements .....	1474
13.4.1. SQL Statements for Controlling Master Servers .....	1474
13.4.2. SQL Statements for Controlling Slave Servers .....	1476
13.5. SQL Syntax for Prepared Statements .....	1486

13.5.1. <a href="#">PREPARE</a> Syntax .....	1489
13.5.2. <a href="#">EXECUTE</a> Syntax .....	1490
13.5.3. <a href="#">DEALLOCATE PREPARE</a> Syntax .....	1490
13.6. MySQL Compound-Statement Syntax .....	1490
13.6.1. <a href="#">BEGIN ... END</a> Compound-Statement Syntax .....	1490
13.6.2. Statement Label Syntax .....	1491
13.6.3. <a href="#">DECLARE</a> Syntax .....	1492
13.6.4. Variables in Stored Programs .....	1492
13.6.5. Flow Control Statements .....	1493
13.6.6. Cursors .....	1498
13.6.7. Condition Handling .....	1499
13.7. Database Administration Statements .....	1521
13.7.1. Account Management Statements .....	1521
13.7.2. Table Maintenance Statements .....	1536
13.7.3. Plugin and User-Defined Function Statements .....	1544
13.7.4. <a href="#">SET</a> Syntax .....	1547
13.7.5. <a href="#">SHOW</a> Syntax .....	1549
13.7.6. Other Administrative Statements .....	1588
13.8. MySQL Utility Statements .....	1596
13.8.1. <a href="#">DESCRIBE</a> Syntax .....	1596
13.8.2. <a href="#">EXPLAIN</a> Syntax .....	1596
13.8.3. <a href="#">HELP</a> Syntax .....	1598
13.8.4. <a href="#">USE</a> Syntax .....	1599

This chapter describes the syntax for the [SQL](#) statements supported by MySQL.

## 13.1. Data Definition Statements

### 13.1.1. [ALTER DATABASE](#) Syntax

```
ALTER {DATABASE | SCHEMA} [db_name]
    alter_specification ...
ALTER {DATABASE | SCHEMA} db_name
    UPGRADE DATA DIRECTORY NAME

alter_specification:
    [DEFAULT] CHARACTER SET [=] charset_name
    | [DEFAULT] COLLATE [=] collation_name
```

[ALTER DATABASE](#) enables you to change the overall characteristics of a database. These characteristics are stored in the `db.opt` file in the database directory. To use [ALTER DATABASE](#), you need the [ALTER \[753\]](#) privilege on the database. [ALTER SCHEMA](#) is a synonym for [ALTER DATABASE](#).

The database name can be omitted from the first syntax, in which case the statement applies to the default database.

### National Language Characteristics

The [CHARACTER SET](#) clause changes the default database character set. The [COLLATE](#) clause changes the default database collation. [Section 10.1, “Character Set Support”](#), discusses character set and collation names.

You can see what character sets and collations are available using, respectively, the [SHOW CHARACTER SET](#) and [SHOW COLLATION](#) statements. See [Section 13.7.5.4, “SHOW CHARACTER SET Syntax”](#), and [Section 13.7.5.5, “SHOW COLLATION Syntax”](#), for more information.

If you change the default character set or collation for a database, stored routines that use the database defaults must be dropped and recreated so that they use the new defaults. (In a stored routine, variables with character data types use the database defaults if the character set or collation are not specified explicitly. See [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#).)

## Upgrading from Versions Older than MySQL 5.1

The syntax that includes the `UPGRADE DATA DIRECTORY NAME` clause updates the name of the directory associated with the database to use the encoding implemented in MySQL 5.1 for mapping database names to database directory names (see [Section 9.2.3, “Mapping of Identifiers to File Names”](#)). This clause is for use under these conditions:

- It is intended when upgrading MySQL to 5.1 or later from older versions.
- It is intended to update a database directory name to the current encoding format if the name contains special characters that need encoding.
- The statement is used by `mysqlcheck` (as invoked by `mysql_upgrade`).

For example, if a database in MySQL 5.0 has the name `a-b-c`, the name contains instances of the `-` (dash) character. In MySQL 5.0, the database directory is also named `a-b-c`, which is not necessarily safe for all file systems. In MySQL 5.1 and later, the same database name is encoded as `a@002db@002dc` to produce a file system-neutral directory name.

When a MySQL installation is upgraded to MySQL 5.1 or later from an older version, the server displays a name such as `a-b-c` (which is in the old format) as `#mysql150#a-b-c`, and you must refer to the name using the `#mysql150#` prefix. Use `UPGRADE DATA DIRECTORY NAME` in this case to explicitly tell the server to re-encode the database directory name to the current encoding format:

```
ALTER DATABASE `#mysql150#a-b-c` UPGRADE DATA DIRECTORY NAME;
```

After executing this statement, you can refer to the database as `a-b-c` without the special `#mysql150#` prefix.

### 13.1.2. ALTER EVENT Syntax

```
ALTER
[DEFINER = { user | CURRENT_USER }]
EVENT event_name
[ON SCHEDULE schedule]
[ON COMPLETION [NOT] PRESERVE]
[RENAME TO new_event_name]
[ENABLE | DISABLE | DISABLE ON SLAVE]
[COMMENT 'comment']
[DO event_body]
```

The `ALTER EVENT` statement changes one or more of the characteristics of an existing event without the need to drop and recreate it. The syntax for each of the `DEFINER`, `ON SCHEDULE`, `ON COMPLETION`, `COMMENT`, `ENABLE / DISABLE`, and `DO` clauses is exactly the same as when used with `CREATE EVENT`. (See [Section 13.1.9, “CREATE EVENT Syntax”](#).)

Any user can alter an event defined on a database for which that user has the `EVENT [754]` privilege. When a user executes a successful `ALTER EVENT` statement, that user becomes the definer for the affected event.

`ALTER EVENT` works only with an existing event:

```
mysql> ALTER EVENT no_such_event
> ON SCHEDULE
> EVERY '2:3' DAY_HOUR;
ERROR 1517 (HY000): Unknown event 'no_such_event'
```

In each of the following examples, assume that the event named `myevent` is defined as shown here:

```
CREATE EVENT myevent
ON SCHEDULE
EVERY 6 HOUR
COMMENT 'A sample comment.'
DO
```

```
UPDATE myschema.mytable SET mycol = mycol + 1;
```

The following statement changes the schedule for `myevent` from once every six hours starting immediately to once every twelve hours, starting four hours from the time the statement is run:

```
ALTER EVENT myevent
ON SCHEDULE
EVERY 12 HOUR
STARTS CURRENT_TIMESTAMP + INTERVAL 4 HOUR;
```

It is possible to change multiple characteristics of an event in a single statement. This example changes the SQL statement executed by `myevent` to one that deletes all records from `mytable`; it also changes the schedule for the event such that it executes once, one day after this `ALTER EVENT` statement is run.

```
ALTER EVENT myevent
ON SCHEDULE
AT CURRENT_TIMESTAMP + INTERVAL 1 DAY
DO
TRUNCATE TABLE myschema.mytable;
```

Specify the options in an `ALTER EVENT` statement only for those characteristics that you want to change; omitted options keep their existing values. This includes any default values for `CREATE EVENT` such as `ENABLE`.

To disable `myevent`, use this `ALTER EVENT` statement:

```
ALTER EVENT myevent
DISABLE;
```

The `ON SCHEDULE` clause may use expressions involving built-in MySQL functions and user variables to obtain any of the *timestamp* or *interval* values which it contains. You cannot use stored routines or user-defined functions in such expressions, and you cannot use any table references; however, you can use `SELECT FROM DUAL`. This is true for both `ALTER EVENT` and `CREATE EVENT` statements. References to stored routines, user-defined functions, and tables in such cases are specifically not permitted, and fail with an error (see Bug #22830).

Although an `ALTER EVENT` statement that contains another `ALTER EVENT` statement in its `DO` clause appears to succeed, when the server attempts to execute the resulting scheduled event, the execution fails with an error.

To rename an event, use the `ALTER EVENT` statement's `RENAME TO` clause. This statement renames the event `myevent` to `yourevent`:

```
ALTER EVENT myevent
RENAME TO yourevent;
```

You can also move an event to a different database using `ALTER EVENT ... RENAME TO ...` and `db_name.event_name` notation, as shown here:

```
ALTER EVENT olddb.myevent
RENAME TO newdb.myevent;
```

To execute the previous statement, the user executing it must have the `EVENT [754]` privilege on both the `olddb` and `newdb` databases.



#### Note

There is no `RENAME EVENT` statement.

The value `DISABLE ON SLAVE` is used on a replication slave instead of `ENABLED` or `DISABLED` to indicate an event that was created on the master and replicated to the slave, but that is not executed on the slave. Normally, `DISABLE ON SLAVE` is set automatically as required; however, there are

some circumstances under which you may want or need to change it manually. See [Section 16.4.1.10, “Replication of Invoked Features”](#), for more information.

### 13.1.3. ALTER FUNCTION Syntax

```
ALTER FUNCTION func_name [characteristic ...]

characteristic:
  COMMENT 'string'
| LANGUAGE SQL
| { CONTAINS SQL | NO SQL | READS SQL DATA | MODIFIES SQL DATA }
| SQL SECURITY { DEFINER | INVOKER }
```

This statement can be used to change the characteristics of a stored function. More than one change may be specified in an `ALTER FUNCTION` statement. However, you cannot change the parameters or body of a stored function using this statement; to make such changes, you must drop and re-create the function using `DROP FUNCTION` and `CREATE FUNCTION`.

You must have the `ALTER ROUTINE` [754] privilege for the function. (That privilege is granted automatically to the function creator.) If binary logging is enabled, the `ALTER FUNCTION` statement might also require the `SUPER` [756] privilege, as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

### 13.1.4. ALTER PROCEDURE Syntax

```
ALTER PROCEDURE proc_name [characteristic ...]

characteristic:
  COMMENT 'string'
| LANGUAGE SQL
| { CONTAINS SQL | NO SQL | READS SQL DATA | MODIFIES SQL DATA }
| SQL SECURITY { DEFINER | INVOKER }
```

This statement can be used to change the characteristics of a stored procedure. More than one change may be specified in an `ALTER PROCEDURE` statement. However, you cannot change the parameters or body of a stored procedure using this statement; to make such changes, you must drop and re-create the procedure using `DROP PROCEDURE` and `CREATE PROCEDURE`.

You must have the `ALTER ROUTINE` [754] privilege for the procedure. By default, that privilege is granted automatically to the procedure creator. This behavior can be changed by disabling the `automatic_sp_privileges` [464] system variable. See [Section 18.2.2, “Stored Routines and MySQL Privileges”](#).

### 13.1.5. ALTER SERVER Syntax

```
ALTER SERVER server_name
  OPTIONS (option [, option] ...)
```

Alters the server information for *server\_name*, adjusting any of the options permitted in the `CREATE SERVER` statement. The corresponding fields in the `mysql.servers` table are updated accordingly. This statement requires the `SUPER` [756] privilege.

For example, to update the `USER` option:

```
ALTER SERVER s OPTIONS (USER 'sally');
```

`ALTER SERVER` does not cause an automatic commit.

In MySQL 5.6, `ALTER SERVER` is not written to the binary log, regardless of the logging format that is in use.

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

### 13.1.6. ALTER TABLE Syntax

```

ALTER [IGNORE] TABLE tbl_name
  [alter_specification [, alter_specification] ...]
  [partition_options]

alter_specification:
  table_options
| ADD [COLUMN] col_name column_definition
  [FIRST | AFTER col_name ]
| ADD [COLUMN] (col_name column_definition,...)
| ADD {INDEX|KEY} [index_name]
  [index_type] (index_col_name,...) [index_option] ...
| ADD [CONSTRAINT [symbol]] PRIMARY KEY
  [index_type] (index_col_name,...) [index_option] ...
| ADD [CONSTRAINT [symbol]]
  UNIQUE [INDEX|KEY] [index_name]
  [index_type] (index_col_name,...) [index_option] ...
| ADD FULLTEXT [INDEX|KEY] [index_name]
  (index_col_name,...) [index_option] ...
| ADD SPATIAL [INDEX|KEY] [index_name]
  (index_col_name,...) [index_option] ...
| ADD [CONSTRAINT [symbol]]
  FOREIGN KEY [index_name] (index_col_name,...)
  reference_definition
| ALGORITHM [=] {DEFAULT|INPLACE|COPY}
| ALTER [COLUMN] col_name {SET DEFAULT literal | DROP DEFAULT}
| CHANGE [COLUMN] old_col_name new_col_name column_definition
  [FIRST|AFTER col_name]
| LOCK [=] {DEFAULT|NONE|SHARED|EXCLUSIVE}
| MODIFY [COLUMN] col_name column_definition
  [FIRST | AFTER col_name]
| DROP [COLUMN] col_name
| DROP PRIMARY KEY
| DROP {INDEX|KEY} index_name
| DROP FOREIGN KEY fk_symbol
| DISABLE KEYS
| ENABLE KEYS
| RENAME [TO|AS] new_tbl_name
| ORDER BY col_name [, col_name] ...
| CONVERT TO CHARACTER SET charset_name [COLLATE collation_name]
| [DEFAULT] CHARACTER SET [=] charset_name [COLLATE [=] collation_name]
| DISCARD TABLESPACE
| IMPORT TABLESPACE
| FORCE
| ADD PARTITION (partition_definition)
| DROP PARTITION partition_names
| TRUNCATE PARTITION {partition_names | ALL}
| COALESCE PARTITION number
| REORGANIZE PARTITION partition_names INTO (partition_definitions)
| EXCHANGE PARTITION partition_name WITH TABLE tbl_name
| ANALYZE PARTITION {partition_names | ALL}
| CHECK PARTITION {partition_names | ALL}
| OPTIMIZE PARTITION {partition_names | ALL}
| REBUILD PARTITION {partition_names | ALL}
| REPAIR PARTITION {partition_names | ALL}
| REMOVE PARTITIONING

index_col_name:
  col_name [(length)] [ASC | DESC]

index_type:
  USING {BTREE | HASH}

index_option:
  KEY_BLOCK_SIZE [=] value
| index_type
| WITH PARSER parser_name
| COMMENT 'string'

table_options:
  table_option [[,] table_option] ... (see CREATE TABLE options)

```



```
partition_options:
  (see CREATE TABLE options)
```

`ALTER TABLE` changes the structure of a table. For example, you can add or delete columns, create or destroy indexes, change the type of existing columns, or rename columns or the table itself. You can also change characteristics such as the storage engine used for the table or the table comment.

Partitioning-related clauses for `ALTER TABLE` can be used with partitioned tables for repartitioning, for adding, dropping, merging, and splitting partitions, and for performing partitioning maintenance. It is possible for an `ALTER TABLE` statement to contain a `PARTITION BY` or `REMOVE PARTITION` clause in an addition to other alter specifications, but the `PARTITION BY` or `REMOVE PARTITION` clause must be specified last after any other specifications. The `ADD PARTITION`, `DROP PARTITION`, `COALESCE PARTITION`, `REORGANIZE PARTITION`, `ANALYZE PARTITION`, `CHECK PARTITION`, and `REPAIR PARTITION` options cannot be combined with other alter specifications in a single `ALTER TABLE`, since the options just listed act on individual partitions. For more information, see [Section 13.1.6.1, “ALTER TABLE Partition Operations”](#).

Following the table name, specify the alterations to be made. If none are given, `ALTER TABLE` does nothing.

The syntax for many of the permissible alterations is similar to clauses of the `CREATE TABLE` statement. See [Section 13.1.14, “CREATE TABLE Syntax”](#), for more information.

Some operations may result in warnings if attempted on a table for which the storage engine does not support the operation. These warnings can be displayed with `SHOW WARNINGS`. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

For information on troubleshooting `ALTER TABLE`, see [Section C.5.7.1, “Problems with ALTER TABLE”](#).

## Storage, Performance, and Concurrency Considerations

In most cases, `ALTER TABLE` makes a temporary copy of the original table. MySQL waits for other operations that are modifying the table, then proceeds. It incorporates the alteration into the copy, deletes the original table, and renames the new one. While `ALTER TABLE` is executing, the original table is readable by other sessions (with the exception noted shortly). Updates and writes to the table that begin after the `ALTER TABLE` operation begins are stalled until the new table is ready, then are automatically redirected to the new table without any failed updates. The temporary table is created in the database directory of the new table. This can differ from the database directory of the original table for `ALTER TABLE` operations that rename the table to a different database.

The exception referred to earlier is that `ALTER TABLE` blocks reads (not just writes) at the point where it is ready to install a new version of the table `.frm` file, discard the old file, and clear outdated table structures from the table and table definition caches. At this point, it must acquire an exclusive lock. To do so, it waits for current readers to finish, and blocks new reads (and writes).

For `MyISAM` tables, you can speed up index re-creation (the slowest part of the alteration process) by setting the `myisam_sort_buffer_size` [513] system variable to a high value.

For some operations, an in-place `ALTER TABLE` is possible that does not require a temporary table:

- For `ALTER TABLE tbl_name RENAME TO new_tbl_name` without any other options, MySQL simply renames any files that correspond to the table `tbl_name` without making a copy. (You can also use the `RENAME TABLE` statement to rename tables. See [Section 13.1.26, “RENAME TABLE Syntax”](#).) Any privileges granted specifically for the renamed table are not migrated to the new name. They must be changed manually.
- Alterations that modify only table metadata and not table data can be made immediately by altering the table's `.frm` file and not touching table contents. The following changes are fast alterations that can be made this way:

- Renaming a column, except for the [InnoDB](#) storage engine before MySQL 5.6.6.
- Changing the default value of a column.
- Changing the definition of an [ENUM](#) or [SET](#) column by adding new enumeration or set members to the *end* of the list of valid member values, as long as the storage size of the data type does not change. For example, adding a member to a [SET](#) column that has 8 members changes the required storage per value from 1 byte to 2 bytes; this will require a table copy. Adding members in the middle of the list causes renumbering of existing members, which requires a table copy.
- [ALTER TABLE](#) with [ADD PARTITION](#), [DROP PARTITION](#), [COALESCE PARTITION](#), [REBUILD PARTITION](#), or [REORGANIZE PARTITION](#) does not create any temporary tables (except when used with [NDB](#) tables); however, these operations can and do create temporary partition files.

[ADD](#) or [DROP](#) operations for [RANGE](#) or [LIST](#) partitions are immediate operations or nearly so. [ADD](#) or [COALESCE](#) operations for [HASH](#) or [KEY](#) partitions copy data between all partitions, unless [LINEAR HASH](#) or [LINEAR KEY](#) was used; this is effectively the same as creating a new table, although the [ADD](#) or [COALESCE](#) operation is performed partition by partition. [REORGANIZE](#) operations copy only changed partitions and do not touch unchanged ones.

- Renaming an index, except for [InnoDB](#).
- Adding or dropping an index, for [InnoDB](#).

You can force an [ALTER TABLE](#) operation that would otherwise not require a table copy to use the temporary table method (as supported in MySQL 5.0) by setting the `old_alter_table` [517] system variable to `ON`, or specifying `ALGORITHM=COPY` as one of the *alter\_specification* clauses. If there is a conflict between the `old_alter_table` setting and an `ALGORITHM` clause with a value other than `DEFAULT`, the `ALGORITHM` clause takes precedence. (`ALGORITHM = DEFAULT` is the same as specifying no `ALGORITHM` clause at all.)

Specifying `ALGORITHM=INPLACE` makes the operation use the in-place technique for clauses and storage engines that support it, and fail with an error otherwise, thus avoiding a lengthy table copy if you try altering a table that uses a different storage engine than you expect. See [Section 5.5, “Online DDL for InnoDB Tables”](#) for information about online DDL for [InnoDB](#) tables.

You can control the level of concurrent reading and writing of the table while it is being altered, using the `LOCK` clause. Specifying a non-default value for this clause lets you require a certain amount of concurrent access or exclusivity during the alter operation, and halts the operation if the requested degree of locking is not available. The parameters for the `LOCK` clause are:

- `LOCK = DEFAULT`

Maximum level of concurrency for the given `ALGORITHM` clause (if any) and [ALTER TABLE](#) operation: Permit concurrent reads and writes if supported. If not, permit concurrent reads if supported. If not, enforce exclusive access.

- `LOCK = NONE`

If supported, permit concurrent reads and writes. Otherwise, return an error message.

- `LOCK = SHARED`

If supported, permit concurrent reads but block writes. Note that writes will be blocked even if concurrent writes are supported by the storage engine for the given `ALGORITHM` clause (if any) and [ALTER TABLE](#) operation. If concurrent reads are not supported, return an error message.

- `LOCK = EXCLUSIVE`

Enforce exclusive access. This will be done even if concurrent reads/writes are supported by the storage engine for the given `ALGORITHM` clause (if any) and [ALTER TABLE](#) operation.

As of MySQL 5.6.3, you can also use `ALTER TABLE tbl_name FORCE` to perform a “null” alter operation that rebuilds the table. Previously the `FORCE` option was recognized but ignored.

## Usage Notes

- To use `ALTER TABLE`, you need `ALTER` [753], `CREATE` [754], and `INSERT` [755] privileges for the table. Renaming a table requires `ALTER` [753] and `DROP` [754] on the old table, `ALTER` [753], `CREATE` [754], and `INSERT` [755] on the new table.
- `IGNORE` is a MySQL extension to standard SQL. It controls how `ALTER TABLE` works if there are duplicates on unique keys in the new table or if warnings occur when strict mode is enabled. If `IGNORE` is not specified, the copy is aborted and rolled back if duplicate-key errors occur. If `IGNORE` is specified, only the first row is used of rows with duplicates on a unique key. The other conflicting rows are deleted. Incorrect values are truncated to the closest matching acceptable value.
- Pending `INSERT DELAYED` statements are lost if a table is write locked and `ALTER TABLE` is used to modify the table structure.
- `table_option` signifies a table option of the kind that can be used in the `CREATE TABLE` statement, such as `ENGINE`, `AUTO_INCREMENT`, or `AVG_ROW_LENGTH`. (Section 13.1.14, “`CREATE TABLE Syntax`”, lists all table options.) However, `ALTER TABLE` ignores the `DATA DIRECTORY` and `INDEX DIRECTORY` table options.

For example, to convert a table to be an `InnoDB` table, use this statement:

```
ALTER TABLE t1 ENGINE = InnoDB;
```

See Section 5.4.3, “`Converting Tables from MyISAM to InnoDB`” for considerations when switching tables to the `InnoDB` storage engine.

When you specify an `ENGINE` clause, `ALTER TABLE` rebuilds the table. This is true even if the table already has the specified storage engine.

The outcome of attempting to change a table's storage engine is affected by whether the desired storage engine is available and the setting of the `NO_ENGINE_SUBSTITUTION` [605] SQL mode, as described in Section 5.1.7, “`Server SQL Modes`”.

To prevent inadvertent loss of data, `ALTER TABLE` cannot be used to change the storage engine of a table to `MERGE` or `BLACKHOLE`.

To change the value of the `AUTO_INCREMENT` counter to be used for new rows, do this:

```
ALTER TABLE t2 AUTO_INCREMENT = value;
```

You cannot reset the counter to a value less than or equal to any that have already been used. For `MyISAM`, if the value is less than or equal to the maximum value currently in the `AUTO_INCREMENT` column, the value is reset to the current maximum plus one. For `InnoDB`, *if the value is less than the current maximum value in the column, no error occurs and the current sequence value is not changed.*

- You can issue multiple `ADD`, `ALTER`, `DROP`, and `CHANGE` clauses in a single `ALTER TABLE` statement, separated by commas. This is a MySQL extension to standard SQL, which permits only one of each clause per `ALTER TABLE` statement. For example, to drop multiple columns in a single statement, do this:

```
ALTER TABLE t2 DROP COLUMN c, DROP COLUMN d;
```

- `CHANGE col_name`, `DROP col_name`, and `DROP INDEX` are MySQL extensions to standard SQL.
- The word `COLUMN` is optional and can be omitted.
- `column_definition` clauses use the same syntax for `ADD` and `CHANGE` as for `CREATE TABLE`. See Section 13.1.14, “`CREATE TABLE Syntax`”.

- You can rename a column using a `CHANGE old_col_name new_col_name column_definition` clause. To do so, specify the old and new column names and the definition that the column currently has. For example, to rename an `INTEGER` column from `a` to `b`, you can do this:

```
ALTER TABLE t1 CHANGE a b INTEGER;
```

To change a column's type but not the name, `CHANGE` syntax still requires an old and new column name, even if they are the same. For example:

```
ALTER TABLE t1 CHANGE b b BIGINT NOT NULL;
```

You can also use `MODIFY` to change a column's type without renaming it:

```
ALTER TABLE t1 MODIFY b BIGINT NOT NULL;
```

`MODIFY` is an extension to `ALTER TABLE` for Oracle compatibility.

When you use `CHANGE` or `MODIFY`, *column\_definition* must include the data type and all attributes that should apply to the new column, other than index attributes such as `PRIMARY KEY` or `UNIQUE`. Attributes present in the original definition but not specified for the new definition are not carried forward. Suppose that a column `col1` is defined as `INT UNSIGNED DEFAULT 1 COMMENT 'my column'` and you modify the column as follows:

```
ALTER TABLE t1 MODIFY col1 BIGINT;
```

The resulting column will be defined as `BIGINT`, but will not include the attributes `UNSIGNED DEFAULT 1 COMMENT 'my column'`. To retain them, the statement should be:

```
ALTER TABLE t1 MODIFY col1 BIGINT UNSIGNED DEFAULT 1 COMMENT 'my column';
```

- When you change a data type using `CHANGE` or `MODIFY`, MySQL tries to convert existing column values to the new type as well as possible.



### Warning

This conversion may result in alteration of data. For example, if you shorten a string column, values may be truncated. To prevent the operation from succeeding if conversions to the new data type would result in loss of data, enable strict SQL mode before using `ALTER TABLE` (see [Section 5.1.7, “Server SQL Modes”](#)).

- To add a column at a specific position within a table row, use `FIRST` or `AFTER col_name`. The default is to add the column last. You can also use `FIRST` and `AFTER` in `CHANGE` or `MODIFY` operations to reorder columns within a table.
- `ALTER ... SET DEFAULT` or `ALTER ... DROP DEFAULT` specify a new default value for a column or remove the old default value, respectively. If the old default is removed and the column can be `NULL`, the new default is `NULL`. If the column cannot be `NULL`, MySQL assigns a default value as described in [Section 11.5, “Data Type Default Values”](#).
- `DROP INDEX` removes an index. This is a MySQL extension to standard SQL. See [Section 13.1.20, “DROP INDEX Syntax”](#). If you are unsure of the index name, use `SHOW INDEX FROM tbl_name`.
- If columns are dropped from a table, the columns are also removed from any index of which they are a part. If all columns that make up an index are dropped, the index is dropped as well. If you use `CHANGE` or `MODIFY` to shorten a column for which an index exists on the column, and the resulting column length is less than the index length, MySQL shortens the index automatically.
- If a table contains only one column, the column cannot be dropped. If what you intend is to remove the table, use `DROP TABLE` instead.

- `DROP PRIMARY KEY` drops the [primary key](#). If there is no primary key, an error occurs. For information about the performance characteristics of primary keys, especially for [InnoDB](#) tables, see [Section 8.3.2, “Using Primary Keys”](#).

If you add a `UNIQUE INDEX` or `PRIMARY KEY` to a table, MySQL stores it before any nonunique index to permit detection of duplicate keys as early as possible.

- Some storage engines permit you to specify an index type when creating an index. The syntax for the `index_type` specifier is `USING type_name`. For details about `USING`, see [Section 13.1.11, “CREATE INDEX Syntax”](#). The preferred position is after the column list. Support for use of the option before the column list will be removed in a future MySQL release.

`index_option` values specify additional options for an index. `USING` is one such option. For details about permissible `index_option` values, see [Section 13.1.11, “CREATE INDEX Syntax”](#).

- After an `ALTER TABLE` statement, it may be necessary to run `ANALYZE TABLE` to update index cardinality information. See [Section 13.7.5.23, “SHOW INDEX Syntax”](#).
- `ORDER BY` enables you to create the new table with the rows in a specific order. Note that the table does not remain in this order after inserts and deletes. This option is useful primarily when you know that you are mostly to query the rows in a certain order most of the time. By using this option after major changes to the table, you might be able to get higher performance. In some cases, it might make sorting easier for MySQL if the table is in order by the column that you want to order it by later.

`ORDER BY` syntax permits one or more column names to be specified for sorting, each of which optionally can be followed by `ASC` or `DESC` to indicate ascending or descending sort order, respectively. The default is ascending order. Only column names are permitted as sort criteria; arbitrary expressions are not permitted. This clause should be given last after any other clauses.

`ORDER BY` does not make sense for [InnoDB](#) tables that contain a user-defined clustered index (`PRIMARY KEY` or `NOT NULL UNIQUE` index). [InnoDB](#) always orders table rows according to such an index if one is present.



#### Note

When used on a partitioned table, `ALTER TABLE ... ORDER BY` orders rows within each partition only.

- If you use `ALTER TABLE` on a [MyISAM](#) table, all nonunique indexes are created in a separate batch (as for `REPAIR TABLE`). This should make `ALTER TABLE` much faster when you have many indexes.

This feature can be activated explicitly for a [MyISAM](#) table. `ALTER TABLE ... DISABLE KEYS` tells MySQL to stop updating nonunique indexes. `ALTER TABLE ... ENABLE KEYS` then should be used to re-create missing indexes. MySQL does this with a special algorithm that is much faster than inserting keys one by one, so disabling keys before performing bulk insert operations should give a considerable speedup. Using `ALTER TABLE ... DISABLE KEYS` requires the [INDEX \[755\]](#) privilege in addition to the privileges mentioned earlier.

While the nonunique indexes are disabled, they are ignored for statements such as `SELECT` and `EXPLAIN` that otherwise would use them.

- If `ALTER TABLE` for an [InnoDB](#) table results in changes to column values (for example, because a column is truncated), [InnoDB](#)'s `FOREIGN KEY` constraint checks do not notice possible violations caused by changing the values.
- The `FOREIGN KEY` and `REFERENCES` clauses are supported by the [InnoDB](#) storage engine, which implements `ADD [CONSTRAINT [symbol]] FOREIGN KEY [index_name] (...)` `REFERENCES ... (...)`. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#). For other storage engines, the clauses are parsed but ignored. The `CHECK` clause is parsed but ignored by all storage engines. See [Section 13.1.14, “CREATE TABLE Syntax”](#). The reason for accepting but

ignoring syntax clauses is for compatibility, to make it easier to port code from other SQL servers, and to run applications that create tables with references. See [Section 1.8.5, “MySQL Differences from Standard SQL”](#).

For `ALTER TABLE`, unlike `CREATE TABLE, ADD FOREIGN KEY` ignores `index_name` if given and uses an automatically generated foreign key name. As a workaround, include the `CONSTRAINT` clause to specify the foreign key name:

```
ADD CONSTRAINT name FOREIGN KEY (...)
```



### Important

The inline `REFERENCES` specifications where the references are defined as part of the column specification are silently ignored by InnoDB. InnoDB only accepts `REFERENCES` clauses defined as part of a separate `FOREIGN KEY` specification.



### Note

Partitioned InnoDB tables do not support foreign keys. See [Section 17.6.2, “Partitioning Limitations Relating to Storage Engines”](#), for more information.

- InnoDB supports the use of `ALTER TABLE` to drop foreign keys:

```
ALTER TABLE tbl_name DROP FOREIGN KEY fk_symbol;
```

For more information, see [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

- You cannot add a foreign key and drop a foreign key in separate clauses of a single `ALTER TABLE` statement. You must use separate statements.
- For an InnoDB table that is created with its own tablespace in an `.ibd` file, that file can be discarded and imported. To discard the `.ibd` file, use this statement:

```
ALTER TABLE tbl_name DISCARD TABLESPACE;
```

This deletes the current `.ibd` file, so be sure that you have a backup first. Attempting to modify the table contents while the tablespace file is discarded results in an error. You can perform the DDL operations listed in [Section 5.5, “Online DDL for InnoDB Tables”](#) while the tablespace file is discarded.

To import the backup `.ibd` file back into the table, copy it into the database directory, and then issue this statement:

```
ALTER TABLE tbl_name IMPORT TABLESPACE;
```

The tablespace file need not necessarily have been created on the server into which it is imported later. In MySQL 5.6, importing a tablespace file from another server works if the both servers have GA (General Availability) status and their versions are within the same series. Otherwise, the file must have been created on the server into which it is imported.

See [Section 5.4.1, “Managing InnoDB Tablespaces”](#).

- To change the table default character set and all character columns (`CHAR`, `VARCHAR`, `TEXT`) to a new character set, use a statement like this:

```
ALTER TABLE tbl_name CONVERT TO CHARACTER SET charset_name;
```

For a column that has a data type of `VARCHAR` or one of the `TEXT` types, `CONVERT TO CHARACTER SET` will change the data type as necessary to ensure that the new column is long enough to store as many characters as the original column. For example, a `TEXT` column has two length bytes, which store the byte-length of values in the column, up to a maximum of 65,535. For a `latin1 TEXT` column, each character requires a single byte, so the column can store up to 65,535 characters. If



the column is converted to `utf8`, each character might require up to three bytes, for a maximum possible length of  $3 \times 65,535 = 196,605$  bytes. That length will not fit in a `TEXT` column's length bytes, so MySQL will convert the data type to `MEDIUMTEXT`, which is the smallest string type for which the length bytes can record a value of 196,605. Similarly, a `VARCHAR` column might be converted to `MEDIUMTEXT`.

To avoid data type changes of the type just described, do not use `CONVERT TO CHARACTER SET`. Instead, use `MODIFY` to change individual columns. For example:

```
ALTER TABLE t MODIFY latin1_text_col TEXT CHARACTER SET utf8;
ALTER TABLE t MODIFY latin1_varchar_col VARCHAR(M) CHARACTER SET utf8;
```

If you specify `CONVERT TO CHARACTER SET binary`, the `CHAR`, `VARCHAR`, and `TEXT` columns are converted to their corresponding binary string types (`BINARY`, `VARBINARY`, `BLOB`). This means that the columns no longer will have a character set and a subsequent `CONVERT TO` operation will not apply to them.

If `charset_name` is `DEFAULT`, the database character set is used.



### Warning

The `CONVERT TO` operation converts column values between the character sets. This is *not* what you want if you have a column in one character set (like `latin1`) but the stored values actually use some other, incompatible character set (like `utf8`). In this case, you have to do the following for each such column:

```
ALTER TABLE t1 CHANGE c1 c1 BLOB;
ALTER TABLE t1 CHANGE c1 c1 TEXT CHARACTER SET utf8;
```

The reason this works is that there is no conversion when you convert to or from `BLOB` columns.

To change only the *default* character set for a table, use this statement:

```
ALTER TABLE tbl_name DEFAULT CHARACTER SET charset_name;
```

The word `DEFAULT` is optional. The default character set is the character set that is used if you do not specify the character set for columns that you add to a table later (for example, with `ALTER TABLE ... ADD column`).

With the `mysql_info()` C API function, you can find out how many rows were copied by `ALTER TABLE`, and (when `IGNORE` is used) how many rows were deleted due to duplication of unique key values. See [Section 21.8.6.35](#), “`mysql_info()`”.

### 13.1.6.1. ALTER TABLE Partition Operations

Partitioning-related clauses for `ALTER TABLE` can be used with partitioned tables for repartitioning, for adding, dropping, merging, and splitting partitions, and for performing partitioning maintenance.

- Simply using a `partition_options` clause with `ALTER TABLE` on a partitioned table repartitions the table according to the partitioning scheme defined by the `partition_options`. This clause always begins with `PARTITION BY`, and follows the same syntax and other rules as apply to the `partition_options` clause for `CREATE TABLE` (see [Section 13.1.14](#), “`CREATE TABLE Syntax`”, for more detailed information), and can also be used to partition an existing table that is not already partitioned. For example, consider a (nonpartitioned) table defined as shown here:

```
CREATE TABLE t1 (
  id INT,
  year_col INT
);
```

This table can be partitioned by `HASH`, using the `id` column as the partitioning key, into 8 partitions by means of this statement:

```
ALTER TABLE t1
PARTITION BY HASH(id)
PARTITIONS 8;
```

MySQL 5.6.11 and later supports an `ALGORITHM` option with `[SUB]PARTITION BY [LINEAR] KEY`. `ALGORITHM=1` causes the server to use the same key-hashing functions as MySQL 5.1 when computing the placement of rows in partitions; `ALGORITHM=2` means that the server employs the key-hashing functions implemented and used by default for new `KEY` partitioned tables in MySQL 5.5 and later. (Partitioned tables created with the key-hashing functions employed in MySQL 5.5 and later cannot be used by a MySQL 5.1 server.) Not specifying the option has the same effect as using `ALGORITHM=2`. This option is intended for use chiefly when upgrading or downgrading `[LINEAR] KEY` partitioned tables between MySQL 5.1 and later MySQL versions, or for creating tables partitioned by `KEY` or `LINEAR KEY` on a MySQL 5.5 or later server which can be used on a MySQL 5.1 server.

To upgrade a `KEY` partitioned table that was created in MySQL 5.1, first execute `SHOW CREATE TABLE` and note the the exact columns and number of partitions shown. Now execute an `ALTER TABLE` statement using exactly the same column list and number of partitions as in the `CREATE TABLE` statement, while adding `ALGORITHM=2` immediately following the `PARTITION BY` keywords. (You should also include the `LINEAR` keyword if it was used for the original table definition.) An example from a session in the `mysql` client is shown here:

```
mysql> SHOW CREATE TABLE p\G
***** 1. row *****
Table: p
Create Table: CREATE TABLE `p` (
  `id` int(11) NOT NULL AUTO_INCREMENT,
  `cd` datetime NOT NULL,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB DEFAULT CHARSET=latin1
/*!50100 PARTITION BY LINEAR KEY (id)
PARTITIONS 32 */
1 row in set (0.00 sec)

mysql> ALTER TABLE p PARTITION BY LINEAR KEY ALGORITHM=2 (id) PARTITIONS 32;
Query OK, 0 rows affected (5.34 sec)
Records: 0 Duplicates: 0 Warnings: 0

mysql> SHOW CREATE TABLE p\G
***** 1. row *****
Table: p
Create Table: CREATE TABLE `p` (
  `id` int(11) NOT NULL AUTO_INCREMENT,
  `cd` datetime NOT NULL,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB DEFAULT CHARSET=latin1
/*!50100 PARTITION BY LINEAR KEY (id)
PARTITIONS 32 */
1 row in set (0.00 sec)
```

Downgrading a table created using the default key-hashing used in MySQL 5.5 and later to enable its use by a MySQL 5.1 server is similar, except in this case you should use `ALGORITHM=1` to force the table's partitions to be rebuilt using the MySQL 5.1 key-hashing functions. It is recommended that you not do this except when necessary for compatibility with a MySQL 5.1 server, as the improved `KEY` hashing functions used by default in MySQL 5.5 and later provide fixes for a number of issues found in the older implementation.



#### Note

A table upgraded by means of `ALTER TABLE ... PARTITION BY ALGORITHM=2 [LINEAR] KEY ...` can no longer be used by a MySQL 5.1



server. (Such a table would need to be downgraded with `ALTER TABLE ... PARTITION BY ALGORITHM=1 [LINEAR] KEY ...` before it could be used again by a MySQL 5.1 server.)

The table that results from using an `ALTER TABLE ... PARTITION BY` statement must follow the same rules as one created using `CREATE TABLE ... PARTITION BY`. This includes the rules governing the relationship between any unique keys (including any primary key) that the table might have, and the column or columns used in the partitioning expression, as discussed in [Section 17.6.1, “Partitioning Keys, Primary Keys, and Unique Keys”](#). The `CREATE TABLE ... PARTITION BY` rules for specifying the number of partitions also apply to `ALTER TABLE ... PARTITION BY`.

The `partition_definition` clause for `ALTER TABLE ADD PARTITION` supports the same options as the clause of the same name for the `CREATE TABLE` statement. (See [Section 13.1.14, “CREATE TABLE Syntax”](#), for the syntax and description.) Suppose that you have the partitioned table created as shown here:

```
CREATE TABLE t1 (
  id INT,
  year_col INT
)
PARTITION BY RANGE (year_col) (
  PARTITION p0 VALUES LESS THAN (1991),
  PARTITION p1 VALUES LESS THAN (1995),
  PARTITION p2 VALUES LESS THAN (1999)
);
```

You can add a new partition `p3` to this table for storing values less than 2002 as follows:

```
ALTER TABLE t1 ADD PARTITION (PARTITION p3 VALUES LESS THAN (2002));
```

`DROP PARTITION` can be used to drop one or more `RANGE` or `LIST` partitions. This statement cannot be used with `HASH` or `KEY` partitions; instead, use `COALESCE PARTITION` (see below). Any data that was stored in the dropped partitions named in the `partition_names` list is discarded. For example, given the table `t1` defined previously, you can drop the partitions named `p0` and `p1` as shown here:

```
ALTER TABLE t1 DROP PARTITION p0, p1;
```

`ADD PARTITION` and `DROP PARTITION` do not currently support `IF [NOT] EXISTS`.

Renames of partitioned table are supported. You can rename individual partitions indirectly using `ALTER TABLE ... REORGANIZE PARTITION`; however, this operation makes a copy of the partition's data..

In MySQL 5.6, it is possible to delete rows from selected partitions using the `TRUNCATE PARTITION` option. This option takes a comma-separated list of one or more partition names. For example, consider the table `t1` as defined here:

```
CREATE TABLE t1 (
  id INT,
  year_col INT
)
PARTITION BY RANGE (year_col) (
  PARTITION p0 VALUES LESS THAN (1991),
  PARTITION p1 VALUES LESS THAN (1995),
  PARTITION p2 VALUES LESS THAN (1999),
  PARTITION p3 VALUES LESS THAN (2003),
  PARTITION p4 VALUES LESS THAN (2007)
);
```

To delete all rows from partition `p0`, you can use the following statement:

```
ALTER TABLE t1 TRUNCATE PARTITION p0;
```

The statement just shown has the same effect as the following `DELETE` statement:

```
DELETE FROM t1 WHERE year_col < 1991;
```

When truncating multiple partitions, the partitions do not have to be contiguous: This can greatly simplify delete operations on partitioned tables that would otherwise require very complex [WHERE](#) conditions if done with [DELETE](#) statements. For example, this statement deletes all rows from partitions [p1](#) and [p3](#):

```
ALTER TABLE t1 TRUNCATE PARTITION p1, p3;
```

An equivalent [DELETE](#) statement is shown here:

```
DELETE FROM t1 WHERE
  (year_col >= 1991 AND year_col < 1995)
  OR
  (year_col >= 2003 AND year_col < 2007);
```

You can also use the [ALL](#) keyword in place of the list of partition names; in this case, the statement acts on all partitions in the table.

[TRUNCATE PARTITION](#) merely deletes rows; it does not alter the definition of the table itself, or of any of its partitions.



#### Note

[TRUNCATE PARTITION](#) does not work with subpartitions.

You can verify that the rows were dropped by checking the [INFORMATION\\_SCHEMA.PARTITIONS](#) table, using a query such as this one:

```
SELECT PARTITION_NAME, TABLE_ROWS
  FROM INFORMATION_SCHEMA.PARTITIONS
  WHERE TABLE_NAME = 't1';
```

[TRUNCATE PARTITION](#) is supported only for partitioned tables that use the [MyISAM](#), [InnoDB](#), or [MEMORY](#) storage engine. It also works on [BLACKHOLE](#) tables (but has no effect). It is not supported for [ARCHIVE](#) tables.

[COALESCE PARTITION](#) can be used with a table that is partitioned by [HASH](#) or [KEY](#) to reduce the number of partitions by *number*. Suppose that you have created table [t2](#) using the following definition:

```
CREATE TABLE t2 (
  name VARCHAR (30),
  started DATE
)
PARTITION BY HASH( YEAR(started) )
PARTITIONS 6;
```

You can reduce the number of partitions used by [t2](#) from 6 to 4 using the following statement:

```
ALTER TABLE t2 COALESCE PARTITION 2;
```

The data contained in the last *number* partitions will be merged into the remaining partitions. In this case, partitions 4 and 5 will be merged into the first 4 partitions (the partitions numbered 0, 1, 2, and 3).

To change some but not all the partitions used by a partitioned table, you can use [REORGANIZE PARTITION](#). This statement can be used in several ways:

- To merge a set of partitions into a single partition. This can be done by naming several partitions in the *partition\_names* list and supplying a single definition for *partition\_definition*.

- To split an existing partition into several partitions. You can accomplish this by naming a single partition for *partition\_names* and providing multiple *partition\_definitions*.
- To change the ranges for a subset of partitions defined using `VALUES LESS THAN` or the value lists for a subset of partitions defined using `VALUES IN`.



### Note

For partitions that have not been explicitly named, MySQL automatically provides the default names `p0`, `p1`, `p2`, and so on. The same is true with regard to subpartitions.

For more detailed information about and examples of `ALTER TABLE ... REORGANIZE PARTITION` statements, see [Section 17.3.1, “Management of RANGE and LIST Partitions”](#).

- It is also possible in MySQL 5.6 to exchange a table partition or subpartition with a table using `ALTER TABLE pt EXCHANGE PARTITION p WITH TABLE nt`, where *pt* is the partitioned table and *p* is the partition or subpartition of *pt* to be exchanged with unpartitioned table *nt*, provided that the following statements are true:
  1. Table *nt* is not itself partitioned.
  2. Table *nt* is not a temporary table.
  3. The structures of tables *pt* and *nt* are otherwise identical.
  4. There are no rows in *nt* that lie outside the boundaries of the partition definition for *p*.
  5. Table *nt* contains no foreign key references, and no other table has any foreign keys that refer to *nt*.

Executing `ALTER TABLE ... EXCHANGE PARTITION` does not invoke any triggers on either the partitioned table or the table to be exchanged.

Any `AUTO_INCREMENT` columns in the table to be exchanged with a partition are reset.

The `IGNORE` keyword has no effect when used with `ALTER TABLE ... EXCHANGE PARTITION`.

For more information about and examples of `ALTER TABLE ... EXCHANGE PARTITION`, see [Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”](#).

- Several additional options provide partition maintenance and repair functionality analogous to that implemented for nonpartitioned tables by statements such as `CHECK TABLE` and `REPAIR TABLE` (which are also supported for partitioned tables; see [Section 13.7.2, “Table Maintenance Statements”](#) for more information). These include `ANALYZE PARTITION`, `CHECK PARTITION`, `OPTIMIZE PARTITION`, `REBUILD PARTITION`, and `REPAIR PARTITION`. Each of these options takes a *partition\_names* clause consisting of one or more names of partitions, separated by commas. The partitions must already exist in the table to be altered. You can also use the `ALL` keyword in place of *partition\_names*, in which case the statement acts on all partitions in the table. For more information and examples, see [Section 17.3.4, “Maintenance of Partitions”](#).

Some MySQL storage engines, such as `InnoDB`, do not support per-partition optimization. For a partitioned table using such a storage engine, `ALTER TABLE ... OPTIMIZE PARTITION` rebuilds the entire table. This is a known issue. (Bug #42822) Beginning with MySQL 5.6.9, running this statement on such a table causes the entire table to be rebuilt and analyzed, and an appropriate warning to be issued. (Bug #11751825)

To work around this problem, use the statements `ALTER TABLE ... REBUILD PARTITION` and `ALTER TABLE ... ANALYZE PARTITION` instead.

The `ANALYZE PARTITION`, `CHECK PARTITION`, `OPTIMIZE PARTITION`, and `REPAIR PARTITION` options are not permitted for tables which are not partitioned.

- `REMOVE PARTITIONING` enables you to remove a table's partitioning without otherwise affecting the table or its data. This option can be combined with other `ALTER TABLE` options such as those used to add, drop, or rename drop columns or indexes.
- Using the `ENGINE` option with `ALTER TABLE` changes the storage engine used by the table without affecting the partitioning.

Prior to MySQL 5.6.6, when `ALTER TABLE ... EXCHANGE PARTITION` or `ALTER TABLE ... TRUNCATE PARTITION` was run against a partitioned table that used `MyISAM` (or another storage engine that makes use of table-level locking), the entire partitioned table was locked; in MySQL 5.6.6 and later, in such cases, only those partitions that are actually read from are locked. This did not (and does not) affect partitioned tables using a storage engine—such as `InnoDB`—that employs row-level locking. See [Section 17.6.4, “Partitioning and Locking”](#).

It is possible for an `ALTER TABLE` statement to contain a `PARTITION BY` or `REMOVE PARTITION` clause in an addition to other alter specifications, but the `PARTITION BY` or `REMOVE PARTITION` clause must be specified last after any other specifications.

The `ADD PARTITION`, `DROP PARTITION`, `COALESCE PARTITION`, `REORGANIZE PARTITION`, `ANALYZE PARTITION`, `CHECK PARTITION`, and `REPAIR PARTITION` options cannot be combined with other alter specifications in a single `ALTER TABLE`, since the options just listed act on individual partitions. For more information, see [Section 13.1.6.1, “ALTER TABLE Partition Operations”](#).

Only a single instance of any one of the following options can be used in a given `ALTER TABLE` statement: `PARTITION BY`, `ADD PARTITION`, `DROP PARTITION`, `TRUNCATE PARTITION`, `EXCHANGE PARTITION`, `REORGANIZE PARTITION`, or `COALESCE PARTITION`, `ANALYZE PARTITION`, `CHECK PARTITION`, `OPTIMIZE PARTITION`, `REBUILD PARTITION`, `REMOVE PARTITIONING`.

For example, the following two statements are invalid:

```
ALTER TABLE t1 ANALYZE PARTITION p1, ANALYZE PARTITION p2;
ALTER TABLE t1 ANALYZE PARTITION p1, CHECK PARTITION p2;
```

In the first case, you can analyze partitions `p1` and `p2` of table `t1` concurrently using a single statement with a single `ANALYZE PARTITION` option that lists both of the partitions to be analyzed, like this:

```
ALTER TABLE t1 ANALYZE PARTITION p1, p2;
```

In the second case, it is not possible to perform `ANALYZE` and `CHECK` operations on different partitions of the same table concurrently. Instead, you must issue two separate statements, like this:

```
ALTER TABLE t1 ANALYZE PARTITION p1;
ALTER TABLE t1 CHECK PARTITION p2;
```

### 13.1.6.2. ALTER TABLE Examples

Begin with a table `t1` that is created as shown here:

```
CREATE TABLE t1 (a INTEGER,b CHAR(10));
```

To rename the table from `t1` to `t2`:

```
ALTER TABLE t1 RENAME t2;
```

To change column `a` from `INTEGER` to `TINYINT NOT NULL` (leaving the name the same), and to change column `b` from `CHAR(10)` to `CHAR(20)` as well as renaming it from `b` to `c`:

```
ALTER TABLE t2 MODIFY a TINYINT NOT NULL, CHANGE b c CHAR(20);
```

To add a new `TIMESTAMP` column named `d`:

```
ALTER TABLE t2 ADD d TIMESTAMP;
```

To add an index on column `d` and a `UNIQUE` index on column `a`:

```
ALTER TABLE t2 ADD INDEX (d), ADD UNIQUE (a);
```

To remove column `c`:

```
ALTER TABLE t2 DROP COLUMN c;
```

To add a new `AUTO_INCREMENT` integer column named `c`:

```
ALTER TABLE t2 ADD c INT UNSIGNED NOT NULL AUTO_INCREMENT,
ADD PRIMARY KEY (c);
```

We indexed `c` (as a `PRIMARY KEY`) because `AUTO_INCREMENT` columns must be indexed, and we declare `c` as `NOT NULL` because primary key columns cannot be `NULL`.

When you add an `AUTO_INCREMENT` column, column values are filled in with sequence numbers automatically. For `MyISAM` tables, you can set the first sequence number by executing `SET INSERT_ID=value` before `ALTER TABLE` or by using the `AUTO_INCREMENT=value` table option. See [Section 5.1.4, “Server System Variables”](#).

With `MyISAM` tables, if you do not change the `AUTO_INCREMENT` column, the sequence number is not affected. If you drop an `AUTO_INCREMENT` column and then add another `AUTO_INCREMENT` column, the numbers are resequenced beginning with 1.

When replication is used, adding an `AUTO_INCREMENT` column to a table might not produce the same ordering of the rows on the slave and the master. This occurs because the order in which the rows are numbered depends on the specific storage engine used for the table and the order in which the rows were inserted. If it is important to have the same order on the master and slave, the rows must be ordered before assigning an `AUTO_INCREMENT` number. Assuming that you want to add an `AUTO_INCREMENT` column to the table `t1`, the following statements produce a new table `t2` identical to `t1` but with an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 (id INT AUTO_INCREMENT PRIMARY KEY)
SELECT * FROM t1 ORDER BY col1, col2;
```

This assumes that the table `t1` has columns `col1` and `col2`.

This set of statements will also produce a new table `t2` identical to `t1`, with the addition of an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 LIKE t1;
ALTER TABLE t2 ADD id INT AUTO_INCREMENT PRIMARY KEY;
INSERT INTO t2 SELECT * FROM t1 ORDER BY col1, col2;
```



### Important

To guarantee the same ordering on both master and slave, *all* columns of `t1` must be referenced in the `ORDER BY` clause.

Regardless of the method used to create and populate the copy having the `AUTO_INCREMENT` column, the final step is to drop the original table and then rename the copy:

```
DROP t1;
ALTER TABLE t2 RENAME t1;
```

## 13.1.7. ALTER VIEW Syntax

```
ALTER
[ALGORITHM = {UNDEFINED | MERGE | TEMPTABLE}]
[DEFINER = { user | CURRENT_USER }]
[SQL SECURITY { DEFINER | INVOKER }]
VIEW view_name [(column_list)]
```

```
AS select_statement
[WITH [CASCADED | LOCAL] CHECK OPTION]
```

This statement changes the definition of a view, which must exist. The syntax is similar to that for [CREATE VIEW](#) and the effect is the same as for [CREATE OR REPLACE VIEW](#). See [Section 13.1.16, “CREATE VIEW Syntax”](#). This statement requires the [CREATE VIEW \[754\]](#) and [DROP \[754\]](#) privileges for the view, and some privilege for each column referred to in the [SELECT](#) statement. [ALTER VIEW](#) is permitted only to the definer or users with the [SUPER \[756\]](#) privilege.

### 13.1.8. CREATE DATABASE Syntax

```
CREATE {DATABASE | SCHEMA} [IF NOT EXISTS] db_name
  [create_specification] ...
```

```
create_specification:
  [DEFAULT] CHARACTER SET [=] charset_name
  | [DEFAULT] COLLATE [=] collation_name
```

[CREATE DATABASE](#) creates a database with the given name. To use this statement, you need the [CREATE \[754\]](#) privilege for the database. [CREATE SCHEMA](#) is a synonym for [CREATE DATABASE](#).

An error occurs if the database exists and you did not specify [IF NOT EXISTS](#).

In MySQL 5.6, [CREATE DATABASE](#) is not permitted within a session that has an active [LOCK TABLES](#) statement.

*create\_specification* options specify database characteristics. Database characteristics are stored in the `db.opt` file in the database directory. The [CHARACTER SET](#) clause specifies the default database character set. The [COLLATE](#) clause specifies the default database collation. [Section 10.1, “Character Set Support”](#), discusses character set and collation names.

A database in MySQL is implemented as a directory containing files that correspond to tables in the database. Because there are no tables in a database when it is initially created, the [CREATE DATABASE](#) statement creates only a directory under the MySQL data directory and the `db.opt` file. Rules for permissible database names are given in [Section 9.2, “Schema Object Names”](#). If a database name contains special characters, the name for the database directory contains encoded versions of those characters as described in [Section 9.2.3, “Mapping of Identifiers to File Names”](#).

If you manually create a directory under the data directory (for example, with `mkdir`), the server considers it a database directory and it shows up in the output of [SHOW DATABASES](#).

You can also use the `mysqladmin` program to create databases. See [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#).

### 13.1.9. CREATE EVENT Syntax

```
CREATE
  [DEFINER = { user | CURRENT_USER }]
  EVENT
  [IF NOT EXISTS]
  event_name
  ON SCHEDULE schedule
  [ON COMPLETION [NOT] PRESERVE]
  [ENABLE | DISABLE | DISABLE ON SLAVE]
  [COMMENT 'comment']
  DO event_body;

schedule:
  AT timestamp [+ INTERVAL interval] ...
  | EVERY interval
  [STARTS timestamp [+ INTERVAL interval] ...]
  [ENDS timestamp [+ INTERVAL interval] ...]

interval:
  quantity {YEAR | QUARTER | MONTH | DAY | HOUR | MINUTE |
            WEEK | SECOND | YEAR_MONTH | DAY_HOUR | DAY_MINUTE |
```

```
DAY_SECOND | HOUR_MINUTE | HOUR_SECOND | MINUTE_SECOND}
```

This statement creates and schedules a new event. The event will not run unless the Event Scheduler is enabled. For information about checking Event Scheduler status and enabling it if necessary, see [Section 18.4.2, “Event Scheduler Configuration”](#).

`CREATE EVENT` requires the `EVENT [754]` privilege for the schema in which the event is to be created. It might also require the `SUPER [756]` privilege, depending on the `DEFINER` value, as described later in this section.

The minimum requirements for a valid `CREATE EVENT` statement are as follows:

- The keywords `CREATE EVENT` plus an event name, which uniquely identifies the event in a database schema.
- An `ON SCHEDULE` clause, which determines when and how often the event executes.
- A `DO` clause, which contains the SQL statement to be executed by an event.

This is an example of a minimal `CREATE EVENT` statement:

```
CREATE EVENT myevent
  ON SCHEDULE AT CURRENT_TIMESTAMP + INTERVAL 1 HOUR
  DO
    UPDATE myschema.mytable SET mycol = mycol + 1;
```

The previous statement creates an event named `myevent`. This event executes once—one hour following its creation—by running an SQL statement that increments the value of the `myschema.mytable` table's `mycol` column by 1.

The `event_name` must be a valid MySQL identifier with a maximum length of 64 characters. Event names are not case sensitive, so you cannot have two events named `myevent` and `MyEvent` in the same schema. In general, the rules governing event names are the same as those for names of stored routines. See [Section 9.2, “Schema Object Names”](#).

An event is associated with a schema. If no schema is indicated as part of `event_name`, the default (current) schema is assumed. To create an event in a specific schema, qualify the event name with a schema using `schema_name.event_name` syntax.

The `DEFINER` clause specifies the MySQL account to be used when checking access privileges at event execution time. If a `user` value is given, it should be a MySQL account specified as `'user_name'@'host_name'` (the same format used in the `GRANT` statement), `CURRENT_USER [1267]`, or `CURRENT_USER() [1267]`. The default `DEFINER` value is the user who executes the `CREATE EVENT` statement. This is the same as specifying `DEFINER = CURRENT_USER` explicitly.

If you specify the `DEFINER` clause, these rules determine the valid `DEFINER` user values:

- If you do not have the `SUPER [756]` privilege, the only permitted `user` value is your own account, either specified literally or by using `CURRENT_USER [1267]`. You cannot set the definer to some other account.
- If you have the `SUPER [756]` privilege, you can specify any syntactically valid account name. If the account does not actually exist, a warning is generated.
- Although it is possible to create an event with a nonexistent `DEFINER` account, an error occurs at event execution time if the account does not exist.

For more information about event security, see [Section 18.6, “Access Control for Stored Programs and Views”](#).

Within an event, the `CURRENT_USER() [1267]` function returns the account used to check privileges at event execution time, which is the `DEFINER` user. For information about user auditing within events, see [Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”](#).



`IF NOT EXISTS` has the same meaning for `CREATE EVENT` as for `CREATE TABLE`: If an event named `event_name` already exists in the same schema, no action is taken, and no error results. (However, a warning is generated in such cases.)

The `ON SCHEDULE` clause determines when, how often, and for how long the `event_body` defined for the event repeats. This clause takes one of two forms:

- `AT timestamp` is used for a one-time event. It specifies that the event executes one time only at the date and time given by `timestamp`, which must include both the date and time, or must be an expression that resolves to a datetime value. You may use a value of either the `DATETIME` or `TIMESTAMP` type for this purpose. If the date is in the past, a warning occurs, as shown here:

```
mysql> SELECT NOW();
+-----+
| NOW() |
+-----+
| 2006-02-10 23:59:01 |
+-----+
1 row in set (0.04 sec)

mysql> CREATE EVENT e_totals
->     ON SCHEDULE AT '2006-02-10 23:59:00'
->     DO INSERT INTO test.totals VALUES (NOW());
Query OK, 0 rows affected, 1 warning (0.00 sec)

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Note
Code: 1588
Message: Event execution time is in the past and ON COMPLETION NOT
PRESERVE is set. The event was dropped immediately after
creation.
```

`CREATE EVENT` statements which are themselves invalid—for whatever reason—fail with an error.

You may use `CURRENT_TIMESTAMP` [1213] to specify the current date and time. In such a case, the event acts as soon as it is created.

To create an event which occurs at some point in the future relative to the current date and time—such as that expressed by the phrase “three weeks from now”—you can use the optional clause `+ INTERVAL interval`. The `interval` portion consists of two parts, a quantity and a unit of time, and follows the same syntax rules that govern intervals used in the `DATE_ADD()` [1213] function (see Section 12.7, “Date and Time Functions”). The units keywords are also the same, except that you cannot use any units involving microseconds when defining an event. With some interval types, complex time units may be used. For example, “two minutes and ten seconds” can be expressed as `+ INTERVAL '2:10' MINUTE_SECOND`.

You can also combine intervals. For example, `AT CURRENT_TIMESTAMP + INTERVAL 3 WEEK + INTERVAL 2 DAY` is equivalent to “three weeks and two days from now”. Each portion of such a clause must begin with `+ INTERVAL`.

- To repeat actions at a regular interval, use an `EVERY` clause. The `EVERY` keyword is followed by an `interval` as described in the previous discussion of the `AT` keyword. (`+ INTERVAL` is *not* used with `EVERY`.) For example, `EVERY 6 WEEK` means “every six weeks”.

Although `+ INTERVAL` clauses are not permitted in an `EVERY` clause, you can use the same complex time units permitted in a `+ INTERVAL`.

An `EVERY` clause may contain an optional `STARTS` clause. `STARTS` is followed by a `timestamp` value that indicates when the action should begin repeating, and may also use `+ INTERVAL interval` to specify an amount of time “from now”. For example, `EVERY 3 MONTH STARTS CURRENT_TIMESTAMP + INTERVAL 1 WEEK` means “every three months, beginning one week from now”. Similarly, you can express “every two weeks, beginning six hours and fifteen minutes from now” as `EVERY 2 WEEK STARTS CURRENT_TIMESTAMP + INTERVAL '6:15'`



`HOUR_MINUTE`. Not specifying `STARTS` is the same as using `STARTS CURRENT_TIMESTAMP`—that is, the action specified for the event begins repeating immediately upon creation of the event.

An `EVERY` clause may contain an optional `ENDS` clause. The `ENDS` keyword is followed by a *timestamp* value that tells MySQL when the event should stop repeating. You may also use `+ INTERVAL interval` with `ENDS`; for instance, `EVERY 12 HOUR STARTS CURRENT_TIMESTAMP + INTERVAL 30 MINUTE ENDS CURRENT_TIMESTAMP + INTERVAL 4 WEEK` is equivalent to “every twelve hours, beginning thirty minutes from now, and ending four weeks from now”. Not using `ENDS` means that the event continues executing indefinitely.

`ENDS` supports the same syntax for complex time units as `STARTS` does.

You may use `STARTS`, `ENDS`, both, or neither in an `EVERY` clause.

If a repeating event does not terminate within its scheduling interval, the result may be multiple instances of the event executing simultaneously. If this is undesirable, you should institute a mechanism to prevent simultaneous instances. For example, you could use the `GET_LOCK()` [1275] function, or row or table locking.

The `ON SCHEDULE` clause may use expressions involving built-in MySQL functions and user variables to obtain any of the *timestamp* or *interval* values which it contains. You may not use stored functions or user-defined functions in such expressions, nor may you use any table references; however, you may use `SELECT FROM DUAL`. This is true for both `CREATE EVENT` and `ALTER EVENT` statements. References to stored functions, user-defined functions, and tables in such cases are specifically not permitted, and fail with an error (see Bug #22830).

Times in the `ON SCHEDULE` clause are interpreted using the current session `time_zone` [557] value. This becomes the event time zone; that is, the time zone that is used for event scheduling and is in effect within the event as it executes. These times are converted to UTC and stored along with the event time zone in the `mysql.event` table. This enables event execution to proceed as defined regardless of any subsequent changes to the server time zone or daylight saving time effects. For additional information about representation of event times, see Section 18.4.4, “Event Metadata”. See also Section 13.7.5.19, “SHOW EVENTS Syntax”, and Section 19.7, “The INFORMATION\_SCHEMA EVENTS Table”.

Normally, once an event has expired, it is immediately dropped. You can override this behavior by specifying `ON COMPLETION PRESERVE`. Using `ON COMPLETION NOT PRESERVE` merely makes the default nonpersistent behavior explicit.

You can create an event but prevent it from being active using the `DISABLE` keyword. Alternatively, you can use `ENABLE` to make explicit the default status, which is active. This is most useful in conjunction with `ALTER EVENT` (see Section 13.1.2, “ALTER EVENT Syntax”).

A third value may also appear in place of `ENABLED` or `DISABLED`; `DISABLE ON SLAVE` is set for the status of an event on a replication slave to indicate that the event was created on the master and replicated to the slave, but is not executed on the slave. See Section 16.4.1.10, “Replication of Invoked Features”.

You may supply a comment for an event using a `COMMENT` clause. *comment* may be any string of up to 64 characters that you wish to use for describing the event. The comment text, being a string literal, must be surrounded by quotation marks.

The `DO` clause specifies an action carried by the event, and consists of an SQL statement. Nearly any valid MySQL statement that can be used in a stored routine can also be used as the action statement for a scheduled event. (See Section E.1, “Restrictions on Stored Programs”.) For example, the following event `e_hourly` deletes all rows from the `sessions` table once per hour, where this table is part of the `site_activity` schema:

```
CREATE EVENT e_hourly
ON SCHEDULE
EVERY 1 HOUR
```

```
COMMENT 'Clears out sessions table each hour.'
DO
DELETE FROM site_activity.sessions;
```

MySQL stores the `sql_mode` [545] system variable setting that is in effect at the time an event is created, and always executes the event with this setting in force, *regardless of the current server SQL mode*.

A `CREATE EVENT` statement that contains an `ALTER EVENT` statement in its `DO` clause appears to succeed; however, when the server attempts to execute the resulting scheduled event, the execution fails with an error.



### Note

Statements such as `SELECT` or `SHOW` that merely return a result set have no effect when used in an event; the output from these is not sent to the MySQL Monitor, nor is it stored anywhere. However, you can use statements such as `SELECT ... INTO` and `INSERT INTO ... SELECT` that store a result. (See the next example in this section for an instance of the latter.)

The schema to which an event belongs is the default schema for table references in the `DO` clause. Any references to tables in other schemas must be qualified with the proper schema name.

As with stored routines, you can use compound-statement syntax in the `DO` clause by using the `BEGIN` and `END` keywords, as shown here:

```
delimiter |
CREATE EVENT e_daily
ON SCHEDULE
EVERY 1 DAY
COMMENT 'Saves total number of sessions then clears the table each day'
DO
BEGIN
INSERT INTO site_activity.totals (time, total)
SELECT CURRENT_TIMESTAMP, COUNT(*)
FROM site_activity.sessions;
DELETE FROM site_activity.sessions;
END |
delimiter ;
```

This example uses the `delimiter` command to change the statement delimiter. See [Section 18.1, “Defining Stored Programs”](#).

More complex compound statements, such as those used in stored routines, are possible in an event. This example uses local variables, an error handler, and a flow control construct:

```
delimiter |
CREATE EVENT e
ON SCHEDULE
EVERY 5 SECOND
DO
BEGIN
DECLARE v INTEGER;
DECLARE CONTINUE HANDLER FOR SQLEXCEPTION BEGIN END;

SET v = 0;

WHILE v < 5 DO
INSERT INTO t1 VALUES (0);
UPDATE t2 SET s1 = s1 + 1;
SET v = v + 1;
END WHILE;
END |
delimiter ;
```

There is no way to pass parameters directly to or from events; however, it is possible to invoke a stored routine with parameters within an event:

```
CREATE EVENT e_call_myproc
  ON SCHEDULE
    AT CURRENT_TIMESTAMP + INTERVAL 1 DAY
  DO CALL myproc(5, 27);
```

If an event's definer has the [SUPER \[756\]](#) privilege, the event can read and write global variables. As granting this privilege entails a potential for abuse, extreme care must be taken in doing so.

Generally, any statements that are valid in stored routines may be used for action statements executed by events. For more information about statements permissible within stored routines, see [Section 18.2.1, “Stored Routine Syntax”](#). You can create an event as part of a stored routine, but an event cannot be created by another event.

### 13.1.10. CREATE FUNCTION Syntax

The `CREATE FUNCTION` statement is used to create stored functions and user-defined functions (UDFs):

- For information about creating stored functions, see [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#).
- For information about creating user-defined functions, see [Section 13.7.3.1, “CREATE FUNCTION Syntax for User-Defined Functions”](#).

### 13.1.11. CREATE INDEX Syntax

```
CREATE [UNIQUE|FULLTEXT|SPATIAL] INDEX index_name
  [index_type]
  ON tbl_name (index_col_name,...)
  [algorithm_option | lock_option] ...

index_col_name:
  col_name [(length)] [ASC | DESC]

index_type:
  USING {BTREE | HASH}

index_option:
  KEY_BLOCK_SIZE [=] value
  | index_type
  | WITH PARSER parser_name
  | COMMENT 'string'

algorithm_option:
  ALGORITHM [=] {DEFAULT|INPLACE|COPY}

lock_option:
  LOCK [=] {DEFAULT|NONE|SHARED|EXCLUSIVE}
```

`CREATE INDEX` is mapped to an `ALTER TABLE` statement to create indexes. See [Section 13.1.6, “ALTER TABLE Syntax”](#). `CREATE INDEX` cannot be used to create a `PRIMARY KEY`; use `ALTER TABLE` instead. For more information about indexes, see [Section 8.3.1, “How MySQL Uses Indexes”](#).

Normally, you create all indexes on a table at the time the table itself is created with `CREATE TABLE`. See [Section 13.1.14, “CREATE TABLE Syntax”](#). This guideline is especially important for InnoDB tables, where the primary key determines the physical layout of rows in the data file. `CREATE INDEX` enables you to add indexes to existing tables.

A column list of the form `(col1,col2,...)` creates a multiple-column index. Index key values are formed by concatenating the values of the given columns.

Indexes can be created that use only the leading part of column values, using `col_name(length)` syntax to specify an index prefix length:

- Prefixes can be specified for `CHAR`, `VARCHAR`, `BINARY`, and `VARBINARY` columns.
- `BLOB` and `TEXT` columns also can be indexed, but a prefix length *must* be given.
- Prefix lengths are given in characters for nonbinary string types and in bytes for binary string types. That is, index entries consist of the first *length* characters of each column value for `CHAR`, `VARCHAR`, and `TEXT` columns, and the first *length* bytes of each column value for `BINARY`, `VARBINARY`, and `BLOB` columns.
- For spatial columns, prefix values cannot be given, as described later in this section.

The statement shown here creates an index using the first 10 characters of the `name` column:

```
CREATE INDEX part_of_name ON customer (name(10));
```

If names in the column usually differ in the first 10 characters, this index should not be much slower than an index created from the entire `name` column. Also, using column prefixes for indexes can make the index file much smaller, which could save a lot of disk space and might also speed up `INSERT` operations.

Prefix support and lengths of prefixes (where supported) are storage engine dependent. For example, a prefix can be up to 1000 bytes long for `MyISAM` tables, and 767 bytes for `InnoDB` tables.



#### Note

Prefix limits are measured in bytes, whereas the prefix length in `CREATE INDEX` statements is interpreted as number of characters for nonbinary data types (`CHAR`, `VARCHAR`, `TEXT`). Take this into account when specifying a prefix length for a column that uses a multi-byte character set.

A `UNIQUE` index creates a constraint such that all values in the index must be distinct. An error occurs if you try to add a new row with a key value that matches an existing row. For all engines, a `UNIQUE` index permits multiple `NULL` values for columns that can contain `NULL`. If you specify a prefix value for a column in a `UNIQUE` index, the column values must be unique within the prefix.

`FULLTEXT` indexes are supported only for `InnoDB` and `MyISAM` tables and can include only `CHAR`, `VARCHAR`, and `TEXT` columns. Indexing always happens over the entire column; column prefix indexing is not supported and any prefix length is ignored if specified. See [Section 12.9, “Full-Text Search Functions”](#), for details of operation.

The `MyISAM`, `InnoDB`, `NDB`, and `ARCHIVE` storage engines support spatial columns such as (`POINT` and `GEOMETRY`). ([Section 12.18, “Spatial Extensions”](#), describes the spatial data types.) However, support for spatial column indexing varies among engines. Spatial and nonspatial indexes are available according to the following rules.

Spatial indexes (created using `SPATIAL INDEX`) have these characteristics:

- Available only for `MyISAM` tables. Specifying `SPATIAL INDEX` for other storage engines results in an error.
- Indexed columns must be `NOT NULL`.
- In MySQL 5.6, column prefix lengths are prohibited. The full width of each column is indexed.

Characteristics of nonspatial indexes (created with `INDEX`, `UNIQUE`, or `PRIMARY KEY`):

- Permitted for any storage engine that supports spatial columns except `ARCHIVE`.
- Columns can be `NULL` unless the index is a primary key.
- For each spatial column in a non-`SPATIAL` index except `POINT` columns, a column prefix length must be specified. (This is the same requirement as for indexed `BLOB` columns.) The prefix length is given in bytes.

- The index type for a non-*SPATIAL* index depends on the storage engine. Currently, B-tree is used.

In MySQL 5.6:

- You can add an index on a column that can have *NULL* values only if you are using the *InnoDB*, *MyISAM*, or *MEMORY* storage engine.
- You can add an index on a *BLOB* or *TEXT* column only if you are using the *InnoDB* or *MyISAM* storage engine.
- When the `innodb_stats_persistent` [1747] setting is enabled, run the `ANALYZE TABLE` statement for an *InnoDB* table after creating an index on that table.

An *index\_col\_name* specification can end with *ASC* or *DESC*. These keywords are permitted for future extensions for specifying ascending or descending index value storage. Currently, they are parsed but ignored; index values are always stored in ascending order.

Following the index column list, index options can be given. An *index\_option* value can be any of the following:

- `KEY_BLOCK_SIZE [=] value`

For *compressed* *InnoDB* tables, optionally specifies the size in bytes to use for *pages*. The value is treated as a hint; a different size could be used if necessary. A value of 0 represents the default compressed page size. See [Section 5.4.6, “Working with InnoDB Compressed Tables”](#) for usage details.



#### Note

Oracle recommends enabling `innodb_strict_mode` [1749] when using the `KEY_BLOCK_SIZE` clause for *InnoDB* tables. See [Section 14.2.5.7, “InnoDB Strict Mode”](#) for details.

- *index\_type*

Some storage engines permit you to specify an index type when creating an index. The permissible index type values supported by different storage engines are shown in the following table. Where multiple index types are listed, the first one is the default when no index type specifier is given.

Storage Engine	Permissible Index Types
<i>InnoDB</i>	<i>BTREE</i>
<i>MyISAM</i>	<i>BTREE</i>
<i>MEMORY/HEAP</i>	<i>HASH</i> , <i>BTREE</i>
<i>NDB</i>	<i>HASH</i> , <i>BTREE</i> (see note in text)

Example:

```
CREATE TABLE lookup (id INT) ENGINE = MEMORY;
CREATE INDEX id_index ON lookup (id) USING BTREE;
```

The *index\_type* clause cannot be used together with *SPATIAL INDEX*.

If you specify an index type that is not valid for a given storage engine, but there is another index type available that the engine can use without affecting query results, the engine uses the available type. The parser recognizes *RTREE* as a type name, but currently this cannot be specified for any storage engine.

Use of this option before the `ON tbl_name` clause is deprecated; support for use of the option in this position will be removed in a future MySQL release. If an *index\_type* option is given in both the earlier and later positions, the final option applies.

`TYPE type_name` is recognized as a synonym for `USING type_name`. However, `USING` is the preferred form.

- `WITH PARSER parser_name`

This option can be used only with `FULLTEXT` indexes. It associates a parser plugin with the index if full-text indexing and searching operations need special handling. See [Section 22.2, “The MySQL Plugin API”](#), for details on creating plugins.

- `COMMENT 'string'`

Index definitions can include an optional comment of up to 1024 characters.

As of MySQL 5.6.6, the `ALGORITHM` and `LOCK` clauses may be given. These influence the table copying method and level of concurrency for reading and writing the table while its indexes are being modified. They have the same meaning as for the `ALTER TABLE` statement. For more information, see [Section 13.1.6, “ALTER TABLE Syntax”](#)

### 13.1.12. CREATE PROCEDURE and CREATE FUNCTION Syntax

```
CREATE
  [DEFINER = { user | CURRENT_USER }]
  PROCEDURE sp_name ([proc_parameter[,...]])
  [characteristic ...] routine_body

CREATE
  [DEFINER = { user | CURRENT_USER }]
  FUNCTION sp_name ([func_parameter[,...]])
  RETURNS type
  [characteristic ...] routine_body

proc_parameter:
  [ IN | OUT | INOUT ] param_name type

func_parameter:
  param_name type

type:
  Any valid MySQL data type

characteristic:
  COMMENT 'string'
  | LANGUAGE SQL
  | [NOT] DETERMINISTIC
  | { CONTAINS SQL | NO SQL | READS SQL DATA | MODIFIES SQL DATA }
  | SQL SECURITY { DEFINER | INVOKER }
```

These statements create stored routines. By default, a routine is associated with the default database. To associate the routine explicitly with a given database, specify the name as `db_name.sp_name` when you create it.

The `CREATE FUNCTION` statement is also used in MySQL to support UDFs (user-defined functions). See [Section 22.3, “Adding New Functions to MySQL”](#). A UDF can be regarded as an external stored function. Stored functions share their namespace with UDFs. See [Section 9.2.4, “Function Name Parsing and Resolution”](#), for the rules describing how the server interprets references to different kinds of functions.

To invoke a stored procedure, use the `CALL` statement (see [Section 13.2.1, “CALL Syntax”](#)). To invoke a stored function, refer to it in an expression. The function returns a value during expression evaluation.

`CREATE PROCEDURE` and `CREATE FUNCTION` require the `CREATE ROUTINE [754]` privilege. They might also require the `SUPER [756]` privilege, depending on the `DEFINER` value, as described later in

this section. If binary logging is enabled, `CREATE FUNCTION` might require the `SUPER` [756] privilege, as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

By default, MySQL automatically grants the `ALTER ROUTINE` [754] and `EXECUTE` [754] privileges to the routine creator. This behavior can be changed by disabling the `automatic_sp_privileges` [464] system variable. See [Section 18.2.2, “Stored Routines and MySQL Privileges”](#).

The `DEFINER` and `SQL SECURITY` clauses specify the security context to be used when checking access privileges at routine execution time, as described later in this section.

If the routine name is the same as the name of a built-in SQL function, a syntax error occurs unless you use a space between the name and the following parenthesis when defining the routine or invoking it later. For this reason, avoid using the names of existing SQL functions for your own stored routines.

The `IGNORE_SPACE` [604] SQL mode applies to built-in functions, not to stored routines. It is always permissible to have spaces after a stored routine name, regardless of whether `IGNORE_SPACE` [604] is enabled.

The parameter list enclosed within parentheses must always be present. If there are no parameters, an empty parameter list of `()` should be used. Parameter names are not case sensitive.

Each parameter is an `IN` parameter by default. To specify otherwise for a parameter, use the keyword `OUT` or `INOUT` before the parameter name.



#### Note

Specifying a parameter as `IN`, `OUT`, or `INOUT` is valid only for a `PROCEDURE`. For a `FUNCTION`, parameters are always regarded as `IN` parameters.

An `IN` parameter passes a value into a procedure. The procedure might modify the value, but the modification is not visible to the caller when the procedure returns. An `OUT` parameter passes a value from the procedure back to the caller. Its initial value is `NULL` within the procedure, and its value is visible to the caller when the procedure returns. An `INOUT` parameter is initialized by the caller, can be modified by the procedure, and any change made by the procedure is visible to the caller when the procedure returns.

For each `OUT` or `INOUT` parameter, pass a user-defined variable in the `CALL` statement that invokes the procedure so that you can obtain its value when the procedure returns. If you are calling the procedure from within another stored procedure or function, you can also pass a routine parameter or local routine variable as an `IN` or `INOUT` parameter.

Routine parameters cannot be referenced in statements prepared within the routine; see [Section E.1, “Restrictions on Stored Programs”](#).

The following example shows a simple stored procedure that uses an `OUT` parameter:

```
mysql> delimiter //
mysql> CREATE PROCEDURE simpleproc (OUT param1 INT)
-> BEGIN
->   SELECT COUNT(*) INTO param1 FROM t;
-> END//
Query OK, 0 rows affected (0.00 sec)

mysql> delimiter ;

mysql> CALL simpleproc(@a);
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @a;
+-----+
| @a    |
+-----+
| 3     |
```



```
+-----+
1 row in set (0.00 sec)
```

The example uses the `mysql` client `delimiter` command to change the statement delimiter from `;` to `//` while the procedure is being defined. This enables the `;` delimiter used in the procedure body to be passed through to the server rather than being interpreted by `mysql` itself. See [Section 18.1, “Defining Stored Programs”](#).

The `RETURNS` clause may be specified only for a `FUNCTION`, for which it is mandatory. It indicates the return type of the function, and the function body must contain a `RETURN value` statement. If the `RETURN` statement returns a value of a different type, the value is coerced to the proper type. For example, if a function specifies an `ENUM` or `SET` value in the `RETURNS` clause, but the `RETURN` statement returns an integer, the value returned from the function is the string for the corresponding `ENUM` member of set of `SET` members.

The following example function takes a parameter, performs an operation using an SQL function, and returns the result. In this case, it is unnecessary to use `delimiter` because the function definition contains no internal `;` statement delimiters:

```
mysql> CREATE FUNCTION hello (s CHAR(20))
mysql> RETURNS CHAR(50) DETERMINISTIC
-> RETURN CONCAT('Hello, ',s,'!');
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT hello('world');
+-----+
| hello('world') |
+-----+
| Hello, world!  |
+-----+
1 row in set (0.00 sec)
```

Parameter types and function return types can be declared to use any valid data type. The `COLLATE` attribute can be used if preceded by the `CHARACTER SET` attribute.

The `routine_body` consists of a valid SQL routine statement. This can be a simple statement such as `SELECT` or `INSERT`, or a compound statement written using `BEGIN` and `END`. Compound statements can contain declarations, loops, and other control structure statements. The syntax for these statements is described in [Section 13.6, “MySQL Compound-Statement Syntax”](#).

MySQL permits routines to contain DDL statements, such as `CREATE` and `DROP`. MySQL also permits stored procedures (but not stored functions) to contain SQL transaction statements such as `COMMIT`. Stored functions may not contain statements that perform explicit or implicit commit or rollback. Support for these statements is not required by the SQL standard, which states that each DBMS vendor may decide whether to permit them.

Statements that return a result set can be used within a stored procedure but not within a stored function. This prohibition includes `SELECT` statements that do not have an `INTO var_list` clause and other statements such as `SHOW`, `EXPLAIN`, and `CHECK TABLE`. For statements that can be determined at function definition time to return a result set, a `Not allowed to return a result set from a function` error occurs (`ER_SP_NO_RETSET [3687]`). For statements that can be determined only at runtime to return a result set, a `PROCEDURE %s can't return a result set in the given context` error occurs (`ER_SP_BADSELECT [3680]`).

`USE` statements within stored routines are not permitted. When a routine is invoked, an implicit `USE db_name` is performed (and undone when the routine terminates). This causes the routine to have the given default database while it executes. References to objects in databases other than the routine default database should be qualified with the appropriate database name.

For additional information about statements that are not permitted in stored routines, see [Section E.1, “Restrictions on Stored Programs”](#).

For information about invoking stored procedures from within programs written in a language that has a MySQL interface, see [Section 13.2.1, “CALL Syntax”](#).



MySQL stores the `sql_mode` [545] system variable setting that is in effect at the time a routine is created, and always executes the routine with this setting in force, *regardless of the server SQL mode in effect when the routine is invoked*.

The switch from the SQL mode of the invoker to that of the routine occurs after evaluation of arguments and assignment of the resulting values to routine parameters. If you define a routine in strict SQL mode but invoke it in nonstrict mode, assignment of arguments to routine parameters does not take place in strict mode. If you require that expressions passed to a routine be assigned in strict SQL mode, you should invoke the routine with strict mode in effect.

The `COMMENT` characteristic is a MySQL extension, and may be used to describe the stored routine. This information is displayed by the `SHOW CREATE PROCEDURE` and `SHOW CREATE FUNCTION` statements.

The `LANGUAGE` characteristic indicates the language in which the routine is written. The server ignores this characteristic; only SQL routines are supported.

A routine is considered “deterministic” if it always produces the same result for the same input parameters, and “not deterministic” otherwise. If neither `DETERMINISTIC` nor `NOT DETERMINISTIC` is given in the routine definition, the default is `NOT DETERMINISTIC`. To declare that a function is deterministic, you must specify `DETERMINISTIC` explicitly.

Assessment of the nature of a routine is based on the “honesty” of the creator: MySQL does not check that a routine declared `DETERMINISTIC` is free of statements that produce nondeterministic results. However, misdeclaring a routine might affect results or affect performance. Declaring a nondeterministic routine as `DETERMINISTIC` might lead to unexpected results by causing the optimizer to make incorrect execution plan choices. Declaring a deterministic routine as `NONDETERMINISTIC` might diminish performance by causing available optimizations not to be used.

If binary logging is enabled, the `DETERMINISTIC` characteristic affects which routine definitions MySQL accepts. See [Section 18.7, “Binary Logging of Stored Programs”](#).

A routine that contains the `NOW()` [1220] function (or its synonyms) or `RAND()` [1206] is nondeterministic, but it might still be replication-safe. For `NOW()` [1220], the binary log includes the timestamp and replicates correctly. `RAND()` [1206] also replicates correctly as long as it is called only a single time during the execution of a routine. (You can consider the routine execution timestamp and random number seed as implicit inputs that are identical on the master and slave.)

Several characteristics provide information about the nature of data use by the routine. In MySQL, these characteristics are advisory only. The server does not use them to constrain what kinds of statements a routine will be permitted to execute.

- `CONTAINS SQL` indicates that the routine does not contain statements that read or write data. This is the default if none of these characteristics is given explicitly. Examples of such statements are `SET @x = 1` or `DO RELEASE_LOCK('abc')`, which execute but neither read nor write data.
- `NO SQL` indicates that the routine contains no SQL statements.
- `READS SQL DATA` indicates that the routine contains statements that read data (for example, `SELECT`), but not statements that write data.
- `MODIFIES SQL DATA` indicates that the routine contains statements that may write data (for example, `INSERT` or `DELETE`).

The `SQL SECURITY` characteristic can be `DEFINER` or `INVOKER` to specify the security context; that is, whether the routine executes using the privileges of the account named in the routine `DEFINER` clause or the user who invokes it. This account must have permission to access the database with which the routine is associated. The default value is `DEFINER`. The user who invokes the routine must have the `EXECUTE` [754] privilege for it, as must the `DEFINER` account if the routine executes in definer security context.

The `DEFINER` clause specifies the MySQL account to be used when checking access privileges at routine execution time for routines that have the `SQL SECURITY DEFINER` characteristic.

If a `user` value is given for the `DEFINER` clause, it should be a MySQL account specified as `'user_name'@'host_name'` (the same format used in the `GRANT` statement), `CURRENT_USER` [1267], or `CURRENT_USER()` [1267]. The default `DEFINER` value is the user who executes the `CREATE PROCEDURE` or `CREATE FUNCTION` or statement. This is the same as specifying `DEFINER = CURRENT_USER` explicitly.

If you specify the `DEFINER` clause, these rules determine the valid `DEFINER` user values:

- If you do not have the `SUPER` [756] privilege, the only permitted `user` value is your own account, either specified literally or by using `CURRENT_USER` [1267]. You cannot set the definer to some other account.
- If you have the `SUPER` [756] privilege, you can specify any syntactically valid account name. If the account does not actually exist, a warning is generated.
- Although it is possible to create a routine with a nonexistent `DEFINER` account, an error occurs at routine execution time if the `SQL SECURITY` value is `DEFINER` but the definer account does not exist.

For more information about stored routine security, see [Section 18.6, “Access Control for Stored Programs and Views”](#).

Within a stored routine that is defined with the `SQL SECURITY DEFINER` characteristic, `CURRENT_USER` [1267] returns the routine's `DEFINER` value. For information about user auditing within stored routines, see [Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”](#).

Consider the following procedure, which displays a count of the number of MySQL accounts listed in the `mysql.user` table:

```
CREATE DEFINER = 'admin'@'localhost' PROCEDURE account_count()
BEGIN
  SELECT 'Number of accounts:', COUNT(*) FROM mysql.user;
END;
```

The procedure is assigned a `DEFINER` account of `'admin'@'localhost'` no matter which user defines it. It executes with the privileges of that account no matter which user invokes it (because the default security characteristic is `DEFINER`). The procedure succeeds or fails depending on whether invoker has the `EXECUTE` [754] privilege for it and `'admin'@'localhost'` has the `SELECT` [755] privilege for the `mysql.user` table.

Now suppose that the procedure is defined with the `SQL SECURITY INVOKER` characteristic:

```
CREATE DEFINER = 'admin'@'localhost' PROCEDURE account_count()
SQL SECURITY INVOKER
BEGIN
  SELECT 'Number of accounts:', COUNT(*) FROM mysql.user;
END;
```

The procedure still has a `DEFINER` of `'admin'@'localhost'`, but in this case, it executes with the privileges of the invoking user. Thus, the procedure succeeds or fails depending on whether the invoker has the `EXECUTE` [754] privilege for it and the `SELECT` [755] privilege for the `mysql.user` table.

The server handles the data type of a routine parameter, local routine variable created with `DECLARE`, or function return value as follows:

- Assignments are checked for data type mismatches and overflow. Conversion and overflow problems result in warnings, or errors in strict SQL mode.
- Only scalar values can be assigned. For example, a statement such as `SET x = (SELECT 1, 2)` is invalid.

- For character data types, if there is a `CHARACTER SET` attribute in the declaration, the specified character set and its default collation is used. If the `COLLATE` attribute is also present, that collation is used rather than the default collation. If there is no `CHARACTER SET` attribute, the database character set and collation in effect at routine creation time are used. (The database character set and collation are given by the value of the `character_set_database` [468] and `collation_database` [469] system variables.)

If you change the database default character set or collation, stored routines that use the database defaults must be dropped and recreated so that they use the new defaults.

### 13.1.13. CREATE SERVER Syntax

```
CREATE SERVER server_name
  FOREIGN DATA WRAPPER wrapper_name
  OPTIONS (option [, option] ...)

option:
{ HOST character-literal
  | DATABASE character-literal
  | USER character-literal
  | PASSWORD character-literal
  | SOCKET character-literal
  | OWNER character-literal
  | PORT numeric-literal }
```

This statement creates the definition of a server for use with the `FEDERATED` storage engine. The `CREATE SERVER` statement creates a new row in the `servers` table in the `mysql` database. This statement requires the `SUPER` [756] privilege.

The `server_name` should be a unique reference to the server. Server definitions are global within the scope of the server, it is not possible to qualify the server definition to a specific database. `server_name` has a maximum length of 64 characters (names longer than 64 characters are silently truncated), and is case insensitive. You may specify the name as a quoted string.

The `wrapper_name` should be `mysql`, and may be quoted with single quotation marks. Other values for `wrapper_name` are not currently supported.

For each `option` you must specify either a character literal or numeric literal. Character literals are UTF-8, support a maximum length of 64 characters and default to a blank (empty) string. String literals are silently truncated to 64 characters. Numeric literals must be a number between 0 and 9999, default value is 0.



#### Note

The `OWNER` option is currently not applied, and has no effect on the ownership or operation of the server connection that is created.

The `CREATE SERVER` statement creates an entry in the `mysql.servers` table that can later be used with the `CREATE TABLE` statement when creating a `FEDERATED` table. The options that you specify will be used to populate the columns in the `mysql.servers` table. The table columns are `Server_name`, `Host`, `Db`, `Username`, `Password`, `Port` and `Socket`.

For example:

```
CREATE SERVER s
  FOREIGN DATA WRAPPER mysql
  OPTIONS (USER 'Remote', HOST '192.168.1.106', DATABASE 'test');
```

Be sure to specify all options necessary to establish a connection to the server. The user name, host name, and database name are mandatory. Other options might be required as well, such as password.

The data stored in the table can be used when creating a connection to a `FEDERATED` table:

```
CREATE TABLE t (s1 INT) ENGINE=FEDERATED CONNECTION='s';
```

For more information, see [Section 14.9, “The FEDERATED Storage Engine”](#).

`CREATE SERVER` causes an automatic commit.

In MySQL 5.6, `CREATE SERVER` is not written to the binary log, regardless of the logging format that is in use.

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

### 13.1.14. CREATE TABLE Syntax

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  (create_definition,...)
  [table_options]
  [partition_options]
```

Or:

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  [(create_definition,...)]
  [table_options]
  [partition_options]
  select_statement
```

Or:

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  { LIKE old_tbl_name | (LIKE old_tbl_name) }
```

```
create_definition:
  col_name column_definition
  | [CONSTRAINT [symbol]] PRIMARY KEY [index_type] (index_col_name,...)
    [index_option] ...
  | {INDEX|KEY} [index_name] [index_type] (index_col_name,...)
    [index_option] ...
  | [CONSTRAINT [symbol]] UNIQUE [INDEX|KEY]
    [index_name] [index_type] (index_col_name,...)
    [index_option] ...
  | {FULLTEXT|SPATIAL} [INDEX|KEY] [index_name] (index_col_name,...)
    [index_option] ...
  | [CONSTRAINT [symbol]] FOREIGN KEY
    [index_name] (index_col_name,...) reference_definition
  | CHECK (expr)
```

```
column_definition:
  data_type [NOT NULL | NULL] [DEFAULT default_value]
  [AUTO_INCREMENT] [UNIQUE [KEY] | [PRIMARY] KEY]
  [COMMENT 'string']
  [COLUMN_FORMAT {FIXED|DYNAMIC|DEFAULT}]
  [reference_definition]
```

```
data_type:
  BIT[(length)]
  | TINYINT[(length)] [UNSIGNED] [ZEROFILL]
  | SMALLINT[(length)] [UNSIGNED] [ZEROFILL]
  | MEDIUMINT[(length)] [UNSIGNED] [ZEROFILL]
  | INT[(length)] [UNSIGNED] [ZEROFILL]
  | INTEGER[(length)] [UNSIGNED] [ZEROFILL]
  | BIGINT[(length)] [UNSIGNED] [ZEROFILL]
  | REAL[(length,decimals)] [UNSIGNED] [ZEROFILL]
  | DOUBLE[(length,decimals)] [UNSIGNED] [ZEROFILL]
  | FLOAT[(length,decimals)] [UNSIGNED] [ZEROFILL]
  | DECIMAL[(length[,decimals])] [UNSIGNED] [ZEROFILL]
  | NUMERIC[(length[,decimals])] [UNSIGNED] [ZEROFILL]
  | DATE
  | TIME
  | TIMESTAMP
  | DATETIME
  | YEAR
  | CHAR[(length)]
```

```

    [CHARACTER SET charset_name] [COLLATE collation_name]
| VARCHAR(length)
| [CHARACTER SET charset_name] [COLLATE collation_name]
| BINARY[(length)]
| VARBINARY(length)
| TINYBLOB
| BLOB
| MEDIUMBLOB
| LONGBLOB
| TINYTEXT [BINARY]
| [CHARACTER SET charset_name] [COLLATE collation_name]
| TEXT [BINARY]
| [CHARACTER SET charset_name] [COLLATE collation_name]
| MEDIUMTEXT [BINARY]
| [CHARACTER SET charset_name] [COLLATE collation_name]
| LONGTEXT [BINARY]
| [CHARACTER SET charset_name] [COLLATE collation_name]
| ENUM(value1,value2,value3,...)
| [CHARACTER SET charset_name] [COLLATE collation_name]
| SET(value1,value2,value3,...)
| [CHARACTER SET charset_name] [COLLATE collation_name]
| spatial_type

index_col_name:
    col_name [(length)] [ASC | DESC]

index_type:
    USING {BTREE | HASH}

index_option:
    KEY_BLOCK_SIZE [=] value
| index_type
| WITH PARSER parser_name
| COMMENT 'string'

reference_definition:
    REFERENCES tbl_name (index_col_name,...)
    [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
    [ON DELETE reference_option]
    [ON UPDATE reference_option]

reference_option:
    RESTRICT | CASCADE | SET NULL | NO ACTION

table_options:
    table_option [[,] table_option] ...

table_option:
    ENGINE [=] engine_name
| AUTO_INCREMENT [=] value
| AVG_ROW_LENGTH [=] value
| [DEFAULT] CHARACTER SET [=] charset_name
| CHECKSUM [=] {0 | 1}
| [DEFAULT] COLLATE [=] collation_name
| COMMENT [=] 'string'
| CONNECTION [=] 'connect_string'
| DATA DIRECTORY [=] 'absolute path to directory'
| DELAY_KEY_WRITE [=] {0 | 1}
| INDEX DIRECTORY [=] 'absolute path to directory'
| INSERT_METHOD [=] { NO | FIRST | LAST }
| KEY_BLOCK_SIZE [=] value
| MAX_ROWS [=] value
| MIN_ROWS [=] value
| PACK_KEYS [=] {0 | 1 | DEFAULT}
| PASSWORD [=] 'string'
| ROW_FORMAT [=] {DEFAULT|DYNAMIC|FIXED|COMPRESSED|REDUNDANT|COMPACT}
| STATS_AUTO_RECALC [=] {DEFAULT|0|1}
| STATS_PERSISTENT [=] {DEFAULT|0|1}
| UNION [=] (tbl_name [, tbl_name] ...)

partition_options:
    PARTITION BY

```

```

    { [[LINEAR] HASH(expr)
      | [[LINEAR] KEY [[ALGORITHM={1|2}] (column_list)
      | RANGE{(expr) | COLUMNS(column_list)}
      | LIST{(expr) | COLUMNS(column_list)} }
  [[PARTITIONS num]
  [[SUBPARTITION BY
    { [[LINEAR] HASH(expr)
      | [[LINEAR] KEY [[ALGORITHM={1|2}] (column_list) }
    [[SUBPARTITIONS num]
  ]
  [(partition_definition [, partition_definition] ...)]

partition_definition:
PARTITION partition_name
  [[VALUES
    {LESS THAN {(expr | value_list) | MAXVALUE}
    |
    IN (value_list)}]
  [[STORAGE] ENGINE [=] engine_name]
  [[COMMENT [=] 'comment_text' ]
  [[DATA DIRECTORY [=] 'data_dir' ]
  [[INDEX DIRECTORY [=] 'index_dir' ]
  [[MAX_ROWS [=] max_number_of_rows]
  [[MIN_ROWS [=] min_number_of_rows]
  [(subpartition_definition [, subpartition_definition] ...)]

subpartition_definition:
SUBPARTITION logical_name
  [[STORAGE] ENGINE [=] engine_name]
  [[COMMENT [=] 'comment_text' ]
  [[DATA DIRECTORY [=] 'data_dir' ]
  [[INDEX DIRECTORY [=] 'index_dir' ]
  [[MAX_ROWS [=] max_number_of_rows]
  [[MIN_ROWS [=] min_number_of_rows]

select_statement:
  [[IGNORE | REPLACE] [[AS] SELECT ... (Some valid select statement)]

```

`CREATE TABLE` creates a table with the given name. You must have the `CREATE [754]` privilege for the table.

Rules for permissible table names are given in [Section 9.2, “Schema Object Names”](#). By default, the table is created in the default database, using the `InnoDB` storage engine. An error occurs if the table exists, if there is no default database, or if the database does not exist.

The table name can be specified as `db_name.tbl_name` to create the table in a specific database. This works regardless of whether there is a default database, assuming that the database exists. If you use quoted identifiers, quote the database and table names separately. For example, write ``mydb`.`mytbl``, not ``mydb.mytbl``.

## Temporary Tables

You can use the `TEMPORARY` keyword when creating a table. A `TEMPORARY` table is visible only to the current connection, and is dropped automatically when the connection is closed. This means that two different connections can use the same temporary table name without conflicting with each other or with an existing non-`TEMPORARY` table of the same name. (The existing table is hidden until the temporary table is dropped.) To create temporary tables, you must have the `CREATE TEMPORARY TABLES [754]` privilege.



### Note

`CREATE TABLE` does not automatically commit the current active transaction if you use the `TEMPORARY` keyword.

## Existing Table with Same Name

The keywords `IF NOT EXISTS` prevent an error from occurring if the table exists. However, there is no verification that the existing table has a structure identical to that indicated by the `CREATE TABLE` statement.

## Physical Representation

MySQL represents each table by an `.frm` table format (definition) file in the database directory. The storage engine for the table might create other files as well.

For `InnoDB` tables, the file storage is controlled by the `innodb_file_per_table` [1718] configuration option. When this option is turned off, all `InnoDB` tables and indexes are stored in the `system tablespace`, represented by one or more `.ibd` files. For each `InnoDB` table created when this option is turned on, the table data and all associated indexes are stored in a `.ibd` file located inside the database directory.

For `MyISAM` tables, the storage engine creates data and index files. Thus, for each `MyISAM` table `tbl_name`, there are three disk files.

File	Purpose
<code>tbl_name.frm</code>	Table format (definition) file
<code>tbl_name.MYD</code>	Data file
<code>tbl_name.MYI</code>	Index file

Chapter 14, *Storage Engines*, describes what files each storage engine creates to represent tables. If a table name contains special characters, the names for the table files contain encoded versions of those characters as described in Section 9.2.3, “Mapping of Identifiers to File Names”.

## Data Types and Attributes for Columns

`data_type` represents the data type in a column definition. `spatial_type` represents a spatial data type. The data type syntax shown is representative only. For a full description of the syntax available for specifying column data types, as well as information about the properties of each type, see Chapter 11, *Data Types*, and Section 12.18, “Spatial Extensions”.

Some attributes do not apply to all data types. `AUTO_INCREMENT` applies only to integer and floating-point types. `DEFAULT` does not apply to the `BLOB` or `TEXT` types.

- If neither `NULL` nor `NOT NULL` is specified, the column is treated as though `NULL` had been specified.
- An integer or floating-point column can have the additional attribute `AUTO_INCREMENT`. When you insert a value of `NULL` (recommended) or `0` into an indexed `AUTO_INCREMENT` column, the column is set to the next sequence value. Typically this is `value+1`, where `value` is the largest value for the column currently in the table. `AUTO_INCREMENT` sequences begin with `1`.

To retrieve an `AUTO_INCREMENT` value after inserting a row, use the `LAST_INSERT_ID()` [1269] SQL function or the `mysql_insert_id()` C API function. See Section 12.14, “Information Functions”, and Section 21.8.6.37, “`mysql_insert_id()`”.

If the `NO_AUTO_VALUE_ON_ZERO` [604] SQL mode is enabled, you can store `0` in `AUTO_INCREMENT` columns as `0` without generating a new sequence value. See Section 5.1.7, “Server SQL Modes”.



### Note

There can be only one `AUTO_INCREMENT` column per table, it must be indexed, and it cannot have a `DEFAULT` value. An `AUTO_INCREMENT` column works properly only if it contains only positive values. Inserting a negative



number is regarded as inserting a very large positive number. This is done to avoid precision problems when numbers “wrap” over from positive to negative and also to ensure that you do not accidentally get an `AUTO_INCREMENT` column that contains 0.

For `MyISAM` tables, you can specify an `AUTO_INCREMENT` secondary column in a multiple-column key. See [Section 3.6.9, “Using AUTO\\_INCREMENT”](#).

To make MySQL compatible with some ODBC applications, you can find the `AUTO_INCREMENT` value for the last inserted row with the following query:

```
SELECT * FROM tbl_name WHERE auto_col IS NULL
```

For information about `InnoDB` and `AUTO_INCREMENT`, see [Section 5.4.4, “AUTO\\_INCREMENT Handling in InnoDB”](#). For information about `AUTO_INCREMENT` and MySQL Replication, see [Section 16.4.1.1, “Replication and AUTO\\_INCREMENT”](#).

- Character data types (`CHAR`, `VARCHAR`, `TEXT`) can include `CHARACTER SET` and `COLLATE` attributes to specify the character set and collation for the column. For details, see [Section 10.1, “Character Set Support”](#). `CHARSET` is a synonym for `CHARACTER SET`. Example:

```
CREATE TABLE t (c CHAR(20) CHARACTER SET utf8 COLLATE utf8_bin);
```

MySQL 5.6 interprets length specifications in character column definitions in characters. (Versions before MySQL 4.1 interpreted them in bytes.) Lengths for `BINARY` and `VARBINARY` are in bytes.

- The `DEFAULT` clause specifies a default value for a column. With one exception, the default value must be a constant; it cannot be a function or an expression. This means, for example, that you cannot set the default for a date column to be the value of a function such as `NOW()` [1220] or `CURRENT_DATE` [1212]. The exception is that you can specify `CURRENT_TIMESTAMP` [1213] as the default for a `TIMESTAMP` or (as of MySQL 5.6.5) `DATETIME` column. See [Section 11.3.5, “Automatic Initialization and Updating for TIMESTAMP and DATETIME”](#).

If a column definition includes no explicit `DEFAULT` value, MySQL determines the default value as described in [Section 11.5, “Data Type Default Values”](#).

`BLOB` and `TEXT` columns cannot be assigned a default value.

If the `NO_ZERO_DATE` [606] or `NO_ZERO_IN_DATE` [606] SQL mode is enabled and a date-valued default is not correct according to that mode, `CREATE TABLE` produces a warning if strict SQL mode is not enabled and an error if strict mode is enabled. For example, with `NO_ZERO_IN_DATE` [606] enabled, `c1 DATE DEFAULT '2010-00-00'` produces a warning. (Before MySQL 5.6.6, the statement produces an error even if strict mode is not enabled.)

- A comment for a column can be specified with the `COMMENT` option, up to 1024 characters long. The comment is displayed by the `SHOW CREATE TABLE` and `SHOW FULL COLUMNS` statements.
- `KEY` is normally a synonym for `INDEX`. The key attribute `PRIMARY KEY` can also be specified as just `KEY` when given in a column definition. This was implemented for compatibility with other database systems.
- A `UNIQUE` index creates a constraint such that all values in the index must be distinct. An error occurs if you try to add a new row with a key value that matches an existing row. For all engines, a `UNIQUE` index permits multiple `NULL` values for columns that can contain `NULL`.
- A `PRIMARY KEY` is a unique index where all key columns must be defined as `NOT NULL`. If they are not explicitly declared as `NOT NULL`, MySQL declares them so implicitly (and silently). A table can have only one `PRIMARY KEY`. The name of a `PRIMARY KEY` is always `PRIMARY`, which thus cannot be used as the name for any other kind of index.

If you do not have a `PRIMARY KEY` and an application asks for the `PRIMARY KEY` in your tables, MySQL returns the first `UNIQUE` index that has no `NULL` columns as the `PRIMARY KEY`.



In `InnoDB` tables, keep the `PRIMARY KEY` short to minimize storage overhead for secondary indexes. Each secondary index entry contains a copy of the primary key columns for the corresponding row. (See [Section 14.2.3.12](#), “`InnoDB Table and Index Structures`”.)

- In the created table, a `PRIMARY KEY` is placed first, followed by all `UNIQUE` indexes, and then the nonunique indexes. This helps the MySQL optimizer to prioritize which index to use and also more quickly to detect duplicated `UNIQUE` keys.
- A `PRIMARY KEY` can be a multiple-column index. However, you cannot create a multiple-column index using the `PRIMARY KEY` key attribute in a column specification. Doing so only marks that single column as primary. You must use a separate `PRIMARY KEY(index_col_name, ...)` clause.
- If a `PRIMARY KEY` or `UNIQUE` index consists of only one column that has an integer type, you can also refer to the column as `_rowid` in `SELECT` statements.
- In MySQL, the name of a `PRIMARY KEY` is `PRIMARY`. For other indexes, if you do not assign a name, the index is assigned the same name as the first indexed column, with an optional suffix (`_2, _3, ...`) to make it unique. You can see index names for a table using `SHOW INDEX FROM tbl_name`. See [Section 13.7.5.23](#), “`SHOW INDEX Syntax`”.
- Some storage engines permit you to specify an index type when creating an index. The syntax for the `index_type` specifier is `USING type_name`.

Example:

```
CREATE TABLE lookup
(id INT, INDEX USING BTREE (id))
ENGINE = MEMORY;
```

The preferred position for `USING` is after the index column list. It can be given before the column list, but support for use of the option in that position is deprecated and will be removed in a future MySQL release.

`index_option` values specify additional options for an index. `USING` is one such option. For details about permissible `index_option` values, see [Section 13.1.11](#), “`CREATE INDEX Syntax`”.

For more information about indexes, see [Section 8.3.1](#), “`How MySQL Uses Indexes`”.

- In MySQL 5.6, only the `InnoDB`, `MyISAM`, and `MEMORY` storage engines support indexes on columns that can have `NULL` values. In other cases, you must declare indexed columns as `NOT NULL` or an error results.
- For `CHAR`, `VARCHAR`, `BINARY`, and `VARBINARY` columns, indexes can be created that use only the leading part of column values, using `col_name(length)` syntax to specify an index prefix length. `BLOB` and `TEXT` columns also can be indexed, but a prefix length *must* be given. Prefix lengths are given in characters for nonbinary string types and in bytes for binary string types. That is, index entries consist of the first `length` characters of each column value for `CHAR`, `VARCHAR`, and `TEXT` columns, and the first `length` bytes of each column value for `BINARY`, `VARBINARY`, and `BLOB` columns. Indexing only a prefix of column values like this can make the index file much smaller. See [Section 8.3.4](#), “`Column Indexes`”.

Only the `InnoDB` and `MyISAM` storage engines support indexing on `BLOB` and `TEXT` columns. For example:

```
CREATE TABLE test (blob_col BLOB, INDEX(blob_col(10)));
```

Prefixes can be up to 1000 bytes long (767 bytes for `InnoDB` tables). Note that prefix limits are measured in bytes, whereas the prefix length in `CREATE TABLE` statements is interpreted as number of characters for nonbinary data types (`CHAR`, `VARCHAR`, `TEXT`). Take this into account when specifying a prefix length for a column that uses a multi-byte character set.

- An *index\_col\_name* specification can end with `ASC` or `DESC`. These keywords are permitted for future extensions for specifying ascending or descending index value storage. Currently, they are parsed but ignored; index values are always stored in ascending order.
- When you use `ORDER BY` or `GROUP BY` on a column in a `SELECT`, the server sorts values using only the initial number of bytes indicated by the `max_sort_length [508]` system variable.
- You can create special `FULLTEXT` indexes, which are used for full-text searches. Only the `InnoDB` and `MyISAM` storage engines support `FULLTEXT` indexes. They can be created only from `CHAR`, `VARCHAR`, and `TEXT` columns. Indexing always happens over the entire column; column prefix indexing is not supported and any prefix length is ignored if specified. See [Section 12.9, “Full-Text Search Functions”](#), for details of operation. A `WITH PARSER` clause can be specified as an *index\_option* value to associate a parser plugin with the index if full-text indexing and searching operations need special handling. This clause is valid only for `FULLTEXT` indexes. See [Section 22.2, “The MySQL Plugin API”](#), for details on creating plugins.
- You can create `SPATIAL` indexes on spatial data types. Spatial types are supported only for `MyISAM` tables and indexed columns must be declared as `NOT NULL`. See [Section 12.18, “Spatial Extensions”](#).
- In MySQL 5.6, index definitions can include an optional comment of up to 1024 characters.
- `InnoDB` tables support checking of foreign key constraints. The columns of the referenced table must always be explicitly named. Both `ON DELETE` and `ON UPDATE` actions on foreign keys. For more detailed information and examples, see [Section 13.1.14.2, “Using FOREIGN KEY Constraints”](#). For information specific to foreign keys in `InnoDB`, see [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

For other storage engines, MySQL Server parses and ignores the `FOREIGN KEY` and `REFERENCES` syntax in `CREATE TABLE` statements. The `CHECK` clause is parsed but ignored by all storage engines. See [Section 1.8.5.4, “Foreign Key Differences”](#).



### Important

For users familiar with the ANSI/ISO SQL Standard, please note that no storage engine, including `InnoDB`, recognizes or enforces the `MATCH` clause used in referential integrity constraint definitions. Use of an explicit `MATCH` clause will not have the specified effect, and also causes `ON DELETE` and `ON UPDATE` clauses to be ignored. For these reasons, specifying `MATCH` should be avoided.

The `MATCH` clause in the SQL standard controls how `NULL` values in a composite (multiple-column) foreign key are handled when comparing to a primary key. `InnoDB` essentially implements the semantics defined by `MATCH SIMPLE`, which permit a foreign key to be all or partially `NULL`. In that case, the (child table) row containing such a foreign key is permitted to be inserted, and does not match any row in the referenced (parent) table. It is possible to implement other semantics using triggers.

Additionally, MySQL requires that the referenced columns be indexed for performance. However, `InnoDB` does not enforce any requirement that the referenced columns be declared `UNIQUE` or `NOT NULL`. The handling of foreign key references to nonunique keys or keys that contain `NULL` values is not well defined for operations such as `UPDATE` or `DELETE CASCADE`. You are advised to use foreign keys that reference only keys that are both `UNIQUE` (or `PRIMARY`) and `NOT NULL`.

MySQL does not recognize or support “inline `REFERENCES` specifications” (as defined in the SQL standard) where the references are defined as part of

the column specification. MySQL accepts [REFERENCES](#) clauses only when specified as part of a separate [FOREIGN KEY](#) specification.



#### Note

Partitioned tables employing the [InnoDB](#) storage engine do not support foreign keys. See [Section 17.6, “Restrictions and Limitations on Partitioning”](#), for more information.

- There is a hard limit of 4096 columns per table, but the effective maximum may be less for a given table and depends on the factors discussed in [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

## Storage Engines

The [ENGINE](#) table option specifies the storage engine for the table, using one of the names shown in the following table. The engine name can be unquoted or quoted. The quoted name 'DEFAULT' is recognized but ignored.

Storage Engine	Description
<a href="#">InnoDB</a>	Transaction-safe tables with row locking and foreign keys. The default storage engine for new tables. See <a href="#">Section 14.2, “The InnoDB Storage Engine”</a> , and in particular <a href="#">Section 14.2.1.1, “InnoDB as the Default MySQL Storage Engine”</a> if you have MySQL experience but are new to <a href="#">InnoDB</a> .
<a href="#">MyISAM</a>	The binary portable storage engine that is primarily used for read-only or read-mostly workloads. See <a href="#">Section 14.3, “The MyISAM Storage Engine”</a> .
<a href="#">MEMORY</a>	The data for this storage engine is stored only in memory. See <a href="#">Section 14.4, “The MEMORY Storage Engine”</a> .
<a href="#">CSV</a>	Tables that store rows in comma-separated values format. See <a href="#">Section 14.5, “The CSV Storage Engine”</a> .
<a href="#">ARCHIVE</a>	The archiving storage engine. See <a href="#">Section 14.6, “The ARCHIVE Storage Engine”</a> .
<a href="#">EXAMPLE</a>	An example engine. See <a href="#">Section 14.10, “The EXAMPLE Storage Engine”</a> .
<a href="#">FEDERATED</a>	Storage engine that accesses remote tables. See <a href="#">Section 14.9, “The FEDERATED Storage Engine”</a> .
<a href="#">HEAP</a>	This is a synonym for <a href="#">MEMORY</a> .
<a href="#">MERGE</a>	A collection of <a href="#">MyISAM</a> tables used as one table. Also known as <a href="#">MRG_MyISAM</a> . See <a href="#">Section 14.8, “The MERGE Storage Engine”</a> .

If a storage engine is specified that is not available, MySQL uses the default engine instead. Normally, this is [MyISAM](#). For example, if a table definition includes the [ENGINE=INNODB](#) option but the MySQL server does not support [INNODB](#) tables, the table is created as a [MyISAM](#) table. This makes it possible to have a replication setup where you have transactional tables on the master but tables created on the slave are nontransactional (to get more speed). In MySQL 5.6, a warning occurs if the storage engine specification is not honored.

Engine substitution can be controlled by the setting of the [NO\\_ENGINE\\_SUBSTITUTION \[605\]](#) SQL mode, as described in [Section 5.1.7, “Server SQL Modes”](#).



#### Note

The older [TYPE](#) option that was synonymous with [ENGINE](#) was removed in MySQL 5.5. *When upgrading to MySQL 5.5 or later, you must convert existing applications that rely on [TYPE](#) to use [ENGINE](#) instead.*

## Optimizing Performance

The other table options are used to optimize the behavior of the table. In most cases, you do not have to specify any of them. These options apply to all storage engines unless otherwise indicated. Options that do not apply to a given storage engine may be accepted and remembered as part of the table definition. Such options then apply if you later use `ALTER TABLE` to convert the table to use a different storage engine.

- `AUTO_INCREMENT`

The initial `AUTO_INCREMENT` value for the table. In MySQL 5.6, this works for `MyISAM`, `MEMORY`, `InnoDB`, and `ARCHIVE` tables. To set the first auto-increment value for engines that do not support the `AUTO_INCREMENT` table option, insert a “dummy” row with a value one less than the desired value after creating the table, and then delete the dummy row.

For engines that support the `AUTO_INCREMENT` table option in `CREATE TABLE` statements, you can also use `ALTER TABLE tbl_name AUTO_INCREMENT = N` to reset the `AUTO_INCREMENT` value. The value cannot be set lower than the maximum value currently in the column.

- `AVG_ROW_LENGTH`

An approximation of the average row length for your table. You need to set this only for large tables with variable-size rows.

When you create a `MyISAM` table, MySQL uses the product of the `MAX_ROWS` and `AVG_ROW_LENGTH` options to decide how big the resulting table is. If you don't specify either option, the maximum size for `MyISAM` data and index files is 256TB by default. (If your operating system does not support files that large, table sizes are constrained by the file size limit.) If you want to keep down the pointer sizes to make the index smaller and faster and you don't really need big files, you can decrease the default pointer size by setting the `myisam_data_pointer_size` [511] system variable. (See [Section 5.1.4, “Server System Variables”](#).) If you want all your tables to be able to grow above the default limit and are willing to have your tables slightly slower and larger than necessary, you can increase the default pointer size by setting this variable. Setting the value to 7 permits table sizes up to 65,536TB.

- `[DEFAULT] CHARACTER SET`

Specify a default character set for the table. `CHARSET` is a synonym for `CHARACTER SET`. If the character set name is `DEFAULT`, the database character set is used.

- `CHECKSUM`

Set this to 1 if you want MySQL to maintain a live checksum for all rows (that is, a checksum that MySQL updates automatically as the table changes). This makes the table a little slower to update, but also makes it easier to find corrupted tables. The `CHECKSUM TABLE` statement reports the checksum. (`MyISAM` only.)

- `[DEFAULT] COLLATE`

Specify a default collation for the table.

- `COMMENT`

A comment for the table, up to 2048 characters long.

- `CONNECTION`

The connection string for a `FEDERATED` table.

**Note**

Older versions of MySQL used a `COMMENT` option for the connection string.

- `DATA DIRECTORY`, `INDEX DIRECTORY`

By using `DATA DIRECTORY='directory'`, you can specify where the InnoDB storage engine puts the `.ibd` tablespace file for a new table. This clause only applies when the `innodb_file_per_table` [1718] configuration option is enabled. The directory must be the full path name to the directory, not a relative path. See [Section 14.2.4.2.34, “Improved Tablespace Management”](#) for details about the performance aspects of tablespace management.

When creating MyISAM tables, you can use the `DATA DIRECTORY='directory'` clause, the `INDEX DIRECTORY='directory'` clause, or both. They specify where to put a MyISAM table's data file and index file respectively.

**Important**

Table-level `DATA DIRECTORY` and `INDEX DIRECTORY` options are ignored for partitioned tables. (Bug #32091)

These options work only when you are not using the `--skip-symbolic-links` [441] option. Your operating system must also have a working, thread-safe `realpath()` call. See [Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”](#), for more complete information.

If a MyISAM table is created with no `DATA DIRECTORY` option, the `.MYD` file is created in the database directory. By default, if MyISAM finds an existing `.MYD` file in this case, it overwrites it. The same applies to `.MYI` files for tables created with no `INDEX DIRECTORY` option. To suppress this behavior, start the server with the `--keep_files_on_create` [491] option, in which case MyISAM will not overwrite existing files and returns an error instead.

If a MyISAM table is created with a `DATA DIRECTORY` or `INDEX DIRECTORY` option and an existing `.MYD` or `.MYI` file is found, MyISAM always returns an error. It will not overwrite a file in the specified directory.

**Important**

You cannot use path names that contain the MySQL data directory with `DATA DIRECTORY` or `INDEX DIRECTORY`. This includes partitioned tables and individual table partitions. (See Bug #32167.)

- `DELAY_KEY_WRITE`

Set this to 1 if you want to delay key updates for the table until the table is closed. See the description of the `delay_key_write` [475] system variable in [Section 5.1.4, “Server System Variables”](#). (MyISAM only.)

- `INSERT_METHOD`

If you want to insert data into a MERGE table, you must specify with `INSERT_METHOD` the table into which the row should be inserted. `INSERT_METHOD` is an option useful for MERGE tables only. Use a value of `FIRST` or `LAST` to have inserts go to the first or last table, or a value of `NO` to prevent inserts. See [Section 14.8, “The MERGE Storage Engine”](#).

- `KEY_BLOCK_SIZE`

For compressed InnoDB tables, optionally specifies the size in bytes to use for pages. The value is treated as a hint; a different size could be used if necessary. A value of 0 represents that the default compressed page size. See [Section 5.4.6, “Working with InnoDB Compressed Tables”](#) for usage details.

Individual index definitions can specify a `KEY_BLOCK_SIZE` value of their own to override the table value.



### Note

Oracle recommends enabling `innodb_strict_mode` [1749] when using the `KEY_BLOCK_SIZE` clause for InnoDB tables. See Section 14.2.5.7, “InnoDB Strict Mode” for details.

- `MAX_ROWS`

The maximum number of rows you plan to store in the table. This is not a hard limit, but rather a hint to the storage engine that the table must be able to store at least this many rows.

The maximum `MAX_ROWS` value is 4294967295; larger values are truncated to this limit.

- `MIN_ROWS`

The minimum number of rows you plan to store in the table. The `MEMORY` storage engine uses this option as a hint about memory use.

- `PACK_KEYS`

`PACK_KEYS` takes effect only with `MyISAM` tables. Set this option to 1 if you want to have smaller indexes. This usually makes updates slower and reads faster. Setting the option to 0 disables all packing of keys. Setting it to `DEFAULT` tells the storage engine to pack only long `CHAR`, `VARCHAR`, `BINARY`, or `VARBINARY` columns.

If you do not use `PACK_KEYS`, the default is to pack strings, but not numbers. If you use `PACK_KEYS=1`, numbers are packed as well.

When packing binary number keys, MySQL uses prefix compression:

- Every key needs one extra byte to indicate how many bytes of the previous key are the same for the next key.
- The pointer to the row is stored in high-byte-first order directly after the key, to improve compression.

This means that if you have many equal keys on two consecutive rows, all following “same” keys usually only take two bytes (including the pointer to the row). Compare this to the ordinary case where the following keys takes `storage_size_for_key + pointer_size` (where the pointer size is usually 4). Conversely, you get a significant benefit from prefix compression only if you have many numbers that are the same. If all keys are totally different, you use one byte more per key, if the key is not a key that can have `NULL` values. (In this case, the packed key length is stored in the same byte that is used to mark if a key is `NULL`.)

- `PASSWORD`

This option is unused. If you have a need to scramble your `.frm` files and make them unusable to any other MySQL server, please contact our sales department.

- `ROW_FORMAT`

Defines the physical format in which the rows are stored. The choices differ depending on the storage engine used for the table.

For `InnoDB` tables:

- Rows are stored in compact format (`ROW_FORMAT=COMPACT`) by default.



- The noncompact format used in older versions of MySQL can still be requested by specifying `ROW_FORMAT=REDUNDANT`.
- To enable compression for InnoDB tables, specify `ROW_FORMAT=COMPRESSED` and follow the procedures in [Section 5.4.6, “Working with InnoDB Compressed Tables”](#).
- For more efficient InnoDB storage of data types, especially BLOB types, specify `ROW_FORMAT=DYNAMIC` and follow the procedures in [Section 5.4.8.3, “DYNAMIC and COMPRESSED Row Formats”](#). Both the `COMPRESSED` and `DYNAMIC` row formats require creating the table with the configuration settings `innodb_file_per_table=1` [1718] and `innodb_file_format=barracuda` [1716].
- When you specify a non-default `ROW_FORMAT` clause, consider also enabling the `innodb_strict_mode` [1749] configuration option. See [Section 14.2.5.7, “InnoDB Strict Mode”](#) for details.

For MyISAM tables, the option value can be `FIXED` or `DYNAMIC` for static or variable-length row format. `myisampack` sets the type to `COMPRESSED`. See [Section 14.3.3, “MyISAM Table Storage Formats”](#).



#### Note

When executing a `CREATE TABLE` statement, if you specify a row format that is not supported by the storage engine that is used for the table, the table is created using that storage engine's default row format. The information reported in this column in response to `SHOW TABLE STATUS` is the actual row format used. This may differ from the value in the `Create_options` column because the original `CREATE TABLE` definition is retained during creation.

- `STATS_AUTO_RECALC`

Specifies whether to automatically recalculate [persistent statistics](#) for an InnoDB table. The value `DEFAULT` causes the persistent statistics setting for the table to be determined by the `innodb_stats_auto_recalc` [1746] configuration option. The value `1` causes statistics to be recalculated when 10% of the data in the table has changed. The value `0` prevents automatic recalculation for this table; with this setting, issue an `ANALYZE TABLE` statement to recalculate the statistics after making substantial changes to the table. For more information about the persistent statistics feature, see [Section 14.2.4.2.10, “Persistent Optimizer Statistics for InnoDB Tables”](#).

- `STATS_PERSISTENT`

Specifies whether to enable [persistent statistics](#) for an InnoDB table. The value `DEFAULT` causes the persistent statistics setting for the table to be determined by the `innodb_stats_persistent` [1747] configuration option. The value `1` enables persistent statistics for the table, while the value `0` turns off this feature. After enabling persistent statistics through a `CREATE TABLE` or `ALTER TABLE` statement, issue an `ANALYZE TABLE` statement to calculate the statistics, after loading representative data into the table. For more information about the persistent statistics feature, see [Section 14.2.4.2.10, “Persistent Optimizer Statistics for InnoDB Tables”](#).

- `UNION`

`UNION` is used when you want to access a collection of identical MyISAM tables as one. This works only with `MERGE` tables. See [Section 14.8, “The MERGE Storage Engine”](#).

You must have `SELECT` [755], `UPDATE` [756], and `DELETE` [754] privileges for the tables you map to a `MERGE` table.

**Note**

Formerly, all tables used had to be in the same database as the `MERGE` table itself. This restriction no longer applies.

**Partitioning**

`partition_options` can be used to control partitioning of the table created with `CREATE TABLE`.

**Important**

Not all options shown in the syntax for `partition_options` at the beginning of this section are available for all partitioning types. Please see the listings for the following individual types for information specific to each type, and see [Chapter 17, Partitioning](#), for more complete information about the workings of and uses for partitioning in MySQL, as well as additional examples of table creation and other statements relating to MySQL partitioning.

If used, a `partition_options` clause begins with `PARTITION BY`. This clause contains the function that is used to determine the partition; the function returns an integer value ranging from 1 to `num`, where `num` is the number of partitions. (The maximum number of user-defined partitions which a table may contain is 1024; the number of subpartitions—discussed later in this section—is included in this maximum.) The choices that are available for this function in MySQL 5.6 are shown in the following list:

- `HASH(expr)`: Hashes one or more columns to create a key for placing and locating rows. *expr* is an expression using one or more table columns. This can be any valid MySQL expression (including MySQL functions) that yields a single integer value. For example, these are both valid `CREATE TABLE` statements using `PARTITION BY HASH`:

```
CREATE TABLE t1 (col1 INT, col2 CHAR(5))
  PARTITION BY HASH(col1);

CREATE TABLE t1 (col1 INT, col2 CHAR(5), col3 DATETIME)
  PARTITION BY HASH ( YEAR(col3) );
```

You may not use either `VALUES LESS THAN` or `VALUES IN` clauses with `PARTITION BY HASH`.

`PARTITION BY HASH` uses the remainder of *expr* divided by the number of partitions (that is, the modulus). For examples and additional information, see [Section 17.2.4, “HASH Partitioning”](#).

The `LINEAR` keyword entails a somewhat different algorithm. In this case, the number of the partition in which a row is stored is calculated as the result of one or more logical `AND [1171]` operations. For discussion and examples of linear hashing, see [Section 17.2.4.1, “LINEAR HASH Partitioning”](#).

- `KEY(column_list)`: This is similar to `HASH`, except that MySQL supplies the hashing function so as to guarantee an even data distribution. The *column\_list* argument is simply a list of 1 or more table columns (maximum: 16). This example shows a simple table partitioned by key, with 4 partitions:

```
CREATE TABLE tk (col1 INT, col2 CHAR(5), col3 DATE)
  PARTITION BY KEY(col3)
  PARTITIONS 4;
```

For tables that are partitioned by key, you can employ linear partitioning by using the `LINEAR` keyword. This has the same effect as with tables that are partitioned by `HASH`. That is, the partition number is found using the `& [1258]` operator rather than the modulus (see [Section 17.2.4.1, “LINEAR HASH Partitioning”](#), and [Section 17.2.5, “KEY Partitioning”](#), for details). This example uses linear partitioning by key to distribute data between 5 partitions:

```
CREATE TABLE tk (col1 INT, col2 CHAR(5), col3 DATE)
  PARTITION BY LINEAR KEY(col3)
  PARTITIONS 5;
```



The `ALGORITHM={1|2}` option is supported with `[SUB]PARTITION BY [LINEAR] KEY` beginning with MySQL 5.6.11. `ALGORITHM=1` causes the server to use the same key-hashing functions as MySQL 5.1; `ALGORITHM=2` means that the server employs the key-hashing functions implemented and used by default for new `KEY` partitioned tables in MySQL 5.5 and later. (Partitioned tables created with the key-hashing functions employed in MySQL 5.5 and later cannot be used by a MySQL 5.1 server.) Not specifying the option has the same effect as using `ALGORITHM=2`. This option is intended for use chiefly when upgrading or downgrading `[LINEAR] KEY` partitioned tables between MySQL 5.1 and later MySQL versions, or for creating tables partitioned by `KEY` or `LINEAR KEY` on a MySQL 5.5 or later server which can be used on a MySQL 5.1 server. For more information, see [Section 13.1.6.1, “ALTER TABLE Partition Operations”](#).

`mysqldump` in MySQL 5.6.11 and later writes this option encased in versioned comments, like this:

```
CREATE TABLE t1 (a INT)
/*!50100 PARTITION BY KEY */ /*!50611 ALGORITHM = 1 */ /*!50100 ()
PARTITIONS 3 */
```

This causes MySQL 5.6.10 and earlier servers to ignore the option, which would otherwise cause a syntax error in those versions. If you plan to load a dump made on a MySQL 5.5.31 or later MySQL 5.5 server where you use tables that are partitioned or subpartitioned by `KEY` into a MySQL 5.6 server previous to version 5.6.11, be sure to consult [Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”](#), before proceeding. (The information found there also applies if you are loading a dump containing `KEY` partitioned or subpartitioned tables made from a MySQL 5.6.11 or later server into a MySQL 5.5.30 or earlier server.)

Also in MySQL 5.6.11 and later, `ALGORITHM=1` is shown when necessary in the output of `SHOW CREATE TABLE` using versioned comments in the same manner as `mysqldump`. `ALGORITHM=2` is always omitted from `SHOW CREATE TABLE` output, even if this option was specified when creating the original table.

You may not use either `VALUES LESS THAN` or `VALUES IN` clauses with `PARTITION BY KEY`.

- `RANGE(expr)`: In this case, `expr` shows a range of values using a set of `VALUES LESS THAN` operators. When using range partitioning, you must define at least one partition using `VALUES LESS THAN`. You cannot use `VALUES IN` with range partitioning.



#### Note

For tables partitioned by `RANGE`, `VALUES LESS THAN` must be used with either an integer literal value or an expression that evaluates to a single integer value. In MySQL 5.6, you can overcome this limitation in a table that is defined using `PARTITION BY RANGE COLUMNS`, as described later in this section.

Suppose that you have a table that you wish to partition on a column containing year values, according to the following scheme.

Partition Number:	Years Range:
0	1990 and earlier
1	1991 to 1994
2	1995 to 1998
3	1999 to 2002
4	2003 to 2005
5	2006 and later

A table implementing such a partitioning scheme can be realized by the `CREATE TABLE` statement shown here:

```
CREATE TABLE t1 (
  year_col INT,
  some_data INT
)
PARTITION BY RANGE (year_col) (
  PARTITION p0 VALUES LESS THAN (1991),
  PARTITION p1 VALUES LESS THAN (1995),
  PARTITION p2 VALUES LESS THAN (1999),
  PARTITION p3 VALUES LESS THAN (2002),
  PARTITION p4 VALUES LESS THAN (2006),
  PARTITION p5 VALUES LESS THAN MAXVALUE
);
```

`PARTITION ... VALUES LESS THAN ...` statements work in a consecutive fashion. `VALUES LESS THAN MAXVALUE` works to specify “leftover” values that are greater than the maximum value otherwise specified.

Note that `VALUES LESS THAN` clauses work sequentially in a manner similar to that of the `case` portions of a `switch ... case` block (as found in many programming languages such as C, Java, and PHP). That is, the clauses must be arranged in such a way that the upper limit specified in each successive `VALUES LESS THAN` is greater than that of the previous one, with the one referencing `MAXVALUE` coming last of all in the list.

- `RANGE COLUMNS(column_list)`: This variant on `RANGE` facilitates partition pruning for queries using range conditions on multiple columns (that is, having conditions such as `WHERE a = 1 AND b < 10` or `WHERE a = 1 AND b = 10 AND c < 10`). It enables you to specify value ranges in multiple columns by using a list of columns in the `COLUMNS` clause and a set of column values in each `PARTITION ... VALUES LESS THAN (value_list)` partition definition clause. (In the simplest case, this set consists of a single column.) The maximum number of columns that can be referenced in the `column_list` and `value_list` is 16.

The `column_list` used in the `COLUMNS` clause may contain only names of columns; each column in the list must be one of the following MySQL data types: the integer types; the string types; and time or date column types. Columns using `BLOB`, `TEXT`, `SET`, `ENUM`, `BIT`, or spatial data types are not permitted; columns that use floating-point number types are also not permitted. You also may not use functions or arithmetic expressions in the `COLUMNS` clause.

The `VALUES LESS THAN` clause used in a partition definition must specify a literal value for each column that appears in the `COLUMNS()` clause; that is, the list of values used for each `VALUES LESS THAN` clause must contain the same number of values as there are columns listed in the `COLUMNS` clause. An attempt to use more or fewer values in a `VALUES LESS THAN` clause than there are in the `COLUMNS` clause causes the statement to fail with the error `Inconsistency in usage of column lists for partitioning...` You cannot use `NULL` for any value appearing in `VALUES LESS THAN`. It is possible to use `MAXVALUE` more than once for a given column other than the first, as shown in this example:

```
CREATE TABLE rc (
  a INT NOT NULL,
  b INT NOT NULL
)
PARTITION BY RANGE COLUMNS(a,b) (
  PARTITION p0 VALUES LESS THAN (10,5),
  PARTITION p1 VALUES LESS THAN (20,10),
  PARTITION p2 VALUES LESS THAN (MAXVALUE,15),
  PARTITION p3 VALUES LESS THAN (MAXVALUE,MAXVALUE)
);
```

Each value used in a `VALUES LESS THAN` value list must match the type of the corresponding column exactly; no conversion is made. For example, you cannot use the string `'1'` for a value that matches a column that uses an integer type (you must use the numeral `1` instead), nor can you use the numeral `1` for a value that matches a column that uses a string type (in such a case, you must use a quoted string: `'1'`).

For more information, see [Section 17.2.1, “RANGE Partitioning”](#), and [Section 17.4, “Partition Pruning”](#).

- **LIST(*expr*)**: This is useful when assigning partitions based on a table column with a restricted set of possible values, such as a state or country code. In such a case, all rows pertaining to a certain state or country can be assigned to a single partition, or a partition can be reserved for a certain set of states or countries. It is similar to **RANGE**, except that only **VALUES IN** may be used to specify permissible values for each partition.

**VALUES IN** is used with a list of values to be matched. For instance, you could create a partitioning scheme such as the following:

```
CREATE TABLE client_firms (
  id INT,
  name VARCHAR(35)
)
PARTITION BY LIST (id) (
  PARTITION r0 VALUES IN (1, 5, 9, 13, 17, 21),
  PARTITION r1 VALUES IN (2, 6, 10, 14, 18, 22),
  PARTITION r2 VALUES IN (3, 7, 11, 15, 19, 23),
  PARTITION r3 VALUES IN (4, 8, 12, 16, 20, 24)
);
```

When using list partitioning, you must define at least one partition using **VALUES IN**. You cannot use **VALUES LESS THAN** with **PARTITION BY LIST**.



#### Note

For tables partitioned by **LIST**, the value list used with **VALUES IN** must consist of integer values only. In MySQL 5.6, you can overcome this limitation using partitioning by **LIST COLUMNS**, which is described later in this section.

- **LIST COLUMNS(*column\_list*)**: This variant on **LIST** facilitates partition pruning for queries using comparison conditions on multiple columns (that is, having conditions such as **WHERE a = 5 AND b = 5** or **WHERE a = 1 AND b = 10 AND c = 5**). It enables you to specify values in multiple columns by using a list of columns in the **COLUMNS** clause and a set of column values in each **PARTITION ... VALUES IN (*value\_list*)** partition definition clause.

The rules governing regarding data types for the column list used in **LIST COLUMNS(*column\_list*)** and the value list used in **VALUES IN(*value\_list*)** are the same as those for the column list used in **RANGE COLUMNS(*column\_list*)** and the value list used in **VALUES LESS THAN(*value\_list*)**, respectively, except that in the **VALUES IN** clause, **MAXVALUE** is not permitted, and you may use **NULL**.

There is one important difference between the list of values used for **VALUES IN** with **PARTITION BY LIST COLUMNS** as opposed to when it is used with **PARTITION BY LIST**. When used with **PARTITION BY LIST COLUMNS**, each element in the **VALUES IN** clause must be a set of column values; the number of values in each set must be the same as the number of columns used in the **COLUMNS** clause, and the data types of these values must match those of the columns (and occur in the same order). In the simplest case, the set consists of a single column. The maximum number of columns that can be used in the *column\_list* and in the elements making up the *value\_list* is 16.

The table defined by the following **CREATE TABLE** statement provides an example of a table using **LIST COLUMNS** partitioning:

```
CREATE TABLE lc (
  a INT NULL,
  b INT NULL
)
PARTITION BY LIST COLUMNS(a,b) (
  PARTITION p0 VALUES IN( (0,0), (NULL,NULL) ),
  PARTITION p1 VALUES IN( (0,1), (0,2), (0,3), (1,1), (1,2) ),
  PARTITION p2 VALUES IN( (1,0), (2,0), (2,1), (3,0), (3,1) ),
```

```
PARTITION p3 VALUES IN( (1,3), (2,2), (2,3), (3,2), (3,3) )
);
```

- The number of partitions may optionally be specified with a `PARTITIONS num` clause, where *num* is the number of partitions. If both this clause *and* any `PARTITION` clauses are used, *num* must be equal to the total number of any partitions that are declared using `PARTITION` clauses.



#### Note

Whether or not you use a `PARTITIONS` clause in creating a table that is partitioned by `RANGE` or `LIST`, you must still include at least one `PARTITION VALUES` clause in the table definition (see below).

- A partition may optionally be divided into a number of subpartitions. This can be indicated by using the optional `SUBPARTITION BY` clause. Subpartitioning may be done by `HASH` or `KEY`. Either of these may be `LINEAR`. These work in the same way as previously described for the equivalent partitioning types. (It is not possible to subpartition by `LIST` or `RANGE`.)

The number of subpartitions can be indicated using the `SUBPARTITIONS` keyword followed by an integer value.

- Rigorous checking of the value used in `PARTITIONS` or `SUBPARTITIONS` clauses is applied and this value must adhere to the following rules:
  - The value must be a positive, nonzero integer.
  - No leading zeros are permitted.
  - The value must be an integer literal, and cannot not be an expression. For example, `PARTITIONS 0.2E+01` is not permitted, even though `0.2E+01` evaluates to 2. (Bug #15890)



#### Note

The expression (*expr*) used in a `PARTITION BY` clause cannot refer to any columns not in the table being created; such references are specifically not permitted and cause the statement to fail with an error. (Bug #29444)

Each partition may be individually defined using a `partition_definition` clause. The individual parts making up this clause are as follows:

- `PARTITION partition_name`: This specifies a logical name for the partition.
- A `VALUES` clause: For range partitioning, each partition must include a `VALUES LESS THAN` clause; for list partitioning, you must specify a `VALUES IN` clause for each partition. This is used to determine which rows are to be stored in this partition. See the discussions of partitioning types in [Chapter 17, Partitioning](#), for syntax examples.
- An optional `COMMENT` clause may be used to specify a string that describes the partition. Example:

```
COMMENT = 'Data for the years previous to 1999'
```

Beginning with MySQL 5.6.6, the maximum length for a partition comment is 1024 characters. (Previously, this limit was not explicitly defined.)

- `DATA DIRECTORY` and `INDEX DIRECTORY` may be used to indicate the directory where, respectively, the data and indexes for this partition are to be stored. Both the `data_dir` and the `index_dir` must be absolute system path names. Example:

```
CREATE TABLE th (id INT, name VARCHAR(30), adate DATE)
PARTITION BY LIST(YEAR(adate))
(
  PARTITION p1999 VALUES IN (1995, 1999, 2003)
  DATA DIRECTORY = '/var/appdata/95/data'
  INDEX DIRECTORY = '/var/appdata/95/idx',
```

```

PARTITION p2000 VALUES IN (1996, 2000, 2004)
  DATA DIRECTORY = '/var/appdata/96/data'
  INDEX DIRECTORY = '/var/appdata/96/idx',
PARTITION p2001 VALUES IN (1997, 2001, 2005)
  DATA DIRECTORY = '/var/appdata/97/data'
  INDEX DIRECTORY = '/var/appdata/97/idx',
PARTITION p2002 VALUES IN (1998, 2002, 2006)
  DATA DIRECTORY = '/var/appdata/98/data'
  INDEX DIRECTORY = '/var/appdata/98/idx'
);

```

`DATA DIRECTORY` and `INDEX DIRECTORY` behave in the same way as in the `CREATE TABLE` statement's `table_option` clause as used for `MyISAM` tables.

One data directory and one index directory may be specified per partition. If left unspecified, the data and indexes are stored by default in the table's database directory.

On Windows, the `DATA DIRECTORY` and `INDEX DIRECTORY` options are not supported for individual partitions or subpartitions of `MyISAM` tables, and the `INDEX DIRECTORY` option is not supported for individual partitions or subpartitions of `InnoDB` tables. These options are ignored on Windows, except that a warning is generated. (Bug #30459)



#### Note

The `DATA DIRECTORY` and `INDEX DIRECTORY` options are ignored for creating partitioned tables if `NO_DIR_IN_CREATE` [604] is in effect. (Bug #24633)

- `MAX_ROWS` and `MIN_ROWS` may be used to specify, respectively, the maximum and minimum number of rows to be stored in the partition. The values for `max_number_of_rows` and `min_number_of_rows` must be positive integers. As with the table-level options with the same names, these act only as “suggestions” to the server and are not hard limits.
- The partitioning handler accepts a `[STORAGE] ENGINE` option for both `PARTITION` and `SUBPARTITION`. Currently, the only way in which this can be used is to set all partitions or all subpartitions to the same storage engine, and an attempt to set different storage engines for partitions or subpartitions in the same table will give rise to the error `ERROR 1469 (HY000): The mix of handlers in the partitions is not permitted in this version of MySQL`. We expect to lift this restriction on partitioning in a future MySQL release.
- The partition definition may optionally contain one or more `subpartition_definition` clauses. Each of these consists at a minimum of the `SUBPARTITION name`, where `name` is an identifier for the subpartition. Except for the replacement of the `PARTITION` keyword with `SUBPARTITION`, the syntax for a subpartition definition is identical to that for a partition definition.

Subpartitioning must be done by `HASH` or `KEY`, and can be done only on `RANGE` or `LIST` partitions. See [Section 17.2.6, “Subpartitioning”](#).

Partitions can be modified, merged, added to tables, and dropped from tables. For basic information about the MySQL statements to accomplish these tasks, see [Section 13.1.6, “ALTER TABLE Syntax”](#). For more detailed descriptions and examples, see [Section 17.3, “Partition Management”](#).



#### Important

The original `CREATE TABLE` statement, including all specifications and table options are stored by MySQL when the table is created. The information is retained so that if you change storage engines, collations or other settings using an `ALTER TABLE` statement, the original table options specified are retained. This enables you to change between `InnoDB` and `MyISAM` table types even though the row formats supported by the two engines are different.

Because the text of the original statement is retained, but due to the way that certain values and options may be silently reconfigured (such as the

`ROW_FORMAT`), the active table definition (accessible through `DESCRIBE` or with `SHOW TABLE STATUS`) and the table creation string (accessible through `SHOW CREATE TABLE`) will report different values.

## Cloning or Copying a Table

You can create one table from another by adding a `SELECT` statement at the end of the `CREATE TABLE` statement:

```
CREATE TABLE new_tbl SELECT * FROM orig_tbl;
```

For more information, see [Section 13.1.14.1, “CREATE TABLE ... SELECT Syntax”](#).

Use `LIKE` to create an empty table based on the definition of another table, including any column attributes and indexes defined in the original table:

```
CREATE TABLE new_tbl LIKE orig_tbl;
```

The copy is created using the same version of the table storage format as the original table. The `SELECT [755]` privilege is required on the original table.

`LIKE` works only for base tables, not for views.



### Important

Beginning with MySQL 5.6.1, you cannot execute `CREATE TABLE` or `CREATE TABLE ... LIKE` while a `LOCK TABLES` statement is in effect.

Also as of MySQL 5.6.1, `CREATE TABLE ... LIKE` makes the same checks as `CREATE TABLE` and does not just copy the `.frm` file. This means that if the current SQL mode is different from the mode in effect when the original table was created, the table definition might be considered invalid for the new mode and the statement will fail.

`CREATE TABLE ... LIKE` does not preserve any `DATA DIRECTORY` or `INDEX DIRECTORY` table options that were specified for the original table, or any foreign key definitions.

If the original table is a `TEMPORARY` table, `CREATE TABLE ... LIKE` does not preserve `TEMPORARY`. To create a `TEMPORARY` destination table, use `CREATE TEMPORARY TABLE ... LIKE`.

### 13.1.14.1. CREATE TABLE ... SELECT Syntax

You can create one table from another by adding a `SELECT` statement at the end of the `CREATE TABLE` statement:

```
CREATE TABLE new_tbl SELECT * FROM orig_tbl;
```

MySQL creates new columns for all elements in the `SELECT`. For example:

```
mysql> CREATE TABLE test (a INT NOT NULL AUTO_INCREMENT,
-> PRIMARY KEY (a), KEY(b))
-> ENGINE=MyISAM SELECT b,c FROM test2;
```

This creates a `MyISAM` table with three columns, `a`, `b`, and `c`. The `ENGINE` option is part of the `CREATE TABLE` statement, and should not be used following the `SELECT`; this would result in a syntax error. The same is true for other `CREATE TABLE` options such as `CHARSET`.

Notice that the columns from the `SELECT` statement are appended to the right side of the table, not overlapped onto it. Take the following example:

```
mysql> SELECT * FROM foo;
+----+
| n  |
+----+
| 1  |
```



```
+----+
mysql> CREATE TABLE bar (m INT) SELECT n FROM foo;
Query OK, 1 row affected (0.02 sec)
Records: 1 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM bar;
+-----+-----+
| m      | n      |
+-----+-----+
| NULL   | 1      |
+-----+-----+
1 row in set (0.00 sec)
```

For each row in table `foo`, a row is inserted in `bar` with the values from `foo` and default values for the new columns.

In a table resulting from `CREATE TABLE ... SELECT`, columns named only in the `CREATE TABLE` part come first. Columns named in both parts or only in the `SELECT` part come after that. The data type of `SELECT` columns can be overridden by also specifying the column in the `CREATE TABLE` part.

If any errors occur while copying the data to the table, it is automatically dropped and not created.

You can precede the `SELECT` by `IGNORE` or `REPLACE` to indicate how to handle rows that duplicate unique key values. With `IGNORE`, new rows that duplicate an existing row on a unique key value are discarded. With `REPLACE`, new rows replace rows that have the same unique key value. If neither `IGNORE` nor `REPLACE` is specified, duplicate unique key values result in an error.

Because the ordering of the rows in the underlying `SELECT` statements cannot always be determined, `CREATE TABLE ... IGNORE SELECT` and `CREATE TABLE ... REPLACE SELECT` statements in MySQL 5.6.4 and later are flagged as unsafe for statement-based replication. With this change, such statements produce a warning in the log when using statement-based mode and are logged using the row-based format when using `MIXED` mode. See also [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

`CREATE TABLE ... SELECT` does not automatically create any indexes for you. This is done intentionally to make the statement as flexible as possible. If you want to have indexes in the created table, you should specify these before the `SELECT` statement:

```
mysql> CREATE TABLE bar (UNIQUE (n)) SELECT n FROM foo;
```

Some conversion of data types might occur. For example, the `AUTO_INCREMENT` attribute is not preserved, and `VARCHAR` columns can become `CHAR` columns. Retained attributes are `NULL` (or `NOT NULL`) and, for those columns that have them, `CHARACTER SET`, `COLLATION`, `COMMENT`, and the `DEFAULT` clause.

When creating a table with `CREATE TABLE ... SELECT`, make sure to alias any function calls or expressions in the query. If you do not, the `CREATE` statement might fail or result in undesirable column names.

```
CREATE TABLE artists_and_works
  SELECT artist.name, COUNT(work.artist_id) AS number_of_works
  FROM artist LEFT JOIN work ON artist.id = work.artist_id
  GROUP BY artist.id;
```

You can also explicitly specify the data type for a generated column:

```
CREATE TABLE foo (a TINYINT NOT NULL) SELECT b+1 AS a FROM bar;
```

For `CREATE TABLE ... SELECT`, if `IF NOT EXISTS` is given and the destination table already exists, the result is version dependent. Before MySQL 5.5.6, MySQL handles the statement as follows:

- The table definition given in the `CREATE TABLE` part is ignored. No error occurs, even if the definition does not match that of the existing table. MySQL attempts to insert the rows from the `SELECT` part anyway.

- If there is a mismatch between the number of columns in the table and the number of columns produced by the `SELECT` part, the selected values are assigned to the rightmost columns. For example, if the table contains  $n$  columns and the `SELECT` produces  $m$  columns, where  $m < n$ , the selected values are assigned to the  $m$  rightmost columns in the table. Each of the initial  $n - m$  columns is assigned its default value, either that specified explicitly in the column definition or the implicit column data type default if the definition contains no default. If the `SELECT` part produces too many columns ( $m > n$ ), an error occurs.
- If strict SQL mode is enabled and any of these initial columns do not have an explicit default value, the statement fails with an error.

The following example illustrates `IF NOT EXISTS` handling:

```
mysql> CREATE TABLE t1 (i1 INT DEFAULT 0, i2 INT, i3 INT, i4 INT);
Query OK, 0 rows affected (0.05 sec)

mysql> CREATE TABLE IF NOT EXISTS t1 (c1 CHAR(10)) SELECT 1, 2;
Query OK, 1 row affected, 1 warning (0.01 sec)
Records: 1 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM t1;
+-----+-----+-----+-----+
| i1   | i2   | i3   | i4   |
+-----+-----+-----+-----+
| 0    | NULL | 1    | 2    |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

As of MySQL 5.5.6, handling of `CREATE TABLE IF NOT EXISTS ... SELECT` statements was changed for the case that the destination table already exists. This change also involves a change in MySQL 5.1 beginning with 5.1.51.

- Previously, for `CREATE TABLE IF NOT EXISTS ... SELECT`, MySQL produced a warning that the table exists, but inserted the rows and wrote the statement to the binary log anyway. By contrast, `CREATE TABLE ... SELECT` (without `IF NOT EXISTS`) failed with an error, but MySQL inserted no rows and did not write the statement to the binary log.
- MySQL now handles both statements the same way when the destination table exists, in that neither statement inserts rows or is written to the binary log. The difference between them is that MySQL produces a warning when `IF NOT EXISTS` is present and an error when it is not.

This change means that, for the preceding example, the `CREATE TABLE IF NOT EXISTS ... SELECT` statement inserts nothing into the destination table as of MySQL 5.5.6.

This change in handling of `IF NOT EXISTS` results in an incompatibility for statement-based replication from a MySQL 5.1 master with the original behavior and a MySQL 5.5 slave with the new behavior. Suppose that `CREATE TABLE IF NOT EXISTS ... SELECT` is executed on the master and the destination table exists. The result is that rows are inserted on the master but not on the slave. (Row-based replication does not have this problem.)

To address this issue, statement-based binary logging for `CREATE TABLE IF NOT EXISTS ... SELECT` is changed in MySQL 5.1 as of 5.1.51:

- If the destination table does not exist, there is no change: The statement is logged as is.
- If the destination table does exist, the statement is logged as the equivalent pair of `CREATE TABLE IF NOT EXISTS` and `INSERT ... SELECT` statements. (If the `SELECT` in the original statement is preceded by `IGNORE` or `REPLACE`, the `INSERT` becomes `INSERT IGNORE` or `REPLACE`, respectively.)

This change provides forward compatibility for statement-based replication from MySQL 5.1 to 5.5 because when the destination table exists, the rows will be inserted on both the master and slave. To take advantage of this compatibility measure, the 5.1 server must be at least 5.1.51 and the 5.5 server must be at least 5.5.6.



To upgrade an existing 5.1-to-5.5 replication scenario, upgrade the master first to 5.1.51 or higher. Note that this differs from the usual replication upgrade advice of upgrading the slave first.

A workaround for applications that wish to achieve the original effect (rows inserted regardless of whether the destination table exists) is to use `CREATE TABLE IF NOT EXISTS` and `INSERT ... SELECT` statements rather than `CREATE TABLE IF NOT EXISTS ... SELECT` statements.

Along with the change just described, the following related change was made: Previously, if an existing view was named as the destination table for `CREATE TABLE IF NOT EXISTS ... SELECT`, rows were inserted into the underlying base table and the statement was written to the binary log. As of MySQL 5.1.51 and 5.5.6, nothing is inserted or logged.

To ensure that the binary log can be used to re-create the original tables, MySQL does not permit concurrent inserts during `CREATE TABLE ... SELECT`.



### Important

You cannot use `FOR UPDATE` as part of the `SELECT` in a statement such as `CREATE TABLE new_table SELECT ... FROM old_table ...`. If you attempt to do so, the statement fails. This represents a change in behavior from MySQL 5.5 and earlier, which permitted `CREATE TABLE ... SELECT` statements to make changes in tables other than the table being created.

This change can also have implications for statement-based replication from an older master to a MySQL 5.6 or newer slave. See [Section 16.4.1.4, “Replication of CREATE TABLE ... SELECT Statements”](#), for more information.

### 13.1.14.2. Using FOREIGN KEY Constraints

MySQL supports foreign keys, which let you cross-reference related data across tables, and [foreign key constraints](#), which help keep this spread-out data consistent. The essential syntax for a foreign key constraint definition in a `CREATE TABLE` or `ALTER TABLE` statement looks like this:

```
[CONSTRAINT [symbol]] FOREIGN KEY
  [index_name] (index_col_name, ...)
REFERENCES tbl_name (index_col_name,...)
[ON DELETE reference_option]
[ON UPDATE reference_option]
```

*reference\_option*:

```
RESTRICT | CASCADE | SET NULL | NO ACTION
```

*index\_name* represents a foreign key ID. If given, this is ignored if an index for the foreign key is defined explicitly. Otherwise, if MySQL creates an index for the foreign key, it uses *index\_name* for the index name.

Foreign keys definitions are subject to the following conditions:

- Foreign key relationships involve a [parent table](#) that holds the central data values, and a [child table](#) with identical values pointing back to its parent. The `FOREIGN KEY` clause is specified in the child table. The parent and child tables must use the same storage engine. They must not be `TEMPORARY` tables.
- Corresponding columns in the foreign key and the referenced key must have similar data types. *The size and sign of integer types must be the same.* The length of string types need not be the same. For nonbinary (character) string columns, the character set and collation must be the same.
- MySQL requires indexes on foreign keys and referenced keys so that foreign key checks can be fast and not require a table scan. In the referencing table, there must be an index where the foreign key columns are listed as the *first* columns in the same order. Such an index is created on the referencing table automatically if it does not exist. This index might be silently dropped later, if you create another index that can be used to enforce the foreign key constraint. *index\_name*, if given, is used as described previously.

- **InnoDB** permits a foreign key to reference any index column or group of columns. However, in the referenced table, there must be an index where the referenced columns are listed as the *first* columns in the same order.
- Index prefixes on foreign key columns are not supported. One consequence of this is that **BLOB** and **TEXT** columns cannot be included in a foreign key because indexes on those columns must always include a prefix length.
- If the **CONSTRAINT** *symbol* clause is given, the *symbol* value, if used, must be unique in the database. If the clause is not given, or a *symbol* is not included following the **CONSTRAINT** keyword, a name for the constraint is created automatically.
- **InnoDB** does not currently support foreign keys for tables with user-defined partitioning. This includes both parent and child tables.

## Referential Actions

This section describes how foreign keys help guarantee [referential integrity](#).

For storage engines supporting foreign keys, MySQL rejects any **INSERT** or **UPDATE** operation that attempts to create a foreign key value in a child table if there is no a matching candidate key value in the parent table.

When an **UPDATE** or **DELETE** operation affects a key value in the parent table that has matching rows in the child table, the result depends on the *referential action* specified using **ON UPDATE** and **ON DELETE** subclauses of the **FOREIGN KEY** clause. MySQL supports five options regarding the action to be taken, listed here:

- **CASCADE**: Delete or update the row from the parent table, and automatically delete or update the matching rows in the child table. Both **ON DELETE CASCADE** and **ON UPDATE CASCADE** are supported. Between two tables, do not define several **ON UPDATE CASCADE** clauses that act on the same column in the parent table or in the child table.



### Note

Currently, cascaded foreign key actions do not activate triggers.

- **SET NULL**: Delete or update the row from the parent table, and set the foreign key column or columns in the child table to **NULL**. Both **ON DELETE SET NULL** and **ON UPDATE SET NULL** clauses are supported.

If you specify a **SET NULL** action, *make sure that you have not declared the columns in the child table as **NOT NULL***.

- **RESTRICT**: Rejects the delete or update operation for the parent table. Specifying **RESTRICT** (or **NO ACTION**) is the same as omitting the **ON DELETE** or **ON UPDATE** clause.
- **NO ACTION**: A keyword from standard SQL. In MySQL, equivalent to **RESTRICT**. The MySQL Server rejects the delete or update operation for the parent table if there is a related foreign key value in the referenced table. Some database systems have deferred checks, and **NO ACTION** is a deferred check. In MySQL, foreign key constraints are checked immediately, so **NO ACTION** is the same as **RESTRICT**.
- **SET DEFAULT**: This action is recognized by the MySQL parser, but **InnoDB** rejects table definitions containing **ON DELETE SET DEFAULT** or **ON UPDATE SET DEFAULT** clauses.

For an **ON DELETE** or **ON UPDATE** that is not specified, the default action is always **RESTRICT**.

MySQL supports foreign key references between one column and another within a table. (A column cannot have a foreign key reference to itself.) In these cases, “child table records” really refers to dependent records within the same table.

## Examples of Foreign Key Clauses

Here is a simple example that relates `parent` and `child` tables through a single-column foreign key:

```
CREATE TABLE parent (
  id INT NOT NULL,
  PRIMARY KEY (id)
) ENGINE=INNODB;

CREATE TABLE child (
  id INT,
  parent_id INT,
  INDEX par_ind (parent_id),
  FOREIGN KEY (parent_id)
    REFERENCES parent(id)
    ON DELETE CASCADE
) ENGINE=INNODB;
```

A more complex example in which a `product_order` table has foreign keys for two other tables. One foreign key references a two-column index in the `product` table. The other references a single-column index in the `customer` table:

```
CREATE TABLE product (
  category INT NOT NULL, id INT NOT NULL,
  price DECIMAL,
  PRIMARY KEY(category, id)
) ENGINE=INNODB;

CREATE TABLE customer (
  id INT NOT NULL,
  PRIMARY KEY (id)
) ENGINE=INNODB;

CREATE TABLE product_order (
  no INT NOT NULL AUTO_INCREMENT,
  product_category INT NOT NULL,
  product_id INT NOT NULL,
  customer_id INT NOT NULL,

  PRIMARY KEY(no),
  INDEX (product_category, product_id),
  INDEX (customer_id),

  FOREIGN KEY (product_category, product_id)
    REFERENCES product(category, id)
    ON UPDATE CASCADE ON DELETE RESTRICT,

  FOREIGN KEY (customer_id)
    REFERENCES customer(id)
) ENGINE=INNODB;
```

## Adding foreign keys

You can add a new foreign key constraint to an existing table by using `ALTER TABLE`. The syntax relating to foreign keys for this statement is shown here:

```
ALTER TABLE tbl_name
  ADD [CONSTRAINT [symbol]] FOREIGN KEY
  [index_name] (index_col_name, ...)
  REFERENCES tbl_name (index_col_name,...)
  [ON DELETE reference_option]
  [ON UPDATE reference_option]
```

The foreign key can be self referential (referring to the same table). When you add a foreign key constraint to a table using `ALTER TABLE`, *remember to create the required indexes first*.

## Dropping Foreign Keys

You can also use `ALTER TABLE` to drop foreign keys, using the syntax shown here:

```
ALTER TABLE tbl_name DROP FOREIGN KEY fk_symbol;
```

If the **FOREIGN KEY** clause included a **CONSTRAINT** name when you created the foreign key, you can refer to that name to drop the foreign key. Otherwise, the *fk\_symbol* value is generated internally when the foreign key is created. To find out the symbol value when you want to drop a foreign key, use a **SHOW CREATE TABLE** statement, as shown here:

```
mysql> SHOW CREATE TABLE ibtest11c\G
***** 1. row *****
      Table: ibtest11c
Create Table: CREATE TABLE `ibtest11c` (
  `A` int(11) NOT NULL auto_increment,
  `D` int(11) NOT NULL default '0',
  `B` varchar(200) NOT NULL default '',
  `C` varchar(175) default NULL,
  PRIMARY KEY (`A`,`D`,`B`),
  KEY `B` (`B`,`C`),
  KEY `C` (`C`),
  CONSTRAINT `_0_38775` FOREIGN KEY (`A`,`D`)
REFERENCES `ibtest11a` (`A`,`D`)
ON DELETE CASCADE ON UPDATE CASCADE,
  CONSTRAINT `_0_38776` FOREIGN KEY (`B`,`C`)
REFERENCES `ibtest11a` (`B`,`C`)
ON DELETE CASCADE ON UPDATE CASCADE
) ENGINE=INNODB CHARSET=latin1
1 row in set (0.01 sec)

mysql> ALTER TABLE ibtest11c DROP FOREIGN KEY `_0_38775`;
```

You cannot add a foreign key and drop a foreign key in the same **ALTER TABLE** statement. A separate statements is required for each operation.

Before MySQL 5.6.7, using **ALTER TABLE** to change the definition of a foreign key column could cause a loss of referential integrity. For example, changing a foreign key column that contained **NULL** values to be **NOT NULL** caused the **NULL** values to be the empty string. Similarly, an **ALTER TABLE IGNORE** that removed rows in a parent table could break referential integrity.

As of 5.6.7, the server prohibits changes to foreign key columns with the potential to cause loss of referential integrity. A workaround is to use **ALTER TABLE ... DROP FOREIGN KEY** before changing the column definition and **ALTER TABLE ... ADD FOREIGN KEY** afterward.

## Foreign Keys and Other MySQL Statements

Table and column identifiers in a **FOREIGN KEY ... REFERENCES ...** clause can be quoted within backticks (```). Alternatively, double quotation marks (`"`) can be used if the **ANSI\_QUOTES** [603] SQL mode is enabled. The setting of the **lower\_case\_table\_names** [501] system variable is also taken into account.

You can view a child table's foreign key definitions as part of the output of the **SHOW CREATE TABLE** statement:

```
SHOW CREATE TABLE tbl_name;
```

You can also obtain information about foreign keys by querying the **INFORMATION\_SCHEMA.KEY\_COLUMN\_USAGE** table.

You can find information about foreign keys used by **InnoDB** tables in the **INNODB\_SYS\_FOREIGN** and **INNODB\_SYS\_FOREIGN\_COLS** tables, also in the **INFORMATION\_SCHEMA** database.

**mysqldump** produces correct definitions of tables in the dump file, including the foreign keys for child tables.

To make it easier to reload dump files for tables that have foreign key relationships, **mysqldump** automatically includes a statement in the dump output to set **foreign\_key\_checks** [483] to 0. This

avoids problems with tables having to be reloaded in a particular order when the dump is reloaded. It is also possible to set this variable manually:

```
mysql> SET foreign_key_checks = 0;
mysql> SOURCE dump_file_name;
mysql> SET foreign_key_checks = 1;
```

This enables you to import the tables in any order if the dump file contains tables that are not correctly ordered for foreign keys. It also speeds up the import operation. Setting `foreign_key_checks` [483] to 0 can also be useful for ignoring foreign key constraints during `LOAD DATA` and `ALTER TABLE` operations. However, even if `foreign_key_checks = 0` [483], MySQL does not permit the creation of a foreign key constraint where a column references a nonmatching column type. Also, if a table has foreign key constraints, `ALTER TABLE` cannot be used to alter the table to use another storage engine. To change the storage engine, you must drop any foreign key constraints first.

You cannot issue `DROP TABLE` for a table that is referenced by a `FOREIGN KEY` constraint, unless you do `SET foreign_key_checks = 0`. When you drop a table, any constraints that were defined in the statement used to create that table are also dropped.

If you re-create a table that was dropped, it must have a definition that conforms to the foreign key constraints referencing it. It must have the correct column names and types, and it must have indexes on the referenced keys, as stated earlier. If these are not satisfied, MySQL returns Error 1005 and refers to Error 150 in the error message, which means that a foreign key constraint was not correctly formed. Similarly, if an `ALTER TABLE` fails due to Error 150, this means that a foreign key definition would be incorrectly formed for the altered table.

For `InnoDB` tables, you can obtain a detailed explanation of the most recent `InnoDB` foreign key error in the MySQL Server, by checking the output of `SHOW ENGINE INNODB STATUS`.



### Important

For users familiar with the ANSI/ISO SQL Standard, please note that no storage engine, including `InnoDB`, recognizes or enforces the `MATCH` clause used in referential-integrity constraint definitions. Use of an explicit `MATCH` clause will not have the specified effect, and also causes `ON DELETE` and `ON UPDATE` clauses to be ignored. For these reasons, specifying `MATCH` should be avoided.

The `MATCH` clause in the SQL standard controls how `NULL` values in a composite (multiple-column) foreign key are handled when comparing to a primary key. MySQL essentially implements the semantics defined by `MATCH SIMPLE`, which permit a foreign key to be all or partially `NULL`. In that case, the (child table) row containing such a foreign key is permitted to be inserted, and does not match any row in the referenced (parent) table. It is possible to implement other semantics using triggers.

Additionally, MySQL requires that the referenced columns be indexed for performance reasons. However, the system does not enforce a requirement that the referenced columns be `UNIQUE` or be declared `NOT NULL`. The handling of foreign key references to nonunique keys or keys that contain `NULL` values is not well defined for operations such as `UPDATE` or `DELETE CASCADE`. You are advised to use foreign keys that reference only `UNIQUE` (including `PRIMARY`) and `NOT NULL` keys.

Furthermore, MySQL does not recognize or support “inline `REFERENCES` specifications” (as defined in the SQL standard) where the references are defined as part of the column specification. MySQL accepts `REFERENCES` clauses only when specified as part of a separate `FOREIGN KEY` specification. For storage engines that do not support foreign keys (such as `MyISAM`), MySQL Server parses and ignores foreign key specifications.

### 13.1.14.3. Silent Column Specification Changes

In some cases, MySQL silently changes column specifications from those given in a [CREATE TABLE](#) or [ALTER TABLE](#) statement. These might be changes to a data type, to attributes associated with a data type, or to an index specification.

All changes are subject to the internal row-size limit of 65,535 bytes, which may cause some attempts at data type changes to fail. See [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

- Columns that are part of a [PRIMARY KEY](#) are made [NOT NULL](#) even if not declared that way.
- Trailing spaces are automatically deleted from [ENUM](#) and [SET](#) member values when the table is created.
- MySQL maps certain data types used by other SQL database vendors to MySQL types. See [Section 11.8, “Using Data Types from Other Database Engines”](#).
- If you include a [USING](#) clause to specify an index type that is not permitted for a given storage engine, but there is another index type available that the engine can use without affecting query results, the engine uses the available type.
- If strict SQL mode is not enabled, a [VARCHAR](#) column with a length specification greater than 65535 is converted to [TEXT](#), and a [VARBINARY](#) column with a length specification greater than 65535 is converted to [BLOB](#). Otherwise, an error occurs in either of these cases.
- Specifying the [CHARACTER SET binary](#) attribute for a character data type causes the column to be created as the corresponding binary data type: [CHAR](#) becomes [BINARY](#), [VARCHAR](#) becomes [VARBINARY](#), and [TEXT](#) becomes [BLOB](#). For the [ENUM](#) and [SET](#) data types, this does not occur; they are created as declared. Suppose that you specify a table using this definition:

```
CREATE TABLE t
(
  c1 VARCHAR(10) CHARACTER SET binary,
  c2 TEXT CHARACTER SET binary,
  c3 ENUM('a','b','c') CHARACTER SET binary
);
```

The resulting table has this definition:

```
CREATE TABLE t
(
  c1 VARBINARY(10),
  c2 BLOB,
  c3 ENUM('a','b','c') CHARACTER SET binary
);
```

To see whether MySQL used a data type other than the one you specified, issue a [DESCRIBE](#) or [SHOW CREATE TABLE](#) statement after creating or altering the table.

Certain other data type changes can occur if you compress a table using [myisampack](#). See [Section 14.3.3.3, “Compressed Table Characteristics”](#).

### 13.1.15. CREATE TRIGGER Syntax

```
CREATE
  [DEFINER = { user | CURRENT_USER }]
  TRIGGER trigger_name trigger_time trigger_event
  ON tbl_name FOR EACH ROW trigger_body
```

This statement creates a new trigger. A trigger is a named database object that is associated with a table, and that activates when a particular event occurs for the table. The trigger becomes associated with the table named *tbl\_name*, which must refer to a permanent table. You cannot associate a trigger with a [TEMPORARY](#) table or a view.

[CREATE TRIGGER](#) requires the [TRIGGER \[756\]](#) privilege for the table associated with the trigger. The statement might also require the [SUPER \[756\]](#) privilege, depending on the [DEFINER](#) value, as described



later in this section. If binary logging is enabled, `CREATE TRIGGER` might require the `SUPER` [756] privilege, as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

The `DEFINER` clause determines the security context to be used when checking access privileges at trigger activation time. See later in this section for more information.

`trigger_time` is the trigger action time. It can be `BEFORE` or `AFTER` to indicate that the trigger activates before or after each row to be modified.

`trigger_event` indicates the kind of statement that activates the trigger. The `trigger_event` can be one of the following:

- `INSERT`: The trigger is activated whenever a new row is inserted into the table; for example, through `INSERT`, `LOAD DATA`, and `REPLACE` statements.
- `UPDATE`: The trigger is activated whenever a row is modified; for example, through `UPDATE` statements.
- `DELETE`: The trigger is activated whenever a row is deleted from the table; for example, through `DELETE` and `REPLACE` statements. However, `DROP TABLE` and `TRUNCATE TABLE` statements on the table do *not* activate this trigger, because they do not use `DELETE`. Dropping a partition does not activate `DELETE` triggers, either. See [Section 13.1.27, “TRUNCATE TABLE Syntax”](#).

It is important to understand that the `trigger_event` does not represent a literal type of SQL statement that activates the trigger so much as it represents a type of table operation. For example, an `INSERT` trigger is activated by not only `INSERT` statements but also `LOAD DATA` statements because both statements insert rows into a table.

A potentially confusing example of this is the `INSERT INTO ... ON DUPLICATE KEY UPDATE ...` syntax: a `BEFORE INSERT` trigger will activate for every row, followed by either an `AFTER INSERT` trigger or both the `BEFORE UPDATE` and `AFTER UPDATE` triggers, depending on whether there was a duplicate key for the row.

There cannot be two triggers for a given table that have the same trigger action time and event. For example, you cannot have two `BEFORE UPDATE` triggers for a table. But you can have a `BEFORE UPDATE` and a `BEFORE INSERT` trigger, or a `BEFORE UPDATE` and an `AFTER UPDATE` trigger.

`trigger_body` is the statement to execute when the trigger activates. If you want to execute multiple statements, use the `BEGIN ... END` compound statement construct. This also enables you to use the same statements that are permissible within stored routines. See [Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”](#). Some statements are not permitted in triggers; see [Section E.1, “Restrictions on Stored Programs”](#).

You can refer to columns in the subject table (the table associated with the trigger) by using the aliases `OLD` and `NEW`. `OLD.col_name` refers to a column of an existing row before it is updated or deleted. `NEW.col_name` refers to the column of a new row to be inserted or an existing row after it is updated.

MySQL stores the `sql_mode` [545] system variable setting that is in effect at the time a trigger is created, and always executes the trigger with this setting in force, *regardless of the server SQL mode in effect when the event begins executing*.



#### Note

Currently, cascaded foreign key actions do not activate triggers.

The `DEFINER` clause specifies the MySQL account to be used when checking access privileges at trigger activation time. If a `user` value is given, it should be a MySQL account specified as `'user_name'@'host_name'` (the same format used in the `GRANT` statement), `CURRENT_USER` [1267], or `CURRENT_USER()` [1267]. The default `DEFINER` value is the user who executes the `CREATE TRIGGER` statement. This is the same as specifying `DEFINER = CURRENT_USER` explicitly.

If you specify the `DEFINER` clause, these rules determine the valid `DEFINER` user values:

- If you do not have the `SUPER` [756] privilege, the only permitted `user` value is your own account, either specified literally or by using `CURRENT_USER` [1267]. You cannot set the definer to some other account.
- If you have the `SUPER` [756] privilege, you can specify any syntactically valid account name. If the account does not actually exist, a warning is generated.
- Although it is possible to create a trigger with a nonexistent `DEFINER` account, it is not a good idea for such triggers to be activated until the account actually does exist. Otherwise, the behavior with respect to privilege checking is undefined.

MySQL takes the `DEFINER` user into account when checking trigger privileges as follows:

- At `CREATE TRIGGER` time, the user who issues the statement must have the `TRIGGER` [756] privilege.
- At trigger activation time, privileges are checked against the `DEFINER` user. This user must have these privileges:
  - The `TRIGGER` [756] privilege.
  - The `SELECT` [755] privilege for the subject table if references to table columns occur using `OLD.col_name` or `NEW.col_name` in the trigger definition.
  - The `UPDATE` [756] privilege for the subject table if table columns are targets of `SET NEW.col_name = value` assignments in the trigger definition.
  - Whatever other privileges normally are required for the statements executed by the trigger.

For more information about trigger security, see [Section 18.6, “Access Control for Stored Programs and Views”](#).

Within a trigger, the `CURRENT_USER()` [1267] function returns the account used to check privileges at trigger activation time. This is the `DEFINER` user, not the user whose actions caused the trigger to be activated. For information about user auditing within triggers, see [Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”](#).

If you use `LOCK TABLES` to lock a table that has triggers, the tables used within the trigger are also locked, as described in [Section 13.3.5.2, “LOCK TABLES and Triggers”](#).

In MySQL 5.6, you can write triggers containing direct references to tables by name, such as the trigger named `testref` shown in this example:

```
CREATE TABLE test1(a1 INT);
CREATE TABLE test2(a2 INT);
CREATE TABLE test3(a3 INT NOT NULL AUTO_INCREMENT PRIMARY KEY);
CREATE TABLE test4(
  a4 INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  b4 INT DEFAULT 0
);

delimiter |

CREATE TRIGGER testref BEFORE INSERT ON test1
FOR EACH ROW BEGIN
  INSERT INTO test2 SET a2 = NEW.a1;
  DELETE FROM test3 WHERE a3 = NEW.a1;
  UPDATE test4 SET b4 = b4 + 1 WHERE a4 = NEW.a1;
END;

delimiter ;

INSERT INTO test3 (a3) VALUES
```



```
(NULL), (NULL), (NULL), (NULL), (NULL),
(NULL), (NULL), (NULL), (NULL), (NULL);
```

```
INSERT INTO test4 (a4) VALUES
(0), (0), (0), (0), (0), (0), (0), (0), (0), (0);
```

Suppose that you insert the following values into table `test1` as shown here:

```
mysql> INSERT INTO test1 VALUES
-> (1), (3), (1), (7), (1), (8), (4), (4);
Query OK, 8 rows affected (0.01 sec)
Records: 8 Duplicates: 0 Warnings: 0
```

As a result, the data in the four tables will be as follows:

```
mysql> SELECT * FROM test1;
```

```
+-----+
| a1 |
+-----+
|  1 |
|  3 |
|  1 |
|  7 |
|  1 |
|  8 |
|  4 |
|  4 |
+-----+
```

```
8 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM test2;
```

```
+-----+
| a2 |
+-----+
|  1 |
|  3 |
|  1 |
|  7 |
|  1 |
|  8 |
|  4 |
|  4 |
+-----+
```

```
8 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM test3;
```

```
+-----+
| a3 |
+-----+
|  2 |
|  5 |
|  6 |
|  9 |
| 10 |
+-----+
```

```
5 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM test4;
```

```
+-----+-----+
| a4 | b4 |
+-----+-----+
|  1 |  3 |
|  2 |  0 |
|  3 |  1 |
|  4 |  2 |
|  5 |  0 |
|  6 |  0 |
|  7 |  1 |
|  8 |  1 |
|  9 |  0 |
| 10 |  0 |
+-----+-----+
```

```
10 rows in set (0.00 sec)
```

### 13.1.16. CREATE VIEW Syntax

```
CREATE
  [OR REPLACE]
  [ALGORITHM = {UNDEFINED | MERGE | TEMPTABLE}]
  [DEFINER = { user | CURRENT_USER }]
  [SQL SECURITY { DEFINER | INVOKER }]
  VIEW view_name [(column_list)]
  AS select_statement
  [WITH [CASCADED | LOCAL] CHECK OPTION]
```

The `CREATE VIEW` statement creates a new view, or replaces an existing one if the `OR REPLACE` clause is given. If the view does not exist, `CREATE OR REPLACE VIEW` is the same as `CREATE VIEW`. If the view does exist, `CREATE OR REPLACE VIEW` is the same as `ALTER VIEW`.

The `select_statement` is a `SELECT` statement that provides the definition of the view. (When you select from the view, you select in effect using the `SELECT` statement.) `select_statement` can select from base tables or other views.

The view definition is “frozen” at creation time, so changes to the underlying tables afterward do not affect the view definition. For example, if a view is defined as `SELECT *` on a table, new columns added to the table later do not become part of the view.

The `ALGORITHM` clause affects how MySQL processes the view. The `DEFINER` and `SQL SECURITY` clauses specify the security context to be used when checking access privileges at view invocation time. The `WITH CHECK OPTION` clause can be given to constrain inserts or updates to rows in tables referenced by the view. These clauses are described later in this section.

The `CREATE VIEW` statement requires the `CREATE VIEW [754]` privilege for the view, and some privilege for each column selected by the `SELECT` statement. For columns used elsewhere in the `SELECT` statement you must have the `SELECT [755]` privilege. If the `OR REPLACE` clause is present, you must also have the `DROP [754]` privilege for the view. `CREATE VIEW` might also require the `SUPER [756]` privilege, depending on the `DEFINER` value, as described later in this section.

When a view is referenced, privilege checking occurs as described later in this section.

A view belongs to a database. By default, a new view is created in the default database. To create the view explicitly in a given database, specify the name as `db_name.view_name` when you create it:

```
mysql> CREATE VIEW test.v AS SELECT * FROM t;
```

Within a database, base tables and views share the same namespace, so a base table and a view cannot have the same name.

Columns retrieved by the `SELECT` statement can be simple references to table columns. They can also be expressions that use functions, constant values, operators, and so forth.

Views must have unique column names with no duplicates, just like base tables. By default, the names of the columns retrieved by the `SELECT` statement are used for the view column names. To define explicit names for the view columns, the optional `column_list` clause can be given as a list of comma-separated identifiers. The number of names in `column_list` must be the same as the number of columns retrieved by the `SELECT` statement.

Unqualified table or view names in the `SELECT` statement are interpreted with respect to the default database. A view can refer to tables or views in other databases by qualifying the table or view name with the proper database name.

A view can be created from many kinds of `SELECT` statements. It can refer to base tables or other views. It can use joins, `UNION`, and subqueries. The `SELECT` need not even refer to any tables. The following example defines a view that selects two columns from another table, as well as an expression calculated from those columns:

```
mysql> CREATE TABLE t (qty INT, price INT);
mysql> INSERT INTO t VALUES(3, 50);
mysql> CREATE VIEW v AS SELECT qty, price, qty*price AS value FROM t;
mysql> SELECT * FROM v;
+-----+-----+-----+
| qty | price | value |
+-----+-----+-----+
| 3 | 50 | 150 |
+-----+-----+-----+
```

A view definition is subject to the following restrictions:

- The `SELECT` statement cannot contain a subquery in the `FROM` clause.
- The `SELECT` statement cannot refer to system or user variables.
- Within a stored program, the definition cannot refer to program parameters or local variables.
- The `SELECT` statement cannot refer to prepared statement parameters.
- Any table or view referred to in the definition must exist. However, after a view has been created, it is possible to drop a table or view that the definition refers to. In this case, use of the view results in an error. To check a view definition for problems of this kind, use the `CHECK TABLE` statement.
- The definition cannot refer to a `TEMPORARY` table, and you cannot create a `TEMPORARY` view.
- Any tables named in the view definition must exist at definition time.
- You cannot associate a trigger with a view.
- Aliases for column names in the `SELECT` statement are checked against the maximum column length of 64 characters (not the maximum alias length of 256 characters).

`ORDER BY` is permitted in a view definition, but it is ignored if you select from a view using a statement that has its own `ORDER BY`.

For other options or clauses in the definition, they are added to the options or clauses of the statement that references the view, but the effect is undefined. For example, if a view definition includes a `LIMIT` clause, and you select from the view using a statement that has its own `LIMIT` clause, it is undefined which limit applies. This same principle applies to options such as `ALL`, `DISTINCT`, or `SQL_SMALL_RESULT` that follow the `SELECT` keyword, and to clauses such as `INTO`, `FOR UPDATE`, `LOCK IN SHARE MODE`, and `PROCEDURE`.

If you create a view and then change the query processing environment by changing system variables, that may affect the results that you get from the view:

```
mysql> CREATE VIEW v (mycol) AS SELECT 'abc';
Query OK, 0 rows affected (0.01 sec)

mysql> SET sql_mode = '';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT "mycol" FROM v;
+-----+
| mycol |
+-----+
| mycol |
+-----+
1 row in set (0.01 sec)

mysql> SET sql_mode = 'ANSI_QUOTES';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT "mycol" FROM v;
+-----+
| mycol |
+-----+
```

```
| abc |
+-----+
1 row in set (0.00 sec)
```

The `DEFINER` and `SQL SECURITY` clauses determine which MySQL account to use when checking access privileges for the view when a statement is executed that references the view. The valid `SQL SECURITY` characteristic values are `DEFINER` and `INVOKER`. These indicate that the required privileges must be held by the user who defined or invoked the view, respectively. The default `SQL SECURITY` value is `DEFINER`.

If a `user` value is given for the `DEFINER` clause, it should be a MySQL account specified as `'user_name'@'host_name'` (the same format used in the `GRANT` statement), `CURRENT_USER [1267]`, or `CURRENT_USER() [1267]`. The default `DEFINER` value is the user who executes the `CREATE VIEW` statement. This is the same as specifying `DEFINER = CURRENT_USER` explicitly.

If you specify the `DEFINER` clause, these rules determine the valid `DEFINER` user values:

- If you do not have the `SUPER [756]` privilege, the only valid `user` value is your own account, either specified literally or by using `CURRENT_USER [1267]`. You cannot set the definer to some other account.
- If you have the `SUPER [756]` privilege, you can specify any syntactically valid account name. If the account does not actually exist, a warning is generated.
- Although it is possible to create a view with a nonexistent `DEFINER` account, an error occurs when the view is referenced if the `SQL SECURITY` value is `DEFINER` but the definer account does not exist.

For more information about view security, see [Section 18.6, “Access Control for Stored Programs and Views”](#).

Within a view definition, `CURRENT_USER [1267]` returns the view's `DEFINER` value by default. For views defined with the `SQL SECURITY INVOKER` characteristic, `CURRENT_USER [1267]` returns the account for the view's invoker. For information about user auditing within views, see [Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”](#).

Within a stored routine that is defined with the `SQL SECURITY DEFINER` characteristic, `CURRENT_USER [1267]` returns the routine's `DEFINER` value. This also affects a view defined within such a routine, if the view definition contains a `DEFINER` value of `CURRENT_USER [1267]`.

View privileges are checked like this:

- At view definition time, the view creator must have the privileges needed to use the top-level objects accessed by the view. For example, if the view definition refers to table columns, the creator must have some privilege for each column in the select list of the definition, and the `SELECT [755]` privilege for each column used elsewhere in the definition. If the definition refers to a stored function, only the privileges needed to invoke the function can be checked. The privileges required at function invocation time can be checked only as it executes: For different invocations, different execution paths within the function might be taken.
- The user who references a view must have appropriate privileges to access it (`SELECT [755]` to select from it, `INSERT [755]` to insert into it, and so forth.)
- When a view has been referenced, privileges for objects accessed by the view are checked against the privileges held by the view `DEFINER` account or invoker, depending on whether the `SQL SECURITY` characteristic is `DEFINER` or `INVOKER`, respectively.
- If reference to a view causes execution of a stored function, privilege checking for statements executed within the function depend on whether the function `SQL SECURITY` characteristic is `DEFINER` or `INVOKER`. If the security characteristic is `DEFINER`, the function runs with the privileges

of the `DEFINER` account. If the characteristic is `INVOKER`, the function runs with the privileges determined by the view's `SQL SECURITY` characteristic.

Example: A view might depend on a stored function, and that function might invoke other stored routines. For example, the following view invokes a stored function `f()`:

```
CREATE VIEW v AS SELECT * FROM t WHERE t.id = f(t.name);
```

Suppose that `f()` contains a statement such as this:

```
IF name IS NULL then
  CALL p1();
ELSE
  CALL p2();
END IF;
```

The privileges required for executing statements within `f()` need to be checked when `f()` executes. This might mean that privileges are needed for `p1()` or `p2()`, depending on the execution path within `f()`. Those privileges must be checked at runtime, and the user who must possess the privileges is determined by the `SQL SECURITY` values of the view `v` and the function `f()`.

The `DEFINER` and `SQL SECURITY` clauses for views are extensions to standard SQL. In standard SQL, views are handled using the rules for `SQL SECURITY DEFINER`. The standard says that the definer of the view, which is the same as the owner of the view's schema, gets applicable privileges on the view (for example, [SELECT \[755\]](#)) and may grant them. MySQL has no concept of a schema “owner”, so MySQL adds a clause to identify the definer. The `DEFINER` clause is an extension where the intent is to have what the standard has; that is, a permanent record of who defined the view. This is why the default `DEFINER` value is the account of the view creator.

The optional `ALGORITHM` clause is a MySQL extension to standard SQL. It affects how MySQL processes the view. `ALGORITHM` takes three values: `MERGE`, `TEMPTABLE`, or `UNDEFINED`. The default algorithm is `UNDEFINED` if no `ALGORITHM` clause is present. For more information, see [Section 18.5.2, “View Processing Algorithms”](#).

Some views are updatable. That is, you can use them in statements such as `UPDATE`, `DELETE`, or `INSERT` to update the contents of the underlying table. For a view to be updatable, there must be a one-to-one relationship between the rows in the view and the rows in the underlying table. There are also certain other constructs that make a view nonupdatable.

The `WITH CHECK OPTION` clause can be given for an updatable view to prevent inserts or updates to rows except those for which the `WHERE` clause in the `select_statement` is true.

In a `WITH CHECK OPTION` clause for an updatable view, the `LOCAL` and `CASCADE` keywords determine the scope of check testing when the view is defined in terms of another view. The `LOCAL` keyword restricts the `CHECK OPTION` only to the view being defined. `CASCADE` causes the checks for underlying views to be evaluated as well. When neither keyword is given, the default is `CASCADE`.

For more information about updatable views and the `WITH CHECK OPTION` clause, see [Section 18.5.3, “Updatable and Insertable Views”](#).

### 13.1.17. DROP DATABASE Syntax

```
DROP {DATABASE | SCHEMA} [IF EXISTS] db_name
```

`DROP DATABASE` drops all tables in the database and deletes the database. Be very careful with this statement! To use `DROP DATABASE`, you need the `DROP [754]` privilege on the database. `DROP SCHEMA` is a synonym for `DROP DATABASE`.



#### Important

When a database is dropped, user privileges on the database are *not* automatically dropped. See [Section 13.7.1.4, “GRANT Syntax”](#).

IF EXISTS is used to prevent an error from occurring if the database does not exist.

If the default database is dropped, the default database is unset (the DATABASE() [1268] function returns NULL).

If you use DROP DATABASE on a symbolically linked database, both the link and the original database are deleted.

DROP DATABASE returns the number of tables that were removed. This corresponds to the number of .frm files removed.

The DROP DATABASE statement removes from the given database directory those files and directories that MySQL itself may create during normal operation:

- All files with the following extensions.

.BAK	.DAT	.HSH	.MRG
.MYD	.MYI	.TRG	.TRN
.db	.frm	.ibd	.ndb
.par			

- The db.opt file, if it exists.

If other files or directories remain in the database directory after MySQL removes those just listed, the database directory cannot be removed. In this case, you must remove any remaining files or directories manually and issue the DROP DATABASE statement again.

You can also drop databases with mysqladmin. See Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”.

### 13.1.18. DROP EVENT Syntax

```
DROP EVENT [IF EXISTS] event_name
```

This statement drops the event named event\_name. The event immediately ceases being active, and is deleted completely from the server.

If the event does not exist, the error ERROR 1517 (HY000): Unknown event 'event\_name' results. You can override this and cause the statement to generate a warning for nonexistent events instead using IF EXISTS.

This statement requires the EVENT [754] privilege for the schema to which the event to be dropped belongs.

### 13.1.19. DROP FUNCTION Syntax

The DROP FUNCTION statement is used to drop stored functions and user-defined functions (UDFs):

- For information about dropping stored functions, see Section 13.1.21, “DROP PROCEDURE and DROP FUNCTION Syntax”.
- For information about dropping user-defined functions, see Section 13.7.3.2, “DROP FUNCTION Syntax”.

### 13.1.20. DROP INDEX Syntax

```
DROP INDEX index_name ON tbl_name
    [algorithm_option | lock_option] ...

algorithm_option:
    ALGORITHM [=] {DEFAULT|INPLACE|COPY}
```

```
lock_option:
  LOCK [=] {DEFAULT|NONE|SHARED|EXCLUSIVE}
```

`DROP INDEX` drops the index named `index_name` from the table `tbl_name`. This statement is mapped to an `ALTER TABLE` statement to drop the index. See [Section 13.1.6, “ALTER TABLE Syntax”](#).

To drop a primary key, the index name is always `PRIMARY`, which must be specified as a quoted identifier because `PRIMARY` is a reserved word:

```
DROP INDEX `PRIMARY` ON t;
```

As of MySQL 5.6.6, the `ALGORITHM` and `LOCK` clauses may be given. These influence the table copying method and level of concurrency for reading and writing the table while its indexes are being modified. They have the same meaning as for the `ALTER TABLE` statement. For more information, see [Section 13.1.6, “ALTER TABLE Syntax”](#)

### 13.1.21. DROP PROCEDURE and DROP FUNCTION Syntax

```
DROP {PROCEDURE | FUNCTION} [IF EXISTS] sp_name
```

This statement is used to drop a stored procedure or function. That is, the specified routine is removed from the server. You must have the `ALTER ROUTINE [754]` privilege for the routine. (If the `automatic_sp_privileges` system variable is enabled, that privilege and `EXECUTE [754]` are granted automatically to the routine creator when the routine is created and dropped from the creator when the routine is dropped. See [Section 18.2.2, “Stored Routines and MySQL Privileges”](#).)

The `IF EXISTS` clause is a MySQL extension. It prevents an error from occurring if the procedure or function does not exist. A warning is produced that can be viewed with `SHOW WARNINGS`.

`DROP FUNCTION` is also used to drop user-defined functions (see [Section 13.7.3.2, “DROP FUNCTION Syntax”](#)).

### 13.1.22. DROP SERVER Syntax

```
DROP SERVER [ IF EXISTS ] server_name
```

Drops the server definition for the server named `server_name`. The corresponding row in the `mysql.servers` table is deleted. This statement requires the `SUPER [756]` privilege.

Dropping a server for a table does not affect any `FEDERATED` tables that used this connection information when they were created. See [Section 13.1.13, “CREATE SERVER Syntax”](#).

`DROP SERVER` does not cause an automatic commit.

In MySQL 5.6, `DROP SERVER` is not written to the binary log, regardless of the logging format that is in use.

Beginning with MySQL 5.6.11, `gtid_next [2022]` must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

### 13.1.23. DROP TABLE Syntax

```
DROP [TEMPORARY] TABLE [IF EXISTS]
  tbl_name [, tbl_name] ...
  [RESTRICT | CASCADE]
```

`DROP TABLE` removes one or more tables. You must have the `DROP [754]` privilege for each table. All table data and the table definition are *removed*, so *be careful* with this statement! If any of the tables named in the argument list do not exist, MySQL returns an error indicating by name which nonexisting tables it was unable to drop, but it also drops all of the tables in the list that do exist.



**Important**

When a table is dropped, user privileges on the table are *not* automatically dropped. See [Section 13.7.1.4, “GRANT Syntax”](#).

Note that for a partitioned table, `DROP TABLE` permanently removes the table definition, all of its partitions, and all of the data which was stored in those partitions. It also removes the partitioning definition (`.par`) file associated with the dropped table.

Use `IF EXISTS` to prevent an error from occurring for tables that do not exist. A `NOTE` is generated for each nonexistent table when using `IF EXISTS`. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

`RESTRICT` and `CASCADE` are permitted to make porting easier. In MySQL 5.6, they do nothing.

**Note**

`DROP TABLE` automatically commits the current active transaction, unless you use the `TEMPORARY` keyword.

The `TEMPORARY` keyword has the following effects:

- The statement drops only `TEMPORARY` tables.
- The statement does not end an ongoing transaction.
- No access rights are checked. (A `TEMPORARY` table is visible only to the session that created it, so no check is necessary.)

Using `TEMPORARY` is a good way to ensure that you do not accidentally drop a non-`TEMPORARY` table.

## 13.1.24. DROP TRIGGER Syntax

```
DROP TRIGGER [IF EXISTS] [schema_name.]trigger_name
```

This statement drops a trigger. The schema (database) name is optional. If the schema is omitted, the trigger is dropped from the default schema. `DROP TRIGGER` requires the `TRIGGER [756]` privilege for the table associated with the trigger.

Use `IF EXISTS` to prevent an error from occurring for a trigger that does not exist. A `NOTE` is generated for a nonexistent trigger when using `IF EXISTS`. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

Triggers for a table are also dropped if you drop the table.

## 13.1.25. DROP VIEW Syntax

```
DROP VIEW [IF EXISTS]
  view_name [, view_name] ...
  [RESTRICT | CASCADE]
```

`DROP VIEW` removes one or more views. You must have the `DROP [754]` privilege for each view. If any of the views named in the argument list do not exist, MySQL returns an error indicating by name which nonexistent views it was unable to drop, but it also drops all of the views in the list that do exist.

The `IF EXISTS` clause prevents an error from occurring for views that don't exist. When this clause is given, a `NOTE` is generated for each nonexistent view. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

`RESTRICT` and `CASCADE`, if given, are parsed and ignored.

## 13.1.26. RENAME TABLE Syntax

```
RENAME TABLE tbl_name TO new_tbl_name
```



```
[, tbl_name2 TO new_tbl_name2] ...
```

This statement renames one or more tables.

The rename operation is done atomically, which means that no other session can access any of the tables while the rename is running. For example, if you have an existing table `old_table`, you can create another table `new_table` that has the same structure but is empty, and then replace the existing table with the empty one as follows (assuming that `backup_table` does not already exist):

```
CREATE TABLE new_table (...);
RENAME TABLE old_table TO backup_table, new_table TO old_table;
```

If the statement renames more than one table, renaming operations are done from left to right. If you want to swap two table names, you can do so like this (assuming that `tmp_table` does not already exist):

```
RENAME TABLE old_table TO tmp_table,
             new_table TO old_table,
             tmp_table TO new_table;
```

As long as two databases are on the same file system, you can use `RENAME TABLE` to move a table from one database to another:

```
RENAME TABLE current_db.tbl_name TO other_db.tbl_name;
```

If there are any triggers associated with a table which is moved to a different database using `RENAME TABLE`, then the statement fails with the error `Trigger in wrong schema`.

`RENAME TABLE` also works for views, as long as you do not try to rename a view into a different database.

Any privileges granted specifically for the renamed table or view are not migrated to the new name. They must be changed manually.

When you execute `RENAME`, you cannot have any locked tables or active transactions. You must also have the `ALTER` [753] and `DROP` [754] privileges on the original table, and the `CREATE` [754] and `INSERT` [755] privileges on the new table.

If MySQL encounters any errors in a multiple-table rename, it does a reverse rename for all renamed tables to return everything to its original state.

You cannot use `RENAME` to rename a `TEMPORARY` table. However, you can use `ALTER TABLE` instead:

```
mysql> ALTER TABLE orig_name RENAME new_name;
```

### 13.1.27. TRUNCATE TABLE Syntax

```
TRUNCATE [TABLE] tbl_name
```

`TRUNCATE TABLE` empties a table completely. It requires the `DROP` [754] privilege.

Logically, `TRUNCATE TABLE` is similar to a `DELETE` statement that deletes all rows, or a sequence of `DROP TABLE` and `CREATE TABLE` statements. To achieve high performance, it bypasses the DML method of deleting data. Thus, it cannot be rolled back, it does not cause `ON DELETE` triggers to fire, and it cannot be performed for `InnoDB` tables with parent-child foreign key relationships.

Although `TRUNCATE TABLE` is similar to `DELETE`, it is classified as a DDL statement rather than a DML statement. It differs from `DELETE` in the following ways in MySQL 5.6:

- Truncate operations drop and re-create the table, which is much faster than deleting rows one by one, particularly for large tables.
- Truncate operations cause an implicit commit, and so cannot be rolled back.
- Truncation operations cannot be performed if the session holds an active table lock.

- `TRUNCATE TABLE` fails for an `InnoDB` table if there are any `FOREIGN KEY` constraints from other tables that reference the table. Foreign key constraints between columns of the same table are permitted.
- Truncation operations do not return a meaningful value for the number of deleted rows. The usual result is “0 rows affected,” which should be interpreted as “no information.”
- As long as the table format file `tbl_name.frm` is valid, the table can be re-created as an empty table with `TRUNCATE TABLE`, even if the data or index files have become corrupted.
- Any `AUTO_INCREMENT` value is reset to its start value. This is true even for `MyISAM` and `InnoDB`, which normally do not reuse sequence values.
- When used with partitioned tables, `TRUNCATE TABLE` preserves the partitioning; that is, the data and index files are dropped and re-created, while the partition definitions (`.par`) file is unaffected.
- The `TRUNCATE TABLE` statement does not invoke `ON DELETE` triggers.

`TRUNCATE TABLE` for a table closes all handlers for the table that were opened with `HANDLER OPEN`.

`TRUNCATE TABLE` is treated for purposes of binary logging and replication as `DROP TABLE` followed by `CREATE TABLE`—that is, as DDL rather than DML. This is due to the fact that, when using `InnoDB` and other transactional storage engines where the transaction isolation level does not permit statement-based logging (`READ COMMITTED` or `READ UNCOMMITTED`), the statement was not logged and replicated when using `STATEMENT` or `MIXED` logging mode. (Bug #36763) However, it is still applied on replication slaves using `InnoDB` in the manner described previously.

`TRUNCATE TABLE` can be used with Performance Schema summary tables, but the effect is to reset the summary columns to 0 or `NULL`, not to remove rows. See [Section 20.9.8, “Performance Schema Summary Tables”](#).

## 13.2. Data Manipulation Statements

### 13.2.1. `CALL` Syntax

```
CALL sp_name([parameter[,...]])
CALL sp_name[()]
```

The `CALL` statement invokes a stored procedure that was defined previously with `CREATE PROCEDURE`.

Stored procedures that take no arguments can be invoked without parentheses. That is, `CALL p()` and `CALL p` are equivalent.

`CALL` can pass back values to its caller using parameters that are declared as `OUT` or `INOUT` parameters. When the procedure returns, a client program can also obtain the number of rows affected for the final statement executed within the routine: At the SQL level, call the `ROW_COUNT()` [1272] function; from the C API, call the `mysql_affected_rows()` function.

To get back a value from a procedure using an `OUT` or `INOUT` parameter, pass the parameter by means of a user variable, and then check the value of the variable after the procedure returns. (If you are calling the procedure from within another stored procedure or function, you can also pass a routine parameter or local routine variable as an `IN` or `INOUT` parameter.) For an `INOUT` parameter, initialize its value before passing it to the procedure. The following procedure has an `OUT` parameter that the procedure sets to the current server version, and an `INOUT` value that the procedure increments by one from its current value:

```
CREATE PROCEDURE p (OUT ver_param VARCHAR(25), INOUT incr_param INT)
BEGIN
  # Set value of OUT parameter
  SELECT VERSION() INTO ver_param;
  # Increment value of INOUT parameter
  SET incr_param = incr_param + 1;
```

```
END;
```

Before calling the procedure, initialize the variable to be passed as the `INOUT` parameter. After calling the procedure, the values of the two variables will have been set or modified:

```
mysql> SET @increment = 10;
mysql> CALL p(@version, @increment);
mysql> SELECT @version, @increment;
+-----+-----+
| @version | @increment |
+-----+-----+
| 5.5.3-m3-log | 11 |
+-----+-----+
```

In prepared `CALL` statements used with `PREPARE` and `EXECUTE`, placeholders can be used for `IN` parameters. For `OUT` and `INOUT` parameters, placeholder support is available as of MySQL 5.5.3. These types of parameters can be used as follows:

```
mysql> SET @increment = 10;
mysql> PREPARE s FROM 'CALL p(?, ?)';
mysql> EXECUTE s USING @version, @increment;
mysql> SELECT @version, @increment;
+-----+-----+
| @version | @increment |
+-----+-----+
| 5.5.3-m3-log | 11 |
+-----+-----+
```

Before MySQL 5.5.3, placeholder support is not available for `OUT` or `INOUT` parameters. To work around this limitation for `OUT` and `INOUT` parameters, forego the use of placeholders; instead, refer to user variables in the `CALL` statement itself and do not specify them in the `EXECUTE` statement:

```
mysql> SET @increment = 10;
mysql> PREPARE s FROM 'CALL p(@version, @increment)';
mysql> EXECUTE s;
mysql> SELECT @version, @increment;
+-----+-----+
| @version | @increment |
+-----+-----+
| 5.5.0-m2-log | 11 |
+-----+-----+
```

To write C programs that use the `CALL` SQL statement to execute stored procedures that produce result sets, the `CLIENT_MULTI_RESULTS` flag must be enabled. This is because each `CALL` returns a result to indicate the call status, in addition to any result sets that might be returned by statements executed within the procedure. `CLIENT_MULTI_RESULTS` must also be enabled if `CALL` is used to execute any stored procedure that contains prepared statements. It cannot be determined when such a procedure is loaded whether those statements will produce result sets, so it is necessary to assume that they will.

`CLIENT_MULTI_RESULTS` can be enabled when you call `mysql_real_connect()`, either explicitly by passing the `CLIENT_MULTI_RESULTS` flag itself, or implicitly by passing `CLIENT_MULTI_STATEMENTS` (which also enables `CLIENT_MULTI_RESULTS`). In MySQL 5.6, `CLIENT_MULTI_RESULTS` is enabled by default.

To process the result of a `CALL` statement executed using `mysql_query()` or `mysql_real_query()`, use a loop that calls `mysql_next_result()` to determine whether there are more results. For an example, see [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

For programs written in a language that provides a MySQL interface, there is no native method prior to MySQL 5.5.3 for directly retrieving the results of `OUT` or `INOUT` parameters from `CALL` statements. To get the parameter values, pass user-defined variables to the procedure in the `CALL` statement and then execute a `SELECT` statement to produce a result set containing the variable values. To handle an `INOUT` parameter, execute a statement prior to the `CALL` that sets the corresponding user variable to the value to be passed to the procedure.

The following example illustrates the technique (without error checking) for the stored procedure `p` described earlier that has an `OUT` parameter and an `INOUT` parameter:

```
mysql_query(mysql, "SET @increment = 10");
mysql_query(mysql, "CALL p(@version, @increment)");
mysql_query(mysql, "SELECT @version, @increment");
result = mysql_store_result(mysql);
row = mysql_fetch_row(result);
mysql_free_result(result);
```

After the preceding code executes, `row[0]` and `row[1]` contain the values of `@version` and `@increment`, respectively.

In MySQL 5.6, C programs can use the prepared-statement interface to execute `CALL` statements and access `OUT` and `INOUT` parameters. This is done by processing the result of a `CALL` statement using a loop that calls `mysql_stmt_next_result()` to determine whether there are more results. For an example, see [Section 21.8.19, “C API Support for Prepared CALL Statements”](#). Languages that provide a MySQL interface can use prepared `CALL` statements to directly retrieve `OUT` and `INOUT` procedure parameters.

As of MySQL 5.6.6, metadata changes to objects referred to by stored programs are detected and cause automatic reparsing of the affected statements when the program is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

## 13.2.2. DELETE Syntax

`DELETE` is a DML statement that removes rows from a table.

### Single-table syntax:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE] FROM tbl_name
  [PARTITION (partition_name,...)]
  [WHERE where_condition]
  [ORDER BY ...]
  [LIMIT row_count]
```

The `DELETE` statement deletes rows from `tbl_name` and returns the number of deleted rows. To check the number of deleted rows, call the `ROW_COUNT()` [\[1272\]](#) function described in [Section 12.14, “Information Functions”](#).

### Main Clauses

The conditions in the optional `WHERE` clause identify which rows to delete. With no `WHERE` clause, all rows are deleted.

`where_condition` is an expression that evaluates to true for each row to be deleted. It is specified as described in [Section 13.2.9, “SELECT Syntax”](#).

If the `ORDER BY` clause is specified, the rows are deleted in the order that is specified. The `LIMIT` clause places a limit on the number of rows that can be deleted. These clauses apply to single-table deletes, but not multi-table deletes.

### Multiple-table syntax:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE]
  tbl_name [...] [, tbl_name [...]] ...
FROM table_references
  [WHERE where_condition]
```

Or:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE]
  FROM tbl_name [...] [, tbl_name [...]] ...
  USING table_references
```

```
[WHERE where_condition]
```

## Privileges

You need the [DELETE \[754\]](#) privilege on a table to delete rows from it. You need only the [SELECT \[755\]](#) privilege for any columns that are only read, such as those named in the [WHERE](#) clause.

## Performance

When you do not need to know the number of deleted rows, the [TRUNCATE TABLE](#) statement is a faster way to empty a table than a [DELETE](#) statement with no [WHERE](#) clause. Unlike [DELETE](#), [TRUNCATE TABLE](#) cannot be used within a transaction or if you have a lock on the table. See [Section 13.1.27, “TRUNCATE TABLE Syntax”](#) and [Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”](#).

The speed of delete operations may also be affected by factors discussed in [Section 8.2.2.3, “Speed of DELETE Statements”](#).

To ensure that a given [DELETE](#) statement does not take too much time, the MySQL-specific [LIMIT row\\_count](#) clause for [DELETE](#) specifies the maximum number of rows to be deleted. If the number of rows to delete is larger than the limit, repeat the [DELETE](#) statement until the number of affected rows is less than the [LIMIT](#) value.

## Subqueries

Currently, you cannot delete from a table and select from the same table in a subquery.

## Partitioned Tables

Beginning with MySQL 5.6.2, [DELETE](#) supports explicit partition selection using the [PARTITION](#) option, which takes a comma-separated list of the names of one or more partitions or subpartitions (or both) from which to select rows to be dropped. Partitions not included in the list are ignored. Given a partitioned table *t* with a partition named *p0*, executing the statement [DELETE FROM t PARTITION \(p0\)](#) has the same effect on the table as executing [ALTER TABLE t TRUNCATE PARTITION \(p0\)](#); in both cases, all rows in partition *p0* are dropped.

[PARTITION](#) can be used along with a [WHERE](#) condition, in which case the condition is tested only on rows in the listed partitions. For example, [DELETE FROM t PARTITION \(p0\) WHERE c < 5](#) deletes rows only from partition *p0* for which the condition *c < 5* is true; rows in any other partitions are not checked and thus not affected by the [DELETE](#).

The [PARTITION](#) option can also be used in multiple-table [DELETE](#) statements. You can use up to one such option per table named in the [FROM](#) option.

See [Section 17.5, “Partition Selection”](#), for more information and examples.

## Auto-Increment Columns

If you delete the row containing the maximum value for an [AUTO\\_INCREMENT](#) column, the value is not reused for a [MyISAM](#) or [InnoDB](#) table. If you delete all rows in the table with [DELETE FROM tbl\\_name](#) (without a [WHERE](#) clause) in [autocommit \[464\]](#) mode, the sequence starts over for all storage engines except [InnoDB](#) and [MyISAM](#). There are some exceptions to this behavior for [InnoDB](#) tables, as discussed in [Section 5.4.4, “AUTO\\_INCREMENT Handling in InnoDB”](#).

For [MyISAM](#) tables, you can specify an [AUTO\\_INCREMENT](#) secondary column in a multiple-column key. In this case, reuse of values deleted from the top of the sequence occurs even for [MyISAM](#) tables. See [Section 3.6.9, “Using AUTO\\_INCREMENT”](#).

## Modifiers

The [DELETE](#) statement supports the following modifiers:

- If you specify `LOW_PRIORITY`, the server delays execution of the `DELETE` until no other clients are reading from the table. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).
- For `MyISAM` tables, if you use the `QUICK` keyword, the storage engine does not merge index leaves during delete, which may speed up some kinds of delete operations.
- The `IGNORE` keyword causes MySQL to ignore all errors during the process of deleting rows. (Errors encountered during the parsing stage are processed in the usual manner.) Errors that are ignored due to the use of `IGNORE` are returned as warnings.

## Order of Deletion

If the `DELETE` statement includes an `ORDER BY` clause, rows are deleted in the order specified by the clause. This is useful primarily in conjunction with `LIMIT`. For example, the following statement finds rows matching the `WHERE` clause, sorts them by `timestamp_column`, and deletes the first (oldest) one:

```
DELETE FROM somelog WHERE user = 'jcole'
ORDER BY timestamp_column LIMIT 1;
```

`ORDER BY` also helps to delete rows in an order required to avoid referential integrity violations.

## InnoDB Tables

If you are deleting many rows from a large table, you may exceed the lock table size for an `InnoDB` table. To avoid this problem, or simply to minimize the time that the table remains locked, the following strategy (which does not use `DELETE` at all) might be helpful:

1. Select the rows *not* to be deleted into an empty table that has the same structure as the original table:

```
INSERT INTO t_copy SELECT * FROM t WHERE ... ;
```

2. Use `RENAME TABLE` to atomically move the original table out of the way and rename the copy to the original name:

```
RENAME TABLE t TO t_old, t_copy TO t;
```

3. Drop the original table:

```
DROP TABLE t_old;
```

No other sessions can access the tables involved while `RENAME TABLE` executes, so the rename operation is not subject to concurrency problems. See [Section 13.1.26, “RENAME TABLE Syntax”](#).

## MyISAM Tables

In `MyISAM` tables, deleted rows are maintained in a linked list and subsequent `INSERT` operations reuse old row positions. To reclaim unused space and reduce file sizes, use the `OPTIMIZE TABLE` statement or the `myisamchk` utility to reorganize tables. `OPTIMIZE TABLE` is easier to use, but `myisamchk` is faster. See [Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#), and [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

The `QUICK` modifier affects whether index leaves are merged for delete operations. `DELETE QUICK` is most useful for applications where index values for deleted rows are replaced by similar index values from rows inserted later. In this case, the holes left by deleted values are reused.

`DELETE QUICK` is not useful when deleted values lead to underfilled index blocks spanning a range of index values for which new inserts occur again. In this case, use of `QUICK` can lead to wasted space in the index that remains unreclaimed. Here is an example of such a scenario:

1. Create a table that contains an indexed `AUTO_INCREMENT` column.

2. Insert many rows into the table. Each insert results in an index value that is added to the high end of the index.
3. Delete a block of rows at the low end of the column range using `DELETE QUICK`.

In this scenario, the index blocks associated with the deleted index values become underfilled but are not merged with other index blocks due to the use of `QUICK`. They remain underfilled when new inserts occur, because new rows do not have index values in the deleted range. Furthermore, they remain underfilled even if you later use `DELETE` without `QUICK`, unless some of the deleted index values happen to lie in index blocks within or adjacent to the underfilled blocks. To reclaim unused index space under these circumstances, use `OPTIMIZE TABLE`.

If you are going to delete many rows from a table, it might be faster to use `DELETE QUICK` followed by `OPTIMIZE TABLE`. This rebuilds the index rather than performing many index block merge operations.

## Multi-Table Deletes

You can specify multiple tables in a `DELETE` statement to delete rows from one or more tables depending on the condition in the `WHERE` clause. You cannot use `ORDER BY` or `LIMIT` in a multiple-table `DELETE`. The `table_references` clause lists the tables involved in the join, as described in [Section 13.2.9.2, "JOIN Syntax"](#).

For the first multiple-table syntax, only matching rows from the tables listed before the `FROM` clause are deleted. For the second multiple-table syntax, only matching rows from the tables listed in the `FROM` clause (before the `USING` clause) are deleted. The effect is that you can delete rows from many tables at the same time and have additional tables that are used only for searching:

```
DELETE t1, t2 FROM t1 INNER JOIN t2 INNER JOIN t3
WHERE t1.id=t2.id AND t2.id=t3.id;
```

Or:

```
DELETE FROM t1, t2 USING t1 INNER JOIN t2 INNER JOIN t3
WHERE t1.id=t2.id AND t2.id=t3.id;
```

These statements use all three tables when searching for rows to delete, but delete matching rows only from tables `t1` and `t2`.

The preceding examples use `INNER JOIN`, but multiple-table `DELETE` statements can use other types of join permitted in `SELECT` statements, such as `LEFT JOIN`. For example, to delete rows that exist in `t1` that have no match in `t2`, use a `LEFT JOIN`:

```
DELETE t1 FROM t1 LEFT JOIN t2 ON t1.id=t2.id WHERE t2.id IS NULL;
```

The syntax permits `.*` after each `tbl_name` for compatibility with `Access`.

If you use a multiple-table `DELETE` statement involving `InnoDB` tables for which there are foreign key constraints, the MySQL optimizer might process tables in an order that differs from that of their parent/child relationship. In this case, the statement fails and rolls back. Instead, you should delete from a single table and rely on the `ON DELETE` capabilities that `InnoDB` provides to cause the other tables to be modified accordingly.



### Note

If you declare an alias for a table, you must use the alias when referring to the table:

```
DELETE t1 FROM test AS t1, test2 WHERE ...
```

Table aliases in a multiple-table `DELETE` should be declared only in the `table_references` part of the statement. Elsewhere, alias references are permitted but not alias declarations.

Correct:



```
DELETE a1, a2 FROM t1 AS a1 INNER JOIN t2 AS a2
WHERE a1.id=a2.id;
```

```
DELETE FROM a1, a2 USING t1 AS a1 INNER JOIN t2 AS a2
WHERE a1.id=a2.id;
```

Incorrect:

```
DELETE t1 AS a1, t2 AS a2 FROM t1 INNER JOIN t2
WHERE a1.id=a2.id;
```

```
DELETE FROM t1 AS a1, t2 AS a2 USING t1 INNER JOIN t2
WHERE a1.id=a2.id;
```

### 13.2.3. DO Syntax

```
DO expr [, expr] ...
```

**DO** executes the expressions but does not return any results. In most respects, **DO** is shorthand for **SELECT *expr*, ...**, but has the advantage that it is slightly faster when you do not care about the result.

**DO** is useful primarily with functions that have side effects, such as [RELEASE\\_LOCK\(\)](#) [1280].

### 13.2.4. HANDLER Syntax

```
HANDLER tbl_name OPEN [ [AS] alias ]
```

```
HANDLER tbl_name READ index_name { = | <= | >= | < | > } (value1,value2,...)
[ WHERE where_condition ] [LIMIT ... ]
```

```
HANDLER tbl_name READ index_name { FIRST | NEXT | PREV | LAST }
[ WHERE where_condition ] [LIMIT ... ]
```

```
HANDLER tbl_name READ { FIRST | NEXT }
[ WHERE where_condition ] [LIMIT ... ]
```

```
HANDLER tbl_name CLOSE
```

The **HANDLER** statement provides direct access to table storage engine interfaces. It is available for **InnoDB** and **MyISAM** tables.

The **HANDLER ... OPEN** statement opens a table, making it accessible using subsequent **HANDLER ... READ** statements. This table object is not shared by other sessions and is not closed until the session calls **HANDLER ... CLOSE** or the session terminates. If you open the table using an alias, further references to the open table with other **HANDLER** statements must use the alias rather than the table name.

The first **HANDLER ... READ** syntax fetches a row where the index specified satisfies the given values and the **WHERE** condition is met. If you have a multiple-column index, specify the index column values as a comma-separated list. Either specify values for all the columns in the index, or specify values for a leftmost prefix of the index columns. Suppose that an index **my\_idx** includes three columns named **col\_a**, **col\_b**, and **col\_c**, in that order. The **HANDLER** statement can specify values for all three columns in the index, or for the columns in a leftmost prefix. For example:

```
HANDLER ... READ my_idx = (col_a_val,col_b_val,col_c_val) ...
```

```
HANDLER ... READ my_idx = (col_a_val,col_b_val) ...
```

```
HANDLER ... READ my_idx = (col_a_val) ...
```

To employ the **HANDLER** interface to refer to a table's **PRIMARY KEY**, use the quoted identifier **`PRIMARY`**:

```
HANDLER tbl_name READ `PRIMARY` ...
```

The second **HANDLER ... READ** syntax fetches a row from the table in index order that matches the **WHERE** condition.



The third `HANDLER ... READ` syntax fetches a row from the table in natural row order that matches the `WHERE` condition. It is faster than `HANDLER tbl_name READ index_name` when a full table scan is desired. Natural row order is the order in which rows are stored in a `MyISAM` table data file. This statement works for `InnoDB` tables as well, but there is no such concept because there is no separate data file.

Without a `LIMIT` clause, all forms of `HANDLER ... READ` fetch a single row if one is available. To return a specific number of rows, include a `LIMIT` clause. It has the same syntax as for the `SELECT` statement. See [Section 13.2.9, “SELECT Syntax”](#).

`HANDLER ... CLOSE` closes a table that was opened with `HANDLER ... OPEN`.

There are several reasons to use the `HANDLER` interface instead of normal `SELECT` statements:

- `HANDLER` is faster than `SELECT`:
  - A designated storage engine handler object is allocated for the `HANDLER ... OPEN`. The object is reused for subsequent `HANDLER` statements for that table; it need not be reinitialized for each one.
  - There is less parsing involved.
  - There is no optimizer or query-checking overhead.
  - The handler interface does not have to provide a consistent look of the data (for example, [dirty reads](#) are permitted), so the storage engine can use optimizations that `SELECT` does not normally permit.
- `HANDLER` makes it easier to port to MySQL applications that use a low-level `ISAM`-like interface. (See [Section 14.2.9, “InnoDB Integration with memcached”](#) for an alternative way to adapt applications that use the key-value store paradigm.)
- `HANDLER` enables you to traverse a database in a manner that is difficult (or even impossible) to accomplish with `SELECT`. The `HANDLER` interface is a more natural way to look at data when working with applications that provide an interactive user interface to the database.

`HANDLER` is a somewhat low-level statement. For example, it does not provide consistency. That is, `HANDLER ... OPEN` does *not* take a snapshot of the table, and does *not* lock the table. This means that after a `HANDLER ... OPEN` statement is issued, table data can be modified (by the current session or other sessions) and these modifications might be only partially visible to `HANDLER ... NEXT` or `HANDLER ... PREV` scans.

An open handler can be closed and marked for reopen, in which case the handler loses its position in the table. This occurs when both of the following circumstances are true:

- Any session executes `FLUSH TABLES` or DDL statements on the handler's table.
- The session in which the handler is open executes non-`HANDLER` statements that use tables.

`TRUNCATE TABLE` for a table closes all handlers for the table that were opened with `HANDLER OPEN`.

If a table is flushed with `FLUSH TABLES tbl_name WITH READ LOCK` was opened with `HANDLER`, the handler is implicitly flushed and loses its position.

### 13.2.5. INSERT Syntax

```
INSERT [LOW_PRIORITY | DELAYED | HIGH_PRIORITY] [IGNORE]
      [INTO] tbl_name
      [PARTITION (partition_name,...)]
      [(col_name,...)]
      {VALUES | VALUE} ({expr | DEFAULT},...),(...),...
      [ ON DUPLICATE KEY UPDATE
```

```
col_name=expr
[, col_name=expr] ... ]
```

Or:

```
INSERT [LOW_PRIORITY | DELAYED | HIGH_PRIORITY] [IGNORE]
[INTO] tbl_name
[PARTITION (partition_name,...)]
SET col_name={expr | DEFAULT}, ...
[ ON DUPLICATE KEY UPDATE
col_name=expr
[, col_name=expr] ... ]
```

Or:

```
INSERT [LOW_PRIORITY | HIGH_PRIORITY] [IGNORE]
[INTO] tbl_name
[PARTITION (partition_name,...)]
[(col_name,...)]
SELECT ...
[ ON DUPLICATE KEY UPDATE
col_name=expr
[, col_name=expr] ... ]
```

`INSERT` inserts new rows into an existing table. The `INSERT ... VALUES` and `INSERT ... SET` forms of the statement insert rows based on explicitly specified values. The `INSERT ... SELECT` form inserts rows selected from another table or tables. `INSERT ... SELECT` is discussed further in [Section 13.2.5.1, “INSERT ... SELECT Syntax”](#).

In MySQL 5.6.2 and later, when inserting into a partitioned table, you can control which partitions and subpartitions accept new rows. The `PARTITION` option takes a comma-separated list of the names of one or more partitions or subpartitions (or both) of the table. If any of the rows to be inserted by a given `INSERT` statement do not match one of the partitions listed, the `INSERT` statement fails with the error `Found a row not matching the given partition set`. See [Section 17.5, “Partition Selection”](#), for more information and examples.

You can use `REPLACE` instead of `INSERT` to overwrite old rows. `REPLACE` is the counterpart to `INSERT IGNORE` in the treatment of new rows that contain unique key values that duplicate old rows: The new rows are used to replace the old rows rather than being discarded. See [Section 13.2.8, “REPLACE Syntax”](#).

`tbl_name` is the table into which rows should be inserted. The columns for which the statement provides values can be specified as follows:

- You can provide a comma-separated list of column names following the table name. In this case, a value for each named column must be provided by the `VALUES` list or the `SELECT` statement.
- If you do not specify a list of column names for `INSERT ... VALUES` or `INSERT ... SELECT`, values for every column in the table must be provided by the `VALUES` list or the `SELECT` statement. If you do not know the order of the columns in the table, use `DESCRIBE tbl_name` to find out.
- The `SET` clause indicates the column names explicitly.

Column values can be given in several ways:

- If you are not running in strict SQL mode, any column not explicitly given a value is set to its default (explicit or implicit) value. For example, if you specify a column list that does not name all the columns in the table, unnamed columns are set to their default values. Default value assignment is described in [Section 11.5, “Data Type Default Values”](#). See also [Section 1.8.6.3, “Constraints on Invalid Data”](#).

If you want an `INSERT` statement to generate an error unless you explicitly specify values for all columns that do not have a default value, you should use strict mode. See [Section 5.1.7, “Server SQL Modes”](#).

- Use the keyword `DEFAULT` to set a column explicitly to its default value. This makes it easier to write `INSERT` statements that assign values to all but a few columns, because it enables you to avoid writing an incomplete `VALUES` list that does not include a value for each column in the table. Otherwise, you would have to write out the list of column names corresponding to each value in the `VALUES` list.

You can also use `DEFAULT(col_name)` [1275] as a more general form that can be used in expressions to produce a given column's default value.

- If both the column list and the `VALUES` list are empty, `INSERT` creates a row with each column set to its default value:

```
INSERT INTO tbl_name () VALUES();
```

In strict mode, an error occurs if any column doesn't have a default value. Otherwise, MySQL uses the implicit default value for any column that does not have an explicitly defined default.

- You can specify an expression `expr` to provide a column value. This might involve type conversion if the type of the expression does not match the type of the column, and conversion of a given value can result in different inserted values depending on the data type. For example, inserting the string `'1999.0e-2'` into an `INT`, `FLOAT`, `DECIMAL(10,6)`, or `YEAR` column results in the values `1999`, `19.9921`, `19.992100`, and `1999` being inserted, respectively. The reason the value stored in the `INT` and `YEAR` columns is `1999` is that the string-to-integer conversion looks only at as much of the initial part of the string as may be considered a valid integer or year. For the floating-point and fixed-point columns, the string-to-floating-point conversion considers the entire string a valid floating-point value.

An expression `expr` can refer to any column that was set earlier in a value list. For example, you can do this because the value for `col2` refers to `col1`, which has previously been assigned:

```
INSERT INTO tbl_name (col1,col2) VALUES(15,col1*2);
```

But the following is not legal, because the value for `col1` refers to `col2`, which is assigned after `col1`:

```
INSERT INTO tbl_name (col1,col2) VALUES(col2*2,15);
```

One exception involves columns that contain `AUTO_INCREMENT` values. Because the `AUTO_INCREMENT` value is generated after other value assignments, any reference to an `AUTO_INCREMENT` column in the assignment returns a `0`.

`INSERT` statements that use `VALUES` syntax can insert multiple rows. To do this, include multiple lists of column values, each enclosed within parentheses and separated by commas. Example:

```
INSERT INTO tbl_name (a,b,c) VALUES(1,2,3),(4,5,6),(7,8,9);
```

The values list for each row must be enclosed within parentheses. The following statement is illegal because the number of values in the list does not match the number of column names:

```
INSERT INTO tbl_name (a,b,c) VALUES(1,2,3,4,5,6,7,8,9);
```

`VALUE` is a synonym for `VALUES` in this context. Neither implies anything about the number of values lists, and either may be used whether there is a single values list or multiple lists.

The affected-rows value for an `INSERT` can be obtained using the `ROW_COUNT()` [1272] function (see Section 12.14, "Information Functions"), or the `mysql_affected_rows()` C API function (see Section 21.8.6.1, "`mysql_affected_rows()`").

If you use an `INSERT ... VALUES` statement with multiple value lists or `INSERT ... SELECT`, the statement returns an information string in this format:

```
Records: 100 Duplicates: 0 Warnings: 0
```

`Records` indicates the number of rows processed by the statement. (This is not necessarily the number of rows actually inserted because `Duplicates` can be nonzero.) `Duplicates` indicates the number of rows that could not be inserted because they would duplicate some existing unique index value. `Warnings` indicates the number of attempts to insert column values that were problematic in some way. Warnings can occur under any of the following conditions:

- Inserting `NULL` into a column that has been declared `NOT NULL`. For multiple-row `INSERT` statements or `INSERT INTO ... SELECT` statements, the column is set to the implicit default value for the column data type. This is `0` for numeric types, the empty string ( `' '` ) for string types, and the “zero” value for date and time types. `INSERT INTO ... SELECT` statements are handled the same way as multiple-row inserts because the server does not examine the result set from the `SELECT` to see whether it returns a single row. (For a single-row `INSERT`, no warning occurs when `NULL` is inserted into a `NOT NULL` column. Instead, the statement fails with an error.)
- Setting a numeric column to a value that lies outside the column's range. The value is clipped to the closest endpoint of the range.
- Assigning a value such as `'10.34 a'` to a numeric column. The trailing nonnumeric text is stripped off and the remaining numeric part is inserted. If the string value has no leading numeric part, the column is set to `0`.
- Inserting a string into a string column (`CHAR`, `VARCHAR`, `TEXT`, or `BLOB`) that exceeds the column's maximum length. The value is truncated to the column's maximum length.
- Inserting a value into a date or time column that is illegal for the data type. The column is set to the appropriate zero value for the type.

If you are using the C API, the information string can be obtained by invoking the `mysql_info()` function. See [Section 21.8.6.35](#), “`mysql_info()`”.

If `INSERT` inserts a row into a table that has an `AUTO_INCREMENT` column, you can find the value used for that column by using the SQL `LAST_INSERT_ID()` [1269] function. From within the C API, use the `mysql_insert_id()` function. However, you should note that the two functions do not always behave identically. The behavior of `INSERT` statements with respect to `AUTO_INCREMENT` columns is discussed further in [Section 12.14](#), “Information Functions”, and [Section 21.8.6.37](#), “`mysql_insert_id()`”.

The `INSERT` statement supports the following modifiers:

- If you use the `DELAYED` keyword, the server puts the row or rows to be inserted into a buffer, and the client issuing the `INSERT DELAYED` statement can then continue immediately. If the table is in use, the server holds the rows. When the table is free, the server begins inserting rows, checking periodically to see whether there are any new read requests for the table. If there are, the delayed row queue is suspended until the table becomes free again. See [Section 13.2.5.2](#), “`INSERT DELAYED Syntax`”.

`DELAYED` is ignored with `INSERT ... SELECT` or `INSERT ... ON DUPLICATE KEY UPDATE`.

`DELAYED` is also disregarded for an `INSERT` that uses functions accessing tables or triggers, or that is called from a function or a trigger.



#### Note

As of MySQL 5.6.6, `INSERT DELAYED` is deprecated, and will be removed in a future release. Use `INSERT` (without `DELAYED`) instead.

- If you use the `LOW_PRIORITY` keyword, execution of the `INSERT` is delayed until no other clients are reading from the table. This includes other clients that began reading while existing clients are reading, and while the `INSERT LOW_PRIORITY` statement is waiting. It is possible, therefore, for a client that issues an `INSERT LOW_PRIORITY` statement to wait for a very long time (or even

forever) in a read-heavy environment. (This is in contrast to `INSERT DELAYED`, which lets the client continue at once.) Note that `LOW_PRIORITY` should normally not be used with `MyISAM` tables because doing so disables concurrent inserts. See [Section 8.10.3, “Concurrent Inserts”](#).

If you specify `HIGH_PRIORITY`, it overrides the effect of the `--low-priority-updates [430]` option if the server was started with that option. It also causes concurrent inserts not to be used. See [Section 8.10.3, “Concurrent Inserts”](#).

`LOW_PRIORITY` and `HIGH_PRIORITY` affect only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).

- If you use the `IGNORE` keyword, errors that occur while executing the `INSERT` statement are treated as warnings instead. For example, without `IGNORE`, a row that duplicates an existing `UNIQUE` index or `PRIMARY KEY` value in the table causes a duplicate-key error and the statement is aborted. With `IGNORE`, the row still is not inserted, but no error is issued.

`IGNORE` has a similar effect on inserts into partitioned tables where no partition matching a given value is found. Without `IGNORE`, such `INSERT` statements are aborted with an error; however, when `INSERT IGNORE` is used, the insert operation fails silently for the row containing the unmatched value, but any rows that are matched are inserted. For an example, see [Section 17.2.2, “LIST Partitioning”](#).

Data conversions that would trigger errors abort the statement if `IGNORE` is not specified. With `IGNORE`, invalid values are adjusted to the closest values and inserted; warnings are produced but the statement does not abort. You can determine with the `mysql_info()` C API function how many rows were actually inserted into the table.

- If you specify `ON DUPLICATE KEY UPDATE`, and a row is inserted that would cause a duplicate value in a `UNIQUE` index or `PRIMARY KEY`, an `UPDATE` of the old row is performed. The affected-rows value per row is 1 if the row is inserted as a new row and 2 if an existing row is updated. See [Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”](#).

Inserting into a table requires the `INSERT [755]` privilege for the table. If the `ON DUPLICATE KEY UPDATE` clause is used and a duplicate key causes an `UPDATE` to be performed instead, the statement requires the `UPDATE [756]` privilege for the columns to be updated. For columns that are read but not modified you need only the `SELECT [755]` privilege (such as for a column referenced only on the right hand side of an `col_name=expr` assignment in an `ON DUPLICATE KEY UPDATE` clause).

Prior to MySQL 5.6.6, an `INSERT` that affected a partitioned table using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the table. This was true even for `INSERT ... PARTITION` statements. (This did not and does not occur with storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, MySQL uses partition lock pruning, so that only partitions into which rows are inserted are actually locked. For more information, see [Section 17.6.4, “Partitioning and Locking”](#).

### 13.2.5.1. INSERT ... SELECT Syntax

```
INSERT [LOW_PRIORITY | HIGH_PRIORITY] [IGNORE]
  [INTO] tbl_name
  [PARTITION (partition_name,...)]
  [(col_name,...)]
  SELECT ...
  [ ON DUPLICATE KEY UPDATE col_name=expr, ... ]
```

With `INSERT ... SELECT`, you can quickly insert many rows into a table from one or many tables. For example:

```
INSERT INTO tbl_temp2 (fld_id)
  SELECT tbl_temp1.fld_order_id
  FROM tbl_temp1 WHERE tbl_temp1.fld_order_id > 100;
```

The following conditions hold for a `INSERT ... SELECT` statements:

- Specify `IGNORE` to ignore rows that would cause duplicate-key violations.
- `DELAYED` is ignored with `INSERT ... SELECT`.
- The target table of the `INSERT` statement may appear in the `FROM` clause of the `SELECT` part of the query. (This was not possible in some older versions of MySQL.) However, you cannot insert into a table and select from the same table in a subquery.

When selecting from and inserting into a table at the same time, MySQL creates a temporary table to hold the rows from the `SELECT` and then inserts those rows into the target table. However, it remains true that you cannot use `INSERT INTO t ... SELECT ... FROM t` when `t` is a `TEMPORARY` table, because `TEMPORARY` tables cannot be referred to twice in the same statement (see [Section C.5.7.2, “TEMPORARY Table Problems”](#)).

- `AUTO_INCREMENT` columns work as usual.
- To ensure that the binary log can be used to re-create the original tables, MySQL does not permit concurrent inserts for `INSERT ... SELECT` statements.
- To avoid ambiguous column reference problems when the `SELECT` and the `INSERT` refer to the same table, provide a unique alias for each table used in the `SELECT` part, and qualify column names in that part with the appropriate alias.

Starting with MySQL 5.6.2, you can explicitly select which partitions or subpartitions (or both) of the source or target table (or both) are to be used with a `PARTITION` option following the name of the table. When `PARTITION` is used with the name of the source table in the `SELECT` portion of the statement, rows are selected only from the partitions or subpartitions named in its partition list. When `PARTITION` is used with the name of the target table for the `INSERT` portion of the statement, then it must be possible to insert all rows selected into the partitions or subpartitions named in the partition list following the option, else the `INSERT ... SELECT` statement fails. For more information and examples, see [Section 17.5, “Partition Selection”](#).

In the values part of `ON DUPLICATE KEY UPDATE`, you can refer to columns in other tables, as long as you do not use `GROUP BY` in the `SELECT` part. One side effect is that you must qualify nonunique column names in the values part.

The order in which rows are returned by a `SELECT` statement with no `ORDER BY` clause is not determined. This means that, when using replication, there is no guarantee that such a `SELECT` returns rows in the same order on the master and the slave; this can lead to inconsistencies between them. To prevent this from occurring, you should always write `INSERT ... SELECT` statements that are to be replicated as `INSERT ... SELECT ... ORDER BY column`. The choice of `column` does not matter as long as the same order for returning the rows is enforced on both the master and the slave. See also [Section 16.4.1.15, “Replication and LIMIT”](#).

Due to this issue, beginning with MySQL 5.6.4, `INSERT ... SELECT ON DUPLICATE KEY UPDATE` and `INSERT IGNORE ... SELECT` statements are flagged as unsafe for statement-based replication. With this change, such statements produce a warning in the log when using statement-based mode and are logged using the row-based format when using `MIXED` mode. (Bug #11758262, Bug #50439)

See also [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

Prior to MySQL 5.6.6, an `INSERT ... SELECT` statement that acted on partitioned tables using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the source and target tables. (This did not and does not occur with tables using storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, only those partitions of the source table that are actually read are locked, although all partitions of the target table are locked. See [Section 17.6.4, “Partitioning and Locking”](#), for more information.

### 13.2.5.2. INSERT DELAYED Syntax

```
INSERT DELAYED ...
```



The `DELAYED` option for the `INSERT` statement is a MySQL extension to standard SQL that can be used for certain kinds of tables (such as `MyISAM`). When a client uses `INSERT DELAYED`, it gets an okay from the server at once, and the row is queued to be inserted when the table is not in use by any other thread.



### Note

`INSERT DELAYED` is slower than a normal `INSERT` if the table is not otherwise in use. There is also the additional overhead for the server to handle a separate thread for each table for which there are delayed rows. This means that you should use `INSERT DELAYED` only when you are really sure that you need it.

As of MySQL 5.6.6, `INSERT DELAYED` is deprecated, and will be removed in a future release. Use `INSERT` (without `DELAYED`) instead.

The queued rows are held only in memory until they are inserted into the table. This means that if you terminate `mysqld` forcibly (for example, with `kill -9`) or if `mysqld` dies unexpectedly, *any queued rows that have not been written to disk are lost*.

There are some constraints on the use of `DELAYED`:

- `INSERT DELAYED` works only with `MyISAM`, `MEMORY`, `ARCHIVE`, and `BLACKHOLE` tables. For engines that do not support `DELAYED`, an error occurs.
- An error occurs for `INSERT DELAYED` if used with a table that has been locked with `LOCK TABLES` because the insert must be handled by a separate thread, not by the session that holds the lock.
- For `MyISAM` tables, if there are no free blocks in the middle of the data file, concurrent `SELECT` and `INSERT` statements are supported. Under these circumstances, you very seldom need to use `INSERT DELAYED` with `MyISAM`.
- `INSERT DELAYED` should be used only for `INSERT` statements that specify value lists. The server ignores `DELAYED` for `INSERT ... SELECT` or `INSERT ... ON DUPLICATE KEY UPDATE` statements.
- Because the `INSERT DELAYED` statement returns immediately, before the rows are inserted, you cannot use `LAST_INSERT_ID()` [1269] to get the `AUTO_INCREMENT` value that the statement might generate.
- `DELAYED` rows are not visible to `SELECT` statements until they actually have been inserted.
- Prior to MySQL 5.6, `INSERT DELAYED` was treated as a normal `INSERT` if the statement inserted multiple rows, binary logging was enabled, and the global logging format was statement-based (that is, whenever `binlog_format` [2010] was set to `STATEMENT`). Beginning with MySQL 5.6, `INSERT DELAYED` is always handled as a simple `INSERT` (that is, without the `DELAYED` option) whenever the value of `binlog_format` [2010] is `STATEMENT` or `MIXED`. (In the latter case, the statement no longer triggers a switch to row-based logging, and so is logged using the statement-based format.)

This does not apply when using row-based binary logging mode (`binlog_format` [2010] set to `ROW`), in which `INSERT DELAYED` statements are always executed using the `DELAYED` option as specified, and logged as row-update events.

- `DELAYED` is ignored on slave replication servers, so that `INSERT DELAYED` is treated as a normal `INSERT` on slaves. This is because `DELAYED` could cause the slave to have different data than the master.
- Pending `INSERT DELAYED` statements are lost if a table is write locked and `ALTER TABLE` is used to modify the table structure.
- `INSERT DELAYED` is not supported for views.
- `INSERT DELAYED` is not supported for partitioned tables.

The following describes in detail what happens when you use the `DELAYED` option to `INSERT` or `REPLACE`. In this description, the “thread” is the thread that received an `INSERT DELAYED` statement and “handler” is the thread that handles all `INSERT DELAYED` statements for a particular table.

- When a thread executes a `DELAYED` statement for a table, a handler thread is created to process all `DELAYED` statements for the table, if no such handler already exists.
- The thread checks whether the handler has previously acquired a `DELAYED` lock; if not, it tells the handler thread to do so. The `DELAYED` lock can be obtained even if other threads have a `READ` or `WRITE` lock on the table. However, the handler waits for all `ALTER TABLE` locks or `FLUSH TABLES` statements to finish, to ensure that the table structure is up to date.
- The thread executes the `INSERT` statement, but instead of writing the row to the table, it puts a copy of the final row into a queue that is managed by the handler thread. Any syntax errors are noticed by the thread and reported to the client program.
- The client cannot obtain from the server the number of duplicate rows or the `AUTO_INCREMENT` value for the resulting row, because the `INSERT` returns before the insert operation has been completed. (If you use the C API, the `mysql_info()` function does not return anything meaningful, for the same reason.)
- The binary log is updated by the handler thread when the row is inserted into the table. In case of multiple-row inserts, the binary log is updated when the first row is inserted.
- Each time that `delayed_insert_limit` [476] rows are written, the handler checks whether any `SELECT` statements are still pending. If so, it permits these to execute before continuing.
- When the handler has no more rows in its queue, the table is unlocked. If no new `INSERT DELAYED` statements are received within `delayed_insert_timeout` [476] seconds, the handler terminates.
- If more than `delayed_queue_size` [477] rows are pending in a specific handler queue, the thread requesting `INSERT DELAYED` waits until there is room in the queue. This is done to ensure that `mysqld` does not use all memory for the delayed memory queue.
- The handler thread shows up in the MySQL process list with `delayed_insert` in the `Command` column. It is killed if you execute a `FLUSH TABLES` statement or kill it with `KILL thread_id`. However, before exiting, it first stores all queued rows into the table. During this time it does not accept any new `INSERT` statements from other threads. If you execute an `INSERT DELAYED` statement after this, a new handler thread is created.

Note that this means that `INSERT DELAYED` statements have higher priority than normal `INSERT` statements if there is an `INSERT DELAYED` handler running. Other update statements have to wait until the `INSERT DELAYED` queue is empty, someone terminates the handler thread (with `KILL thread_id`), or someone executes a `FLUSH TABLES`.

- The following status variables provide information about `INSERT DELAYED` statements.

Status Variable	Meaning
<code>Delayed_insert_threads</code> [588]	Number of handler threads
<code>Delayed_writes</code> [588]	Number of rows written with <code>INSERT DELAYED</code>
<code>Not_flushed_delayed_rows</code> [594]	Number of rows waiting to be written

You can view these variables by issuing a `SHOW STATUS` statement or by executing a `mysqladmin extended-status` command.

### 13.2.5.3. `INSERT ... ON DUPLICATE KEY UPDATE` Syntax

If you specify `ON DUPLICATE KEY UPDATE`, and a row is inserted that would cause a duplicate value in a `UNIQUE` index or `PRIMARY KEY`, MySQL performs an `UPDATE` of the old row. For example, if



column `a` is declared as `UNIQUE` and contains the value `1`, the following two statements have similar effect:

```
INSERT INTO table (a,b,c) VALUES (1,2,3)
ON DUPLICATE KEY UPDATE c=c+1;

UPDATE table SET c=c+1 WHERE a=1;
```

(The effects are not identical for an `InnoDB` table where `a` is an auto-increment column. With an auto-increment column, an `INSERT` statement increases the auto-increment value but `UPDATE` does not.)

The `ON DUPLICATE KEY UPDATE` clause can contain multiple column assignments, separated by commas.

With `ON DUPLICATE KEY UPDATE`, the affected-rows value per row is 1 if the row is inserted as a new row, and 2 if an existing row is updated.

If column `b` is also unique, the `INSERT` is equivalent to this `UPDATE` statement instead:

```
UPDATE table SET c=c+1 WHERE a=1 OR b=2 LIMIT 1;
```

If `a=1 OR b=2` matches several rows, only *one* row is updated. In general, you should try to avoid using an `ON DUPLICATE KEY UPDATE` clause on tables with multiple unique indexes.

You can use the `VALUES(col_name)` [1281] function in the `UPDATE` clause to refer to column values from the `INSERT` portion of the `INSERT ... ON DUPLICATE KEY UPDATE` statement. In other words, `VALUES(col_name)` [1281] in the `ON DUPLICATE KEY UPDATE` clause refers to the value of `col_name` that would be inserted, had no duplicate-key conflict occurred. This function is especially useful in multiple-row inserts. The `VALUES()` [1281] function is meaningful only in `INSERT ... UPDATE` statements and returns `NULL` otherwise. Example:

```
INSERT INTO table (a,b,c) VALUES (1,2,3),(4,5,6)
ON DUPLICATE KEY UPDATE c=VALUES(a)+VALUES(b);
```

That statement is identical to the following two statements:

```
INSERT INTO table (a,b,c) VALUES (1,2,3)
ON DUPLICATE KEY UPDATE c=3;
INSERT INTO table (a,b,c) VALUES (4,5,6)
ON DUPLICATE KEY UPDATE c=9;
```

If a table contains an `AUTO_INCREMENT` column and `INSERT ... ON DUPLICATE KEY UPDATE` inserts or updates a row, the `LAST_INSERT_ID()` [1269] function returns the `AUTO_INCREMENT` value.

The `DELAYED` option is ignored when you use `ON DUPLICATE KEY UPDATE`.

Because the results of `INSERT ... SELECT` statements depend on the ordering of rows from the `SELECT` and this order cannot always be guaranteed, it is possible when logging `INSERT ... SELECT ON DUPLICATE KEY UPDATE` statements for the master and the slave to diverge. Thus, in MySQL 5.6.4 and later, `INSERT ... SELECT ON DUPLICATE KEY UPDATE` statements are flagged as unsafe for statement-based replication. With this change, such statements produce a warning in the log when using statement-based mode and are logged using the row-based format when using `MIXED` mode. In addition, beginning with MySQL 5.6.6, an `INSERT ... ON DUPLICATE KEY UPDATE` statement against a table having more than one unique or primary key is also marked as unsafe. (Bug #11765650, Bug #58637) See also [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

Prior to MySQL 5.6.6, an `INSERT ... ON DUPLICATE KEY UPDATE` on a partitioned table using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the table. (This did not and does not occur with tables using storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, such statements lock only those partitions in which a partitioning key column is updated. See [Section 17.6.4, “Partitioning and Locking”](#), for more information.

## 13.2.6. LOAD DATA INFILE Syntax

```
LOAD DATA [LOW_PRIORITY | CONCURRENT] [LOCAL] INFILE 'file_name'
[REPLACE | IGNORE]
INTO TABLE tbl_name
[PARTITION (partition_name,...)]
[CHARACTER SET charset_name]
[ {FIELDS | COLUMNS}
  [TERMINATED BY 'string']
  [[OPTIONALLY] ENCLOSED BY 'char']
  [ESCAPED BY 'char']
]
[LINES
  [STARTING BY 'string']
  [TERMINATED BY 'string']
]
[IGNORE number {LINES | ROWS}]
[(col_name_or_user_var,...)]
[SET col_name = expr,...]
```

The `LOAD DATA INFILE` statement reads rows from a text file into a table at a very high speed. `LOAD DATA INFILE` is the complement of `SELECT ... INTO OUTFILE`. (See [Section 13.2.9.1, “SELECT ... INTO Syntax”](#).) To write data from a table to a file, use `SELECT ... INTO OUTFILE`. To read the file back into a table, use `LOAD DATA INFILE`. The syntax of the `FIELDS` and `LINES` clauses is the same for both statements. Both clauses are optional, but `FIELDS` must precede `LINES` if both are specified.

You can also load data files by using the `mysqlimport` utility; it operates by sending a `LOAD DATA INFILE` statement to the server. The `--local [300]` option causes `mysqlimport` to read data files from the client host. You can specify the `--compress [299]` option to get better performance over slow networks if the client and server support the compressed protocol. See [Section 4.5.5, “mysqlimport — A Data Import Program”](#).

For more information about the efficiency of `INSERT` versus `LOAD DATA INFILE` and speeding up `LOAD DATA INFILE`, see [Section 8.2.2.1, “Speed of INSERT Statements”](#).

The file name must be given as a literal string. On Windows, specify backslashes in path names as forward slashes or doubled backslashes. The `character_set_filesystem [468]` system variable controls the interpretation of the file name.

In MySQL 5.6.2 and later, `LOAD DATA` supports explicit partition selection using the `PARTITION` option with a comma-separated list of more or more names of partitions, subpartitions, or both. When this option is used, if any rows from the file cannot be inserted into any of the partitions or subpartitions named in the list, the statement fails with the error `Found a row not matching the given partition set`. For more information, see [Section 17.5, “Partition Selection”](#).

For partitioned tables using storage engines that employ table locks, such as `MyISAM`, `LOAD DATA` cannot prune any partition locks. This does not apply to tables using storage engines which employ row-level locking, such as `InnoDB`. For more information, see [Section 17.6.4, “Partitioning and Locking”](#).

The character set indicated by the `character_set_database [468]` system variable is used to interpret the information in the file. `SET NAMES` and the setting of `character_set_client [467]` do not affect interpretation of input. If the contents of the input file use a character set that differs from the default, it is usually preferable to specify the character set of the file by using the `CHARACTER SET` clause. A character set of `binary` specifies “no conversion.”

`LOAD DATA INFILE` interprets all fields in the file as having the same character set, regardless of the data types of the columns into which field values are loaded. For proper interpretation of file contents, you must ensure that it was written with the correct character set. For example, if you write a data file with `mysqldump -T` or by issuing a `SELECT ... INTO OUTFILE` statement in `mysql`, be sure to use a `--default-character-set` option so that output is written in the character set to be used when the file is loaded with `LOAD DATA INFILE`.

**Note**

It is not possible to load data files that use the `ucs2`, `utf16`, `utf16le`, or `utf32` character set.

If you use `LOW_PRIORITY`, execution of the `LOAD DATA` statement is delayed until no other clients are reading from the table. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).

If you specify `CONCURRENT` with a `MyISAM` table that satisfies the condition for concurrent inserts (that is, it contains no free blocks in the middle), other threads can retrieve data from the table while `LOAD DATA` is executing. This option affects the performance of `LOAD DATA` a bit, even if no other thread is using the table at the same time.

With row-based replication, `CONCURRENT` is replicated regardless of MySQL version. With statement-based replication `CONCURRENT` is not replicated prior to MySQL 5.5.1 (see Bug #34628). For more information, see [Section 16.4.1.16, “Replication and LOAD DATA INFILE”](#).

The `LOCAL` keyword affects expected location of the file and error handling, as described later. `LOCAL` works only if your server and your client both have been configured to permit it. For example, if `mysqld` was started with `--local-infile=0` [496], `LOCAL` does not work. See [Section 6.1.6, “Security Issues with LOAD DATA LOCAL”](#).

The `LOCAL` keyword affects where the file is expected to be found:

- If `LOCAL` is specified, the file is read by the client program on the client host and sent to the server. The file can be given as a full path name to specify its exact location. If given as a relative path name, the name is interpreted relative to the directory in which the client program was started.

When using `LOCAL` with `LOAD DATA`, a copy of the file is created in the server's temporary directory. This is *not* the directory determined by the value of `tmpdir` [559] or `slave_load_tmpdir` [1992], but rather the operating system's temporary directory, and is not configurable in the MySQL Server. (Typically the system temporary directory is `/tmp` on Linux systems and `C:\WINDOWS\TEMP` on Windows.) Lack of sufficient space for the copy in this directory can cause the `LOAD DATA LOCAL` statement to fail.

- If `LOCAL` is not specified, the file must be located on the server host and is read directly by the server. The server uses the following rules to locate the file:
  - If the file name is an absolute path name, the server uses it as given.
  - If the file name is a relative path name with one or more leading components, the server searches for the file relative to the server's data directory.
  - If a file name with no leading components is given, the server looks for the file in the database directory of the default database.

In the non-`LOCAL` case, these rules mean that a file named as `./myfile.txt` is read from the server's data directory, whereas the file named as `myfile.txt` is read from the database directory of the default database. For example, if `db1` is the default database, the following `LOAD DATA` statement reads the file `data.txt` from the database directory for `db1`, even though the statement explicitly loads the file into a table in the `db2` database:

```
LOAD DATA INFILE 'data.txt' INTO TABLE db2.my_table;
```

For security reasons, when reading text files located on the server, the files must either reside in the database directory or be readable by all. Also, to use `LOAD DATA INFILE` on server files, you must have the `FILE` [754] privilege. See [Section 6.2.1, “Privileges Provided by MySQL”](#). For non-`LOCAL` load operations, if the `secure_file_priv` [538] system variable is set to a nonempty directory name, the file to be loaded must be located in that directory.

Using `LOCAL` is a bit slower than letting the server access the files directly, because the contents of the file must be sent over the connection by the client to the server. On the other hand, you do not need the `FILE [754]` privilege to load local files.

`LOCAL` also affects error handling:

- With `LOAD DATA INFILE`, data-interpretation and duplicate-key errors terminate the operation.
- With `LOAD DATA LOCAL INFILE`, data-interpretation and duplicate-key errors become warnings and the operation continues because the server has no way to stop transmission of the file in the middle of the operation. For duplicate-key errors, this is the same as if `IGNORE` is specified. `IGNORE` is explained further later in this section.

The `REPLACE` and `IGNORE` keywords control handling of input rows that duplicate existing rows on unique key values:

- If you specify `REPLACE`, input rows replace existing rows. In other words, rows that have the same value for a primary key or unique index as an existing row. See [Section 13.2.8, “REPLACE Syntax”](#).
- If you specify `IGNORE`, input rows that duplicate an existing row on a unique key value are skipped.
- If you do not specify either option, the behavior depends on whether the `LOCAL` keyword is specified. Without `LOCAL`, an error occurs when a duplicate key value is found, and the rest of the text file is ignored. With `LOCAL`, the default behavior is the same as if `IGNORE` is specified; this is because the server has no way to stop transmission of the file in the middle of the operation.

To ignore foreign key constraints during the load operation, issue a `SET foreign_key_checks = 0` statement before executing `LOAD DATA`.

If you use `LOAD DATA INFILE` on an empty `MyISAM` table, all nonunique indexes are created in a separate batch (as for `REPAIR TABLE`). Normally, this makes `LOAD DATA INFILE` much faster when you have many indexes. In some extreme cases, you can create the indexes even faster by turning them off with `ALTER TABLE ... DISABLE KEYS` before loading the file into the table and using `ALTER TABLE ... ENABLE KEYS` to re-create the indexes after loading the file. See [Section 8.2.2.1, “Speed of INSERT Statements”](#).

For both the `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` statements, the syntax of the `FIELDS` and `LINES` clauses is the same. Both clauses are optional, but `FIELDS` must precede `LINES` if both are specified.

If you specify a `FIELDS` clause, each of its subclauses (`TERMINATED BY`, [OPTIONALLY] `ENCLOSED BY`, and `ESCAPED BY`) is also optional, except that you must specify at least one of them.

If you specify no `FIELDS` or `LINES` clause, the defaults are the same as if you had written this:

```
FIELDS TERMINATED BY '\t' ENCLOSED BY '' ESCAPED BY '\\'  
LINES TERMINATED BY '\n' STARTING BY ''
```

(Backslash is the MySQL escape character within strings in SQL statements, so to specify a literal backslash, you must specify two backslashes for the value to be interpreted as a single backslash. The escape sequences `'\t'` and `'\n'` specify tab and newline characters, respectively.)

In other words, the defaults cause `LOAD DATA INFILE` to act as follows when reading input:

- Look for line boundaries at newlines.
- Do not skip over any line prefix.
- Break lines into fields at tabs.
- Do not expect fields to be enclosed within any quoting characters.

- Interpret characters preceded by the escape character “\” as escape sequences. For example, “\t”, “\n”, and “\\” signify tab, newline, and backslash, respectively. See the discussion of [FIELDS ESCAPED BY](#) later for the full list of escape sequences.

Conversely, the defaults cause `SELECT ... INTO OUTFILE` to act as follows when writing output:

- Write tabs between fields.
- Do not enclose fields within any quoting characters.
- Use “\” to escape instances of tab, newline, or “\” that occur within field values.
- Write newlines at the ends of lines.



#### Note

If you have generated the text file on a Windows system, you might have to use `LINES TERMINATED BY '\r\n'` to read the file properly, because Windows programs typically use two characters as a line terminator. Some programs, such as `WordPad`, might use `\r` as a line terminator when writing files. To read such files, use `LINES TERMINATED BY '\r'`.

If all the lines you want to read in have a common prefix that you want to ignore, you can use `LINES STARTING BY 'prefix_string'` to skip over the prefix, *and anything before it*. If a line does not include the prefix, the entire line is skipped. Suppose that you issue the following statement:

```
LOAD DATA INFILE '/tmp/test.txt' INTO TABLE test
  FIELDS TERMINATED BY ',' LINES STARTING BY 'xxx';
```

If the data file looks like this:

```
xxx"abc",1
something xxx"def",2
"ghi",3
```

The resulting rows will be (`"abc",1`) and (`"def",2`). The third row in the file is skipped because it does not contain the prefix.

The `IGNORE number LINES` option can be used to ignore lines at the start of the file. For example, you can use `IGNORE 1 LINES` to skip over an initial header line containing column names:

```
LOAD DATA INFILE '/tmp/test.txt' INTO TABLE test IGNORE 1 LINES;
```

When you use `SELECT ... INTO OUTFILE` in tandem with `LOAD DATA INFILE` to write data from a database into a file and then read the file back into the database later, the field- and line-handling options for both statements must match. Otherwise, `LOAD DATA INFILE` will not interpret the contents of the file properly. Suppose that you use `SELECT ... INTO OUTFILE` to write a file with fields delimited by commas:

```
SELECT * INTO OUTFILE 'data.txt'
  FIELDS TERMINATED BY ','
  FROM table2;
```

To read the comma-delimited file back in, the correct statement would be:

```
LOAD DATA INFILE 'data.txt' INTO TABLE table2
  FIELDS TERMINATED BY ',';
```

If instead you tried to read in the file with the statement shown following, it wouldn't work because it instructs `LOAD DATA INFILE` to look for tabs between fields:

```
LOAD DATA INFILE 'data.txt' INTO TABLE table2
  FIELDS TERMINATED BY '\t';
```

The likely result is that each input line would be interpreted as a single field.

`LOAD DATA INFILE` can be used to read files obtained from external sources. For example, many programs can export data in comma-separated values (CSV) format, such that lines have fields separated by commas and enclosed within double quotation marks, with an initial line of column names. If the lines in such a file are terminated by carriage return/newline pairs, the statement shown here illustrates the field- and line-handling options you would use to load the file:

```
LOAD DATA INFILE 'data.txt' INTO TABLE tbl_name
  FIELDS TERMINATED BY ',' ENCLOSED BY '"'
  LINES TERMINATED BY '\r\n'
  IGNORE 1 LINES;
```

If the input values are not necessarily enclosed within quotation marks, use `OPTIONALLY` before the `ENCLOSED BY` keywords.

Any of the field- or line-handling options can specify an empty string ( `' '` ). If not empty, the `FIELDS [OPTIONALLY] ENCLOSED BY` and `FIELDS ESCAPED BY` values must be a single character. The `FIELDS TERMINATED BY`, `LINES STARTING BY`, and `LINES TERMINATED BY` values can be more than one character. For example, to write lines that are terminated by carriage return/linefeed pairs, or to read a file containing such lines, specify a `LINES TERMINATED BY '\r\n'` clause.

To read a file containing jokes that are separated by lines consisting of `%%`, you can do this

```
CREATE TABLE jokes
  (a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  joke TEXT NOT NULL);
LOAD DATA INFILE '/tmp/jokes.txt' INTO TABLE jokes
  FIELDS TERMINATED BY ''
  LINES TERMINATED BY '\n%%\n' (joke);
```

`FIELDS [OPTIONALLY] ENCLOSED BY` controls quoting of fields. For output (`SELECT ... INTO OUTFILE`), if you omit the word `OPTIONALLY`, all fields are enclosed by the `ENCLOSED BY` character. An example of such output (using a comma as the field delimiter) is shown here:

```
"1","a string","100.20"
"2","a string containing a , comma","102.20"
"3","a string containing a \" quote","102.20"
"4","a string containing a \", quote and comma","102.20"
```

If you specify `OPTIONALLY`, the `ENCLOSED BY` character is used only to enclose values from columns that have a string data type (such as `CHAR`, `BINARY`, `TEXT`, or `ENUM`):

```
1,"a string",100.20
2,"a string containing a , comma",102.20
3,"a string containing a \" quote",102.20
4,"a string containing a \", quote and comma",102.20
```

Note that occurrences of the `ENCLOSED BY` character within a field value are escaped by prefixing them with the `ESCAPED BY` character. Also note that if you specify an empty `ESCAPED BY` value, it is possible to inadvertently generate output that cannot be read properly by `LOAD DATA INFILE`. For example, the preceding output just shown would appear as follows if the escape character is empty. Observe that the second field in the fourth line contains a comma following the quote, which (erroneously) appears to terminate the field:

```
1,"a string",100.20
2,"a string containing a , comma",102.20
3,"a string containing a " quote",102.20
4,"a string containing a ", quote and comma",102.20
```

For input, the `ENCLOSED BY` character, if present, is stripped from the ends of field values. (This is true regardless of whether `OPTIONALLY` is specified; `OPTIONALLY` has no effect on input interpretation.) Occurrences of the `ENCLOSED BY` character preceded by the `ESCAPED BY` character are interpreted as part of the current field value.



If the field begins with the **ENCLOSED BY** character, instances of that character are recognized as terminating a field value only if followed by the field or line **TERMINATED BY** sequence. To avoid ambiguity, occurrences of the **ENCLOSED BY** character within a field value can be doubled and are interpreted as a single instance of the character. For example, if **ENCLOSED BY ' '** is specified, quotation marks are handled as shown here:

```
"The ""BIG"" boss" -> The "BIG" boss
The "BIG" boss     -> The "BIG" boss
The ""BIG"" boss   -> The ""BIG"" boss
```

**FIELDS ESCAPED BY** controls how to read or write special characters:

- For input, if the **FIELDS ESCAPED BY** character is not empty, occurrences of that character are stripped and the following character is taken literally as part of a field value. Some two-character sequences that are exceptions, where the first character is the escape character. These sequences are shown in the following table (using “\” for the escape character). The rules for **NULL** handling are described later in this section.

Character	Escape Sequence
\0	An ASCII NUL (0x00) character
\b	A backspace character
\n	A newline (linefeed) character
\r	A carriage return character
\t	A tab character.
\z	ASCII 26 (Control+Z)
\N	NULL

For more information about “\”-escape syntax, see [Section 9.1.1, “String Literals”](#).

If the **FIELDS ESCAPED BY** character is empty, escape-sequence interpretation does not occur.

- For output, if the **FIELDS ESCAPED BY** character is not empty, it is used to prefix the following characters on output:
  - The **FIELDS ESCAPED BY** character
  - The **FIELDS [OPTIONALLY] ENCLOSED BY** character
  - The first character of the **FIELDS TERMINATED BY** and **LINES TERMINATED BY** values
  - ASCII 0 (what is actually written following the escape character is ASCII “0”, not a zero-valued byte)

If the **FIELDS ESCAPED BY** character is empty, no characters are escaped and **NULL** is output as **NULL**, not **\N**. It is probably not a good idea to specify an empty escape character, particularly if field values in your data contain any of the characters in the list just given.

In certain cases, field- and line-handling options interact:

- If **LINES TERMINATED BY** is an empty string and **FIELDS TERMINATED BY** is nonempty, lines are also terminated with **FIELDS TERMINATED BY**.
- If the **FIELDS TERMINATED BY** and **FIELDS ENCLOSED BY** values are both empty ( ' ' ), a fixed-row (nondelimited) format is used. With fixed-row format, no delimiters are used between fields (but you can still have a line terminator). Instead, column values are read and written using a field width wide enough to hold all values in the field. For **TINYINT**, **SMALLINT**, **MEDIUMINT**, **INT**, and **BIGINT**, the field widths are 4, 6, 8, 11, and 20, respectively, no matter what the declared display width is.



`LINES TERMINATED BY` is still used to separate lines. If a line does not contain all fields, the rest of the columns are set to their default values. If you do not have a line terminator, you should set this to `' '`. In this case, the text file must contain all fields for each row.

Fixed-row format also affects handling of `NULL` values, as described later. Note that fixed-size format does not work if you are using a multi-byte character set.

Handling of `NULL` values varies according to the `FIELDS` and `LINES` options in use:

- For the default `FIELDS` and `LINES` values, `NULL` is written as a field value of `\N` for output, and a field value of `\N` is read as `NULL` for input (assuming that the `ESCAPED BY` character is `"\"`).
- If `FIELDS ENCLOSED BY` is not empty, a field containing the literal word `NULL` as its value is read as a `NULL` value. This differs from the word `NULL` enclosed within `FIELDS ENCLOSED BY` characters, which is read as the string `'NULL'`.
- If `FIELDS ESCAPED BY` is empty, `NULL` is written as the word `NULL`.
- With fixed-row format (which is used when `FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` are both empty), `NULL` is written as an empty string. Note that this causes both `NULL` values and empty strings in the table to be indistinguishable when written to the file because both are written as empty strings. If you need to be able to tell the two apart when reading the file back in, you should not use fixed-row format.

An attempt to load `NULL` into a `NOT NULL` column causes assignment of the implicit default value for the column's data type and a warning, or an error in strict SQL mode. Implicit default values are discussed in [Section 11.5, "Data Type Default Values"](#).

Some cases are not supported by `LOAD DATA INFILE`:

- Fixed-size rows (`FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` both empty) and `BLOB` or `TEXT` columns.
- If you specify one separator that is the same as or a prefix of another, `LOAD DATA INFILE` cannot interpret the input properly. For example, the following `FIELDS` clause would cause problems:

```
FIELDS TERMINATED BY ' ' ENCLOSED BY ' '
```

- If `FIELDS ESCAPED BY` is empty, a field value that contains an occurrence of `FIELDS ENCLOSED BY` or `LINES TERMINATED BY` followed by the `FIELDS TERMINATED BY` value causes `LOAD DATA INFILE` to stop reading a field or line too early. This happens because `LOAD DATA INFILE` cannot properly determine where the field or line value ends.

The following example loads all columns of the `persondata` table:

```
LOAD DATA INFILE 'persondata.txt' INTO TABLE persondata;
```

By default, when no column list is provided at the end of the `LOAD DATA INFILE` statement, input lines are expected to contain a field for each table column. If you want to load only some of a table's columns, specify a column list:

```
LOAD DATA INFILE 'persondata.txt' INTO TABLE persondata (col1,col2,...);
```

You must also specify a column list if the order of the fields in the input file differs from the order of the columns in the table. Otherwise, MySQL cannot tell how to match input fields with table columns.

The column list can contain either column names or user variables. With user variables, the `SET` clause enables you to perform transformations on their values before assigning the result to columns.

User variables in the `SET` clause can be used in several ways. The following example uses the first input column directly for the value of `t1.column1`, and assigns the second input column to a user variable that is subjected to a division operation before being used for the value of `t1.column2`:

```
LOAD DATA INFILE 'file.txt'
INTO TABLE t1
(column1, @var1)
SET column2 = @var1/100;
```

The `SET` clause can be used to supply values not derived from the input file. The following statement sets `column3` to the current date and time:

```
LOAD DATA INFILE 'file.txt'
INTO TABLE t1
(column1, column2)
SET column3 = CURRENT_TIMESTAMP;
```

You can also discard an input value by assigning it to a user variable and not assigning the variable to a table column:

```
LOAD DATA INFILE 'file.txt'
INTO TABLE t1
(column1, @dummy, column2, @dummy, column3);
```

Use of the column/variable list and `SET` clause is subject to the following restrictions:

- Assignments in the `SET` clause should have only column names on the left hand side of assignment operators.
- You can use subqueries in the right hand side of `SET` assignments. A subquery that returns a value to be assigned to a column may be a scalar subquery only. Also, you cannot use a subquery to select from the table that is being loaded.
- Lines ignored by an `IGNORE` clause are not processed for the column/variable list or `SET` clause.
- User variables cannot be used when loading data with fixed-row format because user variables do not have a display width.

When processing an input line, `LOAD DATA` splits it into fields and uses the values according to the column/variable list and the `SET` clause, if they are present. Then the resulting row is inserted into the table. If there are `BEFORE INSERT` or `AFTER INSERT` triggers for the table, they are activated before or after inserting the row, respectively.

If an input line has too many fields, the extra fields are ignored and the number of warnings is incremented.

If an input line has too few fields, the table columns for which input fields are missing are set to their default values. Default value assignment is described in [Section 11.5, “Data Type Default Values”](#).

An empty field value is interpreted different from a missing field:

- For string types, the column is set to the empty string.
- For numeric types, the column is set to 0.
- For date and time types, the column is set to the appropriate “zero” value for the type. See [Section 11.3, “Date and Time Types”](#).

These are the same values that result if you assign an empty string explicitly to a string, numeric, or date or time type explicitly in an `INSERT` or `UPDATE` statement.

Treatment of empty or incorrect field values differs from that just described if the SQL mode is set to a restrictive value. For example, if `sql_mode= 'TRADITIONAL [545]`, conversion of an empty value or a value such as `'x'` for a numeric column results in an error, not conversion to 0. (With `LOCAL`, warnings occur rather than errors, even with a restrictive `sql_mode [545]` value, because the server has no way to stop transmission of the file in the middle of the operation.)

`TIMESTAMP` columns are set to the current date and time only if there is a `NULL` value for the column (that is, `\N`) and the column is not declared to permit `NULL` values, or if the `TIMESTAMP` column's default value is the current timestamp and it is omitted from the field list when a field list is specified.

`LOAD DATA INFILE` regards all input as strings, so you cannot use numeric values for `ENUM` or `SET` columns the way you can with `INSERT` statements. All `ENUM` and `SET` values must be specified as strings.

`BIT` values cannot be loaded using binary notation (for example, `b'011010'`). To work around this, specify the values as regular integers and use the `SET` clause to convert them so that MySQL performs a numeric type conversion and loads them into the `BIT` column properly:

```
shell> cat /tmp/bit_test.txt
2
127
shell> mysql test
mysql> LOAD DATA INFILE '/tmp/bit_test.txt'
-> INTO TABLE bit_test (@var1) SET b = CAST(@var1 AS UNSIGNED);
Query OK, 2 rows affected (0.00 sec)
Records: 2 Deleted: 0 Skipped: 0 Warnings: 0

mysql> SELECT BIN(b+0) FROM bit_test;
+-----+
| bin(b+0) |
+-----+
| 10      |
| 1111111 |
+-----+
2 rows in set (0.00 sec)
```

On Unix, if you need `LOAD DATA` to read from a pipe, you can use the following technique (the example loads a listing of the `/` directory into the table `db1.t1`):

```
mkfifo /mysql/data/db1/ls.dat
chmod 666 /mysql/data/db1/ls.dat
find / -ls > /mysql/data/db1/ls.dat &
mysql -e "LOAD DATA INFILE 'ls.dat' INTO TABLE t1" db1
```

Note that you must run the command that generates the data to be loaded and the `mysql` commands either on separate terminals, or run the data generation process in the background (as shown in the preceding example). If you do not do this, the pipe will block until data is read by the `mysql` process.

When the `LOAD DATA INFILE` statement finishes, it returns an information string in the following format:

```
Records: 1 Deleted: 0 Skipped: 0 Warnings: 0
```

Warnings occur under the same circumstances as when values are inserted using the `INSERT` statement (see [Section 13.2.5, “INSERT Syntax”](#)), except that `LOAD DATA INFILE` also generates warnings when there are too few or too many fields in the input row.

You can use `SHOW WARNINGS` to get a list of the first `max_error_count` [504] warnings as information about what went wrong. See [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

If you are using the C API, you can get information about the statement by calling the `mysql_info()` function. See [Section 21.8.6.35, “mysql\\_info\(\)”](#).

## 13.2.7. LOAD XML Syntax

```
LOAD XML [LOW_PRIORITY | CONCURRENT] [LOCAL] INFILE 'file_name'
[REPLACE | IGNORE]
INTO TABLE [db_name.]tbl_name
[PARTITION (partition_name,...)]
[CHARACTER SET charset_name]
[ROWS IDENTIFIED BY '<tagname>']
[IGNORE number {LINES | ROWS}]
[(column_or_user_var,...)]
```

```
[SET col_name = expr,...]
```

The `LOAD XML` statement reads data from an XML file into a table. The `file_name` must be given as a literal string. The `tagname` in the optional `ROWS IDENTIFIED BY` clause must also be given as a literal string, and must be surrounded by angle brackets (< and >).

`LOAD XML` acts as the complement of running the `mysql` client in XML output mode (that is, starting the client with the `--xml [254]` option). To write data from a table to an XML file, use a command such as the following one from the system shell:

```
shell> mysql --xml -e 'SELECT * FROM mytable' > file.xml
```

To read the file back into a table, use `LOAD XML INFILE`. By default, the `<row>` element is considered to be the equivalent of a database table row; this can be changed using the `ROWS IDENTIFIED BY` clause.

This statement supports three different XML formats:

- Column names as attributes and column values as attribute values:

```
<row column1="value1" column2="value2" .../>
```

- Column names as tags and column values as the content of these tags:

```
<row>
  <column1>value1</column1>
  <column2>value2</column2>
</row>
```

- Column names are the `name` attributes of `<field>` tags, and values are the contents of these tags:

```
<row>
  <field name='column1'>value1</field>
  <field name='column2'>value2</field>
</row>
```

This is the format used by other MySQL tools, such as `mysqldump`.

All 3 formats can be used in the same XML file; the import routine automatically detects the format for each row and interprets it correctly. Tags are matched based on the tag or attribute name and the column name.

The following clauses work essentially the same way for `LOAD XML` as they do for `LOAD DATA`:

- `LOW_PRIORITY` or `CONCURRENT`
- `LOCAL`
- `REPLACE` or `IGNORE`
- `PARTITION`
- `CHARACTER SET`
- `(column_or_user_var,...)`
- `SET`

See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#), for more information about these clauses.

The `IGNORE number LINES` or `IGNORE number ROWS` clause causes the first `number` rows in the XML file to be skipped. It is analogous to the `LOAD DATA` statement's `IGNORE ... LINES` clause.

To illustrate how this statement is used, suppose that we have a table created as follows:

```
USE test;

CREATE TABLE person (
```

```

person_id INT NOT NULL PRIMARY KEY,
fname VARCHAR(40) NULL,
lname VARCHAR(40) NULL,
created TIMESTAMP
);

```

Suppose further that this table is initially empty.

Now suppose that we have a simple XML file `person.xml`, whose contents are as shown here:

```

<?xml version="1.0"?>
<list>
  <person person_id="1" fname="Pekka" lname="Nousiainen"/>
  <person person_id="2" fname="Jonas" lname="Oreland"/>
  <person person_id="3"><fname>Mikael</fname><lname>Ronström</lname></person>
  <person person_id="4"><fname>Lars</fname><lname>Thalman</lname></person>
  <person><field name="person_id">5</field><field name="fname">Tomas</field>
  <field name="lname">Ulin</field></person>
  <person><field name="person_id">6</field><field name="fname">Martin</field>
  <field name="lname">Sköld</field></person>
</list>

```

Each of the permissible XML formats discussed previously is represented in this example file.

To import the data in `person.xml` into the `person` table, you can use this statement:

```

mysql> LOAD XML LOCAL INFILE 'person.xml'
-> INTO TABLE person
-> ROWS IDENTIFIED BY '<person>';

Query OK, 6 rows affected (0.00 sec)
Records: 6 Deleted: 0 Skipped: 0 Warnings: 0

```

Here, we assume that `person.xml` is located in the MySQL data directory. If the file cannot be found, the following error results:

```

ERROR 2 (HY000): File '/person.xml' not found (Errcode: 2)

```

The `ROWS IDENTIFIED BY '<person>'` clause means that each `<person>` element in the XML file is considered equivalent to a row in the table into which the data is to be imported. In this case, this is the `person` table in the `test` database.

As can be seen by the response from the server, 6 rows were imported into the `test.person` table. This can be verified by a simple `SELECT` statement:

```

mysql> SELECT * FROM person;
+-----+-----+-----+-----+
| person_id | fname | lname | created |
+-----+-----+-----+-----+
| 1 | Pekka | Nousiainen | 2007-07-13 16:18:47 |
| 2 | Jonas | Oreland | 2007-07-13 16:18:47 |
| 3 | Mikael | Ronström | 2007-07-13 16:18:47 |
| 4 | Lars | Thalman | 2007-07-13 16:18:47 |
| 5 | Tomas | Ulin | 2007-07-13 16:18:47 |
| 6 | Martin | Sköld | 2007-07-13 16:18:47 |
+-----+-----+-----+-----+
6 rows in set (0.00 sec)

```

This shows, as stated earlier in this section, that any or all of the 3 permitted XML formats may appear in a single file and be read in using `LOAD XML`.

The inverse of the above operation—that is, dumping MySQL table data into an XML file—can be accomplished using the `mysql` client from the system shell, as shown here:



#### Note

The `--xml` [254] option causes the `mysql` client to use XML formatting for its output; the `-e` option causes the client to execute the SQL statement immediately following the option.

```

shell> mysql --xml -e "SELECT * FROM test.person" > person-dump.xml
shell> cat person-dump.xml
<?xml version="1.0"?>

<resultset statement="SELECT * FROM test.person" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <row>
    <field name="person_id">1</field>
    <field name="fname">Pekka</field>
    <field name="lname">Nousiainen</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>

  <row>
    <field name="person_id">2</field>
    <field name="fname">Jonas</field>
    <field name="lname">Oreland</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>

  <row>
    <field name="person_id">3</field>
    <field name="fname">Mikael</field>
    <field name="lname">Ronström</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>

  <row>
    <field name="person_id">4</field>
    <field name="fname">Lars</field>
    <field name="lname">Thalmann</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>

  <row>
    <field name="person_id">5</field>
    <field name="fname">Tomas</field>
    <field name="lname">Ulin</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>

  <row>
    <field name="person_id">6</field>
    <field name="fname">Martin</field>
    <field name="lname">Sköld</field>
    <field name="created">2007-07-13 16:18:47</field>
  </row>
</resultset>

```

You can verify that the dump is valid by creating a copy of the `person` and then importing the dump file into the new table, like this:

```

mysql> USE test;
mysql> CREATE TABLE person2 LIKE person;
Query OK, 0 rows affected (0.00 sec)

mysql> LOAD XML LOCAL INFILE 'person-dump.xml'
-> INTO TABLE person2;
Query OK, 6 rows affected (0.01 sec)
Records: 6 Deleted: 0 Skipped: 0 Warnings: 0

mysql> SELECT * FROM person2;
+-----+-----+-----+-----+
| person_id | fname | lname | created |
+-----+-----+-----+-----+
| 1 | Pekka | Nousiainen | 2007-07-13 16:18:47 |
| 2 | Jonas | Oreland | 2007-07-13 16:18:47 |
| 3 | Mikael | Ronström | 2007-07-13 16:18:47 |
| 4 | Lars | Thalmann | 2007-07-13 16:18:47 |
| 5 | Tomas | Ulin | 2007-07-13 16:18:47 |
| 6 | Martin | Sköld | 2007-07-13 16:18:47 |
+-----+-----+-----+-----+
6 rows in set (0.00 sec)

```

Using a `ROWS IDENTIFIED BY '<tagname>'` clause, it is possible to import data from the same XML file into database tables with different definitions. For this example, suppose that you have a file named `address.xml` which contains the following XML:

```
<?xml version="1.0"?>
<list>
  <person person_id="1">
    <fname>Robert</fname>
    <lname>Jones</lname>
    <address address_id="1" street="Mill Creek Road" zip="45365" city="Sidney"/>
    <address address_id="2" street="Main Street" zip="28681" city="Taylorsville"/>
  </person>

  <person person_id="2">
    <fname>Mary</fname>
    <lname>Smith</lname>
    <address address_id="3" street="River Road" zip="80239" city="Denver"/>
    <!-- <address address_id="4" street="North Street" zip="37920" city="Knoxville"/> -->
  </person>
</list>
```

You can again use the `test.person` table as defined previously in this section, after clearing all the existing records from the table and then showing its structure as shown here:

```
mysql> TRUNCATE person;
Query OK, 0 rows affected (0.04 sec)

mysql> SHOW CREATE TABLE person\G
***** 1. row *****
      Table: person
Create Table: CREATE TABLE `person` (
  `person_id` int(11) NOT NULL,
  `fname` varchar(40) DEFAULT NULL,
  `lname` varchar(40) DEFAULT NULL,
  `created` timestamp NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP,
  PRIMARY KEY (`person_id`)
) ENGINE=MyISAM DEFAULT CHARSET=latin1
1 row in set (0.00 sec)
```

Now create an `address` table in the `test` database using the following `CREATE TABLE` statement:

```
CREATE TABLE address (
  address_id INT NOT NULL PRIMARY KEY,
  person_id INT NULL,
  street VARCHAR(40) NULL,
  zip INT NULL,
  city VARCHAR(40) NULL,
  created TIMESTAMP
);
```

To import the data from the XML file into the `person` table, execute the following `LOAD XML` statement, which specifies that rows are to be specified by the `<person>` element, as shown here:

```
mysql> LOAD XML LOCAL INFILE 'address.xml'
-> INTO TABLE person
-> ROWS IDENTIFIED BY '<person>';
Query OK, 2 rows affected (0.00 sec)
Records: 2 Deleted: 0 Skipped: 0 Warnings: 0
```

You can verify that the records were imported using a `SELECT` statement:

```
mysql> SELECT * FROM person;
+-----+-----+-----+-----+
| person_id | fname | lname | created          |
+-----+-----+-----+-----+
|          1 | Robert | Jones | 2007-07-24 17:37:06 |
|          2 | Mary  | Smith | 2007-07-24 17:37:06 |
+-----+-----+-----+-----+
```



```
2 rows in set (0.00 sec)
```

Since the `<address>` elements in the XML file have no corresponding columns in the `person` table, they are skipped.

To import the data from the `<address>` elements into the `address` table, use the `LOAD XML` statement shown here:

```
mysql> LOAD XML LOCAL INFILE 'address.xml'
-> INTO TABLE address
-> ROWS IDENTIFIED BY '<address>';
Query OK, 3 rows affected (0.00 sec)
Records: 3 Deleted: 0 Skipped: 0 Warnings: 0
```

You can see that the data was imported using a `SELECT` statement such as this one:

```
mysql> SELECT * FROM address;
+-----+-----+-----+-----+-----+-----+
| address_id | person_id | street          | zip   | city          | created          |
+-----+-----+-----+-----+-----+-----+
|          1 |          1 | Mill Creek Road | 45365 | Sidney       | 2007-07-24 17:37:37 |
|          2 |          1 | Main Street     | 28681 | Taylorsville | 2007-07-24 17:37:37 |
|          3 |          2 | River Road      | 80239 | Denver       | 2007-07-24 17:37:37 |
+-----+-----+-----+-----+-----+-----+
3 rows in set (0.00 sec)
```

The data from the `<address>` element that is enclosed in XML comments is not imported. However, since there is a `person_id` column in the `address` table, the value of the `person_id` attribute from the parent `<person>` element for each `<address>` is imported into the `address` table.

**Security Considerations.** As with the `LOAD DATA` statement, the transfer of the XML file from the client host to the server host is initiated by the MySQL server. In theory, a patched server could be built that would tell the client program to transfer a file of the server's choosing rather than the file named by the client in the `LOAD XML` statement. Such a server could access any file on the client host to which the client user has read access.

In a Web environment, clients usually connect to MySQL from a Web server. A user that can run any command against the MySQL server can use `LOAD XML LOCAL` to read any files to which the Web server process has read access. In this environment, the client with respect to the MySQL server is actually the Web server, not the remote program being run by the user who connects to the Web server.

You can disable loading of XML files from clients by starting the server with `--local-infile=0` [496] or `--local-infile=OFF` [496]. This option can also be used when starting the `mysql` client to disable `LOAD XML` for the duration of the client session.

To prevent a client from loading XML files from the server, do not grant the `FILE` [754] privilege to the corresponding MySQL user account, or revoke this privilege if the client user account already has it.



### Important

Revoking the `FILE` [754] privilege (or not granting it in the first place) keeps the user only from executing the `LOAD XML INFILE` statement (as well as the `LOAD_FILE()` [1181] function; it does *not* prevent the user from executing `LOAD XML LOCAL INFILE`. To disallow this statement, you must start the server or the client with `--local-infile=OFF`.

In other words, the `FILE` [754] privilege affects only whether the client can read files on the server; it has no bearing on whether the client can read files on the local file system.

For partitioned tables using storage engines that employ table locks, such as `MyISAM`, `LOAD XML` cannot prune any partition locks. This does not apply to tables using storage engines which employ

row-level locking, such as [InnoDB](#). For more information, see [Section 17.6.4, “Partitioning and Locking”](#).

### 13.2.8. REPLACE Syntax

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name
  [PARTITION (partition_name,...)]
  [(col_name,...)]
  {VALUES | VALUE} ({expr | DEFAULT},...),(...),...
```

Or:

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name
  [PARTITION (partition_name,...)]
  SET col_name={expr | DEFAULT}, ...
```

Or:

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name
  [PARTITION (partition_name,...)]
  [(col_name,...)]
  SELECT ...
```

**REPLACE** works exactly like **INSERT**, except that if an old row in the table has the same value as a new row for a **PRIMARY KEY** or a **UNIQUE** index, the old row is deleted before the new row is inserted. See [Section 13.2.5, “INSERT Syntax”](#).

**REPLACE** is a MySQL extension to the SQL standard. It either inserts, or *deletes* and inserts. For another MySQL extension to standard SQL—that either inserts or *updates*—see [Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”](#).

Note that unless the table has a **PRIMARY KEY** or **UNIQUE** index, using a **REPLACE** statement makes no sense. It becomes equivalent to **INSERT**, because there is no index to be used to determine whether a new row duplicates another.

Values for all columns are taken from the values specified in the **REPLACE** statement. Any missing columns are set to their default values, just as happens for **INSERT**. You cannot refer to values from the current row and use them in the new row. If you use an assignment such as `SET col_name = col_name + 1`, the reference to the column name on the right hand side is treated as `DEFAULT(col_name)` [1275], so the assignment is equivalent to `SET col_name = DEFAULT(col_name) + 1`.

To use **REPLACE**, you must have both the **INSERT** [755] and **DELETE** [754] privileges for the table.

Beginning with MySQL 5.6.2, **REPLACE** supports explicit partition selection using the **PARTITION** option with a comma-separated list of names of partitions, subpartitions, or both. As with **INSERT**, if it is not possible to insert the new row into any of these partitions or subpartitions, the **REPLACE** statement fails with the error `Found a row not matching the given partition set`. See [Section 17.5, “Partition Selection”](#), for more information.

The **REPLACE** statement returns a count to indicate the number of rows affected. This is the sum of the rows deleted and inserted. If the count is 1 for a single-row **REPLACE**, a row was inserted and no rows were deleted. If the count is greater than 1, one or more old rows were deleted before the new row was inserted. It is possible for a single row to replace more than one old row if the table contains multiple unique indexes and the new row duplicates values for different old rows in different unique indexes.

The affected-rows count makes it easy to determine whether **REPLACE** only added a row or whether it also replaced any rows: Check whether the count is 1 (added) or greater (replaced).

If you are using the C API, the affected-rows count can be obtained using the `mysql_affected_rows()` function.

Currently, you cannot replace into a table and select from the same table in a subquery.

MySQL uses the following algorithm for `REPLACE` (and `LOAD DATA ... REPLACE`):

1. Try to insert the new row into the table
2. While the insertion fails because a duplicate-key error occurs for a primary key or unique index:
  - a. Delete from the table the conflicting row that has the duplicate key value
  - b. Try again to insert the new row into the table

It is possible that in the case of a duplicate-key error, a storage engine may perform the `REPLACE` as an update rather than a delete plus insert, but the semantics are the same. There are no user-visible effects other than a possible difference in how the storage engine increments `Handler_xxx` status variables.

Because the results of `REPLACE ... SELECT` statements depend on the ordering of rows from the `SELECT` and this order cannot always be guaranteed, it is possible when logging these statements for the master and the slave to diverge. For this reason, in MySQL 5.6.4 and later, `REPLACE ... SELECT` statements are flagged as unsafe for statement-based replication. With this change, such statements produce a warning in the log when using the `STATEMENT` binary logging mode, and are logged using the row-based format when using `MIXED` mode. See also [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

Prior to MySQL 5.6.6, a `REPLACE` that affected a partitioned table using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the table. This was true even for `REPLACE ... PARTITION` statements. (This did not and does not occur with storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, MySQL uses partition lock pruning, so that only partitions containing rows matching the `REPLACE` statement's `WHERE` clause are actually locked, as long as none of the table's partitioning columns are updated; otherwise the entire table is locked. For more information, see [Section 17.6.4, “Partitioning and Locking”](#).

### 13.2.9. SELECT Syntax

```
SELECT
  [ALL | DISTINCT | DISTINCTROW ]
  [HIGH_PRIORITY]
  [STRAIGHT_JOIN]
  [SQL_SMALL_RESULT] [SQL_BIG_RESULT] [SQL_BUFFER_RESULT]
  [SQL_CACHE | SQL_NO_CACHE] [SQL_CALC_FOUND_ROWS]
  select_expr [, select_expr ...]
  [FROM table_references
  [PARTITION partition_list]
  [WHERE where_condition]
  [GROUP BY {col_name | expr | position}
  [ASC | DESC], ... [WITH ROLLUP]]
  [HAVING where_condition]
  [ORDER BY {col_name | expr | position}
  [ASC | DESC], ...]
  [LIMIT {[offset,] row_count | row_count OFFSET offset}]
  [PROCEDURE procedure_name(argument_list)]
  [INTO OUTFILE 'file_name'
  [CHARACTER SET charset_name]
  export_options
  | INTO DUMPFILE 'file_name'
  | INTO var_name [, var_name]]
  [FOR UPDATE | LOCK IN SHARE MODE]]
```

`SELECT` is used to retrieve rows selected from one or more tables, and can include `UNION` statements and subqueries. See [Section 13.2.9.4, “UNION Syntax”](#), and [Section 13.2.10, “Subquery Syntax”](#).

The most commonly used clauses of `SELECT` statements are these:

- Each `select_expr` indicates a column that you want to retrieve. There must be at least one `select_expr`.
- `table_references` indicates the table or tables from which to retrieve rows. Its syntax is described in [Section 13.2.9.2, “JOIN Syntax”](#).
- Starting in MySQL 5.6.2, `SELECT` supports explicit partition selection using the `PARTITION` keyword with a list of partitions or subpartitions (or both) following the name of the table in a `table_reference` (see [Section 13.2.9.2, “JOIN Syntax”](#)). In this case, rows are selected only from the partitions listed, and any other partitions of the table are ignored. For more information and examples, see [Section 17.5, “Partition Selection”](#).

In MySQL 5.6.6 and later, `SELECT ... PARTITION` from tables using storage engines such as `MyISAM` that perform table-level locks (and thus partition locks) lock only the partitions or subpartitions named by the `PARTITION` option.

See [Section 17.6.4, “Partitioning and Locking”](#), for more information.

- The `WHERE` clause, if given, indicates the condition or conditions that rows must satisfy to be selected. `where_condition` is an expression that evaluates to true for each row to be selected. The statement selects all rows if there is no `WHERE` clause.

In the `WHERE` expression, you can use any of the functions and operators that MySQL supports, except for aggregate (summary) functions. See [Section 9.5, “Expression Syntax”](#), and [Chapter 12, Functions and Operators](#).

`SELECT` can also be used to retrieve rows computed without reference to any table.

For example:

```
mysql> SELECT 1 + 1;
-> 2
```

You are permitted to specify `DUAL` as a dummy table name in situations where no tables are referenced:

```
mysql> SELECT 1 + 1 FROM DUAL;
-> 2
```

`DUAL` is purely for the convenience of people who require that all `SELECT` statements should have `FROM` and possibly other clauses. MySQL may ignore the clauses. MySQL does not require `FROM DUAL` if no tables are referenced.

In general, clauses used must be given in exactly the order shown in the syntax description. For example, a `HAVING` clause must come after any `GROUP BY` clause and before any `ORDER BY` clause. The exception is that the `INTO` clause can appear either as shown in the syntax description or immediately following the `select_expr` list. For more information about `INTO`, see [Section 13.2.9.1, “SELECT ... INTO Syntax”](#).

The list of `select_expr` terms comprises the select list that indicates which columns to retrieve. Terms specify a column or expression or can use `*`-shorthand:

- A select list consisting only of a single unqualified `*` can be used as shorthand to select all columns from all tables:

```
SELECT * FROM t1 INNER JOIN t2 ...
```

- `tbl_name.*` can be used as a qualified shorthand to select all columns from the named table:

```
SELECT t1.*, t2.* FROM t1 INNER JOIN t2 ...
```

- Use of an unqualified `*` with other items in the select list may produce a parse error. To avoid this problem, use a qualified `tbl_name.*` reference

```
SELECT AVG(score), t1.* FROM t1 ...
```

The following list provides additional information about other `SELECT` clauses:

- A `select_expr` can be given an alias using `AS alias_name`. The alias is used as the expression's column name and can be used in `GROUP BY`, `ORDER BY`, or `HAVING` clauses. For example:

```
SELECT CONCAT(last_name, ', ', first_name) AS full_name
FROM mytable ORDER BY full_name;
```

The `AS` keyword is optional when aliasing a `select_expr` with an identifier. The preceding example could have been written like this:

```
SELECT CONCAT(last_name, ', ', first_name) full_name
FROM mytable ORDER BY full_name;
```

However, because the `AS` is optional, a subtle problem can occur if you forget the comma between two `select_expr` expressions: MySQL interprets the second as an alias name. For example, in the following statement, `columnb` is treated as an alias name:

```
SELECT columna columnb FROM mytable;
```

For this reason, it is good practice to be in the habit of using `AS` explicitly when specifying column aliases.

It is not permissible to refer to a column alias in a `WHERE` clause, because the column value might not yet be determined when the `WHERE` clause is executed. See [Section C.5.5.4, “Problems with Column Aliases”](#).

- The `FROM table_references` clause indicates the table or tables from which to retrieve rows. If you name more than one table, you are performing a join. For information on join syntax, see [Section 13.2.9.2, “JOIN Syntax”](#). For each table specified, you can optionally specify an alias.

```
tbl_name [[AS] alias] [index_hint]
```

The use of index hints provides the optimizer with information about how to choose indexes during query processing. For a description of the syntax for specifying these hints, see [Section 13.2.9.3, “Index Hint Syntax”](#).

You can use `SET max_seeks_for_key=value` as an alternative way to force MySQL to prefer key scans instead of table scans. See [Section 5.1.4, “Server System Variables”](#).

- You can refer to a table within the default database as `tbl_name`, or as `db_name.tbl_name` to specify a database explicitly. You can refer to a column as `col_name`, `tbl_name.col_name`, or `db_name.tbl_name.col_name`. You need not specify a `tbl_name` or `db_name.tbl_name` prefix for a column reference unless the reference would be ambiguous. See [Section 9.2.1, “Identifier Qualifiers”](#), for examples of ambiguity that require the more explicit column reference forms.
- A table reference can be aliased using `tbl_name AS alias_name` or `tbl_name alias_name`:

```
SELECT t1.name, t2.salary FROM employee AS t1, info AS t2
WHERE t1.name = t2.name;

SELECT t1.name, t2.salary FROM employee t1, info t2
WHERE t1.name = t2.name;
```

- Columns selected for output can be referred to in `ORDER BY` and `GROUP BY` clauses using column names, column aliases, or column positions. Column positions are integers and begin with 1:

```
SELECT college, region, seed FROM tournament
```

```
ORDER BY region, seed;

SELECT college, region AS r, seed AS s FROM tournament
ORDER BY r, s;

SELECT college, region, seed FROM tournament
ORDER BY 2, 3;
```

To sort in reverse order, add the `DESC` (descending) keyword to the name of the column in the `ORDER BY` clause that you are sorting by. The default is ascending order; this can be specified explicitly using the `ASC` keyword.

If `ORDER BY` occurs within a subquery and also is applied in the outer query, the outermost `ORDER BY` takes precedence. For example, results for the following statement are sorted in descending order, not ascending order:

```
(SELECT ... ORDER BY a) ORDER BY a DESC;
```

Use of column positions is deprecated because the syntax has been removed from the SQL standard.

- If you use `GROUP BY`, output rows are sorted according to the `GROUP BY` columns as if you had an `ORDER BY` for the same columns. To avoid the overhead of sorting that `GROUP BY` produces, add `ORDER BY NULL`:

```
SELECT a, COUNT(b) FROM test_table GROUP BY a ORDER BY NULL;
```

Relying on implicit `GROUP BY` sorting in MySQL 5.6 is deprecated. To achieve a specific sort order of grouped results, it is preferable to use an explicit `ORDER BY` clause. `GROUP BY` sorting is a MySQL extension that may change in a future release; for example, to make it possible for the optimizer to order groupings in whatever manner it deems most efficient and to avoid the sorting overhead.

- MySQL extends the `GROUP BY` clause so that you can also specify `ASC` and `DESC` after columns named in the clause:

```
SELECT a, COUNT(b) FROM test_table GROUP BY a DESC;
```

- MySQL extends the use of `GROUP BY` to permit selecting fields that are not mentioned in the `GROUP BY` clause. If you are not getting the results that you expect from your query, please read the description of `GROUP BY` found in [Section 12.17, “Functions and Modifiers for Use with GROUP BY Clauses”](#).
- `GROUP BY` permits a `WITH ROLLUP` modifier. See [Section 12.17.2, “GROUP BY Modifiers”](#).
- The `HAVING` clause is applied nearly last, just before items are sent to the client, with no optimization. (`LIMIT` is applied after `HAVING`.)

The SQL standard requires that `HAVING` must reference only columns in the `GROUP BY` clause or columns used in aggregate functions. However, MySQL supports an extension to this behavior, and permits `HAVING` to refer to columns in the `SELECT` list and columns in outer subqueries as well.

If the `HAVING` clause refers to a column that is ambiguous, a warning occurs. In the following statement, `col2` is ambiguous because it is used as both an alias and a column name:

```
SELECT COUNT(col1) AS col2 FROM t GROUP BY col2 HAVING col2 = 2;
```

Preference is given to standard SQL behavior, so if a `HAVING` column name is used both in `GROUP BY` and as an aliased column in the output column list, preference is given to the column in the `GROUP BY` column.

- Do not use `HAVING` for items that should be in the `WHERE` clause. For example, do not write the following:

```
SELECT col_name FROM tbl_name HAVING col_name > 0;
```

Write this instead:

```
SELECT col_name FROM tbl_name WHERE col_name > 0;
```

- The **HAVING** clause can refer to aggregate functions, which the **WHERE** clause cannot:

```
SELECT user, MAX(salary) FROM users
GROUP BY user HAVING MAX(salary) > 10;
```

(This did not work in some older versions of MySQL.)

- MySQL permits duplicate column names. That is, there can be more than one *select\_expr* with the same name. This is an extension to standard SQL. Because MySQL also permits **GROUP BY** and **HAVING** to refer to *select\_expr* values, this can result in an ambiguity:

```
SELECT 12 AS a, a FROM t GROUP BY a;
```

In that statement, both columns have the name *a*. To ensure that the correct column is used for grouping, use different names for each *select\_expr*.

- MySQL resolves unqualified column or alias references in **ORDER BY** clauses by searching in the *select\_expr* values, then in the columns of the tables in the **FROM** clause. For **GROUP BY** or **HAVING** clauses, it searches the **FROM** clause before searching in the *select\_expr* values. (For **GROUP BY** and **HAVING**, this differs from the pre-MySQL 5.0 behavior that used the same rules as for **ORDER BY**.)
- The **LIMIT** clause can be used to constrain the number of rows returned by the **SELECT** statement. **LIMIT** takes one or two numeric arguments, which must both be nonnegative integer constants, with these exceptions:
  - Within prepared statements, **LIMIT** parameters can be specified using `?` placeholder markers.
  - Within stored programs, **LIMIT** parameters can be specified using integer-valued routine parameters or local variables.

With two arguments, the first argument specifies the offset of the first row to return, and the second specifies the maximum number of rows to return. The offset of the initial row is 0 (not 1):

```
SELECT * FROM tbl LIMIT 5,10; # Retrieve rows 6-15
```

To retrieve all rows from a certain offset up to the end of the result set, you can use some large number for the second parameter. This statement retrieves all rows from the 96th row to the last:

```
SELECT * FROM tbl LIMIT 95,18446744073709551615;
```

With one argument, the value specifies the number of rows to return from the beginning of the result set:

```
SELECT * FROM tbl LIMIT 5; # Retrieve first 5 rows
```

In other words, **LIMIT row\_count** is equivalent to **LIMIT 0, row\_count**.

For prepared statements, you can use placeholders. The following statements will return one row from the *tbl* table:

```
SET @a=1;
PREPARE STMT FROM 'SELECT * FROM tbl LIMIT ?';
EXECUTE STMT USING @a;
```

The following statements will return the second to sixth row from the *tbl* table:

```
SET @skip=1; SET @numrows=5;
PREPARE STMT FROM 'SELECT * FROM tbl LIMIT ?, ?';
EXECUTE STMT USING @skip, @numrows;
```



For compatibility with PostgreSQL, MySQL also supports the `LIMIT row_count OFFSET offset` syntax.

If `LIMIT` occurs within a subquery and also is applied in the outer query, the outermost `LIMIT` takes precedence. For example, the following statement produces two rows, not one:

```
(SELECT ... LIMIT 1) LIMIT 2;
```

- A `PROCEDURE` clause names a procedure that should process the data in the result set. For an example, see [Section 8.4.2.4, “Using PROCEDURE ANALYSE”](#), which describes `ANALYSE`, a procedure that can be used to obtain suggestions for optimal column data types that may help reduce table sizes.
- The `SELECT ... INTO` form of `SELECT` enables the query result to be written to a file or stored in variables. For more information, see [Section 13.2.9.1, “SELECT ... INTO Syntax”](#).
- If you use `FOR UPDATE` with a storage engine that uses page or row locks, rows examined by the query are write-locked until the end of the current transaction. Using `LOCK IN SHARE MODE` sets a shared lock that permits other transactions to read the examined rows but not to update or delete them. See [Section 14.2.3.4, “Locking Reads \(SELECT ... FOR UPDATE and SELECT ... LOCK IN SHARE MODE\)”](#).

In addition, you cannot use `FOR UPDATE` as part of the `SELECT` in a statement such as `CREATE TABLE new_table SELECT ... FROM old_table ...` (If you attempt to do so, the statement is rejected with the error `Can't update table 'old_table' while 'new_table' is being created`.) This is a change in behavior from MySQL 5.5 and earlier, which permitted `CREATE TABLE ... SELECT` statements to make changes in tables other than the table being created.

Following the `SELECT` keyword, you can use a number of options that affect the operation of the statement. `HIGH_PRIORITY`, `STRAIGHT_JOIN`, and options beginning with `SQL_` are MySQL extensions to standard SQL.

- The `ALL` and `DISTINCT` options specify whether duplicate rows should be returned. `ALL` (the default) specifies that all matching rows should be returned, including duplicates. `DISTINCT` specifies removal of duplicate rows from the result set. It is an error to specify both options. `DISTINCTROW` is a synonym for `DISTINCT`.
- `HIGH_PRIORITY` gives the `SELECT` higher priority than a statement that updates a table. You should use this only for queries that are very fast and must be done at once. A `SELECT HIGH_PRIORITY` query that is issued while the table is locked for reading runs even if there is an update statement waiting for the table to be free. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).

`HIGH_PRIORITY` cannot be used with `SELECT` statements that are part of a `UNION`.

- `STRAIGHT_JOIN` forces the optimizer to join the tables in the order in which they are listed in the `FROM` clause. You can use this to speed up a query if the optimizer joins the tables in nonoptimal order. `STRAIGHT_JOIN` also can be used in the `table_references` list. See [Section 13.2.9.2, “JOIN Syntax”](#).

`STRAIGHT_JOIN` does not apply to any table that the optimizer treats as a `const` [898] or `system` [898] table. Such a table produces a single row, is read during the optimization phase of query execution, and references to its columns are replaced with the appropriate column values before query execution proceeds. These tables will appear first in the query plan displayed by `EXPLAIN`. See [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#). This exception may not apply to `const` [898] or `system` [898] tables that are used on the `NULL`-complemented side of an outer join (that is, the right-side table of a `LEFT JOIN` or the left-side table of a `RIGHT JOIN`).

- `SQL_BIG_RESULT` or `SQL_SMALL_RESULT` can be used with `GROUP BY` or `DISTINCT` to tell the optimizer that the result set has many rows or is small, respectively. For `SQL_BIG_RESULT`, MySQL

directly uses disk-based temporary tables if needed, and prefers sorting to using a temporary table with a key on the `GROUP BY` elements. For `SQL_SMALL_RESULT`, MySQL uses fast temporary tables to store the resulting table instead of using sorting. This should not normally be needed.

- `SQL_BUFFER_RESULT` forces the result to be put into a temporary table. This helps MySQL free the table locks early and helps in cases where it takes a long time to send the result set to the client. This option can be used only for top-level `SELECT` statements, not for subqueries or following `UNION`.
- `SQL_CALC_FOUND_ROWS` tells MySQL to calculate how many rows there would be in the result set, disregarding any `LIMIT` clause. The number of rows can then be retrieved with `SELECT FOUND_ROWS()`. See [Section 12.14, “Information Functions”](#).
- The `SQL_CACHE` and `SQL_NO_CACHE` options affect caching of query results in the query cache (see [Section 8.9.3, “The MySQL Query Cache”](#)). `SQL_CACHE` tells MySQL to store the result in the query cache if it is cacheable and the value of the `query_cache_type` [529] system variable is `2` or `DEMAND`. With `SQL_NO_CACHE`, the server does not use the query cache. It neither checks the query cache to see whether the result is already cached, nor does it cache the query result. (Due to a limitation in the parser, a space character must precede and follow the `SQL_NO_CACHE` keyword; a nonspace such as a newline causes the server to check the query cache to see whether the result is already cached.)

For views, `SQL_NO_CACHE` applies if it appears in any `SELECT` in the query. For a cacheable query, `SQL_CACHE` applies if it appears in the first `SELECT` of a view referred to by the query.

In MySQL 5.6, these two options are mutually exclusive and an error occurs if they are both specified. Also, these options are not permitted in subqueries (including subqueries in the `FROM` clause), and `SELECT` statements in unions other than the first `SELECT`.

Prior to MySQL 5.6.6, a `SELECT` from a partitioned table using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the table. This was true even for `SELECT ... PARTITION` queries. (This did not and does not occur with storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, MySQL uses partition lock pruning, so that only partitions containing rows matching the `SELECT` statement's `WHERE` clause are actually locked. For more information, see [Section 17.6.4, “Partitioning and Locking”](#).

### 13.2.9.1. `SELECT ... INTO` Syntax

The `SELECT ... INTO` form of `SELECT` enables a query result to be stored in variables or written to a file:

- `SELECT ... INTO var_list` selects column values and into variables.
- `SELECT ... INTO OUTFILE` writes the selected rows to a file. Column and line terminators can be specified to produce a specific output format.
- `SELECT ... INTO DUMPFILE` writes a single row to a file without any formatting.

The `SELECT` syntax description (see [Section 13.2.9, “SELECT Syntax”](#)) shows the `INTO` clause near the end of the statement. It is also possible to use `INTO` immediately following the `select_expr` list.

An `INTO` clause should not be used in a nested `SELECT` because such a `SELECT` must return its result to the outer context.

The `INTO` clause can name a list of one or more variables, which can be user-defined variables, stored procedure or function parameters, or stored program local variables. (Within a prepared `SELECT ... INTO OUTFILE` statement, only user-defined variables are permitted; see [Section 13.6.4.2, “Local Variable Scope and Resolution”](#).)

The selected values are assigned to the variables. The number of variables must match the number of columns. The query should return a single row. If the query returns no rows, a warning with error code 1329 occurs (`No data`), and the variable values remain unchanged. If the query returns multiple

rows, error 1172 occurs (`Result consisted of more than one row`). If it is possible that the statement may retrieve multiple rows, you can use `LIMIT 1` to limit the result set to a single row.

```
SELECT id, data INTO @x, @y FROM test.t1 LIMIT 1;
```

User variable names are not case sensitive. See [Section 9.4, “User-Defined Variables”](#).

The `SELECT ... INTO OUTFILE 'file_name'` form of `SELECT` writes the selected rows to a file. The file is created on the server host, so you must have the `FILE [754]` privilege to use this syntax. `file_name` cannot be an existing file, which among other things prevents files such as `/etc/passwd` and database tables from being destroyed. The `character_set_filesystem [468]` system variable controls the interpretation of the file name.

The `SELECT ... INTO OUTFILE` statement is intended primarily to let you very quickly dump a table to a text file on the server machine. If you want to create the resulting file on some other host than the server host, you normally cannot use `SELECT ... INTO OUTFILE` since there is no way to write a path to the file relative to the server host's file system.

However, if the MySQL client software is installed on the remote machine, you can instead use a client command such as `mysql -e "SELECT ..." > file_name` to generate the file on the client host.

It is also possible to create the resulting file on a different host other than the server host, if the location of the file on the remote host can be accessed using a network-mapped path on the server's file system. In this case, the presence of `mysql` (or some other MySQL client program) is not required on the target host.

`SELECT ... INTO OUTFILE` is the complement of `LOAD DATA INFILE`. Column values are written converted to the character set specified in the `CHARACTER SET` clause. If no such clause is present, values are dumped using the `binary` character set. In effect, there is no character set conversion. If a result set contains columns in several character sets, the output data file will as well and you may not be able to reload the file correctly.

The syntax for the `export_options` part of the statement consists of the same `FIELDS` and `LINES` clauses that are used with the `LOAD DATA INFILE` statement. See [Section 13.2.6, “LOAD DATA INFILE Syntax”](#), for information about the `FIELDS` and `LINES` clauses, including their default values and permissible values.

`FIELDS ESCAPED BY` controls how to write special characters. If the `FIELDS ESCAPED BY` character is not empty, it is used when necessary to avoid ambiguity as a prefix that precedes following characters on output:

- The `FIELDS ESCAPED BY` character
- The `FIELDS [OPTIONALLY] ENCLOSED BY` character
- The first character of the `FIELDS TERMINATED BY` and `LINES TERMINATED BY` values
- ASCII `NUL` (the zero-valued byte; what is actually written following the escape character is ASCII “0”, not a zero-valued byte)

The `FIELDS TERMINATED BY`, `ENCLOSED BY`, `ESCAPED BY`, or `LINES TERMINATED BY` characters *must* be escaped so that you can read the file back in reliably. ASCII `NUL` is escaped to make it easier to view with some pagers.

The resulting file does not have to conform to SQL syntax, so nothing else need be escaped.

If the `FIELDS ESCAPED BY` character is empty, no characters are escaped and `NULL` is output as `NULL`, not `\N`. It is probably not a good idea to specify an empty escape character, particularly if field values in your data contain any of the characters in the list just given.

Here is an example that produces a file in the comma-separated values (CSV) format used by many programs:

```
SELECT a,b,a+b INTO OUTFILE '/tmp/result.txt'
```

```

FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"'
LINES TERMINATED BY '\n'
FROM test_table;

```

If you use `INTO DUMPFILE` instead of `INTO OUTFILE`, MySQL writes only one row into the file, without any column or line termination and without performing any escape processing. This is useful if you want to store a `BLOB` value in a file.



#### Note

Any file created by `INTO OUTFILE` or `INTO DUMPFILE` is writable by all users on the server host. The reason for this is that the MySQL server cannot create a file that is owned by anyone other than the user under whose account it is running. (You should *never* run `mysqld` as `root` for this and other reasons.) The file thus must be world-writable so that you can manipulate its contents.

If the `secure_file_priv` [538] system variable is set to a nonempty directory name, the file to be written must be located in that directory.

In the context of `SELECT ... INTO` statements that occur as part of events executed by the Event Scheduler, diagnostics messages (not only errors, but also warnings) are written to the error log, and, on Windows, to the application event log. For additional information, see [Section 18.4.5, “Event Scheduler Status”](#).

### 13.2.9.2. JOIN Syntax

MySQL supports the following `JOIN` syntaxes for the `table_references` part of `SELECT` statements and multiple-table `DELETE` and `UPDATE` statements:

```

table_references:
    escaped_table_reference [, escaped_table_reference] ...

escaped_table_reference:
    table_reference
    | { OJ table_reference }

table_reference:
    table_factor
    | join_table

table_factor:
    tbl_name [PARTITION (partition_names)]
    [[AS] alias] [index_hint_list]
    | table_subquery [AS] alias
    | ( table_references )

join_table:
    table_reference [INNER | CROSS] JOIN table_factor [join_condition]
    | table_reference STRAIGHT_JOIN table_factor
    | table_reference STRAIGHT_JOIN table_factor ON conditional_expr
    | table_reference {LEFT|RIGHT} [OUTER] JOIN table_reference join_condition
    | table_reference NATURAL [{LEFT|RIGHT} [OUTER]] JOIN table_factor

join_condition:
    ON conditional_expr
    | USING (column_list)

index_hint_list:
    index_hint [, index_hint] ...

index_hint:
    USE {INDEX|KEY}
    [FOR {JOIN|ORDER BY|GROUP BY}] ([index_list])
    | IGNORE {INDEX|KEY}
    [FOR {JOIN|ORDER BY|GROUP BY}] (index_list)
    | FORCE {INDEX|KEY}
    [FOR {JOIN|ORDER BY|GROUP BY}] (index_list)

```

```
index_list:
    index_name [, index_name] ...
```

A table reference is also known as a join expression.

In MySQL 5.6.2 and later, a table reference (when it refers to a partitioned table) may contain a [PARTITION](#) option, including a comma-separated list of partitions, subpartitions, or both. This option follows the name of the table and precedes any alias declaration. The effect of this option is that rows are selected only from the listed partitions or subpartitions—in other words, any partitions or subpartitions not named in the list are ignored. For more information, see [Section 17.5, “Partition Selection”](#).

The syntax of *table\_factor* is extended in comparison with the SQL Standard. The latter accepts only *table\_reference*, not a list of them inside a pair of parentheses.

This is a conservative extension if we consider each comma in a list of *table\_reference* items as equivalent to an inner join. For example:

```
SELECT * FROM t1 LEFT JOIN (t2, t3, t4)
    ON (t2.a=t1.a AND t3.b=t1.b AND t4.c=t1.c)
```

is equivalent to:

```
SELECT * FROM t1 LEFT JOIN (t2 CROSS JOIN t3 CROSS JOIN t4)
    ON (t2.a=t1.a AND t3.b=t1.b AND t4.c=t1.c)
```

In MySQL, [JOIN](#), [CROSS JOIN](#), and [INNER JOIN](#) are syntactic equivalents (they can replace each other). In standard SQL, they are not equivalent. [INNER JOIN](#) is used with an [ON](#) clause, [CROSS JOIN](#) is used otherwise.

In general, parentheses can be ignored in join expressions containing only inner join operations. MySQL also supports nested joins (see [Section 8.13.9, “Nested Join Optimization”](#)).

Index hints can be specified to affect how the MySQL optimizer makes use of indexes. For more information, see [Section 13.2.9.3, “Index Hint Syntax”](#).

The following list describes general factors to take into account when writing joins.

- A table reference can be aliased using *tbl\_name AS alias\_name* or *tbl\_name alias\_name*:

```
SELECT t1.name, t2.salary
    FROM employee AS t1 INNER JOIN info AS t2 ON t1.name = t2.name;

SELECT t1.name, t2.salary
    FROM employee t1 INNER JOIN info t2 ON t1.name = t2.name;
```

- A *table\_subquery* is also known as a subquery in the [FROM](#) clause. Such subqueries *must* include an alias to give the subquery result a table name. A trivial example follows; see also [Section 13.2.10.8, “Subqueries in the FROM Clause”](#).

```
SELECT * FROM (SELECT 1, 2, 3) AS t1;
```

- [INNER JOIN](#) and [,](#) (comma) are semantically equivalent in the absence of a join condition: both produce a Cartesian product between the specified tables (that is, each and every row in the first table is joined to each and every row in the second table).

However, the precedence of the comma operator is less than of [INNER JOIN](#), [CROSS JOIN](#), [LEFT JOIN](#), and so on. If you mix comma joins with the other join types when there is a join condition, an error of the form `Unknown column 'col_name' in 'on clause'` may occur. Information about dealing with this problem is given later in this section.

- The *conditional\_expr* used with [ON](#) is any conditional expression of the form that can be used in a [WHERE](#) clause. Generally, you should use the [ON](#) clause for conditions that specify how to join tables, and the [WHERE](#) clause to restrict which rows you want in the result set.

- If there is no matching row for the right table in the `ON` or `USING` part in a `LEFT JOIN`, a row with all columns set to `NULL` is used for the right table. You can use this fact to find rows in a table that have no counterpart in another table:

```
SELECT left_tbl.*
FROM left_tbl LEFT JOIN right_tbl ON left_tbl.id = right_tbl.id
WHERE right_tbl.id IS NULL;
```

This example finds all rows in `left_tbl` with an `id` value that is not present in `right_tbl` (that is, all rows in `left_tbl` with no corresponding row in `right_tbl`). This assumes that `right_tbl.id` is declared `NOT NULL`. See [Section 8.13.7, “LEFT JOIN and RIGHT JOIN Optimization”](#).

- The `USING(column_list)` clause names a list of columns that must exist in both tables. If tables `a` and `b` both contain columns `c1`, `c2`, and `c3`, the following join compares corresponding columns from the two tables:

```
a LEFT JOIN b USING (c1,c2,c3)
```

- The `NATURAL [LEFT] JOIN` of two tables is defined to be semantically equivalent to an `INNER JOIN` or a `LEFT JOIN` with a `USING` clause that names all columns that exist in both tables.
- `RIGHT JOIN` works analogously to `LEFT JOIN`. To keep code portable across databases, it is recommended that you use `LEFT JOIN` instead of `RIGHT JOIN`.
- The `{ OJ ... }` syntax shown in the join syntax description exists only for compatibility with ODBC. The curly braces in the syntax should be written literally; they are not metasyntax as used elsewhere in syntax descriptions.

```
SELECT left_tbl.*
FROM { OJ left_tbl LEFT OUTER JOIN right_tbl ON left_tbl.id = right_tbl.id }
WHERE right_tbl.id IS NULL;
```

You can use other types of joins within `{ OJ ... }`, such as `INNER JOIN` or `RIGHT OUTER JOIN`. This helps with compatibility with some third-party applications, but is not official ODBC syntax.

- `STRAIGHT_JOIN` is similar to `JOIN`, except that the left table is always read before the right table. This can be used for those (few) cases for which the join optimizer puts the tables in the wrong order.

Some join examples:

```
SELECT * FROM table1, table2;

SELECT * FROM table1 INNER JOIN table2 ON table1.id=table2.id;

SELECT * FROM table1 LEFT JOIN table2 ON table1.id=table2.id;

SELECT * FROM table1 LEFT JOIN table2 USING (id);

SELECT * FROM table1 LEFT JOIN table2 ON table1.id=table2.id
LEFT JOIN table3 ON table2.id=table3.id;
```

### Join Processing Changes in MySQL 5.0.12



#### Note

Natural joins and joins with `USING`, including outer join variants, are processed according to the SQL:2003 standard. The goal was to align the syntax and semantics of MySQL with respect to `NATURAL JOIN` and `JOIN ... USING` according to SQL:2003. However, these changes in join processing can result in different output columns for some joins. Also, some queries that appeared to work correctly in older versions (prior to 5.0.12) must be rewritten to comply with the standard.

These changes have five main aspects:

- The way that MySQL determines the result columns of `NATURAL` or `USING` join operations (and thus the result of the entire `FROM` clause).
- Expansion of `SELECT *` and `SELECT tbl_name.*` into a list of selected columns.
- Resolution of column names in `NATURAL` or `USING` joins.
- Transformation of `NATURAL` or `USING` joins into `JOIN ... ON`.
- Resolution of column names in the `ON` condition of a `JOIN ... ON`.

The following list provides more detail about several effects of current join processing versus join processing in older versions. The term “previously” means “prior to MySQL 5.0.12.”

- The columns of a `NATURAL` join or a `USING` join may be different from previously. Specifically, redundant output columns no longer appear, and the order of columns for `SELECT *` expansion may be different from before.

Consider this set of statements:

```
CREATE TABLE t1 (i INT, j INT);
CREATE TABLE t2 (k INT, j INT);
INSERT INTO t1 VALUES(1,1);
INSERT INTO t2 VALUES(1,1);
SELECT * FROM t1 NATURAL JOIN t2;
SELECT * FROM t1 JOIN t2 USING (j);
```

Previously, the statements produced this output:

```
+-----+-----+-----+
| i   | j   | k   | j   |
+-----+-----+-----+
|  1  |  1  |  1  |  1  |
+-----+-----+-----+
| i   | j   | k   | j   |
+-----+-----+-----+
|  1  |  1  |  1  |  1  |
+-----+-----+-----+
```

In the first `SELECT` statement, column `j` appears in both tables and thus becomes a join column, so, according to standard SQL, it should appear only once in the output, not twice. Similarly, in the second `SELECT` statement, column `j` is named in the `USING` clause and should appear only once in the output, not twice. But in both cases, the redundant column is not eliminated. Also, the order of the columns is not correct according to standard SQL.

Now the statements produce this output:

```
+-----+-----+-----+
| j   | i   | k   |
+-----+-----+-----+
|  1  |  1  |  1  |
+-----+-----+-----+
| j   | i   | k   |
+-----+-----+-----+
|  1  |  1  |  1  |
+-----+-----+-----+
```

The redundant column is eliminated and the column order is correct according to standard SQL:

- First, coalesced common columns of the two joined tables, in the order in which they occur in the first table
- Second, columns unique to the first table, in order in which they occur in that table
- Third, columns unique to the second table, in order in which they occur in that table



The single result column that replaces two common columns is defined using the coalesce operation. That is, for two `t1.a` and `t2.a` the resulting single join column `a` is defined as `a = COALESCE(t1.a, t2.a)`, where:

```
COALESCE(x, y) = (CASE WHEN V1 IS NOT NULL THEN V1 ELSE V2 END)
```

If the join operation is any other join, the result columns of the join consists of the concatenation of all columns of the joined tables. This is the same as previously.

A consequence of the definition of coalesced columns is that, for outer joins, the coalesced column contains the value of the non-NULL column if one of the two columns is always NULL. If neither or both columns are NULL, both common columns have the same value, so it doesn't matter which one is chosen as the value of the coalesced column. A simple way to interpret this is to consider that a coalesced column of an outer join is represented by the common column of the inner table of a JOIN. Suppose that the tables `t1(a,b)` and `t2(a,c)` have the following contents:

```
t1    t2
----  ----
1 x   2 z
2 y   3 w
```

Then:

```
mysql> SELECT * FROM t1 NATURAL LEFT JOIN t2;
+-----+-----+-----+
| a     | b     | c     |
+-----+-----+-----+
| 1     | x     | NULL  |
| 2     | y     | z     |
+-----+-----+-----+
```

Here column `a` contains the values of `t1.a`.

```
mysql> SELECT * FROM t1 NATURAL RIGHT JOIN t2;
+-----+-----+-----+
| a     | c     | b     |
+-----+-----+-----+
| 2     | z     | y     |
| 3     | w     | NULL  |
+-----+-----+-----+
```

Here column `a` contains the values of `t2.a`.

Compare these results to the otherwise equivalent queries with `JOIN ... ON`:

```
mysql> SELECT * FROM t1 LEFT JOIN t2 ON (t1.a = t2.a);
+-----+-----+-----+-----+
| a     | b     | a     | c     |
+-----+-----+-----+-----+
| 1     | x     | NULL  | NULL  |
| 2     | y     | 2     | z     |
+-----+-----+-----+-----+
```

```
mysql> SELECT * FROM t1 RIGHT JOIN t2 ON (t1.a = t2.a);
+-----+-----+-----+-----+
| a     | b     | a     | c     |
+-----+-----+-----+-----+
| 2     | y     | 2     | z     |
| NULL  | NULL  | 3     | w     |
+-----+-----+-----+-----+
```

- Previously, a `USING` clause could be rewritten as an `ON` clause that compares corresponding columns. For example, the following two clauses were semantically identical:

```
a LEFT JOIN b USING (c1,c2,c3)
a LEFT JOIN b ON a.c1=b.c1 AND a.c2=b.c2 AND a.c3=b.c3
```

Now the two clauses no longer are quite the same:

- With respect to determining which rows satisfy the join condition, both joins remain semantically identical.
- With respect to determining which columns to display for `SELECT *` expansion, the two joins are not semantically identical. The `USING` join selects the coalesced value of corresponding columns, whereas the `ON` join selects all columns from all tables. For the preceding `USING` join, `SELECT *` selects these values:

```
COALESCE(a.c1,b.c1), COALESCE(a.c2,b.c2), COALESCE(a.c3,b.c3)
```

For the `ON` join, `SELECT *` selects these values:

```
a.c1, a.c2, a.c3, b.c1, b.c2, b.c3
```

With an inner join, `COALESCE(a.c1,b.c1)` [1169] is the same as either `a.c1` or `b.c1` because both columns will have the same value. With an outer join (such as `LEFT JOIN`), one of the two columns can be `NULL`. That column will be omitted from the result.

- The evaluation of multi-way natural joins differs in a very important way that affects the result of `NATURAL` or `USING` joins and that can require query rewriting. Suppose that you have three tables `t1(a,b)`, `t2(c,b)`, and `t3(a,c)` that each have one row: `t1(1,2)`, `t2(10,2)`, and `t3(7,10)`. Suppose also that you have this `NATURAL JOIN` on the three tables:

```
SELECT ... FROM t1 NATURAL JOIN t2 NATURAL JOIN t3;
```

Previously, the left operand of the second join was considered to be `t2`, whereas it should be the nested join (`t1 NATURAL JOIN t2`). As a result, the columns of `t3` are checked for common columns only in `t2`, and, if `t3` has common columns with `t1`, these columns are not used as equi-join columns. Thus, previously, the preceding query was transformed to the following equi-join:

```
SELECT ... FROM t1, t2, t3
WHERE t1.b = t2.b AND t2.c = t3.c;
```

That join is missing one more equi-join predicate (`t1.a = t3.a`). As a result, it produces one row, not the empty result that it should. The correct equivalent query is this:

```
SELECT ... FROM t1, t2, t3
WHERE t1.b = t2.b AND t2.c = t3.c AND t1.a = t3.a;
```

If you require the same query result in current versions of MySQL as in older versions, rewrite the natural join as the first equi-join.

- Previously, the comma operator (`,`) and `JOIN` both had the same precedence, so the join expression `t1, t2 JOIN t3` was interpreted as `((t1, t2) JOIN t3)`. Now `JOIN` has higher precedence, so the expression is interpreted as `(t1, (t2 JOIN t3))`. This change affects statements that use an `ON` clause, because that clause can refer only to columns in the operands of the join, and the change in precedence changes interpretation of what those operands are.

Example:

```
CREATE TABLE t1 (i1 INT, j1 INT);
CREATE TABLE t2 (i2 INT, j2 INT);
CREATE TABLE t3 (i3 INT, j3 INT);
INSERT INTO t1 VALUES(1,1);
INSERT INTO t2 VALUES(1,1);
INSERT INTO t3 VALUES(1,1);
SELECT * FROM t1, t2 JOIN t3 ON (t1.i1 = t3.i3);
```

Previously, the `SELECT` was legal due to the implicit grouping of `t1,t2` as `(t1,t2)`. Now the `JOIN` takes precedence, so the operands for the `ON` clause are `t2` and `t3`. Because `t1.i1` is not a column in either of the operands, the result is an `Unknown column 't1.i1' in 'on clause'`

error. To allow the join to be processed, group the first two tables explicitly with parentheses so that the operands for the `ON` clause are `(t1,t2)` and `t3`:

```
SELECT * FROM (t1, t2) JOIN t3 ON (t1.i1 = t3.i3);
```

Alternatively, avoid the use of the comma operator and use `JOIN` instead:

```
SELECT * FROM t1 JOIN t2 JOIN t3 ON (t1.i1 = t3.i3);
```

This change also applies to statements that mix the comma operator with `INNER JOIN`, `CROSS JOIN`, `LEFT JOIN`, and `RIGHT JOIN`, all of which now have higher precedence than the comma operator.

- Previously, the `ON` clause could refer to columns in tables named to its right. Now an `ON` clause can refer only to its operands.

Example:

```
CREATE TABLE t1 (i1 INT);
CREATE TABLE t2 (i2 INT);
CREATE TABLE t3 (i3 INT);
SELECT * FROM t1 JOIN t2 ON (i1 = i3) JOIN t3;
```

Previously, the `SELECT` statement was legal. Now the statement fails with an `Unknown column 'i3' in 'on clause'` error because `i3` is a column in `t3`, which is not an operand of the `ON` clause. The statement should be rewritten as follows:

```
SELECT * FROM t1 JOIN t2 JOIN t3 ON (i1 = i3);
```

- Resolution of column names in `NATURAL` or `USING` joins is different than previously. For column names that are outside the `FROM` clause, MySQL now handles a superset of the queries compared to previously. That is, in cases when MySQL formerly issued an error that some column is ambiguous, the query now is handled correctly. This is due to the fact that MySQL now treats the common columns of `NATURAL` or `USING` joins as a single column, so when a query refers to such columns, the query compiler does not consider them as ambiguous.

Example:

```
SELECT * FROM t1 NATURAL JOIN t2 WHERE b > 1;
```

Previously, this query would produce an error `ERROR 1052 (23000): Column 'b' in where clause is ambiguous`. Now the query produces the correct result:

```
+-----+-----+-----+
| b    | c    | y    |
+-----+-----+-----+
|    4 |    2 |    3 |
+-----+-----+-----+
```

One extension of MySQL compared to the SQL:2003 standard is that MySQL enables you to qualify the common (coalesced) columns of `NATURAL` or `USING` joins (just as previously), while the standard disallows that.

### 13.2.9.3. Index Hint Syntax

You can provide hints to give the optimizer information about how to choose indexes during query processing. [Section 13.2.9.2, "JOIN Syntax"](#), describes the general syntax for specifying tables in a `SELECT` statement. The syntax for an individual table, including that for index hints, looks like this:

```
tbl_name [[AS] alias] [index_hint_list]

index_hint_list:
    index_hint [, index_hint] ...

index_hint:
```

```

USE {INDEX|KEY}
  [FOR {JOIN|ORDER BY|GROUP BY}] ([index_list])
| IGNORE {INDEX|KEY}
  [FOR {JOIN|ORDER BY|GROUP BY}] (index_list)
| FORCE {INDEX|KEY}
  [FOR {JOIN|ORDER BY|GROUP BY}] (index_list)

```

```

index_list:
  index_name [, index_name] ...

```

By specifying `USE INDEX (index_list)`, you can tell MySQL to use only one of the named indexes to find rows in the table. The alternative syntax `IGNORE INDEX (index_list)` can be used to tell MySQL to not use some particular index or indexes. These hints are useful if `EXPLAIN` shows that MySQL is using the wrong index from the list of possible indexes.

You can also use `FORCE INDEX`, which acts like `USE INDEX (index_list)` but with the addition that a table scan is assumed to be *very* expensive. In other words, a table scan is used only if there is no way to use one of the given indexes to find rows in the table.

Each hint requires the names of *indexes*, not the names of columns. The name of a `PRIMARY KEY` is `PRIMARY`. To see the index names for a table, use `SHOW INDEX`.

An *index\_name* value need not be a full index name. It can be an unambiguous prefix of an index name. If a prefix is ambiguous, an error occurs.

Examples:

```

SELECT * FROM table1 USE INDEX (col1_index,col2_index)
  WHERE col1=1 AND col2=2 AND col3=3;

SELECT * FROM table1 IGNORE INDEX (col3_index)
  WHERE col1=1 AND col2=2 AND col3=3;

```

The syntax for index hints has the following characteristics:

- It is syntactically valid to specify an empty *index\_list* for `USE INDEX`, which means “use no indexes.” Specifying an empty *index\_list* for `FORCE INDEX` or `IGNORE INDEX` is a syntax error.
- You can specify the scope of a index hint by adding a `FOR` clause to the hint. This provides more fine-grained control over the optimizer's selection of an execution plan for various phases of query processing. To affect only the indexes used when MySQL decides how to find rows in the table and how to process joins, use `FOR JOIN`. To influence index usage for sorting or grouping rows, use `FOR ORDER BY` or `FOR GROUP BY`. (However, if there is a covering index for the table and it is used to access the table, the optimizer will ignore `IGNORE INDEX FOR {ORDER BY|GROUP BY}` hints that disable that index.)
- You can specify multiple index hints:

```

SELECT * FROM t1 USE INDEX (i1) IGNORE INDEX FOR ORDER BY (i2) ORDER BY a;

```

It is not a error to name the same index in several hints (even within the same hint):

```

SELECT * FROM t1 USE INDEX (i1) USE INDEX (i1,i1);

```

However, it is an error to mix `USE INDEX` and `FORCE INDEX` for the same table:

```

SELECT * FROM t1 USE INDEX FOR JOIN (i1) FORCE INDEX FOR JOIN (i2);

```

if you specify no `FOR` clause for an index hint, the hint by default applies to all parts of the statement. For example, this hint:

```

IGNORE INDEX (i1)

```

is equivalent to this combination of hints:

```

IGNORE INDEX FOR JOIN (i1)

```

```
IGNORE INDEX FOR ORDER BY (i1)
IGNORE INDEX FOR GROUP BY (i1)
```

To cause the server to use the older behavior for hint scope when no `FOR` clause is present (so that hints apply only to row retrieval), enable the `old [517]` system variable at server startup. Take care about enabling this variable in a replication setup. With statement-based binary logging, having different modes for the master and slaves might lead to replication errors.

When index hints are processed, they are collected in a single list by type (`USE`, `FORCE`, `IGNORE`) and by scope (`FOR JOIN`, `FOR ORDER BY`, `FOR GROUP BY`). For example:

```
SELECT * FROM t1
  USE INDEX () IGNORE INDEX (i2) USE INDEX (i1) USE INDEX (i2);
```

is equivalent to:

```
SELECT * FROM t1
  USE INDEX (i1,i2) IGNORE INDEX (i2);
```

The index hints then are applied for each scope in the following order:

1. `{USE|FORCE} INDEX` is applied if present. (If not, the optimizer-determined set of indexes is used.)
2. `IGNORE INDEX` is applied over the result of the previous step. For example, the following two queries are equivalent:

```
SELECT * FROM t1 USE INDEX (i1) IGNORE INDEX (i2) USE INDEX (i2);
SELECT * FROM t1 USE INDEX (i1);
```

For `FULLTEXT` searches, index hints work as follows:

- For natural language mode searches, index hints are silently ignored. For example, `IGNORE INDEX(i)` is ignored with no warning and the index is still used.

For boolean mode searches, index hints with `FOR ORDER BY` or `FOR GROUP BY` are silently ignored. Index hints with `FOR JOIN` or no `FOR` modifier are honored. In contrast to how hints apply for non-`FULLTEXT` searches, the hint is used for all phases of query execution (finding rows and retrieval, grouping, and ordering). This is true even if the hint is given for a non-`FULLTEXT` index.

For example, the following two queries are equivalent:

```
SELECT * FROM t
  USE INDEX (index1)
  IGNORE INDEX (index1) FOR ORDER BY
  IGNORE INDEX (index1) FOR GROUP BY
  WHERE ... IN BOOLEAN MODE ... ;

SELECT * FROM t
  USE INDEX (index1)
  WHERE ... IN BOOLEAN MODE ... ;
```

#### 13.2.9.4. UNION Syntax

```
SELECT ...
UNION [ALL | DISTINCT] SELECT ...
[UNION [ALL | DISTINCT] SELECT ...]
```

`UNION` is used to combine the result from multiple `SELECT` statements into a single result set.

The column names from the first `SELECT` statement are used as the column names for the results returned. Selected columns listed in corresponding positions of each `SELECT` statement should have the same data type. (For example, the first column selected by the first statement should have the same type as the first column selected by the other statements.)

If the data types of corresponding `SELECT` columns do not match, the types and lengths of the columns in the `UNION` result take into account the values retrieved by all of the `SELECT` statements. For example, consider the following:

```
mysql> SELECT REPEAT('a',1) UNION SELECT REPEAT('b',10);
+-----+
| REPEAT('a',1) |
+-----+
| a              |
| bbbbbbbbbbb   |
+-----+
```

The `SELECT` statements are normal select statements, but with the following restrictions:

- Only the last `SELECT` statement can use `INTO OUTFILE`. (However, the entire `UNION` result is written to the file.)
- `HIGH_PRIORITY` cannot be used with `SELECT` statements that are part of a `UNION`. If you specify it for the first `SELECT`, it has no effect. If you specify it for any subsequent `SELECT` statements, a syntax error results.

The default behavior for `UNION` is that duplicate rows are removed from the result. The optional `DISTINCT` keyword has no effect other than the default because it also specifies duplicate-row removal. With the optional `ALL` keyword, duplicate-row removal does not occur and the result includes all matching rows from all the `SELECT` statements.

You can mix `UNION ALL` and `UNION DISTINCT` in the same query. Mixed `UNION` types are treated such that a `DISTINCT` union overrides any `ALL` union to its left. A `DISTINCT` union can be produced explicitly by using `UNION DISTINCT` or implicitly by using `UNION` with no following `DISTINCT` or `ALL` keyword.

To apply `ORDER BY` or `LIMIT` to an individual `SELECT`, place the clause inside the parentheses that enclose the `SELECT`:

```
(SELECT a FROM t1 WHERE a=10 AND B=1 ORDER BY a LIMIT 10)
UNION
(SELECT a FROM t2 WHERE a=11 AND B=2 ORDER BY a LIMIT 10);
```

However, use of `ORDER BY` for individual `SELECT` statements implies nothing about the order in which the rows appear in the final result because `UNION` by default produces an unordered set of rows. Therefore, the use of `ORDER BY` in this context is typically in conjunction with `LIMIT`, so that it is used to determine the subset of the selected rows to retrieve for the `SELECT`, even though it does not necessarily affect the order of those rows in the final `UNION` result. If `ORDER BY` appears without `LIMIT` in a `SELECT`, it is optimized away because it will have no effect anyway.

To use an `ORDER BY` or `LIMIT` clause to sort or limit the entire `UNION` result, parenthesize the individual `SELECT` statements and place the `ORDER BY` or `LIMIT` after the last one. The following example uses both clauses:

```
(SELECT a FROM t1 WHERE a=10 AND B=1)
UNION
(SELECT a FROM t2 WHERE a=11 AND B=2)
ORDER BY a LIMIT 10;
```

A statement without parentheses is equivalent to one parenthesized as just shown.

This kind of `ORDER BY` cannot use column references that include a table name (that is, names in `tbl_name.col_name` format). Instead, provide a column alias in the first `SELECT` statement and refer to the alias in the `ORDER BY`. (Alternatively, refer to the column in the `ORDER BY` using its column position. However, use of column positions is deprecated.)

Also, if a column to be sorted is aliased, the `ORDER BY` clause *must* refer to the alias, not the column name. The first of the following statements will work, but the second will fail with an `Unknown column 'a' in 'order clause'` error:

```
(SELECT a AS b FROM t) UNION (SELECT ...) ORDER BY b;
(SELECT a AS b FROM t) UNION (SELECT ...) ORDER BY a;
```

To cause rows in a **UNION** result to consist of the sets of rows retrieved by each **SELECT** one after the other, select an additional column in each **SELECT** to use as a sort column and add an **ORDER BY** following the last **SELECT**:

```
(SELECT 1 AS sort_col, col1a, col1b, ... FROM t1)
UNION
(SELECT 2, col2a, col2b, ... FROM t2) ORDER BY sort_col;
```

To additionally maintain sort order within individual **SELECT** results, add a secondary column to the **ORDER BY** clause:

```
(SELECT 1 AS sort_col, col1a, col1b, ... FROM t1)
UNION
(SELECT 2, col2a, col2b, ... FROM t2) ORDER BY sort_col, col1a;
```

Use of an additional column also enables you to determine which **SELECT** each row comes from. Extra columns can provide other identifying information as well, such as a string that indicates a table name.

### 13.2.10. Subquery Syntax

A subquery is a **SELECT** statement within another statement.

Starting with MySQL 4.1, all subquery forms and operations that the SQL standard requires are supported, as well as a few features that are MySQL-specific.

Here is an example of a subquery:

```
SELECT * FROM t1 WHERE column1 = (SELECT column1 FROM t2);
```

In this example, **SELECT \* FROM t1 ...** is the *outer query* (or *outer statement*), and **(SELECT column1 FROM t2)** is the *subquery*. We say that the subquery is *nested* within the outer query, and in fact it is possible to nest subqueries within other subqueries, to a considerable depth. A subquery must always appear within parentheses.

The main advantages of subqueries are:

- They allow queries that are *structured* so that it is possible to isolate each part of a statement.
- They provide alternative ways to perform operations that would otherwise require complex joins and unions.
- Many people find subqueries more readable than complex joins or unions. Indeed, it was the innovation of subqueries that gave people the original idea of calling the early SQL “Structured Query Language.”

Here is an example statement that shows the major points about subquery syntax as specified by the SQL standard and supported in MySQL:

```
DELETE FROM t1
WHERE s11 > ANY
  (SELECT COUNT(*) /* no hint */ FROM t2
  WHERE NOT EXISTS
    (SELECT * FROM t3
     WHERE ROW(5*t2.s1,77)=
      (SELECT 50,11*s1 FROM t4 UNION SELECT 50,77 FROM
       (SELECT * FROM t5) AS t5)));
```

A subquery can return a scalar (a single value), a single row, a single column, or a table (one or more rows of one or more columns). These are called scalar, column, row, and table subqueries. Subqueries that return a particular kind of result often can be used only in certain contexts, as described in the following sections.



There are few restrictions on the type of statements in which subqueries can be used. A subquery can contain many of the keywords or clauses that an ordinary `SELECT` can contain: `DISTINCT`, `GROUP BY`, `ORDER BY`, `LIMIT`, joins, index hints, `UNION` constructs, comments, functions, and so on.

A subquery's outer statement can be any one of: `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `SET`, or `DO`.

In MySQL, you cannot modify a table and select from the same table in a subquery. This applies to statements such as `DELETE`, `INSERT`, `REPLACE`, `UPDATE`, and (because subqueries can be used in the `SET` clause) `LOAD DATA INFILE`.

For information about how the optimizer handles subqueries, see [Section 8.13.16, "Subquery Optimization"](#). For a discussion of restrictions on subquery use, including performance issues for certain forms of subquery syntax, see [Section E.4, "Restrictions on Subqueries"](#).

### 13.2.10.1. The Subquery as Scalar Operand

In its simplest form, a subquery is a scalar subquery that returns a single value. A scalar subquery is a simple operand, and you can use it almost anywhere a single column value or literal is legal, and you can expect it to have those characteristics that all operands have: a data type, a length, an indication that it can be `NULL`, and so on. For example:

```
CREATE TABLE t1 (s1 INT, s2 CHAR(5) NOT NULL);
INSERT INTO t1 VALUES(100, 'abcde');
SELECT (SELECT s2 FROM t1);
```

The subquery in this `SELECT` returns a single value ('abcde') that has a data type of `CHAR`, a length of 5, a character set and collation equal to the defaults in effect at `CREATE TABLE` time, and an indication that the value in the column can be `NULL`. Nullability of the value selected by a scalar subquery is not copied because if the subquery result is empty, the result is `NULL`. For the subquery just shown, if `t1` were empty, the result would be `NULL` even though `s2` is `NOT NULL`.

There are a few contexts in which a scalar subquery cannot be used. If a statement permits only a literal value, you cannot use a subquery. For example, `LIMIT` requires literal integer arguments, and `LOAD DATA INFILE` requires a literal string file name. You cannot use subqueries to supply these values.

When you see examples in the following sections that contain the rather spartan construct (`SELECT column1 FROM t1`), imagine that your own code contains much more diverse and complex constructions.

Suppose that we make two tables:

```
CREATE TABLE t1 (s1 INT);
INSERT INTO t1 VALUES (1);
CREATE TABLE t2 (s1 INT);
INSERT INTO t2 VALUES (2);
```

Then perform a `SELECT`:

```
SELECT (SELECT s1 FROM t2) FROM t1;
```

The result is `2` because there is a row in `t2` containing a column `s1` that has a value of `2`.

A scalar subquery can be part of an expression, but remember the parentheses, even if the subquery is an operand that provides an argument for a function. For example:

```
SELECT UPPER((SELECT s1 FROM t1)) FROM t2;
```

### 13.2.10.2. Comparisons Using Subqueries

The most common use of a subquery is in the form:

```
non_subquery_operand comparison_operator (subquery)
```

Where *comparison\_operator* is one of these operators:

```
= > < >= <= <> != <=>
```

For example:

```
... WHERE 'a' = (SELECT column1 FROM t1)
```

MySQL also permits this construct:

```
non_subquery_operand LIKE (subquery)
```

At one time the only legal place for a subquery was on the right side of a comparison, and you might still find some old DBMSs that insist on this.

Here is an example of a common-form subquery comparison that you cannot do with a join. It finds all the rows in table *t1* for which the *column1* value is equal to a maximum value in table *t2*:

```
SELECT * FROM t1
WHERE column1 = (SELECT MAX(column2) FROM t2);
```

Here is another example, which again is impossible with a join because it involves aggregating for one of the tables. It finds all rows in table *t1* containing a value that occurs twice in a given column:

```
SELECT * FROM t1 AS t
WHERE 2 = (SELECT COUNT(*) FROM t1 WHERE t1.id = t.id);
```

For a comparison of the subquery to a scalar, the subquery must return a scalar. For a comparison of the subquery to a row constructor, the subquery must be a row subquery that returns a row with the same number of values as the row constructor. See [Section 13.2.10.5, “Row Subqueries”](#).

### 13.2.10.3. Subqueries with **ANY**, **IN**, or **SOME**

Syntax:

```
operand comparison_operator ANY (subquery)
operand IN (subquery)
operand comparison_operator SOME (subquery)
```

Where *comparison\_operator* is one of these operators:

```
= > < >= <= <> !=
```

The **ANY** keyword, which must follow a comparison operator, means “return **TRUE** if the comparison is **TRUE** for **ANY** of the values in the column that the subquery returns.” For example:

```
SELECT s1 FROM t1 WHERE s1 > ANY (SELECT s1 FROM t2);
```

Suppose that there is a row in table *t1* containing (10). The expression is **TRUE** if table *t2* contains (21,14,7) because there is a value 7 in *t2* that is less than 10. The expression is **FALSE** if table *t2* contains (20,10), or if table *t2* is empty. The expression is *unknown* (that is, **NULL**) if table *t2* contains (NULL,NULL,NULL).

When used with a subquery, the word **IN** is an alias for = **ANY**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 = ANY (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 IN (SELECT s1 FROM t2);
```

**IN** and = **ANY** are not synonyms when used with an expression list. **IN** can take an expression list, but = **ANY** cannot. See [Section 12.3.2, “Comparison Functions and Operators”](#).

**NOT IN** is not an alias for <> **ANY**, but for <> **ALL**. See [Section 13.2.10.4, “Subqueries with ALL”](#).

The word **SOME** is an alias for **ANY**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 <> ANY (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 <> SOME (SELECT s1 FROM t2);
```

Use of the word **SOME** is rare, but this example shows why it might be useful. To most people, the English phrase “a is not equal to any b” means “there is no b which is equal to a,” but that is not what is meant by the SQL syntax. The syntax means “there is some b to which a is not equal.” Using **<> SOME** instead helps ensure that everyone understands the true meaning of the query.

#### 13.2.10.4. Subqueries with **ALL**

Syntax:

```
operand comparison_operator ALL (subquery)
```

The word **ALL**, which must follow a comparison operator, means “return **TRUE** if the comparison is **TRUE** for **ALL** of the values in the column that the subquery returns.” For example:

```
SELECT s1 FROM t1 WHERE s1 > ALL (SELECT s1 FROM t2);
```

Suppose that there is a row in table **t1** containing (10). The expression is **TRUE** if table **t2** contains (-5, 0, +5) because 10 is greater than all three values in **t2**. The expression is **FALSE** if table **t2** contains (12, 6, NULL, -100) because there is a single value 12 in table **t2** that is greater than 10. The expression is *unknown* (that is, **NULL**) if table **t2** contains (0, NULL, 1).

Finally, the expression is **TRUE** if table **t2** is empty. So, the following expression is **TRUE** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > ALL (SELECT s1 FROM t2);
```

But this expression is **NULL** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > (SELECT s1 FROM t2);
```

In addition, the following expression is **NULL** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > ALL (SELECT MAX(s1) FROM t2);
```

In general, *tables containing NULL values* and *empty tables* are “edge cases.” When writing subqueries, always consider whether you have taken those two possibilities into account.

**NOT IN** is an alias for **<> ALL**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 <> ALL (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 NOT IN (SELECT s1 FROM t2);
```

#### 13.2.10.5. Row Subqueries

The discussion to this point has been of scalar or column subqueries; that is, subqueries that return a single value or a column of values. A *row subquery* is a subquery variant that returns a single row and can thus return more than one column value. Legal operators for row subquery comparisons are:

```
= > < >= <= <> != <=>
```

Here are two examples:

```
SELECT * FROM t1
WHERE (col1,col2) = (SELECT col3, col4 FROM t2 WHERE id = 10);
SELECT * FROM t1
WHERE ROW(col1,col2) = (SELECT col3, col4 FROM t2 WHERE id = 10);
```

For both queries, if the table **t2** contains a single row with **id = 10**, the subquery returns a single row. If this row has **col3** and **col4** values equal to the **col1** and **col2** values of any rows in **t1**, the **WHERE** expression is **TRUE** and each query returns those **t1** rows. If the **t2** row **col3** and **col4** values are not equal the **col1** and **col2** values of any **t1** row, the expression is **FALSE** and the query returns an empty result set. The expression is *unknown* (that is, **NULL**) if the subquery produces no rows. An

error occurs if the subquery produces multiple rows because a row subquery can return at most one row.

The expressions `(1, 2)` and `ROW(1, 2)` are sometimes called *row constructors*. The two are equivalent. The row constructor and the row returned by the subquery must contain the same number of values.

A row constructor is used for comparisons with subqueries that return two or more columns. When a subquery returns a single column, this is regarded as a scalar value and not as a row, so a row constructor cannot be used with a subquery that does not return at least two columns. Thus, the following query fails with a syntax error:

```
SELECT * FROM t1 WHERE ROW(1) = (SELECT column1 FROM t2)
```

Row constructors are legal in other contexts. For example, the following two statements are semantically equivalent (and are handled in the same way by the optimizer):

```
SELECT * FROM t1 WHERE (column1, column2) = (1, 1);
SELECT * FROM t1 WHERE column1 = 1 AND column2 = 1;
```

The following query answers the request, “find all rows in table `t1` that also exist in table `t2`”:

```
SELECT column1, column2, column3
FROM t1
WHERE (column1, column2, column3) IN
      (SELECT column1, column2, column3 FROM t2);
```

### 13.2.10.6. Subqueries with **EXISTS** or **NOT EXISTS**

If a subquery returns any rows at all, `EXISTS subquery` is `TRUE`, and `NOT EXISTS subquery` is `FALSE`. For example:

```
SELECT column1 FROM t1 WHERE EXISTS (SELECT * FROM t2);
```

Traditionally, an `EXISTS` subquery starts with `SELECT *`, but it could begin with `SELECT 5` or `SELECT column1` or anything at all. MySQL ignores the `SELECT` list in such a subquery, so it makes no difference.

For the preceding example, if `t2` contains any rows, even rows with nothing but `NULL` values, the `EXISTS` condition is `TRUE`. This is actually an unlikely example because a `[NOT] EXISTS` subquery almost always contains correlations. Here are some more realistic examples:

- What kind of store is present in one or more cities?

```
SELECT DISTINCT store_type FROM stores
WHERE EXISTS (SELECT * FROM cities_stores
              WHERE cities_stores.store_type = stores.store_type);
```

- What kind of store is present in no cities?

```
SELECT DISTINCT store_type FROM stores
WHERE NOT EXISTS (SELECT * FROM cities_stores
                  WHERE cities_stores.store_type = stores.store_type);
```

- What kind of store is present in all cities?

```
SELECT DISTINCT store_type FROM stores s1
WHERE NOT EXISTS (
  SELECT * FROM cities WHERE NOT EXISTS (
    SELECT * FROM cities_stores
    WHERE cities_stores.city = cities.city
    AND cities_stores.store_type = stores.store_type));
```

The last example is a double-nested `NOT EXISTS` query. That is, it has a `NOT EXISTS` clause within a `NOT EXISTS` clause. Formally, it answers the question “does a city exist with a store that is not in `Stores`”? But it is easier to say that a nested `NOT EXISTS` answers the question “is `x TRUE` for all `y`?”

### 13.2.10.7. Correlated Subqueries

A *correlated subquery* is a subquery that contains a reference to a table that also appears in the outer query. For example:

```
SELECT * FROM t1
WHERE column1 = ANY (SELECT column1 FROM t2
                     WHERE t2.column2 = t1.column2);
```

Notice that the subquery contains a reference to a column of `t1`, even though the subquery's `FROM` clause does not mention a table `t1`. So, MySQL looks outside the subquery, and finds `t1` in the outer query.

Suppose that table `t1` contains a row where `column1 = 5` and `column2 = 6`; meanwhile, table `t2` contains a row where `column1 = 5` and `column2 = 7`. The simple expression `... WHERE column1 = ANY (SELECT column1 FROM t2)` would be `TRUE`, but in this example, the `WHERE` clause within the subquery is `FALSE` (because `(5,6)` is not equal to `(5,7)`), so the expression as a whole is `FALSE`.

**Scoping rule:** MySQL evaluates from inside to outside. For example:

```
SELECT column1 FROM t1 AS x
WHERE x.column1 = (SELECT column1 FROM t2 AS x
                  WHERE x.column1 = (SELECT column1 FROM t3
                                    WHERE x.column2 = t3.column1));
```

In this statement, `x.column2` must be a column in table `t2` because `SELECT column1 FROM t2 AS x ...` renames `t2`. It is not a column in table `t1` because `SELECT column1 FROM t1 ...` is an outer query that is *farther out*.

For subqueries in `HAVING` or `ORDER BY` clauses, MySQL also looks for column names in the outer select list.

For certain cases, a correlated subquery is optimized. For example:

```
val IN (SELECT key_val FROM tbl_name WHERE correlated_condition)
```

Otherwise, they are inefficient and likely to be slow. Rewriting the query as a join might improve performance.

Aggregate functions in correlated subqueries may contain outer references, provided the function contains nothing but outer references, and provided the function is not contained in another function or expression.

### 13.2.10.8. Subqueries in the `FROM` Clause

Subqueries are legal in a `SELECT` statement's `FROM` clause. The actual syntax is:

```
SELECT ... FROM (subquery) [AS] name ...
```

The `[AS] name` clause is mandatory, because every table in a `FROM` clause must have a name. Any columns in the `subquery` select list must have unique names.

For the sake of illustration, assume that you have this table:

```
CREATE TABLE t1 (s1 INT, s2 CHAR(5), s3 FLOAT);
```

Here is how to use a subquery in the `FROM` clause, using the example table:

```
INSERT INTO t1 VALUES (1,'1',1.0);
INSERT INTO t1 VALUES (2,'2',2.0);
SELECT sb1,sb2,sb3
FROM (SELECT s1 AS sb1, s2 AS sb2, s3*2 AS sb3 FROM t1) AS sb
WHERE sb1 > 1;
```

Result: 2, '2', 4.0.

Here is another example: Suppose that you want to know the average of a set of sums for a grouped table. This does not work:

```
SELECT AVG(SUM(column1)) FROM t1 GROUP BY column1;
```

However, this query provides the desired information:

```
SELECT AVG(sum_column1)
FROM (SELECT SUM(column1) AS sum_column1
      FROM t1 GROUP BY column1) AS t1;
```

Notice that the column name used within the subquery (`sum_column1`) is recognized in the outer query.

Subqueries in the `FROM` clause can return a scalar, column, row, or table. Subqueries in the `FROM` clause cannot be correlated subqueries, unless used within the `ON` clause of a `JOIN` operation.

Before MySQL 5.6.3, subqueries in the `FROM` clause are executed even for the `EXPLAIN` statement (that is, derived temporary tables are materialized). This occurs because upper-level queries need information about all tables during the optimization phase, and the table represented by a subquery in the `FROM` clause is unavailable unless the subquery is executed. As of MySQL 5.6.3, the optimizer determines information about derived tables in a different way and materialization of them does not occur for `EXPLAIN`. See [Section 8.13.16.3, “Optimizing Subqueries in the `FROM` Clause \(Derived Tables\)”](#).

It is possible under certain circumstances to modify table data using `EXPLAIN SELECT`. This can occur if the outer query accesses any tables and an inner query invokes a stored function that changes one or more rows of a table. Suppose that there are two tables `t1` and `t2` in database `d1`, created as shown here:

```
mysql> CREATE DATABASE d1;
Query OK, 1 row affected (0.00 sec)

mysql> USE d1;
Database changed

mysql> CREATE TABLE t1 (c1 INT);
Query OK, 0 rows affected (0.15 sec)

mysql> CREATE TABLE t2 (c1 INT);
Query OK, 0 rows affected (0.08 sec)
```

Now we create a stored function `f1` which modifies `t2`:

```
mysql> DELIMITER //
mysql> CREATE FUNCTION f1(p1 INT) RETURNS INT
mysql> BEGIN
mysql>     INSERT INTO t2 VALUES (p1);
mysql>     RETURN p1;
mysql> END //
Query OK, 0 rows affected (0.01 sec)

mysql> DELIMITER ;
```

Referencing the function directly in an `EXPLAIN SELECT` does not have any effect on `t2`, as shown here:

```
mysql> SELECT * FROM t2;
Empty set (0.00 sec)

mysql> EXPLAIN SELECT f1(5);
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | NULL | NULL | NULL | NULL | NULL | NULL | NULL | No tables used |
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT * FROM t2;
Empty set (0.00 sec)

```

This is because the `SELECT` statement did not reference any tables, as can be seen in the `table` and `Extra` columns of the output. This is also true of the following nested `SELECT`:

```

mysql> EXPLAIN SELECT NOW() AS a1, (SELECT f1(5)) AS a2;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | PRIMARY | NULL | NULL | NULL | NULL | NULL | NULL | NULL | No tables used |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Note | 1249 | Select 2 was reduced during optimization |
+-----+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT * FROM t2;
Empty set (0.00 sec)

```

However, if the outer `SELECT` references any tables, the optimizer executes the statement in the subquery as well:

```

mysql> EXPLAIN SELECT * FROM t1 AS a1, (SELECT f1(5)) AS a2;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | PRIMARY | a1 | system | NULL | NULL | NULL | NULL | 0 | const row not found |
| 1 | PRIMARY | <derived2> | system | NULL | NULL | NULL | NULL | 1 | |
| 2 | DERIVED | NULL | NULL | NULL | NULL | NULL | NULL | NULL | No tables used |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
3 rows in set (0.00 sec)

mysql> SELECT * FROM t2;
+-----+
| c1 |
+-----+
| 5 |
+-----+
1 row in set (0.00 sec)

```

This also means that an `EXPLAIN SELECT` statement such as the one shown here may take a long time to execute because the `BENCHMARK()` [1266] function is executed once for each row in `t1`:

```
EXPLAIN SELECT * FROM t1 AS a1, (SELECT BENCHMARK(1000000, MD5(NOW())));
```

### 13.2.10.9. Subquery Errors

There are some errors that apply only to subqueries. This section describes them.

- Unsupported subquery syntax:

```

ERROR 1235 (ER_NOT_SUPPORTED_YET)
SQLSTATE = 42000
Message = "This version of MySQL doesn't yet support
'LIMIT & IN/ALL/ANY/SOME subquery'"

```

This means that MySQL does not support statements of the following form:

```
SELECT * FROM t1 WHERE s1 IN (SELECT s2 FROM t2 ORDER BY s1 LIMIT 1)
```

- Incorrect number of columns from subquery:



```
ERROR 1241 (ER_OPERAND_COL)
SQLSTATE = 21000
Message = "Operand should contain 1 column(s)"
```

This error occurs in cases like this:

```
SELECT (SELECT column1, column2 FROM t2) FROM t1;
```

You may use a subquery that returns multiple columns, if the purpose is row comparison. In other contexts, the subquery must be a scalar operand. See [Section 13.2.10.5, "Row Subqueries"](#).

- Incorrect number of rows from subquery:

```
ERROR 1242 (ER_SUBSELECT_NO_1_ROW)
SQLSTATE = 21000
Message = "Subquery returns more than 1 row"
```

This error occurs for statements where the subquery must return at most one row but returns multiple rows. Consider the following example:

```
SELECT * FROM t1 WHERE column1 = (SELECT column1 FROM t2);
```

If `SELECT column1 FROM t2` returns just one row, the previous query will work. If the subquery returns more than one row, error 1242 will occur. In that case, the query should be rewritten as:

```
SELECT * FROM t1 WHERE column1 = ANY (SELECT column1 FROM t2);
```

- Incorrectly used table in subquery:

```
Error 1093 (ER_UPDATE_TABLE_USED)
SQLSTATE = HY000
Message = "You can't specify target table 'x'
for update in FROM clause"
```

This error occurs in cases such as the following, which attempts to modify a table and select from the same table in the subquery:

```
UPDATE t1 SET column2 = (SELECT MAX(column1) FROM t1);
```

You can use a subquery for assignment within an `UPDATE` statement because subqueries are legal in `UPDATE` and `DELETE` statements as well as in `SELECT` statements. However, you cannot use the same table (in this case, table `t1`) for both the subquery `FROM` clause and the update target.

For transactional storage engines, the failure of a subquery causes the entire statement to fail. For nontransactional storage engines, data modifications made before the error was encountered are preserved.

### 13.2.10.10. Optimizing Subqueries

Development is ongoing, so no optimization tip is reliable for the long term. The following list provides some interesting tricks that you might want to play with:

- Use subquery clauses that affect the number or order of the rows in the subquery. For example:

```
SELECT * FROM t1 WHERE t1.column1 IN
  (SELECT column1 FROM t2 ORDER BY column1);
SELECT * FROM t1 WHERE t1.column1 IN
  (SELECT DISTINCT column1 FROM t2);
SELECT * FROM t1 WHERE EXISTS
  (SELECT * FROM t2 LIMIT 1);
```

- Replace a join with a subquery. For example, try this:

```
SELECT DISTINCT column1 FROM t1 WHERE t1.column1 IN (
  SELECT column1 FROM t2);
```

Instead of this:

```
SELECT DISTINCT t1.column1 FROM t1, t2
WHERE t1.column1 = t2.column1;
```

- Some subqueries can be transformed to joins for compatibility with older versions of MySQL that do not support subqueries. However, in some cases, converting a subquery to a join may improve performance. See [Section 13.2.10.11, “Rewriting Subqueries as Joins”](#).
- Move clauses from outside to inside the subquery. For example, use this query:

```
SELECT * FROM t1
WHERE s1 IN (SELECT s1 FROM t1 UNION ALL SELECT s1 FROM t2);
```

Instead of this query:

```
SELECT * FROM t1
WHERE s1 IN (SELECT s1 FROM t1) OR s1 IN (SELECT s1 FROM t2);
```

For another example, use this query:

```
SELECT (SELECT column1 + 5 FROM t1) FROM t2;
```

Instead of this query:

```
SELECT (SELECT column1 FROM t1) + 5 FROM t2;
```

- Use a row subquery instead of a correlated subquery. For example, use this query:

```
SELECT * FROM t1
WHERE (column1,column2) IN (SELECT column1,column2 FROM t2);
```

Instead of this query:

```
SELECT * FROM t1
WHERE EXISTS (SELECT * FROM t2 WHERE t2.column1=t1.column1
AND t2.column2=t1.column2);
```

- Use `NOT (a = ANY (...))` rather than `a <> ALL (...)`.
- Use `x = ANY (table containing (1,2))` rather than `x=1 OR x=2`.
- Use `= ANY` rather than `EXISTS`.
- For uncorrelated subqueries that always return one row, `IN` is always slower than `=`. For example, use this query:

```
SELECT * FROM t1
WHERE t1.col_name = (SELECT a FROM t2 WHERE b = some_const);
```

Instead of this query:

```
SELECT * FROM t1
WHERE t1.col_name IN (SELECT a FROM t2 WHERE b = some_const);
```

These tricks might cause programs to go faster or slower. Using MySQL facilities like the `BENCHMARK()` [\[1266\]](#) function, you can get an idea about what helps in your own situation. See [Section 12.14, “Information Functions”](#).

Some optimizations that MySQL itself makes are:

- MySQL executes uncorrelated subqueries only once. Use `EXPLAIN` to make sure that a given subquery really is uncorrelated.
- MySQL rewrites `IN`, `ALL`, `ANY`, and `SOME` subqueries in an attempt to take advantage of the possibility that the select-list columns in the subquery are indexed.

- MySQL replaces subqueries of the following form with an index-lookup function, which `EXPLAIN` describes as a special join type (`unique_subquery` [899] or `index_subquery` [899]):

```
... IN (SELECT indexed_column FROM single_table ...)
```

- MySQL enhances expressions of the following form with an expression involving `MIN()` [1284] or `MAX()` [1284], unless `NULL` values or empty sets are involved:

```
value {ALL|ANY|SOME} {> | < | >= | <=} (uncorrelated subquery)
```

For example, this `WHERE` clause:

```
WHERE 5 > ALL (SELECT x FROM t)
```

might be treated by the optimizer like this:

```
WHERE 5 > (SELECT MAX(x) FROM t)
```

See also [MySQL Internals: How MySQL Transforms Subqueries](#).

### 13.2.10.11. Rewriting Subqueries as Joins

Sometimes there are other ways to test membership in a set of values than by using a subquery. Also, on some occasions, it is not only possible to rewrite a query without a subquery, but it can be more efficient to make use of some of these techniques rather than to use subqueries. One of these is the `IN()` construct:

For example, this query:

```
SELECT * FROM t1 WHERE id IN (SELECT id FROM t2);
```

Can be rewritten as:

```
SELECT DISTINCT t1.* FROM t1, t2 WHERE t1.id=t2.id;
```

The queries:

```
SELECT * FROM t1 WHERE id NOT IN (SELECT id FROM t2);
SELECT * FROM t1 WHERE NOT EXISTS (SELECT id FROM t2 WHERE t1.id=t2.id);
```

Can be rewritten as:

```
SELECT table1.*
FROM table1 LEFT JOIN table2 ON table1.id=table2.id
WHERE table2.id IS NULL;
```

A `LEFT [OUTER] JOIN` can be faster than an equivalent subquery because the server might be able to optimize it better—a fact that is not specific to MySQL Server alone. Prior to SQL-92, outer joins did not exist, so subqueries were the only way to do certain things. Today, MySQL Server and many other modern database systems offer a wide range of outer join types.

MySQL Server supports multiple-table `DELETE` statements that can be used to efficiently delete rows based on information from one table or even from many tables at the same time. Multiple-table `UPDATE` statements are also supported. See [Section 13.2.2, “DELETE Syntax”](#), and [Section 13.2.11, “UPDATE Syntax”](#).

## 13.2.11. UPDATE Syntax

Single-table syntax:

```
UPDATE [LOW_PRIORITY] [IGNORE] table_reference
SET col_name1={expr1|DEFAULT} [, col_name2={expr2|DEFAULT}] ...
[WHERE where_condition]
[ORDER BY ...]
[LIMIT row_count]
```

Multiple-table syntax:

```
UPDATE [LOW_PRIORITY] [IGNORE] table_references
  SET col_name1={expr1|DEFAULT} [, col_name2={expr2|DEFAULT}] ...
  [WHERE where_condition]
```

For the single-table syntax, the `UPDATE` statement updates columns of existing rows in the named table with new values. The `SET` clause indicates which columns to modify and the values they should be given. Each value can be given as an expression, or the keyword `DEFAULT` to set a column explicitly to its default value. The `WHERE` clause, if given, specifies the conditions that identify which rows to update. With no `WHERE` clause, all rows are updated. If the `ORDER BY` clause is specified, the rows are updated in the order that is specified. The `LIMIT` clause places a limit on the number of rows that can be updated.

For the multiple-table syntax, `UPDATE` updates rows in each table named in *table\_references* that satisfy the conditions. In this case, `ORDER BY` and `LIMIT` cannot be used.

For partitioned tables, both the single-table and multiple-table forms of this statement support the use of a `PARTITION` option as part of a table reference. This option takes a list of one or more partitions or subpartitions (or both). Only the partitions (or subpartitions) listed are checked for matches, and a row that is not in any of these partitions or subpartitions is not updated, whether it satisfies the *where\_condition* or not.



#### Note

Unlike the case when using `PARTITION` with an `INSERT` or `REPLACE` statement, an otherwise valid `UPDATE ... PARTITION` statement is considered successful even if no rows in the listed partitions (or subpartitions) match the *where\_condition*.

See [Section 17.5, “Partition Selection”](#), for more information and examples.

*where\_condition* is an expression that evaluates to true for each row to be updated. For expression syntax, see [Section 9.5, “Expression Syntax”](#).

*table\_references* and *where\_condition* are specified as described in [Section 13.2.9, “SELECT Syntax”](#).

You need the `UPDATE [756]` privilege only for columns referenced in an `UPDATE` that are actually updated. You need only the `SELECT [755]` privilege for any columns that are read but not modified.

The `UPDATE` statement supports the following modifiers:

- With the `LOW_PRIORITY` keyword, execution of the `UPDATE` is delayed until no other clients are reading from the table. This affects only storage engines that use only table-level locking (such as `MyISAM`, `MEMORY`, and `MERGE`).
- With the `IGNORE` keyword, the update statement does not abort even if errors occur during the update. Rows for which duplicate-key conflicts occur are not updated. Rows for which columns are updated to values that would cause data conversion errors are updated to the closest valid values instead.

In MySQL 5.6.4 and later, `UPDATE IGNORE` statements, including those having an `ORDER BY` clause, are flagged as unsafe for statement-based replication. (This is because the order in which the rows are updated determines which rows are ignored.) With this change, such statements produce a warning in the log when using statement-based mode and are logged using the row-based format when using `MIXED` mode. (Bug #11758262, Bug #50439) See [Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”](#), for more information.

If you access a column from the table to be updated in an expression, `UPDATE` uses the current value of the column. For example, the following statement sets `col1` to one more than its current value:

```
UPDATE t1 SET col1 = col1 + 1;
```

The second assignment in the following statement sets `col2` to the current (updated) `col1` value, not the original `col1` value. The result is that `col1` and `col2` have the same value. This behavior differs from standard SQL.

```
UPDATE t1 SET col1 = col1 + 1, col2 = col1;
```

Single-table `UPDATE` assignments are generally evaluated from left to right. For multiple-table updates, there is no guarantee that assignments are carried out in any particular order.

If you set a column to the value it currently has, MySQL notices this and does not update it.

If you update a column that has been declared `NOT NULL` by setting to `NULL`, an error occurs if strict SQL mode is enabled; otherwise, the column is set to the implicit default value for the column data type and the warning count is incremented. The implicit default value is `0` for numeric types, the empty string (`' '`) for string types, and the “zero” value for date and time types. See [Section 11.5, “Data Type Default Values”](#).

`UPDATE` returns the number of rows that were actually changed. The `mysql_info()` C API function returns the number of rows that were matched and updated and the number of warnings that occurred during the `UPDATE`.

You can use `LIMIT row_count` to restrict the scope of the `UPDATE`. A `LIMIT` clause is a row-matched restriction. The statement stops as soon as it has found `row_count` rows that satisfy the `WHERE` clause, whether or not they actually were changed.

If an `UPDATE` statement includes an `ORDER BY` clause, the rows are updated in the order specified by the clause. This can be useful in certain situations that might otherwise result in an error. Suppose that a table `t` contains a column `id` that has a unique index. The following statement could fail with a duplicate-key error, depending on the order in which rows are updated:

```
UPDATE t SET id = id + 1;
```

For example, if the table contains 1 and 2 in the `id` column and 1 is updated to 2 before 2 is updated to 3, an error occurs. To avoid this problem, add an `ORDER BY` clause to cause the rows with larger `id` values to be updated before those with smaller values:

```
UPDATE t SET id = id + 1 ORDER BY id DESC;
```

You can also perform `UPDATE` operations covering multiple tables. However, you cannot use `ORDER BY` or `LIMIT` with a multiple-table `UPDATE`. The `table_references` clause lists the tables involved in the join. Its syntax is described in [Section 13.2.9.2, “JOIN Syntax”](#). Here is an example:

```
UPDATE items,month SET items.price=month.price
WHERE items.id=month.id;
```

The preceding example shows an inner join that uses the comma operator, but multiple-table `UPDATE` statements can use any type of join permitted in `SELECT` statements, such as `LEFT JOIN`.

If you use a multiple-table `UPDATE` statement involving `InnoDB` tables for which there are foreign key constraints, the MySQL optimizer might process tables in an order that differs from that of their parent/child relationship. In this case, the statement fails and rolls back. Instead, update a single table and rely on the `ON UPDATE` capabilities that `InnoDB` provides to cause the other tables to be modified accordingly. See [Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#).

Currently, you cannot update a table and select from the same table in a subquery.

Index hints (see [Section 13.2.9.3, “Index Hint Syntax”](#)) are accepted but ignored for `UPDATE` statements.

Prior to MySQL 5.6.6, an `UPDATE` on a partitioned table using a storage engine such as `MyISAM` that employs table-level locks locked all partitions of the table. This was true even for `UPDATE ... PARTITION` queries. (This did not and does not occur with storage engines such as `InnoDB` that employ row-level locking.) In MySQL 5.6.6 and later, MySQL uses partition lock pruning, so that only

partitions containing rows matching the `UPDATE` statement's `WHERE` clause are actually locked, as long as none of the table's partitioning columns are updated. For more information, see [Section 17.6.4, "Partitioning and Locking"](#).

## 13.3. MySQL Transactional and Locking Statements

MySQL supports local transactions (within a given client session) through statements such as `SET autocommit`, `START TRANSACTION`, `COMMIT`, and `ROLLBACK`. See [Section 13.3.1, "START TRANSACTION, COMMIT, and ROLLBACK Syntax"](#). XA transaction support enables MySQL to participate in distributed transactions as well. See [Section 13.3.7, "XA Transactions"](#).

### 13.3.1. START TRANSACTION, COMMIT, and ROLLBACK Syntax

```
START TRANSACTION
    [transaction_characteristic [, transaction_characteristic] ...]

transaction_characteristic:
    WITH CONSISTENT SNAPSHOT
    | READ WRITE
    | READ ONLY

BEGIN [WORK]
COMMIT [WORK] [AND [NO] CHAIN] [[NO] RELEASE]
ROLLBACK [WORK] [AND [NO] CHAIN] [[NO] RELEASE]
SET autocommit = {0 | 1}
```

These statements provide control over use of [transactions](#):

- `START TRANSACTION` or `BEGIN` start a new transaction.
- `COMMIT` commits the current transaction, making its changes permanent.
- `ROLLBACK` rolls back the current transaction, canceling its changes.
- `SET autocommit` disables or enables the default autocommit mode for the current session.

By default, MySQL runs with [autocommit](#) mode enabled. This means that as soon as you execute a statement that updates (modifies) a table, MySQL stores the update on disk to make it permanent. The change cannot be rolled back.

To disable autocommit mode implicitly for a single series of statements, use the `START TRANSACTION` statement:

```
START TRANSACTION;
SELECT @A:=SUM(salary) FROM table1 WHERE type=1;
UPDATE table2 SET summary=@A WHERE type=1;
COMMIT;
```

With `START TRANSACTION`, autocommit remains disabled until you end the transaction with `COMMIT` or `ROLLBACK`. The autocommit mode then reverts to its previous state.

`START TRANSACTION` permits several modifiers that control transaction characteristics. To specify multiple modifiers, separate them by commas.

- The `WITH CONSISTENT SNAPSHOT` modifier starts a [consistent read](#) for storage engines that are capable of it. This applies only to `InnoDB`. The effect is the same as issuing a `START TRANSACTION` followed by a `SELECT` from any `InnoDB` table. See [Section 14.2.3.3, "Consistent Nonlocking Reads"](#). The `WITH CONSISTENT SNAPSHOT` modifier does not change the current transaction [isolation level](#), so it provides a consistent snapshot only if the current isolation level is one that permits consistent read (`REPEATABLE READ` [1469] or `SERIALIZABLE` [1470]).
- The `READ WRITE` and `READ ONLY` modifiers set the transaction access mode. They permit or prohibit changes to tables used in the transaction. The `READ ONLY` restriction prevents the

transaction from modifying or locking both transactional and nontransactional tables that are visible to other transactions; the transaction can still modify or lock temporary tables. These modifiers are available as of MySQL 5.6.5.

MySQL enables extra optimizations for queries on [InnoDB](#) tables when the transaction is known to be read-only. Specifying [READ ONLY](#) ensures these optimizations are applied in cases where the read-only status cannot be determined automatically. See [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#) for more information.

If no access mode is specified, the default mode applies. Unless the default has been changed, it is read/write. It is not permitted to specify both [READ WRITE](#) and [READ ONLY](#) in the same statement.

In read-only mode, it remains possible to change tables created with the [TEMPORARY](#) keyword using DML statements. Changes made with DDL statements are not permitted, just as with permanent tables.

For additional information about transaction access mode, including ways to change the default mode, see [Section 13.3.6, “SET TRANSACTION Syntax”](#).

If the [read\\_only \[532\]](#) system variable is enabled, explicitly starting a transaction with [START TRANSACTION READ WRITE](#) requires the [SUPER \[756\]](#) privilege.



### Important

Many APIs used for writing MySQL client applications (such as JDBC) provide their own methods for starting transactions that can (and sometimes should) be used instead of sending a [START TRANSACTION](#) statement from the client. See [Chapter 21, Connectors and APIs](#), or the documentation for your API, for more information.

To disable autocommit mode explicitly, use the following statement:

```
SET autocommit=0;
```

After disabling autocommit mode by setting the [autocommit \[464\]](#) variable to zero, changes to transaction-safe tables (such as those for [InnoDB](#) or [NDB](#)) are not made permanent immediately. You must use [COMMIT](#) to store your changes to disk or [ROLLBACK](#) to ignore the changes.

[autocommit \[464\]](#) is a session variable and must be set for each session. To disable autocommit mode for each new connection, see the description of the [autocommit \[464\]](#) system variable at [Section 5.1.4, “Server System Variables”](#).

[BEGIN](#) and [BEGIN WORK](#) are supported as aliases of [START TRANSACTION](#) for initiating a transaction. [START TRANSACTION](#) is standard SQL syntax, is the recommended way to start an ad-hoc transaction, and permits modifiers that [BEGIN](#) does not.

The [BEGIN](#) statement differs from the use of the [BEGIN](#) keyword that starts a [BEGIN ... END](#) compound statement. The latter does not begin a transaction. See [Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”](#).



### Note

Within all stored programs (stored procedures and functions, triggers, and events), the parser treats [BEGIN \[WORK\]](#) as the beginning of a [BEGIN ... END](#) block. Begin a transaction in this context with [START TRANSACTION](#) instead.

The optional [WORK](#) keyword is supported for [COMMIT](#) and [ROLLBACK](#), as are the [CHAIN](#) and [RELEASE](#) clauses. [CHAIN](#) and [RELEASE](#) can be used for additional control over transaction completion. The value of the [completion\\_type \[470\]](#) system variable determines the default completion behavior. See [Section 5.1.4, “Server System Variables”](#).



The `AND CHAIN` clause causes a new transaction to begin as soon as the current one ends, and the new transaction has the same isolation level as the just-terminated transaction. The `RELEASE` clause causes the server to disconnect the current client session after terminating the current transaction. Including the `NO` keyword suppresses `CHAIN` or `RELEASE` completion, which can be useful if the `completion_type` [470] system variable is set to cause chaining or release completion by default.

Beginning a transaction causes any pending transaction to be committed. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#), for more information.

Beginning a transaction also causes table locks acquired with `LOCK TABLES` to be released, as though you had executed `UNLOCK TABLES`. Beginning a transaction does not release a global read lock acquired with `FLUSH TABLES WITH READ LOCK`.

For best results, transactions should be performed using only tables managed by a single transaction-safe storage engine. Otherwise, the following problems can occur:

- If you use tables from more than one transaction-safe storage engine (such as `InnoDB`), and the transaction isolation level is not `SERIALIZABLE` [1470], it is possible that when one transaction commits, another ongoing transaction that uses the same tables will see only some of the changes made by the first transaction. That is, the atomicity of transactions is not guaranteed with mixed engines and inconsistencies can result. (If mixed-engine transactions are infrequent, you can use `SET TRANSACTION ISOLATION LEVEL` to set the isolation level to `SERIALIZABLE` [1470] on a per-transaction basis as necessary.)
- If you use tables that are not transaction-safe within a transaction, changes to those tables are stored at once, regardless of the status of autocommit mode.
- If you issue a `ROLLBACK` statement after updating a nontransactional table within a transaction, an `ER_WARNING_NOT_COMPLETE_ROLLBACK` [3673] warning occurs. Changes to transaction-safe tables are rolled back, but not changes to nontransaction-safe tables.

Each transaction is stored in the binary log in one chunk, upon `COMMIT`. Transactions that are rolled back are not logged. (**Exception:** Modifications to nontransactional tables cannot be rolled back. If a transaction that is rolled back includes modifications to nontransactional tables, the entire transaction is logged with a `ROLLBACK` statement at the end to ensure that modifications to the nontransactional tables are replicated.) See [Section 5.2.4, “The Binary Log”](#).

You can change the isolation level or access mode for transactions with the `SET TRANSACTION` statement. See [Section 13.3.6, “SET TRANSACTION Syntax”](#).

Rolling back can be a slow operation that may occur implicitly without the user having explicitly asked for it (for example, when an error occurs). Because of this, `SHOW PROCESSLIST` displays `Rolling back` in the `State` column for the session, not only for explicit rollbacks performed with the `ROLLBACK` statement but also for implicit rollbacks.



#### Note

In MySQL 5.6, `BEGIN`, `COMMIT`, and `ROLLBACK` are not affected by `--replicate-do-db` [1976] or `--replicate-ignore-db` [1977] rules.

## 13.3.2. Statements That Cannot Be Rolled Back

Some statements cannot be rolled back. In general, these include data definition language (DDL) statements, such as those that create or drop databases, those that create, drop, or alter tables or stored routines.

You should design your transactions not to include such statements. If you issue a statement early in a transaction that cannot be rolled back, and then another statement later fails, the full effect of the transaction cannot be rolled back in such cases by issuing a `ROLLBACK` statement.

### 13.3.3. Statements That Cause an Implicit Commit

The statements listed in this section (and any synonyms for them) implicitly end any transaction active in the current session, as if you had done a `COMMIT` before executing the statement. As of MySQL 5.5.3, most of these statements also cause an implicit commit after executing; for additional details, see the end of this section.

- **Data definition language (DDL) statements that define or modify database objects.** `ALTER DATABASE ... UPGRADE DATA DIRECTORY NAME`, `ALTER EVENT`, `ALTER PROCEDURE`, `ALTER SERVER`, `ALTER TABLE`, `ALTER VIEW`, `CREATE DATABASE`, `CREATE EVENT`, `CREATE INDEX`, `CREATE PROCEDURE`, `CREATE SERVER`, `CREATE TABLE`, `CREATE TRIGGER`, `CREATE VIEW`, `DROP DATABASE`, `DROP EVENT`, `DROP INDEX`, `DROP PROCEDURE`, `DROP SERVER`, `DROP TABLE`, `DROP TRIGGER`, `DROP VIEW`, `RENAME TABLE`, `TRUNCATE TABLE`.

`ALTER FUNCTION`, `CREATE FUNCTION` and `DROP FUNCTION` also cause an implicit commit when used with stored functions, but not with UDFs. (`ALTER FUNCTION` can only be used with stored functions.)

`ALTER TABLE`, `CREATE TABLE`, and `DROP TABLE` do not commit a transaction if the `TEMPORARY` keyword is used. (This does not apply to other operations on temporary tables such as `CREATE INDEX`, which do cause a commit.) However, although no implicit commit occurs, neither can the statement be rolled back. Therefore, use of such statements will violate transaction atomicity: For example, if you use `CREATE TEMPORARY TABLE` and then roll back the transaction, the table remains in existence.

The `CREATE TABLE` statement in InnoDB is processed as a single transaction. This means that a `ROLLBACK` from the user does not undo `CREATE TABLE` statements the user made during that transaction.

`CREATE TABLE ... SELECT` causes an implicit commit before and after the statement is executed when you are creating nontemporary tables. (No commit occurs for `CREATE TEMPORARY TABLE ... SELECT`.) This is to prevent an issue during replication where the table could be created on the master after a rollback, but fail to be recorded in the binary log, and therefore not replicated to the slave. For more information, see Bug #22865.

- **Statements that implicitly use or modify tables in the `mysql` database.** `CREATE USER`, `DROP USER`, `GRANT`, `RENAME USER`, `REVOKE`, `SET PASSWORD`.
- **Transaction-control and locking statements.** `BEGIN`, `LOCK TABLES`, `SET autocommit = 1` (if the value is not already 1), `START TRANSACTION`, `UNLOCK TABLES`.

`UNLOCK TABLES` commits a transaction only if any tables currently have been locked with `LOCK TABLES` to acquire nontransactional table locks. A commit does not occur for `UNLOCK TABLES` following `FLUSH TABLES WITH READ LOCK` because the latter statement does not acquire table-level locks.

Transactions cannot be nested. This is a consequence of the implicit commit performed for any current transaction when you issue a `START TRANSACTION` statement or one of its synonyms.

Statements that cause an implicit commit cannot be used in an XA transaction while the transaction is in an `ACTIVE` state.

The `BEGIN` statement differs from the use of the `BEGIN` keyword that starts a `BEGIN ... END` compound statement. The latter does not cause an implicit commit. See Section 13.6.1, “`BEGIN ... END` Compound-Statement Syntax”.

- **Data loading statements.** `LOAD DATA INFILE`. `LOAD DATA INFILE` causes an implicit commit only for tables using the `NDB` storage engine. For more information, see Bug #11151.
- **Administrative statements.** `ANALYZE TABLE`, `CACHE INDEX`, `CHECK TABLE`, `LOAD INDEX INTO CACHE`, `OPTIMIZE TABLE`, `REPAIR TABLE`.

- **Replication control statements.** Beginning with MySQL 5.6.7: `START SLAVE`, `STOP SLAVE`, `RESET SLAVE`, `CHANGE MASTER TO`. (Bug #13858841)

As of MySQL 5.5.3, most statements that previously caused an implicit commit before executing also do so after executing. The intent is to handle each such statement in its own special transaction because it cannot be rolled back anyway. The following list provides additional details pertaining to this change:

- The `CREATE TABLE` variants (`CREATE TABLE` for InnoDB tables and `CREATE TABLE ... SELECT`) that previously were special cases no longer are so because `CREATE TABLE` uniformly causes an implicit commit before and after executing.
- The `FLUSH` and `RESET` statements cause an implicit commit.
- Transaction-control and locking statements behave as before.

### 13.3.4. SAVEPOINT, ROLLBACK TO SAVEPOINT, and RELEASE SAVEPOINT Syntax

```
SAVEPOINT identifier
ROLLBACK [WORK] TO [SAVEPOINT] identifier
RELEASE SAVEPOINT identifier
```

InnoDB supports the SQL statements `SAVEPOINT`, `ROLLBACK TO SAVEPOINT`, `RELEASE SAVEPOINT` and the optional `WORK` keyword for `ROLLBACK`.

The `SAVEPOINT` statement sets a named transaction savepoint with a name of *identifier*. If the current transaction has a savepoint with the same name, the old savepoint is deleted and a new one is set.

The `ROLLBACK TO SAVEPOINT` statement rolls back a transaction to the named savepoint without terminating the transaction. Modifications that the current transaction made to rows after the savepoint was set are undone in the rollback, but InnoDB does *not* release the row locks that were stored in memory after the savepoint. (For a new inserted row, the lock information is carried by the transaction ID stored in the row; the lock is not separately stored in memory. In this case, the row lock is released in the undo.) Savepoints that were set at a later time than the named savepoint are deleted.

If the `ROLLBACK TO SAVEPOINT` statement returns the following error, it means that no savepoint with the specified name exists:

```
ERROR 1305 (42000): SAVEPOINT identifier does not exist
```

The `RELEASE SAVEPOINT` statement removes the named savepoint from the set of savepoints of the current transaction. No commit or rollback occurs. It is an error if the savepoint does not exist.

All savepoints of the current transaction are deleted if you execute a `COMMIT`, or a `ROLLBACK` that does not name a savepoint.

A new savepoint level is created when a stored function is invoked or a trigger is activated. The savepoints on previous levels become unavailable and thus do not conflict with savepoints on the new level. When the function or trigger terminates, any savepoints it created are released and the previous savepoint level is restored.

### 13.3.5. LOCK TABLES and UNLOCK TABLES Syntax

```
LOCK TABLES
  tbl_name [[AS] alias] lock_type
  [, tbl_name [[AS] alias] lock_type] ...

lock_type:
  READ [LOCAL]
  | [LOW_PRIORITY] WRITE
```

## UNLOCK TABLES

MySQL enables client sessions to acquire table locks explicitly for the purpose of cooperating with other sessions for access to tables, or to prevent other sessions from modifying tables during periods when a session requires exclusive access to them. A session can acquire or release locks only for itself. One session cannot acquire locks for another session or release locks held by another session.

Locks may be used to emulate transactions or to get more speed when updating tables. This is explained in more detail later in this section.

`LOCK TABLES` explicitly acquires table locks for the current client session. Table locks can be acquired for base tables or views. You must have the `LOCK TABLES [755]` privilege, and the `SELECT [755]` privilege for each object to be locked.

For view locking, `LOCK TABLES` adds all base tables used in the view to the set of tables to be locked and locks them automatically. If you lock a table explicitly with `LOCK TABLES`, any tables used in triggers are also locked implicitly, as described in [Section 13.3.5.2, “LOCK TABLES and Triggers”](#).

`UNLOCK TABLES` explicitly releases any table locks held by the current session. `LOCK TABLES` implicitly releases any table locks held by the current session before acquiring new locks.

Another use for `UNLOCK TABLES` is to release the global read lock acquired with the `FLUSH TABLES WITH READ LOCK` statement, which enables you to lock all tables in all databases. See [Section 13.7.6.3, “FLUSH Syntax”](#). (This is a very convenient way to get backups if you have a file system such as Veritas that can take snapshots in time.)

A table lock protects only against inappropriate reads or writes by other sessions. The session holding the lock, even a read lock, can perform table-level operations such as `DROP TABLE`. Truncate operations are not transaction-safe, so an error occurs if the session attempts one during an active transaction or while holding a table lock.

The following discussion applies only to non-`TEMPORARY` tables. `LOCK TABLES` is permitted (but ignored) for a `TEMPORARY` table. The table can be accessed freely by the session within which it was created, regardless of what other locking may be in effect. No lock is necessary because no other session can see the table.

For information about other conditions on the use of `LOCK TABLES` and statements that cannot be used while `LOCK TABLES` is in effect, see [Section 13.3.5.3, “Table-Locking Restrictions and Conditions”](#)

### Rules for Lock Acquisition

To acquire table locks within the current session, use the `LOCK TABLES` statement. The following lock types are available:

`READ [LOCAL]` lock:

- The session that holds the lock can read the table (but not write it).
- Multiple sessions can acquire a `READ` lock for the table at the same time.
- Other sessions can read the table without explicitly acquiring a `READ` lock.
- The `LOCAL` modifier enables nonconflicting `INSERT` statements (concurrent inserts) by other sessions to execute while the lock is held. (See [Section 8.10.3, “Concurrent Inserts”](#).) However, `READ LOCAL` cannot be used if you are going to manipulate the database using processes external to the server while you hold the lock. For `InnoDB` tables, `READ LOCAL` is the same as `READ`.

`[LOW_PRIORITY] WRITE` lock:

- The session that holds the lock can read and write the table.

- Only the session that holds the lock can access the table. No other session can access it until the lock is released.
- Lock requests for the table by other sessions block while the `WRITE` lock is held.
- The `LOW_PRIORITY` modifier has no effect. In previous versions of MySQL, it affected locking behavior, but this is no longer true. As of MySQL 5.6.5, it is deprecated and its use produces a warning. Use `WRITE` without `LOW_PRIORITY` instead.

If the `LOCK TABLES` statement must wait due to locks held by other sessions on any of the tables, it blocks until all locks can be acquired.

A session that requires locks must acquire all the locks that it needs in a single `LOCK TABLES` statement. While the locks thus obtained are held, the session can access only the locked tables. For example, in the following sequence of statements, an error occurs for the attempt to access `t2` because it was not locked in the `LOCK TABLES` statement:

```
mysql> LOCK TABLES t1 READ;
mysql> SELECT COUNT(*) FROM t1;
+-----+
| COUNT(*) |
+-----+
|          3 |
+-----+
mysql> SELECT COUNT(*) FROM t2;
ERROR 1100 (HY000): Table 't2' was not locked with LOCK TABLES
```

Tables in the `INFORMATION_SCHEMA` database are an exception. They can be accessed without being locked explicitly even while a session holds table locks obtained with `LOCK TABLES`.

You cannot refer to a locked table multiple times in a single query using the same name. Use aliases instead, and obtain a separate lock for the table and each alias:

```
mysql> LOCK TABLE t WRITE, t AS t1 READ;
mysql> INSERT INTO t SELECT * FROM t;
ERROR 1100: Table 't' was not locked with LOCK TABLES
mysql> INSERT INTO t SELECT * FROM t AS t1;
```

The error occurs for the first `INSERT` because there are two references to the same name for a locked table. The second `INSERT` succeeds because the references to the table use different names.

If your statements refer to a table by means of an alias, you must lock the table using that same alias. It does not work to lock the table without specifying the alias:

```
mysql> LOCK TABLE t READ;
mysql> SELECT * FROM t AS myalias;
ERROR 1100: Table 'myalias' was not locked with LOCK TABLES
```

Conversely, if you lock a table using an alias, you must refer to it in your statements using that alias:

```
mysql> LOCK TABLE t AS myalias READ;
mysql> SELECT * FROM t;
ERROR 1100: Table 't' was not locked with LOCK TABLES
mysql> SELECT * FROM t AS myalias;
```

`WRITE` locks normally have higher priority than `READ` locks to ensure that updates are processed as soon as possible. This means that if one session obtains a `READ` lock and then another session requests a `WRITE` lock, subsequent `READ` lock requests wait until the session that requested the `WRITE` lock has obtained the lock and released it.

`LOCK TABLES` acquires locks as follows:

1. Sort all tables to be locked in an internally defined order. From the user standpoint, this order is undefined.

2. If a table is to be locked with a read and a write lock, put the write lock request before the read lock request.
3. Lock one table at a time until the session gets all locks.

This policy ensures that table locking is deadlock free.



#### Note

`LOCK TABLES` or `UNLOCK TABLES`, when applied to a partitioned table, always locks or unlocks the entire table; these statements do not support partition lock pruning. See [Section 17.6.4, “Partitioning and Locking”](#).

### Rules for Lock Release

When the table locks held by a session are released, they are all released at the same time. A session can release its locks explicitly, or locks may be released implicitly under certain conditions.

- A session can release its locks explicitly with `UNLOCK TABLES`.
- If a session issues a `LOCK TABLES` statement to acquire a lock while already holding locks, its existing locks are released implicitly before the new locks are granted.
- If a session begins a transaction (for example, with `START TRANSACTION`), an implicit `UNLOCK TABLES` is performed, which causes existing locks to be released. (For additional information about the interaction between table locking and transactions, see [Section 13.3.5.1, “Interaction of Table Locking and Transactions”](#).)

If the connection for a client session terminates, whether normally or abnormally, the server implicitly releases all table locks held by the session (transactional and nontransactional). If the client reconnects, the locks will no longer be in effect. In addition, if the client had an active transaction, the server rolls back the transaction upon disconnect, and if reconnect occurs, the new session begins with autocommit enabled. For this reason, clients may wish to disable auto-reconnect. With auto-reconnect in effect, the client is not notified if reconnect occurs but any table locks or current transaction will have been lost. With auto-reconnect disabled, if the connection drops, an error occurs for the next statement issued. The client can detect the error and take appropriate action such as reacquiring the locks or redoing the transaction. See [Section 21.8.15, “Controlling Automatic Reconnection Behavior”](#).



#### Note

If you use `ALTER TABLE` on a locked table, it may become unlocked. For example, if you attempt a second `ALTER TABLE` operation, the result may be an error `Table 'tbl_name' was not locked with LOCK TABLES`. To handle this, lock the table again prior to the second alteration. See also [Section C.5.7.1, “Problems with ALTER TABLE”](#).

### 13.3.5.1. Interaction of Table Locking and Transactions

`LOCK TABLES` and `UNLOCK TABLES` interact with the use of transactions as follows:

- `LOCK TABLES` is not transaction-safe and implicitly commits any active transaction before attempting to lock the tables.
- `UNLOCK TABLES` implicitly commits any active transaction, but only if `LOCK TABLES` has been used to acquire table locks. For example, in the following set of statements, `UNLOCK TABLES` releases the global read lock but does not commit the transaction because no table locks are in effect:

```
FLUSH TABLES WITH READ LOCK;
START TRANSACTION;
SELECT ... ;
UNLOCK TABLES;
```



- Beginning a transaction (for example, with `START TRANSACTION`) implicitly commits any current transaction and releases existing table locks.
- `FLUSH TABLES WITH READ LOCK` acquires a global read lock and not table locks, so it is not subject to the same behavior as `LOCK TABLES` and `UNLOCK TABLES` with respect to table locking and implicit commits. For example, `START TRANSACTION` does not release the global read lock. See [Section 13.7.6.3, “FLUSH Syntax”](#).
- Other statements that implicitly cause transactions to be committed do not release existing table locks. For a list of such statements, see [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).
- The correct way to use `LOCK TABLES` and `UNLOCK TABLES` with transactional tables, such as `InnoDB` tables, is to begin a transaction with `SET autocommit = 0` (not `START TRANSACTION`) followed by `LOCK TABLES`, and to not call `UNLOCK TABLES` until you commit the transaction explicitly. For example, if you need to write to table `t1` and read from table `t2`, you can do this:

```
SET autocommit=0;
LOCK TABLES t1 WRITE, t2 READ, ...;... do something with tables t1 and t2 here ...
COMMIT;
UNLOCK TABLES;
```

When you call `LOCK TABLES`, `InnoDB` internally takes its own table lock, and MySQL takes its own table lock. `InnoDB` releases its internal table lock at the next commit, but for MySQL to release its table lock, you have to call `UNLOCK TABLES`. You should not have `autocommit = 1` [464], because then `InnoDB` releases its internal table lock immediately after the call of `LOCK TABLES`, and deadlocks can very easily happen. `InnoDB` does not acquire the internal table lock at all if `autocommit = 1` [464], to help old applications avoid unnecessary deadlocks.

- `ROLLBACK` does not release table locks.

### 13.3.5.2. LOCK TABLES and Triggers

If you lock a table explicitly with `LOCK TABLES`, any tables used in triggers are also locked implicitly:

- The locks are taken at the same time as those acquired explicitly with the `LOCK TABLES` statement.
- The lock on a table used in a trigger depends on whether the table is used only for reading. If so, a read lock suffices. Otherwise, a write lock is used.
- If a table is locked explicitly for reading with `LOCK TABLES`, but needs to be locked for writing because it might be modified within a trigger, a write lock is taken rather than a read lock. (That is, an implicit write lock needed due to the table's appearance within a trigger causes an explicit read lock request for the table to be converted to a write lock request.)

Suppose that you lock two tables, `t1` and `t2`, using this statement:

```
LOCK TABLES t1 WRITE, t2 READ;
```

If `t1` or `t2` have any triggers, tables used within the triggers will also be locked. Suppose that `t1` has a trigger defined like this:

```
CREATE TRIGGER t1_a_ins AFTER INSERT ON t1 FOR EACH ROW
BEGIN
  UPDATE t4 SET count = count+1
    WHERE id = NEW.id AND EXISTS (SELECT a FROM t3);
  INSERT INTO t2 VALUES(1, 2);
END;
```

The result of the `LOCK TABLES` statement is that `t1` and `t2` are locked because they appear in the statement, and `t3` and `t4` are locked because they are used within the trigger:

- `t1` is locked for writing per the `WRITE` lock request.
- `t2` is locked for writing, even though the request is for a `READ` lock. This occurs because `t2` is inserted into within the trigger, so the `READ` request is converted to a `WRITE` request.



- `t3` is locked for reading because it is only read from within the trigger.
- `t4` is locked for writing because it might be updated within the trigger.

### 13.3.5.3. Table-Locking Restrictions and Conditions

You can safely use `KILL` to terminate a session that is waiting for a table lock. See [Section 13.7.6.4, “KILL Syntax”](#).

Do *not* lock any tables that you are using with `INSERT DELAYED`. An `INSERT DELAYED` in this case results in an error because the insert must be handled by a separate thread, not by the session which holds the lock.

`LOCK TABLES` and `UNLOCK TABLES` cannot be used within stored programs.

Tables in the `performance_schema` database cannot be locked with `LOCK TABLES`, except the `setup_XXX` tables.

The following statements are prohibited while a `LOCK TABLES` statement is in effect: `CREATE TABLE`, `CREATE TABLE ... LIKE`, `CREATE VIEW`, `DROP VIEW`, and DDL statements on stored functions and procedures and events.

For some operations, system tables in the `mysql` database must be accessed. For example, the `HELP` statement requires the contents of the server-side help tables, and `CONVERT_TZ()` [1212] might need to read the time zone tables. The server implicitly locks the system tables for reading as necessary so that you need not lock them explicitly. These tables are treated as just described:

```
mysql.help_category
mysql.help_keyword
mysql.help_relation
mysql.help_topic
mysql.proc
mysql.time_zone
mysql.time_zone_leap_second
mysql.time_zone_name
mysql.time_zone_transition
mysql.time_zone_transition_type
```

If you want to explicitly place a `WRITE` lock on any of those tables with a `LOCK TABLES` statement, the table must be the only one locked; no other table can be locked with the same statement.

Normally, you do not need to lock tables, because all single `UPDATE` statements are atomic; no other session can interfere with any other currently executing SQL statement. However, there are a few cases when locking tables may provide an advantage:

- If you are going to run many operations on a set of `MyISAM` tables, it is much faster to lock the tables you are going to use. Locking `MyISAM` tables speeds up inserting, updating, or deleting on them because MySQL does not flush the key cache for the locked tables until `UNLOCK TABLES` is called. Normally, the key cache is flushed after each SQL statement.

The downside to locking the tables is that no session can update a `READ`-locked table (including the one holding the lock) and no session can access a `WRITE`-locked table other than the one holding the lock.

- If you are using tables for a nontransactional storage engine, you must use `LOCK TABLES` if you want to ensure that no other session modifies the tables between a `SELECT` and an `UPDATE`. The example shown here requires `LOCK TABLES` to execute safely:

```
LOCK TABLES trans READ, customer WRITE;
SELECT SUM(value) FROM trans WHERE customer_id=some_id;
UPDATE customer
  SET total_value=sum_from_previous_statement
  WHERE customer_id=some_id;
UNLOCK TABLES;
```

Without `LOCK TABLES`, it is possible that another session might insert a new row in the `trans` table between execution of the `SELECT` and `UPDATE` statements.

You can avoid using `LOCK TABLES` in many cases by using relative updates (`UPDATE customer SET value=value+new_value`) or the `LAST_INSERT_ID()` [1269] function. See Section 1.8.5.3, “Transaction and Atomic Operation Differences”.

You can also avoid locking tables in some cases by using the user-level advisory lock functions `GET_LOCK()` [1275] and `RELEASE_LOCK()` [1280]. These locks are saved in a hash table in the server and implemented with `pthread_mutex_lock()` and `pthread_mutex_unlock()` for high speed. See Section 12.16, “Miscellaneous Functions”.

See Section 8.10.1, “Internal Locking Methods”, for more information on locking policy.

### 13.3.6. SET TRANSACTION Syntax

```
SET [GLOBAL | SESSION] TRANSACTION
    transaction_characteristic [, transaction_characteristic] ...

transaction_characteristic:
    ISOLATION LEVEL level
    | READ WRITE
    | READ ONLY

level:
    REPEATABLE READ
    | READ COMMITTED
    | READ UNCOMMITTED
    | SERIALIZABLE
```

This statement specifies `transaction` characteristics. It takes a list of one or more characteristic values separated by commas. These characteristics set the transaction `isolation level` or access mode. The isolation level is used for operations on `InnoDB` tables. The access mode may be specified as of MySQL 5.6.5 and indicates whether transactions operate in read/write or read-only mode.

In addition, `SET TRANSACTION` can include an optional `GLOBAL` or `SESSION` keyword to indicate the scope of the statement.

#### Scope of Transaction Characteristics

You can set transaction characteristics globally, for the current session, or for the next transaction:

- With the `GLOBAL` keyword, the statement applies globally for all subsequent sessions. Existing sessions are unaffected.
- With the `SESSION` keyword, the statement applies to all subsequent transactions performed within the current session.
- Without any `SESSION` or `GLOBAL` keyword, the statement applies to the next (not started) transaction performed within the current session.

A global change to transaction characteristics requires the `SUPER` [756] privilege. Any session is free to change its session characteristics (even in the middle of a transaction), or the characteristics for its next transaction.

`SET TRANSACTION` without `GLOBAL` or `SESSION` is not permitted while there is an active transaction:

```
mysql> START TRANSACTION;
Query OK, 0 rows affected (0.02 sec)

mysql> SET TRANSACTION ISOLATION LEVEL SERIALIZABLE;
ERROR 1568 (25001): Transaction characteristics can't be changed
while a transaction is in progress
```

To set the global default isolation level at server startup, use the `--transaction-isolation=level` [445] option to `mysqld` on the command line or in an option file. Values of `level` for this option use dashes rather than spaces, so the permissible values are `READ-UNCOMMITTED` [1470], `READ-COMMITTED` [1469], `REPEATABLE-READ` [1469], or `SERIALIZABLE` [1470]. For example, to set the default isolation level to `REPEATABLE READ` [1469], use these lines in the `[mysqld]` section of an option file:

```
[mysqld]
transaction-isolation = REPEATABLE-READ
```

It is possible to check or set the global and session transaction isolation levels at runtime by using the `tx_isolation` [560] system variable:

```
SELECT @@GLOBAL.tx_isolation, @@tx_isolation;
SET GLOBAL tx_isolation='REPEATABLE-READ';
SET SESSION tx_isolation='SERIALIZABLE';
```

Similarly, to set the transaction access mode at server startup or at runtime, use the `--transaction-read-only` [445] option or `tx_read_only` [561] system variable. By default, these are `OFF` (the mode is read/write) but can be set to `ON` for a default mode of read only.

Setting the global or session value of `tx_isolation` [560] or `tx_read_only` [561] is equivalent to setting the isolation level or access mode with `SET GLOBAL TRANSACTION` or `SET SESSION TRANSACTION`.

## Details and Usage of Isolation Levels

`InnoDB` supports each of the transaction isolation levels described here using different [locking strategies](#). You can enforce a high degree of consistency with the default `REPEATABLE READ` [1469] level, for operations on crucial data where `ACID` compliance is important. Or you can relax the consistency rules with `READ COMMITTED` [1469] or even `READ UNCOMMITTED` [1470], in situations such as bulk reporting where precise consistency and repeatable results are less important than minimizing the amount of overhead for locking. `SERIALIZABLE` [1470] enforces even stricter rules than `REPEATABLE READ` [1469], and is used mainly in specialized situations, such as with `XA` transactions and for troubleshooting issues with concurrency and [deadlocks](#).

For full information about how these isolation levels work with `InnoDB` transactions, see [Section 14.2.3.1, “The `InnoDB` Transaction Model and Locking”](#). In particular, for additional information about `InnoDB` record-level locks and how it uses them to execute various types of statements, see [Section 14.2.3.5, “`InnoDB` Record, Gap, and Next-Key Locks”](#) and [Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”](#).

The following list describes how MySQL supports the different transaction levels. The list goes from the most commonly used level to the least used.

- `REPEATABLE READ`

This is the default isolation level for `InnoDB`. For [consistent reads](#), there is an important difference from the `READ COMMITTED` [1469] isolation level: All consistent reads within the same transaction read the snapshot established by the first read. This convention means that if you issue several plain (nonlocking) `SELECT` statements within the same transaction, these `SELECT` statements are consistent also with respect to each other. See [Section 14.2.3.3, “Consistent Nonlocking Reads”](#).

For [locking reads](#) (`SELECT` with `FOR UPDATE` or `LOCK IN SHARE MODE`), `UPDATE`, and `DELETE` statements, locking depends on whether the statement uses a unique index with a unique search condition, or a range-type search condition. For a unique index with a unique search condition, `InnoDB` locks only the index record found, not the [gap](#) before it. For other search conditions, `InnoDB` locks the index range scanned, using [gap locks](#) or [next-key locks](#) to block insertions by other sessions into the gaps covered by the range.

- `READ COMMITTED`

A somewhat Oracle-like isolation level with respect to consistent (nonlocking) reads: Each [consistent read](#), even within the same transaction, sets and reads its own fresh [snapshot](#). See [Section 14.2.3.3, “Consistent Nonlocking Reads”](#).

For [locking reads](#) ([SELECT](#) with [FOR UPDATE](#) or [LOCK IN SHARE MODE](#)), [InnoDB](#) locks only index records, not the [gaps](#) before them, and thus permits the free insertion of new records next to locked records. For [UPDATE](#) and [DELETE](#) statements, locking depends on whether the statement uses a unique index with a unique search condition (such as [WHERE id = 100](#)), or a range-type search condition (such as [WHERE id > 100](#)). For a unique index with a unique search condition, [InnoDB](#) locks only the index record found, not the gap before it. For range-type searches, [InnoDB](#) locks the index range scanned, using [gap locks](#) or [next-key locks](#) to block insertions by other sessions into the gaps covered by the range. This is necessary because “phantom rows” must be blocked for MySQL replication and recovery to work.



#### Note

In MySQL 5.6, if the [READ COMMITTED](#) isolation level is used, or the deprecated [innodb\\_locks\\_unsafe\\_for\\_binlog \[1729\]](#) system variable is enabled, there is no [InnoDB](#) gap locking except for foreign-key constraint checking and duplicate-key checking. Also, record locks for nonmatching rows are released after MySQL has evaluated the [WHERE](#) condition.

If you use [READ COMMITTED](#) or enable [innodb\\_locks\\_unsafe\\_for\\_binlog \[1729\]](#), you *must* use row-based binary logging.

- [READ UNCOMMITTED](#)

[SELECT](#) statements are performed in a nonlocking fashion, but a possible earlier version of a row might be used. Thus, using this isolation level, such reads are not consistent. This is also called a [dirty read](#). Otherwise, this isolation level works like [READ COMMITTED \[1469\]](#).

- [SERIALIZABLE](#)

This level is like [REPEATABLE READ \[1469\]](#), but [InnoDB](#) implicitly converts all plain [SELECT](#) statements to [SELECT ... LOCK IN SHARE MODE](#) if autocommit is disabled. If autocommit is enabled, the [SELECT](#) is its own transaction. It therefore is known to be read only and can be serialized if performed as a consistent (nonlocking) read and need not block for other transactions. (To force a plain [SELECT](#) to block if other transactions have modified the selected rows, disable autocommit.)

## Transaction Access Mode

As of MySQL 5.6.5, the transaction access mode may be specified with [SET TRANSACTION](#). By default, a transaction takes place in read/write mode, with both reads and writes permitted to tables used in the transaction. This mode may be specified explicitly using an access mode of [READ WRITE](#).

If the transaction access mode is set to [READ ONLY](#), changes to tables are prohibited. This may enable storage engines to make performance improvements that are possible when writes are not permitted.

It is not permitted to specify both [READ WRITE](#) and [READ ONLY](#) in the same statement.

In read-only mode, it remains possible to change tables created with the [TEMPORARY](#) keyword using DML statements. Changes made with DDL statements are not permitted, just as with permanent tables.

The [READ WRITE](#) and [READ ONLY](#) access modes also may be specified for an individual transaction using the [START TRANSACTION](#) statement.

## 13.3.7. XA Transactions

Support for XA transactions is available for the [InnoDB](#) storage engine. The MySQL XA implementation is based on the X/Open CAE document *Distributed Transaction Processing: The XA Specification*. This document is published by The Open Group and available at <http://www.opengroup.org/public/pubs/catalog/c193.htm>. Limitations of the current XA implementation are described in [Section E.6, “Restrictions on XA Transactions”](#).

On the client side, there are no special requirements. The XA interface to a MySQL server consists of SQL statements that begin with the `XA` keyword. MySQL client programs must be able to send SQL statements and to understand the semantics of the XA statement interface. They do not need be linked against a recent client library. Older client libraries also will work.

Currently, among the MySQL Connectors, MySQL Connector/J 5.0.0 and higher supports XA directly, by means of a class interface that handles the Xan SQL statement interface for you.

XA supports distributed transactions; that is, the ability to permit multiple separate transactional resources to participate in a global transaction. Transactional resources often are RDBMSs but may be other kinds of resources.

A global transaction involves several actions that are transactional in themselves, but that all must either complete successfully as a group, or all be rolled back as a group. In essence, this extends ACID properties “up a level” so that multiple ACID transactions can be executed in concert as components of a global operation that also has ACID properties. (However, for a distributed transaction, you must use the [SERIALIZABLE \[1470\]](#) isolation level to achieve ACID properties. It is enough to use [REPEATABLE READ \[1469\]](#) for a nondistributed transaction, but not for a distributed transaction.)

Some examples of distributed transactions:

- An application may act as an integration tool that combines a messaging service with an RDBMS. The application makes sure that transactions dealing with message sending, retrieval, and processing that also involve a transactional database all happen in a global transaction. You can think of this as “transactional email.”
- An application performs actions that involve different database servers, such as a MySQL server and an Oracle server (or multiple MySQL servers), where actions that involve multiple servers must happen as part of a global transaction, rather than as separate transactions local to each server.
- A bank keeps account information in an RDBMS and distributes and receives money through automated teller machines (ATMs). It is necessary to ensure that ATM actions are correctly reflected in the accounts, but this cannot be done with the RDBMS alone. A global transaction manager integrates the ATM and database resources to ensure overall consistency of financial transactions.

Applications that use global transactions involve one or more Resource Managers and a Transaction Manager:

- A Resource Manager (RM) provides access to transactional resources. A database server is one kind of resource manager. It must be possible to either commit or roll back transactions managed by the RM.
- A Transaction Manager (TM) coordinates the transactions that are part of a global transaction. It communicates with the RMs that handle each of these transactions. The individual transactions within a global transaction are “branches” of the global transaction. Global transactions and their branches are identified by a naming scheme described later.

The MySQL implementation of XA MySQL enables a MySQL server to act as a Resource Manager that handles XA transactions within a global transaction. A client program that connects to the MySQL server acts as the Transaction Manager.

To carry out a global transaction, it is necessary to know which components are involved, and bring each component to a point when it can be committed or rolled back. Depending on what each component reports about its ability to succeed, they must all commit or roll back as an atomic group. That is, either all components must commit, or all components must roll back. To manage a global

transaction, it is necessary to take into account that any component or the connecting network might fail.

The process for executing a global transaction uses two-phase commit (2PC). This takes place after the actions performed by the branches of the global transaction have been executed.

1. In the first phase, all branches are prepared. That is, they are told by the TM to get ready to commit. Typically, this means each RM that manages a branch records the actions for the branch in stable storage. The branches indicate whether they are able to do this, and these results are used for the second phase.
2. In the second phase, the TM tells the RMs whether to commit or roll back. If all branches indicated when they were prepared that they will be able to commit, all branches are told to commit. If any branch indicated when it was prepared that it will not be able to commit, all branches are told to roll back.

In some cases, a global transaction might use one-phase commit (1PC). For example, when a Transaction Manager finds that a global transaction consists of only one transactional resource (that is, a single branch), that resource can be told to prepare and commit at the same time.

### 13.3.7.1. XA Transaction SQL Syntax

To perform XA transactions in MySQL, use the following statements:

```
XA {START|BEGIN} xid [JOIN|RESUME]
XA END xid [SUSPEND [FOR MIGRATE]]
XA PREPARE xid
XA COMMIT xid [ONE PHASE]
XA ROLLBACK xid
XA RECOVER
```

For `XA START`, the `JOIN` and `RESUME` clauses are not supported.

For `XA END` the `SUSPEND [FOR MIGRATE]` clause is not supported.

Each XA statement begins with the `XA` keyword, and most of them require an `xid` value. An `xid` is an XA transaction identifier. It indicates which transaction the statement applies to. `xid` values are supplied by the client, or generated by the MySQL server. An `xid` value has from one to three parts:

```
xid: gtrid [, bqual [, formatID ]]
```

`gtrid` is a global transaction identifier, `bqual` is a branch qualifier, and `formatID` is a number that identifies the format used by the `gtrid` and `bqual` values. As indicated by the syntax, `bqual` and `formatID` are optional. The default `bqual` value is `' '` if not given. The default `formatID` value is 1 if not given.

`gtrid` and `bqual` must be string literals, each up to 64 bytes (not characters) long. `gtrid` and `bqual` can be specified in several ways. You can use a quoted string (`'ab'`), hex string (`0x6162`, `X'ab'`), or bit value (`b'nnnn'`).

`formatID` is an unsigned integer.

The `gtrid` and `bqual` values are interpreted in bytes by the MySQL server's underlying XA support routines. However, while an SQL statement containing an XA statement is being parsed, the server works with some specific character set. To be safe, write `gtrid` and `bqual` as hex strings.

`xid` values typically are generated by the Transaction Manager. Values generated by one TM must be different from values generated by other TMs. A given TM must be able to recognize its own `xid` values in a list of values returned by the `XA RECOVER` statement.



`XA START xid` starts an XA transaction with the given *xid* value. Each XA transaction must have a unique *xid* value, so the value must not currently be used by another XA transaction. Uniqueness is assessed using the *gtrid* and *bqual* values. All following XA statements for the XA transaction must be specified using the same *xid* value as that given in the `XA START` statement. If you use any of those statements but specify an *xid* value that does not correspond to some existing XA transaction, an error occurs.

One or more XA transactions can be part of the same global transaction. All XA transactions within a given global transaction must use the same *gtrid* value in the *xid* value. For this reason, *gtrid* values must be globally unique so that there is no ambiguity about which global transaction a given XA transaction is part of. The *bqual* part of the *xid* value must be different for each XA transaction within a global transaction. (The requirement that *bqual* values be different is a limitation of the current MySQL XA implementation. It is not part of the XA specification.)

The `XA RECOVER` statement returns information for those XA transactions on the MySQL server that are in the `PREPARED` state. (See [Section 13.3.7.2, “XA Transaction States”](#).) The output includes a row for each such XA transaction on the server, regardless of which client started it.

`XA RECOVER` output rows look like this (for an example *xid* value consisting of the parts 'abc', 'def', and 7):

```
mysql> XA RECOVER;
+-----+-----+-----+-----+
| formatID | gtrid_length | bqual_length | data |
+-----+-----+-----+-----+
|          7 |             3 |             3 | abcdef |
+-----+-----+-----+-----+
```

The output columns have the following meanings:

- `formatID` is the *formatID* part of the transaction *xid*
- `gtrid_length` is the length in bytes of the *gtrid* part of the *xid*
- `bqual_length` is the length in bytes of the *bqual* part of the *xid*
- `data` is the concatenation of the *gtrid* and *bqual* parts of the *xid*

### 13.3.7.2. XA Transaction States

An XA transaction progresses through the following states:

1. Use `XA START` to start an XA transaction and put it in the `ACTIVE` state.
2. For an `ACTIVE` XA transaction, issue the SQL statements that make up the transaction, and then issue an `XA END` statement. `XA END` puts the transaction in the `IDLE` state.
3. For an `IDLE` XA transaction, you can issue either an `XA PREPARE` statement or an `XA COMMIT ... ONE PHASE` statement:
  - `XA PREPARE` puts the transaction in the `PREPARED` state. An `XA RECOVER` statement at this point will include the transaction's *xid* value in its output, because `XA RECOVER` lists all XA transactions that are in the `PREPARED` state.
  - `XA COMMIT ... ONE PHASE` prepares and commits the transaction. The *xid* value will not be listed by `XA RECOVER` because the transaction terminates.
4. For a `PREPARED` XA transaction, you can issue an `XA COMMIT` statement to commit and terminate the transaction, or `XA ROLLBACK` to roll back and terminate the transaction.

Here is a simple XA transaction that inserts a row into a table as part of a global transaction:

```
mysql> XA START 'xatest';
```



```

Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO mytable (i) VALUES(10);
Query OK, 1 row affected (0.04 sec)

mysql> XA END 'xatest';
Query OK, 0 rows affected (0.00 sec)

mysql> XA PREPARE 'xatest';
Query OK, 0 rows affected (0.00 sec)

mysql> XA COMMIT 'xatest';
Query OK, 0 rows affected (0.00 sec)

```

Within the context of a given client connection, XA transactions and local (non-XA) transactions are mutually exclusive. For example, if `XA START` has been issued to begin an XA transaction, a local transaction cannot be started until the XA transaction has been committed or rolled back. Conversely, if a local transaction has been started with `START TRANSACTION`, no XA statements can be used until the transaction has been committed or rolled back.

Note that if an XA transaction is in the `ACTIVE` state, you cannot issue any statements that cause an implicit commit. That would violate the XA contract because you could not roll back the XA transaction. You will receive the following error if you try to execute such a statement:

```

ERROR 1399 (XAE07): XAER_RMFAIL: The command cannot be executed
when global transaction is in the ACTIVE state

```

Statements to which the preceding remark applies are listed at [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

## 13.4. Replication Statements

Replication can be controlled through the SQL interface using the statements described in this section. One group of statements controls master servers, the other controls slave servers.

### 13.4.1. SQL Statements for Controlling Master Servers

This section discusses statements for managing master replication servers. [Section 13.4.2, “SQL Statements for Controlling Slave Servers”](#), discusses statements for managing slave servers.

In addition to the statements described here, the following `SHOW` statements are used with master servers in replication. For information about these statements, see [Section 13.7.5, “SHOW Syntax”](#).

- `SHOW BINARY LOGS`
- `SHOW BINLOG EVENTS`
- `SHOW MASTER STATUS`
- `SHOW SLAVE HOSTS`

#### 13.4.1.1. `PURGE BINARY LOGS` Syntax

```

PURGE { BINARY | MASTER } LOGS
      { TO 'log_name' | BEFORE datetime_expr }

```

The binary log is a set of files that contain information about data modifications made by the MySQL server. The log consists of a set of binary log files, plus an index file (see [Section 5.2.4, “The Binary Log”](#)).

The `PURGE BINARY LOGS` statement deletes all the binary log files listed in the log index file prior to the specified log file name or date. `BINARY` and `MASTER` are synonyms. Deleted log files also are removed from the list recorded in the index file, so that the given log file becomes the first in the list.

This statement has no effect if the server was not started with the `--log-bin [2002]` option to enable binary logging.

Examples:

```
PURGE BINARY LOGS TO 'mysql-bin.010';
PURGE BINARY LOGS BEFORE '2008-04-02 22:46:26';
```

The `BEFORE` variant's `datetime_expr` argument should evaluate to a `DATETIME` value (a value in `'YYYY-MM-DD hh:mm:ss'` format).

This statement is safe to run while slaves are replicating. You need not stop them. If you have an active slave that currently is reading one of the log files you are trying to delete, this statement does nothing and fails with an error. However, if a slave is not connected and you happen to purge one of the log files it has yet to read, the slave will be unable to replicate after it reconnects.

To safely purge binary log files, follow this procedure:

1. On each slave server, use `SHOW SLAVE STATUS` to check which log file it is reading.
2. Obtain a listing of the binary log files on the master server with `SHOW BINARY LOGS`.
3. Determine the earliest log file among all the slaves. This is the target file. If all the slaves are up to date, this is the last log file on the list.
4. Make a backup of all the log files you are about to delete. (This step is optional, but always advisable.)
5. Purge all log files up to but not including the target file.

You can also set the `expire_logs_days [480]` system variable to expire binary log files automatically after a given number of days (see [Section 5.1.4, “Server System Variables”](#)). If you are using replication, you should set the variable no lower than the maximum number of days your slaves might lag behind the master.

`PURGE BINARY LOGS TO` and `PURGE BINARY LOGS BEFORE` both fail with an error when binary log files listed in the `.index` file had been removed from the system by some other means (such as using `rm` on Linux). (Bug #18199, Bug #18453) To handle such errors, edit the `.index` file (which is a simple text file) manually to ensure that it lists only the binary log files that are actually present, then run again the `PURGE BINARY LOGS` statement that failed.

### 13.4.1.2. RESET MASTER Syntax

```
RESET MASTER
```

Deletes all binary log files listed in the index file, resets the binary log index file to be empty, and creates a new binary log file.

In MySQL 5.6.5 and later, `RESET MASTER` also clears the values of the `gtid_purged [2023]` system variable (known as `gtid_lost [2022]` in MySQL 5.6.8 and earlier) as well as the global value of the `gtid_executed [2021]` (`gtid_done [2020]`, prior to MySQL 5.6.9) system variable (but not its session value); that is, executing this statement sets each of these values to an empty string ('').

This statement is intended to be used only when the master is started for the first time.



#### Important

The effects of `RESET MASTER` differ from those of `PURGE BINARY LOGS` in 2 key ways:

1. `RESET MASTER` removes *all* binary log files that are listed in the index file, leaving only a single, empty binary log file with a numeric suffix of `.000001`, whereas the numbering is not reset by `PURGE BINARY LOGS`.

2. `RESET MASTER` is *not* intended to be used while any replication slaves are running. The behavior of `RESET MASTER` when used while slaves are running is undefined (and thus unsupported), whereas `PURGE BINARY LOGS` may be safely used while replication slaves are running.

See also [Section 13.4.1.1](#), “`PURGE BINARY LOGS Syntax`”.

`RESET MASTER` can prove useful when you first set up the master and the slave, so that you can verify the setup as follows:

1. Start the master and slave, and start replication (see [Section 16.1.1](#), “`How to Set Up Replication`”).
2. Execute a few test queries on the master.
3. Check that the queries were replicated to the slave.
4. When replication is running correctly, issue `STOP SLAVE` followed by `RESET SLAVE` on the slave, then verify that any unwanted data no longer exists on the slave.
5. Issue `RESET MASTER` on the master to clean up the test queries.

After verifying the setup and getting rid of any unwanted and log files generated by testing, you can start the slave and begin replicating.

### 13.4.1.3. `SET sql_log_bin Syntax`

```
SET sql_log_bin = {0|1}
```

The `sql_log_bin` [544] variable controls whether logging to the binary log is done. The default value is 1 (do logging). To change logging for the current session, change the session value of this variable. The session user must have the `SUPER` [756] privilege to set this variable.

In MySQL 5.6, it is not possible to set `@session.sql_log_bin` within a transaction or subquery. (Bug #53437)

## 13.4.2. SQL Statements for Controlling Slave Servers

This section discusses statements for managing slave replication servers. [Section 13.4.1](#), “`SQL Statements for Controlling Master Servers`”, discusses statements for managing master servers.

In addition to the statements described here, `SHOW SLAVE STATUS` and `SHOW RELAYLOG EVENTS` are also used with replication slaves. For information about these statements, see [Section 13.7.5.35](#), “`SHOW SLAVE STATUS Syntax`”, and [Section 13.7.5.33](#), “`SHOW RELAYLOG EVENTS Syntax`”.

### 13.4.2.1. `CHANGE MASTER TO Syntax`

```
CHANGE MASTER TO option [, option] ...
```

*option*:

```
MASTER_BIND = 'interface_name'
| MASTER_HOST = 'host_name'
| MASTER_USER = 'user_name'
| MASTER_PASSWORD = 'password'
| MASTER_PORT = port_num
| MASTER_CONNECT_RETRY = interval
| MASTER_RETRY_COUNT = count
| MASTER_DELAY = interval
| MASTER_HEARTBEAT_PERIOD = interval
| MASTER_LOG_FILE = 'master_log_name'
| MASTER_LOG_POS = master_log_pos
| MASTER_AUTO_POSITION = {0|1}
| RELAY_LOG_FILE = 'relay_log_name'
| RELAY_LOG_POS = relay_log_pos
| MASTER_SSL = {0|1}
| MASTER_SSL_CA = 'ca_file_name'
```

```

MASTER_SSL_CAPATH = 'ca_directory_name'
MASTER_SSL_CERT = 'cert_file_name'
MASTER_SSL_CRL = 'crl_file_name'
MASTER_SSL_CRLPATH = 'crl_directory_name'
MASTER_SSL_KEY = 'key_file_name'
MASTER_SSL_CIPHER = 'cipher_list'
MASTER_SSL_VERIFY_SERVER_CERT = {0|1}
IGNORE_SERVER_IDS = (server_id_list)

server_id_list:
[server_id [, server_id] ... ]

```

**CHANGE MASTER TO** changes the parameters that the slave server uses for connecting to the master server, for reading the master binary log, and reading the slave relay log. It also updates the contents of the `master.info` and `relay-log.info` files. To use **CHANGE MASTER TO**, the slave replication threads must be stopped (use **STOP SLAVE** if necessary). In MySQL 5.6.11 and later, `gtid_next` [2022] must also be set to **AUTOMATIC** (BUG#16062608).

Options not specified retain their value, except as indicated in the following discussion. Thus, in most cases, there is no need to specify options that do not change. For example, if the password to connect to your MySQL master has changed, you just need to issue these statements to tell the slave about the new password:

```

STOP SLAVE; -- if replication was running
CHANGE MASTER TO MASTER_PASSWORD='new3cret';
START SLAVE; -- if you want to restart replication

```

**MASTER\_HOST**, **MASTER\_USER**, **MASTER\_PASSWORD**, and **MASTER\_PORT** provide information to the slave about how to connect to its master:

- **MASTER\_HOST** and **MASTER\_PORT** are the host name (or IP address) of the master host and its TCP/IP port.



#### Note

Replication cannot use Unix socket files. You must be able to connect to the master MySQL server using TCP/IP.

If you specify the **MASTER\_HOST** or **MASTER\_PORT** option, the slave assumes that the master server is different from before (even if the option value is the same as its current value.) In this case, the old values for the master binary log file name and position are considered no longer applicable, so if you do not specify **MASTER\_LOG\_FILE** and **MASTER\_LOG\_POS** in the statement, **MASTER\_LOG\_FILE=''** and **MASTER\_LOG\_POS=4** are silently appended to it.

Setting **MASTER\_HOST=''** (that is, setting its value explicitly to an empty string) is *not* the same as not setting **MASTER\_HOST** at all. Beginning with MySQL 5.5, trying to set **MASTER\_HOST** to an empty string fails with an error. Previously, setting **MASTER\_HOST** to an empty string caused **START SLAVE** subsequently to fail. (Bug #28796)

In MySQL 5.6.5 and later, values used for **MASTER\_HOST** and other **CHANGE MASTER TO** options are checked for linefeed (`\n` or `0x0A`) characters; the presence of such characters in these values causes the statement to fail with **ER\_MASTER\_INFO**. (Bug #11758581, Bug #50801)

- **MASTER\_USER** and **MASTER\_PASSWORD** are the user name and password of the account to use for connecting to the master.

In MySQL 5.6.4 and later, **MASTER\_USER** cannot be made empty; setting **MASTER\_USER = ''** or leaving it unset when setting a value for **MASTER\_PASSWORD** causes an error (Bug #13427949).

Currently, a password used for a replication slave account is effectively limited to 32 characters in length; the password can be longer, but any excess characters are truncated. This is not due to any limit imposed by the MySQL Server generally, but rather is an issue specific to MySQL Replication. (For more information, see Bug #43439.)

The text of a running `CHANGE MASTER TO` statement, including values for `MASTER_USER` and `MASTER_PASSWORD`, can be seen in the output of a concurrent `SHOW PROCESSLIST` statement. (The complete text of a `START SLAVE` statement is also visible to `SHOW PROCESSLIST`.)

The `MASTER_SSL_XXX` options provide information about using SSL for the connection. They correspond to the `--ssl-xxx` options described in [Section 6.3.9.4, “SSL Command Options”](#), and [Section 16.3.7, “Setting Up Replication Using SSL”](#). These options can be changed even on slaves that are compiled without SSL support. They are saved to the `master.info` file, but are ignored if the slave does not have SSL support enabled. `MASTER_SSL_CRL` and `MASTER_SSL_CRLPATH` were added in MySQL 5.6.3.

`MASTER_CONNECT_RETRY` specifies how many seconds to wait between connect retries. The default is 60.

`MASTER_RETRY_COUNT`, added in MySQL 5.6.1, limits the *number* of reconnection attempts and updates the value of the `Master_Retry_Count` column in the output of `SHOW SLAVE STATUS` (also added in MySQL 5.6.1). The default value is  $24 * 3600 = 86400$ . `MASTER_RETRY_COUNT` is intended to replace the older `--master-retry-count [1973]` server option, and is now the preferred method for setting this limit. You are encouraged not to rely on `--master-retry-count [1973]` in new applications and, when upgrading to MySQL 5.6.1 or later from earlier versions of MySQL, to update any existing applications that rely on it, so that they use `CHANGE MASTER TO ... MASTER_RETRY_COUNT` instead.

`MASTER_DELAY` specifies how many seconds behind the master the slave must lag. An event received from the master is not executed until at least *interval* seconds later than its execution on the master. The default is 0. An error occurs if *interval* is not a nonnegative integer in the range from 0 to  $2^{31} - 1$ . For more information, see [Section 16.3.9, “Delayed Replication”](#). This option was added in MySQL 5.6.0.

`MASTER_BIND` is for use on replication slaves having multiple network interfaces, and determines which of the slave's network interfaces is chosen for connecting to the master.

The address configured with this option, if any, can be seen in the `Master_Bind` column of the output from `SHOW SLAVE STATUS`. If you are using slave status log tables (server started with `--master-info-repository=TABLE [2000]`), the value can also be seen as the `Master_bind` column of the `mysql.slave_master_info` table.

The ability to bind a replication slave to a specific network interface was added in MySQL 5.6.2.

`MASTER_HEARTBEAT_PERIOD` sets the interval in seconds between replication heartbeats. Whenever the master's binary log is updated with an event, the waiting period for the next heartbeat is reset. *interval* is a decimal value having the range 0 to 4294967 seconds and a resolution in milliseconds; the smallest nonzero value is 0.001. Heartbeats are sent by the master only if there are no unsent events in the binary log file for a period longer than *interval*.

If you are logging master connection information to tables, `MASTER_HEARTBEAT_PERIOD` can be seen as the value of the `Heartbeat` column of the `mysql.slave_master_info` table.

Setting *interval* to 0 disables heartbeats altogether. The default value for *interval* is equal to the value of `slave_net_timeout [1993]` divided by 2.

Setting `@@global.slave_net_timeout` to a value less than that of the current heartbeat interval results in a warning being issued. The effect of issuing `RESET SLAVE` on the heartbeat interval is to reset it to the default value.

`MASTER_LOG_FILE` and `MASTER_LOG_POS` are the coordinates at which the slave I/O thread should begin reading from the master the next time the thread starts. `RELAY_LOG_FILE` and `RELAY_LOG_POS` are the coordinates at which the slave SQL thread should begin reading from the relay log the next time the thread starts. If you specify either of `MASTER_LOG_FILE` or

`MASTER_LOG_POS`, you cannot specify `RELAY_LOG_FILE` or `RELAY_LOG_POS`. In MySQL 5.6.5 and later, if you specify either of `MASTER_LOG_FILE` or `MASTER_LOG_POS`, you also cannot specify `MASTER_AUTO_POSITION = 1` (described later in this section). If neither of `MASTER_LOG_FILE` or `MASTER_LOG_POS` is specified, the slave uses the last coordinates of the *slave SQL thread* before `CHANGE MASTER TO` was issued. This ensures that there is no discontinuity in replication, even if the slave SQL thread was late compared to the slave I/O thread, when you merely want to change, say, the password to use.

`MASTER_AUTO_POSITION` was added in MySQL 5.6.5. If `MASTER_AUTO_POSITION = 1` is used with `CHANGE MASTER TO`, the slave attempts to connect to the master using the GTID-based replication protocol. In this case, the coordinates represented by `MASTER_LOG_FILE` and `MASTER_LOG_POS` are not used, and global transaction identifiers are used instead. Thus the use of either or both of these options together with `MASTER_AUTO_POSITION` causes an error.

Beginning with MySQL 5.6.10, you can see whether replication is running with autopositioning enabled by checking the output of `SHOW SLAVE STATUS`. (Bug #15992220)

`gtid_mode` [2022] must also be enabled before issuing `CHANGE MASTER TO ... MASTER_AUTO_POSITION = 1`. Otherwise, the statement fails with an error.

To revert to the older file-based replication protocol after using GTIDs, you can issue a new `CHANGE MASTER TO` statement that specifies `MASTER_AUTO_POSITION = 0`, as well as at least one of `MASTER_LOG_FILE` or `MASTER_LOG_POSITION`.

`CHANGE MASTER TO` deletes all relay log files and starts a new one, unless you specify `RELAY_LOG_FILE` or `RELAY_LOG_POS`. In that case, relay log files are kept; the `relay_log_purge` [534] global variable is set silently to 0.

Prior to MySQL 5.6.2, `RELAY_LOG_FILE` required an absolute path. Beginning with MySQL 5.6.2, the path can be relative, in which case it is assumed to be relative to the slave's data directory. (Bug #12190)

`IGNORE_SERVER_IDS` takes a comma-separated list of 0 or more server IDs. Events originating from the corresponding servers are ignored, with the exception of log rotation and deletion events, which are still recorded in the relay log.

In circular replication, the originating server normally acts as the terminator of its own events, so that they are not applied more than once. Thus, this option is useful in circular replication when one of the servers in the circle is removed. Suppose that you have a circular replication setup with 4 servers, having server IDs 1, 2, 3, and 4, and server 3 fails. When bridging the gap by starting replication from server 2 to server 4, you can include `IGNORE_SERVER_IDS = (3)` in the `CHANGE MASTER TO` statement that you issue on server 4 to tell it to use server 2 as its master instead of server 3. Doing so causes it to ignore and not to propagate any statements that originated with the server that is no longer in use.

If a `CHANGE MASTER TO` statement is issued without any `IGNORE_SERVER_IDS` option, any existing list is preserved; `RESET SLAVE` also has no effect on the server ID list. To clear the list of ignored servers, it is necessary to use the option with an empty list:

```
CHANGE MASTER TO IGNORE_SERVER_IDS = ( );
```

If `IGNORE_SERVER_IDS` contains the server's own ID and the server was started with the `--replicate-same-server-id` [1979] option enabled, an error results.

In MySQL 5.6, the `master.info` file and the output of `SHOW SLAVE STATUS` provide the list of servers that are currently ignored. For more information, see [Section 16.2.2.2, "Slave Status Logs"](#), and [Section 13.7.5.35, "SHOW SLAVE STATUS Syntax"](#).

In MySQL 5.6, invoking `CHANGE MASTER TO` causes the previous values for `MASTER_HOST`, `MASTER_PORT`, `MASTER_LOG_FILE`, and `MASTER_LOG_POS` to be written to the error log, along with other information about the slave's state prior to execution.



In MySQL 5.6.7 and later, `CHANGE MASTER TO` causes an implicit commit of an ongoing transaction. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

`CHANGE MASTER TO` is useful for setting up a slave when you have the snapshot of the master and have recorded the master binary log coordinates corresponding to the time of the snapshot. After loading the snapshot into the slave to synchronize it to the slave, you can run `CHANGE MASTER TO MASTER_LOG_FILE='log_name', MASTER_LOG_POS=log_pos` on the slave to specify the coordinates at which the slave should begin reading the master binary log.

The following example changes the master server the slave uses and establishes the master binary log coordinates from which the slave begins reading. This is used when you want to set up the slave to replicate the master:

```
CHANGE MASTER TO
  MASTER_HOST='master2.mycompany.com',
  MASTER_USER='replication',
  MASTER_PASSWORD='big3cret',
  MASTER_PORT=3306,
  MASTER_LOG_FILE='master2-bin.001',
  MASTER_LOG_POS=4,
  MASTER_CONNECT_RETRY=10;
```

The next example shows an operation that is less frequently employed. It is used when the slave has relay log files that you want it to execute again for some reason. To do this, the master need not be reachable. You need only use `CHANGE MASTER TO` and start the SQL thread (`START SLAVE SQL_THREAD`):

```
CHANGE MASTER TO
  RELAY_LOG_FILE='slave-relay-bin.006',
  RELAY_LOG_POS=4025;
```

You can even use the second operation in a nonreplication setup with a standalone, nonslave server for recovery following a crash. Suppose that your server has crashed and you have restored it from a backup. You want to replay the server's own binary log files (not relay log files, but regular binary log files), named (for example) `myhost-bin.*`. First, make a backup copy of these binary log files in some safe place, in case you don't exactly follow the procedure below and accidentally have the server purge the binary log. Use `SET GLOBAL relay_log_purge=0` for additional safety. Then start the server without the `--log-bin [2002]` option. Instead, use the `--replicate-same-server-id [1979]`, `--relay-log=myhost-bin [1974]` (to make the server believe that these regular binary log files are relay log files) and `--skip-slave-start [1984]` options. After the server starts, issue these statements:

```
CHANGE MASTER TO
  RELAY_LOG_FILE='myhost-bin.153',
  RELAY_LOG_POS=410,
  MASTER_HOST='some_dummy_string';
START SLAVE SQL_THREAD;
```

The server reads and executes its own binary log files, thus achieving crash recovery. Once the recovery is finished, run `STOP SLAVE`, shut down the server, delete the `master.info` and `relay-log.info` files, and restart the server with its original options.

Specifying the `MASTER_HOST` option (even with a dummy value) is required to make the server think it is a slave.

The following table shows the maximum permissible length for the string-valued options.

Option	Maximum Length
<code>MASTER_HOST</code>	60
<code>MASTER_USER</code>	16
<code>MASTER_PASSWORD</code>	32
<code>MASTER_LOG_FILE</code>	255



Option	Maximum Length
<code>RELAY_LOG_FILE</code>	255
<code>MASTER_SSL_CA</code>	255
<code>MASTER_SSL_CAPATH</code>	255
<code>MASTER_SSL_CERT</code>	255
<code>MASTER_SSL_CRL</code>	255
<code>MASTER_SSL_CRLPATH</code>	255
<code>MASTER_SSL_KEY</code>	255
<code>MASTER_SSL_CIPHER</code>	511

### 13.4.2.2. `MASTER_POS_WAIT()` Syntax

```
SELECT MASTER_POS_WAIT('master_log_file', master_log_pos [, timeout])
```

This is actually a function, not a statement. It is used to ensure that the slave has read and executed events up to a given position in the master's binary log. See [Section 12.16, “Miscellaneous Functions”](#), for a full description.

### 13.4.2.3. `RESET SLAVE` Syntax

```
RESET SLAVE [ALL]
```

`RESET SLAVE` makes the slave forget its replication position in the master's binary log. This statement is meant to be used for a clean start: It deletes the `master.info` and `relay-log.info` files, all the relay log files, and starts a new relay log file. It also resets to 0 the replication delay specified with the `MASTER_DELAY` option to `CHANGE MASTER TO`. To use `RESET SLAVE`, the slave replication threads must be stopped (use `STOP SLAVE` if necessary).



#### Note

All relay log files are deleted, even if they have not been completely executed by the slave SQL thread. (This is a condition likely to exist on a replication slave if you have issued a `STOP SLAVE` statement or if the slave is highly loaded.)

In MySQL 5.6 (unlike the case in MySQL 5.1 and earlier), `RESET SLAVE` does not change any replication connection parameters such as master host, master port, master user, or master password, which are retained in memory. This means that `START SLAVE` can be issued without requiring a `CHANGE MASTER TO` statement following `RESET SLAVE`.

Connection parameters are reset if the slave `mysqld` is shut down following `RESET SLAVE`. In MySQL 5.6.3 and later, you can instead use `RESET SLAVE ALL` to reset these connection parameters (Bug #11809016).

In MySQL 5.6.7 and later, `RESET SLAVE` causes an implicit commit of an ongoing transaction. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

If the slave SQL thread was in the middle of replicating temporary tables when it was stopped, and `RESET SLAVE` is issued, these replicated temporary tables are deleted on the slave.

### 13.4.2.4. `SET GLOBAL sql_slave_skip_counter` Syntax

```
SET GLOBAL sql_slave_skip_counter = N
```

This statement skips the next `N` events from the master. This is useful for recovering from replication stops caused by a statement.

This statement is valid only when the slave threads are not running. Otherwise, it produces an error.

When using this statement, it is important to understand that the binary log is actually organized as a sequence of groups known as *event groups*. Each event group consists of a sequence of events.

- For transactional tables, an event group corresponds to a transaction.
- For nontransactional tables, an event group corresponds to a single SQL statement.



#### Note

A single transaction can contain changes to both transactional and nontransactional tables.

When you use `SET GLOBAL sql_slave_skip_counter` to skip events and the result is in the middle of a group, the slave continues to skip events until it reaches the end of the group. Execution then starts with the next event group.

In MySQL 5.6, issuing this statement causes the previous values of `RELAY_LOG_FILE`, `RELAY_LOG_POS`, and `sql_slave_skip_counter` [1996] to be written to the error log.

### 13.4.2.5. START SLAVE Syntax

```
START SLAVE [thread_types] [until_option] [connection_options]

thread_types:
  [thread_type [, thread_type] ... ]

thread_type:
  IO_THREAD | SQL_THREAD

until_option:
  UNTIL { {SQL_BEFORE_GTIDS | SQL_AFTER_GTIDS} = gtid_set
          | MASTER_LOG_FILE = 'log_name', MASTER_LOG_POS = log_pos
          | RELAY_LOG_FILE = 'log_name', RELAY_LOG_POS = log_pos
          | SQL_AFTER_MTS_GAPS }

connection_options:
  [USER='user_name'] [PASSWORD='user_pass'] [DEFAULT_AUTH='plugin_name'] [PLUGIN_DIR='plugin_dir']

gtid_set:
  uuid_set [, uuid_set] ...
  | ''

uuid_set:
  uuid:interval[:interval]...

uuid:
  hhhhhhhh-hhhh-hhhh-hhhh-hhhhhhhhhhhh

h:
  [0-9,A-F]

interval:
  n[-n]

  (n >= 1)
```

`START SLAVE` with no *thread\_type* options starts both of the slave threads. The I/O thread reads events from the master server and stores them in the relay log. The SQL thread reads events from the relay log and executes them. `START SLAVE` requires the `SUPER` [756] privilege.

If `START SLAVE` succeeds in starting the slave threads, it returns without any error. However, even in that case, it might be that the slave threads start and then later stop (for example, because they do not manage to connect to the master or read its binary log, or some other problem). `START SLAVE` does not warn you about this. You must check the slave's error log for error messages generated by the slave threads, or check that they are running satisfactorily with `SHOW SLAVE STATUS`.

In MySQL 5.6.7 and later, `START SLAVE` causes an implicit commit of an ongoing transaction. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

MySQL 5.6.4 and later supports pluggable user-password authentication with `START SLAVE` with the `USER`, `PASSWORD`, `DEFAULT_AUTH` and `PLUGIN_DIR` options, as described in the following list:

- `USER`: User name. Cannot be set to an empty or null string, or left unset if `PASSWORD` is used.
- `PASSWORD`: Password.
- `DEFAULT_AUTH`: Name of plugin; default is MySQL native authentication.
- `PLUGIN_DIR`: Location of plugin.

Starting with MySQL 5.6.4, you cannot use the `SQL_THREAD` option when specifying `USER`, `PASSWORD`, or both. (Bug #13083642).

See [Section 6.3.7, “Pluggable Authentication”](#), for more information.

If an insecure connection is used with any these options, the server issues the warning `Sending passwords in plain text without SSL/TLS is extremely insecure`.

Starting with MySQL 5.6.6, `START SLAVE ... UNTIL` supports two additional options for use with global transaction identifiers (GTIDs) (see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#)). Each of these takes a set of one or more global transaction identifiers `gtid_set` as an argument (see [GTID sets](#), for more information).

When no `thread_type` is specified, `START SLAVE UNTIL SQL_BEFORE_GTIDS` causes both the slave SQL thread to process and the slave I/O thread to fetch transactions until both of them have reached the *first* transaction whose GTID is listed in the `gtid_set`. `START SLAVE UNTIL SQL_AFTER_GTIDS` causes the slave threads to process all transactions until the *last* transaction in the `gtid_set` has been processed by both threads. In other words, `START SLAVE UNTIL SQL_BEFORE_GTIDS` causes the slave SQL to process and I/O threads to fetch all transactions occurring before the first GTID in the `gtid_set` is reached, and `START SLAVE UNTIL SQL_AFTER_GTIDS` causes the slave threads to handle all transactions, including those whose GTIDs are found in `gtid_set`, until each has encountered a transaction whose GTID is not part of the set. `SQL_BEFORE_GTIDS` and `SQL_AFTER_GTIDS` each support the `SQL_THREAD` and `IO_THREAD` options.

For example, `START SLAVE SQL_THREAD UNTIL SQL_BEFORE_GTIDS = 3E11FA47-71CA-11E1-9E33-C80AA9429562:11-56` causes the slave SQL thread to process all transactions originating from the master whose `server_uuid` [1955] is `3E11FA47-71CA-11E1-9E33-C80AA9429562` until it encounters the transaction having sequence number 11; it then stops without processing this transaction. In other words, all transactions up to and including the transaction with sequence number 10 are processed. Executing `START SLAVE IO_THREAD UNTIL SQL_AFTER_GTIDS = 3E11FA47-71CA-11E1-9E33-C80AA9429562:11-56`, on the other hand, would cause the slave I/O thread to obtain all transactions just mentioned from the master, including all of the transactions having the sequence numbers 11 through 56, and then to stop without processing any additional transactions; that is, the transaction having sequence number 56 would be the last transaction fetched by the slave I/O thread.

With neither the `SQL_THREAD` option nor the `IO_THREAD` option, the previous statement would cause the slave SQL thread to execute all transactions originating from this master, including all transactions with sequence numbers 11 through 56, and then to stop without processing any additional transactions. The same command would also cause the slave I/O thread to start. When the SQL thread reaches the condition, it is stopped. In other words, `START SLAVE UNTIL SQL_BEFORE_GTIDS` has the same effect as `START SLAVE SQL_THREAD, IO_THREAD UNTIL SQL_BEFORE_GTIDS`; the slave SQL

thread and slave I/O thread are each started, and the SQL thread continues executing transactions until the stop condition for that thread is met. (Similarly, `START SLAVE UNTIL SQL_AFTER_GTIDS` is effectively the same as `START SLAVE SQL_THREAD, IO_THREAD UNTIL SQL_AFTER_GTIDS`.)



#### Note

The `SQL_BEFORE_GTIDS` and `SQL_AFTER_GTIDS` keywords are present in the MySQL 5.6.5 server; however, neither of them functioned correctly as options with `START SLAVE [SQL_THREAD | IO_THREAD] UNTIL` in that version, and are therefore supported beginning only with MySQL 5.6.6. (Bug#13810456)

`START SLAVE UNTIL SQL_AFTER_MTS_GAPS` is available in MySQL 5.6.6 or later. This statement causes a multi-threaded slave's SQL threads to run until no more gaps are found in the relay log, and then to stop. This statement can take an `SQL_THREAD` option, but the effects of the statement remain unchanged. It has no effect on the slave I/O thread (and cannot be used with the `IO_THREAD` option). `START SLAVE UNTIL SQL_AFTER_MTS_GAPS` should be used before switching the slave from multi-threaded mode to single-threaded mode (that is, when resetting `slave_parallel_workers` [1993] back to 0 from a positive, nonzero value) after slave has failed with errors in multi-threaded mode.

To change a failed multi-threaded slave to single-threaded mode, you can issue the following series of statements, in the order shown:

```
START SLAVE UNTIL SQL_AFTER_MTS_GAPS;

SET @@GLOBAL.slave_parallel_workers = 0;

START SLAVE SQL_THREAD;
```

If you were running the failed multi-threaded slave with `relay_log_recovery` [1990] enabled, then you must issue `START SLAVE UNTIL SQL_AFTER_MTS_GAPS` prior to executing `CHANGE MASTER TO`. Otherwise the latter statement fails.



#### Note

It is possible to view the entire text of a running `START SLAVE ...` statement, including any `USER` or `PASSWORD` values used, in the output of `SHOW PROCESSLIST`. This is also true for the text of a running `CHANGE MASTER TO` statement, including any values it employs for `MASTER_USER` or `MASTER_PASSWORD`.

`START SLAVE` sends an acknowledgment to the user after both the I/O thread and the SQL thread have started. However, the I/O thread may not yet have connected. For this reason, a successful `START SLAVE` causes `SHOW SLAVE STATUS` to show `Slave_SQL_Running=Yes`, but this does not guarantee that `Slave_IO_Running=Yes` (because `Slave_IO_Running=Yes` only if the I/O thread is running *and connected*). For more information, see [Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”](#), and [Section 16.1.5.1, “Checking Replication Status”](#).

You can add `IO_THREAD` and `SQL_THREAD` options to the statement to name which of the threads to start. In MySQL 5.6.4 and later, the `SQL_THREAD` option is disallowed when specifying `USER`, `PASSWORD`, or both (Bug #13083642).

An `UNTIL` clause (*until\_option*, in the preceding grammar) may be added to specify that the slave should start and run until the SQL thread reaches a given point in the master binary log or in the slave relay log. When the SQL thread reaches that point, it stops. If the `SQL_THREAD` option is specified in the statement, it starts only the SQL thread. Otherwise, it starts both slave threads. If the SQL thread is running, the `UNTIL` clause is ignored and a warning is issued. You cannot use an `UNTIL` clause with the `IO_THREAD` option.

In MySQL 5.6.6 and later, it is also possible with `START SLAVE UNTIL` to specify a stop point relative to a given GTID or set of GTIDs using one of the options `SQL_BEFORE_GTIDS` or `SQL_AFTER_GTIDS`, as explained previously in this section. When using one of these options, you can

specify `SQL_THREAD`, `IO_THREAD`, both of these, or neither of them. If you specify only `SQL_THREAD`, then only the slave SQL thread is affected by the statement; if only `IO_THREAD` is used, then only the slave I/O is affected. If both `SQL_THREAD` and `IO_THREAD` are used, or if neither of them is used, then both the SQL and I/O threads are affected by the statement.

The `UNTIL` clause is not supported for multi-threaded slaves except when also using `SQL_AFTER_MTS_GAPS`. Prior to MySQL 5.6.6, `UNTIL` was not supported at all for multi-threaded slaves.

For an `UNTIL` clause, you must specify any one of the following:

- *Both* a log file name and a position in that file
- (MySQL 5.6.6 or later:) *Either* of `SQL_BEFORE_GTIDS` or `SQL_AFTER_GTIDS`
- (MySQL 5.6.6 or later:) `SQL_AFTER_MTS_GAPS`

Do not mix master and relay log options. In MySQL 5.6.6 and later, do not mix log file options with GTID options.

Any `UNTIL` condition is reset by a subsequent `STOP SLAVE` statement, a `START SLAVE` statement that includes no `UNTIL` clause, or a server restart.

When specifying a log file and position, you can use the `IO_THREAD` option with `START SLAVE ... UNTIL` even though only the SQL thread is affected by this statement. The `IO_THREAD` option is ignored in such cases. The preceding restriction does not apply when using one of the GTID options (`SQL_BEFORE_GTIDS` and `SQL_AFTER_GTIDS`) introduced in MySQL 5.6.6; the GTID options support both `SQL_THREAD` and `IO_THREAD`, as explained previously in this section.

The `UNTIL` clause can be useful for debugging replication, or to cause replication to proceed until just before the point where you want to avoid having the slave replicate an event. For example, if an unwise `DROP TABLE` statement was executed on the master, you can use `UNTIL` to tell the slave to execute up to that point but no farther. To find what the event is, use `mysqlbinlog` with the master binary log or slave relay log, or by using a `SHOW BINLOG EVENTS` statement.

If you are using `UNTIL` to have the slave process replicated queries in sections, it is recommended that you start the slave with the `--skip-slave-start [1984]` option to prevent the SQL thread from running when the slave server starts. It is probably best to use this option in an option file rather than on the command line, so that an unexpected server restart does not cause it to be forgotten.

The `SHOW SLAVE STATUS` statement includes output fields that display the current values of the `UNTIL` condition.

In very old versions of MySQL (before 4.0.5), this statement was called `SLAVE START`. That syntax is no longer accepted as of MySQL 5.6.1.

### 13.4.2.6. `STOP SLAVE` Syntax

```
STOP SLAVE [thread_types]

thread_types:
    [thread_type [, thread_type] ... ]

thread_type: IO_THREAD | SQL_THREAD
```

Stops the slave threads. `STOP SLAVE` requires the `SUPER [756]` privilege. Recommended best practice is to execute `STOP SLAVE` on the slave before stopping the slave server (see [Section 5.1.12, “The Shutdown Process”](#), for more information).

*When using the row-based logging format.* You should execute `STOP SLAVE` or `STOP SLAVE SQL_THREAD` on the slave prior to shutting down the slave server if you are replicating any tables that use a nontransactional storage engine (see the *Note* later in this section).

Like `START SLAVE`, this statement may be used with the `IO_THREAD` and `SQL_THREAD` options to name the thread or threads to be stopped.

In MySQL 5.6.7 and later, `STOP SLAVE` causes an implicit commit of an ongoing transaction. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).



#### Note

In MySQL 5.6, `STOP SLAVE` waits until the current replication event group affecting one or more nontransactional tables has finished executing (if there is any such replication group), or until the user issues a `KILL QUERY` or `KILL CONNECTION` statement. (Bug #319, Bug #38205)

In old versions of MySQL (before 4.0.5), this statement was called `SLAVE STOP`. That syntax is no longer accepted as of MySQL 5.6.1.

## 13.5. SQL Syntax for Prepared Statements

MySQL 5.6 provides support for server-side prepared statements. This support takes advantage of the efficient client/server binary protocol available since MySQL 4.1. Using prepared statements with placeholders for parameter values has the following benefits:

- Less overhead for parsing the statement each time it is executed. Typically, database applications process large volumes of almost-identical statements, with only changes to literal or variable values in clauses such as `WHERE` for queries and deletes, `SET` for updates, and `VALUES` for inserts.
- Protection against SQL injection attacks. The parameter values can contain unescaped SQL quote and delimiter characters.

### Prepared Statements in Application Programs

You can use server-side prepared statements through client programming interfaces, including the [MySQL C API client library](#) or [MySQL Connector/C](#) for C programs, [MySQL Connector/J](#) for Java programs, and [MySQL Connector/Net](#) for programs using .NET technologies. For example, the C API provides a set of function calls that make up its prepared statement API. See [Section 21.8.7, “C API Prepared Statements”](#). Other language interfaces can provide support for prepared statements that use the binary protocol by linking in the C client library, one example being the [mysql\\_i extension](#), available in PHP 5.0 and later.

### Prepared Statements in SQL Scripts

An alternative SQL interface to prepared statements is available. This interface is not as efficient as using the binary protocol through a prepared statement API, but requires no programming because it is available directly at the SQL level:

- You can use it when no programming interface is available to you.
- You can use it from any program that can send SQL statements to the server to be executed, such as the `mysql` client program.
- You can use it even if the client is using an old version of the client library, as long as you connect to a server running MySQL 4.1 or higher.

SQL syntax for prepared statements is intended to be used for situations such as these:

- To test how prepared statements work in your application before coding it.

- To use prepared statements when you do not have access to a programming API that supports them.
- To interactively troubleshoot application issues with prepared statements.
- To create a test case that reproduces a problem with prepared statements, so that you can file a bug report.

## PREPARE, EXECUTE, and DEALLOCATE PREPARE Statements

SQL syntax for prepared statements is based on three SQL statements:

- `PREPARE` prepares a statement for execution (see [Section 13.5.1, “PREPARE Syntax”](#)).
- `EXECUTE` executes a prepared statement (see [Section 13.5.2, “EXECUTE Syntax”](#)).
- `DEALLOCATE PREPARE` releases a prepared statement (see [Section 13.5.3, “DEALLOCATE PREPARE Syntax”](#)).

The following examples show two equivalent ways of preparing a statement that computes the hypotenuse of a triangle given the lengths of the two sides.

The first example shows how to create a prepared statement by using a string literal to supply the text of the statement:

```
mysql> PREPARE stmt1 FROM 'SELECT SQRT(POW(?,2) + POW(?,2)) AS hypotenuse';
mysql> SET @a = 3;
mysql> SET @b = 4;
mysql> EXECUTE stmt1 USING @a, @b;
+-----+
| hypotenuse |
+-----+
|          5 |
+-----+
mysql> DEALLOCATE PREPARE stmt1;
```

The second example is similar, but supplies the text of the statement as a user variable:

```
mysql> SET @s = 'SELECT SQRT(POW(?,2) + POW(?,2)) AS hypotenuse';
mysql> PREPARE stmt2 FROM @s;
mysql> SET @a = 6;
mysql> SET @b = 8;
mysql> EXECUTE stmt2 USING @a, @b;
+-----+
| hypotenuse |
+-----+
|          10 |
+-----+
mysql> DEALLOCATE PREPARE stmt2;
```

Here is an additional example that demonstrates how to choose the table on which to perform a query at runtime, by storing the name of the table as a user variable:

```
mysql> USE test;
mysql> CREATE TABLE t1 (a INT NOT NULL);
mysql> INSERT INTO t1 VALUES (4), (8), (11), (32), (80);

mysql> SET @table = 't1';
mysql> SET @s = CONCAT('SELECT * FROM ', @table);

mysql> PREPARE stmt3 FROM @s;
mysql> EXECUTE stmt3;
+----+
| a  |
+----+
|  4 |
|  8 |
```



```
| 11 |
| 32 |
| 80 |
+-----+
mysql> DEALLOCATE PREPARE stmt3;
```

A prepared statement is specific to the session in which it was created. If you terminate a session without deallocating a previously prepared statement, the server deallocates it automatically.

A prepared statement is also global to the session. If you create a prepared statement within a stored routine, it is not deallocated when the stored routine ends.

To guard against too many prepared statements being created simultaneously, set the `max_prepared_stmt_count` [506] system variable. To prevent the use of prepared statements, set the value to 0.

## SQL Syntax Allowed in Prepared Statements

The following SQL statements can be used as prepared statements:

```
ALTER TABLE
ALTER USER (as of MySQL 5.6.8)
ANALYZE TABLE
CACHE INDEX
CALL
CHANGE MASTER
CHECKSUM {TABLE | TABLES}
COMMIT
{CREATE | RENAME | DROP} DATABASE
{CREATE | DROP} INDEX
{CREATE | RENAME | DROP} TABLE
{CREATE | RENAME | DROP} USER
DELETE
DO
FLUSH {TABLE | TABLES | TABLES WITH READ LOCK | HOSTS | PRIVILEGES
      | LOGS | STATUS | MASTER | SLAVE | DES_KEY_FILE | USER_RESOURCES}
GRANT
INSERT
INSTALL PLUGIN
KILL
LOAD INDEX INTO CACHE
OPTIMIZE TABLE
REPAIR TABLE
REPLACE
RESET {MASTER | SLAVE | QUERY CACHE}
REVOKE
SELECT
SET
SHOW {AUTHORS | CONTRIBUTORS | WARNINGS | ERRORS}
SHOW BINLOG EVENTS
SHOW CREATE {PROCEDURE | FUNCTION | EVENT | TABLE | VIEW}
SHOW {MASTER | BINARY} LOGS
SHOW {MASTER | SLAVE} STATUS
SLAVE {START | STOP}
TRUNCATE TABLE
UNINSTALL PLUGIN
UPDATE
```

Other statements are not supported in MySQL 5.6.

Generally, statements not permitted in SQL prepared statements are also not permitted in stored programs. Exceptions are noted in [Section E.1, “Restrictions on Stored Programs”](#).

Metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

Placeholders can be used for the arguments of the `LIMIT` clause when using prepared statements. See [Section 13.2.9, “SELECT Syntax”](#).

In prepared `CALL` statements used with `PREPARE` and `EXECUTE`, placeholder support for `OUT` and `INOUT` parameters is available beginning with MySQL 5.6. See [Section 13.2.1, “CALL Syntax”](#), for an example and a workaround for earlier versions. Placeholders can be used for `IN` parameters regardless of version.

SQL syntax for prepared statements cannot be used in nested fashion. That is, a statement passed to `PREPARE` cannot itself be a `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE` statement.

SQL syntax for prepared statements is distinct from using prepared statement API calls. For example, you cannot use the `mysql_stmt_prepare()` C API function to prepare a `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE` statement.

SQL syntax for prepared statements can be used within stored procedures, but not in stored functions or triggers. However, a cursor cannot be used for a dynamic statement that is prepared and executed with `PREPARE` and `EXECUTE`. The statement for a cursor is checked at cursor creation time, so the statement cannot be dynamic.

SQL syntax for prepared statements does not support multi-statements (that is, multiple statements within a single string separated by “;” characters).

Prepared statements use the query cache under the conditions described in [Section 8.9.3.1, “How the Query Cache Operates”](#).

To write C programs that use the `CALL` SQL statement to execute stored procedures that contain prepared statements, the `CLIENT_MULTI_RESULTS` flag must be enabled. This is because each `CALL` returns a result to indicate the call status, in addition to any result sets that might be returned by statements executed within the procedure.

`CLIENT_MULTI_RESULTS` can be enabled when you call `mysql_real_connect()`, either explicitly by passing the `CLIENT_MULTI_RESULTS` flag itself, or implicitly by passing `CLIENT_MULTI_STATEMENTS` (which also enables `CLIENT_MULTI_RESULTS`). For additional information, see [Section 13.2.1, “CALL Syntax”](#).

### 13.5.1. PREPARE Syntax

```
PREPARE stmt_name FROM preparable_stmt
```

The `PREPARE` statement prepares a SQL statement and assigns it a name, *stmt\_name*, by which to refer to the statement later. The prepared statement is executed with `EXECUTE` and released with `DEALLOCATE PREPARE`. For examples, see [Section 13.5, “SQL Syntax for Prepared Statements”](#).

Statement names are not case sensitive. *preparable\_stmt* is either a string literal or a user variable that contains the text of the SQL statement. The text must represent a single statement, not multiple statements. Within the statement, `?` characters can be used as parameter markers to indicate where data values are to be bound to the query later when you execute it. The `?` characters should not be enclosed within quotation marks, even if you intend to bind them to string values. Parameter markers can be used only where data values should appear, not for SQL keywords, identifiers, and so forth.

If a prepared statement with the given name already exists, it is deallocated implicitly before the new statement is prepared. This means that if the new statement contains an error and cannot be prepared, an error is returned and no statement with the given name exists.

The scope of a prepared statement is the session within which it is created, which has several implications:

- A prepared statement created in one session is not available to other sessions.

- When a session ends, whether normally or abnormally, its prepared statements no longer exist. If auto-reconnect is enabled, the client is not notified that the connection was lost. For this reason, clients may wish to disable auto-reconnect. See [Section 21.8.15, “Controlling Automatic Reconnection Behavior”](#).
- A prepared statement created within a stored program continues to exist after the program finishes executing and can be executed outside the program later.
- A statement prepared in stored program context cannot refer to stored procedure or function parameters or local variables because they go out of scope when the program ends and would be unavailable were the statement to be executed later outside the program. As a workaround, refer instead to user-defined variables, which also have session scope; see [Section 9.4, “User-Defined Variables”](#).

## 13.5.2. EXECUTE Syntax

```
EXECUTE stmt_name
    [USING @var_name [, @var_name] ...]
```

After preparing a statement with `PREPARE`, you execute it with an `EXECUTE` statement that refers to the prepared statement name. If the prepared statement contains any parameter markers, you must supply a `USING` clause that lists user variables containing the values to be bound to the parameters. Parameter values can be supplied only by user variables, and the `USING` clause must name exactly as many variables as the number of parameter markers in the statement.

You can execute a given prepared statement multiple times, passing different variables to it or setting the variables to different values before each execution.

For examples, see [Section 13.5, “SQL Syntax for Prepared Statements”](#).

## 13.5.3. DEALLOCATE PREPARE Syntax

```
{DEALLOCATE | DROP} PREPARE stmt_name
```

To deallocate a prepared statement produced with `PREPARE`, use a `DEALLOCATE PREPARE` statement that refers to the prepared statement name. Attempting to execute a prepared statement after deallocating it results in an error. If too many prepared statements are created and not deallocated by either the `DEALLOCATE PREPARE` statement or the end of the session, you might encounter the upper limit enforced by the `max_prepared_stmt_count` [506] system variable.

For examples, see [Section 13.5, “SQL Syntax for Prepared Statements”](#).

## 13.6. MySQL Compound-Statement Syntax

This section describes the syntax for the `BEGIN ... END` compound statement and other statements that can be used in the body of stored programs: Stored procedures and functions, triggers, and events. These objects are defined in terms of SQL code that is stored on the server for later invocation (see [Chapter 18, \*Stored Programs and Views\*](#)).

A compound statement is a block that can contain other blocks; declarations for variables, condition handlers, and cursors; and flow control constructs such as loops and conditional tests.

### 13.6.1. BEGIN ... END Compound-Statement Syntax

```
[begin_label:] BEGIN
    [statement_list]
END [end_label]
```

`BEGIN ... END` syntax is used for writing compound statements, which can appear within stored programs (stored procedures and functions, triggers, and events). A compound statement can contain

multiple statements, enclosed by the `BEGIN` and `END` keywords. `statement_list` represents a list of one or more statements, each terminated by a semicolon (`;`) statement delimiter. The `statement_list` itself is optional, so the empty compound statement (`BEGIN END`) is legal.

`BEGIN ... END` blocks can be nested.

Use of multiple statements requires that a client is able to send statement strings containing the `;` statement delimiter. In the `mysql` command-line client, this is handled with the `delimiter` command. Changing the `;` end-of-statement delimiter (for example, to `//`) permit `;` to be used in a program body. For an example, see [Section 18.1, “Defining Stored Programs”](#).

A `BEGIN ... END` block can be labeled. See [Section 13.6.2, “Statement Label Syntax”](#).

The optional `[NOT] ATOMIC` clause is not supported. This means that no transactional savepoint is set at the start of the instruction block and the `BEGIN` clause used in this context has no effect on the current transaction.



#### Note

Within all stored programs, the parser treats `BEGIN [WORK]` as the beginning of a `BEGIN ... END` block. To begin a transaction in this context, use `START TRANSACTION` instead.

## 13.6.2. Statement Label Syntax

```
[begin_label:] BEGIN
    [statement_list]
END [end_label]

[begin_label:] LOOP
    statement_list
END LOOP [end_label]

[begin_label:] REPEAT
    statement_list
UNTIL search_condition
END REPEAT [end_label]

[begin_label:] WHILE search_condition DO
    statement_list
END WHILE [end_label]
```

Labels are permitted for `BEGIN ... END` blocks and for the `LOOP`, `REPEAT`, and `WHILE` statements. Label use for those statements follows these rules:

- `begin_label` must be followed by a colon.
- `begin_label` can be given without `end_label`. If `end_label` is present, it must be the same as `begin_label`.
- `end_label` cannot be given without `begin_label`.
- Labels at the same nesting level must be distinct.
- Labels can be up to 16 characters long.

To refer to a label within the labeled construct, use an `ITERATE` or `LEAVE` statement. The following example uses those statements to continue iterating or terminate the loop:

```
CREATE PROCEDURE doiterate(p1 INT)
BEGIN
    label1: LOOP
        SET p1 = p1 + 1;
        IF p1 < 10 THEN ITERATE label1; END IF;
```

```
LEAVE label1;
END LOOP label1;
END;
```

The scope of a block label does not include the code for handlers declared within the block. For details, see [Section 13.6.7.2, “DECLARE ... HANDLER Syntax”](#).

### 13.6.3. DECLARE Syntax

The `DECLARE` statement is used to define various items local to a program:

- Local variables. See [Section 13.6.4, “Variables in Stored Programs”](#).
- Conditions and handlers. See [Section 13.6.7, “Condition Handling”](#).
- Cursors. See [Section 13.6.6, “Cursors”](#).

`DECLARE` is permitted only inside a `BEGIN ... END` compound statement and must be at its start, before any other statements.

Declarations must follow a certain order. Cursor declarations must appear before handler declarations. Variable and condition declarations must appear before cursor or handler declarations.

### 13.6.4. Variables in Stored Programs

System variables and user-defined variables can be used in stored programs, just as they can be used outside stored-program context. In addition, stored programs can use `DECLARE` to define local variables, and stored routines (procedures and functions) can be declared to take parameters that communicate values between the routine and its caller.

- To declare local variables, use the `DECLARE` statement, as described in [Section 13.6.4.1, “Local Variable DECLARE Syntax”](#).
- Variables can be set directly with the `SET` statement. See [Section 13.7.4, “SET Syntax”](#).
- Results from queries can be retrieved into local variables using `SELECT ... INTO var_list` or by opening a cursor and using `FETCH ... INTO var_list`. See [Section 13.2.9.1, “SELECT ... INTO Syntax”](#), and [Section 13.6.6, “Cursors”](#).

For information about the scope of local variables and how MySQL resolves ambiguous names, see [Section 13.6.4.2, “Local Variable Scope and Resolution”](#).

It is not permitted to assign the value `DEFAULT` to stored procedure or function parameters or stored program local variables (for example with a `SET var_name = DEFAULT` statement). As of MySQL 5.6.6, this results in a syntax error.

#### 13.6.4.1. Local Variable DECLARE Syntax

```
DECLARE var_name [, var_name] ... type [DEFAULT value]
```

This statement declares local variables within stored programs. To provide a default value for a variable, include a `DEFAULT` clause. The value can be specified as an expression; it need not be a constant. If the `DEFAULT` clause is missing, the initial value is `NULL`.

Local variables are treated like stored routine parameters with respect to data type and overflow checking. See [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#).

Variable declarations must appear before cursor or handler declarations.

Local variable names are not case sensitive. Permissible characters and quoting rules are the same as for other identifiers, as described in [Section 9.2, “Schema Object Names”](#).

The scope of a local variable is the `BEGIN ... END` block within which it is declared. The variable can be referred to in blocks nested within the declaring block, except those blocks that declare a variable with the same name.

### 13.6.4.2. Local Variable Scope and Resolution

The scope of a local variable is the `BEGIN ... END` block within which it is declared. The variable can be referred to in blocks nested within the declaring block, except those blocks that declare a variable with the same name.

Because local variables are in scope only during stored program execution, references to them are not permitted in prepared statements created within a stored program. Prepared statement scope is the current session, not the stored program, so the statement could be executed after the program ends, at which point the variables would no longer be in scope. For example, `SELECT ... INTO local_var` cannot be used as a prepared statement. This restriction also applies to stored procedure and function parameters. See [Section 13.5.1, “PREPARE Syntax”](#).

A local variable should not have the same name as a table column. If an SQL statement, such as a `SELECT ... INTO` statement, contains a reference to a column and a declared local variable with the same name, MySQL currently interprets the reference as the name of a variable. Consider the following procedure definition:

```
CREATE PROCEDURE sp1 (x VARCHAR(5))
BEGIN
  DECLARE xname VARCHAR(5) DEFAULT 'bob';
  DECLARE newname VARCHAR(5);
  DECLARE xid INT;

  SELECT xname, id INTO newname, xid
  FROM table1 WHERE xname = xname;
  SELECT newname;
END;
```

MySQL interprets `xname` in the `SELECT` statement as a reference to the `xname` variable rather than the `xname` column. Consequently, when the procedure `sp1()` is called, the `newname` variable returns the value `'bob'` regardless of the value of the `table1.xname` column.

Similarly, the cursor definition in the following procedure contains a `SELECT` statement that refers to `xname`. MySQL interprets this as a reference to the variable of that name rather than a column reference.

```
CREATE PROCEDURE sp2 (x VARCHAR(5))
BEGIN
  DECLARE xname VARCHAR(5) DEFAULT 'bob';
  DECLARE newname VARCHAR(5);
  DECLARE xid INT;
  DECLARE done TINYINT DEFAULT 0;
  DECLARE cur1 CURSOR FOR SELECT xname, id FROM table1;
  DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = 1;

  OPEN cur1;
  read_loop: LOOP
    FETCH FROM cur1 INTO newname, xid;
    IF done THEN LEAVE read_loop; END IF;
    SELECT newname;
  END LOOP;
  CLOSE cur1;
END;
```

See also [Section E.1, “Restrictions on Stored Programs”](#).

## 13.6.5. Flow Control Statements

MySQL supports the `IF`, `CASE`, `ITERATE`, `LEAVE LOOP`, `WHILE`, and `REPEAT` constructs for flow control within stored programs. It also supports `RETURN` within stored functions.

Many of these constructs contain other statements, as indicated by the grammar specifications in the following sections. Such constructs may be nested. For example, an `IF` statement might contain a `WHILE` loop, which itself contains a `CASE` statement.

MySQL does not support `FOR` loops.

### 13.6.5.1. `CASE` Syntax

```
CASE case_value
  WHEN when_value THEN statement_list
  [WHEN when_value THEN statement_list] ...
  [ELSE statement_list]
END CASE
```

Or:

```
CASE
  WHEN search_condition THEN statement_list
  [WHEN search_condition THEN statement_list] ...
  [ELSE statement_list]
END CASE
```

The `CASE` statement for stored programs implements a complex conditional construct.



#### Note

There is also a `CASE [1174] expression`, which differs from the `CASE statement` described here. See [Section 12.4, “Control Flow Functions”](#). The `CASE statement` cannot have an `ELSE NULL` clause, and it is terminated with `END CASE` instead of `END`.

For the first syntax, `case_value` is an expression. This value is compared to the `when_value` expression in each `WHEN` clause until one of them is equal. When an equal `when_value` is found, the corresponding `THEN` clause `statement_list` executes. If no `when_value` is equal, the `ELSE` clause `statement_list` executes, if there is one.

This syntax cannot be used to test for equality with `NULL` because `NULL = NULL` is false. See [Section 3.3.4.6, “Working with NULL Values”](#).

For the second syntax, each `WHEN` clause `search_condition` expression is evaluated until one is true, at which point its corresponding `THEN` clause `statement_list` executes. If no `search_condition` is equal, the `ELSE` clause `statement_list` executes, if there is one.

If no `when_value` or `search_condition` matches the value tested and the `CASE` statement contains no `ELSE` clause, a `Case not found for CASE statement` error results.

Each `statement_list` consists of one or more SQL statements; an empty `statement_list` is not permitted.

To handle situations where no value is matched by any `WHEN` clause, use an `ELSE` containing an empty `BEGIN ... END` block, as shown in this example. (The indentation used here in the `ELSE` clause is for purposes of clarity only, and is not otherwise significant.)

```
DELIMITER |

CREATE PROCEDURE p()
  BEGIN
    DECLARE v INT DEFAULT 1;

    CASE v
      WHEN 2 THEN SELECT v;
      WHEN 3 THEN SELECT 0;
      ELSE
        BEGIN

```



```

        END;
    END CASE;
END;
|

```

### 13.6.5.2. IF Syntax

```

IF search_condition THEN statement_list
  [ELSEIF search_condition THEN statement_list] ...
  [ELSE statement_list]
END IF

```

The `IF` statement for stored programs implements a basic conditional construct.



#### Note

There is also an `IF()` [1174] *function*, which differs from the `IF statement` described here. See [Section 12.4, “Control Flow Functions”](#). The `IF` statement can have `THEN`, `ELSE`, and `ELSEIF` clauses, and it is terminated with `END IF`.

If the `search_condition` evaluates to true, the corresponding `THEN` or `ELSEIF` clause `statement_list` executes. If no `search_condition` matches, the `ELSE` clause `statement_list` executes.

Each `statement_list` consists of one or more SQL statements; an empty `statement_list` is not permitted.

An `IF ... END IF` block, like all other flow-control blocks used within stored programs, must be terminated with a semicolon, as shown in this example:

```

DELIMITER //

CREATE FUNCTION SimpleCompare(n INT, m INT)
  RETURNS VARCHAR(20)

BEGIN
  DECLARE s VARCHAR(20);

  IF n > m THEN SET s = '>';
  ELSEIF n = m THEN SET s = '=';
  ELSE SET s = '<';
  END IF;

  SET s = CONCAT(n, ' ', s, ' ', m);

  RETURN s;
END //

DELIMITER ;

```

As with other flow-control constructs, `IF ... END IF` blocks may be nested within other flow-control constructs, including other `IF` statements. Each `IF` must be terminated by its own `END IF` followed by a semicolon. You can use indentation to make nested flow-control blocks more easily readable by humans (although this is not required by MySQL), as shown here:

```

DELIMITER //

CREATE FUNCTION VerboseCompare (n INT, m INT)
  RETURNS VARCHAR(50)

BEGIN
  DECLARE s VARCHAR(50);

  IF n = m THEN SET s = 'equals';
  ELSE
    IF n > m THEN SET s = 'greater';
    ELSE SET s = 'less';
  END IF;
END

```

```

END IF;

SET s = CONCAT('is ', s, ' than');
END IF;

SET s = CONCAT(n, ' ', s, ' ', m, '.');

RETURN s;
END //
DELIMITER ;

```

In this example, the inner `IF` is evaluated only if `n` is not equal to `m`.

### 13.6.5.3. ITERATE Syntax

```
ITERATE label
```

`ITERATE` can appear only within `LOOP`, `REPEAT`, and `WHILE` statements. `ITERATE` means “start the loop again.”

For an example, see [Section 13.6.5.5, “LOOP Syntax”](#).

### 13.6.5.4. LEAVE Syntax

```
LEAVE label
```

This statement is used to exit the flow control construct that has the given label. If the label is for the outermost stored program block, `LEAVE` exits the program.

`LEAVE` can be used within `BEGIN ... END` or loop constructs (`LOOP`, `REPEAT`, `WHILE`).

For an example, see [Section 13.6.5.5, “LOOP Syntax”](#).

### 13.6.5.5. LOOP Syntax

```
[begin_label:] LOOP
    statement_list
END LOOP [end_label]
```

`LOOP` implements a simple loop construct, enabling repeated execution of the statement list, which consists of one or more statements, each terminated by a semicolon (`;`) statement delimiter. The statements within the loop are repeated until the loop is terminated. Usually, this is accomplished with a `LEAVE` statement. Within a stored function, `RETURN` can also be used, which exits the function entirely.

Neglecting to include a loop-termination statement results in an infinite loop.

A `LOOP` statement can be labeled. For the rules regarding label use, see [Section 13.6.2, “Statement Label Syntax”](#).

Example:

```

CREATE PROCEDURE doiterate(p1 INT)
BEGIN
    label1: LOOP
        SET p1 = p1 + 1;
        IF p1 < 10 THEN
            ITERATE label1;
        END IF;
        LEAVE label1;
    END LOOP label1;
    SET @x = p1;
END;

```

### 13.6.5.6. REPEAT Syntax

```
[begin_label:] REPEAT
    statement_list
UNTIL search_condition
END REPEAT [end_label]
```

The statement list within a `REPEAT` statement is repeated until the `search_condition` expression is true. Thus, a `REPEAT` always enters the loop at least once. `statement_list` consists of one or more statements, each terminated by a semicolon (;) statement delimiter.

A `REPEAT` statement can be labeled. For the rules regarding label use, see [Section 13.6.2, “Statement Label Syntax”](#).

Example:

```
mysql> delimiter //

mysql> CREATE PROCEDURE dorepeat(p1 INT)
-> BEGIN
->   SET @x = 0;
->   REPEAT
->     SET @x = @x + 1;
->   UNTIL @x > p1 END REPEAT;
-> END
-> //
Query OK, 0 rows affected (0.00 sec)

mysql> CALL dorepeat(1000)//
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @x//
+-----+
| @x    |
+-----+
| 1001  |
+-----+
1 row in set (0.00 sec)
```

### 13.6.5.7. RETURN Syntax

```
RETURN expr
```

The `RETURN` statement terminates execution of a stored function and returns the value `expr` to the function caller. There must be at least one `RETURN` statement in a stored function. There may be more than one if the function has multiple exit points.

This statement is not used in stored procedures, triggers, or events. The `LEAVE` statement can be used to exit a stored program of those types.

### 13.6.5.8. WHILE Syntax

```
[begin_label:] WHILE search_condition DO
    statement_list
END WHILE [end_label]
```

The statement list within a `WHILE` statement is repeated as long as the `search_condition` expression is true. `statement_list` consists of one or more SQL statements, each terminated by a semicolon (;) statement delimiter.

A `WHILE` statement can be labeled. For the rules regarding label use, see [Section 13.6.2, “Statement Label Syntax”](#).

Example:

```
CREATE PROCEDURE dowhile()
BEGIN
```

```

DECLARE v1 INT DEFAULT 5;

WHILE v1 > 0 DO
    ...
    SET v1 = v1 - 1;
END WHILE;
END;

```

## 13.6.6. Cursors

MySQL supports cursors inside stored programs. The syntax is as in embedded SQL. Cursors have these properties:

- Asensitive: The server may or may not make a copy of its result table
- Read only: Not updatable
- Nonscrollable: Can be traversed only in one direction and cannot skip rows

Cursor declarations must appear before handler declarations and after variable and condition declarations.

Example:

```

CREATE PROCEDURE curdemo()
BEGIN
    DECLARE done INT DEFAULT FALSE;
    DECLARE a CHAR(16);
    DECLARE b, c INT;
    DECLARE cur1 CURSOR FOR SELECT id,data FROM test.t1;
    DECLARE cur2 CURSOR FOR SELECT i FROM test.t2;
    DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE;

    OPEN cur1;
    OPEN cur2;

    read_loop: LOOP
        FETCH cur1 INTO a, b;
        FETCH cur2 INTO c;
        IF done THEN
            LEAVE read_loop;
        END IF;
        IF b < c THEN
            INSERT INTO test.t3 VALUES (a,b);
        ELSE
            INSERT INTO test.t3 VALUES (a,c);
        END IF;
    END LOOP;

    CLOSE cur1;
    CLOSE cur2;
END;

```

### 13.6.6.1. Cursor **CLOSE** Syntax

```
CLOSE cursor_name
```

This statement closes a previously opened cursor. For an example, see [Section 13.6.6, “Cursors”](#).

An error occurs if the cursor is not open.

If not closed explicitly, a cursor is closed at the end of the `BEGIN ... END` block in which it was declared.

### 13.6.6.2. Cursor **DECLARE** Syntax

```
DECLARE cursor_name CURSOR FOR select_statement
```

This statement declares a cursor and associates it with a `SELECT` statement that retrieves the rows to be traversed by the cursor. To fetch the rows later, use a `FETCH` statement. The number of columns retrieved by the `SELECT` statement must match the number of output variables specified in the `FETCH` statement.

The `SELECT` statement cannot have an `INTO` clause.

Cursor declarations must appear before handler declarations and after variable and condition declarations.

A stored program may contain multiple cursor declarations, but each cursor declared in a given block must have a unique name. For an example, see [Section 13.6.6, “Cursors”](#).

For information available through `SHOW` statements, it is possible in many cases to obtain equivalent information by using a cursor with an `INFORMATION_SCHEMA` table.

### 13.6.6.3. Cursor `FETCH` Syntax

```
FETCH [[NEXT] FROM] cursor_name INTO var_name [, var_name] ...
```

This statement fetches the next row for the `SELECT` statement associated with the specified cursor (which must be open), and advances the cursor pointer. If a row exists, the fetched columns are stored in the named variables. The number of columns retrieved by the `SELECT` statement must match the number of output variables specified in the `FETCH` statement.

If no more rows are available, a No Data condition occurs with SQLSTATE value `'02000'`. To detect this condition, you can set up a handler for it (or for a `NOT FOUND` condition). For an example, see [Section 13.6.6, “Cursors”](#).

### 13.6.6.4. Cursor `OPEN` Syntax

```
OPEN cursor_name
```

This statement opens a previously declared cursor. For an example, see [Section 13.6.6, “Cursors”](#).

## 13.6.7. Condition Handling

Conditions may arise during stored program execution that require special handling, such as exiting the current program block or continuing execution. Handlers can be defined for general conditions such as warnings or exceptions, or for specific conditions such as a particular error code. Specific conditions can be assigned names and referred to that way in handlers.

To name a condition, use the `DECLARE ... CONDITION` statement. To declare a handler, use the `DECLARE ... HANDLER` statement. See [Section 13.6.7.1, “DECLARE ... CONDITION Syntax”](#), and [Section 13.6.7.2, “DECLARE ... HANDLER Syntax”](#). For information about how the server chooses handlers when a condition occurs, see [Section 13.6.7.6, “Scope Rules for Handlers”](#).

To raise a condition, use the `SIGNAL` statement. To modify condition information within a condition handler, use `RESIGNAL`. See [Section 13.6.7.1, “DECLARE ... CONDITION Syntax”](#), and [Section 13.6.7.2, “DECLARE ... HANDLER Syntax”](#).

To retrieve information from the diagnostics area, use the `GET DIAGNOSTICS` statement (see [Section 13.6.7.3, “GET DIAGNOSTICS Syntax”](#)). For information about the diagnostics area, see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

### 13.6.7.1. `DECLARE ... CONDITION` Syntax

```
DECLARE condition_name CONDITION FOR condition_value
condition_value:
```

```
mysql_error_code
| SQLSTATE [VALUE] sqlstate_value
```

The `DECLARE ... CONDITION` statement declares a named error condition, associating a name with a condition that needs specific handling. The name can be referred to in a subsequent `DECLARE ... HANDLER` statement (see [Section 13.6.7.2, “DECLARE ... HANDLER Syntax”](#)).

Condition declarations must appear before cursor or handler declarations.

The *condition\_value* for `DECLARE ... CONDITION` can be a MySQL error code (a number) or an SQLSTATE value (a 5-character string literal). You should not use MySQL error code 0 or SQLSTATE values that begin with '00', because those indicate success rather than an error condition. For a list of MySQL error codes and SQLSTATE values, see [Section C.3, “Server Error Codes and Messages”](#).

Using names for conditions can help make stored program code clearer. For example, this handler applies to attempts to drop a nonexistent table, but that is apparent only if you know the meaning of MySQL error code 1051:

```
DECLARE CONTINUE HANDLER FOR 1051
BEGIN
  -- body of handler
END;
```

By declaring a name for the condition, the purpose of the handler is more readily seen:

```
DECLARE no_such_table CONDITION FOR 1051;
DECLARE CONTINUE HANDLER FOR no_such_table
BEGIN
  -- body of handler
END;
```

Here is a named condition for the same condition, but based on the corresponding SQLSTATE value rather than the MySQL error code:

```
DECLARE no_such_table CONDITION FOR SQLSTATE '42S02';
DECLARE CONTINUE HANDLER FOR no_such_table
BEGIN
  -- body of handler
END;
```

Condition names referred to in `SIGNAL` or use `RESIGNAL` statements must be associated with SQLSTATE values, not MySQL error codes.

### 13.6.7.2. DECLARE ... HANDLER Syntax

```
DECLARE handler_action HANDLER
  FOR condition_value [, condition_value] ...
  statement

handler_action:
  CONTINUE
| EXIT
| UNDO

condition_value:
  mysql_error_code
| SQLSTATE [VALUE] sqlstate_value
| condition_name
| SQLWARNING
| NOT FOUND
| SQLEXCEPTION
```

The `DECLARE ... HANDLER` statement specifies a handler that deals with one or more conditions. If one of these conditions occurs, the specified *statement* executes. *statement* can be a simple statement such as `SET var_name = value`, or a compound statement written using `BEGIN` and `END` (see [Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”](#)).

Handler declarations must appear after variable or condition declarations.

The `handler_action` value indicates what action the handler takes after execution of the handler statement:

- `CONTINUE`: Execution of the current program continues.
- `EXIT`: Execution terminates for the `BEGIN ... END` compound statement in which the handler is declared. This is true even if the condition occurs in an inner block.
- `UNDO`: Not supported.

The `condition_value` for `DECLARE ... HANDLER` indicates the specific condition or class of conditions that activates the handler:

- A MySQL error code (a number) or an SQLSTATE value (a 5-character string literal). You should not use MySQL error code 0 or SQLSTATE values that begin with `'00'`, because those indicate success rather than an error condition. For a list of MySQL error codes and SQLSTATE values, see [Section C.3, “Server Error Codes and Messages”](#).
- A condition name previously specified with `DECLARE ... CONDITION`. A condition name can be associated with a MySQL error code or SQLSTATE value. See [Section 13.6.7.1, “DECLARE ... CONDITION Syntax”](#).
- `SQLWARNING` is shorthand for the class of SQLSTATE values that begin with `'01'`.
- `NOT FOUND` is shorthand for the class of SQLSTATE values that begin with `'02'`. This is relevant within the context of cursors and is used to control what happens when a cursor reaches the end of a data set. If no more rows are available, a No Data condition occurs with SQLSTATE value `'02000'`. To detect this condition, you can set up a handler for it (or for a `NOT FOUND` condition). For an example, see [Section 13.6.6, “Cursors”](#). This condition also occurs for `SELECT ... INTO var_list` statements that retrieve no rows.
- `SQLEXCEPTION` is shorthand for the class of SQLSTATE values that do not begin with `'00'`, `'01'`, or `'02'`.

For information about how the server chooses handlers when a condition occurs, see [Section 13.6.7.6, “Scope Rules for Handlers”](#).

If a condition occurs for which no handler has been declared, the action taken depends on the condition class:

- For `SQLEXCEPTION` conditions, the stored program terminates at the statement that raised the condition, as if there were an `EXIT` handler. If the program was called by another stored program, the calling program handles the condition using the handler selection rules applied to its own handlers.
- For `SQLWARNING` conditions, the program continues executing, as if there were a `CONTINUE` handler.
- For `NOT FOUND` conditions, if the condition was raised normally, the action is `CONTINUE`. If it was raised by `SIGNAL` or `RESIGNAL`, the action is `EXIT`.

The following example uses a handler for `SQLSTATE '23000'`, which occurs for a duplicate-key error:

```
mysql> CREATE TABLE test.t (s1 INT, PRIMARY KEY (s1));
Query OK, 0 rows affected (0.00 sec)

mysql> delimiter //

mysql> CREATE PROCEDURE handlerdemo ()
-> BEGIN
->   DECLARE CONTINUE HANDLER FOR SQLSTATE '23000' SET @x2 = 1;
->   SET @x = 1;
```



```

-> INSERT INTO test.t VALUES (1);
-> SET @x = 2;
-> INSERT INTO test.t VALUES (1);
-> SET @x = 3;
-> END;
-> //
Query OK, 0 rows affected (0.00 sec)

mysql> CALL handlerdemo();//
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @x//
+-----+
| @x    |
+-----+
| 3     |
+-----+
1 row in set (0.00 sec)

```

Notice that `@x` is 3 after the procedure executes, which shows that execution continued to the end of the procedure after the error occurred. If the `DECLARE ... HANDLER` statement had not been present, MySQL would have taken the default action (`EXIT`) after the second `INSERT` failed due to the `PRIMARY KEY` constraint, and `SELECT @x` would have returned 2.

To ignore a condition, declare a `CONTINUE` handler for it and associate it with an empty block. For example:

```
DECLARE CONTINUE HANDLER FOR SQLWARNING BEGIN END;
```

The scope of a block label does not include the code for handlers declared within the block. Therefore, the statement associated with a handler cannot use `ITERATE` or `LEAVE` to refer to labels for blocks that enclose the handler declaration. Consider the following example, where the `REPEAT` block has a label of `retry`:

```

CREATE PROCEDURE p ()
BEGIN
  DECLARE i INT DEFAULT 3;
  retry:
  REPEAT
  BEGIN
    DECLARE CONTINUE HANDLER FOR SQLWARNING
    BEGIN
      ITERATE retry;      # illegal
    END;
    IF i < 0 THEN
      LEAVE retry;       # legal
    END IF;
    SET i = i - 1;
  END;
  UNTIL FALSE END REPEAT;
END;

```

The `retry` label is in scope for the `IF` statement within the block. It is not in scope for the `CONTINUE` handler, so the reference there is invalid and results in an error:

```
ERROR 1308 (42000): LEAVE with no matching label: retry
```

To avoid references to outer labels in handlers, use one of these strategies:

- To leave the block, use an `EXIT` handler. If no block cleanup is required, the `BEGIN ... END` handler body can be empty:

```
DECLARE EXIT HANDLER FOR SQLWARNING BEGIN END;
```

Otherwise, put the cleanup statements in the handler body:

```
DECLARE EXIT HANDLER FOR SQLWARNING
BEGIN
```

```

    block cleanup statements
END;

```

- To continue execution, set a status variable in a `CONTINUE` handler that can be checked in the enclosing block to determine whether the handler was invoked. The following example uses the variable `done` for this purpose:

```

CREATE PROCEDURE p ()
BEGIN
  DECLARE i INT DEFAULT 3;
  DECLARE done INT DEFAULT FALSE;
  retry:
  REPEAT
  BEGIN
    DECLARE CONTINUE HANDLER FOR SQLWARNING
    BEGIN
      SET done = TRUE;
    END;
    IF done OR i < 0 THEN
      LEAVE retry;
    END IF;
    SET i = i - 1;
  END;
  UNTIL FALSE END REPEAT;
END;

```

### 13.6.7.3. GET DIAGNOSTICS Syntax

```

GET [CURRENT] DIAGNOSTICS
{
  statement_information_item
  [, statement_information_item] ...
  | CONDITION condition_number
  condition_information_item
  [, condition_information_item] ...
}

statement_information_item:
  target = statement_information_item_name

condition_information_item:
  target = condition_information_item_name

statement_information_item_name:
  NUMBER
  | ROW_COUNT

condition_information_item_name:
  CLASS_ORIGIN
  | SUBCLASS_ORIGIN
  | RETURNED_SQLSTATE
  | MESSAGE_TEXT
  | MYSQL_ERRNO
  | CONSTRAINT_CATALOG
  | CONSTRAINT_SCHEMA
  | CONSTRAINT_NAME
  | CATALOG_NAME
  | SCHEMA_NAME
  | TABLE_NAME
  | COLUMN_NAME
  | CURSOR_NAME

condition_number, target:
  (see following discussion)

```

SQL statements produce diagnostic information that populates the diagnostics area. The `GET DIAGNOSTICS` statement enables applications to inspect this information. It is available as of MySQL 5.6.4. (You can also use `SHOW WARNINGS` or `SHOW ERRORS` to see all conditions or all errors.)

No special privileges are required to execute `GET DIAGNOSTICS`.

The keyword `CURRENT` means to retrieve information from the current diagnostics area. In MySQL, it has no effect because that is the default behavior.

For a description of the diagnostics area, see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#). Briefly, it contains two kinds of information:

- Statement information, such as the affected-rows count or the number of conditions that occurred.
- Condition information, such as the error code and message. If a statement raises multiple conditions, this part of the diagnostics area has a condition area for each one. If a statement raises no conditions, this part of the diagnostics area is empty.

For a statement that produces three conditions, the diagnostics area contains statement and condition information like this:

```
Statement information:
  row count
  ... other statement information items ...
Condition area list:
  Condition area 1:
    error code for condition 1
    error message for condition 1
    ... other condition information items ...
  Condition area 2:
    error code for condition 2:
    error message for condition 2
    ... other condition information items ...
  Condition area 3:
    error code for condition 3
    error message for condition 3
    ... other condition information items ...
```

`GET DIAGNOSTICS` can obtain either statement or condition information, but not both in the same statement:

- To obtain statement information, retrieve the desired statement items into target variables:

```
GET DIAGNOSTICS @p1 = ROW_COUNT, @p2 = NUMBER;
```

- To obtain condition information, specify the condition number and retrieve the desired condition items into target variables:

```
GET DIAGNOSTICS CONDITION 1
  @p1 = RETURNED_SQLSTATE, @p2 = MESSAGE_TEXT;
```

The retrieval list specifies one or more `target = item_name` assignments, separated by commas. Each assignment names a target variable and either a `statement_information_item_name` or `condition_information_item_name` designator, depending on whether the statement retrieves statement or condition information.

Valid `target` designators for storing item information can be stored procedure or function parameters, stored program local variables declared with `DECLARE`, or user-defined variables.

Valid `condition_number` designators can be stored procedure or function parameters, stored program local variables declared with `DECLARE`, user-defined variables, system variables, or literals. A character literal may include a `_charset` introducer. A warning occurs if the condition number is not in the range from 1 to the number of condition areas that have information. In this case, the warning is added to the diagnostics area without clearing it.

`GET DIAGNOSTICS` is typically used in a handler within a stored program, but it is a MySQL extension that it is permitted outside that context to check the execution of any SQL statement. For example, if you invoke the `mysql` client program, you can enter these statements at the prompt:

```
mysql> DROP TABLE test.no_such_table;
```

```

ERROR 1051 (42S02): Unknown table 'test.no_such_table'
mysql> GET DIAGNOSTICS CONDITION 1
      -> @p1 = RETURNED_SQLSTATE, @p2 = MESSAGE_TEXT;
mysql> SELECT @p1, @p2;
+-----+-----+
| @p1   | @p2                                     |
+-----+-----+
| 42S02 | Unknown table 'test.no_such_table' |
+-----+-----+

```

Currently, not all condition items recognized by `GET DIAGNOSTICS` are populated when a condition occurs. For example:

```

mysql> GET DIAGNOSTICS CONDITION 1
      -> @p3 = SCHEMA_NAME, @p4 = TABLE_NAME;
mysql> SELECT @p3, @p4;
+-----+-----+
| @p3   | @p4   |
+-----+-----+
|       |       |
+-----+-----+

```

For information about permissible statement and condition information items, and which ones are populated when a condition occurs, see [Section 13.6.7.7.2, “Diagnostics Area Information Items”](#).

Here is an example that uses `GET DIAGNOSTICS` and an exception handler in stored procedure context to assess the outcome of an insert operation. If the insert was successful, the procedure also uses `GET DIAGNOSTICS` to get the rows-affected count. This shows that you can use `GET DIAGNOSTICS` multiple times to retrieve information about a statement as long as the diagnostics area has not been cleared.

```

CREATE PROCEDURE do_insert(value INT)
BEGIN
  -- declare variables to hold diagnostics area information
  DECLARE code CHAR(5) DEFAULT '00000';
  DECLARE msg TEXT;
  DECLARE rows INT;
  DECLARE result TEXT;
  -- declare exception handler for failed insert
  DECLARE CONTINUE HANDLER FOR SQLEXCEPTION
  BEGIN
    GET DIAGNOSTICS CONDITION 1
      code = RETURNED_SQLSTATE, msg = MESSAGE_TEXT;
  END;

  -- perform the insert
  INSERT INTO t1 (int_col) VALUES(value);
  -- check whether the insert was successful
  IF code = '00000' THEN
    GET DIAGNOSTICS rows = ROW_COUNT;
    SET result = CONCAT('insert succeeded, row count = ',rows);
  ELSE
    SET result = CONCAT('insert failed, error = ',code,', message = ',msg);
  END IF;
  -- say what happened
  SELECT result;
END;

```

Suppose that `t1.int_col` is an integer column that is declared as `NOT NULL`. The procedure produces these results:

```

mysql> CALL do_insert(1);
+-----+-----+
| result                                     |
+-----+-----+
| insert succeeded, row count = 1 |
+-----+-----+

mysql> CALL do_insert(NULL);

```

```

+-----+
| result |
+-----+
| insert failed, error = 23000, message = Column 'int_col' cannot be null |
+-----+

```

Within a condition handler, `GET DIAGNOSTICS` should be used before other statements that might clear the diagnostics area and cause information to be lost about the condition that activated the handler. For information about when the diagnostics area is set and cleared, see [Section 13.6.7.7.3, “How the Diagnostics Area is Populated”](#).

In standard SQL, the first condition relates to the `SQLSTATE` value returned for the previous SQL statement. In MySQL, this is not guaranteed, so to get the main error, you cannot do this:

```
GET DIAGNOSTICS CONDITION 1 @errno = MYSQL_ERRNO;
```

Instead, do this:

```
GET DIAGNOSTICS @cno = NUMBER;
GET DIAGNOSTICS CONDITION @cno @errno = MYSQL_ERRNO;
```

### 13.6.7.4. RESIGNAL Syntax

```

RESIGNAL [condition_value]
  [SET signal_information_item
    [, signal_information_item] ...]

condition_value:
  SQLSTATE [VALUE] sqlstate_value
  | condition_name

signal_information_item:
  condition_information_item_name = simple_value_specification

condition_information_item_name:
  CLASS_ORIGIN
  | SUBCLASS_ORIGIN
  | MESSAGE_TEXT
  | MYSQL_ERRNO
  | CONSTRAINT_CATALOG
  | CONSTRAINT_SCHEMA
  | CONSTRAINT_NAME
  | CATALOG_NAME
  | SCHEMA_NAME
  | TABLE_NAME
  | COLUMN_NAME
  | CURSOR_NAME

condition_name, simple_value_specification:
  (see following discussion)

```

`RESIGNAL` passes on the error condition information that is available during execution of a condition handler within a compound statement inside a stored procedure or function, trigger, or event. `RESIGNAL` may change some or all information before passing it on. `RESIGNAL` is related to `SIGNAL`, but instead of originating a condition as `SIGNAL` does, `RESIGNAL` relays existing condition information, possibly after modifying it.

`RESIGNAL` makes it possible to both handle an error and return the error information. Otherwise, by executing an SQL statement within the handler, information that caused the handler's activation is destroyed. `RESIGNAL` also can make some procedures shorter if a given handler can handle part of a situation, then pass the condition “up the line” to another handler.

All forms of `RESIGNAL` require that a handler be active. Otherwise, `RESIGNAL` is illegal and a `RESIGNAL when handler not active` error occurs.

No special privileges are required to execute the `RESIGNAL` statement.

To retrieve information from the diagnostics area, use the `GET DIAGNOSTICS` statement (see [Section 13.6.7.3, “GET DIAGNOSTICS Syntax”](#)). For information about the diagnostics area, see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

For *condition\_value* and *signal\_information\_item*, the definitions and rules are the same for `RESIGNAL` as for `SIGNAL` (see [Section 13.6.7.5, “SIGNAL Syntax”](#)).

The `RESIGNAL` statement takes *condition\_value* and `SET` clauses, both of which are optional. This leads to several possible uses:

- `RESIGNAL` alone:

```
RESIGNAL;
```

- `RESIGNAL` with new signal information:

```
RESIGNAL SET signal_information_item [, signal_information_item] ...;
```

- `RESIGNAL` with a condition value and possibly new signal information:

```
RESIGNAL condition_value
  [SET signal_information_item [, signal_information_item] ...];
```

These use cases all cause changes to the diagnostics and condition areas:

- A diagnostics area contains one or more condition areas.
- A condition area contains condition information items, such as the `SQLSTATE` value, `MYSQL_ERRNO`, or `MESSAGE_TEXT`.

The maximum number of condition areas in a diagnostics area is determined by the value of the `max_error_count` [504] system variable. See [Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”](#).

#### 13.6.7.4.1. `RESIGNAL` Alone

A simple `RESIGNAL` alone means “pass on the error with no change.” It restores the last diagnostics area and makes it the current diagnostics area. That is, it “pops” the diagnostics area stack.

Within a condition handler that catches a condition, one use for `RESIGNAL` alone is to perform some other actions, and then pass on without change the original condition information (the information that existed before entry into the handler).

Example:

```
DROP TABLE IF EXISTS xx;
delimiter //
CREATE PROCEDURE p ()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION
  BEGIN
    SET @error_count = @error_count + 1;
    IF @a = 0 THEN RESIGNAL; END IF;
  END;
  DROP TABLE xx;
END//
delimiter ;
SET @error_count = 0;
SET @a = 0;
CALL p();
```

Suppose that the `DROP TABLE xx` statement fails. The diagnostics area stack looks like this:

```
DA 1. ERROR 1051 (42S02): Unknown table 'xx'
```

Then execution enters the `EXIT` handler. It starts by pushing the top of the diagnostics area stack, which now looks like this:

```
DA 1. ERROR 1051 (42S02): Unknown table 'xx'
DA 2. ERROR 1051 (42S02): Unknown table 'xx'
```

At this point, the contents of the first (current) and second (stacked) diagnostics areas are the same. The first diagnostics area may be modified by statements executing subsequently within the handler.

Usually a procedure statement clears the first diagnostics area. `BEGIN` is an exception, it does not clear, it does nothing. `SET` is not an exception, it clears, performs the operation, and then produces a result of “success.” The diagnostics area stack now looks like this:

```
DA 1. ERROR 0000 (00000): Successful operation
DA 2. ERROR 1051 (42S02): Unknown table 'xx'
```

At this point, if `@a = 0`, `RESIGNAL` pops the diagnostics area stack, which now looks like this:

```
DA 1. ERROR 1051 (42S02): Unknown table 'xx'
```

And that is what the caller sees.

If `@a` is not 0, the handler simply ends, which means that there is no more use for the current diagnostics area (it has been “handled”), so it can be thrown away, causing the stacked diagnostics area to become the current diagnostics area again. The diagnostics area stack looks like this:

```
DA 1. ERROR 0000 (00000): Successful operation
```

The details make it look complex, but the end result is quite useful: Handlers can execute without destroying information about the condition that caused activation of the handler.

#### 13.6.7.4.2. `RESIGNAL` with New Signal Information

`RESIGNAL` with a `SET` clause provides new signal information, so the statement means “pass on the error with changes”:

```
RESIGNAL SET signal_information_item [, signal_information_item] ...;
```

As with `RESIGNAL` alone, the idea is to pop the diagnostics area stack so that the original information will go out. Unlike `RESIGNAL` alone, anything specified in the `SET` clause changes.

Example:

```
DROP TABLE IF EXISTS xx;
delimiter //
CREATE PROCEDURE p ()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION
  BEGIN
    SET @error_count = @error_count + 1;
    IF @a = 0 THEN RESIGNAL SET MYSQL_ERRNO = 5; END IF;
  END;
  DROP TABLE xx;
END//
delimiter ;
SET @error_count = 0;
SET @a = 0;
CALL p();
```

Remember from the previous discussion that `RESIGNAL` alone results in a diagnostics area stack like this:

```
DA 1. ERROR 1051 (42S02): Unknown table 'xx'
```

The `RESIGNAL SET MYSQL_ERRNO = 5` statement results in this stack instead:

```
DA 1. ERROR 5 (42S02): Unknown table 'xx'
```



In other words, it changes the error number, and nothing else.

The `RESIGNAL` statement can change any or all of the signal information items, making the first condition area of the diagnostics area look quite different.

### 13.6.7.4.3. `RESIGNAL` with a Condition Value and Optional New Signal Information

`RESIGNAL` with a condition value means “push a condition into the current diagnostics area.” If the `SET` clause is present, it also changes the error information.

```
RESIGNAL condition_value
  [SET signal_information_item [, signal_information_item] ...];
```

This form of `RESIGNAL` restores the last diagnostics area and makes it the current diagnostics area. That is, it “pops” the diagnostics area stack, which is the same as what a simple `RESIGNAL` alone would do. However, it also changes the diagnostics area depending on the condition value or signal information.

Example:

```
DROP TABLE IF EXISTS xx;
delimiter //
CREATE PROCEDURE p ()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION
  BEGIN
    SET @error_count = @error_count + 1;
    IF @a = 0 THEN RESIGNAL SQLSTATE '45000' SET MYSQL_ERRNO=5; END IF;
  END;
  DROP TABLE xx;
END//
delimiter ;
SET @error_count = 0;
SET @a = 0;
SET @max_error_count = 2;
CALL p();
SHOW ERRORS;
```

This is similar to the previous example, and the effects are the same, except that if `RESIGNAL` happens the current condition area looks different at the end. (The reason the condition is added rather than replaced is the use of a condition value.)

The `RESIGNAL` statement includes a condition value (`SQLSTATE '45000'`), so it “pushes” a new condition area, resulting in a diagnostics area stack that looks like this:

```
DA 1. (condition 1) ERROR 5 (45000) Unknown table 'xx'
      (condition 2) ERROR 1051 (42S02): Unknown table 'xx'
```

The result of `CALL p()` and `SHOW ERRORS` for this example is:

```
mysql> CALL p();
ERROR 5 (45000): Unknown table 'xx'
mysql> SHOW ERRORS;
+-----+-----+-----+
| Level | Code | Message                               |
+-----+-----+-----+
| Error | 5    | Unknown table 'xx'                   |
| Error | 1051 | Unknown table 'xx'                   |
+-----+-----+-----+
```

### 13.6.7.4.4. `RESIGNAL` Requires an Active Handler

All forms of `RESIGNAL` require that a handler be active. Otherwise, `RESIGNAL` is illegal and a `RESIGNAL when handler not active` error occurs. For example:

```
mysql> CREATE PROCEDURE p () RESIGNAL;
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> CALL p();
ERROR 1739 (0K000): RESIGNAL when handler not active
```

Here is a more difficult example:

```
delimiter //
CREATE FUNCTION f () RETURNS INT
BEGIN
  RESIGNAL;
  RETURN 5;
END//
CREATE PROCEDURE p ()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION SET @a=f();
  SIGNAL SQLSTATE '55555';
END//
delimiter ;
CALL p();
```

At the time the `RESIGNAL` executes, there is a handler, even though the `RESIGNAL` is not defined inside the handler.

A statement such as the one following may appear bizarre because `RESIGNAL` apparently is not in a handler:

```
CREATE TRIGGER t_bi BEFORE INSERT ON t FOR EACH ROW RESIGNAL;
```

But it does not matter. `RESIGNAL` does not have to be technically “in” (that is, contained in), a handler declaration. The requirement is that a handler must be active. If that is true when the trigger activates, the `RESIGNAL` is legal.

### 13.6.7.5. SIGNAL Syntax

```
SIGNAL condition_value
  [SET signal_information_item
  [, signal_information_item] ...]

condition_value:
  SQLSTATE [VALUE] sqlstate_value
  | condition_name

signal_information_item:
  condition_information_item_name = simple_value_specification

condition_information_item_name:
  CLASS_ORIGIN
  | SUBCLASS_ORIGIN
  | MESSAGE_TEXT
  | MYSQL_ERRNO
  | CONSTRAINT_CATALOG
  | CONSTRAINT_SCHEMA
  | CONSTRAINT_NAME
  | CATALOG_NAME
  | SCHEMA_NAME
  | TABLE_NAME
  | COLUMN_NAME
  | CURSOR_NAME

condition_name, simple_value_specification:
  (see following discussion)
```

`SIGNAL` is the way to “return” an error. `SIGNAL` provides error information to a handler, to an outer portion of the application, or to the client. Also, it provides control over the error's characteristics (error number, `SQLSTATE` value, message). Without `SIGNAL`, it is necessary to resort to workarounds such as deliberately referring to a nonexistent table to cause a routine to return an error.

No special privileges are required to execute the `SIGNAL` statement.

To retrieve information from the diagnostics area, use the `GET DIAGNOSTICS` statement (see [Section 13.6.7.3, “GET DIAGNOSTICS Syntax”](#)). For information about the diagnostics area, see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

The *condition\_value* in a `SIGNAL` statement indicates the error value to be returned. It can be an `SQLSTATE` value (a 5-character string literal) or a *condition\_name* that refers to a named condition previously defined with `DECLARE ... CONDITION` (see [Section 13.6.7.1, “DECLARE ... CONDITION Syntax”](#)).

An `SQLSTATE` value can indicate errors, warnings, or “not found.” The first two characters of the value indicate its error class, as discussed in [Section 13.6.7.5.1, “Signal Condition Information Items”](#). Some signal values cause statement termination; see [Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”](#).

The `SQLSTATE` value for a `SIGNAL` statement should not start with `'00'` because such values indicate success and are not valid for signaling an error. This is true whether the `SQLSTATE` value is specified directly in the `SIGNAL` statement or in a named condition referred to in the statement. If the value is invalid, a `Bad SQLSTATE` error occurs.

To signal a generic `SQLSTATE` value, use `'45000'`, which means “unhandled user-defined exception.”

The `SIGNAL` statement optionally includes a `SET` clause that contains multiple signal items, in a comma-separated list of *condition\_information\_item\_name* = *simple\_value\_specification* assignments.

Each *condition\_information\_item\_name* may be specified only once in the `SET` clause. Otherwise, a `Duplicate condition information item` error occurs.

Valid *simple\_value\_specification* designators can be specified using stored procedure or function parameters, stored program local variables declared with `DECLARE`, user-defined variables, system variables, or literals. A character literal may include a `_charset` introducer.

For information about permissible *condition\_information\_item\_name* values, see [Section 13.6.7.5.1, “Signal Condition Information Items”](#).

The following procedure signals an error or warning depending on the value of `pval`, its input parameter:

```
CREATE PROCEDURE p (pval INT)
BEGIN
  DECLARE specialty CONDITION FOR SQLSTATE '45000';
  IF pval = 0 THEN
    SIGNAL SQLSTATE '01000';
  ELSEIF pval = 1 THEN
    SIGNAL SQLSTATE '45000'
    SET MESSAGE_TEXT = 'An error occurred';
  ELSEIF pval = 2 THEN
    SIGNAL specialty
    SET MESSAGE_TEXT = 'An error occurred';
  ELSE
    SIGNAL SQLSTATE '01000'
    SET MESSAGE_TEXT = 'A warning occurred', MYSQL_ERRNO = 1000;
    SIGNAL SQLSTATE '45000'
    SET MESSAGE_TEXT = 'An error occurred', MYSQL_ERRNO = 1001;
  END IF;
END;
```

If `pval` is 0, `p()` signals a warning because `SQLSTATE` values that begin with `'01'` are signals in the warning class. The warning does not terminate the procedure, and can be seen with `SHOW WARNINGS` after the procedure returns.

If `pval` is 1, `p()` signals an error and sets the `MESSAGE_TEXT` condition information item. The error terminates the procedure, and the text is returned with the error information.

If `pval` is 2, the same error is signaled, although the `SQLSTATE` value is specified using a named condition in this case.

If `pval` is anything else, `p()` first signals a warning and sets the message text and error number condition information items. This warning does not terminate the procedure, so execution continues and `p()` then signals an error. The error does terminate the procedure. The message text and error number set by the warning are replaced by the values set by the error, which are returned with the error information.

`SIGNAL` is typically used within stored programs, but it is a MySQL extension that it is permitted outside that context. For example, if you invoke the `mysql` client program, you can enter any of these statements at the prompt:

```
mysql> SIGNAL SQLSTATE '77777';
mysql> CREATE TRIGGER t_bi BEFORE INSERT ON t
-> FOR EACH ROW SIGNAL SQLSTATE '77777';
mysql> CREATE EVENT e ON SCHEDULE EVERY 1 SECOND
-> DO SIGNAL SQLSTATE '77777';
```

`SIGNAL` executes according to the following rules:

If the `SIGNAL` statement indicates a particular `SQLSTATE` value, that value is used to signal the condition specified. Example:

```
CREATE PROCEDURE p (divisor INT)
BEGIN
  IF divisor = 0 THEN
    SIGNAL SQLSTATE '22012';
  END IF;
END;
```

If the `SIGNAL` statement uses a named condition, the condition must be declared in some scope that applies to the `SIGNAL` statement, and must be defined using an `SQLSTATE` value, not a MySQL error number. Example:

```
CREATE PROCEDURE p (divisor INT)
BEGIN
  DECLARE divide_by_zero CONDITION FOR SQLSTATE '22012';
  IF divisor = 0 THEN
    SIGNAL divide_by_zero;
  END IF;
END;
```

If the named condition does not exist in the scope of the `SIGNAL` statement, an `Undefined CONDITION` error occurs.

If `SIGNAL` refers to a named condition that is defined with a MySQL error number rather than an `SQLSTATE` value, a `SIGNAL/RESIGNAL` can only use a `CONDITION` defined with `SQLSTATE` error occurs. The following statements cause that error because the named condition is associated with a MySQL error number:

```
DECLARE no_such_table CONDITION FOR 1051;
SIGNAL no_such_table;
```

If a condition with a given name is declared multiple times in different scopes, the declaration with the most local scope applies. Consider the following procedure:

```
CREATE PROCEDURE p (divisor INT)
BEGIN
  DECLARE my_error CONDITION FOR SQLSTATE '45000';
  IF divisor = 0 THEN
    BEGIN
      DECLARE my_error CONDITION FOR SQLSTATE '22012';
      SIGNAL my_error;
    END;
END;
```

```
END IF;
  SIGNAL my_error;
END;
```

If `divisor` is 0, the first `SIGNAL` statement executes. The innermost `my_error` condition declaration applies, raising `SQLSTATE '22012'`.

If `divisor` is not 0, the second `SIGNAL` statement executes. The outermost `my_error` condition declaration applies, raising `SQLSTATE '45000'`.

For information about how the server chooses handlers when a condition occurs, see [Section 13.6.7.6, “Scope Rules for Handlers”](#).

Signals can be raised within exception handlers:

```
CREATE PROCEDURE p ()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION
  BEGIN
    SIGNAL SQLSTATE VALUE '99999'
      SET MESSAGE_TEXT = 'An error occurred';
  END;
  DROP TABLE no_such_table;
END;
```

`CALL p()` reaches the `DROP TABLE` statement. There is no table named `no_such_table`, so the error handler is activated. The error handler destroys the original error (“no such table”) and makes a new error with `SQLSTATE '99999'` and message `An error occurred`.

### 13.6.7.5.1. Signal Condition Information Items

The following table lists the names of diagnostics area condition information items that can be set in a `SIGNAL` (or `RESIGNAL`) statement. All items are standard SQL except `MYSQL_ERRNO`, which is a MySQL extension. For more information about these items see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

Item Name	Definition
CLASS_ORIGIN	VARCHAR(64)
SUBCLASS_ORIGIN	VARCHAR(64)
CONSTRAINT_CATALOG	VARCHAR(64)
CONSTRAINT_SCHEMA	VARCHAR(64)
CONSTRAINT_NAME	VARCHAR(64)
CATALOG_NAME	VARCHAR(64)
SCHEMA_NAME	VARCHAR(64)
TABLE_NAME	VARCHAR(64)
COLUMN_NAME	VARCHAR(64)
CURSOR_NAME	VARCHAR(64)
MESSAGE_TEXT	VARCHAR(128)
MYSQL_ERRNO	SMALLINT UNSIGNED

The character set for character items is UTF-8.

It is illegal to assign `NULL` to a condition information item in a `SIGNAL` statement.

A `SIGNAL` statement always specifies an `SQLSTATE` value, either directly, or indirectly by referring to a named condition defined with an `SQLSTATE` value. The first two characters of an `SQLSTATE` value are its class, and the class determines the default value for the condition information items:

- Class = '00' (success)

Illegal. `SQLSTATE` values that begin with '00' indicate success and are not valid for `SIGNAL`.

- Class = '01' (warning)

```
MESSAGE_TEXT = 'Unhandled user-defined warning';
MYSQL_ERRNO = ER_SIGNAL_WARN
```

- Class = '02' (not found)

```
MESSAGE_TEXT = 'Unhandled user-defined not found';
MYSQL_ERRNO = ER_SIGNAL_NOT_FOUND
```

- Class > '02' (exception)

```
MESSAGE_TEXT = 'Unhandled user-defined exception';
MYSQL_ERRNO = ER_SIGNAL_EXCEPTION
```

For legal classes, the other condition information items are set as follows:

```
CLASS_ORIGIN = SUBCLASS_ORIGIN = '';
CONSTRAINT_CATALOG = CONSTRAINT_SCHEMA = CONSTRAINT_NAME = '';
CATALOG_NAME = SCHEMA_NAME = TABLE_NAME = COLUMN_NAME = '';
CURSOR_NAME = '';
```

The error values that are accessible after `SIGNAL` executes are the `SQLSTATE` value raised by the `SIGNAL` statement and the `MESSAGE_TEXT` and `MYSQL_ERRNO` items. These values are available from the C API:

- `SQLSTATE` value: Call `mysql_sqlstate()`
- `MYSQL_ERRNO` value: Call `mysql_errno()`
- `MESSAGE_TEXT` value: Call `mysql_error()`

From SQL, the output from `SHOW WARNINGS` and `SHOW ERRORS` indicates the `MYSQL_ERRNO` and `MESSAGE_TEXT` values in the `Code` and `Message` columns.

To retrieve information from the diagnostics area, use the `GET DIAGNOSTICS` statement (see [Section 13.6.7.3, “GET DIAGNOSTICS Syntax”](#)). For information about the diagnostics area, see [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

### 13.6.7.5.2. Effect of Signals on Handlers, Cursors, and Statements

Signals have different effects on statement execution depending on the signal class. The class determines how severe an error is. MySQL ignores the value of the `sql_mode` [545] system variable; in particular, strict SQL mode does not matter. MySQL also ignores `IGNORE`: The intent of `SIGNAL` is to raise a user-generated error explicitly, so a signal is never ignored.

In the following descriptions, “unhandled” means that no handler for the signaled `SQLSTATE` value has been defined with `DECLARE ... HANDLER`.

- Class = '00' (success)

Illegal. `SQLSTATE` values that begin with '00' indicate success and are not valid for `SIGNAL`.

- Class = '01' (warning)

The value of the `warning_count` [564] system variable goes up. `SHOW WARNINGS` shows the signal. `SQLWARNING` handlers catch the signal. If the signal is unhandled in a function, statements do not end.

- Class = '02' (not found)

`NOT FOUND` handlers catch the signal. There is no effect on cursors. If the signal is unhandled in a function, statements end.

- Class > '02' (exception)

`SQLWARNING` handlers catch the signal. If the signal is unhandled in a function, statements end.

- `Class = '40'`

Treated as an ordinary exception.

Example:

```
mysql> delimiter //
mysql> CREATE FUNCTION f () RETURNS INT
-> BEGIN
->   SIGNAL SQLSTATE '01234'; -- signal a warning
->   RETURN 5;
-> END//
mysql> delimiter ;
mysql> CREATE TABLE t (s1 INT);
mysql> INSERT INTO t VALUES (f());
```

The result is that a row containing 5 is inserted into table `t`. The warning that is signaled can be viewed with `SHOW WARNINGS`.

### 13.6.7.6. Scope Rules for Handlers

A stored program may include handlers to be invoked when certain conditions occur within the program. The applicability of each handler depends on its location within the program definition and on the condition or conditions that it handles:

- A handler declared in a `BEGIN ... END` block is in scope only for the SQL statements following the handler declarations in the block. If the handler itself raises a condition, it cannot handle that condition, nor can any other handlers declared in the block. In the following example, handlers `H1` and `H2` are in scope for conditions raised by statements `stmt1` and `stmt2`. But neither `H1` nor `H2` are in scope for conditions raised in the body of `H1` or `H2`.

```
BEGIN -- outer block
  DECLARE EXIT HANDLER FOR ...; -- handler H1
  DECLARE EXIT HANDLER FOR ...; -- handler H2
  stmt1;
  stmt2;
END;
```

- A handler is in scope only for the block in which it is declared, and cannot be activated for conditions occurring outside that block. In the following example, handler `H1` is in scope for `stmt1` in the inner block, but not for `stmt2` in the outer block:

```
BEGIN -- outer block
  BEGIN -- inner block
    DECLARE EXIT HANDLER FOR ...; -- handler H1
    stmt1;
  END;
  stmt2;
END;
```

- A handler can be specific or general. A specific handler is for a MySQL error code, `SQLSTATE` value, or condition name. A general handler is for a condition in the `SQLWARNING`, `SQLWARNING`, or `NOT FOUND` class. Condition specificity is related to condition precedence, as described later.

Multiple handlers can be declared in different scopes and with different specificities. For example, there might be a specific MySQL error code handler in an outer block, and a general `SQLWARNING` handler in an inner block. Or there might be handlers for a specific MySQL error code and the general `SQLWARNING` class in the same block.

Whether a handler is activated depends not only on its own scope and condition value, but on what other handlers are present. When a condition occurs in a stored program, the server searches for applicable handlers in the current scope (current `BEGIN ... END` block). If there are no applicable



handlers, the search continues outward with the handlers in each successive containing scope (block). When the server finds one or more applicable handlers at a given scope, it chooses among them based on condition precedence:

- A MySQL error code handler takes precedence over an `SQLSTATE` value handler.
- An `SQLSTATE` value handler takes precedence over general `SQLWARNING`, `SQLEXCEPTION`, or `NOT FOUND` handlers.
- An `SQLEXCEPTION` handler takes precedence over an `SQLWARNING` handler.
- The precedence of `NOT FOUND` depends on how the condition is raised:
  - Normally, a condition in the `NOT FOUND` class can be handled by an `SQLWARNING` or `NOT FOUND` handler, with the `SQLWARNING` handler taking precedence if both are present. Normal occurrence of `NOT FOUND` takes place when a cursor used to fetch a set of rows reaches the end of the data set, or for instances of `SELECT ... INTO var_list` such that the `WHERE` clause finds no rows.
  - If a `NOT FOUND` condition is raised by a `SIGNAL` (or `RESIGNAL`) statement, the condition can be handled by a `NOT FOUND` handler but not an `SQLWARNING` handler.
- It is possible to have several applicable handlers with the same precedence. For example, a statement could generate multiple warnings with different error codes, for each of which an error-specific handler exists. In this case, the choice of which handler the server activates is indeterminate, and may change depending on the circumstances under which the condition occurs.

One implication of the handler selection rules is that if multiple applicable handlers occur in different scopes, handlers with the most local scope take precedence over handlers in outer scopes, even over those for more specific conditions.

If there is no appropriate handler when a condition occurs, the action taken depends on the class of the condition:

- For `SQLEXCEPTION` conditions, the stored program terminates at the statement that raised the condition, as if there were an `EXIT` handler. If the program was called by another stored program, the calling program handles the condition using the handler selection rules applied to its own handlers.
- For `SQLWARNING` conditions, the program continues executing, as if there were a `CONTINUE` handler.
- For `NOT FOUND` conditions, if the condition was raised normally, the action is `CONTINUE`. If it was raised by `SIGNAL` or `RESIGNAL`, the action is `EXIT`.

The following examples demonstrate how MySQL applies the handler selection rules.

This procedure contains two handlers, one for the specific `SQLSTATE` value (`'42S02'`) that occurs for attempts to drop a nonexistent table, and one for the general `SQLEXCEPTION` class:

```
CREATE PROCEDURE p1()
BEGIN
  DECLARE CONTINUE HANDLER FOR SQLSTATE '42S02'
    SELECT 'SQLSTATE handler was activated' AS msg;
  DECLARE CONTINUE HANDLER FOR SQLEXCEPTION
    SELECT 'SQLEXCEPTION handler was activated' AS msg;

  DROP TABLE test.t;
END;
```

Both handlers are declared in the same block and have the same scope. However, `SQLSTATE` handlers take precedence over `SQLEXCEPTION` handlers, so if the table `t` is nonexistent, the `DROP TABLE` statement raises a condition that activates the `SQLSTATE` handler:

```
mysql> CALL p1();
```

```
+-----+
| msg           |
+-----+
| SQLSTATE handler was activated |
+-----+
```

This procedure contains the same two handlers. But this time, the `DROP TABLE` statement and `SQLEXCEPTION` handler are in an inner block relative to the `SQLSTATE` handler:

```
CREATE PROCEDURE p2()
BEGIN -- outer block
    DECLARE CONTINUE HANDLER FOR SQLSTATE '42S02'
        SELECT 'SQLSTATE handler was activated' AS msg;
    BEGIN -- inner block
        DECLARE CONTINUE HANDLER FOR SQLEXCEPTION
            SELECT 'SQLEXCEPTION handler was activated' AS msg;

        DROP TABLE test.t; -- occurs within inner block
    END;
END;
```

In this case, the handler that is more local to where the condition occurs takes precedence. The `SQLEXCEPTION` handler activates, even though it is more general than the `SQLSTATE` handler:

```
mysql> CALL p2();
+-----+
| msg           |
+-----+
| SQLEXCEPTION handler was activated |
+-----+
```

In this procedure, one of the handlers is declared in a block inner to the scope of the `DROP TABLE` statement:

```
CREATE PROCEDURE p3()
BEGIN -- outer block
    DECLARE CONTINUE HANDLER FOR SQLEXCEPTION
        SELECT 'SQLEXCEPTION handler was activated' AS msg;
    BEGIN -- inner block
        DECLARE CONTINUE HANDLER FOR SQLSTATE '42S02'
            SELECT 'SQLSTATE handler was activated' AS msg;
    END;

    DROP TABLE test.t; -- occurs within outer block
END;
```

Only the `SQLEXCEPTION` handler applies because the other one is not in scope for the condition raised by the `DROP TABLE`:

```
mysql> CALL p3();
+-----+
| msg           |
+-----+
| SQLEXCEPTION handler was activated |
+-----+
```

In this procedure, both handlers are declared in a block inner to the scope of the `DROP TABLE` statement:

```
CREATE PROCEDURE p4()
BEGIN -- outer block
    BEGIN -- inner block
        DECLARE CONTINUE HANDLER FOR SQLEXCEPTION
            SELECT 'SQLEXCEPTION handler was activated' AS msg;
        DECLARE CONTINUE HANDLER FOR SQLSTATE '42S02'
            SELECT 'SQLSTATE handler was activated' AS msg;
    END;

    DROP TABLE test.t; -- occurs within outer block
END;
```

```
END;
```

Neither handler applies because they are not in scope for the `DROP TABLE`. The condition raised by the statement goes unhandled and terminates the procedure with an error:

```
mysql> CALL p4();
ERROR 1051 (42S02): Unknown table 'test.t'
```

### 13.6.7.7. The MySQL Diagnostics Area

SQL statements produce diagnostic information that populates the diagnostics area. This section describes the structure of the diagnostics area in MySQL and how it differs from standard SQL. It also discusses the information items recognized by MySQL and how statements clear and set the diagnostics area.

#### 13.6.7.7.1. Diagnostics Area Structure

The diagnostics area contains two kinds of information:

- Statement information, such as the affected-rows count or the number of conditions that occurred.
- Condition information, such as the error code and message. If a statement raises multiple conditions, this part of the diagnostics area has a condition area for each one. If a statement raises no conditions, this part of the diagnostics area is empty.

For a statement that produces three conditions, the diagnostics area contains statement and condition information like this:

```
Statement information:
  row count
  ... other statement information items ...
Condition area list:
  Condition area 1:
    error code for condition 1
    error message for condition 1
    ... other condition information items ...
  Condition area 2:
    error code for condition 2:
    error message for condition 2
    ... other condition information items ...
  Condition area 3:
    error code for condition 3
    error message for condition 3
    ... other condition information items ...
```

Standard SQL has a diagnostics area stack, containing a diagnostics area for each nested execution context. Standard SQL syntax includes `GET STACKED DIAGNOSTICS` for referring to the second diagnostics area. MySQL does not support the `STACKED` keyword because there is a single diagnostics area containing information from the most recent statement that wrote to it.

#### 13.6.7.7.2. Diagnostics Area Information Items

The diagnostics area contains statement and condition information items. Numeric items are integers. The character set for character items is UTF-8. No item can be `NULL`. If a statement or condition item is not set by a statement that populates the diagnostics area, its value will be 0 or the empty string, depending on the item data type.

The statement information part of the diagnostics area contains these items:

- **NUMBER**: An integer indicating the number of condition areas that have information.
- **ROW\_COUNT**: An integer indicating the number of rows affected by the statement. `ROW_COUNT` has the same value as the `ROW_COUNT( ) [1272]` function (see [Section 12.14, “Information Functions”](#)).

The condition information part of the diagnostics area contains a condition area for each condition. Condition areas are numbered from 1 to the value of the `NUMBER` statement condition item. If `NUMBER` is 0, there are no condition areas.

Each condition area contains the items in the following list. All items are standard SQL except `MYSQL_ERRNO`, which is a MySQL extension. The definitions apply for conditions generated other than by a signal (that is, by a `SIGNAL` or `RESIGNAL` statement). For nonsignal conditions, MySQL populates only those condition items not described as always empty. The effects of signals on the condition area are described later.

- `CLASS_ORIGIN`: A string containing the class of the `RETURNED_SQLSTATE` value. If the `RETURNED_SQLSTATE` value begins with a class value defined in SQL standards document ISO 9075-2 (section 24.1, `SQLSTATE`), `CLASS_ORIGIN` is `'ISO 9075'`. Otherwise, `CLASS_ORIGIN` is `'MySQL'`.
- `SUBCLASS_ORIGIN`: A string containing the subclass of the `RETURNED_SQLSTATE` value. If `CLASS_ORIGIN` is `'ISO 9075'` or `RETURNED_SQLSTATE` ends with `'000'`, `SUBCLASS_ORIGIN` is `'ISO 9075'`. Otherwise, `SUBCLASS_ORIGIN` is `'MySQL'`.
- `RETURNED_SQLSTATE`: A string that indicates the `SQLSTATE` value for the condition.
- `MESSAGE_TEXT`: A string that indicates the error message for the condition.
- `MYSQL_ERRNO`: An integer that indicates the MySQL error code for the condition.
- `CONSTRAINT_CATALOG`, `CONSTRAINT_SCHEMA`, `CONSTRAINT_NAME`: Strings that indicate the catalog, schema, and name for a violated constraint. They are always empty.
- `CATALOG_NAME`, `SCHEMA_NAME`, `TABLE_NAME`, `COLUMN_NAME`: Strings that indicate the catalog, schema, table, and column related to the condition. They are always empty.
- `CURSOR_NAME`: A string that indicates the cursor name. This is always empty.

For the `RETURNED_SQLSTATE`, `MESSAGE_TEXT`, and `MYSQL_ERRNO` values for particular errors, see [Section C.3, “Server Error Codes and Messages”](#).

If a `SIGNAL` (or `RESIGNAL`) statement populates the diagnostics area, its `SET` clause can assign to any condition information item except `RETURNED_SQLSTATE` any value that is legal for the item data type. `SIGNAL` also sets the `RETURNED_SQLSTATE` value, but not directly in its `SET` clause. That value comes from the `SIGNAL` statement `SQLSTATE` argument.

`SIGNAL` also sets statement information items. It sets `NUMBER` to 1, and `ROW_COUNT` to `-1` for errors and 0 otherwise.

### 13.6.7.7.3. How the Diagnostics Area is Populated

Most SQL statements populate the diagnostics area automatically, and its contents can be set explicitly with the `SIGNAL` and `RESIGNAL` statements. The diagnostics area can be examined with `GET DIAGNOSTICS` to extract specific items, or with `SHOW WARNINGS` or `SHOW ERRORS` to see all conditions or all errors.

SQL statements clear and set the diagnostics area as follows:

- When the server starts executing a statement after parsing it, it clears the diagnostics area for nondiagnostic statements that use tables. Diagnostic statements are `SHOW WARNINGS`, `SHOW ERRORS`, and `GET DIAGNOSTICS`.
- If a statement raises a condition, the diagnostics area is cleared of conditions that belong to earlier statements. The exception is that conditions raised by `GET DIAGNOSTICS` and `RESIGNAL` are added to the diagnostics area without clearing it.

Thus, even a statement that does not normally clear the diagnostics area when it begins executing clears it if the statement raises a condition.

The following example shows the effect of various statements on the diagnostics area, using `SHOW WARNINGS` to display information about all conditions stored there.

This `DROP TABLE` statement uses a table, so it clears the diagnostics area and populates it when the condition occurs:

```
mysql> DROP TABLE IF EXISTS test.no_such_table;
Query OK, 0 rows affected, 1 warning (0.01 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                               |
+-----+-----+-----+
| Note  | 1051 | Unknown table 'test.no_such_table' |
+-----+-----+-----+
1 row in set (0.00 sec)
```

This `SET` statement does not use tables, and `SET` does not generate warnings, so it leaves the diagnostics area unchanged:

```
mysql> SET @x = 1;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                               |
+-----+-----+-----+
| Note  | 1051 | Unknown table 'test.no_such_table' |
+-----+-----+-----+
1 row in set (0.00 sec)
```

This `SET` statement generates an error, so it clears and populates the diagnostics area:

```
mysql> SET @x = @@x;
ERROR 1193 (HY000): Unknown system variable 'x'

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                               |
+-----+-----+-----+
| Error | 1193 | Unknown system variable 'x' |
+-----+-----+-----+
1 row in set (0.00 sec)
```

The previous `SET` statement produced a single condition, so 1 is the only valid condition number for `GET DIAGNOSTICS` at this point. The following statement uses a condition number of 2, which produces a warning that is added to the diagnostics area without clearing it:

```
mysql> GET DIAGNOSTICS CONDITION 2 @p = MESSAGE_TEXT;
Query OK, 0 rows affected, 1 warning (0.00 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                               |
+-----+-----+-----+
| Error | 1193 | Unknown system variable 'xx' |
| Error | 1753 | Invalid condition number |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

Now there are two conditions in the diagnostics area, so the same `GET DIAGNOSTICS` statement succeeds:

```
mysql> GET DIAGNOSTICS CONDITION 2 @p = MESSAGE_TEXT;
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT @p;
+-----+
| @p    |
+-----+
| Invalid condition number |
+-----+
1 row in set (0.01 sec)
```

#### 13.6.7.7.4. Diagnostics Area-Related System Variables

Certain system variables control or are related to some aspects of the diagnostics area:

- `max_error_count` [504] controls the number of condition areas in the diagnostics area. If more conditions than this occur, MySQL silently discards information for the excess conditions. (Conditions added by `RESIGNAL` are always added, with older conditions being discarded as necessary to make room.)
- `warning_count` [564] indicates the number of conditions that occurred. This includes errors, warnings, and notes. Normally, `NUMBER` and `warning_count` [564] are the same. However, as the number of conditions generated exceeds `max_error_count` [504], the value of `warning_count` [564] continues to rise whereas `NUMBER` remains capped at `max_error_count` [504] because no additional conditions are stored in the diagnostics area.
- `error_count` [480] indicates the number of errors that occurred. This value includes “not found” and exception conditions, but excludes warnings and notes. Like `warning_count` [564], its value can exceed `max_error_count` [504].
- If the `sql_notes` [546] system variable is set to 0, notes are not recorded and do not increment `warning_count` [564].

Example: If `max_error_count` [504] is 10, the diagnostics area can contain a maximum of 10 condition areas. Suppose that a statement raises 20 conditions, 12 of which are errors. In that case, the diagnostics area contains the first 10 conditions, `NUMBER` is 10, `warning_count` [564] is 20, and `error_count` [480] is 12.

Changes to the value of `max_error_count` [504] have no effect until the next attempt to modify the diagnostics area. If the diagnostics area contains 10 condition areas and `max_error_count` [504] is set to 5, that has no immediate effect on the size or content of the diagnostics area.

Before MySQL 5.6, statement information items are not available directly. `ROW_COUNT` can be obtained by calling the `ROW_COUNT()` function. `NUMBER` is approximated by the value of the `warning_count` system variable. However, whereas `NUMBER` is capped to the value of `max_error_count`, `warning_count` is not.

## 13.7. Database Administration Statements

### 13.7.1. Account Management Statements

MySQL account information is stored in the tables of the `mysql` database. This database and the access control system are discussed extensively in [Chapter 5, MySQL Server Administration](#), which you should consult for additional details.



#### Important

Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

### 13.7.1.1. ALTER USER Syntax

```
ALTER USER user_specification
    [, user_specification] ...
```

```
user_specification:
    user PASSWORD EXPIRE
```

The `ALTER USER` statement alters MySQL accounts. To use it, you must have the global `CREATE USER` [754] privilege or the `INSERT` [755] privilege for the `mysql` database. This statement was added in MySQL 5.6.6.

For each account, `ALTER USER` expires its password. For example:

```
ALTER USER 'jeffrey'@'localhost' PASSWORD EXPIRE;
```

The account name uses the format described in [Section 6.2.3, “Specifying Account Names”](#). If you specify only the user name part of the account name, a host name part of `'%'` is used.

Password expiration for an account affects the corresponding row of the `mysql.user` table: The server sets the `password_expired` column to `'Y'`.



#### Warning

In MySQL 5.6.6, `ALTER USER` also sets the `Password` column to the empty string, so do not use this statement until 5.6.7.

After an account's password has been expired, all operations performed in subsequent connections to the server using the account result in an error until the user issues a `SET PASSWORD` statement to establish a new account password:

```
mysql> SELECT 1;
ERROR 1820 (HY000): You must SET PASSWORD before executing this statement

mysql> SET PASSWORD = PASSWORD('new_password');
Query OK, 0 rows affected (0.01 sec)

mysql> SELECT 1;
+----+
| 1 |
+----+
| 1 |
+----+
1 row in set (0.00 sec)
```

As of MySQL 5.6.8, this restricted mode of operation also permits `SET` statements, which is useful if the account password uses a hashing format that requires `old_passwords` [517] to be set to a value different from its default.

It is also possible for an administrative user to reset the account password, but any existing connections for the account remain under restriction. Clients using the account must disconnect and reconnect before statements can be executed successfully.

It is possible after password expiration to “reset” a password by using `SET PASSWORD` to set it to its current value. As a matter of good policy, it is preferable to choose a different password.

### 13.7.1.2. CREATE USER Syntax

```
CREATE USER user_specification
    [, user_specification] ...
```

```
user_specification:
    user
    [
        IDENTIFIED BY [PASSWORD] 'password'
```



```
| IDENTIFIED WITH auth_plugin [AS 'auth_string']  
]
```

The `CREATE USER` statement creates new MySQL accounts. To use it, you must have the global `CREATE USER` [754] privilege or the `INSERT` [755] privilege for the `mysql` database. For each account, `CREATE USER` creates a new row in the `mysql.user` table and assigns the account no privileges. An error occurs if the account already exists.

For `CREATE USER` statements that do not specify an `IDENTIFIED WITH auth_plugin` clause, the server associates the account with the default authentication plugin. As of MySQL 5.6.6, this is the plugin named by the `--default-authentication-plugin` [419] option at server startup, or `mysql_native_password` if that option is not used. Before 5.6.6, the default plugin is `mysql_native_password`. For information about authentication plugins, see [Section 6.3.7, “Pluggable Authentication”](#).

Each account name uses the format described in [Section 6.2.3, “Specifying Account Names”](#). For example:

```
CREATE USER 'jeffrey'@'localhost' IDENTIFIED BY 'mypass';
```

If you specify only the user name part of the account name, a host name part of `'%'` is used.

The user specification may indicate how the user should authenticate when connecting to the server:

- To enable the user to connect with no password (which is *insecure*), include no `IDENTIFIED BY` clause:

```
CREATE USER 'jeffrey'@'localhost';
```

In this case, the account uses the default authentication plugin and clients must provide no password.

- To assign a password, use `IDENTIFIED BY` with the literal plaintext password value:

```
CREATE USER 'jeffrey'@'localhost' IDENTIFIED BY 'mypass';
```

The account uses the default authentication plugin and clients must match the given password.

- To avoid specifying the plaintext password if you know its hash value (the value that `PASSWORD()` [1263] would return for the password), specify the hash value preceded by the keyword `PASSWORD`:

```
CREATE USER 'jeffrey'@'localhost'  
IDENTIFIED BY PASSWORD '*90E462C37378CED12064BB3388827D2BA3A9B689';
```

The account uses the default authentication plugin and the has must be in the format required by that plugin. Clients must match the given password.

- To authenticate the account using a specific authentication plugin, use `IDENTIFIED WITH`, where `auth_plugin` is the plugin name. It can be an unquoted name or a quoted string literal. `'auth_string'` is an optional quoted string literal to pass to the plugin. The plugin interprets the meaning of the string, so its format is plugin specific. Consult the documentation for a given plugin for information about the authentication string values it accepts.

```
CREATE USER 'jeffrey'@'localhost'  
IDENTIFIED WITH my_auth_plugin;
```

For connections that use this account, the server invokes the named plugin and clients must provide credentials as required for the authentication method that the plugin implements. If the server cannot find the plugin, either at account-creation time or connect time, an error occurs.

The `IDENTIFIED BY` and `IDENTIFIED WITH` clauses are mutually exclusive, so at most one of them can be specified for a given user.

For additional information about setting passwords, see [Section 6.3.5, “Assigning Account Passwords”](#).



#### Important

`CREATE USER` may be recorded in server logs or in a history file such as `~/.mysql_history`, which means that cleartext passwords may be read by anyone having read access to that information. See [Section 6.1.2, “Keeping Passwords Secure”](#).



#### Important

Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

### 13.7.1.3. DROP USER Syntax

```
DROP USER user [, user] ...
```

The `DROP USER` statement removes one or more MySQL accounts and their privileges. It removes privilege rows for the account from all grant tables. To use this statement, you must have the global `CREATE USER [754]` privilege or the `DELETE [754]` privilege for the `mysql` database. Each account name uses the format described in [Section 6.2.3, “Specifying Account Names”](#). For example:

```
DROP USER 'jeffrey'@'localhost';
```

If you specify only the user name part of the account name, a host name part of `'%'` is used.



#### Important

`DROP USER` does not automatically close any open user sessions. Rather, in the event that a user with an open session is dropped, the statement does not take effect until that user's session is closed. Once the session is closed, the user is dropped, and that user's next attempt to log in will fail. *This is by design.*

`DROP USER` does not automatically drop or invalidate databases or objects within them that the old user created. This includes stored programs or views for which the `DEFINER` attribute names the dropped user. Attempts to access such objects may produce an error if they execute in definer security context. (For information about security context, see [Section 18.6, “Access Control for Stored Programs and Views”](#).)

### 13.7.1.4. GRANT Syntax

```
GRANT
  priv_type [(column_list)]
  [, priv_type [(column_list)] ...]
ON [object_type] priv_level
TO user_specification [, user_specification] ...
[REQUIRE {NONE | ssl_option [[AND] ssl_option] ...}]
[WITH with_option ...]

GRANT PROXY ON user_specification
TO user_specification [, user_specification] ...
[WITH GRANT OPTION]

object_type:
TABLE
| FUNCTION
| PROCEDURE

priv_level:
*
```

```

| *.*
| db_name.*
| db_name.tbl_name
| tbl_name
| db_name.routine_name

user_specification:
  user
  [
    IDENTIFIED BY [PASSWORD] 'password'
    | IDENTIFIED WITH auth_plugin [AS 'auth_string']
  ]

ssl_option:
  SSL
  | X509
  | CIPHER 'cipher'
  | ISSUER 'issuer'
  | SUBJECT 'subject'

with_option:
  GRANT OPTION
  | MAX_QUERIES_PER_HOUR count
  | MAX_UPDATES_PER_HOUR count
  | MAX_CONNECTIONS_PER_HOUR count
  | MAX_USER_CONNECTIONS count

```

The [GRANT](#) statement grants privileges to MySQL user accounts. [GRANT](#) also serves to specify other account characteristics such as use of secure connections and limits on access to server resources. To use [GRANT](#), you must have the [GRANT OPTION \[755\]](#) privilege, and you must have the privileges that you are granting.

Normally, a database administrator first uses [CREATE USER](#) to create an account, then [GRANT](#) to define its privileges and characteristics. For example:

```

CREATE USER 'jeffrey'@'localhost' IDENTIFIED BY 'mypass';
GRANT ALL ON db1.* TO 'jeffrey'@'localhost';
GRANT SELECT ON db2.invoice TO 'jeffrey'@'localhost';
GRANT USAGE ON *.* TO 'jeffrey'@'localhost' WITH MAX_QUERIES_PER_HOUR 90;

```

However, if an account named in a [GRANT](#) statement does not already exist, [GRANT](#) may create it under the conditions described later in the discussion of the [NO\\_AUTO\\_CREATE\\_USER \[604\]](#) SQL mode.

The [REVOKE](#) statement is related to [GRANT](#) and enables administrators to remove account privileges. See [Section 13.7.1.6, “REVOKE Syntax”](#).

When successfully executed from the `mysql` program, [GRANT](#) responds with `Query OK, 0 rows affected`. To determine what privileges result from the operation, use [SHOW GRANTS](#). See [Section 13.7.5.22, “SHOW GRANTS Syntax”](#).

There are several aspects to the [GRANT](#) statement, described under the following topics in this section:

- [Privileges Supported by MySQL](#)
- [Global Privileges](#)
- [Database Privileges](#)
- [Table Privileges](#)
- [Column Privileges](#)
- [Stored Routine Privileges](#)
- [Proxy User Privileges](#)

- [Account Names and Passwords](#)
- [Other Account Characteristics](#)
- [MySQL and Standard SQL Versions of GRANT](#)



### Important

Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

## Privileges Supported by MySQL

The following table summarizes the permissible *priv\_type* privilege types that can be specified for the [GRANT](#) and [REVOKE](#) statements. For additional information about these privileges, see [Section 6.2.1, “Privileges Provided by MySQL”](#).

**Table 13.1. Permissible Privileges for [GRANT](#) and [REVOKE](#)**

Privilege	Meaning
<a href="#">ALL [PRIVILEGES] [753]</a>	Grant all privileges at specified access level except <a href="#">GRANT OPTION [755]</a>
<a href="#">ALTER [753]</a>	Enable use of <a href="#">ALTER TABLE</a>
<a href="#">ALTER ROUTINE [754]</a>	Enable stored routines to be altered or dropped
<a href="#">CREATE [754]</a>	Enable database and table creation
<a href="#">CREATE ROUTINE [754]</a>	Enable stored routine creation
<a href="#">CREATE TABLESPACE [754]</a>	Enable tablespaces and log file groups to be created, altered, or dropped
<a href="#">CREATE TEMPORARY TABLES [754]</a>	Enable use of <a href="#">CREATE TEMPORARY TABLE</a>
<a href="#">CREATE USER [754]</a>	Enable use of <a href="#">CREATE USER</a> , <a href="#">DROP USER</a> , <a href="#">RENAME USER</a> , and <a href="#">REVOKE ALL PRIVILEGES</a>
<a href="#">CREATE VIEW [754]</a>	Enable views to be created or altered
<a href="#">DELETE [754]</a>	Enable use of <a href="#">DELETE</a>
<a href="#">DROP [754]</a>	Enable databases, tables, and views to be dropped
<a href="#">EVENT [754]</a>	Enable use of events for the Event Scheduler
<a href="#">EXECUTE [754]</a>	Enable the user to execute stored routines
<a href="#">FILE [754]</a>	Enable the user to cause the server to read or write files
<a href="#">GRANT OPTION [755]</a>	Enable privileges to be granted to or removed from other accounts
<a href="#">INDEX [755]</a>	Enable indexes to be created or dropped
<a href="#">INSERT [755]</a>	Enable use of <a href="#">INSERT</a>
<a href="#">LOCK TABLES [755]</a>	Enable use of <a href="#">LOCK TABLES</a> on tables for which you have the <a href="#">SELECT</a> privilege
<a href="#">PROCESS [755]</a>	Enable the user to see all processes with <a href="#">SHOW PROCESSLIST</a>
<a href="#">PROXY [755]</a>	Enable user proxying
<a href="#">REFERENCES [755]</a>	Not implemented
<a href="#">RELOAD [755]</a>	Enable use of <a href="#">FLUSH</a> operations

Privilege	Meaning
REPLICATION CLIENT [755]	Enable the user to ask where master or slave servers are
REPLICATION SLAVE [755]	Enable replication slaves to read binary log events from the master
SELECT [755]	Enable use of <code>SELECT</code>
SHOW DATABASES [756]	Enable <code>SHOW DATABASES</code> to show all databases
SHOW VIEW [756]	Enable use of <code>SHOW CREATE VIEW</code>
SHUTDOWN [756]	Enable use of <code>mysqladmin shutdown</code>
SUPER [756]	Enable use of other administrative operations such as <code>CHANGE MASTER TO</code> , <code>KILL</code> , <code>PURGE BINARY LOGS</code> , <code>SET GLOBAL</code> , and <code>mysqladmin debug</code> command
TRIGGER [756]	Enable trigger operations
UPDATE [756]	Enable use of <code>UPDATE</code>
USAGE [756]	Synonym for “no privileges”

A trigger is associated with a table, so to create or drop a trigger, you must have the `TRIGGER` [756] privilege for the table, not the trigger.

In `GRANT` statements, the `ALL [PRIVILEGES]` [753] or `PROXY` [755] privilege must be named by itself and cannot be specified along with other privileges. `ALL [PRIVILEGES]` [753] stands for all privileges available for the level at which privileges are to be granted except for the `GRANT OPTION` [755] and `PROXY` [755] privileges.

`USAGE` [756] can be specified to create a user that has no privileges, or to specify the `REQUIRE` or `WITH` clauses for an account without changing its existing privileges.

MySQL account information is stored in the tables of the `mysql` database. This database and the access control system are discussed extensively in Section 6.2, “The MySQL Access Privilege System”, which you should consult for additional details.

If the grant tables hold privilege rows that contain mixed-case database or table names and the `lower_case_table_names` [501] system variable is set to a nonzero value, `REVOKE` cannot be used to revoke these privileges. It will be necessary to manipulate the grant tables directly. (`GRANT` will not create such rows when `lower_case_table_names` [501] is set, but such rows might have been created prior to setting that variable.)

Privileges can be granted at several levels, depending on the syntax used for the `ON` clause. For `REVOKE`, the same `ON` syntax specifies which privileges to take away. The examples shown here include no `IDENTIFIED BY 'password'` clause for brevity, but you should include one if the account does not already exist, to avoid creating an insecure account that has no password.

## Global Privileges

Global privileges are administrative or apply to all databases on a given server. To assign global privileges, use `ON *.*` syntax:

```
GRANT ALL ON *.* TO 'someuser'@'somehost';
GRANT SELECT, INSERT ON *.* TO 'someuser'@'somehost';
```

The `CREATE TABLESPACE` [754], `CREATE USER` [754], `FILE` [754], `PROCESS` [755], `RELOAD` [755], `REPLICATION CLIENT` [755], `REPLICATION SLAVE` [755], `SHOW DATABASES` [756], `SHUTDOWN` [756], and `SUPER` [756] privileges are administrative and can only be granted globally.

Other privileges can be granted globally or at more specific levels.

MySQL stores global privileges in the `mysql.user` table.

## Database Privileges

Database privileges apply to all objects in a given database. To assign database-level privileges, use `ON db_name.*` syntax:

```
GRANT ALL ON mydb.* TO 'someuser'@'somehost';
GRANT SELECT, INSERT ON mydb.* TO 'someuser'@'somehost';
```

If you use `ON *` syntax (rather than `ON *.*`) and you have selected a default database, privileges are assigned at the database level for the default database. An error occurs if there is no default database.

The `CREATE` [754], `DROP` [754], `EVENT` [754], `GRANT OPTION` [755], and `LOCK TABLES` [755] privileges can be specified at the database level. Table or routine privileges also can be specified at the database level, in which case they apply to all tables or routines in the database.

MySQL stores database privileges in the `mysql.db` table.

## Table Privileges

Table privileges apply to all columns in a given table. To assign table-level privileges, use `ON db_name.tbl_name` syntax:

```
GRANT ALL ON mydb.mytbl TO 'someuser'@'somehost';
GRANT SELECT, INSERT ON mydb.mytbl TO 'someuser'@'somehost';
```

If you specify `tbl_name` rather than `db_name.tbl_name`, the statement applies to `tbl_name` in the default database. An error occurs if there is no default database.

The permissible `priv_type` values at the table level are `ALTER` [753], `CREATE VIEW` [754], `CREATE` [754], `DELETE` [754], `DROP` [754], `GRANT OPTION` [755], `INDEX` [755], `INSERT` [755], `SELECT` [755], `SHOW VIEW` [756], `TRIGGER` [756], and `UPDATE` [756].

MySQL stores table privileges in the `mysql.tables_priv` table.

## Column Privileges

Column privileges apply to single columns in a given table. Each privilege to be granted at the column level must be followed by the column or columns, enclosed within parentheses.

```
GRANT SELECT (col1), INSERT (col1,col2) ON mydb.mytbl TO 'someuser'@'somehost';
```

The permissible `priv_type` values for a column (that is, when you use a `column_list` clause) are `INSERT` [755], `SELECT` [755], and `UPDATE` [756].

MySQL stores column privileges in the `mysql.columns_priv` table.

## Stored Routine Privileges

The `ALTER ROUTINE` [754], `CREATE ROUTINE` [754], `EXECUTE` [754], and `GRANT OPTION` [755] privileges apply to stored routines (procedures and functions). They can be granted at the global and database levels. Except for `CREATE ROUTINE` [754], these privileges can be granted at the routine level for individual routines.

```
GRANT CREATE ROUTINE ON mydb.* TO 'someuser'@'somehost';
GRANT EXECUTE ON PROCEDURE mydb.myproc TO 'someuser'@'somehost';
```

The permissible `priv_type` values at the routine level are `ALTER ROUTINE` [754], `EXECUTE` [754], and `GRANT OPTION` [755]. `CREATE ROUTINE` [754] is not a routine-level privilege because you must have this privilege to create a routine in the first place.

MySQL stores routine-level privileges in the `mysql.procs_priv` table.

## Proxy User Privileges

The [PROXY \[755\]](#) privilege enables one user to be a proxy for another. The proxy user impersonates or takes the identity of the proxied user.

```
GRANT PROXY ON 'localuser'@'localhost' TO 'externaluser'@'somehost';
```

When [PROXY \[755\]](#) is granted, it must be the only privilege named in the [GRANT](#) statement, the [REQUIRE](#) clause cannot be given, and the only permitted [WITH](#) option is [WITH GRANT OPTION](#).

Proxying requires that the proxy user authenticate through a plugin that returns the name of the proxied user to the server when the proxy user connects, and that the proxy user have the [PROXY](#) privilege for the proxied user. For details and examples, see [Section 6.3.8, “Proxy Users”](#).

MySQL stores proxy privileges in the `mysql.proxies_priv` table.

For the global, database, table, and routine levels, [GRANT ALL](#) assigns only the privileges that exist at the level you are granting. For example, [GRANT ALL ON db\\_name.\\*](#) is a database-level statement, so it does not grant any global-only privileges such as [FILE \[754\]](#). Granting [ALL \[753\]](#) does not assign the [PROXY \[755\]](#) privilege.

The `object_type` clause, if present, should be specified as [TABLE](#), [FUNCTION](#), or [PROCEDURE](#) when the following object is a table, a stored function, or a stored procedure.

The privileges for a database, table, column, or routine are formed additively as the logical [OR \[1172\]](#) of the privileges at each of the privilege levels. For example, if a user has a global [SELECT \[755\]](#) privilege, the privilege cannot be denied by an absence of the privilege at the database, table, or column level. Details of the privilege-checking procedure are presented in [Section 6.2.5, “Access Control, Stage 2: Request Verification”](#).

MySQL enables you to grant privileges on databases or tables that do not exist. For tables, the privileges to be granted must include the [CREATE \[754\]](#) privilege. *This behavior is by design*, and is intended to enable the database administrator to prepare user accounts and privileges for databases or tables that are to be created at a later time.



### Important

MySQL does not automatically revoke any privileges when you drop a database or table. However, if you drop a routine, any routine-level privileges granted for that routine are revoked.

## Account Names and Passwords

The `user` value indicates the MySQL account to which the [GRANT](#) statement applies. To accommodate granting rights to users from arbitrary hosts, MySQL supports specifying the `user` value in the form `user_name@host_name`. If a `user_name` or `host_name` value is legal as an unquoted identifier, you need not quote it. However, quotation marks are necessary to specify a `user_name` string containing special characters (such as “-”), or a `host_name` string containing special characters or wildcard characters (such as “%”); for example, `'test-user'@'%.com'`. Quote the user name and host name separately.

You can specify wildcards in the host name. For example, `user_name@'%.example.com'` applies to `user_name` for any host in the `example.com` domain, and `user_name@'192.168.1.%'` applies to `user_name` for any host in the `192.168.1` class C subnet.

The simple form `user_name` is a synonym for `user_name@'%'`.

MySQL does not support wildcards in user names. To refer to an anonymous user, specify an account with an empty user name with the [GRANT](#) statement:

```
GRANT ALL ON test.* TO ''@'localhost' ...
```



In this case, any user who connects from the local host with the correct password for the anonymous user will be permitted access, with the privileges associated with the anonymous-user account.

For additional information about user name and host name values in account names, see [Section 6.2.3, “Specifying Account Names”](#).

To specify quoted values, quote database, table, column, and routine names as identifiers. Quote user names and host names as identifiers or as strings. Quote passwords as strings. For string-quoting and identifier-quoting guidelines, see [Section 9.1.1, “String Literals”](#), and [Section 9.2, “Schema Object Names”](#).

The “\_” and “%” wildcards are permitted when specifying database names in `GRANT` statements that grant privileges at the global or database levels. This means, for example, that if you want to use a “\_” character as part of a database name, you should specify it as “\\_” in the `GRANT` statement, to prevent the user from being able to access additional databases matching the wildcard pattern; for example,

```
GRANT ... ON `foo\_bar`.* TO ....
```



### Warning

If you permit anonymous users to connect to the MySQL server, you should also grant privileges to all local users as `user_name@localhost`. Otherwise, the anonymous user account for `localhost` in the `mysql.user` table (created during MySQL installation) is used when named users try to log in to the MySQL server from the local machine. For details, see [Section 6.2.4, “Access Control, Stage 1: Connection Verification”](#).

To determine whether the preceding warning applies to you, execute the following query, which lists any anonymous users:

```
SELECT Host, User FROM mysql.user WHERE User='';
```

To avoid the problem just described, delete the local anonymous user account using this statement:

```
DROP USER ''@'localhost';
```

`GRANT` supports host names up to 60 characters long. Database, table, column, and routine names can be up to 64 characters. User names can be up to 16 characters.



### Warning

*The permissible length for user names cannot be changed by altering the `mysql.user` table. Attempting to do so results in unpredictable behavior which may even make it impossible for users to log in to the MySQL server.* You should never alter any of the tables in the `mysql` database in any manner whatsoever except by means of the procedure described in [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

The user specification may indicate how the user should authenticate when connecting to the server, through inclusion of an `IDENTIFIED BY` or `IDENTIFIED WITH` clause. The syntax is the same as for the `CREATE USER` statement. See [Section 13.7.1.2, “CREATE USER Syntax”](#).

When the `IDENTIFIED BY` clause is present and you have global grant privileges, the password becomes the new password for the account, even if the account exists and already has a password. With no `IDENTIFIED BY` clause, the account password remains unchanged.

If the account named in a `GRANT` statement does not exist in the `mysql.user` table, `GRANT` creates it if the `NO_AUTO_CREATE_USER [604]` SQL mode is not enabled. *This is very insecure* unless you specify a nonempty password using `IDENTIFIED BY` or an authentication plugin using `IDENTIFIED WITH`.

If the account does not exist and `NO_AUTO_CREATE_USER` [604] is enabled, `GRANT` fails and does not create the account unless you specify a nonempty password with `IDENTIFIED BY` or specify an `IDENTIFIED WITH` clause to name an authentication plugin.



### Important

`GRANT` may be recorded in server logs or in a history file such as `~/mysql_history`, which means that cleartext passwords may be read by anyone having read access to that information. See [Section 6.1.2, “Keeping Passwords Secure”](#).

## Other Account Characteristics

The `WITH` clause is used for several purposes:

- To enable a user to grant privileges to other users
- To specify resource limits for a user
- To specify whether and how a user must use secure connections to the server

The `WITH GRANT OPTION` clause gives the user the ability to give to other users any privileges the user has at the specified privilege level. You should be careful to whom you give the `GRANT OPTION` [755] privilege because two users with different privileges may be able to combine privileges!

You cannot grant another user a privilege which you yourself do not have; the `GRANT OPTION` [755] privilege enables you to assign only those privileges which you yourself possess.

Be aware that when you grant a user the `GRANT OPTION` [755] privilege at a particular privilege level, any privileges the user possesses (or may be given in the future) at that level can also be granted by that user to other users. Suppose that you grant a user the `INSERT` [755] privilege on a database. If you then grant the `SELECT` [755] privilege on the database and specify `WITH GRANT OPTION`, that user can give to other users not only the `SELECT` [755] privilege, but also `INSERT` [755]. If you then grant the `UPDATE` [756] privilege to the user on the database, the user can grant `INSERT` [755], `SELECT` [755], and `UPDATE` [756].

For a nonadministrative user, you should not grant the `ALTER` [753] privilege globally or for the `mysql` database. If you do that, the user can try to subvert the privilege system by renaming tables!

For additional information about security risks associated with particular privileges, see [Section 6.2.1, “Privileges Provided by MySQL”](#).

Several `WITH` clause options specify limits on use of server resources by an account:

- The `MAX_QUERIES_PER_HOUR count`, `MAX_UPDATES_PER_HOUR count`, and `MAX_CONNECTIONS_PER_HOUR count` limits restrict the number of queries, updates, and connections to the server permitted to this account during any given one-hour period. (Queries for which results are served from the query cache do not count against the `MAX_QUERIES_PER_HOUR` limit.) If `count` is 0 (the default), this means that there is no limitation for the account.
- The `MAX_USER_CONNECTIONS count` limit restricts the maximum number of simultaneous connections to the server by the account. A nonzero `count` specifies the limit for the account explicitly. If `count` is 0 (the default), the server determines the number of simultaneous connections for the account from the global value of the `max_user_connections` [509] system variable. If `max_user_connections` [509] is also zero, there is no limit for the account.

To specify resource limits for an existing user without affecting existing privileges, use `GRANT USAGE` at the global level (`ON *.*`) and name the limits to be changed. For example:

```
GRANT USAGE ON *.* TO ...
WITH MAX_QUERIES_PER_HOUR 500 MAX_UPDATES_PER_HOUR 100;
```

Limits not specified retain their current values.

For more information on restricting access to server resources, see [Section 6.3.4, “Setting Account Resource Limits”](#).

MySQL can check X509 certificate attributes in addition to the usual authentication that is based on the user name and password. To specify SSL-related options for a MySQL account, use the `REQUIRE` clause of the `GRANT` statement. (For background information on the use of SSL with MySQL, see [Section 6.3.9, “Using SSL for Secure Connections”](#).)

There are a number of different possibilities for limiting connection types for a given account:

- `REQUIRE NONE` indicates that the account has no SSL or X509 requirements. This is the default if no SSL-related `REQUIRE` options are specified. Unencrypted connections are permitted if the user name and password are valid. However, encrypted connections can also be used, at the client's option, if the client has the proper certificate and key files. That is, the client need not specify any SSL command options, in which case the connection will be unencrypted. To use an encrypted connection, the client must specify either the `--ssl-ca [810]` option, or all three of the `--ssl-ca [810]`, `--ssl-key [811]`, and `--ssl-cert [811]` options.
- The `REQUIRE SSL` option tells the server to permit only SSL-encrypted connections for the account.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret' REQUIRE SSL;
```

To connect, the client must specify the `--ssl-ca [810]` option to authenticate the server certificate, and may additionally specify the `--ssl-key [811]` and `--ssl-cert [811]` options. If neither `--ssl-ca [810]` option nor `--ssl-capath [810]` option is specified, the client does not authenticate the server certificate.

- `REQUIRE X509` means that the client must have a valid certificate but that the exact certificate, issuer, and subject do not matter. The only requirement is that it should be possible to verify its signature with one of the CA certificates.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret' REQUIRE X509;
```

To connect, the client must specify the `--ssl-ca [810]`, `--ssl-key [811]`, and `--ssl-cert [811]` options. This is also true for `ISSUER` and `SUBJECT` because those `REQUIRE` options imply `X509`.

- `REQUIRE ISSUER 'issuer'` places the restriction on connection attempts that the client must present a valid X509 certificate issued by CA `'issuer'`. If the client presents a certificate that is valid but has a different issuer, the server rejects the connection. Use of X509 certificates always implies encryption, so the `SSL` option is unnecessary in this case.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
  REQUIRE ISSUER '/C=FI/ST=Some-State/L=Helsinki/
  O=MySQL Finland AB/CN=Tonu Samuel/emailAddress=tonu@example.com';
```

The `'issuer'` value should be entered as a single string.



#### Note

If MySQL is linked against a version of OpenSSL older than 0.9.6h, use `Email` rather than `emailAddress` in the `'issuer'` value.

- `REQUIRE SUBJECT 'subject'` places the restriction on connection attempts that the client must present a valid X509 certificate containing the subject `subject`. If the client presents a certificate that is valid but has a different subject, the server rejects the connection.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
```

```

REQUIRE SUBJECT '/C=EE/ST=Some-State/L=Tallinn/
O=MySQL demo client certificate/
CN=Tonu Samuel/emailAddress=tonu@example.com';

```

The '*subject*' value should be entered as a single string. MySQL does a simple string comparison of this value to the value in the certificate, so lettercase and component ordering must be given exactly as present in the certificate.



### Note

Regarding `emailAddress`, see the note in the description of `REQUIRE ISSUER`.

- `REQUIRE CIPHER 'cipher'` is needed to ensure that ciphers and key lengths of sufficient strength are used. SSL itself can be weak if old algorithms using short encryption keys are used. Using this option, you can ask that a specific cipher method is used for a connection.

```

GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
IDENTIFIED BY 'goodsecret'
REQUIRE CIPHER 'EDH-RSA-DES-CBC3-SHA';

```

The `SUBJECT`, `ISSUER`, and `CIPHER` options can be combined in the `REQUIRE` clause like this:

```

GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
IDENTIFIED BY 'goodsecret'
REQUIRE SUBJECT '/C=EE/ST=Some-State/L=Tallinn/
O=MySQL demo client certificate/
CN=Tonu Samuel/emailAddress=tonu@example.com'
AND ISSUER '/C=FI/ST=Some-State/L=Helsinki/
O=MySQL Finland AB/CN=Tonu Samuel/emailAddress=tonu@example.com'
AND CIPHER 'EDH-RSA-DES-CBC3-SHA';

```

The order of the options does not matter, but no option can be specified twice. The `AND` keyword is optional between `REQUIRE` options.

If you are using table, column, or routine privileges for even one user, the server examines table, column, and routine privileges for all users and this slows down MySQL a bit. Similarly, if you limit the number of queries, updates, or connections for any users, the server must monitor these values.

## MySQL and Standard SQL Versions of `GRANT`

The biggest differences between the MySQL and standard SQL versions of `GRANT` are:

- MySQL associates privileges with the combination of a host name and user name and not with only a user name.
- Standard SQL does not have global or database-level privileges, nor does it support all the privilege types that MySQL supports.
- MySQL does not support the standard SQL `UNDER` privilege.
- Standard SQL privileges are structured in a hierarchical manner. If you remove a user, all privileges the user has been granted are revoked. This is also true in MySQL if you use `DROP USER`. See [Section 13.7.1.3, “DROP USER Syntax”](#).
- In standard SQL, when you drop a table, all privileges for the table are revoked. In standard SQL, when you revoke a privilege, all privileges that were granted based on that privilege are also revoked. In MySQL, privileges can be dropped only with explicit `DROP USER` or `REVOKE` statements or by manipulating the MySQL grant tables directly.
- In MySQL, it is possible to have the `INSERT [755]` privilege for only some of the columns in a table. In this case, you can still execute `INSERT` statements on the table, provided that you insert values only for those columns for which you have the `INSERT [755]` privilege. The omitted columns are

set to their implicit default values if strict SQL mode is not enabled. In strict mode, the statement is rejected if any of the omitted columns have no default value. (Standard SQL requires you to have the [INSERT \[755\]](#) privilege on all columns.) [Section 5.1.7, “Server SQL Modes”](#), discusses strict mode. [Section 11.5, “Data Type Default Values”](#), discusses implicit default values.

### 13.7.1.5. RENAME USER Syntax

```
RENAME USER old_user TO new_user
[, old_user TO new_user] ...
```

The `RENAME USER` statement renames existing MySQL accounts. To use it, you must have the global [CREATE USER \[754\]](#) privilege or the [UPDATE \[756\]](#) privilege for the `mysql` database. An error occurs if any old account does not exist or any new account exists. Each account name uses the format described in [Section 6.2.3, “Specifying Account Names”](#). For example:

```
RENAME USER 'jeffrey'@'localhost' TO 'jeff'@'127.0.0.1';
```

If you specify only the user name part of the account name, a host name part of `'%'` is used.

`RENAME USER` causes the privileges held by the old user to be those held by the new user. However, `RENAME USER` does not automatically drop or invalidate databases or objects within them that the old user created. This includes stored programs or views for which the `DEFINER` attribute names the old user. Attempts to access such objects may produce an error if they execute in definer security context. (For information about security context, see [Section 18.6, “Access Control for Stored Programs and Views”](#).)

The privilege changes take effect as indicated in [Section 6.2.6, “When Privilege Changes Take Effect”](#).

### 13.7.1.6. REVOKE Syntax

```
REVOKE
  priv_type [(column_list)]
  [, priv_type [(column_list))] ...
ON [object_type] priv_level
FROM user [, user] ...

REVOKE ALL PRIVILEGES, GRANT OPTION
FROM user [, user] ...

REVOKE PROXY ON user
FROM user [, user] ...
```

The `REVOKE` statement enables system administrators to revoke privileges from MySQL accounts. Each account name uses the format described in [Section 6.2.3, “Specifying Account Names”](#). For example:

```
REVOKE INSERT ON *.* FROM 'jeffrey'@'localhost';
```

If you specify only the user name part of the account name, a host name part of `'%'` is used.

For details on the levels at which privileges exist, the permissible `priv_type` and `priv_level` values, and the syntax for specifying users and passwords, see [Section 13.7.1.4, “GRANT Syntax”](#)

To use the first `REVOKE` syntax, you must have the [GRANT OPTION \[755\]](#) privilege, and you must have the privileges that you are revoking.

To revoke all privileges, use the second syntax, which drops all global, database, table, column, and routine privileges for the named user or users:

```
REVOKE ALL PRIVILEGES, GRANT OPTION FROM user [, user] ...
```

To use this `REVOKE` syntax, you must have the global [CREATE USER \[754\]](#) privilege or the [UPDATE \[756\]](#) privilege for the `mysql` database.

`REVOKE` removes privileges, but does not drop `mysql.user` table entries. To remove a user account entirely, use `DROP USER` (see [Section 13.7.1.3](#), “`DROP USER Syntax`”) or `DELETE`.

If the grant tables hold privilege rows that contain mixed-case database or table names and the `lower_case_table_names` [501] system variable is set to a nonzero value, `REVOKE` cannot be used to revoke these privileges. It will be necessary to manipulate the grant tables directly. (`GRANT` will not create such rows when `lower_case_table_names` [501] is set, but such rows might have been created prior to setting the variable.)

When successfully executed from the `mysql` program, `REVOKE` responds with `Query OK, 0 rows affected`. To determine what privileges result from the operation, use `SHOW GRANTS`. See [Section 13.7.5.22](#), “`SHOW GRANTS Syntax`”.

### 13.7.1.7. `SET PASSWORD Syntax`

```
SET PASSWORD [FOR user] =
{
  PASSWORD('cleartext password')
  | OLD_PASSWORD('cleartext password')
  | 'encrypted password'
}
```

The `SET PASSWORD` statement assigns a password to an existing MySQL user account. When the `read_only` [532] system variable is enabled, the `SUPER` [756] privilege is required to use `SET PASSWORD`, in addition to whatever other privileges might be required.

If the password is specified using the `PASSWORD()` [1263] or `OLD_PASSWORD()` [1263] function, the cleartext (unencrypted) password should be given as the argument to the function, which hashes the password and returns the encrypted password string. If the password is specified without using either function, it should be the already encrypted password value as a literal string. In all cases, the encrypted password string must be in the format required by the authentication method used for the account.

With no `FOR user` clause, this statement sets the password for the current user. (To see which account the server authenticated you as, invoke the `CURRENT_USER()` [1267] function.) Any client who successfully connects to the server using a nonanonymous account can change the password for that account.

With a `FOR user` clause, this statement sets the password for the named user. You must have the `UPDATE` privilege for the `mysql` database to do this. The user account name uses the format described in [Section 6.2.3](#), “`Specifying Account Names`”. The `user` value should be given as `'user_name'@'host_name'`, where `'user_name'` and `'host_name'` are exactly as listed in the `User` and `Host` columns of the `mysql.user` table row. (If you specify only a user name, a host name of `'%'` is used.) For example, to set the password for an account with `User` and `Host` column values of `'bob'` and `'%.example.org'`, write the statement like this:

```
SET PASSWORD FOR 'bob'@'%.example.org' = PASSWORD('cleartext password');
```

That is equivalent to the following statements:

```
UPDATE mysql.user SET Password=PASSWORD('cleartext password')
  WHERE User='bob' AND Host='%.example.org';
FLUSH PRIVILEGES;
```

Another way to set the password is to use `GRANT`:

```
GRANT USAGE ON *.* TO 'bob'@'%.example.org' IDENTIFIED BY 'cleartext password';
```

The `old_passwords` [517] system variable value determines the hashing method used by `PASSWORD()` [1263]. If you specify the password using that function and `SET PASSWORD` rejects the password as not being in the correct format, it may be necessary to set `old_passwords` [517]



to change the hashing method. For descriptions of the permitted values, see [Section 5.1.4, “Server System Variables”](#).

For more information about setting passwords, see [Section 6.3.5, “Assigning Account Passwords”](#)



#### Important

`SET PASSWORD` may be recorded in server logs or in a history file such as `~/.mysql_history`, which means that cleartext passwords may be read by anyone having read access to that information. See [Section 6.1.2, “Keeping Passwords Secure”](#).



#### Caution

If you are connecting to a MySQL 4.1 or later server using a pre-4.1 client program, do not change your password without first reading [Section 6.1.2.4, “Password Hashing in MySQL”](#). The default password hashing format changed in MySQL 4.1, and if you change your password, it might be stored using a hashing format that pre-4.1 clients cannot generate, thus preventing you from connecting to the server afterward.

If you are using MySQL Replication, be aware that, currently, a password used by a replication slave as part of a `CHANGE MASTER TO` statement is effectively limited to 32 characters in length; the password can be longer, but any excess characters are truncated. This is not due to any limit imposed by the MySQL Server generally, but rather is an issue specific to MySQL Replication. (For more information, see Bug #43439.)

## 13.7.2. Table Maintenance Statements

### 13.7.2.1. `ANALYZE TABLE` Syntax

```
ANALYZE [NO_WRITE_TO_BINLOG | LOCAL] TABLE
      tbl_name [, tbl_name] ...
```

`ANALYZE TABLE` analyzes and stores the key distribution for a table. During the analysis, the table is locked with a read lock for `InnoDB` and `MyISAM`. This statement works with `InnoDB`, `NDB`, and `MyISAM` tables. For `MyISAM` tables, this statement is equivalent to using `myisamchk --analyze`.

For more information on how the analysis works within `InnoDB`, see [Section 14.2.4.2.10, “Persistent Optimizer Statistics for InnoDB Tables”](#) and [Section 14.2.7, “Limits on InnoDB Tables”](#). In particular, when you enable the `innodb_stats_persistent [1747]` option, you must run `ANALYZE TABLE` after loading substantial data into an `InnoDB` table, or creating a new index for one.

MySQL uses the stored key distribution to decide the order in which tables should be joined when you perform a join on something other than a constant. In addition, key distributions can be used when deciding which indexes to use for a specific table within a query.

This statement requires `SELECT [755]` and `INSERT [755]` privileges for the table.

`ANALYZE TABLE` is supported for partitioned tables, and you can use `ALTER TABLE ... ANALYZE PARTITION` to analyze one or more partitions; for more information, see [Section 13.1.6, “ALTER TABLE Syntax”](#), and [Section 17.3.4, “Maintenance of Partitions”](#).

Beginning with MySQL 5.6.11, `gtid_next [2022]` must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

`ANALYZE TABLE` returns a result set with the following columns.

Column	Value
Table	The table name



Column	Value
Op	Always <code>analyze</code>
Msg_type	<code>status</code> , <code>error</code> , <code>info</code> , <code>note</code> , or <code>warning</code>
Msg_text	An informational message

You can check the stored key distribution with the `SHOW INDEX` statement. See [Section 13.7.5.23](#), “`SHOW INDEX Syntax`”.

If the table has not changed since the last `ANALYZE TABLE` statement, the table is not analyzed again.

By default, the server writes `ANALYZE TABLE` statements to the binary log so that they replicate to replication slaves. To suppress logging, use the optional `NO_WRITE_TO_BINLOG` keyword or its alias `LOCAL`.

### 13.7.2.2. CHECK TABLE Syntax

```
CHECK TABLE tbl_name [, tbl_name] ... [option] ...
```

```
option = {FOR UPGRADE | QUICK | FAST | MEDIUM | EXTENDED | CHANGED}
```

`CHECK TABLE` checks a table or tables for errors. `CHECK TABLE` works for InnoDB, MyISAM, ARCHIVE, and CSV tables. For MyISAM tables, the key statistics are updated as well.

To check a table, you must have some privilege for it.

`CHECK TABLE` can also check views for problems, such as tables that are referenced in the view definition that no longer exist.

`CHECK TABLE` is supported for partitioned tables, and you can use `ALTER TABLE ... CHECK PARTITION` to check one or more partitions; for more information, see [Section 13.1.6](#), “`ALTER TABLE Syntax`”, and [Section 17.3.4](#), “`Maintenance of Partitions`”.

Beginning with MySQL 5.6.11, `gtid_next [2022]` must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

### Output

`CHECK TABLE` returns a result set with the following columns.

Column	Value
Table	The table name
Op	Always <code>check</code>
Msg_type	<code>status</code> , <code>error</code> , <code>info</code> , <code>note</code> , or <code>warning</code>
Msg_text	An informational message

Note that the statement might produce many rows of information for each checked table. The last row has a `Msg_type` value of `status` and the `Msg_text` normally should be `OK`. If you don't get `OK`, or `Table is already up to date` you should normally run a repair of the table. See [Section 7.6](#), “`MyISAM Table Maintenance and Crash Recovery`”. `Table is already up to date` means that the storage engine for the table indicated that there was no need to check the table.

### Checking Version Compatibility

The `FOR UPGRADE` option checks whether the named tables are compatible with the current version of MySQL. With `FOR UPGRADE`, the server checks each table to determine whether there have been any incompatible changes in any of the table's data types or indexes since the table was created. If not,

the check succeeds. Otherwise, if there is a possible incompatibility, the server runs a full check on the table (which might take some time). If the full check succeeds, the server marks the table's `.frm` file with the current MySQL version number. Marking the `.frm` file ensures that further checks for the table with the same version of the server will be fast.

Incompatibilities might occur because the storage format for a data type has changed or because its sort order has changed. Our aim is to avoid these changes, but occasionally they are necessary to correct problems that would be worse than an incompatibility between releases.

Currently, `FOR UPGRADE` discovers these incompatibilities:

- The indexing order for end-space in `TEXT` columns for `InnoDB` and `MyISAM` tables changed between MySQL 4.1 and 5.0.
- The storage method of the new `DECIMAL` data type changed between MySQL 5.0.3 and 5.0.5.
- If your table was created by a different version of the MySQL server than the one you are currently running, `FOR UPGRADE` indicates that the table has an `.frm` file with an incompatible version. In this case, the result set returned by `CHECK TABLE` contains a line with a `Msg_type` value of `error` and a `Msg_text` value of `Table upgrade required. Please do "REPAIR TABLE `tbl_name`" to fix it!`
- Changes are sometimes made to character sets or collations that require table indexes to be rebuilt. For details about these changes and when `FOR UPGRADE` detects them, see [Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”](#).
- The `YEAR(2)` is deprecated as of MySQL 5.6.6. `CHECK TABLE` recommends `REPAIR TABLE` for tables containing this data type. `REPAIR TABLE` converts `YEAR(2)` to `YEAR(4)`.

## Checking Data Consistency

The other check options that can be given are shown in the following table. These options are passed to the storage engine, which may use them or not.

Type	Meaning
<code>QUICK</code>	Do not scan the rows to check for incorrect links. Applies to <code>InnoDB</code> and <code>MyISAM</code> tables and views.
<code>FAST</code>	Check only tables that have not been closed properly. Applies only to <code>MyISAM</code> tables and views; ignored for <code>InnoDB</code> .
<code>CHANGED</code>	Check only tables that have been changed since the last check or that have not been closed properly. Applies only to <code>MyISAM</code> tables and views; ignored for <code>InnoDB</code> .
<code>MEDIUM</code>	Scan rows to verify that deleted links are valid. This also calculates a key checksum for the rows and verifies this with a calculated checksum for the keys. Applies only to <code>MyISAM</code> tables and views; ignored for <code>InnoDB</code> .
<code>EXTENDED</code>	Do a full key lookup for all keys for each row. This ensures that the table is 100% consistent, but takes a long time. Applies only to <code>MyISAM</code> tables and views; ignored for <code>InnoDB</code> .

If none of the options `QUICK`, `MEDIUM`, or `EXTENDED` are specified, the default check type for dynamic-format `MyISAM` tables is `MEDIUM`. This has the same result as running `myisamchk --medium-check tbl_name` on the table. The default check type also is `MEDIUM` for static-format `MyISAM` tables, unless `CHANGED` or `FAST` is specified. In that case, the default is `QUICK`. The row scan is skipped for `CHANGED` and `FAST` because the rows are very seldom corrupted.

You can combine check options, as in the following example that does a quick check on the table to determine whether it was closed properly:

```
CHECK TABLE test_table FAST QUICK;
```

**Note**

In some cases, `CHECK TABLE` changes the table. This happens if the table is marked as “corrupted” or “not closed properly” but `CHECK TABLE` does not find any problems in the table. In this case, `CHECK TABLE` marks the table as okay.

If a table is corrupted, it is most likely that the problem is in the indexes and not in the data part. All of the preceding check types check the indexes thoroughly and should thus find most errors.

If you just want to check a table that you assume is okay, you should use no check options or the `QUICK` option. The latter should be used when you are in a hurry and can take the very small risk that `QUICK` does not find an error in the data file. (In most cases, under normal usage, MySQL should find any error in the data file. If this happens, the table is marked as “corrupted” and cannot be used until it is repaired.)

`FAST` and `CHANGED` are mostly intended to be used from a script (for example, to be executed from `cron`) if you want to check tables from time to time. In most cases, `FAST` is to be preferred over `CHANGED`. (The only case when it is not preferred is when you suspect that you have found a bug in the `MyISAM` code.)

`EXTENDED` is to be used only after you have run a normal check but still get strange errors from a table when MySQL tries to update a row or find a row by key. This is very unlikely if a normal check has succeeded.

Use of `CHECK TABLE ... EXTENDED` might influence the execution plan generated by the query optimizer.

Some problems reported by `CHECK TABLE` cannot be corrected automatically:

- `Found row where the auto_increment column has the value 0.`

This means that you have a row in the table where the `AUTO_INCREMENT` index column contains the value 0. (It is possible to create a row where the `AUTO_INCREMENT` column is 0 by explicitly setting the column to 0 with an `UPDATE` statement.)

This is not an error in itself, but could cause trouble if you decide to dump the table and restore it or do an `ALTER TABLE` on the table. In this case, the `AUTO_INCREMENT` column changes value according to the rules of `AUTO_INCREMENT` columns, which could cause problems such as a duplicate-key error.

To get rid of the warning, simply execute an `UPDATE` statement to set the column to some value other than 0.

## InnoDB Tables

The following notes apply to `InnoDB` tables:

- If `CHECK TABLE` finds a problem for an `InnoDB` table, the server shuts down to prevent error propagation. Details of the error will be written to the error log.
- When an `InnoDB` table is stored in its own `.ibd` file in `file-per-table` mode, the first 3 pages of the `.ibd` contain header information rather than table or index data. The `CHECK TABLE` statement does not detect inconsistencies that only affect the header data. To verify the entire contents of an `InnoDB .ibd` file, use the `innochecksum` command.

### 13.7.2.3. CHECKSUM TABLE Syntax

```
CHECKSUM TABLE tbl_name [, tbl_name] ... [ QUICK | EXTENDED ]
```

`CHECKSUM TABLE` reports a `checksum` for the contents of a table. You can use this statement to verify that the contents are the same before and after a backup, rollback, or other operation that is intended to put the data back to a known state. This statement requires the `SELECT [755]` privilege for the table.

## Performance Considerations

By default, the entire table is read row by row and the checksum is calculated. For large tables, this could take a long time, thus you would only perform this operation occasionally. This row-by-row calculation is what you get with the `EXTENDED` clause, with `InnoDB` and all other storage engines other than `MyISAM`, and with `MyISAM` tables not created with the `CHECKSUM=1` clause.

For `MyISAM` tables created with the `CHECKSUM=1` clause, `CHECKSUM TABLE` or `CHECKSUM TABLE ... QUICK` returns the “live” table checksum that can be returned very fast. If the table does not meet all these conditions, the `QUICK` method returns `NULL`. See [Section 13.1.14, “CREATE TABLE Syntax”](#) for the syntax of the `CHECKSUM` clause.

For a nonexistent table, `CHECKSUM TABLE` returns `NULL` and generates a warning.

Prior to MySQL 5.6.4, `CHECKSUM TABLE` returned 0 for partitioned tables unless the `EXTENDED` option was used. (Bug #11933226, Bug #60681)

The checksum value depends on the table row format. If the row format changes, the checksum also changes. For example, the storage format for `VARCHAR` changed between MySQL 4.1 and 5.0, so if a 4.1 table is upgraded to MySQL 5.0, the checksum value may change.



### Important

If the checksums for two tables are different, then it is almost certain that the tables are different in some way. However, because the hashing function used by `CHECKSUM TABLE` is not guaranteed to be collision-free, there is a slight chance that two tables which are not identical can produce the same checksum.

### 13.7.2.4. OPTIMIZE TABLE Syntax

```
OPTIMIZE [NO_WRITE_TO_BINLOG | LOCAL] TABLE
tbl_name [, tbl_name] ...
```

Reorganizes the physical storage of table data and associated index data, to reduce storage space and improve I/O efficiency when accessing the table. The exact changes made to each table depend on the [storage engine](#) used by that table.

Use `OPTIMIZE TABLE` in these cases, depending on the type of table:

- After doing substantial insert, update, or delete operations on an `InnoDB` table that has its own `.ibd` file because it was created with the `innodb_file_per_table` [1718] option enabled. The table and indexes are reorganized, and disk space can be reclaimed for use by the operating system.
- After doing substantial insert, update, or delete operations on columns that are part of a `FULLTEXT` index in an `InnoDB` table. Set the configuration option `innodb_optimize_fulltext_only=1` [1740] first. To keep the index maintenance period to a reasonable time, set the `innodb_ft_num_word_optimize` [1725] option to specify how many words to update in the search index, and run a sequence of `OPTIMIZE TABLE` statements until the search index is fully updated.
- After deleting a large part of a `MyISAM` or `ARCHIVE` table, or making many changes to a `MyISAM` or `ARCHIVE` table with variable-length rows (tables that have `VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT` columns). Deleted rows are maintained in a linked list and subsequent `INSERT` operations reuse old row positions. You can use `OPTIMIZE TABLE` to reclaim the unused space and to defragment the data file. After extensive changes to a table, this statement may also improve performance of statements that use the table, sometimes significantly.

This statement requires `SELECT` [755] and `INSERT` [755] privileges for the table.

`OPTIMIZE TABLE` is also supported for partitioned tables. For information about using this statement with partitioned tables and table partitions, see [Section 17.3.4, “Maintenance of Partitions”](#).

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

`OPTIMIZE TABLE` works *only* for `InnoDB`, `MyISAM`, and `ARCHIVE` tables. By default, it does *not* work for tables created using any other storage engine. You can make `OPTIMIZE TABLE` work on other storage engines by starting `mysqld` with the `--skip-new` option. In this case, `OPTIMIZE TABLE` is just mapped to `ALTER TABLE`.

## InnoDB Details

For `InnoDB` tables, `OPTIMIZE TABLE` is mapped to `ALTER TABLE`, which rebuilds the table to update index statistics and free unused space in the clustered index. This is displayed in the output of `OPTIMIZE TABLE` when you run it on an `InnoDB` table, as shown here:

```
mysql> OPTIMIZE TABLE foo;
+-----+-----+-----+-----+
| Table      | Op       | Msg_type | Msg_text                                     |
+-----+-----+-----+-----+
| test.foo   | optimize | note     | Table does not support optimize, doing recreate + analyze instead |
| test.foo   | optimize | status   | OK                                           |
+-----+-----+-----+-----+
```

`InnoDB` stores data using a page-allocation method and does not suffer from fragmentation in the same way that legacy storage engines (such as `MyISAM`) will. When considering whether or not to run `optimize`, consider the workload of transactions that your server will process:

- Some level of fragmentation is expected. `InnoDB` only fills `pages` 93% full, to leave room for updates without having to split pages.
- Delete operations might leave gaps that leave pages less filled than desired, which could make it worthwhile to optimize the table.
- Updates to rows usually rewrite the data within the same page, depending on the data type and row format, when sufficient space is available. See [Section 5.4.6.5, “How Compression Works for InnoDB Tables”](#) and [Section 5.4.8.1, “Overview of InnoDB Row Storage”](#).
- High-concurrency workloads might leave gaps in indexes over time, as `InnoDB` retains multiple versions of the same data due through its `MVCC` mechanism. See [Section 14.2.3.11, “InnoDB Multi-Versioning”](#).

## MyISAM Details

For `MyISAM` tables, `OPTIMIZE TABLE` works as follows:

1. If the table has deleted or split rows, repair the table.
2. If the index pages are not sorted, sort them.
3. If the table's statistics are not up to date (and the repair could not be accomplished by sorting the index), update them.

## Result Set

`OPTIMIZE TABLE` returns a result set with the following columns.

Column	Value
Table	The table name
Op	Always <code>optimize</code>
Msg_type	<code>status</code> , <code>error</code> , <code>info</code> , <code>note</code> , or <code>warning</code>

Column	Value
<code>Msg_text</code>	An informational message

Note that MySQL [locks the table](#) during the time `OPTIMIZE TABLE` is running.

By default, the server writes `OPTIMIZE TABLE` statements to the binary log so that they replicate to replication slaves. To suppress logging, specify the optional `NO_WRITE_TO_BINLOG` clause or its alias `LOCAL`.

`OPTIMIZE TABLE` does not sort R-tree indexes, such as spatial indexes on `POINT` columns. (Bug #23578)

`OPTIMIZE TABLE` catches and throws any errors that occur while copying table statistics from the old file to the newly created file. For example, if the user ID of the owner of the `.frm`, `.MYD`, or `.MYI` file is different from the user ID of the `mysqld` process, `OPTIMIZE TABLE` generates a "cannot change ownership of the file" error unless `mysqld` is started by the `root` user.

### 13.7.2.5. REPAIR TABLE Syntax

```
REPAIR [NO_WRITE_TO_BINLOG | LOCAL] TABLE
    tbl_name [, tbl_name] ...
    [QUICK] [EXTENDED] [USE_FRM]
```

`REPAIR TABLE` repairs a possibly corrupted table, for certain storage engines only. By default, it has the same effect as `myisamchk --recover tbl_name`.



#### Note

`REPAIR TABLE` only applies to `MyISAM`, `ARCHIVE`, and `CSV` tables. See [Section 14.3, “The MyISAM Storage Engine”](#), and [Section 14.6, “The ARCHIVE Storage Engine”](#), and [Section 14.5, “The CSV Storage Engine”](#)

This statement requires [SELECT \[755\]](#) and [INSERT \[755\]](#) privileges for the table.

`REPAIR TABLE` is supported for partitioned tables. However, the `USE_FRM` option cannot be used with this statement on a partitioned table.

Beginning with MySQL 5.6.11, `gtid_next [2022]` must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

You can use `ALTER TABLE ... REPAIR PARTITION` to repair one or more partitions; for more information, see [Section 13.1.6, “ALTER TABLE Syntax”](#), and [Section 17.3.4, “Maintenance of Partitions”](#).

Although normally you should never have to run `REPAIR TABLE`, if disaster strikes, this statement is very likely to get back all your data from a `MyISAM` table. If your tables become corrupted often, try to find the reason for it, to eliminate the need to use `REPAIR TABLE`. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#), and [Section 14.3.4, “MyISAM Table Problems”](#).



#### Caution

Make a backup of a table before performing a table repair operation; under some circumstances the operation might cause data loss. Possible causes include but are not limited to file system errors. See [Chapter 7, Backup and Recovery](#).



#### Warning

If the server crashes during a `REPAIR TABLE` operation, it is essential after restarting it that you immediately execute another `REPAIR TABLE` statement

for the table before performing any other operations on it. In the worst case, you might have a new clean index file without information about the data file, and then the next operation you perform could overwrite the data file. This is an unlikely but possible scenario that underscores the value of making a backup first.

`REPAIR TABLE` returns a result set with the following columns.

Column	Value
Table	The table name
Op	Always <code>repair</code>
Msg_type	<code>status</code> , <code>error</code> , <code>info</code> , <code>note</code> , or <code>warning</code>
Msg_text	An informational message

The `REPAIR TABLE` statement might produce many rows of information for each repaired table. The last row has a `Msg_type` value of `status` and `Msg_text` normally should be `OK`. If you do not get `OK` for a `MyISAM` table, you should try repairing it with `myisamchk --safe-recover`. (`REPAIR TABLE` does not implement all the options of `myisamchk`.) With `myisamchk --safe-recover`, you can also use options that `REPAIR TABLE` does not support, such as `--max-record-length` [322].

If you use the `QUICK` option, `REPAIR TABLE` tries to repair only the index file, and not the data file. This type of repair is like that done by `myisamchk --recover --quick`.

If you use the `EXTENDED` option, MySQL creates the index row by row instead of creating one index at a time with sorting. This type of repair is like that done by `myisamchk --safe-recover`.

The `USE_FRM` option is available for use if the `.MYI` index file is missing or if its header is corrupted. This option tells MySQL not to trust the information in the `.MYI` file header and to re-create it using information from the `.frm` file. This kind of repair cannot be done with `myisamchk`.



#### Note

Use the `USE_FRM` option *only* if you cannot use regular `REPAIR` modes! Telling the server to ignore the `.MYI` file makes important table metadata stored in the `.MYI` unavailable to the repair process, which can have deleterious consequences:

- The current `AUTO_INCREMENT` value is lost.
- The link to deleted records in the table is lost, which means that free space for deleted records will remain unoccupied thereafter.
- The `.MYI` header indicates whether the table is compressed. If the server ignores this information, it cannot tell that a table is compressed and repair can cause change or loss of table contents. This means that `USE_FRM` should not be used with compressed tables. That should not be necessary, anyway: Compressed tables are read only, so they should not become corrupt.



#### Caution

If you use `USE_FRM` for a table that was created by a different version of the MySQL server than the one you are currently running, `REPAIR TABLE` will not attempt to repair the table. In this case, the result set returned by `REPAIR TABLE` contains a line with a `Msg_type` value of `error` and a `Msg_text` value of `Failed repairing incompatible .FRM file`.

If `USE_FRM` is *not* used, `REPAIR TABLE` checks the table to see whether an upgrade is required. If so, it performs the upgrade, following the same rules as `CHECK TABLE ... FOR UPGRADE`. See



Section 13.7.2.2, “[CHECK TABLE Syntax](#)”, for more information. `REPAIR TABLE` without `USE_FRM` upgrades the `.frm` file to the current version.

By default, the server writes `REPAIR TABLE` statements to the binary log so that they replicate to replication slaves. To suppress logging, use the optional `NO_WRITE_TO_BINLOG` keyword or its alias `LOCAL`.



### Important

In the event that a table on the master becomes corrupted and you run `REPAIR TABLE` on it, any resulting changes to the original table are *not* propagated to slaves.

You may be able to increase `REPAIR TABLE` performance by setting certain system variables. See Section 8.6.3, “[Speed of REPAIR TABLE Statements](#)”.

`REPAIR TABLE` catches and throws any errors that occur while copying table statistics from the old corrupted file to the newly created file. For example, if the user ID of the owner of the `.frm`, `.MYD`, or `.MYI` file is different from the user ID of the `mysqld` process, `REPAIR TABLE` generates a "cannot change ownership of the file" error unless `mysqld` is started by the `root` user.

## 13.7.3. Plugin and User-Defined Function Statements

### 13.7.3.1. CREATE FUNCTION Syntax for User-Defined Functions

```
CREATE [AGGREGATE] FUNCTION function_name RETURNS {STRING|INTEGER|REAL|DECIMAL}
SONAME shared_library_name
```

A user-defined function (UDF) is a way to extend MySQL with a new function that works like a native (built-in) MySQL function such as `ABS()` [1202] or `CONCAT()` [1178].

*function\_name* is the name that should be used in SQL statements to invoke the function. The `RETURNS` clause indicates the type of the function's return value. `DECIMAL` is a legal value after `RETURNS`, but currently `DECIMAL` functions return string values and should be written like `STRING` functions.

*shared\_library\_name* is the basename of the shared object file that contains the code that implements the function. The file must be located in the plugin directory. This directory is given by the value of the `plugin_dir` [525] system variable. For more information, see Section 22.3.2.5, “[Compiling and Installing User-Defined Functions](#)”.

To create a function, you must have the `INSERT` [755] privilege for the `mysql` database. This is necessary because `CREATE FUNCTION` adds a row to the `mysql.func` system table that records the function's name, type, and shared library name. If you do not have this table, you should run the `mysql_upgrade` command to create it. See Section 4.4.7, “[mysql\\_upgrade — Check and Upgrade MySQL Tables](#)”.

An active function is one that has been loaded with `CREATE FUNCTION` and not removed with `DROP FUNCTION`. All active functions are reloaded each time the server starts, unless you start `mysqld` with the `--skip-grant-tables` [439] option. In this case, UDF initialization is skipped and UDFs are unavailable.

For instructions on writing user-defined functions, see Section 22.3.2, “[Adding a New User-Defined Function](#)”. For the UDF mechanism to work, functions must be written in C or C++ (or another language that can use C calling conventions), your operating system must support dynamic loading and you must have compiled `mysqld` dynamically (not statically).

An `AGGREGATE` function works exactly like a native MySQL aggregate (summary) function such as `SUM` or `COUNT()` [1283]. For `AGGREGATE` to work, your `mysql.func` table must contain a `type` column. If

your `mysql.func` table does not have this column, you should run the `mysql_upgrade` program to create it (see [Section 4.4.7](#), “`mysql_upgrade` — Check and Upgrade MySQL Tables”).



#### Note

To upgrade the shared library associated with a UDF, issue a `DROP FUNCTION` statement, upgrade the shared library, and then issue a `CREATE FUNCTION` statement. If you upgrade the shared library first and then use `DROP FUNCTION`, the server may crash.

### 13.7.3.2. `DROP FUNCTION` Syntax

```
DROP FUNCTION function_name
```

This statement drops the user-defined function (UDF) named *function\_name*.

To drop a function, you must have the `DELETE [754]` privilege for the `mysql` database. This is because `DROP FUNCTION` removes a row from the `mysql.func` system table that records the function's name, type, and shared library name.



#### Note

To upgrade the shared library associated with a UDF, issue a `DROP FUNCTION` statement, upgrade the shared library, and then issue a `CREATE FUNCTION` statement. If you upgrade the shared library first and then use `DROP FUNCTION`, the server may crash.

`DROP FUNCTION` is also used to drop stored functions (see [Section 13.1.21](#), “`DROP PROCEDURE and DROP FUNCTION Syntax`”).

### 13.7.3.3. `INSTALL PLUGIN` Syntax

```
INSTALL PLUGIN plugin_name SONAME 'shared_library_name'
```

This statement installs a server plugin. It requires the `INSERT privilege` for the `mysql.plugin` table.

*plugin\_name* is the name of the plugin as defined in the plugin descriptor structure contained in the library file (see [Section 22.2.4.2](#), “`Plugin Data Structures`”). Plugin names are not case sensitive. For maximal compatibility, plugin names should be limited to ASCII letters, digits, and underscore because they are used in C source files, shell command lines, M4 and Bourne shell scripts, and SQL environments.

*shared\_library\_name* is the name of the shared library that contains the plugin code. The name includes the file name extension (for example, `libmyplugin.so`, `libmyplugin.dll`, or `libmyplugin.dylib`).

The shared library must be located in the plugin directory (the directory named by the `plugin_dir [525]` system variable). The library must be in the plugin directory itself, not in a subdirectory. By default, `plugin_dir [525]` is the `plugin` directory under the directory named by the `pkglibdir` configuration variable, but it can be changed by setting the value of `plugin_dir [525]` at server startup. For example, set its value in a `my.cnf` file:

```
[mysqld]
plugin_dir=/path/to/plugin/directory
```

If the value of `plugin_dir [525]` is a relative path name, it is taken to be relative to the MySQL base directory (the value of the `basedir [466]` system variable).

`INSTALL PLUGIN` loads and initializes the plugin code to make the plugin available for use. A plugin is initialized by executing its initialization function, which handles any setup that the plugin must perform

before it can be used. When the server shuts down, it executes the deinitialization function for each plugin that is loaded so that the plugin has a chance to perform any final cleanup.

`INSTALL PLUGIN` also registers the plugin by adding a line that indicates the plugin name and library file name to the `mysql.plugin` table. At server startup, the server loads and initializes any plugin that is listed in the `mysql.plugin` table. This means that a plugin is installed with `INSTALL PLUGIN` only once, not every time the server starts. Plugin loading at startup does not occur if the server is started with the `--skip-grant-tables` [439] option.

A plugin library can contain multiple plugins. For each of them to be installed, use a separate `INSTALL PLUGIN` statement. Each statement names a different plugin, but all of them specify the same library name.

`INSTALL PLUGIN` causes the server to read option (`my.cnf`) files just as during server startup. This enables the plugin to pick up any relevant options from those files. It is possible to add plugin options to an option file even before loading a plugin (if the `loose` prefix is used). It is also possible to uninstall a plugin, edit `my.cnf`, and install the plugin again. Restarting the plugin this way enables it to the new option values without a server restart.

For options that control individual plugin loading at server startup, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#). If you need to load plugins for a single server startup when the `--skip-grant-tables` [439] option is given (which tells the server not to read system tables), use the `--plugin-load` [435] option. See [Section 5.1.3, “Server Command Options”](#).

To remove a plugin, use the `UNINSTALL PLUGIN` statement.

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

To see what plugins are installed, use the `SHOW PLUGINS` statement or query the `INFORMATION_SCHEMA.PLUGINS` table.

If you recompile a plugin library and need to reinstall it, you can use either of the following methods:

- Use `UNINSTALL PLUGIN` to uninstall all plugins in the library, install the new plugin library file in the plugin directory, and then use `INSTALL PLUGIN` to install all plugins in the library. This procedure has the advantage that it can be used without stopping the server. However, if the plugin library contains many plugins, you must issue many `INSTALL PLUGIN` and `UNINSTALL PLUGIN` statements.
- Stop the server, install the new plugin library file in the plugin directory, and restart the server.

#### 13.7.3.4. `UNINSTALL PLUGIN` Syntax

```
UNINSTALL PLUGIN plugin_name
```

This statement removes an installed server plugin. It requires the `DELETE` [754] privilege for the `mysql.plugin` table.

*plugin\_name* must be the name of some plugin that is listed in the `mysql.plugin` table. The server executes the plugin's deinitialization function and removes the row for the plugin from the `mysql.plugin` table, so that subsequent server restarts will not load and initialize the plugin. `UNINSTALL PLUGIN` does not remove the plugin's shared library file.

You cannot uninstall a plugin if any table that uses it is open.

Plugin removal has implications for the use of associated tables. For example, if a full-text parser plugin is associated with a `FULLTEXT` index on the table, uninstalling the plugin makes the table unusable. Any attempt to access the table results in an error. The table cannot even be opened, so you cannot drop an index for which the plugin is used. This means that uninstalling a plugin is something to do with

care unless you do not care about the table contents. If you are uninstalling a plugin with no intention of reinstalling it later and you care about the table contents, you should dump the table with `mysqldump` and remove the `WITH PARSER` clause from the dumped `CREATE TABLE` statement so that you can reload the table later. If you do not care about the table, `DROP TABLE` can be used even if any plugins associated with the table are missing.

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

## 13.7.4. SET Syntax

```
SET variable_assignment [, variable_assignment] ...

variable_assignment:
    user_var_name = expr
  | [GLOBAL | SESSION] system_var_name = expr
  | [@@global. | @@session. | @@] system_var_name = expr
```

The `SET` statement assigns values to different types of variables that affect the operation of the server or your client.

This section describes use of `SET` for assigning values to variables. The `SET` statement can be used to assign values to these types of variables:

- System variables. See [Section 5.1.4, “Server System Variables”](#). System variables also can be set at server startup, as described in [Section 5.1.5, “Using System Variables”](#).

User-defined variables. See [Section 9.4, “User-Defined Variables”](#).

- Stored procedure and function parameters, and stored program local variables. See [Section 13.6.4, “Variables in Stored Programs”](#).

Some variants of `SET` syntax are used in other contexts:

- `SET CHARACTER SET` and `SET NAMES` assign values to character set and collation variables associated with the connection to the server. `SET ONE_SHOT` is used for replication. These variants are described later in this section.
- `SET PASSWORD` assigns account passwords. See [Section 13.7.1.7, “SET PASSWORD Syntax”](#).
- `SET TRANSACTION ISOLATION LEVEL` sets the isolation level for transaction processing. See [Section 13.3.6, “SET TRANSACTION Syntax”](#).

The following discussion shows the different `SET` syntaxes that you can use to set variables. The examples use the `= [1173]` assignment operator, but you can also use the `:= [1172]` assignment operator for this purpose.

A user variable is written as `@var_name` and can be set as follows:

```
SET @var_name = expr;
```

Many system variables are dynamic and can be changed while the server runs by using the `SET` statement. For a list, see [Section 5.1.5.2, “Dynamic System Variables”](#). To change a system variable with `SET`, refer to it as `var_name`, optionally preceded by a modifier:

- To indicate explicitly that a variable is a global variable, precede its name by `GLOBAL` or `@@global..` The `SUPER [756]` privilege is required to set global variables.
- To indicate explicitly that a variable is a session variable, precede its name by `SESSION`, `@@session.`, or `@@`. Setting a session variable requires no special privilege, but a client can change only its own session variables, not those of any other client.

- `LOCAL` and `@@local.` are synonyms for `SESSION` and `@@session..`
- If no modifier is present, `SET` changes the session variable.

A `SET` statement can contain multiple variable assignments, separated by commas. For example, the statement can assign values to a user-defined variable and a system variable. If you set several system variables, the most recent `GLOBAL` or `SESSION` modifier in the statement is used for following variables that have no modifier specified.

Examples:

```
SET sort_buffer_size=10000;
SET @@local.sort_buffer_size=10000;
SET GLOBAL sort_buffer_size=1000000, SESSION sort_buffer_size=1000000;
SET @@sort_buffer_size=1000000;
SET @@global.sort_buffer_size=1000000, @@local.sort_buffer_size=1000000;
```

The `@@var_name` syntax for system variables is supported for compatibility with some other database systems.

If you change a session system variable, the value remains in effect until your session ends or until you change the variable to a different value. The change is not visible to other clients.

If you change a global system variable, the value is remembered and used for new connections until the server restarts. (To make a global system variable setting permanent, you should set it in an option file.) The change is visible to any client that accesses that global variable. However, the change affects the corresponding session variable only for clients that connect after the change. The global variable change does not affect the session variable for any client that is currently connected (not even that of the client that issues the `SET GLOBAL` statement).

To prevent incorrect usage, MySQL produces an error if you use `SET GLOBAL` with a variable that can only be used with `SET SESSION` or if you do not specify `GLOBAL` (or `@@global.`) when setting a global variable.

To set a `SESSION` variable to the `GLOBAL` value or a `GLOBAL` value to the compiled-in MySQL default value, use the `DEFAULT` keyword. For example, the following two statements are identical in setting the session value of `max_join_size` [505] to the global value:

```
SET max_join_size=DEFAULT;
SET @@session.max_join_size=@@global.max_join_size;
```

Not all system variables can be set to `DEFAULT`. In such cases, use of `DEFAULT` results in an error.

It is not permitted to assign the value `DEFAULT` to user-defined variables, stored procedure or function parameters, or stored program local variables. This results in a syntax error for user-defined variables, and also for parameters or local variables as of MySQL 5.6.6.

You can refer to the values of specific global or session system variables in expressions by using one of the `@@`-modifiers. For example, you can retrieve values in a `SELECT` statement like this:

```
SELECT @@global.sql_mode, @@session.sql_mode, @@sql_mode;
```

When you refer to a system variable in an expression as `@@var_name` (that is, when you do not specify `@@global.` or `@@session.`), MySQL returns the session value if it exists and the global value otherwise. (This differs from `SET @@var_name = value`, which always refers to the session value.)



#### Note

Some variables displayed by `SHOW VARIABLES` may not be available using `SELECT @@var_name` syntax; an `Unknown system variable` occurs. As a workaround in such cases, you can use `SHOW VARIABLES LIKE 'var_name'`.

Suffixes for specifying a value multiplier can be used when setting a variable at server startup, but not to set the value with `SET` at runtime. On the other hand, with `SET` you can assign a variable's value using an expression, which is not true when you set a variable at server startup. For example, the first of the following lines is legal at server startup, but the second is not:

```
shell> mysql --max_allowed_packet=16M
shell> mysql --max_allowed_packet=16*1024*1024
```

Conversely, the second of the following lines is legal at runtime, but the first is not:

```
mysql> SET GLOBAL max_allowed_packet=16M;
mysql> SET GLOBAL max_allowed_packet=16*1024*1024;
```

To display system variables names and values, use the `SHOW VARIABLES` statement. (See [Section 13.7.5.40, "SHOW VARIABLES Syntax"](#).)

The following list describes `SET` options that have nonstandard syntax (that is, options that are not set with `name = value` syntax).

- `CHARACTER SET {charset_name | DEFAULT}`

This maps all strings from and to the client with the given mapping. You can add new mappings by editing `sql/convert.cc` in the MySQL source distribution. `SET CHARACTER SET` sets three session system variables: `character_set_client` [467] and `character_set_results` [468] are set to the given character set, and `character_set_connection` [467] to the value of `character_set_database` [468]. See [Section 10.1.4, "Connection Character Sets and Collations"](#).

The default mapping can be restored by using the value `DEFAULT`. The default depends on the server configuration.

`ucs2`, `utf16`, and `utf32` cannot be used as a client character set, which means that they do not work for `SET CHARACTER SET`.

- `NAMES {'charset_name' [COLLATE 'collation_name']} | DEFAULT}`

`SET NAMES` sets the three session system variables `character_set_client` [467], `character_set_connection` [467], and `character_set_results` [468] to the given character set. Setting `character_set_connection` [467] to `charset_name` also sets `collation_connection` [469] to the default collation for `charset_name`. The optional `COLLATE` clause may be used to specify a collation explicitly. See [Section 10.1.4, "Connection Character Sets and Collations"](#).

The default mapping can be restored by using a value of `DEFAULT`. The default depends on the server configuration.

`ucs2`, `utf16`, and `utf32` cannot be used as a client character set, which means that they do not work for `SET NAMES`.

- `ONE_SHOT`

`ONE_SHOT` is for internal use only, has been deprecated since MySQL 5.0, and was removed in MySQL 5.6.1.

## 13.7.5. SHOW Syntax

`SHOW` has many forms that provide information about databases, tables, columns, or status information about the server. This section describes those following:

```
SHOW AUTHORS
SHOW {BINARY | MASTER} LOGS
SHOW BINLOG EVENTS [IN 'log_name'] [FROM pos] [LIMIT [offset,] row_count]
```



```

SHOW CHARACTER SET [like_or_where]
SHOW COLLATION [like_or_where]
SHOW [FULL] COLUMNS FROM tbl_name [FROM db_name] [like_or_where]
SHOW CONTRIBUTORS
SHOW CREATE DATABASE db_name
SHOW CREATE EVENT event_name
SHOW CREATE FUNCTION func_name
SHOW CREATE PROCEDURE proc_name
SHOW CREATE TABLE tbl_name
SHOW CREATE TRIGGER trigger_name
SHOW CREATE VIEW view_name
SHOW DATABASES [like_or_where]
SHOW ENGINE engine_name {STATUS | MUTEX}
SHOW [STORAGE] ENGINES
SHOW ERRORS [LIMIT [offset,] row_count]
SHOW EVENTS
SHOW FUNCTION CODE func_name
SHOW FUNCTION STATUS [like_or_where]
SHOW GRANTS FOR user
SHOW INDEX FROM tbl_name [FROM db_name]
SHOW MASTER STATUS
SHOW OPEN TABLES [FROM db_name] [like_or_where]
SHOW PLUGINS
SHOW PROCEDURE CODE proc_name
SHOW PROCEDURE STATUS [like_or_where]
SHOW PRIVILEGES
SHOW [FULL] PROCESSLIST
SHOW PROFILE [types] [FOR QUERY n] [OFFSET n] [LIMIT n]
SHOW PROFILES
SHOW SLAVE HOSTS
SHOW SLAVE STATUS
SHOW [GLOBAL | SESSION] STATUS [like_or_where]
SHOW TABLE STATUS [FROM db_name] [like_or_where]
SHOW [FULL] TABLES [FROM db_name] [like_or_where]
SHOW TRIGGERS [FROM db_name] [like_or_where]
SHOW [GLOBAL | SESSION] VARIABLES [like_or_where]
SHOW WARNINGS [LIMIT [offset,] row_count]

like_or_where:
    LIKE 'pattern'
    | WHERE expr

```

If the syntax for a given `SHOW` statement includes a `LIKE 'pattern'` [1190] part, '*pattern*' is a string that can contain the SQL “%” and “\_” wildcard characters. The pattern is useful for restricting statement output to matching values.

Several `SHOW` statements also accept a `WHERE` clause that provides more flexibility in specifying which rows to display. See [Section 19.32, “Extensions to SHOW Statements”](#).

Many MySQL APIs (such as PHP) enable you to treat the result returned from a `SHOW` statement as you would a result set from a `SELECT`; see [Chapter 21, Connectors and APIs](#), or your API documentation for more information. In addition, you can work in SQL with results from queries on tables in the `INFORMATION_SCHEMA` database, which you cannot easily do with results from `SHOW` statements. See [Chapter 19, INFORMATION\\_SCHEMA Tables](#).

### 13.7.5.1. SHOW AUTHORS Syntax

```
SHOW AUTHORS
```

The `SHOW AUTHORS` statement displays information about the people who work on MySQL. For each author, it displays `Name`, `Location`, and `Comment` values.

This statement is removed as of MySQL 5.6.8.

### 13.7.5.2. SHOW BINARY LOGS Syntax

```
SHOW BINARY LOGS
SHOW MASTER LOGS
```



Lists the binary log files on the server. This statement is used as part of the procedure described in [Section 13.4.1.1, “PURGE BINARY LOGS Syntax”](#), that shows how to determine which logs can be purged.

```
mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name      | File_size |
+-----+-----+
| binlog.000015 | 724935    |
| binlog.000016 | 733481    |
+-----+-----+
```

`SHOW MASTER LOGS` is equivalent to `SHOW BINARY LOGS`.

In MySQL 5.6.5 and earlier, the `SUPER` [756] privilege was required to use this statement. Starting with MySQL 5.6.6, a user with the `REPLICATION CLIENT` [755] privilege may also execute this statement.

### 13.7.5.3. SHOW BINLOG EVENTS Syntax

```
SHOW BINLOG EVENTS
  [IN 'log_name'] [FROM pos] [LIMIT [offset,] row_count]
```

Shows the events in the binary log. If you do not specify `'log_name'`, the first binary log is displayed.

The `LIMIT` clause has the same syntax as for the `SELECT` statement. See [Section 13.2.9, “SELECT Syntax”](#).



#### Note

Issuing a `SHOW BINLOG EVENTS` with no `LIMIT` clause could start a very time- and resource-consuming process because the server returns to the client the complete contents of the binary log (which includes all statements executed by the server that modify data). As an alternative to `SHOW BINLOG EVENTS`, use the `mysqlbinlog` utility to save the binary log to a text file for later examination and analysis. See [Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#).



#### Note

Some events relating to the setting of user and system variables are not included in the output from `SHOW BINLOG EVENTS`. To get complete coverage of events within a binary log, use `mysqlbinlog`.



#### Note

`SHOW BINLOG EVENTS` does *not* work with relay log files. You can use `SHOW RELAYLOG EVENTS` for this purpose.

### 13.7.5.4. SHOW CHARACTER SET Syntax

```
SHOW CHARACTER SET
  [LIKE 'pattern' | WHERE expr]
```

The `SHOW CHARACTER SET` statement shows all available character sets. The `LIKE` [1190] clause, if present, indicates which character set names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#). For example:

```
mysql> SHOW CHARACTER SET LIKE 'latin%';
+-----+-----+-----+-----+
| Charset | Description          | Default collation | Maxlen |
+-----+-----+-----+-----+
| latin1  | cp1252 West European | latin1_swedish_ci | 1      |
+-----+-----+-----+-----+
```

latin2	ISO 8859-2 Central European	latin2_general_ci	1
latin5	ISO 8859-9 Turkish	latin5_turkish_ci	1
latin7	ISO 8859-13 Baltic	latin7_general_ci	1

The `Maxlen` column shows the maximum number of bytes required to store one character.

### 13.7.5.5. SHOW COLLATION Syntax

```
SHOW COLLATION
[LIKE 'pattern' | WHERE expr]
```

This statement lists collations supported by the server. By default, the output from `SHOW COLLATION` includes all available collations. The `LIKE [1190]` clause, if present, indicates which collation names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#). For example:

```
mysql> SHOW COLLATION LIKE 'latin1%';
+-----+-----+-----+-----+-----+-----+
| Collation          | Charset | Id | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| latin1_german1_ci | latin1  | 5  |         |         | 0       |
| latin1_swedish_ci| latin1  | 8  | Yes     | Yes     | 0       |
| latin1_danish_ci  | latin1  | 15 |         |         | 0       |
| latin1_german2_ci | latin1  | 31 |         | Yes     | 2       |
| latin1_bin        | latin1  | 47 |         | Yes     | 0       |
| latin1_general_ci | latin1  | 48 |         |         | 0       |
| latin1_general_cs | latin1  | 49 |         |         | 0       |
| latin1_spanish_ci | latin1  | 94 |         |         | 0       |
+-----+-----+-----+-----+-----+-----+
```

The `Collation` and `Charset` columns indicate the names of the collation and the character set with which it is associated. `Id` is the collation ID. `Default` indicates whether the collation is the default for its character set. `Compiled` indicates whether the character set is compiled into the server. `Sortlen` is related to the amount of memory required to sort strings expressed in the character set.

To see the default collation for each character set, use the following statement. `Default` is a reserved word, so to use it as an identifier, it must be quoted as such:

```
mysql> SHOW COLLATION WHERE `Default` = 'Yes';
+-----+-----+-----+-----+-----+-----+
| Collation          | Charset | Id | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| big5_chinese_ci   | big5    | 1  | Yes     | Yes     | 1       |
| dec8_swedish_ci   | dec8    | 3  | Yes     | Yes     | 1       |
| cp850_general_ci  | cp850   | 4  | Yes     | Yes     | 1       |
| hp8_english_ci    | hp8     | 6  | Yes     | Yes     | 1       |
| koi8r_general_ci  | koi8r   | 7  | Yes     | Yes     | 1       |
| latin1_swedish_ci | latin1  | 8  | Yes     | Yes     | 1       |
...

```

### 13.7.5.6. SHOW COLUMNS Syntax

```
SHOW [FULL] COLUMNS {FROM | IN} tbl_name [{FROM | IN} db_name]
[LIKE 'pattern' | WHERE expr]
```

`SHOW COLUMNS` displays information about the columns in a given table. It also works for views. The `LIKE [1190]` clause, if present, indicates which column names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

`SHOW COLUMNS` displays information only for those columns for which you have some privilege.

```
mysql> SHOW COLUMNS FROM City;
+-----+-----+-----+-----+-----+-----+
| Field      | Type      | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+-----+-----+
| Id      | int(11) | NO   | PRI | NULL | auto_increment |
| Name    | char(35) | NO   |     |     |                 |
| Country | char(3)  | NO   | UNI |     |                 |
| District | char(20) | YES  | MUL |     |                 |
| Population | int(11) | NO   |     | 0   |                 |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)

```

If the data types differ from what you expect them to be based on a `CREATE TABLE` statement, note that MySQL sometimes changes data types when you create or alter a table. The conditions under which this occurs are described in [Section 13.1.14.3, “Silent Column Specification Changes”](#).

The `FULL` keyword causes the output to include the column collation and comments, as well as the privileges you have for each column.

You can use `db_name.tbl_name` as an alternative to the `tbl_name FROM db_name` syntax. In other words, these two statements are equivalent:

```

mysql> SHOW COLUMNS FROM mytable FROM mydb;
mysql> SHOW COLUMNS FROM mydb.mytable;

```

`SHOW COLUMNS` displays the following values for each table column:

`Field` indicates the column name.

`Type` indicates the column data type.

`Collation` indicates the collation for nonbinary string columns, or `NULL` for other columns. This value is displayed only if you use the `FULL` keyword.

The `Null` field contains `YES` if `NULL` values can be stored in the column, `NO` if not.

The `Key` field indicates whether the column is indexed:

- If `Key` is empty, the column either is not indexed or is indexed only as a secondary column in a multiple-column, nonunique index.
- If `Key` is `PRI`, the column is a `PRIMARY KEY` or is one of the columns in a multiple-column `PRIMARY KEY`.
- If `Key` is `UNI`, the column is the first column of a `UNIQUE` index. (A `UNIQUE` index permits multiple `NULL` values, but you can tell whether the column permits `NULL` by checking the `Null` field.)
- If `Key` is `MUL`, the column is the first column of a nonunique index in which multiple occurrences of a given value are permitted within the column.

If more than one of the `Key` values applies to a given column of a table, `Key` displays the one with the highest priority, in the order `PRI`, `UNI`, `MUL`.

A `UNIQUE` index may be displayed as `PRI` if it cannot contain `NULL` values and there is no `PRIMARY KEY` in the table. A `UNIQUE` index may display as `MUL` if several columns form a composite `UNIQUE` index; although the combination of the columns is unique, each column can still hold multiple occurrences of a given value.

The `Default` field indicates the default value that is assigned to the column. This is `NULL` if the column has an explicit default of `NULL`, or if the column definition has no `DEFAULT` clause.

The `Extra` field contains any additional information that is available about a given column. The value is nonempty in these cases: `auto_increment` for columns that have the `AUTO_INCREMENT` attribute; `on update CURRENT_TIMESTAMP` for `TIMESTAMP` or `DATETIME` columns that have the `ON UPDATE CURRENT_TIMESTAMP` attribute.

`Privileges` indicates the privileges you have for the column. This value is displayed only if you use the `FULL` keyword.

`Comment` indicates any comment the column has. This value is displayed only if you use the `FULL` keyword.

`SHOW FIELDS` is a synonym for `SHOW COLUMNS`. You can also list a table's columns with the `mysqlshow db_name tbl_name` command.

The `DESCRIBE` statement provides information similar to `SHOW COLUMNS`. See [Section 13.8.1, “DESCRIBE Syntax”](#).

The `SHOW CREATE TABLE`, `SHOW TABLE STATUS`, and `SHOW INDEX` statements also provide information about tables. See [Section 13.7.5, “SHOW Syntax”](#).

### 13.7.5.7. SHOW CONTRIBUTORS Syntax

```
SHOW CONTRIBUTORS
```

The `SHOW CONTRIBUTORS` statement displays information about the people who contribute to MySQL source or to causes that we support. For each contributor, it displays `Name`, `Location`, and `Comment` values.

This statement is removed as of MySQL 5.6.8.

### 13.7.5.8. SHOW CREATE DATABASE Syntax

```
SHOW CREATE {DATABASE | SCHEMA} [IF NOT EXISTS] db_name
```

Shows the `CREATE DATABASE` statement that creates the given database. If the `SHOW` statement includes an `IF NOT EXISTS` clause, the output too includes such a clause. `SHOW CREATE SCHEMA` is a synonym for `SHOW CREATE DATABASE`.

```
mysql> SHOW CREATE DATABASE test\G
***** 1. row *****
      Database: test
Create Database: CREATE DATABASE `test`
                  /*!40100 DEFAULT CHARACTER SET latin1 */

mysql> SHOW CREATE SCHEMA test\G
***** 1. row *****
      Database: test
Create Database: CREATE DATABASE `test`
                  /*!40100 DEFAULT CHARACTER SET latin1 */
```

`SHOW CREATE DATABASE` quotes table and column names according to the value of the `sql_quote_show_create` [546] option. See [Section 5.1.4, “Server System Variables”](#).

### 13.7.5.9. SHOW CREATE EVENT Syntax

```
SHOW CREATE EVENT event_name
```

This statement displays the `CREATE EVENT` statement needed to re-create a given event. It requires the `EVENT` [754] privilege for the database from which the event is to be shown. For example (using the same event `e_daily` defined and then altered in [Section 13.7.5.19, “SHOW EVENTS Syntax”](#)):

```
mysql> SHOW CREATE EVENT test.e_daily\G
***** 1. row *****
      Event: e_daily
      sql_mode:
      time_zone: SYSTEM
Create Event: CREATE EVENT `e_daily`
              ON SCHEDULE EVERY 1 DAY
              STARTS CURRENT_TIMESTAMP + INTERVAL 6 HOUR
```

```

ON COMPLETION NOT PRESERVE
ENABLE
COMMENT 'Saves total number of sessions then
        clears the table each day'
DO BEGIN
    INSERT INTO site_activity.totals (time, total)
        SELECT CURRENT_TIMESTAMP, COUNT(*)
        FROM site_activity.sessions;
    DELETE FROM site_activity.sessions;
END
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci

```

`character_set_client` [467] is the session value of the `character_set_client` [467] system variable when the event was created. `collation_connection` [469] is the session value of the `collation_connection` [469] system variable when the event was created. `Database Collation` is the collation of the database with which the event is associated.

Note that the output reflects the current status of the event (`ENABLE`) rather than the status with which it was created.

### 13.7.5.10. SHOW CREATE FUNCTION Syntax

```
SHOW CREATE FUNCTION func_name
```

This statement is similar to `SHOW CREATE PROCEDURE` but for stored functions. See [Section 13.7.5.11, “SHOW CREATE PROCEDURE Syntax”](#).

### 13.7.5.11. SHOW CREATE PROCEDURE Syntax

```
SHOW CREATE PROCEDURE proc_name
```

This statement is a MySQL extension. It returns the exact string that can be used to re-create the named stored procedure. A similar statement, `SHOW CREATE FUNCTION`, displays information about stored functions (see [Section 13.7.5.10, “SHOW CREATE FUNCTION Syntax”](#)).

Both statements require that you be the owner of the routine or have `SELECT` access to the `mysql.proc` table. If you do not have privileges for the routine itself, the value displayed for the `Create Procedure` or `Create Function` field will be `NULL`.

```

mysql> SHOW CREATE PROCEDURE test.simpleproc\G
***** 1. row *****
      Procedure: simpleproc
      sql_mode:
      Create Procedure: CREATE PROCEDURE `simpleproc` (OUT param1 INT)
                        BEGIN
                        SELECT COUNT(*) INTO param1 FROM t;
                        END
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci

mysql> SHOW CREATE FUNCTION test.hello\G
***** 1. row *****
      Function: hello
      sql_mode:
      Create Function: CREATE FUNCTION `hello` (s CHAR(20))
                       RETURNS CHAR(50)
                       RETURN CONCAT('Hello, ',s, '!')
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci

```

`character_set_client` [467] is the session value of the `character_set_client` [467] system variable when the routine was created. `collation_connection` [469] is the session value of

the `collation_connection` [469] system variable when the routine was created. `Database Collation` is the collation of the database with which the routine is associated.

### 13.7.5.12. SHOW CREATE TABLE Syntax

```
SHOW CREATE TABLE tbl_name
```

Shows the `CREATE TABLE` statement that creates the given table. To use this statement, you must have some privilege for the table. This statement also works with views.

```
mysql> SHOW CREATE TABLE t\G
***** 1. row *****
      Table: t
Create Table: CREATE TABLE t (
  id INT(11) default NULL auto_increment,
  s char(60) default NULL,
  PRIMARY KEY (id)
) ENGINE=MyISAM
```

`SHOW CREATE TABLE` quotes table and column names according to the value of the `sql_quote_show_create` [546] option. See Section 5.1.4, “Server System Variables”.

### 13.7.5.13. SHOW CREATE TRIGGER Syntax

```
SHOW CREATE TRIGGER trigger_name
```

This statement shows a `CREATE TRIGGER` statement that creates the given trigger.

```
mysql> SHOW CREATE TRIGGER ins_sum\G
***** 1. row *****
      Trigger: ins_sum
      sql_mode:
SQL Original Statement: CREATE DEFINER=`bob`@`localhost` TRIGGER ins_sum
                        BEFORE INSERT ON account
                        FOR EACH ROW SET @sum = @sum + NEW.amount
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci
```

You can also obtain information about trigger objects from `INFORMATION_SCHEMA`, which contains a `TRIGGERS` table. See Section 19.27, “The `INFORMATION_SCHEMA TRIGGERS` Table”.

### 13.7.5.14. SHOW CREATE VIEW Syntax

```
SHOW CREATE VIEW view_name
```

This statement shows a `CREATE VIEW` statement that creates the given view.

```
mysql> SHOW CREATE VIEW v\G
***** 1. row *****
      View: v
      Create View: CREATE ALGORITHM=UNDEFINED
                  DEFINER=`bob`@`localhost`
                  SQL SECURITY DEFINER VIEW
                  `v` AS select 1 AS `a`,2 AS `b`
character_set_client: latin1
collation_connection: latin1_swedish_ci
```

`character_set_client` [467] is the session value of the `character_set_client` [467] system variable when the view was created. `collation_connection` [469] is the session value of the `collation_connection` [469] system variable when the view was created.

Use of `SHOW CREATE VIEW` requires the `SHOW VIEW` [756] privilege and the `SELECT` [755] privilege for the view in question.

You can also obtain information about view objects from `INFORMATION_SCHEMA`, which contains a `VIEWS` table. See [Section 19.29](#), “The `INFORMATION_SCHEMA VIEWS` Table”.

MySQL lets you use different `sql_mode` [545] settings to tell the server the type of SQL syntax to support. For example, you might use the `ANSI` [608] SQL mode to ensure MySQL correctly interprets the standard SQL concatenation operator, the double bar (`||`), in your queries. If you then create a view that concatenates items, you might worry that changing the `sql_mode` [545] setting to a value different from `ANSI` [608] could cause the view to become invalid. But this is not the case. No matter how you write out a view definition, MySQL always stores it the same way, in a canonical form. Here is an example that shows how the server changes a double bar concatenation operator to a `CONCAT()` [1178] function:

```
mysql> SET sql_mode = 'ANSI';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE VIEW test.v AS SELECT 'a' || 'b' as coll;
Query OK, 0 rows affected (0.01 sec)

mysql> SHOW CREATE VIEW test.v\G
***** 1. row *****
          View: v
          Create View: CREATE VIEW "v" AS select concat('a','b') AS "coll"
...
1 row in set (0.00 sec)
```

The advantage of storing a view definition in canonical form is that changes made later to the value of `sql_mode` [545] will not affect the results from the view. However an additional consequence is that comments prior to `SELECT` are stripped from the definition by the server.

### 13.7.5.15. SHOW DATABASES Syntax

```
SHOW {DATABASES | SCHEMAS}
     [LIKE 'pattern' | WHERE expr]
```

`SHOW DATABASES` lists the databases on the MySQL server host. `SHOW SCHEMAS` is a synonym for `SHOW DATABASES`. The `LIKE` [1190] clause, if present, indicates which database names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32](#), “Extensions to `SHOW` Statements”.

You see only those databases for which you have some kind of privilege, unless you have the global `SHOW DATABASES` privilege. You can also get this list using the `mysqlshow` command.

If the server was started with the `--skip-show-database` [441] option, you cannot use this statement at all unless you have the `SHOW DATABASES` [756] privilege.

MySQL implements databases as directories in the data directory, so this statement simply lists directories in that location. However, the output may include names of directories that do not correspond to actual databases.

### 13.7.5.16. SHOW ENGINE Syntax

```
SHOW ENGINE engine_name {STATUS | MUTEX}
```

`SHOW ENGINE` displays operational information about a storage engine. The following statements currently are supported:

```
SHOW ENGINE INNODB STATUS
SHOW ENGINE INNODB MUTEX
SHOW ENGINE PERFORMANCE_SCHEMA STATUS
```

`SHOW ENGINE INNODB STATUS` displays extensive information from the standard `InnoDB` Monitor about the state of the `InnoDB` storage engine. For information about the standard monitor and other `InnoDB` Monitors that provide information about `InnoDB` processing, see [Section 14.2.4.4](#), “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”.



`SHOW ENGINE INNODB MUTEX` displays InnoDB mutex statistics. The statement displays the following fields:

- `Type`

Always `InnoDB`.

- `Name`

The source file where the mutex is implemented, and the line number in the file where the mutex is created. The line number may change depending on your version of MySQL.

- `Status`

The mutex status. This field displays several values if `UNIV_DEBUG` was defined at MySQL compilation time (for example, in `include/univ.i` in the `InnoDB` part of the MySQL source tree). If `UNIV_DEBUG` was not defined, the statement displays only the `os_waits` value. In the latter case (without `UNIV_DEBUG`), the information on which the output is based is insufficient to distinguish regular mutexes and mutexes that protect rw-locks (which permit multiple readers or a single writer). Consequently, the output may appear to contain multiple rows for the same mutex.

- `count` indicates how many times the mutex was requested.
- `spin_waits` indicates how many times the spinlock had to run.
- `spin_rounds` indicates the number of spinlock rounds. (`spin_rounds` divided by `spin_waits` provides the average round count.)
- `os_waits` indicates the number of operating system waits. This occurs when the spinlock did not work (the mutex was not locked during the spinlock and it was necessary to yield to the operating system and wait).
- `os_yields` indicates the number of times a the thread trying to lock a mutex gave up its timeslice and yielded to the operating system (on the presumption that permitting other threads to run will free the mutex so that it can be locked).
- `os_wait_times` indicates the amount of time (in ms) spent in operating system waits, if the `timed_mutexes` [557] system variable is 1 (`ON`). If `timed_mutexes` [557] is 0 (`OFF`), timing is disabled, so `os_wait_times` is 0. `timed_mutexes` [557] is off by default.

Information from this statement can be used to diagnose system problems. For example, large values of `spin_waits` and `spin_rounds` may indicate scalability problems.

Use `SHOW ENGINE PERFORMANCE_SCHEMA STATUS` to inspect the internal operation of the Performance Schema code:

```
mysql> SHOW ENGINE PERFORMANCE_SCHEMA STATUS\G
...
***** 3. row *****
  Type: performance_schema
  Name: events_waits_history.row_size
  Status: 76
***** 4. row *****
  Type: performance_schema
  Name: events_waits_history.row_count
  Status: 10000
***** 5. row *****
  Type: performance_schema
  Name: events_waits_history.memory
  Status: 760000
...
***** 57. row *****
  Type: performance_schema
  Name: performance_schema.memory
  Status: 26459600
```

...

This statement is intended to help the DBA understand the effects that different Performance Schema options have on memory requirements.

`Name` values consist of two parts, which name an internal buffer and a buffer attribute, respectively. Interpret buffer names as follows:

- An internal buffer that is not exposed as a table is named within parentheses. Examples: `(pfs_cond_class).row_size`, `(pfs_mutex_class).memory`.
- An internal buffer that is exposed as a table in the `performance_schema` database is named after the table, without parentheses. Examples: `events_waits_history.row_size`, `mutex_instances.row_count`.
- A value that applies to the Performance Schema as a whole begins with `performance_schema`. Example: `performance_schema.memory`.

Buffer attributes have these meanings:

- `row_size` is the size of the internal record used by the implementation, such as the size of a row in a table. `row_size` values cannot be changed.
- `row_count` is the number of internal records, such as the number of rows in a table. `row_count` values can be changed using Performance Schema configuration options.
- For a table, `tbl_name.memory` is the product of `row_size` and `row_count`. For the Performance Schema as a whole, `performance_schema.memory` is the sum of all the memory used (the sum of all other `memory` values).

In some cases, there is a direct relationship between a Performance Schema configuration parameter and a `SHOW ENGINE` value. For example, `events_waits_history_long.row_count` corresponds to `performance_schema_events_waits_history_long_size` [2310]. In other cases, the relationship is more complex. For example, `events_waits_history.row_count` corresponds to `performance_schema_events_waits_history_size` [2310] (the number of rows per thread) multiplied by `performance_schema_max_thread_instances` [2316] (the number of threads).

### 13.7.5.17. SHOW ENGINES Syntax

```
SHOW [STORAGE] ENGINES
```

`SHOW ENGINES` displays status information about the server's storage engines. This is particularly useful for checking whether a storage engine is supported, or to see what the default engine is.

```
mysql> SHOW ENGINES\G
***** 1. row *****
  Engine: MEMORY
  Support: YES
  Comment: Hash based, stored in memory, useful for temporary tables
  Transactions: NO
    XA: NO
  Savepoints: NO
***** 2. row *****
  Engine: MyISAM
  Support: YES
  Comment: MyISAM storage engine
  Transactions: NO
    XA: NO
  Savepoints: NO
***** 3. row *****
  Engine: InnoDB
  Support: DEFAULT
  Comment: Supports transactions, row-level locking, and foreign keys
  Transactions: YES
```

```

      XA: YES
      Savepoints: YES
***** 4. row *****
      Engine: EXAMPLE
      Support: YES
      Comment: Example storage engine
Transactions: NO
      XA: NO
      Savepoints: NO
***** 5. row *****
      Engine: ARCHIVE
      Support: YES
      Comment: Archive storage engine
Transactions: NO
      XA: NO
      Savepoints: NO
***** 6. row *****
      Engine: CSV
      Support: YES
      Comment: CSV storage engine
Transactions: NO
      XA: NO
      Savepoints: NO
***** 7. row *****
      Engine: BLACKHOLE
      Support: YES
      Comment: /dev/null storage engine (anything you write >
to it disappears)
Transactions: NO
      XA: NO
      Savepoints: NO
***** 8. row *****
      Engine: FEDERATED
      Support: YES
      Comment: Federated MySQL storage engine
Transactions: NO
      XA: NO
      Savepoints: NO
***** 9. row *****
      Engine: MRG_MYISAM
      Support: YES
      Comment: Collection of identical MyISAM tables
Transactions: NO
      XA: NO
      Savepoints: NO

```

The output from `SHOW ENGINES` may vary according to the MySQL version used and other factors. The values shown in the `Support` column indicate the server's level of support for the storage engine, as shown in the following table.

Value	Meaning
YES	The engine is supported and is active
DEFAULT	Like YES, plus this is the default engine
NO	The engine is not supported
DISABLED	The engine is supported but has been disabled

A value of `NO` means that the server was compiled without support for the engine, so it cannot be enabled at runtime.

A value of `DISABLED` occurs either because the server was started with an option that disables the engine, or because not all options required to enable it were given. In the latter case, the error log file should contain a reason indicating why the option is disabled. See [Section 5.2.2, "The Error Log"](#).

You might also see `DISABLED` for a storage engine if the server was compiled to support it, but was started with a `--skip-engine_name` option.

All MySQL servers support [MyISAM](#) tables, because [MyISAM](#) is the default storage engine. It is not possible to disable [MyISAM](#).

The [Transactions](#), [XA](#), and [Savepoints](#) columns indicate whether the storage engine supports transactions, XA transactions, and savepoints, respectively.

### 13.7.5.18. SHOW ERRORS Syntax

```
SHOW ERRORS [LIMIT [offset,] row_count]
SHOW COUNT(*) ERRORS
```

This statement is similar to [SHOW WARNINGS](#), except that it displays information only for errors, rather than for errors, warnings, and notes.

The [LIMIT](#) clause has the same syntax as for the [SELECT](#) statement. See [Section 13.2.9, “SELECT Syntax”](#).

The [SHOW COUNT\(\\*\) ERRORS](#) statement displays the number of errors. You can also retrieve this number from the [error\\_count](#) [480] variable:

```
SHOW COUNT(*) ERRORS;
SELECT @@error_count;
```

[SHOW ERRORS](#) and [error\\_count](#) [480] apply only to errors, not warnings or notes. In other respects, they are similar to [SHOW WARNINGS](#) and [warning\\_count](#) [564]. In particular, [SHOW ERRORS](#) cannot display information for more than [max\\_error\\_count](#) [504] messages, and [error\\_count](#) [480] can exceed the value of [max\\_error\\_count](#) [504] if the number of errors exceeds [max\\_error\\_count](#) [504].

For more information, see [Section 13.7.5.41, “SHOW WARNINGS Syntax”](#).

### 13.7.5.19. SHOW EVENTS Syntax

```
SHOW EVENTS [{FROM | IN} schema_name]
           [LIKE 'pattern' | WHERE expr]
```

This statement displays information about Event Manager events. It requires the [EVENT](#) [754] privilege for the database from which the events are to be shown.

In its simplest form, [SHOW EVENTS](#) lists all of the events in the current schema:

```
mysql> SELECT CURRENT_USER(), SCHEMA();
+-----+-----+
| CURRENT_USER() | SCHEMA() |
+-----+-----+
| jon@ghidora   | myschema |
+-----+-----+
1 row in set (0.00 sec)

mysql> SHOW EVENTS\G
***** 1. row *****
      Db: myschema
      Name: e_daily
      Definer: jon@ghidora
      Time zone: SYSTEM
      Type: RECURRING
      Execute at: NULL
      Interval value: 10
      Interval field: SECOND
      Starts: 2006-02-09 10:41:23
      Ends: NULL
      Status: ENABLED
      Originator: 0
character_set_client: latin1
collation_connection: latin1_swedish_ci
```

```
Database Collation: latin1_swedish_ci
```

To see events for a specific schema, use the `FROM` clause. For example, to see events for the `test` schema, use the following statement:

```
SHOW EVENTS FROM test;
```

The `LIKE` [1190] clause, if present, indicates which event names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in Section 19.32, “Extensions to `SHOW Statements`”.

`SHOW EVENTS` output has the following columns:

- `Db`: The schema (database) on which the event is defined.
- `Name`: The name of the event.
- `Time zone`: The event time zone, which is the time zone used for scheduling the event and that is in effect within the event as it executes. The default value is `SYSTEM`.
- `Definer`: The account of the user who created the event, in `'user_name'@'host_name'` format.
- `Type`: The event repetition type, either `ONE TIME` (transient) or `RECURRING` (repeating).
- `Execute At`: The date and time when a transient event is set to execute. Shown as a `DATETIME` value.

For a recurring event, the value of this column is always `NULL`.

- `Interval Value`: For a recurring event, the number of intervals to wait between event executions.

For a transient event, the value of this column is always `NULL`.

- `Interval Field`: The time units used for the interval which a recurring event waits before repeating.

For a transient event, the value of this column is always `NULL`.

- `Starts`: The start date and time for a recurring event. This is displayed as a `DATETIME` value, and is `NULL` if no start date and time are defined for the event.

For a transient event, this column is always `NULL`.

- `Ends`: The end date and time for a recurring event. This is displayed as a `DATETIME` value, and defaults to `NULL` if no end date and time is defined for the event.

For a transient event, this column is always `NULL`.

- `Status`: The event status. One of `ENABLED`, `DISABLED`, or `SLAVESIDE_DISABLED`.

`SLAVESIDE_DISABLED` indicates that the creation of the event occurred on another MySQL server acting as a replication master and replicated to the current MySQL server which is acting as a slave, but the event is not presently being executed on the slave.

- `Originator`: The server ID of the MySQL server on which the event was created. Defaults to 0.
- `character_set_client` [467] is the session value of the `character_set_client` [467] system variable when the routine was created. `collation_connection` [469] is the session value of the `collation_connection` [469] system variable when the routine was created. `Database Collation` is the collation of the database with which the routine is associated.

For more information about `SLAVE_DISABLED` and the `Originator` column, see Section 16.4.1.10, “Replication of Invoked Features”.

The event action statement is not shown in the output of `SHOW EVENTS`. Use `SHOW CREATE EVENT` or the `INFORMATION_SCHEMA.EVENTS` table.

Times displayed by `SHOW EVENTS` are given in the event time zone, as discussed in [Section 18.4.4, “Event Metadata”](#).

The columns in the output of `SHOW EVENTS` are similar to, but not identical to the columns in the `INFORMATION_SCHEMA.EVENTS` table. See [Section 19.7, “The INFORMATION\\_SCHEMA EVENTS Table”](#).

### 13.7.5.20. SHOW FUNCTION CODE Syntax

```
SHOW FUNCTION CODE func_name
```

This statement is similar to `SHOW PROCEDURE CODE` but for stored functions. See [Section 13.7.5.28, “SHOW PROCEDURE CODE Syntax”](#).

### 13.7.5.21. SHOW FUNCTION STATUS Syntax

```
SHOW FUNCTION STATUS
  [LIKE 'pattern' | WHERE expr]
```

This statement is similar to `SHOW PROCEDURE STATUS` but for stored functions. See [Section 13.7.5.29, “SHOW PROCEDURE STATUS Syntax”](#).

### 13.7.5.22. SHOW GRANTS Syntax

```
SHOW GRANTS [FOR user]
```

This statement lists the `GRANT` statement or statements that must be issued to duplicate the privileges that are granted to a MySQL user account. The account is named using the same format as for the `GRANT` statement; for example, `'jeffrey'@'localhost'`. If you specify only the user name part of the account name, a host name part of `'%'` is used. For additional information about specifying account names, see [Section 13.7.1.4, “GRANT Syntax”](#).

```
mysql> SHOW GRANTS FOR 'root'@'localhost';
+-----+-----+
| Grants for root@localhost |
+-----+-----+
| GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' WITH GRANT OPTION |
+-----+-----+
```

To list the privileges granted to the account that you are using to connect to the server, you can use any of the following statements:

```
SHOW GRANTS;
SHOW GRANTS FOR CURRENT_USER;
SHOW GRANTS FOR CURRENT_USER();
```

If `SHOW GRANTS FOR CURRENT_USER` (or any of the equivalent syntaxes) is used in `DEFINER` context, such as within a stored procedure that is defined with `SQL SECURITY DEFINER`, the grants displayed are those of the definer and not the invoker.

`SHOW GRANTS` displays only the privileges granted explicitly to the named account. Other privileges might be available to the account, but they are not displayed. For example, if an anonymous account exists, the named account might be able to use its privileges, but `SHOW GRANTS` will not display them.

`SHOW GRANTS` requires the `SELECT [755]` privilege for the `mysql` database, except to see the privileges for the current user.

### 13.7.5.23. SHOW INDEX Syntax

```
SHOW {INDEX | INDEXES | KEYS}
  {FROM | IN} tbl_name
  [{FROM | IN} db_name]
```

```
[WHERE expr]
```

`SHOW INDEX` returns table index information. The format resembles that of the `SQLStatistics` call in ODBC. This statement requires some privilege for any column in the table.

`SHOW INDEX` returns the following fields:

- `Table`

The name of the table.

- `Non_unique`

0 if the index cannot contain duplicates, 1 if it can.

- `Key_name`

The name of the index. If the index is the primary key, the name is always `PRIMARY`.

- `Seq_in_index`

The column sequence number in the index, starting with 1.

- `Column_name`

The column name.

- `Collation`

How the column is sorted in the index. In MySQL, this can have values "A" (Ascending) or `NULL` (Not sorted).

- `Cardinality`

An estimate of the number of unique values in the index. This is updated by running `ANALYZE TABLE` or `myisamchk -a`. `Cardinality` is counted based on statistics stored as integers, so the value is not necessarily exact even for small tables. The higher the cardinality, the greater the chance that MySQL uses the index when doing joins.

- `Sub_part`

The number of indexed characters if the column is only partly indexed, `NULL` if the entire column is indexed.

- `Packed`

Indicates how the key is packed. `NULL` if it is not.

- `Null`

Contains `YES` if the column may contain `NULL` values and `' '` if not.

- `Index_type`

The index method used (`BTREE`, `FULLTEXT`, `HASH`, `RTREE`).

- `Comment`

Information about the index not described in its own column, such as `disabled` if the index is disabled.

- `Index_comment`

Any comment provided for the index with a `COMMENT` attribute when the index was created.



You can use `db_name.tbl_name` as an alternative to the `tbl_name FROM db_name` syntax. These two statements are equivalent:

```
SHOW INDEX FROM mytable FROM mydb;
SHOW INDEX FROM mydb.mytable;
```

The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

You can also list a table's indexes with the `mysqlshow -k db_name tbl_name` command.

### 13.7.5.24. SHOW MASTER STATUS Syntax

```
SHOW MASTER STATUS
```

This statement provides status information about the binary log files of the master. It requires either the `SUPER` [756] or `REPLICATION CLIENT` [755] privilege.

Example:

```
mysql> SHOW MASTER STATUS\G
***** 1. row *****
      File: master-bin.000002
      Position: 1307
      Binlog_Do_DB: test
      Binlog_Ignore_DB: manual, mysql
      Executed_Gtid_Set: 3E11FA47-71CA-11E1-9E33-C80AA9429562:1-5
1 row in set (0.00 sec)
```

The `Executed_Gtid_Set` column was added in MySQL 5.6.5. When global transaction IDs are in use, this column shows the set of GTIDs for transactions that have been executed on the master. This is the same as the master's value for the global `gtid_executed` [2021] system variable (named `gtid_done` [2020] prior to MySQL 5.6.9).

### 13.7.5.25. SHOW OPEN TABLES Syntax

```
SHOW OPEN TABLES [{FROM | IN} db_name]
  [LIKE 'pattern' | WHERE expr]
```

`SHOW OPEN TABLES` lists the non-`TEMPORARY` tables that are currently open in the table cache. See [Section 8.4.3.1, “How MySQL Opens and Closes Tables”](#). The `FROM` clause, if present, restricts the tables shown to those present in the `db_name` database. The `LIKE` [1190] clause, if present, indicates which table names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

`SHOW OPEN TABLES` returns the following columns:

- `Database`

The database containing the table.

- `Table`

The table name.

- `In_use`

The number of table locks or lock requests there are for the table. For example, if one client acquires a lock for a table using `LOCK TABLE t1 WRITE`, `In_use` will be 1. If another client issues `LOCK TABLE t1 WRITE` while the table remains locked, the client will block waiting for the lock, but the lock request causes `In_use` to be 2. If the count is zero, the table is open but not currently being used. `In_use` is also increased by the `HANDLER ... OPEN` statement and decreased by `HANDLER ... CLOSE`.

- `Name_locked`

Whether the table name is locked. Name locking is used for operations such as dropping or renaming tables.

If you have no privileges for a table, it does not show up in the output from `SHOW OPEN TABLES`.

### 13.7.5.26. SHOW PLUGINS Syntax

```
SHOW PLUGINS
```

`SHOW PLUGINS` displays information about server plugins. Plugin information is also available in the `INFORMATION_SCHEMA.PLUGINS` table. See [Section 19.15](#), “The `INFORMATION_SCHEMA.PLUGINS` Table”.

Example of `SHOW PLUGINS` output:

```
mysql> SHOW PLUGINS\G
***** 1. row *****
  Name: binlog
  Status: ACTIVE
  Type: STORAGE ENGINE
  Library: NULL
  License: GPL
***** 2. row *****
  Name: CSV
  Status: ACTIVE
  Type: STORAGE ENGINE
  Library: NULL
  License: GPL
***** 3. row *****
  Name: MEMORY
  Status: ACTIVE
  Type: STORAGE ENGINE
  Library: NULL
  License: GPL
***** 4. row *****
  Name: MyISAM
  Status: ACTIVE
  Type: STORAGE ENGINE
  Library: NULL
  License: GPL
...
```

`SHOW PLUGINS` returns the following columns:

- **Name:** The name used to refer to the plugin in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`.
- **Status:** The plugin status, one of `ACTIVE`, `INACTIVE`, `DISABLED`, or `DELETED`.
- **Type:** The type of plugin, such as `STORAGE ENGINE`, `INFORMATION_SCHEMA`, or `AUTHENTICATION`.
- **Library:** The name of the plugin shared object file. This is the name used to refer to the plugin file in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`. This file is located in the directory named by the `plugin_dir` [525] system variable. If the library name is `NULL`, the plugin is compiled in and cannot be uninstalled with `UNINSTALL PLUGIN`.
- **License:** How the plugin is licensed; for example, `GPL`.

For plugins installed with `INSTALL PLUGIN`, the `Name` and `Library` values are also registered in the `mysql.plugin` table.

For information about plugin data structures that form the basis of the information displayed by `SHOW PLUGINS`, see [Section 22.2](#), “The MySQL Plugin API”.

### 13.7.5.27. SHOW PRIVILEGES Syntax

## SHOW PRIVILEGES

`SHOW PRIVILEGES` shows the list of system privileges that the MySQL server supports. The exact list of privileges depends on the version of your server.

```
mysql> SHOW PRIVILEGES\G
***** 1. row *****
Privilege: Alter
  Context: Tables
  Comment: To alter the table
***** 2. row *****
Privilege: Alter routine
  Context: Functions,Procedures
  Comment: To alter or drop stored functions/procedures
***** 3. row *****
Privilege: Create
  Context: Databases,Tables,Indexes
  Comment: To create new databases and tables
***** 4. row *****
Privilege: Create routine
  Context: Databases
  Comment: To use CREATE FUNCTION/PROCEDURE
***** 5. row *****
Privilege: Create temporary tables
  Context: Databases
  Comment: To use CREATE TEMPORARY TABLE
...

```

Privileges belonging to a specific user are displayed by the `SHOW GRANTS` statement. See [Section 13.7.5.22, "SHOW GRANTS Syntax"](#), for more information.

### 13.7.5.28. SHOW PROCEDURE CODE Syntax

```
SHOW PROCEDURE CODE proc_name
```

This statement is a MySQL extension that is available only for servers that have been built with debugging support. It displays a representation of the internal implementation of the named stored procedure. A similar statement, `SHOW FUNCTION CODE`, displays information about stored functions (see [Section 13.7.5.20, "SHOW FUNCTION CODE Syntax"](#)).

Both statements require that you be the owner of the routine or have `SELECT` access to the `mysql.proc` table.

If the named routine is available, each statement produces a result set. Each row in the result set corresponds to one "instruction" in the routine. The first column is `Pos`, which is an ordinal number beginning with 0. The second column is `Instruction`, which contains an SQL statement (usually changed from the original source), or a directive which has meaning only to the stored-routine handler.

```
mysql> DELIMITER //
mysql> CREATE PROCEDURE p1 ()
-> BEGIN
->   DECLARE fanta INT DEFAULT 55;
->   DROP TABLE t2;
->   LOOP
->     INSERT INTO t3 VALUES (fanta);
->   END LOOP;
-> END//
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW PROCEDURE CODE p1//
+-----+-----+
| Pos | Instruction |
+-----+-----+
| 0 | set fanta@0 55 |
| 1 | stmt 9 "DROP TABLE t2" |
| 2 | stmt 5 "INSERT INTO t3 VALUES (fanta)" |
| 3 | jump 2 |
+-----+-----+
4 rows in set (0.00 sec)

```

In this example, the nonexecutable `BEGIN` and `END` statements have disappeared, and for the `DECLARE variable_name` statement, only the executable part appears (the part where the default is assigned). For each statement that is taken from source, there is a code word `stmt` followed by a type (9 means `DROP`, 5 means `INSERT`, and so on). The final row contains an instruction `jump 2`, meaning `GOTO instruction #2`.

### 13.7.5.29. SHOW PROCEDURE STATUS Syntax

```
SHOW PROCEDURE STATUS
  [LIKE 'pattern' | WHERE expr]
```

This statement is a MySQL extension. It returns characteristics of a stored procedure, such as the database, name, type, creator, creation and modification dates, and character set information. A similar statement, `SHOW FUNCTION STATUS`, displays information about stored functions (see [Section 13.7.5.21, “SHOW FUNCTION STATUS Syntax”](#)).

The `LIKE [1190]` clause, if present, indicates which procedure or function names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

```
mysql> SHOW PROCEDURE STATUS LIKE 'spl'\G
***** 1. row *****
      Db: test
      Name: spl
      Type: PROCEDURE
      Definer: testuser@localhost
      Modified: 2004-08-03 15:29:37
      Created: 2004-08-03 15:29:37
      Security_type: DEFINER
      Comment:
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci
```

`character_set_client [467]` is the session value of the `character_set_client [467]` system variable when the routine was created. `collation_connection [469]` is the session value of the `collation_connection [469]` system variable when the routine was created. `Database Collation` is the collation of the database with which the routine is associated.

You can also get information about stored routines from the `ROUTINES` table in `INFORMATION_SCHEMA`. See [Section 19.19, “The INFORMATION\\_SCHEMA ROUTINES Table”](#).

### 13.7.5.30. SHOW PROCESSLIST Syntax

```
SHOW [FULL] PROCESSLIST
```

`SHOW PROCESSLIST` shows you which threads are running. You can also get this information from the `INFORMATION_SCHEMA PROCESSLIST` table or the `mysqladmin processlist` command. If you have the `PROCESS [755]` privilege, you can see all threads. Otherwise, you can see only your own threads (that is, threads associated with the MySQL account that you are using). If you do not use the `FULL` keyword, only the first 100 characters of each statement are shown in the `Info` field.

Process information is also available from the `performance_schema.threads` table. However, access to `threads` does not require a mutex and has minimal impact on server performance. `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` have negative performance consequences because they require a mutex. `threads` also shows information about background threads, which `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` do not. This means that `threads` can be used to monitor activity the other thread information sources cannot.

The `SHOW PROCESSLIST` statement is very useful if you get the “too many connections” error message and want to find out what is going on. MySQL reserves one extra connection to be used by accounts that have the `SUPER [756]` privilege, to ensure that administrators should always be able to connect and check the system (assuming that you are not giving this privilege to all your users).

Threads can be killed with the `KILL` statement. See [Section 13.7.6.4, “KILL Syntax”](#).

Here is an example of `SHOW PROCESSLIST` output:

```
mysql> SHOW FULL PROCESSLIST\G
***** 1. row *****
Id: 1
User: system user
Host:
db: NULL
Command: Connect
Time: 1030455
State: Waiting for master to send event
Info: NULL
***** 2. row *****
Id: 2
User: system user
Host:
db: NULL
Command: Connect
Time: 1004
State: Has read all relay log; waiting for the slave
      I/O thread to update it
Info: NULL
***** 3. row *****
Id: 3112
User: replikator
Host: artemis:2204
db: NULL
Command: Binlog Dump
Time: 2144
State: Has sent all binlog to slave; waiting for binlog to be updated
Info: NULL
***** 4. row *****
Id: 3113
User: replikator
Host: iconnect2:45781
db: NULL
Command: Binlog Dump
Time: 2086
State: Has sent all binlog to slave; waiting for binlog to be updated
Info: NULL
***** 5. row *****
Id: 3123
User: stefan
Host: localhost
db: apollon
Command: Query
Time: 0
State: NULL
Info: SHOW FULL PROCESSLIST
5 rows in set (0.00 sec)
```

The columns produced by `SHOW PROCESSLIST` have the following meanings:

- `Id`

The connection identifier.

- `User`

The MySQL user who issued the statement. If this is `system user`, it refers to a nonclient thread spawned by the server to handle tasks internally. This could be the I/O or SQL thread used on replication slaves or a delayed-row handler. `unauthenticated user` refers to a thread that has become associated with a client connection but for which authentication of the client user has not yet been done. `event_scheduler` refers to the thread that monitors scheduled events. For `system user`, there is no host specified in the `Host` column.

- `Host`

The host name of the client issuing the statement (except for `system user` where there is no host). `SHOW PROCESSLIST` reports the host name for TCP/IP connections in `host_name:client_port` format to make it easier to determine which client is doing what.

- `db`

The default database, if one is selected, otherwise `NULL`.

- `Command`

The type of command the thread is executing. For descriptions for thread commands, see [Section 8.12.5, “Examining Thread Information”](#). The value of this column corresponds to the `COM_xxx` commands of the client/server protocol and `Com_xxx` status variables. See [Section 5.1.6, “Server Status Variables”](#)

- `Time`

The time in seconds that the thread has been in its current state.

- `State`

An action, event, or state that indicates what the thread is doing. Descriptions for `State` values can be found at [Section 8.12.5, “Examining Thread Information”](#).

Most states correspond to very quick operations. If a thread stays in a given state for many seconds, there might be a problem that needs to be investigated.

For the `SHOW PROCESSLIST` statement, the value of `State` is `NULL`.

- `Info`

The statement the thread is executing, or `NULL` if it is not executing any statement. The statement might be the one sent to the server, or an innermost statement if the statement executes other statements. For example, if a `CALL` statement executes a stored procedure that is executing a `SELECT` statement, the `Info` value shows the `SELECT` statement.

### 13.7.5.31. SHOW PROFILE Syntax

```
SHOW PROFILE [type [, type] ... ]
  [FOR QUERY n]
  [LIMIT row_count [OFFSET offset]]

type:
  ALL
  | BLOCK IO
  | CONTEXT SWITCHES
  | CPU
  | IPC
  | MEMORY
  | PAGE FAULTS
  | SOURCE
  | SWAPS
```

The `SHOW PROFILE` and `SHOW PROFILES` statements display profiling information that indicates resource usage for statements executed during the course of the current session.



#### Note

These statements are deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release. Use the Performance Schema instead; see [Chapter 20, MySQL Performance Schema](#).

Profiling is controlled by the `profiling` [526] session variable, which has a default value of 0 (`OFF`). Profiling is enabled by setting `profiling` [526] to 1 or `ON`:

```
mysql> SET profiling = 1;
```

`SHOW PROFILES` displays a list of the most recent statements sent to the server. The size of the list is controlled by the `profiling_history_size` [526] session variable, which has a default value of 15. The maximum value is 100. Setting the value to 0 has the practical effect of disabling profiling.

All statements are profiled except `SHOW PROFILE` and `SHOW PROFILES`, so you will find neither of those statements in the profile list. Malformed statements are profiled. For example, `SHOW PROFILING` is an illegal statement, and a syntax error occurs if you try to execute it, but it will show up in the profiling list.

`SHOW PROFILE` displays detailed information about a single statement. Without the `FOR QUERY n` clause, the output pertains to the most recently executed statement. If `FOR QUERY n` is included, `SHOW PROFILE` displays information for statement `n`. The values of `n` correspond to the `Query_ID` values displayed by `SHOW PROFILES`.

The `LIMIT row_count` clause may be given to limit the output to `row_count` rows. If `LIMIT` is given, `OFFSET offset` may be added to begin the output `offset` rows into the full set of rows.

By default, `SHOW PROFILE` displays `Status` and `Duration` columns. The `Status` values are like the `State` values displayed by `SHOW PROCESSLIST`, although there might be some minor differences in interpretation for the two statements for some status values (see [Section 8.12.5, “Examining Thread Information”](#)).

Optional `type` values may be specified to display specific additional types of information:

- `ALL` displays all information
- `BLOCK IO` displays counts for block input and output operations
- `CONTEXT SWITCHES` displays counts for voluntary and involuntary context switches
- `CPU` displays user and system CPU usage times
- `IPC` displays counts for messages sent and received
- `MEMORY` is not currently implemented
- `PAGE FAULTS` displays counts for major and minor page faults
- `SOURCE` displays the names of functions from the source code, together with the name and line number of the file in which the function occurs
- `SWAPS` displays swap counts

Profiling is enabled per session. When a session ends, its profiling information is lost.

```
mysql> SELECT @@profiling;
+-----+
| @@profiling |
+-----+
|           0 |
+-----+
1 row in set (0.00 sec)

mysql> SET profiling = 1;
Query OK, 0 rows affected (0.00 sec)

mysql> DROP TABLE IF EXISTS t1;
Query OK, 0 rows affected, 1 warning (0.00 sec)

mysql> CREATE TABLE T1 (id INT);
Query OK, 0 rows affected (0.01 sec)
```



```
mysql> SHOW PROFILES;
```

Query_ID	Duration	Query
0	0.000088	SET PROFILING = 1
1	0.000136	DROP TABLE IF EXISTS t1
2	0.011947	CREATE TABLE t1 (id INT)

```
3 rows in set (0.00 sec)
```

```
mysql> SHOW PROFILE;
```

Status	Duration
checking permissions	0.000040
creating table	0.000056
After create	0.011363
query end	0.000375
freeing items	0.000089
logging slow query	0.000019
cleaning up	0.000005

```
7 rows in set (0.00 sec)
```

```
mysql> SHOW PROFILE FOR QUERY 1;
```

Status	Duration
query end	0.000107
freeing items	0.000008
logging slow query	0.000015
cleaning up	0.000006

```
4 rows in set (0.00 sec)
```

```
mysql> SHOW PROFILE CPU FOR QUERY 2;
```

Status	Duration	CPU_user	CPU_system
checking permissions	0.000040	0.000038	0.000002
creating table	0.000056	0.000028	0.000028
After create	0.011363	0.000217	0.001571
query end	0.000375	0.000013	0.000028
freeing items	0.000089	0.000010	0.000014
logging slow query	0.000019	0.000009	0.000010
cleaning up	0.000005	0.000003	0.000002

```
7 rows in set (0.00 sec)
```



### Note

Profiling is only partially functional on some architectures. For values that depend on the `getrusage()` system call, `NULL` is returned on systems such as Windows that do not support the call. In addition, profiling is per process and not per thread. This means that activity on threads within the server other than your own may affect the timing information that you see.

You can also get profiling information from the `PROFILING` table in `INFORMATION_SCHEMA`. See [Section 19.17, “The `INFORMATION\_SCHEMA PROFILING Table`”](#). For example, the following queries produce the same result:

```
SHOW PROFILE FOR QUERY 2;
```

```
SELECT STATE, FORMAT(DURATION, 6) AS DURATION
FROM INFORMATION_SCHEMA.PROFILING
WHERE QUERY_ID = 2 ORDER BY SEQ;
```

### 13.7.5.32. SHOW PROFILES Syntax

```
SHOW PROFILES
```

The `SHOW PROFILES` statement, together with `SHOW PROFILE`, displays profiling information that indicates resource usage for statements executed during the course of the current session. For more information, see [Section 13.7.5.31, “SHOW PROFILE Syntax”](#).



#### Note

These statements are deprecated as of MySQL 5.6.7 and will be removed in a future MySQL release. Use the Performance Schema instead; see [Chapter 20, MySQL Performance Schema](#).

### 13.7.5.33. SHOW RELAYLOG EVENTS Syntax

```
SHOW RELAYLOG EVENTS
  [IN 'log_name'] [FROM pos] [LIMIT [offset,] row_count]
```

Shows the events in the relay log of a replication slave. If you do not specify `'log_name'`, the first relay log is displayed. This statement has no effect on the master.

The `LIMIT` clause has the same syntax as for the `SELECT` statement. See [Section 13.2.9, “SELECT Syntax”](#).



#### Note

Issuing a `SHOW RELAYLOG EVENTS` with no `LIMIT` clause could start a very time- and resource-consuming process because the server returns to the client the complete contents of the relay log (including all statements modifying data that have been received by the slave).



#### Note

Some events relating to the setting of user and system variables are not included in the output from `SHOW RELAYLOG EVENTS`. To get complete coverage of events within a relay log, use `mysqlbinlog`.

### 13.7.5.34. SHOW SLAVE HOSTS Syntax

```
SHOW SLAVE HOSTS
```

Displays a list of replication slaves currently registered with the master.

The list is displayed on any server (not just the master server). The output looks like this:

```
mysql> SHOW SLAVE HOSTS;
+-----+-----+-----+-----+-----+
| Server_id | Host      | Port | Master_id | Slave_UUID |
+-----+-----+-----+-----+-----+
| 192168010 | iconnect2 | 3306 | 192168011 | 14cb6624-7f93-11e0-b2c0-c80aa9429562 |
| 1921680101 | athena    | 3306 | 192168011 | 07af4990-f41f-11df-a566-7ac56fdaf645 |
+-----+-----+-----+-----+-----+
```

- `Server_id`: The unique server ID of the slave server, as configured in the server's option file, or on the command line with `--server-id=value` [1955].
- `Host`: The host name of the slave server, as configured in the server's option file, or on the command line with `--report-host=host_name` [1980]. Note that this can differ from the machine name as configured in the operating system.
- `Port`: The port the slave server is listening on.

In MySQL 5.6.5 and later, a zero in this column means that the slave port (`--report-port` [1981]) was not set. Prior to MySQL 5.6.5, 3306 was used as the default in such cases (Bug #13333431).

- `Master_id`: The unique server ID of the master server that the slave server is replicating from.

- `Slave_UUID`: The globally unique ID of this slave, as generated on the slave and found in the slave's `auto.cnf` file.

This column was added in MySQL 5.6.0.

### 13.7.5.35. SHOW SLAVE STATUS Syntax

```
SHOW SLAVE STATUS
```

This statement provides status information on essential parameters of the slave threads. It requires either the `SUPER` [756] or `REPLICATION CLIENT` [755] privilege.

If you issue this statement using the `mysql` client, you can use a `\G` statement terminator rather than a semicolon to obtain a more readable vertical layout:

```
mysql> SHOW SLAVE STATUS\G
***** 1. row *****
      Slave_IO_State: Waiting for master to send event
      Master_Host: localhost
      Master_User: root
      Master_Port: 13000
      Connect_Retry: 60
      Master_Log_File: master-bin.000002
      Read_Master_Log_Pos: 1307
      Relay_Log_File: slave-relay-bin.000003
      Relay_Log_Pos: 1508
      Relay_Master_Log_File: master-bin.000002
      Slave_IO_Running: Yes
      Slave_SQL_Running: Yes
      Replicate_Do_DB:
      Replicate_Ignore_DB:
      Replicate_Do_Table:
      Replicate_Ignore_Table:
      Replicate_Wild_Do_Table:
      Replicate_Wild_Ignore_Table:
      Last_Errno: 0
      Last_Error:
      Skip_Counter: 0
      Exec_Master_Log_Pos: 1307
      Relay_Log_Space: 1858
      Until_Condition: None
      Until_Log_File:
      Until_Log_Pos: 0
      Master_SSL_Allowed: No
      Master_SSL_CA_File:
      Master_SSL_CA_Path:
      Master_SSL_Cert:
      Master_SSL_Cipher:
      Master_SSL_Key:
      Seconds_Behind_Master: 0
      Master_SSL_Verify_Server_Cert: No
      Last_IO_Errno: 0
      Last_IO_Error:
      Last_SQL_Errno: 0
      Last_SQL_Error:
      Replicate_Ignore_Server_Ids:
      Master_Server_Id: 1
      Master_UUID: 3e11fa47-71ca-11e1-9e33-c80aa9429562
      Master_Info_File: /var/mysql/2/data/master.info
      SQL_Delay: 0
      SQL_Remaining_Delay: NULL
      Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it
      Master_Retry_Count: 10
      Master_Bind:
      Last_IO_Error_Timestamp:
      Last_SQL_Error_Timestamp:
      Master_SSL_Crl:
      Master_SSL_Crlpath:
      Retrieved_Gtid_Set: 3e11fa47-71ca-11e1-9e33-c80aa9429562:1-5
      Executed_Gtid_Set: 3e11fa47-71ca-11e1-9e33-c80aa9429562:1-5
      Auto_Position: 1
```

```
1 row in set (0.00 sec)
```

The following list describes the fields returned by `SHOW SLAVE STATUS`. For additional information about interpreting their meanings, see [Section 8.12.5.6, “Replication Slave I/O Thread States”](#).

- `Slave_IO_State`

A copy of the `State` field of the `SHOW PROCESSLIST` output for the slave I/O thread. This tells you what the thread is doing: trying to connect to the master, waiting for events from the master, reconnecting to the master, and so on. For a listing of possible states, see [Section 8.12.5.6, “Replication Slave I/O Thread States”](#).

- `Master_Host`

The master host that the slave is connected to.

- `Master_User`

The user name of the account used to connect to the master.

- `Master_Port`

The port used to connect to the master.

- `Connect_Retry`

The number of seconds between connect retries (default 60). This can be set with the `CHANGE MASTER TO` statement.

- `Master_Log_File`

The name of the master binary log file from which the I/O thread is currently reading.

- `Read_Master_Log_Pos`

The position in the current master binary log file up to which the I/O thread has read.

- `Relay_Log_File`

The name of the relay log file from which the SQL thread is currently reading and executing.

- `Relay_Log_Pos`

The position in the current relay log file up to which the SQL thread has read and executed.

- `Relay_Master_Log_File`

The name of the master binary log file containing the most recent event executed by the SQL thread.

- `Slave_IO_Running`

Whether the I/O thread is started and has connected successfully to the master. Internally, the state of this thread is represented by one of the following three values:

- **`MYSQL_SLAVE_NOT_RUN.`** The slave I/O thread is not running. For this state, `Slave_IO_Running` is `No`.
- **`MYSQL_SLAVE_RUN_NOT_CONNECT.`** The slave I/O thread is running, but is not connected to a replication master. For this state, `Slave_IO_Running` depends on the server version as shown in the following table.

MySQL Version	Slave_IO_Running
4.1 (4.1.13 and earlier); 5.0 (5.0.11 and earlier)	Yes

MySQL Version	Slave_IO_Running
4.1 (4.1.14 and later); 5.0 (5.0.12 and later)	No
5.1 (5.1.45 and earlier); 5.4	No
5.1 (5.1.46 and later); 5.5; 5.6	Connecting

- **MYSQL\_SLAVE\_RUN\_CONNECT.** The slave I/O thread is running, and is connected to a replication master. For this state, `Slave_IO_Running` is `Yes`.

The value of the `Slave_running` [598] system status variable corresponds with this value.

- `Slave_SQL_Running`

Whether the SQL thread is started.

- `Replicate_Do_DB`, `Replicate_Ignore_DB`

The lists of databases that were specified with the `--replicate-do-db` [1976] and `--replicate-ignore-db` [1977] options, if any.

- `Replicate_Do_Table`, `Replicate_Ignore_Table`, `Replicate_Wild_Do_Table`, `Replicate_Wild_Ignore_Table`

The lists of tables that were specified with the `--replicate-do-table` [1978], `--replicate-ignore-table` [1978], `--replicate-wild-do-table` [1979], and `--replicate-wild-ignore-table` [1980] options, if any.

- `Last_Errno`, `Last_Error`

These columns are aliases for `Last_SQL_Errno` and `Last_SQL_Error`.

Issuing `RESET MASTER` or `RESET SLAVE` resets the values shown in these columns.



#### Note

When the slave SQL thread receives an error, it reports the error first, then stops the SQL thread. This means that there is a small window of time during which `SHOW SLAVE STATUS` shows a nonzero value for `Last_SQL_Errno` even though `Slave_SQL_Running` still displays `Yes`.

- `Skip_Counter`

The current value of the `sql_slave_skip_counter` [1996] system variable. See [Section 13.4.2.4, “SET GLOBAL sql\\_slave\\_skip\\_counter Syntax”](#).

- `Exec_Master_Log_Pos`

The position in the current master binary log file to which the SQL thread has read and executed, marking the start of the next transaction or event to be processed. You can use this value with the `CHANGE MASTER TO` statement's `MASTER_LOG_POS` option when starting a new slave from an existing slave, so that the new slave reads from this point. The coordinates given by (`Relay_Master_Log_File`, `Exec_Master_Log_Pos`) in the master's binary log correspond to the coordinates given by (`Relay_Log_File`, `Relay_Log_Pos`) in the relay log.

When using a multi-threaded slave (by setting `slave_parallel_workers` [1993] to a nonzero value in MySQL 5.6.3 and later), the value in this column actually represents a “low-water” mark, before which no uncommitted transactions remain. Because the current implementation allows execution of transactions on different databases in a different order on the slave than on the master, this is not necessarily the position of the most recently executed transaction.

- `Relay_Log_Space`

The total combined size of all existing relay log files.

- `Until_Condition`, `Until_Log_File`, `Until_Log_Pos`

The values specified in the `UNTIL` clause of the `START SLAVE` statement.

`Until_Condition` has these values:

- `None` if no `UNTIL` clause was specified
- `Master` if the slave is reading until a given position in the master's binary log
- `Relay` if the slave is reading until a given position in its relay log

`Until_Log_File` and `Until_Log_Pos` indicate the log file name and position that define the coordinates at which the SQL thread stops executing.

- `Master_SSL_Allowed`, `Master_SSL_CA_File`, `Master_SSL_CA_Path`, `Master_SSL_Cert`, `Master_SSL_Cipher`, `Master_SSL_CRL_File`, `Master_SSL_CRL_Path`, `Master_SSL_Key`, `Master_SSL_Verify_Server_Cert`

These fields show the SSL parameters used by the slave to connect to the master, if any.

`Master_SSL_Allowed` has these values:

- `Yes` if an SSL connection to the master is permitted
- `No` if an SSL connection to the master is not permitted
- `Ignored` if an SSL connection is permitted but the slave server does not have SSL support enabled

The values of the other SSL-related fields correspond to the values of the `MASTER_SSL_CA`, `MASTER_SSL_CAPATH`, `MASTER_SSL_CERT`, `MASTER_SSL_CIPHER`, `MASTER_SSL_CRL`, `MASTER_SSL_CRLPATH`, `MASTER_SSL_KEY`, and `MASTER_SSL_VERIFY_SERVER_CERT` options to the `CHANGE MASTER TO` statement. See [Section 13.4.2.1, "CHANGE MASTER TO Syntax"](#).

`Master_SSL_CRL_File` and `Master_SSL_CRL_Path` were added in MySQL 5.6.3.

- `Seconds_Behind_Master`

This field is an indication of how "late" the slave is:

- When the slave SQL thread is actively processing updates, this field is the number of seconds that have elapsed since the timestamp of the most recent event on the master executed by that thread.
- When the SQL thread has caught up to the slave I/O thread and is idle waiting for more events from the I/O thread, this field is zero.

In essence, this field measures the time difference in seconds between the slave SQL thread and the slave I/O thread.

If the network connection between master and slave is fast, the slave I/O thread is very close to the master, so this field is a good approximation of how late the slave SQL thread is compared to the master. If the network is slow, this is *not* a good approximation; the slave SQL thread may quite often be caught up with the slow-reading slave I/O thread, so `Seconds_Behind_Master` often shows a value of 0, even if the I/O thread is late compared to the master. In other words, *this column is useful only for fast networks*.

This time difference computation works even if the master and slave do not have identical clock times, provided that the difference, computed when the slave I/O thread starts, remains constant

from then on. Any changes—including NTP updates—can lead to clock skews that can make calculation of `Seconds_Behind_Master` less reliable.

In MySQL 5.6.9 and later, this field is `NULL` (undefined or unknown) if the slave SQL thread is not running, or if the SQL thread has consumed all of the relay log and the slave I/O thread is not running. Previously, this field was `NULL` if the slave SQL thread or the slave I/O thread was not running or was not connected to the master. (Bug #12946333) For example, if (prior to MySQL 5.6.9) the slave I/O thread was running but was not connected to the master and was sleeping for the number of seconds given by the `CHANGE MASTER TO` statement or `--master-connect-retry [1987]` option (default 60) before reconnecting, the value was `NULL`. Now in such cases, the connection to the master is not tested; instead, if the I/O thread is running but the relay log is exhausted, `Seconds_Behind_Master` is set to 0.

The value of `Seconds_Behind_Master` is based on the timestamps stored in events, which are preserved through replication. This means that if a master M1 is itself a slave of M0, any event from M1's binary log that originates from M0's binary log has M0's timestamp for that event. This enables MySQL to replicate `TIMESTAMP` successfully. However, the problem for `Seconds_Behind_Master` is that if M1 also receives direct updates from clients, the `Seconds_Behind_Master` value randomly fluctuates because sometimes the last event from M1 originates from M0 and sometimes is the result of a direct update on M1.

When using a multi-threaded slave (MySQL 5.6.3 and later), you should keep in mind that this value is based on `Exec_Master_Log_Pos`, and so may not reflect the position of the most recently committed transaction.

- `Last_IO_Errno`, `Last_IO_Error`

The error number and error message of the last error that caused the I/O thread to stop. An error number of 0 and message of the empty string mean “no error.” If the `Last_IO_Error` value is not empty, the error values also appear in the slave's error log.

Beginning with MySQL 5.6.1, I/O error information includes a timestamp showing when the most recent I/O thread error occurred. This timestamp uses the format `YYMMDD HH:MM:SS`.

Previous to MySQL 5.6.3, this timestamp was prefixed to the error message text shown in the `Last_IO_Error` column. In MySQL 5.6.3 and later, the timestamp appears instead in the `Last_SQL_Error_Timestamp` column.

Issuing `RESET MASTER` or `RESET SLAVE` resets the values shown in these columns.

- `Last_SQL_Errno`, `Last_SQL_Error`

The error number and error message of the last error that caused the SQL thread to stop. An error number of 0 and message of the empty string mean “no error.” If the `Last_SQL_Error` value is not empty, the error values also appear in the slave's error log.

Beginning with MySQL 5.6.1, SQL error information includes a timestamp showing when the most recent SQL thread error occurred. This timestamp uses the format `YYMMDD HH:MM:SS`.

Previous to MySQL 5.6.3, this timestamp was prefixed to the error message text shown in the `Last_SQL_Error` column. In MySQL 5.6.3 and later, the timestamp appears instead in the `Last_SQL_Error_Timestamp` column.

Issuing `RESET MASTER` or `RESET SLAVE` resets the values shown in these columns.

- `Replicate_Ignore_Server_Ids`

In MySQL 5.6, you can tell a slave to ignore events from 0 or more masters using the `IGNORE_SERVER_IDS` option in a `CHANGE MASTER TO` statement. This is normally of interest only when using a circular or other multi-master replication setup.



The message shown for `Replicate_Ignore_Server_Ids` consists of a space-delimited list of one or more numbers, the first value indicating the number of servers to be ignored; if not 0 (the default), this server-count value is followed by the actual server IDs. For example, if a `CHANGE MASTER TO` statement containing the `IGNORE_SERVER_IDS = (2,6,9)` option has been issued to tell a slave to ignore masters having the server ID 2, 6, or 9, that information appears as shown here:

```
Replicate_Ignore_Server_Ids: 3 2 6 9
```

- `Master_Server_Id`

The `server_id` [539] value from the master.

- `Master_UUID`

The `server_uuid` [1955] value from the master. This field was added in MySQL 5.6.0.

- `Master_Info_File`

The location of the `master.info` file. This field was added in MySQL 5.6.0.

- `SQL_Delay`

The number of seconds that the slave must lag the master. This field was added in MySQL 5.6.0.

- `SQL_Remaining_Delay`

When `Slave_SQL_Running_State` is `Waiting until MASTER_DELAY seconds after master executed event`, this field contains the number of seconds left of the delay. At other times, this field is `NULL`. This field was added in MySQL 5.6.0.

- `Slave_SQL_Running_State`

The state of the SQL thread (analogous to `Slave_IO_State`). The value is identical to the `State` value of the SQL thread as displayed by `SHOW PROCESSLIST`; [Section 8.12.5.7, “Replication Slave SQL Thread States”](#), provides a listing of possible states. This field was added in MySQL 5.6.0.

- `Master_Retry_Count`

The number of times the slave can attempt to reconnect to the master in the event of a lost connection. This value can be set using the `MASTER_RETRY_COUNT` option in a `CHANGE MASTER TO` statement (preferred) or the older `--master-retry-count` [1973] server option (still supported for backward compatibility). This field was added in MySQL 5.6.1.

- `Master_Bind`

Shows the network interface that the slave is bound to, if any, set using the `MASTER_BIND` option for the `CHANGE MASTER TO` statement.

This column was added in MySQL 5.6.2.

- `Last_IO_Error_Timestamp`

A timestamp in `YYMMDD HH:MM:SS` format that shows when the most recent I/O error took place.

This column was added in MySQL 5.6.3. Previously in MySQL 5.6, this timestamp was prepended to the error text shown in `Last_IO_Error`.

- `Last_SQL_Error_Timestamp`

A timestamp in `YYMMDD HH:MM:SS` format that shows when the last SQL error occurred.

This column was added in MySQL 5.6.3. Previously in MySQL 5.6, this timestamp was prepended to the error text shown in `Last_SQL_Error`.

- `Retrieved_Gtid_Set`

The set of global transaction IDs corresponding to all transactions received by this slave. Empty if GTIDs are not in use.

This column was added in MySQL 5.6.5.

Prior to MySQL 5.6.10, this value was printed using uppercase. In MySQL 5.6.10 and later, it is always printed using lowercase. (Bug #15869441)

- `Executed_Gtid_Set`

The set of global transaction IDs for all received transactions subsequently executed on this slave. Empty if GTIDs are not in use.

This column was added in MySQL 5.6.5.

Prior to MySQL 5.6.10, this value was printed using uppercase. In MySQL 5.6.10 and later, it is always printed using lowercase. (Bug #15869441)

- `Auto_Position`

The value of this column is 1 if autopositioning is in use; otherwise it is 0.

This column was added in MySQL 5.6.10. (Bug #15992220)

### 13.7.5.36. SHOW STATUS Syntax

```
SHOW [GLOBAL | SESSION] STATUS
     [LIKE 'pattern' | WHERE expr]
```

`SHOW STATUS` provides server status information. This information also can be obtained using the `mysqladmin extended-status` command. The `LIKE [1190]` clause, if present, indicates which variable names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#). This statement does not require any privilege. It requires only the ability to connect to the server.

Partial output is shown here. The list of names and values may be different for your server. The meaning of each variable is given in [Section 5.1.6, “Server Status Variables”](#).

```
mysql> SHOW STATUS;
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Aborted_clients | 0 |
| Aborted_connects | 0 |
| Bytes_received | 155372598 |
| Bytes_sent | 1176560426 |
| Connections | 30023 |
| Created_tmp_disk_tables | 0 |
| Created_tmp_tables | 8340 |
| Created_tmp_files | 60 |
| ... | ... |
| Open_tables | 1 |
| Open_files | 2 |
| Open_streams | 0 |
| Opened_tables | 44600 |
| Questions | 2026873 |
| ... | ... |
| Table_locks_immediate | 1920382 |
| Table_locks_waited | 0 |
| Threads_cached | 0 |
| Threads_created | 30022 |
| Threads_connected | 1 |
| Threads_running | 1 |
| Uptime | 80380 |
```

With a [LIKE \[1190\]](#) clause, the statement displays only rows for those variables with names that match the pattern:

```
mysql> SHOW STATUS LIKE 'Key%';
+-----+-----+
| Variable_name      | Value      |
+-----+-----+
| Key_blocks_used    | 14955      |
| Key_read_requests  | 96854827   |
| Key_reads          | 162040     |
| Key_write_requests | 7589728    |
| Key_writes         | 3813196    |
+-----+-----+
```

With the [GLOBAL](#) modifier, `SHOW STATUS` displays the status values for all connections to MySQL. With [SESSION](#), it displays the status values for the current connection. If no modifier is present, the default is [SESSION](#). [LOCAL](#) is a synonym for [SESSION](#).

Some status variables have only a global value. For these, you get the same value for both [GLOBAL](#) and [SESSION](#). The scope for each status variable is listed at [Section 5.1.6, "Server Status Variables"](#).

Each invocation of the `SHOW STATUS` statement uses an internal temporary table and increments the global `Created_tmp_tables [588]` value.

### 13.7.5.37. SHOW TABLE STATUS Syntax

```
SHOW TABLE STATUS [{FROM | IN} db_name]
                   [LIKE 'pattern' | WHERE expr]
```

`SHOW TABLE STATUS` works like `SHOW TABLES`, but provides a lot of information about each non-[TEMPORARY](#) table. You can also get this list using the `mysqlshow --status db_name` command. The [LIKE \[1190\]](#) clause, if present, indicates which table names to match. The [WHERE](#) clause can be given to select rows using more general conditions, as discussed in [Section 19.32, "Extensions to SHOW Statements"](#).

This statement also displays information about views.

`SHOW TABLE STATUS` returns the following fields:

- [Name](#)  
The name of the table.
- [Engine](#)  
The storage engine for the table. See [Chapter 14, Storage Engines](#).
- [Version](#)  
The version number of the table's `.frm` file.
- [Row\\_format](#)  
The row-storage format ([Fixed](#), [Dynamic](#), [Compressed](#), [Redundant](#), [Compact](#)). For [MyISAM](#) tables, ([Dynamic](#) corresponds to what `myisamchk -dv` reports as [Packed](#). The format of [InnoDB](#) tables is reported as [Redundant](#) or [Compact](#). For the [Barracuda](#) file format of the [InnoDB Plugin](#), the format may be [Compressed](#) or [Dynamic](#).
- [Rows](#)  
The number of rows. Some storage engines, such as [MyISAM](#), store the exact count. For other storage engines, such as [InnoDB](#), this value is an approximation, and may vary from the actual value by as much as 40 to 50%. In such cases, use `SELECT COUNT(*)` to obtain an accurate count.

The `Rows` value is `NULL` for tables in the `INFORMATION_SCHEMA` database.

- `Avg_row_length`

The average row length.

- `Data_length`

The length of the data file.

- `Max_data_length`

The maximum length of the data file. This is the total number of bytes of data that can be stored in the table, given the data pointer size used.

- `Index_length`

The length of the index file.

- `Data_free`

The number of allocated but unused bytes.

This information is also shown for `InnoDB` tables (previously, it was in the `Comment` value). `InnoDB` tables report the free space of the tablespace to which the table belongs. For a table located in the shared tablespace, this is the free space of the shared tablespace. If you are using multiple tablespaces and the table has its own tablespace, the free space is for only that table. Free space means the number of completely free 1MB extents minus a safety margin. Even if free space displays as 0, it may be possible to insert rows as long as new extents need not be allocated.

For partitioned tables, this value is only an estimate and may not be absolutely correct. A more accurate method of obtaining this information in such cases is to query the `INFORMATION_SCHEMA.PARTITIONS` table, as shown in this example:

```
SELECT      SUM(DATA_FREE)
FROM        INFORMATION_SCHEMA.PARTITIONS
WHERE      TABLE_SCHEMA = 'mydb'
AND        TABLE_NAME   = 'mytable';
```

For more information, see [Section 19.14, “The `INFORMATION\_SCHEMA.PARTITIONS` Table”](#).

- `Auto_increment`

The next `AUTO_INCREMENT` value.

- `Create_time`

When the table was created.

- `Update_time`

When the data file was last updated. For some storage engines, this value is `NULL`. For example, `InnoDB` stores multiple tables in its `system tablespace` and the data file timestamp does not apply. Even with `file-per-table` mode with each `InnoDB` table in a separate `.ibd` file, `change buffering` can delay the write to the data file, so the file modification time is different from the time of the last insert, update, or delete. For `MyISAM`, the data file timestamp is used; however, on Windows the timestamp is not updated by updates so the value is inaccurate.

- `Check_time`

When the table was last checked. Not all storage engines update this time, in which case the value is always `NULL`.

- [Collation](#)

The table's character set and collation.

- [Checksum](#)

The live checksum value (if any).

- [Create\\_options](#)

Extra options used with `CREATE TABLE`. The original options supplied when `CREATE TABLE` is called are retained and the options reported here may differ from the active table settings and options.

- [Comment](#)

The comment used when creating the table (or information as to why MySQL could not access the table information).

For `MEMORY` tables, the `Data_length`, `Max_data_length`, and `Index_length` values approximate the actual amount of allocated memory. The allocation algorithm reserves memory in large amounts to reduce the number of allocation operations.

For views, all the fields displayed by `SHOW TABLE STATUS` are `NULL` except that `Name` indicates the view name and `Comment` says `view`.

### 13.7.5.38. SHOW TABLES Syntax

```
SHOW [FULL] TABLES [{FROM | IN} db_name]
                [LIKE 'pattern' | WHERE expr]
```

`SHOW TABLES` lists the non-`TEMPORARY` tables in a given database. You can also get this list using the `mysqlshow db_name` command. The `LIKE` [1190] clause, if present, indicates which table names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

Matching performed by the `LIKE` clause is dependent on the setting of the `lower_case_table_names` [501] system variable.

This statement also lists any views in the database. The `FULL` modifier is supported such that `SHOW FULL TABLES` displays a second output column. Values for the second column are `BASE TABLE` for a table and `VIEW` for a view.

If you have no privileges for a base table or view, it does not show up in the output from `SHOW TABLES` or `mysqlshow db_name`.

### 13.7.5.39. SHOW TRIGGERS Syntax

```
SHOW TRIGGERS [{FROM | IN} db_name]
                [LIKE 'pattern' | WHERE expr]
```

`SHOW TRIGGERS` lists the triggers currently defined for tables in a database (the default database unless a `FROM` clause is given). This statement returns results only for databases and tables for which you have the `TRIGGER` [756] privilege. The `LIKE` [1190] clause, if present, indicates which table names to match and causes the statement to display triggers for those tables. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#).

For the trigger `ins_sum` as defined in [Section 18.3, “Using Triggers”](#), the output of this statement is as shown here:

```
mysql> SHOW TRIGGERS LIKE 'acc%'\G
***** 1. row *****
      Trigger: ins_sum
```

```

      Event: INSERT
      Table: account
      Statement: SET @sum = @sum + NEW.amount
      Timing: BEFORE
      Created: NULL
      sql_mode:
      Definer: myname@localhost
character_set_client: latin1
collation_connection: latin1_swedish_ci
Database Collation: latin1_swedish_ci

```

`character_set_client` [467] is the session value of the `character_set_client` [467] system variable when the trigger was created. `collation_connection` [469] is the session value of the `collation_connection` [469] system variable when the trigger was created. `Database Collation` is the collation of the database with which the trigger is associated.



#### Note

When using a `LIKE` [1190] clause with `SHOW TRIGGERS`, the expression to be matched (*expr*) is compared with the name of the table on which the trigger is declared, and not with the name of the trigger:

```

mysql> SHOW TRIGGERS LIKE 'ins%';
Empty set (0.01 sec)

```

A brief explanation of the columns in the output of this statement is shown here:

- `Trigger`

The name of the trigger.

- `Event`

The event that causes trigger activation: one of `'INSERT'`, `'UPDATE'`, or `'DELETE'`.

- `Table`

The table for which the trigger is defined.

- `Statement`

The statement to be executed when the trigger is activated. This is the same as the text shown in the `ACTION_STATEMENT` column of `INFORMATION_SCHEMA.TRIGGERS`.

- `Timing`

One of the two values `'BEFORE'` or `'AFTER'`.

- `Created`

Currently, the value of this column is always `NULL`.

- `sql_mode` [545]

The SQL mode in effect when the trigger executes.

- `Definer`

The account that created the trigger.

See also [Section 19.27, “The INFORMATION\\_SCHEMA TRIGGERS Table”](#).

### 13.7.5.40. SHOW VARIABLES Syntax

```

SHOW [GLOBAL | SESSION] VARIABLES
     [LIKE 'pattern' | WHERE expr]

```

`SHOW VARIABLES` shows the values of MySQL system variables. This information also can be obtained using the `mysqladmin variables` command. The `LIKE [1190]` clause, if present, indicates which variable names to match. The `WHERE` clause can be given to select rows using more general conditions, as discussed in [Section 19.32, “Extensions to SHOW Statements”](#). This statement does not require any privilege. It requires only the ability to connect to the server.

With the `GLOBAL` modifier, `SHOW VARIABLES` displays the values that are used for new connections to MySQL. In MySQL 5.6, if a variable has no global value, no value is displayed. With `SESSION`, `SHOW VARIABLES` displays the values that are in effect for the current connection. If no modifier is present, the default is `SESSION`. `LOCAL` is a synonym for `SESSION`.

`SHOW VARIABLES` is subject to a version-dependent display-width limit. For variables with very long values that are not completely displayed, use `SELECT` as a workaround. For example:

```
SELECT @@GLOBAL.innodb_data_file_path;
```

If the default system variable values are unsuitable, you can set them using command options when `mysqld` starts, and most can be changed at runtime with the `SET` statement. See [Section 5.1.5, “Using System Variables”](#), and [Section 13.7.4, “SET Syntax”](#).

Partial output is shown here. The list of names and values may be different for your server. [Section 5.1.4, “Server System Variables”](#), describes the meaning of each variable, and [Section 8.11.2, “Tuning Server Parameters”](#), provides information about tuning them.

```
mysql> SHOW VARIABLES;
```

Variable_name	Value
auto_increment_increment	1
auto_increment_offset	1
autocommit	ON
automatic_sp_privileges	ON
back_log	50
basedir	/home/jon/bin/mysql-5.5
big_tables	OFF
binlog_cache_size	32768
binlog_direct_non_transactional_updates	OFF
binlog_format	STATEMENT
binlog_stmt_cache_size	32768
bulk_insert_buffer_size	8388608
...	
max_allowed_packet	1048576
max_binlog_cache_size	18446744073709547520
max_binlog_size	1073741824
max_binlog_stmt_cache_size	18446744073709547520
max_connect_errors	10
max_connections	151
max_delayed_threads	20
max_error_count	64
max_heap_table_size	16777216
max_insert_delayed_threads	20
max_join_size	18446744073709551615
...	
thread_handling	one-thread-per-connection
thread_stack	262144
time_format	%H:%i:%s
time_zone	SYSTEM
timed_mutexes	OFF
timestamp	1316689732
tmp_table_size	16777216
tmpdir	/tmp
transaction_alloc_block_size	8192
transaction_prealloc_size	4096
tx_isolation	REPEATABLE-READ
unique_checks	ON
updatable_views_with_limit	YES
version	5.5.17-log
version_comment	Source distribution



version_compile_machine	x86_64
version_compile_os	Linux
wait_timeout	28800
warning_count	0

With a [LIKE \[1190\]](#) clause, the statement displays only rows for those variables with names that match the pattern. To obtain the row for a specific variable, use a [LIKE \[1190\]](#) clause as shown:

```
SHOW VARIABLES LIKE 'max_join_size';
SHOW SESSION VARIABLES LIKE 'max_join_size';
```

To get a list of variables whose name match a pattern, use the “%” wildcard character in a [LIKE \[1190\]](#) clause:

```
SHOW VARIABLES LIKE '%size%';
SHOW GLOBAL VARIABLES LIKE '%size%';
```

Wildcard characters can be used in any position within the pattern to be matched. Strictly speaking, because “\_” is a wildcard that matches any single character, you should escape it as “\\_” to match it literally. In practice, this is rarely necessary.

### 13.7.5.41. SHOW WARNINGS Syntax

```
SHOW WARNINGS [LIMIT [offset,] row_count]
SHOW COUNT(*) WARNINGS
```

[SHOW WARNINGS](#) shows information about the conditions (errors, warnings, and notes) that resulted from the last statement in the current session that generated messages. It shows nothing if the last statement used a table and generated no messages. (That is, a statement that uses a table but generates no messages clears the message list.) Statements that do not use tables and do not generate messages have no effect on the message list.

Warnings are generated for DML statements such as [INSERT](#), [UPDATE](#), and [LOAD DATA INFILE](#) as well as DDL statements such as [CREATE TABLE](#) and [ALTER TABLE](#).

[SHOW WARNINGS](#) is also used following [EXPLAIN EXTENDED](#), to display the extra information generated by [EXPLAIN](#) when the [EXTENDED](#) keyword is used. See [Section 8.8.3, “EXPLAIN EXTENDED Output Format”](#).

The [LIMIT](#) clause has the same syntax as for the [SELECT](#) statement. See [Section 13.2.9, “SELECT Syntax”](#).

A related statement, [SHOW ERRORS](#), shows only the error conditions (it excludes warnings and notes). See [Section 13.7.5.18, “SHOW ERRORS Syntax”](#). [GET DIAGNOSTICS](#) can be used to examine information for individual conditions. See [Section 13.6.7.3, “GET DIAGNOSTICS Syntax”](#).

The [SHOW COUNT\(\\*\) WARNINGS](#) statement displays the total number of errors, warnings, and notes. You can also retrieve this number from the [warning\\_count \[564\]](#) system variable:

```
SHOW COUNT(*) WARNINGS;
SELECT @@warning_count;
```

Here is a simple example that shows a syntax warning for [CREATE TABLE](#) and conversion warnings for [INSERT](#):

```
mysql> CREATE TABLE t1
> (a TINYINT NOT NULL, b CHAR(4))
> TYPE=MyISAM;
Query OK, 0 rows affected, 1 warning (0.00 sec)
mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1287
Message: 'TYPE=storage_engine' is deprecated, use
'ENGINE=storage_engine' instead
1 row in set (0.00 sec)
```

```
mysql> INSERT INTO t1 VALUES(10,'mysql'),
-> (NULL,'test'), (300,'Open Source');
Query OK, 3 rows affected, 4 warnings (0.01 sec)
Records: 3 Duplicates: 0 Warnings: 4

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1265
Message: Data truncated for column 'b' at row 1
***** 2. row *****
Level: Warning
Code: 1263
Message: Data truncated, NULL supplied to NOT NULL column 'a' at row 2
***** 3. row *****
Level: Warning
Code: 1264
Message: Data truncated, out of range for column 'a' at row 3
***** 4. row *****
Level: Warning
Code: 1265
Message: Data truncated for column 'b' at row 3
4 rows in set (0.00 sec)
```

The `max_error_count` [504] system variable controls the maximum number of error, warning, and note messages for which the server stores information, and thus the number of messages that `SHOW WARNINGS` displays. By default, `max_error_count` [504] is 64. To change the number of messages the server can store, change the value of `max_error_count` [504].

The value of `warning_count` [564] is not limited by `max_error_count` [504] if the number of messages generated exceeds `max_error_count` [504].

In the following example, the `ALTER TABLE` statement produces three warning messages (as shown by the value of `warning_count` [564]), but only one is stored because `max_error_count` [504] has been set to 1:

```
mysql> SHOW VARIABLES LIKE 'max_error_count';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| max_error_count | 64 |
+-----+-----+
1 row in set (0.00 sec)

mysql> SET max_error_count=1;
Query OK, 0 rows affected (0.00 sec)

mysql> ALTER TABLE t1 MODIFY b CHAR;
Query OK, 3 rows affected, 3 warnings (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 3

mysql> SELECT @@warning_count;
+-----+
| @@warning_count |
+-----+
| 3 |
+-----+
1 row in set (0.01 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Warning | 1263 | Data truncated for column 'b' at row 1 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

To disable warnings, set `max_error_count` [504] to 0. In this case, `warning_count` [564] still indicates how many warnings have occurred, but none of the messages are stored.

The following `DROP TABLE` statement results in a note:

```
mysql> DROP TABLE IF EXISTS test.no_such_table;
Query OK, 0 rows affected, 1 warning (0.01 sec)
mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                                |
+-----+-----+-----+
| Note  | 1051 | Unknown table 'test.no_such_table'    |
+-----+-----+-----+
```

If the `sql_notes` [546] system variable is set to 0, notes do not increment `warning_count` [564] and the server does not record them.

The MySQL server sends back a count indicating the total number of errors, warnings, and notes resulting from the last statement. From the C API, this value can be obtained by calling `mysql_warning_count()`. See Section 21.8.6.73, “`mysql_warning_count()`”.

## 13.7.6. Other Administrative Statements

### 13.7.6.1. BINLOG Syntax

```
BINLOG 'str'
```

`BINLOG` is an internal-use statement. It is generated by the `mysqlbinlog` program as the printable representation of certain events in binary log files. (See Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”.) The `'str'` value is a base 64-encoded string that the server decodes to determine the data change indicated by the corresponding event. This statement requires the `SUPER` [756] privilege.

As of MySQL 5.6, this statement can execute only format description events and row events. Previously it could execute all types of events.

### 13.7.6.2. CACHE INDEX Syntax

```
CACHE INDEX
  tbl_index_list [, tbl_index_list] ...
  [PARTITION (partition_list | ALL)]
  IN key_cache_name

tbl_index_list:
  tbl_name [[INDEX|KEY] (index_name[, index_name] ...)]

partition_list:
  partition_name[, partition_name][, ...]
```

The `CACHE INDEX` statement assigns table indexes to a specific key cache. It is used only for `MyISAM` tables. After the indexes have been assigned, they can be preloaded into the cache if desired with `LOAD INDEX INTO CACHE`.

The following statement assigns indexes from the tables `t1`, `t2`, and `t3` to the key cache named `hot_cache`:

```
mysql> CACHE INDEX t1, t2, t3 IN hot_cache;
+-----+-----+-----+
| Table  | Op                | Msg_type | Msg_text |
+-----+-----+-----+
| test.t1 | assign_to_keycache | status   | OK       |
| test.t2 | assign_to_keycache | status   | OK       |
| test.t3 | assign_to_keycache | status   | OK       |
+-----+-----+-----+
```

The syntax of `CACHE INDEX` enables you to specify that only particular indexes from a table should be assigned to the cache. The current implementation assigns all the table's indexes to the cache, so there is no reason to specify anything other than the table name.

The key cache referred to in a `CACHE INDEX` statement can be created by setting its size with a parameter setting statement or in the server parameter settings. For example:

```
mysql> SET GLOBAL keycache1.key_buffer_size=128*1024;
```

Key cache parameters can be accessed as members of a structured system variable. See [Section 5.1.5.1, “Structured System Variables”](#).

A key cache must exist before you can assign indexes to it:

```
mysql> CACHE INDEX t1 IN non_existent_cache;
ERROR 1284 (HY000): Unknown key cache 'non_existent_cache'
```

By default, table indexes are assigned to the main (default) key cache created at the server startup. When a key cache is destroyed, all indexes assigned to it become assigned to the default key cache again.

Index assignment affects the server globally: If one client assigns an index to a given cache, this cache is used for all queries involving the index, no matter which client issues the queries.

In MySQL 5.6, this statement is also supported for partitioned `MyISAM` tables. You can assign one or more indexes for one, several, or all partitions to a given key cache. For example, you can do the following:

```
CREATE TABLE pt (c1 INT, c2 VARCHAR(50), INDEX i(c1))
  PARTITION BY HASH(c1)
  PARTITIONS 4;

SET GLOBAL kc_fast.key_buffer_size = 128 * 1024;
SET GLOBAL kc_slow.key_buffer_size = 128 * 1024;

CACHE INDEX pt PARTITION (p0) IN kc_fast;
CACHE INDEX pt PARTITION (p1, p3) IN kc_slow;
```

The previous set of statements performs the following actions:

- Creates a partitioned table with 4 partitions; these partitions are automatically named `p0`, ..., `p3`; this table has an index named `i` on column `c1`.
- Creates 2 key caches named `kc_fast` and `kc_slow`
- Assigns the index for partition `p0` to the `kc_fast` key cache and the index for partitions `p1` and `p3` to the `kc_slow` key cache; the index for the remaining partition (`p2`) uses the server's default key cache.

If you wish instead to assign the indexes for all partitions in table `pt` to a single key cache named `kc_all`, you can use either one of the following 2 statements:

```
CACHE INDEX pt PARTITION (ALL) IN kc_all;

CACHE INDEX pt IN kc_all;
```

The two statements just shown are equivalent, and issuing either one of them has exactly the same effect. In other words, if you wish to assign indexes for all partitions of a partitioned table to the same key cache, then the `PARTITION (ALL)` clause is optional.

When assigning indexes for multiple partitions to a key cache, the partitions do not have to be contiguous, and you are not required to list their names in any particular order. Indexes for any partitions that are not explicitly assigned to a key cache automatically use the server's default key cache.

In MySQL 5.6, index preloading is also supported for partitioned `MyISAM` tables. For more information, see [Section 13.7.6.5, “LOAD INDEX INTO CACHE Syntax”](#).

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

### 13.7.6.3. FLUSH Syntax

```
FLUSH [NO_WRITE_TO_BINLOG | LOCAL]
      flush_option [, flush_option] ...
```

The `FLUSH` statement has several variant forms that clear or reload various internal caches, flush tables, or acquire locks. To execute `FLUSH`, you must have the `RELOAD [755]` privilege. Specific flush options might require additional privileges, as described later.

By default, the server writes `FLUSH` statements to the binary log so that they replicate to replication slaves. To suppress logging, use the optional `NO_WRITE_TO_BINLOG` keyword or its alias `LOCAL`.



#### Note

`FLUSH LOGS`, `FLUSH TABLES WITH READ LOCK` (with or without a table list), and `FLUSH TABLES tbl_name ... FOR EXPORT` are not written to the binary log in any case because they would cause problems if replicated to a slave.

The `FLUSH` statement causes an implicit commit. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

The `RESET` statement is similar to `FLUSH`. See [Section 13.7.6.6, “RESET Syntax”](#), for information about using the `RESET` statement with replication.

*flush\_option* can be any of the following items.

- `DES_KEY_FILE`

Reloads the DES keys from the file that was specified with the `--des-key-file [421]` option at server startup time.

- `HOSTS`

Empties the host cache. You should flush the host cache if some of your hosts change IP address or if the error message `Host 'host_name' is blocked` occurs. (See [Section C.5.2.6, “Host 'host\\_name' is blocked”](#).) When more than `max_connect_errors [502]` errors occur successively for a given host while connecting to the MySQL server, MySQL assumes that something is wrong and blocks the host from further connection requests. Flushing the host cache enables further connection attempts from the host. The default value of `max_connect_errors [502]` is 10. To avoid this error message, start the server with `max_connect_errors [502]` set to a large value.

- [*log\_type*] `LOGS`

With no *log\_type* option, `FLUSH LOGS` closes and reopens all log files. If binary logging is enabled, the sequence number of the binary log file is incremented by one relative to the previous file. On Unix, this is the same thing as sending a `SIGHUP` signal to the `mysqld` server (except on some Mac OS X 10.3 versions where `mysqld` ignores `SIGHUP` and `SIGQUIT`).

With a *log\_type* option, only the specified log type is flushed. These *log\_type* options are permitted:

- `BINARY` closes and reopens the binary log files.
- `ENGINE` closes and reopens any flushable logs for installed storage engines. Currently, this causes `InnoDB` to flush its logs to disk.
- `ERROR` closes and reopens the error log file.
- `GENERAL` closes and reopens the general query log file.

- [RELAY](#) closes and reopens the relay log files.
- [SLOW](#) closes and reopens the slow query log file.
- [PRIVILEGES](#)

Reloads the privileges from the grant tables in the `mysql` database. On Unix, this also occurs if the server receives a `SIGHUP` signal.

The server caches information in memory as a result of `GRANT`, `CREATE USER`, `CREATE SERVER`, and `INSTALL PLUGIN` statements. This memory is not released by the corresponding `REVOKE`, `DROP USER`, `DROP SERVER`, and `UNINSTALL PLUGIN` statements, so for a server that executes many instances of the statements that cause caching, there will be an increase in memory use. This cached memory can be freed with `FLUSH PRIVILEGES`.

- [QUERY CACHE](#)

Defragment the query cache to better utilize its memory. `FLUSH QUERY CACHE` does not remove any queries from the cache, unlike `FLUSH TABLES` or `RESET QUERY CACHE`.

- [STATUS](#)

This option adds the current thread's session status variable values to the global values and resets the session values to zero. Some global variables may be reset to zero as well. It also resets the counters for key caches (default and named) to zero and sets `Max_used_connections` [594] to the current number of open connections. This is something you should use only when debugging a query. See [Section 1.7, "How to Report Bugs or Problems"](#).

- [TABLES](#)

`FLUSH TABLES` flushes tables, and, depending on the variant used, acquires locks. The permitted syntax is discussed later in this section.

- [USER\\_RESOURCES](#)

Resets all per-hour user resources to zero. This enables clients that have reached their hourly connection, query, or update limits to resume activity immediately. `FLUSH USER_RESOURCES` does not apply to the limit on maximum simultaneous connections. See [Section 6.3.4, "Setting Account Resource Limits"](#).

The `mysqladmin` utility provides a command-line interface to some flush operations, using commands such as `flush-hosts`, `flush-logs`, `flush-privileges`, `flush-status`, and `flush-tables`. See [Section 4.5.2, "mysqladmin — Client for Administering a MySQL Server"](#).

**Note**

It is not possible to issue `FLUSH` statements within stored functions or triggers. However, you may use `FLUSH` in stored procedures, so long as these are not called from stored functions or triggers. See [Section E.1, "Restrictions on Stored Programs"](#).

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing `FLUSH` statements (BUG#16062608).

## **FLUSH TABLES Syntax**

`FLUSH TABLES` has several forms, described following. If any variant of the `TABLES` option is used in a `FLUSH` statement, it must be the only option used. `FLUSH TABLE` is a synonym for `FLUSH TABLES`.

- [FLUSH TABLES](#)

Closes all open tables, forces all tables in use to be closed, and flushes the query cache. `FLUSH TABLES` also removes all query results from the query cache, like the `RESET QUERY CACHE` statement.

In MySQL 5.6, `FLUSH TABLES` is not permitted when there is an active `LOCK TABLES ... READ`. To flush and lock tables, use `FLUSH TABLES tbl_name ... WITH READ LOCK` instead.

- `FLUSH TABLES tbl_name [, tbl_name] ...`

With a list of one or more comma-separated table names, this statement is like `FLUSH TABLES` with no names except that the server flushes only the named tables. No error occurs if a named table does not exist.

- `FLUSH TABLES WITH READ LOCK`

Closes all open tables and locks all tables for all databases with a global read lock. This is a very convenient way to get backups if you have a file system such as Veritas or ZFS that can take snapshots in time. Use `UNLOCK TABLES` to release the lock.

`FLUSH TABLES WITH READ LOCK` acquires a global read lock and not table locks, so it is not subject to the same behavior as `LOCK TABLES` and `UNLOCK TABLES` with respect to table locking and implicit commits:

- `UNLOCK TABLES` implicitly commits any active transaction only if any tables currently have been locked with `LOCK TABLES`. The commit does not occur for `UNLOCK TABLES` following `FLUSH TABLES WITH READ LOCK` because the latter statement does not acquire table locks.
- Beginning a transaction causes table locks acquired with `LOCK TABLES` to be released, as though you had executed `UNLOCK TABLES`. Beginning a transaction does not release a global read lock acquired with `FLUSH TABLES WITH READ LOCK`.

`FLUSH TABLES WITH READ LOCK` does not prevent the server from inserting rows into the log tables (see [Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”](#)).

- `FLUSH TABLES tbl_name [, tbl_name] ... WITH READ LOCK`

This statement flushes and acquires read locks for the named tables. The statement first acquires exclusive metadata locks for the tables, so it waits for transactions that have those tables open to complete. Then the statement flushes the tables from the table cache, reopens the tables, acquires table locks (like `LOCK TABLES ... READ`), and downgrades the metadata locks from exclusive to shared. After the statement acquires locks and downgrades the metadata locks, other sessions can read but not modify the tables.

Because this statement acquires table locks, you must have the `LOCK TABLES [755]` privilege for each table, in addition to the `RELOAD [755]` privilege that is required to use any `FLUSH` statement.

This statement applies only to existing base tables. If a name refers to a base table, that table is used. If it refers to a `TEMPORARY` table, it is ignored. If a name applies to a view, an `ER_WRONG_OBJECT [3682]` error occurs. Otherwise, an `ER_NO_SUCH_TABLE [3670]` error occurs.

Use `UNLOCK TABLES` to release the locks, `LOCK TABLES` to release the locks and acquire other locks, or `START TRANSACTION` to release the locks and begin a new transaction.

This variant of `FLUSH` enables tables to be flushed and locked in a single operation. It provides a workaround for the restriction in MySQL 5.6 that `FLUSH TABLES` is not permitted when there is an active `LOCK TABLES ... READ`.

This statement does not perform an implicit `UNLOCK TABLES`, so an error results if you use the statement while there is any active `LOCK TABLES` or use it a second time without first releasing the locks acquired.



If a flushed table was opened with `HANDLER`, the handler is implicitly flushed and loses its position.

- `FLUSH TABLES tbl_name [, tbl_name] ... FOR EXPORT`

This `FLUSH TABLES` variant applies to `InnoDB` tables. It is available as of MySQL 5.6.6. The statement ensures that changes to the named tables have been flushed to disk so that binary table copies can be made while the server is running.

The statement works like this:

1. It acquires shared metadata locks for the named tables. The statement blocks as long as other sessions have active transactions that have modified those tables or hold table locks for them. When the locks have been acquired, the statement blocks transactions that attempt to update the tables while permitting read-only operations to continue.
2. It checks whether all storage engines for the tables support `FOR EXPORT`. If any do not, an `ER_ILLEGAL HA [3662]` error occurs and the statement fails.
3. The statement notifies the storage engine for each table to make the table ready for export. The storage engine must ensure that any pending changes are written to disk.
4. The statement puts the session in lock-tables mode so that the metadata locks acquired earlier are not released when the `FOR EXPORT` statement completes.

Because this statement acquires table locks, you must have the `LOCK TABLES [755]` privilege for each table, in addition to the `RELOAD [755]` privilege that is required to use any `FLUSH` statement.

This statement applies only to existing base tables. If a name refers to a base table, that table is used. If it refers to a `TEMPORARY` table, it is ignored. If a name applies to a view, an `ER_WRONG_OBJECT [3682]` error occurs. Otherwise, an `ER_NO_SUCH_TABLE [3670]` error occurs.

`InnoDB` supports `FOR EXPORT` for tables that have their own `.ibd` file (that is, tables that were created with the `innodb_file_per_table [1718]` setting enabled). `InnoDB` ensures when notified by the `FOR EXPORT` statement that any changes have been flushed to disk. This permits a binary copy of table contents to be made while the `FOR EXPORT` statement is in effect because the `.ibd` file is transaction consistent and can be copied while the server is running. `FOR EXPORT` does not apply to `InnoDB` system tablespace files, or to `InnoDB` tables that have any `FULLTEXT` indexes.

When notified by `FOR EXPORT`, `InnoDB` writes to disk certain kinds of data that is normally held in memory or in separate disk buffers outside the tablespace files. For each table, `InnoDB` also produces a file named `table_name.cfg` in the same database directory as the table. The `.cfg` file contains metadata needed to reimport the tablespace files later, into the same or different server.

When the `FOR EXPORT` statement completes, `InnoDB` will have flushed all `dirty pages` to the table data files. Any `change buffer` entries are merged prior to flushing. At this point, the tables are locked and quiescent: The tables are in a transactionally consistent state on disk and you can copy the `.ibd` tablespace files along with the corresponding `.cfg` files to get a consistent snapshot of those tables.

For the procedure to reimport the copied table data into a MySQL instance, see [Section 14.2.4.2.34, “Improved Tablespace Management”](#).

After you are done with the tables, use `UNLOCK TABLES` to release the locks, `LOCK TABLES` to release the locks and acquire other locks, or `START TRANSACTION` to release the locks and begin a new transaction.

While any of these statements is in effect within the session, attempts to use `FLUSH TABLES ... FOR EXPORT` produce an error:

```
FLUSH TABLES ... WITH READ LOCK
```

```
FLUSH TABLES ... FOR EXPORT
LOCK TABLES ... READ
LOCK TABLES ... WRITE
```

While `FLUSH TABLES ... FOR EXPORT` is in effect within the session, attempts to use any of these statements produce an error:

```
FLUSH TABLES WITH READ LOCK
FLUSH TABLES ... WITH READ LOCK
FLUSH TABLES ... FOR EXPORT
```

### 13.7.6.4. `KILL` Syntax

```
KILL [CONNECTION | QUERY] thread_id
```

Each connection to `mysqld` runs in a separate thread. You can see which threads are running with the `SHOW PROCESSLIST` statement and kill a thread with the `KILL thread_id` statement.

`KILL` permits an optional `CONNECTION` or `QUERY` modifier:

- `KILL CONNECTION` is the same as `KILL` with no modifier: It terminates the connection associated with the given *thread\_id*.
- `KILL QUERY` terminates the statement that the connection is currently executing, but leaves the connection itself intact.

If you have the `PROCESS [755]` privilege, you can see all threads. If you have the `SUPER [756]` privilege, you can kill all threads and statements. Otherwise, you can see and kill only your own threads and statements.

You can also use the `mysqladmin processlist` and `mysqladmin kill` commands to examine and kill threads.



#### Note

You cannot use `KILL` with the Embedded MySQL Server library because the embedded server merely runs inside the threads of the host application. It does not create any connection threads of its own.

When you use `KILL`, a thread-specific kill flag is set for the thread. In most cases, it might take some time for the thread to die because the kill flag is checked only at specific intervals:

- In `SELECT`, `ORDER BY` and `GROUP BY` loops, the flag is checked after reading a block of rows. If the kill flag is set, the statement is aborted.
- During `ALTER TABLE`, the kill flag is checked before each block of rows are read from the original table. If the kill flag was set, the statement is aborted and the temporary table is deleted.
- During `UPDATE` or `DELETE` operations, the kill flag is checked after each block read and after each updated or deleted row. If the kill flag is set, the statement is aborted. Note that if you are not using transactions, the changes are not rolled back.
- `GET_LOCK( ) [1275]` aborts and returns `NULL`.
- An `INSERT DELAYED` thread quickly flushes (inserts) all rows it has in memory and then terminates.
- If the thread is in the table lock handler (state: `Locked`), the table lock is quickly aborted.
- If the thread is waiting for free disk space in a write call, the write is aborted with a “disk full” error message.

**Warning**

Killing a `REPAIR TABLE` or `OPTIMIZE TABLE` operation on a `MyISAM` table results in a table that is corrupted and unusable. Any reads or writes to such a table fail until you optimize or repair it again (without interruption).

**13.7.6.5. LOAD INDEX INTO CACHE Syntax**

```
LOAD INDEX INTO CACHE
  tbl_index_list [, tbl_index_list] ...

tbl_index_list:
  tbl_name
  [PARTITION (partition_list | ALL)]
  [[INDEX|KEY] (index_name[, index_name] ...)]
  [IGNORE LEAVES]

partition_list:
  partition_name[, partition_name][, ...]
```

The `LOAD INDEX INTO CACHE` statement preloads a table index into the key cache to which it has been assigned by an explicit `CACHE INDEX` statement, or into the default key cache otherwise.

`LOAD INDEX INTO CACHE` is used only for `MyISAM` tables. In MySQL 5.6, it is also supported for partitioned `MyISAM` tables; in addition, indexes on partitioned tables can be preloaded for one, several, or all partitions.

The `IGNORE LEAVES` modifier causes only blocks for the nonleaf nodes of the index to be preloaded.

`IGNORE LEAVES` is also supported for partitioned `MyISAM` tables.

The following statement preloads nodes (index blocks) of indexes for the tables `t1` and `t2`:

```
mysql> LOAD INDEX INTO CACHE t1, t2 IGNORE LEAVES;
+-----+-----+-----+-----+
| Table | Op          | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.t1 | preload_keys | status   | OK       |
| test.t2 | preload_keys | status   | OK       |
+-----+-----+-----+-----+
```

This statement preloads all index blocks from `t1`. It preloads only blocks for the nonleaf nodes from `t2`.

The syntax of `LOAD INDEX INTO CACHE` enables you to specify that only particular indexes from a table should be preloaded. The current implementation preloads all the table's indexes into the cache, so there is no reason to specify anything other than the table name.

Beginning with MySQL 5.6.11, `gtid_next` [2022] must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

In MySQL 5.6, it is possible to preload indexes on specific partitions of partitioned `MyISAM` tables. For example, of the following 2 statements, the first preloads indexes for partition `p0` of a partitioned table `pt`, while the second preloads the indexes for partitions `p1` and `p3` of the same table:

```
LOAD INDEX INTO CACHE pt PARTITION (p0);
LOAD INDEX INTO CACHE pt PARTITION (p1, p3);
```

To preload the indexes for all partitions in table `pt`, you can use either one of the following 2 statements:

```
LOAD INDEX INTO CACHE pt PARTITION (ALL);
LOAD INDEX INTO CACHE pt;
```

The two statements just shown are equivalent, and issuing either one of them has exactly the same effect. In other words, if you wish to preload indexes for all partitions of a partitioned table, then the `PARTITION (ALL)` clause is optional.

When preloading indexes for multiple partitions, the partitions do not have to be contiguous, and you are not required to list their names in any particular order.

`LOAD INDEX INTO CACHE ... IGNORE LEAVES` fails unless all indexes in a table have the same block size. You can determine index block sizes for a table by using `myisamchk -dv` and checking the `Blocksize` column.

### 13.7.6.6. RESET Syntax

```
RESET reset_option [, reset_option] ...
```

The `RESET` statement is used to clear the state of various server operations. You must have the `RELOAD [755]` privilege to execute `RESET`.

`RESET` acts as a stronger version of the `FLUSH` statement. See [Section 13.7.6.3, “FLUSH Syntax”](#).

The `RESET` statement causes an implicit commit. See [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

Beginning with MySQL 5.6.11, `gtid_next [2022]` must be set to `AUTOMATIC` before issuing this statement (BUG#16062608).

*reset\_option* can be any of the following:

- `MASTER`

Deletes all binary logs listed in the index file, resets the binary log index file to be empty, and creates a new binary log file.

- `QUERY CACHE`

Removes all query results from the query cache.

- `SLAVE`

Makes the slave forget its replication position in the master binary logs. Also resets the relay log by deleting any existing relay log files and beginning a new one.

## 13.8. MySQL Utility Statements

### 13.8.1. DESCRIBE Syntax

The `DESCRIBE` and `EXPLAIN` statements are synonyms, used either to obtain information about table structure or query execution plans. For more information, see [Section 13.8.2, “EXPLAIN Syntax”](#).

### 13.8.2. EXPLAIN Syntax

```
{EXPLAIN | DESCRIBE | DESC}
  tbl_name [col_name | wild]

{EXPLAIN | DESCRIBE | DESC}
  [explain_type] explainable_stmt

explain_type:
  EXTENDED
  | PARTITIONS
  | FORMAT = format_name

format_name:
  TRADITIONAL
  | JSON

explainable_stmt:
  SELECT statement
  | DELETE statement
```

```

| INSERT statement
| REPLACE statement
| UPDATE statement

```

The `DESCRIBE` and `EXPLAIN` statements are synonyms. In practice, the `DESCRIBE` keyword is more often used to obtain information about table structure, whereas `EXPLAIN` is used to obtain a query execution plan (that is, an explanation of how MySQL would execute a query). The following discussion uses the `DESCRIBE` and `EXPLAIN` keywords in accordance with those uses, but the MySQL parser treats them as completely synonymous.

## Obtaining Table Structure Information

`DESCRIBE` provides information about the columns in a table. It is a shortcut for `SHOW COLUMNS FROM`. These statements also display information for views. (See [Section 13.7.5.6, “SHOW COLUMNS Syntax”](#).)

`col_name` can be a column name, or a string containing the SQL “%” and “\_” wildcard characters to obtain output only for the columns with names matching the string. There is no need to enclose the string within quotation marks unless it contains spaces or other special characters.

```

mysql> DESCRIBE City;
+-----+-----+-----+-----+-----+-----+
| Field | Type   | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| Id    | int(11)| NO   | PRI | NULL    | auto_increment |
| Name  | char(35)| NO   |     |         |              |
| Country | char(3) | NO   | UNI |         |              |
| District | char(20) | YES | MUL |         |              |
| Population | int(11) | NO   |     | 0       |              |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)

```

The description for `SHOW COLUMNS` provides more information about the output columns (see [Section 13.7.5.6, “SHOW COLUMNS Syntax”](#)).

If the data types differ from what you expect them to be based on a `CREATE TABLE` statement, note that MySQL sometimes changes data types when you create or alter a table. The conditions under which this occurs are described in [Section 13.1.14.3, “Silent Column Specification Changes”](#).

The `DESCRIBE` statement is provided for compatibility with Oracle.

The `SHOW CREATE TABLE`, `SHOW TABLE STATUS`, and `SHOW INDEX` statements also provide information about tables. See [Section 13.7.5, “SHOW Syntax”](#).

## Obtaining Query Execution Plan Information

The `EXPLAIN` statement provides a way to obtain information about how MySQL executes a statement:

- When you precede a statement with the keyword `EXPLAIN`, MySQL displays information from the optimizer about the statement execution plan. That is, MySQL explains how it would process the statement, including information about how tables are joined and in which order. `EXPLAIN EXTENDED` can be used to obtain additional information.

As of MySQL 5.6.3, `EXPLAIN` provides information about `SELECT`, `DELETE`, `INSERT`, `REPLACE`, and `UPDATE` statements. Before MySQL 5.6.3, `EXPLAIN` provides information only about `SELECT` statements.

For information about using `EXPLAIN` and `EXPLAIN EXTENDED` to obtain query execution plan information, see [Section 8.8.1, “Optimizing Queries with EXPLAIN”](#).

- `EXPLAIN PARTITIONS` is useful only when examining queries involving partitioned tables. For details, see [Section 17.3.5, “Obtaining Information About Partitions”](#).
- As of MySQL 5.6.5, the `FORMAT` option can be used to select the output format. `TRADITIONAL` presents the output in tabular format. This is the default if no `FORMAT` option is present. `JSON` format

displays the information in JSON format. With `FORMAT = JSON`, the output includes extended and partition information.

### 13.8.3. HELP Syntax

```
HELP 'search_string'
```

The `HELP` statement returns online information from the MySQL Reference manual. Its proper operation requires that the help tables in the `mysql` database be initialized with help topic information (see [Section 5.1.10, “Server-Side Help”](#)).

The `HELP` statement searches the help tables for the given search string and displays the result of the search. The search string is not case sensitive.

The search string can contain the the wildcard characters “%” and “\_”. These have the same meaning as for pattern-matching operations performed with the `LIKE` [\[1190\]](#) operator. For example, `HELP 'rep %'` returns a list of topics that begin with `rep`.

The `HELP` statement understands several types of search strings:

- At the most general level, use `contents` to retrieve a list of the top-level help categories:

```
HELP 'contents'
```

- For a list of topics in a given help category, such as `Data Types`, use the category name:

```
HELP 'data types'
```

- For help on a specific help topic, such as the `ASCII()` [\[1177\]](#) function or the `CREATE TABLE` statement, use the associated keyword or keywords:

```
HELP 'ascii'
HELP 'create table'
```

In other words, the search string matches a category, many topics, or a single topic. You cannot necessarily tell in advance whether a given search string will return a list of items or the help information for a single help topic. However, you can tell what kind of response `HELP` returned by examining the number of rows and columns in the result set.

The following descriptions indicate the forms that the result set can take. Output for the example statements is shown using the familiar “tabular” or “vertical” format that you see when using the `mysql` client, but note that `mysql` itself reformats `HELP` result sets in a different way.

- Empty result set

No match could be found for the search string.

- Result set containing a single row with three columns

This means that the search string yielded a hit for the help topic. The result has three columns:

- `name`: The topic name.
- `description`: Descriptive help text for the topic.
- `example`: Usage example or examples. This column might be blank.

Example: `HELP 'replace'`

Yields:

```
name: REPLACE
description: Syntax:
REPLACE(str,from_str,to_str)
```

```
Returns the string str with all occurrences of the string from_str
replaced by the string to_str. REPLACE() performs a case-sensitive
match when searching for from_str.
example: mysql> SELECT REPLACE('www.mysql.com', 'w', 'Ww');
-> 'WwWwWw.mysql.com'
```

- Result set containing multiple rows with two columns

This means that the search string matched many help topics. The result set indicates the help topic names:

- **name**: The help topic name.
- **is\_it\_category**: **Y** if the name represents a help category, **N** if it does not. If it does not, the **name** value when specified as the argument to the **HELP** statement should yield a single-row result set containing a description for the named item.

Example: **HELP 'status'**

Yields:

name	is_it_category
SHOW	N
SHOW ENGINE	N
SHOW MASTER STATUS	N
SHOW PROCEDURE STATUS	N
SHOW SLAVE STATUS	N
SHOW STATUS	N
SHOW TABLE STATUS	N

- Result set containing multiple rows with three columns

This means the search string matches a category. The result set contains category entries:

- **source\_category\_name**: The help category name.
- **name**: The category or topic name
- **is\_it\_category**: **Y** if the name represents a help category, **N** if it does not. If it does not, the **name** value when specified as the argument to the **HELP** statement should yield a single-row result set containing a description for the named item.

Example: **HELP 'functions'**

Yields:

source_category_name	name	is_it_category
Functions	CREATE FUNCTION	N
Functions	DROP FUNCTION	N
Functions	Bit Functions	Y
Functions	Comparison operators	Y
Functions	Control flow functions	Y
Functions	Date and Time Functions	Y
Functions	Encryption Functions	Y
Functions	Information Functions	Y
Functions	Logical operators	Y
Functions	Miscellaneous Functions	Y
Functions	Numeric Functions	Y
Functions	String Functions	Y

### 13.8.4. USE Syntax



```
USE db_name
```

The `USE db_name` statement tells MySQL to use the *db\_name* database as the default (current) database for subsequent statements. The database remains the default until the end of the session or another `USE` statement is issued:

```
USE db1;
SELECT COUNT(*) FROM mytable;    # selects from db1.mytable
USE db2;
SELECT COUNT(*) FROM mytable;    # selects from db2.mytable
```

Making a particular database the default by means of the `USE` statement does not preclude you from accessing tables in other databases. The following example accesses the `author` table from the `db1` database and the `editor` table from the `db2` database:

```
USE db1;
SELECT author_name,editor_name FROM author,db2.editor
WHERE author.editor_id = db2.editor.editor_id;
```

---

# Chapter 14. Storage Engines

## Table of Contents

14.1. Setting the Storage Engine .....	1604
14.2. The <a href="#">InnoDB</a> Storage Engine .....	1605
14.2.1. Getting Started with <a href="#">InnoDB</a> Tables .....	1606
14.2.2. Administering <a href="#">InnoDB</a> .....	1613
14.2.3. <a href="#">InnoDB</a> Concepts and Architecture .....	1622
14.2.4. <a href="#">InnoDB</a> Performance Tuning and Troubleshooting .....	1644
14.2.5. <a href="#">InnoDB</a> Features for Flexibility, Ease of Use and Reliability .....	1686
14.2.6. <a href="#">InnoDB</a> Startup Options and System Variables .....	1692
14.2.7. Limits on <a href="#">InnoDB</a> Tables .....	1756
14.2.8. MySQL and the ACID Model .....	1759
14.2.9. <a href="#">InnoDB</a> Integration with memcached .....	1761
14.3. The <a href="#">MyISAM</a> Storage Engine .....	1787
14.3.1. <a href="#">MyISAM</a> Startup Options .....	1789
14.3.2. Space Needed for Keys .....	1791
14.3.3. <a href="#">MyISAM</a> Table Storage Formats .....	1791
14.3.4. <a href="#">MyISAM</a> Table Problems .....	1794
14.4. The <a href="#">MEMORY</a> Storage Engine .....	1795
14.5. The <a href="#">CSV</a> Storage Engine .....	1799
14.5.1. Repairing and Checking CSV Tables .....	1800
14.5.2. CSV Limitations .....	1801
14.6. The <a href="#">ARCHIVE</a> Storage Engine .....	1801
14.7. The <a href="#">BLACKHOLE</a> Storage Engine .....	1802
14.8. The <a href="#">MERGE</a> Storage Engine .....	1804
14.8.1. <a href="#">MERGE</a> Table Advantages and Disadvantages .....	1806
14.8.2. <a href="#">MERGE</a> Table Problems .....	1807
14.9. The <a href="#">FEDERATED</a> Storage Engine .....	1809
14.9.1. <a href="#">FEDERATED</a> Storage Engine Overview .....	1809
14.9.2. How to Create <a href="#">FEDERATED</a> Tables .....	1810
14.9.3. <a href="#">FEDERATED</a> Storage Engine Notes and Tips .....	1813
14.9.4. <a href="#">FEDERATED</a> Storage Engine Resources .....	1814
14.10. The <a href="#">EXAMPLE</a> Storage Engine .....	1814
14.11. Other Storage Engines .....	1815
14.12. Overview of MySQL Storage Engine Architecture .....	1815
14.12.1. Pluggable Storage Engine Architecture .....	1815
14.12.2. The Common Database Server Layer .....	1816

Storage engines are MySQL components that handle the SQL operations for different table types. [InnoDB](#) is the most general-purpose storage engine, and Oracle recommends using it for tables except for specialized use cases. (The [CREATE TABLE](#) statement in MySQL 5.6 creates [InnoDB](#) tables by default.)

MySQL Server uses a pluggable storage engine architecture that enables storage engines to be loaded into and unloaded from a running MySQL server.

To determine which storage engines your server supports, use the [SHOW ENGINES](#) statement. The value in the [Support](#) column indicates whether an engine can be used. A value of [YES](#), [NO](#), or [DEFAULT](#) indicates that an engine is available, not available, or available and currently set as the default storage engine.

This chapter primarily describes the features and performance characteristics of [InnoDB](#) tables. It also covers the use cases for the special-purpose MySQL storage engines, except for [NDB](#) which is covered in [MySQL Cluster NDB 7.2](#). For advanced users, it also contains a description of the pluggable storage engine architecture (see [Section 14.12, "Overview of MySQL Storage Engine Architecture"](#)).

For information about storage engine support offered in commercial MySQL Server binaries, see [MySQL Enterprise Server 5.6](#), on the MySQL Web site. The storage engines available might depend on which edition of Enterprise Server you are using.

For answers to some commonly asked questions about MySQL storage engines, see [Section B.2, “MySQL 5.6 FAQ: Storage Engines”](#).

## MySQL 5.6 Supported storage Engines

- **InnoDB**: A transaction-safe (ACID compliant) storage engine for MySQL that has commit, rollback, and crash-recovery capabilities to protect user data. **InnoDB** row-level locking (without escalation to coarser granularity locks) and Oracle-style consistent nonlocking reads increase multi-user concurrency and performance. **InnoDB** stores user data in clustered indexes to reduce I/O for common queries based on primary keys. To maintain data integrity, **InnoDB** also supports **FOREIGN KEY** referential-integrity constraints. **InnoDB** is the default storage engine in MySQL 5.6.
- **MyISAM**: These tables have a small footprint. **Table-level locking** limits the performance in read/write workloads, so it is often used in read-only or read-mostly workloads in Web and data warehousing configurations.
- **Memory**: Stores all data in RAM, for fast access in environments that require quick lookups of non-critical data. This engine was formerly known as the **HEAP** engine. Its use cases are decreasing; **InnoDB** with its buffer pool memory area provides a general-purpose and durable way to keep most or all data in memory, and **NDBCLUSTER** provides fast key-value lookups for huge distributed data sets.
- **CSV**: Its tables are really text files with comma-separated values. CSV tables let you import or dump data in CSV format, to exchange data with scripts and applications that read and write that same format. Because CSV tables are not indexed, you typically keep the data in **InnoDB** tables during normal operation, and only use CSV tables during the import or export stage.
- **Archive**: These compact, unindexed tables are intended for storing and retrieving large amounts of seldom-referenced historical, archived, or security audit information.
- **Blackhole**: The Blackhole storage engine accepts but does not store data, similar to the Unix `/dev/null` device. Queries always return an empty set. These tables can be used in replication configurations where DML statements are sent to slave servers, but the master server does not keep its own copy of the data.
- **Merge**: Enables a MySQL DBA or developer to logically group a series of identical **MyISAM** tables and reference them as one object. Good for VLDB environments such as data warehousing.
- **Federated**: Offers the ability to link separate MySQL servers to create one logical database from many physical servers. Very good for distributed or data mart environments.
- **Example**: This engine serves as an example in the MySQL source code that illustrates how to begin writing new storage engines. It is primarily of interest to developers. The storage engine is a “stub” that does nothing. You can create tables with this engine, but no data can be stored in them or retrieved from them.

You are not restricted to using the same storage engine for an entire server or schema. You can specify the storage engine for any table. For example, an application might use mostly **InnoDB** tables, with one **CSV** table for exporting data to a spreadsheet and a few **MEMORY** tables for temporary workspaces.

### Choosing a Storage Engine

The various storage engines provided with MySQL are designed with different use cases in mind. The following table provides an overview of some storage engines provided with MySQL:

**Table 14.1. Storage Engines Feature Summary**

<b>Feature</b>	<b>MyISAM</b>	<b>Memory</b>	<b>InnoDB</b>	<b>Archive</b>	<b>NDB</b>
Storage limits	256TB	RAM	64TB	None	384EB
Transactions	No	No	Yes	No	Yes
Locking granularity	Table	Table	Row	Table	Row
MVCC	No	No	Yes	No	No
Geospatial data type support	Yes	No	Yes	Yes	Yes
Geospatial indexing support	Yes	No	No	No	No
B-tree indexes	Yes	Yes	Yes	No	No
T-tree indexes	No	No	No	No	Yes
Hash indexes	No	Yes	No <sup>a</sup>	No	Yes
Full-text search indexes	Yes	No	Yes <sup>b</sup>	No	No
Clustered indexes	No	No	Yes	No	No
Data caches	No	N/A	Yes	No	Yes
Index caches	Yes	N/A	Yes	No	Yes
Compressed data	Yes <sup>c</sup>	No	Yes <sup>d</sup>	Yes	No
Encrypted data <sup>e</sup>	Yes	Yes	Yes	Yes	Yes
Cluster database support	No	No	No	No	Yes
Replication support <sup>f</sup>	Yes	Yes	Yes	Yes	Yes
Foreign key support	No	No	Yes	No	No
Backup / point-in-time recovery <sup>g</sup>	Yes	Yes	Yes	Yes	Yes
Query cache support	Yes	Yes	Yes	Yes	Yes

Feature	MyISAM	Memory	InnoDB	Archive	NDB
Update statistics for data dictionary	Yes	Yes	Yes	Yes	Yes

<sup>a</sup>InnoDB utilizes hash indexes internally for its Adaptive Hash Index feature.

<sup>b</sup>InnoDB support for FULLTEXT indexes is available in MySQL 5.6.4 and higher.

<sup>c</sup>Compressed MyISAM tables are supported only when using the compressed row format. Tables using the compressed row format with MyISAM are read only.

<sup>d</sup>Compressed InnoDB tables require the InnoDB Barracuda file format.

<sup>e</sup>Implemented in the server (via encryption functions), rather than in the storage engine.

<sup>f</sup>Implemented in the server, rather than in the storage engine.

<sup>g</sup>Implemented in the server, rather than in the storage engine.

## 14.1. Setting the Storage Engine

When you create a new table, you can specify a special-purpose storage engine to use by adding an `ENGINE` table option to the `CREATE TABLE` statement:

```
-- ENGINE=INNODB not needed unless you have set a different default storage engine.
CREATE TABLE t1 (i INT) ENGINE = INNODB;
-- Simple table definitions can be switched from one to another.
CREATE TABLE t2 (i INT) ENGINE = CSV;
-- Some storage engines have their own specific clauses in CREATE TABLE syntax.
CREATE TABLE t3 (i INT) ENGINE = MEMORY USING BTREE;
```

When you omit the `ENGINE` option, the default storage engine is used. The default engine is `InnoDB` in MySQL 5.6. You can specify the default engine by using the `--default-storage-engine` [420] server startup option, or by setting the `default-storage-engine` [420] option in the `my.cnf` configuration file.

You can set the default storage engine for the current session by setting the `default_storage_engine` [474] variable:

```
SET default_storage_engine=NDBCLUSTER;
```

As of MySQL 5.6.3, the storage engine for `TEMPORARY` tables created with `CREATE TEMPORARY TABLE` can be set separately from the engine for permanent tables by setting the `default_tmp_storage_engine` [474], either at startup or at runtime. Before MySQL 5.6.3, `default_storage_engine` [474] sets the engine for both permanent and `TEMPORARY` tables.

To convert a table from one storage engine to another, use an `ALTER TABLE` statement that indicates the new engine:

```
ALTER TABLE t ENGINE = InnoDB;
```

See [Section 13.1.14, “CREATE TABLE Syntax”](#), and [Section 13.1.6, “ALTER TABLE Syntax”](#).

If you try to use a storage engine that is not compiled in or that is compiled in but deactivated, MySQL instead creates a table using the default storage engine. For example, in a replication setup, perhaps your master server uses `InnoDB` tables for maximum safety, but the slave servers use other storage engines for speed at the expense of durability or concurrency.

By default, a warning is generated whenever `CREATE TABLE` or `ALTER TABLE` cannot use the default storage engine. To prevent confusing, unintended behavior if the desired engine is unavailable, enable the `NO_ENGINE_SUBSTITUTION` [605] SQL mode. If the desired engine is unavailable, this setting produces an error instead of a warning, and the table is not created or altered. See [Section 5.1.7, “Server SQL Modes”](#).

For new tables, MySQL always creates an `.frm` file to hold the table and column definitions. The table's index and data may be stored in one or more other files, depending on the storage engine. The server creates the `.frm` file above the storage engine level. Individual storage engines create

any additional files required for the tables that they manage. If a table name contains special characters, the names for the table files contain encoded versions of those characters as described in Section 9.2.3, “Mapping of Identifiers to File Names”.

## 14.2. The InnoDB Storage Engine

InnoDB is a general-purpose storage engine that balances high reliability and high performance. In MySQL 5.6, issuing the `CREATE TABLE` statement with no `ENGINE=` clause creates an InnoDB table.

Key advantages of InnoDB tables include:

- Its DML operations follow the ACID model, with transactions featuring commit, rollback, and crash-recovery capabilities to protect user data.
- Row-level locking and Oracle-style consistent reads increase multi-user concurrency and performance.
- InnoDB tables arrange your data on disk to optimize queries based on primary keys.
- To maintain data integrity, InnoDB also supports FOREIGN KEY constraints. Inserts, updates, and deletes are all checked to ensure they do not result in inconsistencies across different tables.
- You can freely mix InnoDB tables with tables from other MySQL storage engines, even within the same statement. For example, you can use a join operation to combine data from InnoDB and MEMORY tables in a single query.
- The latest InnoDB offers significant new features over MySQL 5.1 and earlier. These features focus on performance and scalability, reliability, flexibility, and usability:
  - **Fast index creation:** add or drop indexes without copying the data.
  - **Data compression:** shrink tables, to significantly reduce storage and I/O.
  - **More efficient storage for large column values:** fully off-page storage of long BLOB, TEXT, and VARCHAR columns.
  - **Barracuda file format:** enables new features while protecting upward and downward compatibility
  - **INFORMATION\_SCHEMA tables:** information about compression and locking
  - **Performance and scalability enhancements:** includes features such as multiple background I/O threads, multiple buffer pools, and group commit.
  - **Other changes:** for flexibility, ease of use and reliability.

**Table 14.2. InnoDB Storage Engine Features**

<b>Storage limits</b>	64TB	<b>Transactions</b>	Yes	<b>Locking granularity</b>	Row
<b>MVCC</b>	Yes	<b>Geospatial data type support</b>	Yes	<b>Geospatial indexing support</b>	No
<b>B-tree indexes</b>	Yes	<b>T-tree indexes</b>	No	<b>Hash indexes</b>	No <sup>a</sup>
<b>Full-text search indexes</b>	Yes <sup>b</sup>	<b>Clustered indexes</b>	Yes	<b>Data caches</b>	Yes
<b>Index caches</b>	Yes	<b>Compressed data</b>	Yes <sup>c</sup>	<b>Encrypted data<sup>d</sup></b>	Yes
<b>Cluster database support</b>	No	<b>Replication support<sup>e</sup></b>	Yes	<b>Foreign key support</b>	Yes
<b>Backup / point-in-time recovery<sup>f</sup></b>	Yes	<b>Query cache support</b>	Yes	<b>Update statistics for data dictionary</b>	Yes

<sup>a</sup>InnoDB utilizes hash indexes internally for its Adaptive Hash Index feature.

<sup>b</sup>InnoDB support for FULLTEXT indexes is available in MySQL 5.6.4 and higher.

<sup>c</sup>Compressed InnoDB tables require the InnoDB Barracuda file format.

<sup>d</sup>Implemented in the server (via encryption functions), rather than in the storage engine.

<sup>e</sup>Implemented in the server, rather than in the storage engine.

<sup>f</sup>Implemented in the server, rather than in the storage engine.

InnoDB has been designed for maximum performance when processing large data volumes. Its CPU efficiency is probably not matched by any other disk-based relational database engine.

The InnoDB storage engine maintains its own [buffer pool](#) for caching data and indexes in main memory. By default, with the [innodb\\_file\\_per\\_table \[1718\]](#) setting enabled, each new InnoDB table and its associated indexes are stored in a separate file. When the [innodb\\_file\\_per\\_table \[1718\]](#) option is disabled, InnoDB stores all its tables and indexes in the single [system tablespace](#), which may consist of several files (or raw disk partitions). InnoDB tables can handle large quantities of data, even on operating systems where file size is limited to 2GB.

InnoDB is published under the same GNU GPL License Version 2 (of June 1991) as MySQL. For more information on MySQL licensing, see <http://www.mysql.com/company/legal/licensing/>.

## Additional Resources

- For InnoDB-related terms and definitions, see [MySQL Glossary](#).
- A forum dedicated to the InnoDB storage engine is available at <http://forums.mysql.com/list.php?22>.

### 14.2.1. Getting Started with InnoDB Tables

#### 14.2.1.1. InnoDB as the Default MySQL Storage Engine

MySQL has a well-earned reputation for being easy-to-use and delivering performance and scalability. In previous versions of MySQL, MyISAM was the default storage engine. In our experience, most users never changed the default settings. With MySQL 5.5, InnoDB becomes the default storage engine. Again, we expect most users will not change the default settings. But, because of InnoDB, the default settings deliver the benefits users expect from their RDBMS: ACID Transactions, Referential Integrity, and Crash Recovery. Let's explore how using InnoDB tables improves your life as a MySQL user, DBA, or developer.

#### Trends in Storage Engine Usage

In the first years of MySQL growth, early web-based applications didn't push the limits of concurrency and availability. In 2010, hard drive and memory capacity and the performance/price ratio have all gone through the roof. Users pushing the performance boundaries of MySQL care a lot about reliability and crash recovery. MySQL databases are big, busy, robust, distributed, and important.

InnoDB hits the sweet spot of these top user priorities. The trend of storage engine usage has shifted in favor of the more scalable InnoDB. Thus MySQL 5.5 is the logical transition release to make InnoDB the default storage engine.

MySQL continues to work on addressing use cases that formerly required MyISAM tables. In MySQL 5.6 and higher:

- InnoDB can perform full-text search using the [FULLTEXT](#) index type. See [Section 14.2.3.12.3, "FULLTEXT Indexes"](#) for details.
- InnoDB now performs better with read-only or read-mostly workloads. Automatic optimizations apply to InnoDB queries in [autocommit](#) mode, and you can explicitly mark transactions as read-only with the syntax [START TRANSACTION READ ONLY](#). See [Section 14.2.4.2.3, "Optimizations for Read-Only Transactions"](#) for details.
- Applications distributed on read-only media can now use InnoDB tables. See [Section 14.2.5.1, "Support for Read-Only Media"](#) for details.



## Consequences of InnoDB as Default MySQL Storage Engine

Starting from MySQL 5.5.5, the default storage engine for new tables is InnoDB. This change applies to newly created tables that don't specify a storage engine with a clause such as `ENGINE=MyISAM`. (Given this change of default behavior, MySQL 5.5 might be a logical point to evaluate whether your tables that do use MyISAM could benefit from switching to InnoDB.)

The `mysql` and `information_schema` databases, that implement some of the MySQL internals, still use MyISAM. In particular, you cannot switch the grant tables to use InnoDB.

## Benefits of InnoDB Tables

If you use `MyISAM` tables but aren't tied to them for technical reasons, you'll find many things more convenient when you use `InnoDB` tables in MySQL 5.5:

- If your server crashes because of a hardware or software issue, regardless of what was happening in the database at the time, you don't need to do anything special after restarting the database. InnoDB [crash recovery](#) automatically finalizes any changes that were committed before the time of the crash, and undoes any changes that were in process but not committed. Just restart and continue where you left off. This process is now much faster than in MySQL 5.1 and earlier.
- The InnoDB [buffer pool](#) caches table and index data as the data is accessed. Frequently used data is processed directly from memory. This cache applies to so many types of information, and speeds up processing so much, that dedicated database servers assign up to 80% of their physical memory to the InnoDB buffer pool.
- If you split up related data into different tables, you can set up [foreign keys](#) that enforce [referential integrity](#). Update or delete data, and the related data in other tables is updated or deleted automatically. Try to insert data into a secondary table without corresponding data in the primary table, and the bad data gets kicked out automatically.
- If data becomes corrupted on disk or in memory, a [checksum](#) mechanism alerts you to the bogus data before you use it.
- When you design your database with appropriate [primary key](#) columns for each table, operations involving those columns are automatically optimized. It is very fast to reference the primary key columns in `WHERE` clauses, `ORDER BY` clauses, `GROUP BY` clauses, and `join` operations.
- Inserts, updates, deletes are optimized by an automatic mechanism called [change buffering](#). InnoDB not only allows concurrent read and write access to the same table, it caches changed data to streamline disk I/O.
- Performance benefits are not limited to giant tables with long-running queries. When the same rows are accessed over and over from a table, a feature called the [Adaptive Hash Index](#) takes over to make these lookups even faster, as if they came out of a hash table.

## Best Practices for InnoDB Tables

If you have been using `InnoDB` for a long time, you already know about features like transactions and foreign keys. If not, read about them throughout this chapter. To make a long story short:

- Specify a [primary key](#) for every table using the most frequently queried column or columns, or an [auto-increment](#) value if there is no obvious primary key.
- Embrace the idea of [joins](#), where data is pulled from multiple tables based on identical ID values from those tables. For fast join performance, define [foreign keys](#) on the join columns, and declare those columns with the same data type in each table. The foreign keys also propagate deletes or updates to all affected tables, and prevent insertion of data in a child table if the corresponding IDs are not present in the parent table.
- Turn off [autocommit](#). Committing hundreds of times a second puts a cap on performance (limited by the write speed of your storage device).

- Group sets of related [DML](#) operations into [transactions](#), by bracketing them with [START TRANSACTION](#) and [COMMIT](#) statements. While you don't want to commit too often, you also don't want to issue huge batches of [INSERT](#), [UPDATE](#), or [DELETE](#) statements that run for hours without committing.
- Stop using [LOCK TABLE](#) statements. InnoDB can handle multiple sessions all reading and writing to the same table at once, without sacrificing reliability or high performance. To get exclusive write access to a set of rows, use the [SELECT ... FOR UPDATE](#) syntax to lock just the rows you intend to update.
- Enable the [innodb\\_file\\_per\\_table](#) option to put the data and indexes for individual tables into separate files, instead of in a single giant [system tablespace](#). (This setting is required to use some of the other features, such as table [compression](#) and fast [truncation](#).)
- Evaluate whether your data and access patterns benefit from the new InnoDB table [compression](#) feature ([ROW\\_FORMAT=COMPRESSED](#) on the [CREATE TABLE](#) statement. You can compress InnoDB tables without sacrificing read/write capability.
- Run your server with the option [--sql\\_mode=NO\\_ENGINE\\_SUBSTITUTION](#) to prevent tables being created with a different storage engine if there is an issue with the one specified in the [ENGINE=](#) clause of [CREATE TABLE](#).

## Recent Improvements for InnoDB Tables

If you have experience with InnoDB, but from MySQL 5.1 or earlier, read about the latest InnoDB enhancements in [Section 14.2.4.2, “InnoDB Performance and Scalability Enhancements”](#) and [Section 14.2.5, “InnoDB Features for Flexibility, Ease of Use and Reliability”](#). To make a long story short:

- You can compress tables and associated indexes.
- You can create and drop indexes with much less performance or availability impact than before.
- Truncating a table is very fast, and can free up disk space for the operating system to reuse, rather than freeing up space within the system tablespace that only InnoDB could reuse.
- The storage layout for table data is more efficient for BLOBs and long text fields, with the [DYNAMIC](#) row format.
- You can monitor the internal workings of the storage engine by querying [INFORMATION\\_SCHEMA](#) tables.
- You can monitor the performance details of the storage engine by querying [performance\\_schema](#) tables.
- There are many many performance improvements. In particular, crash recovery, the automatic process that makes all data consistent when the database is restarted, is fast and reliable. (Now much much faster than long-time InnoDB users are used to.) The bigger the database, the more dramatic the speedup.

Most new performance features are automatic, or at most require setting a value for a configuration option. For details, see [Section 14.2.4.2, “InnoDB Performance and Scalability Enhancements”](#). For InnoDB-specific tuning techniques you can apply in your application code, see [Section 8.5, “Optimizing for InnoDB Tables”](#). Advanced users can review [Section 14.2.6, “InnoDB Startup Options and System Variables”](#).

## Testing and Benchmarking with InnoDB as Default Storage Engine

Even before completing your upgrade from MySQL 5.1 or earlier to MySQL 5.5 or higher, you can preview whether your database server or application works correctly with InnoDB as the default storage engine. To set up InnoDB as the default storage engine with an earlier MySQL release, either specify

on the command line `--default-storage-engine=InnoDB`, or add to your `my.cnf` file `default-storage-engine=innodb` in the `[mysqld]` section, then restart the server.

Since changing the default storage engine only affects new tables as they are created, run all your application installation and setup steps to confirm that everything installs properly. Then exercise all the application features to make sure all the data loading, editing, and querying features work. If a table relies on some MyISAM-specific feature, you'll receive an error; add the `ENGINE=MyISAM` clause to the `CREATE TABLE` statement to avoid the error. (For example, tables that rely on full-text search must be MyISAM tables rather than InnoDB ones.)

If you did not make a deliberate decision about the storage engine, and you just want to preview how certain tables work when they're created under InnoDB, issue the command `ALTER TABLE table_name ENGINE=InnoDB;` for each table. Or, to run test queries and other statements without disturbing the original table, make a copy like so:

```
CREATE TABLE InnoDB_Table (...) ENGINE=InnoDB AS SELECT * FROM MyISAM_Table;
```

Since there are so many performance enhancements in InnoDB in MySQL 5.5 and higher, to get a true idea of the performance with a full application under a realistic workload, install the latest MySQL server and run benchmarks.

Test the full application lifecycle, from installation, through heavy usage, and server restart. Kill the server process while the database is busy to simulate a power failure, and verify that the data is recovered successfully when you restart the server.

Test any replication configurations, especially if you use different MySQL versions and options on the master and the slaves.

## Verifying that InnoDB is the Default Storage Engine

To know what the status of InnoDB is, whether you're doing what-if testing with an older MySQL or comprehensive testing with the latest MySQL:

- Issue the command `SHOW ENGINES;` to see all the different MySQL storage engines. Look for `DEFAULT` in the InnoDB line.
- If InnoDB is not present at all, you have a `mysqld` binary that was compiled without InnoDB support and you need to get a different one.
- If InnoDB is present but disabled, go back through your startup options and configuration file and get rid of any `skip-innodb` option.

### 14.2.1.2. Configuring InnoDB

The first decisions to make about InnoDB configuration involve how to lay out InnoDB data files, and how much memory to allocate for the InnoDB storage engine. You record these choices either by recording them in a configuration file that MySQL reads at startup, or by specifying them as command-line options in a startup script. The full list of options, descriptions, and allowed parameter values is at [Section 14.2.6, "InnoDB Startup Options and System Variables"](#).

## Overview of InnoDB Tablespace and Log Files

Two important disk-based resources managed by the InnoDB storage engine are its tablespace data files and its log files. If you specify no InnoDB configuration options, MySQL creates an auto-extending data file, slightly larger than 12MB, named `ibdata1` and two log files named `ib_logfile0` and `ib_logfile1` in the MySQL data directory. Their size is given by the size of the `innodb_log_file_size [1732]` system variable. To get good performance, explicitly provide InnoDB parameters as discussed in the following examples. Naturally, edit the settings to suit your hardware and requirements.

The examples shown here are representative. See [Section 14.2.6, "InnoDB Startup Options and System Variables"](#) for additional information about InnoDB-related configuration parameters.

## Considerations for Storage Devices

In some cases, database performance improves if the data is not all placed on the same physical disk. Putting log files on a different disk from data is very often beneficial for performance. The example illustrates how to do this. It places the two data files on different disks and places the log files on the third disk. InnoDB fills the tablespace beginning with the first data file. You can also use raw disk partitions (raw devices) as InnoDB data files, which may speed up I/O. See [Section 14.2.2.3, “Using Raw Disk Partitions for the Shared Tablespace”](#).



### Caution

InnoDB is a transaction-safe (ACID compliant) storage engine for MySQL that has commit, rollback, and crash-recovery capabilities to protect user data. **However, it cannot do so** if the underlying operating system or hardware does not work as advertised. Many operating systems or disk subsystems may delay or reorder write operations to improve performance. On some operating systems, the very `fsync()` system call that should wait until all unwritten data for a file has been flushed might actually return before the data has been flushed to stable storage. Because of this, an operating system crash or a power outage may destroy recently committed data, or in the worst case, even corrupt the database because of write operations having been reordered. If data integrity is important to you, perform some “pull-the-plug” tests before using anything in production. On Mac OS X 10.3 and up, InnoDB uses a special `fcntl()` file flush method. Under Linux, it is advisable to **disable the write-back cache**.

On ATA/SATA disk drives, a command such `hdparm -W0 /dev/hda` may work to disable the write-back cache. **Beware that some drives or disk controllers may be unable to disable the write-back cache.**



### Caution

If reliability is a consideration for your data, do not configure InnoDB to use data files or log files on NFS volumes. Potential problems vary according to OS and version of NFS, and include such issues as lack of protection from conflicting writes, and limitations on maximum file sizes.

## Specifying the Location and Size for InnoDB Tablespace Files

To set up the InnoDB tablespace files, use the `innodb_data_file_path` [1714] option in the `[mysqld]` section of the `my.cnf` option file. On Windows, you can use `my.ini` instead. The value of `innodb_data_file_path` [1714] should be a list of one or more data file specifications. If you name more than one data file, separate them by semicolon (“;”) characters:

```
innodb_data_file_path=datafile_spec1[:datafile_spec2]...
```

For example, the following setting explicitly creates a minimally sized system tablespace:

```
[mysqld]
innodb_data_file_path=ibdata1:12M:autoextend
```

This setting configures a single 12MB data file named `ibdata1` that is auto-extending. No location for the file is given, so by default, InnoDB creates it in the MySQL data directory.

Sizes are specified using `K`, `M`, or `G` suffix letters to indicate units of KB, MB, or GB.

A tablespace containing a fixed-size 50MB data file named `ibdata1` and a 50MB auto-extending file named `ibdata2` in the data directory can be configured like this:

```
[mysqld]
innodb_data_file_path=ibdata1:50M;ibdata2:50M:autoextend
```

The full syntax for a data file specification includes the file name, its size, and several optional attributes:

```
file_name:file_size[:autoextend[:max:max_file_size]]
```

The `autoextend` and `max` attributes can be used only for the last data file in the `innodb_data_file_path` [1714] line.

If you specify the `autoextend` option for the last data file, InnoDB extends the data file if it runs out of free space in the tablespace. The increment is 8MB at a time by default. To modify the increment, change the `innodb_autoextend_increment` [1704] system variable.

If the disk becomes full, you might want to add another data file on another disk. For tablespace reconfiguration instructions, see [Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”](#).

InnoDB is not aware of the file system maximum file size, so be cautious on file systems where the maximum file size is a small value such as 2GB. To specify a maximum size for an auto-extending data file, use the `max` attribute following the `autoextend` attribute. Use the `max` attribute only in cases where constraining disk usage is of critical importance, because exceeding the maximum size causes a fatal error, possibly including a crash. The following configuration permits `ibdata1` to grow up to a limit of 500MB:

```
[mysqld]
innodb_data_file_path=ibdata1:12M:autoextend:max:500M
```

InnoDB creates tablespace files in the MySQL data directory by default. To specify a location explicitly, use the `innodb_data_home_dir` [1715] option. For example, to use two files named `ibdata1` and `ibdata2` but create them in the `/ibdata` directory, configure InnoDB like this:

```
[mysqld]
innodb_data_home_dir = /ibdata
innodb_data_file_path=ibdata1:50M;ibdata2:50M:autoextend
```



#### Note

InnoDB does not create directories, so make sure that the `/ibdata` directory exists before you start the server. This is also true of any log file directories that you configure. Use the Unix or DOS `mkdir` command to create any necessary directories.

Make sure that the MySQL server has the proper access rights to create files in the data directory. More generally, the server must have access rights in any directory where it needs to create data files or log files.

InnoDB forms the directory path for each data file by textually concatenating the value of `innodb_data_home_dir` [1715] to the data file name, adding a path name separator (slash or backslash) between values if necessary. If the `innodb_data_home_dir` [1715] option is not specified in `my.cnf` at all, the default value is the “dot” directory `./`, which means the MySQL data directory. (The MySQL server changes its current working directory to its data directory when it begins executing.)

If you specify `innodb_data_home_dir` [1715] as an empty string, you can specify absolute paths for the data files listed in the `innodb_data_file_path` [1714] value. The following example is equivalent to the preceding one:

```
[mysqld]
innodb_data_home_dir =
innodb_data_file_path=/ibdata/ibdata1:50M;/ibdata/ibdata2:50M:autoextend
```

## Specifying InnoDB Configuration Options

**Sample `my.cnf` file for small systems.** Suppose that you have a computer with 512MB RAM and one hard disk. The following example shows possible configuration parameters in `my.cnf` or `my.ini`

for InnoDB, including the `autoextend` attribute. The example suits most users, both on Unix and Windows, who do not want to distribute InnoDB data files and log files onto several disks. It creates an auto-extending data file `ibdata1` and two InnoDB log files `ib_logfile0` and `ib_logfile1` in the MySQL data directory.

```
[mysqld]
# You can write your other MySQL server options here
# ...
# Data files must be able to hold your data and indexes.
# Make sure that you have enough free disk space.
innodb_data_file_path = ibdata1:12M:autoextend
#
# Set buffer pool size to 50-80% of your computer's memory
innodb_buffer_pool_size=256M
innodb_additional_mem_pool_size=20M
#
# Set the log file size to about 25% of the buffer pool size
innodb_log_file_size=64M
innodb_log_buffer_size=8M
#
innodb_flush_log_at_trx_commit=1
```

Note that data files must be less than 2GB in some file systems. The combined size of the log files can be up to 512GB. The combined size of data files must be at least slightly larger than 10MB.

## Setting Up the InnoDB System Tablespace

When you create an InnoDB system tablespace for the first time, it is best that you start the MySQL server from the command prompt. InnoDB then prints the information about the database creation to the screen, so you can see what is happening. For example, on Windows, if `mysqld` is located in `C:\Program Files\MySQL\MySQL Server 5.6\bin`, you can start it like this:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld" --console
```

If you do not send server output to the screen, check the server's error log to see what InnoDB prints during the startup process.

For an example of what the information displayed by InnoDB should look like, see [Section 14.2.2.1, "Creating the InnoDB Tablespace"](#).

## Editing the MySQL Configuration File

You can place InnoDB options in the `[mysqld]` group of any option file that your server reads when it starts. The locations for option files are described in [Section 4.2.3.3, "Using Option Files"](#).

If you installed MySQL on Windows using the installation and configuration wizards, the option file will be the `my.ini` file located in your MySQL installation directory. See [Section 2.3.3, "Installing MySQL on Microsoft Windows Using MySQL Installer"](#).

If your PC uses a boot loader where the `C:` drive is not the boot drive, your only option is to use the `my.ini` file in your Windows directory (typically `C:\WINDOWS`). You can use the `SET` command at the command prompt in a console window to print the value of `WINDIR`:

```
C:\> SET WINDIR
windir=C:\WINDOWS
```

To make sure that `mysqld` reads options only from a specific file, use the `--defaults-file` [\[219\]](#) option as the first option on the command line when starting the server:

```
mysqld --defaults-file=your_path_to_my_cnf
```

**Sample `my.cnf` file for large systems.** Suppose that you have a Linux computer with 2GB RAM and three 60GB hard disks at directory paths `/`, `/dr2` and `/dr3`. The following example shows possible configuration parameters in `my.cnf` for InnoDB.

```
[mysqld]
```



```
# You can write your other MySQL server options here
# ...
innodb_data_home_dir =
#
# Data files must be able to hold your data and indexes
innodb_data_file_path = /db/ibdata1:2000M:/dr2/db/ibdata2:2000M:autoextend
#
# Set buffer pool size to 50-80% of your computer's memory,
# but make sure on Linux x86 total memory usage is < 2GB
innodb_buffer_pool_size=1G
innodb_additional_mem_pool_size=20M
innodb_log_group_home_dir = /dr3/iblogs
#
# Set the log file size to about 25% of the buffer pool size
innodb_log_file_size=250M
innodb_log_buffer_size=8M
#
innodb_flush_log_at_trx_commit=1
innodb_lock_wait_timeout=50
#
# Uncomment the next line if you want to use it
#innodb_thread_concurrency=5
```

## Determining the Maximum Memory Allocation for InnoDB



### Warning

On 32-bit GNU/Linux x86, be careful not to set memory usage too high. `glibc` may permit the process heap to grow over thread stacks, which crashes your server. It is a risk if the value of the following expression is close to or exceeds 2GB:

```
innodb_buffer_pool_size
+ key_buffer_size
+ max_connections*(sort_buffer_size+read_buffer_size+binlog_cache_size)
+ max_connections*2MB
```

Each thread uses a stack (often 2MB, but only 256KB in MySQL binaries provided by Oracle Corporation.) and in the worst case also uses `sort_buffer_size + read_buffer_size` additional memory.

**Tuning other `mysqld` server parameters.** The following values are typical and suit most users:

```
[mysqld]
skip-external-locking
max_connections=200
read_buffer_size=1M
sort_buffer_size=1M
#
# Set key_buffer to 5 - 50% of your RAM depending on how much
# you use MyISAM tables, but keep key_buffer_size + InnoDB
# buffer pool size < 80% of your RAM
key_buffer_size=value
```

On Linux, if the kernel is enabled for large page support, `InnoDB` can use large pages to allocate memory for its buffer pool and additional memory pool. See [Section 8.11.4.2, "Enabling Large Page Support"](#).

## 14.2.2. Administering InnoDB

Administration tasks related to `InnoDB` mainly involve these aspects:

- Managing the [data files](#) that represent the [system tablespace](#), `InnoDB` tables, and their associated indexes. You can change the way these files are laid out and divided, which affects both performance and the features available for specific tables.
- Managing the [redo log](#) files that are used for [crash recovery](#). You can specify the size of these files.



- Making sure that [InnoDB](#) is used for the tables where it is intended, rather than a different storage engine.
- General administrative tasks related to performance. You might consult with application developers during the application design phase, monitor performance on an ongoing basis to ensure the system settings are working well, and diagnose and help fix performance and capacity issues that arise suddenly.

Since [InnoDB](#) tables are now the default for MySQL, much of the associated administration material is now in the main “Administration” chapter, [Chapter 5, MySQL Server Administration](#).

### 14.2.2.1. Creating the [InnoDB](#) Tablespace

Suppose that you have installed MySQL and have edited your option file so that it contains the necessary [InnoDB](#) configuration parameters. Before starting MySQL, verify that the directories you have specified for [InnoDB](#) data files and log files exist and that the MySQL server has access rights to those directories. [InnoDB](#) does not create directories, only files. Check also that you have enough disk space for the data and log files.

It is best to run the MySQL server `mysqld` from the command prompt when you first start the server with [InnoDB](#) enabled, not from `mysqld_safe` or as a Windows service. When you run from a command prompt you see what `mysqld` prints and what is happening. On Unix, just invoke `mysqld`. On Windows, start `mysqld` with the `--console` [417] option to direct the output to the console window.

When you start the MySQL server after initially configuring [InnoDB](#) in your option file, [InnoDB](#) creates your data files and log files, and prints something like this:

```
InnoDB: The first specified datafile /home/heikki/data/ibdata1
did not exist:
InnoDB: a new database to be created!
InnoDB: Setting file /home/heikki/data/ibdata1 size to 134217728
InnoDB: Database physically writes the file full: wait...
InnoDB: datafile /home/heikki/data/ibdata2 did not exist:
new to be created
InnoDB: Setting file /home/heikki/data/ibdata2 size to 262144000
InnoDB: Database physically writes the file full: wait...
InnoDB: Log file /home/heikki/data/logs/ib_logfile0 did not exist:
new to be created
InnoDB: Setting log file /home/heikki/data/logs/ib_logfile0 size
to 5242880
InnoDB: Log file /home/heikki/data/logs/ib_logfile1 did not exist:
new to be created
InnoDB: Setting log file /home/heikki/data/logs/ib_logfile1 size
to 5242880
InnoDB: Doublewrite buffer not found: creating new
InnoDB: Doublewrite buffer created
InnoDB: Creating foreign key constraint system tables
InnoDB: Foreign key constraint system tables created
InnoDB: Started
mysqld: ready for connections
```

At this point [InnoDB](#) has initialized its tablespace and log files. You can connect to the MySQL server with the usual MySQL client programs like `mysql`. When you shut down the MySQL server with `mysqladmin shutdown`, the output is like this:

```
010321 18:33:34 mysqld: Normal shutdown
010321 18:33:34 mysqld: Shutdown Complete
InnoDB: Starting shutdown...
InnoDB: Shutdown completed
```

You can look at the data file and log directories and you see the files created there. When MySQL is started again, the data files and log files have been created already, so the output is much briefer:

```
InnoDB: Started
mysqld: ready for connections
```

If you add the `innodb_file_per_table` [1718] option to `my.cnf`, InnoDB stores each table in its own `.ibd` file, in the same MySQL database directory where the `.frm` file is created. See Section 5.4.1, “Managing InnoDB Tablespaces”.

#### 14.2.2.2. Adding, Removing, or Resizing InnoDB Data and Log Files

This section describes what you can do when your InnoDB system tablespace runs out of room or when you want to change the size of the redo log files.

The easiest way to increase the size of the InnoDB system tablespace is to configure it from the beginning to be auto-extending. Specify the `autoextend` attribute for the last data file in the tablespace definition. Then InnoDB increases the size of that file automatically in 8MB increments when it runs out of space. The increment size can be changed by setting the value of the `innodb_autoextend_increment` [1704] system variable, which is measured in megabytes.

You can expand the system tablespace by a defined amount by adding another data file:

- Shut down the MySQL server.
- If the previous last data file is defined with the keyword `autoextend`, change its definition to use a fixed size, based on how large it has actually grown. Check the size of the data file, round it down to the closest multiple of 1024 × 1024 bytes (= 1MB), and specify this rounded size explicitly in `innodb_data_file_path` [1714].
- Add a new data file to the end of `innodb_data_file_path` [1714], optionally making that file auto-extending. Only the last data file in the `innodb_data_file_path` [1714] can be specified as auto-extending.
- Start the MySQL server again.

For example, this tablespace has just one auto-extending data file `ibdata1`:

```
innodb_data_home_dir =
innodb_data_file_path = /ibdata/ibdata1:10M:autoextend
```

Suppose that this data file, over time, has grown to 988MB. Here is the configuration line after modifying the original data file to use a fixed size and adding a new auto-extending data file:

```
innodb_data_home_dir =
innodb_data_file_path = /ibdata/ibdata1:988M;/disk2/ibdata2:50M:autoextend
```

When you add a new data file to the system tablespace configuration, make sure that the filename does not refer to an existing file. InnoDB creates and initializes the file when you restart the server.

Currently, you cannot remove a data file from the system tablespace. To decrease the system tablespace size, use this procedure:

1. Use `mysqldump` to dump all your InnoDB tables.
2. Stop the server.
3. Remove all the existing tablespace files, including the `ibdata` and `ib_log` files. If you want to keep a backup copy of the information, then copy all the `ib*` files to another location before the removing the files in your MySQL installation.
4. Remove any `.frm` files for InnoDB tables.
5. Configure a new tablespace.
6. Restart the server.
7. Import the dump files.

If you want to change the number or the size of your InnoDB log files, use the following instructions. The procedure to use depends on the value of `innodb_fast_shutdown` [1716], which determines whether or not to bring the system tablespace fully up-to-date before a shutdown operation:

- If `innodb_fast_shutdown` [1716] is not set to 2: Stop the MySQL server and make sure that it shuts down without errors, to ensure that there is no information for outstanding transactions in the redo log. Copy the old redo log files to a safe place, in case something went wrong during the shutdown and you need them to recover the tablespace. Delete the old log files from the log file directory, edit `my.cnf` to change the log file configuration, and start the MySQL server again. `mysqld` sees that no InnoDB log files exist at startup and creates new ones.
- If `innodb_fast_shutdown` [1716] is set to 2: Set `innodb_fast_shutdown` [1716] to 1:

```
mysql> SET GLOBAL innodb_fast_shutdown = 1;
```

Then follow the instructions in the previous item.

### 14.2.2.3. Using Raw Disk Partitions for the Shared Tablespace

You can use raw disk partitions as data files in the InnoDB [system tablespace](#). This technique enables nonbuffered I/O on Windows and on some Linux and Unix systems without file system overhead. Perform tests with and without raw partitions to verify whether this change actually improves performance on your system.

When you create a new data file, put the keyword `newraw` immediately after the data file size in `innodb_data_file_path` [1714]. The partition must be at least as large as the size that you specify. Note that 1MB in InnoDB is 1024 x 1024 bytes, whereas 1MB in disk specifications usually means 1,000,000 bytes.

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=/dev/hdd1:3Gnewraw:/dev/hdd2:2Gnewraw
```

The next time you start the server, InnoDB notices the `newraw` keyword and initializes the new partition. However, do not create or change any InnoDB tables yet. Otherwise, when you next restart the server, InnoDB reinitializes the partition and your changes are lost. (As a safety measure InnoDB prevents users from modifying data when any partition with `newraw` is specified.)

After InnoDB has initialized the new partition, stop the server, change `newraw` in the data file specification to `raw`:

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=/dev/hdd1:3Graw:/dev/hdd2:2Graw
```

Then restart the server and InnoDB permits changes to be made.

On Windows, you can allocate a disk partition as a data file like this:

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=//./D:10Gnewraw
```

The `//./` corresponds to the Windows syntax of `\\.\` for accessing physical drives.

When you use a raw disk partition, ensure that the user ID that runs the MySQL server has read and write privileges for that partition. For example, if you run the server as the `mysql` user, the partition must be readable and writeable by `mysql`. If you run the server with the `--memlock` [431] option, the server must be run as `root`, so the partition must be readable and writeable by `root`.

### 14.2.2.4. Backing Up and Recovering an InnoDB Database

The key to safe database management is making regular backups. Depending on your data volume, number of MySQL servers, and database workload, you can use these techniques, alone or in combination: [hot backup](#) with MySQL Enterprise Backup; [cold backup](#) by copying files while the

MySQL server is shut down; [physical backup](#) for fast operation (especially for restore); [logical backup](#) with `mysqldump` for smaller data volumes or to record the structure of schema objects.

## Hot Backups

The `mysqlbackup` command, part of the MySQL Enterprise Backup component, lets you back up a running MySQL instance, including [InnoDB](#) and [MyISAM](#) tables, with minimal disruption to operations while producing a consistent snapshot of the database. When `mysqlbackup` is copying [InnoDB](#) tables, reads and writes to both [InnoDB](#) and [MyISAM](#) tables can continue. During the copying of [MyISAM](#) tables, reads (but not writes) to those tables are permitted. MySQL Enterprise Backup can also create compressed backup files, and back up subsets of tables and databases. In conjunction with MySQL's binary log, users can perform point-in-time recovery. MySQL Enterprise Backup is part of the MySQL Enterprise subscription. For more details, see [Section 23.2, "MySQL Enterprise Backup"](#).

## Cold Backups

If you can shut down your MySQL server, you can make a binary backup that consists of all files used by [InnoDB](#) to manage its tables. Use the following procedure:

1. Do a [slow shutdown](#) of the MySQL server and make sure that it stops without errors.
2. Copy all [InnoDB](#) data files (`ibdata` files and `.ibd` files) into a safe place.
3. Copy all the `.frm` files for [InnoDB](#) tables to a safe place.
4. Copy all [InnoDB](#) log files (`ib_logfile` files) to a safe place.
5. Copy your `my.cnf` configuration file or files to a safe place.

## Alternative Backup Types

In addition to making binary backups as just described, regularly make dumps of your tables with `mysqldump`. A binary file might be corrupted without you noticing it. Dumped tables are stored into text files that are human-readable, so spotting table corruption becomes easier. Also, because the format is simpler, the chance for serious data corruption is smaller. `mysqldump` also has a `--single-transaction` [\[295\]](#) option for making a consistent snapshot without locking out other clients. See [Section 7.3.1, "Establishing a Backup Policy"](#).

Replication works with [InnoDB](#) tables, so you can use MySQL replication capabilities to keep a copy of your database at database sites requiring high availability.

## Performing Recovery

To recover your [InnoDB](#) database to the present from the time at which the binary backup was made, you must run your MySQL server with binary logging turned on, even before taking the backup. To achieve point-in-time recovery after restoring a backup, you can apply changes from the binary log that occurred after the backup was made. See [Section 7.5, "Point-in-Time \(Incremental\) Recovery Using the Binary Log"](#).

To recover from a crash of your MySQL server, the only requirement is to restart it. [InnoDB](#) automatically checks the logs and performs a roll-forward of the database to the present. [InnoDB](#) automatically rolls back uncommitted transactions that were present at the time of the crash. During recovery, `mysqld` displays output something like this:

```
InnoDB: Database was not shut down normally.
InnoDB: Starting recovery from log files...
InnoDB: Starting log scan based on checkpoint at
InnoDB: log sequence number 0 13674004
InnoDB: Doing recovery: scanned up to log sequence number 0 13739520
InnoDB: Doing recovery: scanned up to log sequence number 0 13805056
InnoDB: Doing recovery: scanned up to log sequence number 0 13870592
InnoDB: Doing recovery: scanned up to log sequence number 0 13936128
...
InnoDB: Doing recovery: scanned up to log sequence number 0 20555264
```

```
InnoDB: Doing recovery: scanned up to log sequence number 0 20620800
InnoDB: Doing recovery: scanned up to log sequence number 0 20664692
InnoDB: 1 uncommitted transaction(s) which must be rolled back
InnoDB: Starting rollback of uncommitted transactions
InnoDB: Rolling back trx no 16745
InnoDB: Rolling back of trx no 16745 completed
InnoDB: Rollback of uncommitted transactions completed
InnoDB: Starting an apply batch of log records to the database...
InnoDB: Apply batch completed
InnoDB: Started
mysqld: ready for connections
```

If your database becomes corrupted or disk failure occurs, you must perform the recovery using a backup. In the case of corruption, first find a backup that is not corrupted. After restoring the base backup, do a point-in-time recovery from the binary log files using `mysqlbinlog` and `mysql` to restore the changes that occurred after the backup was made.

In some cases of database corruption, it is enough just to dump, drop, and re-create one or a few corrupt tables. You can use the `CHECK TABLE` SQL statement to check whether a table is corrupt, although `CHECK TABLE` naturally cannot detect every possible kind of corruption. You can use the Tablespace Monitor to check the integrity of the file space management inside the tablespace files.

In some cases, apparent database page corruption is actually due to the operating system corrupting its own file cache, and the data on disk may be okay. It is best first to try restarting your computer. Doing so may eliminate errors that appeared to be database page corruption. If MySQL still has trouble starting because of InnoDB consistency problems, see [Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”](#) for steps to start the instance in a diagnostic mode where you can dump the data.

#### 14.2.2.5. Moving or Copying InnoDB Tables to Another Machine

This section explains various techniques for moving or copying some or all InnoDB tables to a different server. For example, you might move an entire MySQL instance to a larger, faster server; you might clone an entire MySQL instance to a new replication slave server; you might copy individual tables to another server to development and test an application, or to a data warehouse server to produce reports.

#### Using Lowercase Table Names

On Windows, InnoDB always stores database and table names internally in lowercase. To move databases in a binary format from Unix to Windows or from Windows to Unix, create all databases and tables using lowercase names. A convenient way to accomplish this is to add the following line to the `[mysqld]` section of your `my.cnf` or `my.ini` file before creating any databases or tables:

```
[mysqld]
lower_case_table_names=1
```

#### Using MySQL Enterprise Backup

The MySQL Enterprise Backup product lets you back up a running MySQL database, including InnoDB and MyISAM tables, with minimal disruption to operations while producing a consistent snapshot of the database. When MySQL Enterprise Backup is copying InnoDB tables, reads and writes to both InnoDB and MyISAM tables can continue. During the copying of MyISAM and other non-InnoDB tables, reads (but not writes) to those tables are permitted. In addition, MySQL Enterprise Backup can create compressed backup files, and back up subsets of InnoDB tables. In conjunction with the MySQL binary log, you can perform point-in-time recovery. MySQL Enterprise Backup is included as part of the MySQL Enterprise subscription.

For more details about MySQL Enterprise Backup, see [MySQL Enterprise Backup User's Guide \(Version 3.8.1\)](#).

#### Copying Data Files

Like MyISAM data files, InnoDB data and log files are binary-compatible on all platforms having the same floating-point number format. You can move an InnoDB database simply by copying all the

relevant files listed in [Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”](#). If the floating-point formats differ but you have not used `FLOAT` or `DOUBLE` data types in your tables, then the procedure is the same: simply copy the relevant files.

## Portability Considerations for `.ibd` Files

When you move or copy `.ibd` files, the database directory name must be the same on the source and destination systems. The table definition stored in the InnoDB shared tablespace includes the database name. The transaction IDs and log sequence numbers stored in the tablespace files also differ between databases.

To move an `.ibd` file and the associated table from one database to another, use a `RENAME TABLE` statement:

```
RENAME TABLE db1.tbl_name TO db2.tbl_name;
```

If you have a “clean” backup of an `.ibd` file, you can restore it to the MySQL installation from which it originated as follows:

1. The table must not have been dropped or truncated since you copied the `.ibd` file, because doing so changes the table ID stored inside the tablespace.
2. Issue this `ALTER TABLE` statement to delete the current `.ibd` file:

```
ALTER TABLE tbl_name DISCARD TABLESPACE;
```

3. Copy the backup `.ibd` file to the proper database directory.
4. Issue this `ALTER TABLE` statement to tell InnoDB to use the new `.ibd` file for the table:

```
ALTER TABLE tbl_name IMPORT TABLESPACE;
```

In this context, a “clean” `.ibd` file backup is one for which the following requirements are satisfied:

- There are no uncommitted modifications by transactions in the `.ibd` file.
- There are no unmerged insert buffer entries in the `.ibd` file.
- Purge has removed all delete-marked index records from the `.ibd` file.
- `mysqld` has flushed all modified pages of the `.ibd` file from the buffer pool to the file.

You can make a clean backup `.ibd` file using the following method:

1. Stop all activity from the `mysqld` server and commit all transactions.
2. Wait until `SHOW ENGINE INNODB STATUS` shows that there are no active transactions in the database, and the main thread status of InnoDB is `Waiting for server activity`. Then you can make a copy of the `.ibd` file.

Another method for making a clean copy of an `.ibd` file is to use the MySQL Enterprise Backup product:

1. Use MySQL Enterprise Backup to back up the InnoDB installation.
2. Start a second `mysqld` server on the backup and let it clean up the `.ibd` files in the backup.

## Export and Import

If you use `mysqldump` to dump your tables on one machine and then import the dump files on the other machine, it does not matter whether the formats differ or your tables contain floating-point data.

One way to increase performance is to switch off `autocommit` mode when importing data, assuming that the tablespace has enough space for the big rollback segment that the import transactions generate. Do the commit only after importing a whole table or a segment of a table.



### 14.2.2.6. InnoDB and MySQL Replication

MySQL replication works for InnoDB tables as it does for MyISAM tables. It is also possible to use replication in a way where the storage engine on the slave is not the same as the original storage engine on the master. For example, you can replicate modifications to an InnoDB table on the master to a MyISAM table on the slave.

To set up a new slave for a master, make a copy of the InnoDB tablespace and the log files, as well as the .frm files of the InnoDB tables, and move the copies to the slave. If the `innodb_file_per_table` [1718] option is enabled, copy the .ibd files as well. For the proper procedure to do this, see Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”.

To make a new slave without taking down the master or an existing slave, use the MySQL Enterprise Backup product. If you can shut down the master or an existing slave, take a cold backup of the InnoDB tablespaces and log files and use that to set up a slave.

Transactions that fail on the master do not affect replication at all. MySQL replication is based on the binary log where MySQL writes SQL statements that modify data. A transaction that fails (for example, because of a foreign key violation, or because it is rolled back) is not written to the binary log, so it is not sent to slaves. See Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”.

**Replication and CASCADE.** Cascading actions for InnoDB tables on the master are replicated on the slave *only* if the tables sharing the foreign key relation use InnoDB on both the master and slave. This is true whether you are using statement-based or row-based replication. Suppose that you have started replication, and then create two tables on the master using the following CREATE TABLE statements:

```
CREATE TABLE fc1 (
  i INT PRIMARY KEY,
  j INT
) ENGINE = InnoDB;

CREATE TABLE fc2 (
  m INT PRIMARY KEY,
  n INT,
  FOREIGN KEY ni (n) REFERENCES fc1 (i)
  ON DELETE CASCADE
) ENGINE = InnoDB;
```

Suppose that the slave does not have InnoDB support enabled. If this is the case, then the tables on the slave are created, but they use the MyISAM storage engine, and the FOREIGN KEY option is ignored. Now we insert some rows into the tables on the master:

```
master> INSERT INTO fc1 VALUES (1, 1), (2, 2);
Query OK, 2 rows affected (0.09 sec)
Records: 2 Duplicates: 0 Warnings: 0

master> INSERT INTO fc2 VALUES (1, 1), (2, 2), (3, 1);
Query OK, 3 rows affected (0.19 sec)
Records: 3 Duplicates: 0 Warnings: 0
```

At this point, on both the master and the slave, table fc1 contains 2 rows, and table fc2 contains 3 rows, as shown here:

```
master> SELECT * FROM fc1;
+----+-----+
| i | j |
+----+-----+
| 1 | 1 |
| 2 | 2 |
+----+-----+
2 rows in set (0.00 sec)

master> SELECT * FROM fc2;
+----+-----+
| m | n |
+----+-----+
```



```

| 1 | 1 |
| 2 | 2 |
| 3 | 1 |
+----+-----+
3 rows in set (0.00 sec)

slave> SELECT * FROM fc1;
+----+-----+
| i | j |
+----+-----+
| 1 | 1 |
| 2 | 2 |
+----+-----+
2 rows in set (0.00 sec)

slave> SELECT * FROM fc2;
+----+-----+
| m | n |
+----+-----+
| 1 | 1 |
| 2 | 2 |
| 3 | 1 |
+----+-----+
3 rows in set (0.00 sec)

```

Now suppose that you perform the following `DELETE` statement on the master:

```

master> DELETE FROM fc1 WHERE i=1;
Query OK, 1 row affected (0.09 sec)

```

Due to the cascade, table `fc2` on the master now contains only 1 row:

```

master> SELECT * FROM fc2;
+----+-----+
| m | n |
+----+-----+
| 2 | 2 |
+----+-----+
1 row in set (0.00 sec)

```

However, the cascade does not propagate on the slave because on the slave the `DELETE` for `fc1` deletes no rows from `fc2`. The slave's copy of `fc2` still contains all of the rows that were originally inserted:

```

slave> SELECT * FROM fc2;
+----+-----+
| m | n |
+----+-----+
| 1 | 1 |
| 3 | 1 |
| 2 | 2 |
+----+-----+
3 rows in set (0.00 sec)

```

This difference is due to the fact that the cascading deletes are handled internally by the `InnoDB` storage engine, which means that none of the changes are logged.

### 14.2.2.7. Checking InnoDB Availability

To determine whether your server supports `InnoDB`, use the `SHOW ENGINES` statement. (Now that `InnoDB` is the default MySQL storage engine, only very specialized environments might not support it.)

### 14.2.2.8. Turning Off InnoDB

Oracle recommends `InnoDB` as the preferred storage engine for typical database applications, from single-user wikis and blogs running on a local system, to high-end applications pushing the limits of performance. In MySQL 5.6, `InnoDB` is the default storage engine for new tables.

If you do not want to use `InnoDB` tables:

- Start the server with the `--innodb=OFF` [1698] or `--skip-innodb` [1698] option to disable the InnoDB storage engine.
- Because the default storage engine is InnoDB, the server will not start unless you also use `--default-storage-engine` [474] and `--default-tmp-storage-engine` [474] to set the default to some other engine for both permanent and `TEMPORARY` tables.
- To prevent the server from crashing when the InnoDB-related `information_schema` tables are queried, also disable the plugins associated with those tables. Specify in the `[mysqld]` section of the MySQL configuration file:

```
loose-innodb-trx=0
loose-innodb-locks=0
loose-innodb-lock-waits=0
loose-innodb-cmp=0
loose-innodb-cmp-per-index=0
loose-innodb-cmp-per-index-reset=0
loose-innodb-cmp-reset=0
loose-innodb-cmpmem=0
loose-innodb-cmpmem-reset=0
loose-innodb-buffer-page=0
loose-innodb-buffer-page-lru=0
loose-innodb-buffer-pool-stats=0
loose-innodb-metrics=0
loose-innodb-ft-default-stopword=0
loose-innodb-ft-inserted=0
loose-innodb-ft-deleted=0
loose-innodb-ft-being-deleted=0
loose-innodb-ft-config=0
loose-innodb-ft-index-cache=0
loose-innodb-ft-index-table=0
loose-innodb-sys-tables=0
loose-innodb-sys-tablestats=0
loose-innodb-sys-indexes=0
loose-innodb-sys-columns=0
loose-innodb-sys-fields=0
loose-innodb-sys-foreign=0
loose-innodb-sys-foreign-cols=0
```

### 14.2.3. InnoDB Concepts and Architecture

The information in this section provides background to help you get the most performance and functionality from using InnoDB tables. It is intended for:

- Anyone switching to MySQL from another database system, to explain what things might seem familiar and which might be all-new.
- Anyone moving from MyISAM tables to InnoDB, now that InnoDB is the default MySQL storage engine.
- Anyone considering their application architecture or software stack, to understand the design considerations, performance characteristics, and scalability of InnoDB tables at a detailed level.

In this section, you will learn:

- How InnoDB implements transactions, and how the inner workings of transactions compare with other database systems you might be familiar with.
- How InnoDB implements row-level locking to allow queries and DML statements to read and write the same table simultaneously.
- How multi-version concurrency control (MVCC) keeps transactions from viewing or modifying each others' data before the appropriate time.
- The physical layout of InnoDB-related objects on disk, such as tables, indexes, tablespaces, undo logs, and the redo log.

### 14.2.3.1. The InnoDB Transaction Model and Locking

To implement a large-scale, busy, or highly reliable database application, to port substantial code from a different database system, or to push MySQL performance to the limits of the laws of physics, you must understand the notions of [transactions](#) and [locking](#) as they relate to the InnoDB storage engine.

In the InnoDB transaction model, the goal is to combine the best properties of a multi-versioning database with traditional two-phase locking. InnoDB does locking on the row level and runs queries as nonlocking consistent reads by default, in the style of Oracle. The lock information in InnoDB is stored so space-efficiently that lock escalation is not needed: Typically, several users are permitted to lock every row in InnoDB tables, or any random subset of the rows, without causing InnoDB memory exhaustion.

In InnoDB, all user activity occurs inside a transaction. If autocommit mode is enabled, each SQL statement forms a single transaction on its own. By default, MySQL starts the session for each new connection with autocommit enabled, so MySQL does a commit after each SQL statement if that statement did not return an error. If a statement returns an error, the commit or rollback behavior depends on the error. See [Section 14.2.3.14, “InnoDB Error Handling”](#).

A session that has autocommit enabled can perform a multiple-statement transaction by starting it with an explicit `START TRANSACTION` or `BEGIN` statement and ending it with a `COMMIT` or `ROLLBACK` statement. See [Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#).

If autocommit mode is disabled within a session with `SET autocommit = 0`, the session always has a transaction open. A `COMMIT` or `ROLLBACK` statement ends the current transaction and a new one starts.

A `COMMIT` means that the changes made in the current transaction are made permanent and become visible to other sessions. A `ROLLBACK` statement, on the other hand, cancels all modifications made by the current transaction. Both `COMMIT` and `ROLLBACK` release all InnoDB locks that were set during the current transaction.

In terms of the SQL:1992 transaction [isolation levels](#), the default InnoDB level is `REPEATABLE READ` [1469]. InnoDB offers all four transaction isolation levels described by the SQL standard: `READ UNCOMMITTED` [1470], `READ COMMITTED` [1469], `REPEATABLE READ` [1469], and `SERIALIZABLE` [1470].

A user can change the isolation level for a single session or for all subsequent connections with the `SET TRANSACTION` statement. To set the server's default isolation level for all connections, use the `--transaction-isolation` [445] option on the command line or in an option file. For detailed information about isolation levels and level-setting syntax, see [Section 13.3.6, “SET TRANSACTION Syntax”](#).

In [row-level locking](#), InnoDB normally uses next-key locking. That means that besides index records, InnoDB can also lock the [gap](#) preceding an index record to block insertions by other sessions where the indexed values would be inserted in that gap within the tree data structure. A next-key lock refers to a lock that locks an index record and the gap before it. A gap lock refers to a lock that locks only the gap before some index record.

For more information about row-level locking, and the circumstances under which gap locking is disabled, see [Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#).

### 14.2.3.2. InnoDB Lock Modes

InnoDB implements standard [row-level locking](#) where there are two types of locks:

- A [shared \(S\) lock](#) permits the transaction that holds the lock to read a row.
- An [exclusive \(X\) lock](#) permits the transaction that holds the lock to update or delete a row.

If transaction `T1` holds a shared (S) lock on row `r`, then requests from some distinct transaction `T2` for a lock on row `r` are handled as follows:

- A request by **T2** for an *S* lock can be granted immediately. As a result, both **T1** and **T2** hold an *S* lock on *r*.
- A request by **T2** for an *X* lock cannot be granted immediately.

If a transaction **T1** holds an exclusive (*X*) lock on row *r*, a request from some distinct transaction **T2** for a lock of either type on *r* cannot be granted immediately. Instead, transaction **T2** has to wait for transaction **T1** to release its lock on row *r*.

Additionally, InnoDB supports *multiple granularity locking* which permits coexistence of record locks and locks on entire tables. To make locking at multiple granularity levels practical, additional types of locks called *intention locks* are used. Intention locks are table locks in InnoDB that indicate which type of lock (shared or exclusive) a transaction will require later for a row in that table. There are two types of intention locks used in InnoDB (assume that transaction **T** has requested a lock of the indicated type on table *t*):

- **Intention shared** (*IS*): Transaction **T** intends to set *S* locks on individual rows in table *t*.
- **Intention exclusive** (*IX*): Transaction **T** intends to set *X* locks on those rows.

For example, `SELECT ... LOCK IN SHARE MODE` sets an *IS* lock and `SELECT ... FOR UPDATE` sets an *IX* lock.

The intention locking protocol is as follows:

- Before a transaction can acquire an *S* lock on a row in table *t*, it must first acquire an *IS* or stronger lock on *t*.
- Before a transaction can acquire an *X* lock on a row, it must first acquire an *IX* lock on *t*.

These rules can be conveniently summarized by means of the following *lock type compatibility matrix*.

	<i>X</i>	<i>IX</i>	<i>S</i>	<i>IS</i>
<i>X</i>	Conflict	Conflict	Conflict	Conflict
<i>IX</i>	Conflict	Compatible	Conflict	Compatible
<i>S</i>	Conflict	Conflict	Compatible	Compatible
<i>IS</i>	Conflict	Compatible	Compatible	Compatible

A lock is granted to a requesting transaction if it is compatible with existing locks, but not if it conflicts with existing locks. A transaction waits until the conflicting existing lock is released. If a lock request conflicts with an existing lock and cannot be granted because it would cause *deadlock*, an error occurs.

Thus, intention locks do not block anything except full table requests (for example, `LOCK TABLES ... WRITE`). The main purpose of *IX* and *IS* locks is to show that someone is locking a row, or going to lock a row in the table.

The following example illustrates how an error can occur when a lock request would cause a deadlock. The example involves two clients, A and B.

First, client A creates a table containing one row, and then begins a transaction. Within the transaction, A obtains an *S* lock on the row by selecting it in share mode:

```
mysql> CREATE TABLE t (i INT) ENGINE = InnoDB;
Query OK, 0 rows affected (1.07 sec)

mysql> INSERT INTO t (i) VALUES(1);
Query OK, 1 row affected (0.09 sec)

mysql> START TRANSACTION;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT * FROM t WHERE i = 1 LOCK IN SHARE MODE;
```

```
+-----+
| i     |
+-----+
| 1     |
+-----+
1 row in set (0.10 sec)
```

Next, client B begins a transaction and attempts to delete the row from the table:

```
mysql> START TRANSACTION;
Query OK, 0 rows affected (0.00 sec)

mysql> DELETE FROM t WHERE i = 1;
```

The delete operation requires an *X* lock. The lock cannot be granted because it is incompatible with the *S* lock that client A holds, so the request goes on the queue of lock requests for the row and client B blocks.

Finally, client A also attempts to delete the row from the table:

```
mysql> DELETE FROM t WHERE i = 1;
ERROR 1213 (40001): Deadlock found when trying to get lock;
try restarting transaction
```

Deadlock occurs here because client A needs an *X* lock to delete the row. However, that lock request cannot be granted because client B already has a request for an *X* lock and is waiting for client A to release its *S* lock. Nor can the *S* lock held by A be upgraded to an *X* lock because of the prior request by B for an *X* lock. As a result, InnoDB generates an error for one of the clients and releases its locks. The client returns this error:

```
ERROR 1213 (40001): Deadlock found when trying to get lock;
try restarting transaction
```

At that point, the lock request for the other client can be granted and it deletes the row from the table.

### 14.2.3.3. Consistent Nonlocking Reads

A **consistent read** means that InnoDB uses multi-versioning to present to a query a snapshot of the database at a point in time. The query sees the changes made by transactions that committed before that point of time, and no changes made by later or uncommitted transactions. The exception to this rule is that the query sees the changes made by earlier statements within the same transaction. This exception causes the following anomaly: If you update some rows in a table, a **SELECT** sees the latest version of the updated rows, but it might also see older versions of any rows. If other sessions simultaneously update the same table, the anomaly means that you might see the table in a state that never existed in the database.

If the transaction **isolation level** is **REPEATABLE READ** [1469] (the default level), all consistent reads within the same transaction read the snapshot established by the first such read in that transaction. You can get a fresher snapshot for your queries by committing the current transaction and after that issuing new queries.

With **READ COMMITTED** [1469] isolation level, each consistent read within a transaction sets and reads its own fresh snapshot.

Consistent read is the default mode in which InnoDB processes **SELECT** statements in **READ COMMITTED** [1469] and **REPEATABLE READ** [1469] isolation levels. A consistent read does not set any locks on the tables it accesses, and therefore other sessions are free to modify those tables at the same time a consistent read is being performed on the table.

Suppose that you are running in the default **REPEATABLE READ** [1469] isolation level. When you issue a consistent read (that is, an ordinary **SELECT** statement), InnoDB gives your transaction a timepoint according to which your query sees the database. If another transaction deletes a row and commits after your timepoint was assigned, you do not see the row as having been deleted. Inserts and updates are treated similarly.



### Note

The snapshot of the database state applies to [SELECT](#) statements within a transaction, not necessarily to [DML](#) statements. If you insert or modify some rows and then commit that transaction, a [DELETE](#) or [UPDATE](#) statement issued from another concurrent [REPEATABLE READ](#) transaction could affect those just-committed rows, even though the session could not query them. If a transaction does update or delete rows committed by a different transaction, those changes do become visible to the current transaction. For example, you might encounter a situation like the following:

```
SELECT COUNT(c1) FROM t1 WHERE c1 = 'xyz'; -- Returns 0: no rows match.
DELETE FROM t1 WHERE c1 = 'xyz'; -- Deletes several rows recently committed by other tran
SELECT COUNT(c2) FROM t1 WHERE c2 = 'abc'; -- Returns 0: no rows match.
UPDATE t1 SET c2 = 'cba' WHERE c2 = 'abc'; -- Affects 10 rows: another txn just committed
SELECT COUNT(c2) FROM t1 WHERE c2 = 'cba'; -- Returns 10: this txn can now see the rows i
```

You can advance your timepoint by committing your transaction and then doing another [SELECT](#) or [START TRANSACTION WITH CONSISTENT SNAPSHOT](#).

This is called *multi-versioned concurrency control*.

In the following example, session A sees the row inserted by B only when B has committed the insert and A has committed as well, so that the timepoint is advanced past the commit of B.

	Session A	Session B
time	SET autocommit=0;	SET autocommit=0;
	SELECT * FROM t;	
	empty set	
		INSERT INTO t VALUES (1, 2);
v	SELECT * FROM t;	
	empty set	COMMIT;
	SELECT * FROM t;	
	empty set	
	COMMIT;	
	SELECT * FROM t;	
	-----	
	1   2	
	-----	
	1 row in set	

If you want to see the “freshest” state of the database, use either the [READ COMMITTED \[1469\]](#) isolation level or a [locking read](#):

```
SELECT * FROM t LOCK IN SHARE MODE;
```

With [READ COMMITTED \[1469\]](#) isolation level, each consistent read within a transaction sets and reads its own fresh snapshot. With [LOCK IN SHARE MODE](#), a locking read occurs instead: A [SELECT](#) blocks until the transaction containing the freshest rows ends (see [Section 14.2.3.4, “Locking Reads \(SELECT ... FOR UPDATE and SELECT ... LOCK IN SHARE MODE\)”](#)).

Consistent read does not work over certain DDL statements:

- Consistent read does not work over [DROP TABLE](#), because MySQL cannot use a table that has been dropped and [InnoDB](#) destroys the table.
- Consistent read does not work over [ALTER TABLE](#), because that statement makes a temporary copy of the original table and deletes the original table when the temporary copy is built. When you

reissue a consistent read within a transaction, rows in the new table are not visible because those rows did not exist when the transaction's snapshot was taken. In this case, the transaction returns an error as of MySQL 5.6.6: `ER_TABLE_DEF_CHANGED` [3687], "Table definition has changed, please retry transaction".

The type of read varies for selects in clauses like `INSERT INTO ... SELECT`, `UPDATE ... (SELECT)`, and `CREATE TABLE ... SELECT` that do not specify `FOR UPDATE` or `LOCK IN SHARE MODE`:

- By default, InnoDB uses stronger locks and the `SELECT` part acts like `READ COMMITTED` [1469], where each consistent read, even within the same transaction, sets and reads its own fresh snapshot.
- To use a consistent read in such cases, enable the `innodb_locks_unsafe_for_binlog` [1729] option and set the isolation level of the transaction to `READ UNCOMMITTED` [1470], `READ COMMITTED` [1469], or `REPEATABLE READ` [1469] (that is, anything other than `SERIALIZABLE` [1470]). In this case, no locks are set on rows read from the selected table.

#### 14.2.3.4. Locking Reads (`SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE`)

If you query data and then insert or update related data within the same transaction, the regular `SELECT` statement does not give enough protection. Other transactions can update or delete the same rows you just queried. InnoDB supports two types of [locking reads](#) that offer extra safety:

- `SELECT ... LOCK IN SHARE MODE` sets a shared mode lock on any rows that are read. Other sessions can read the rows, but cannot modify them until your transaction commits. If any of these rows were changed by another transaction that has not yet committed, your query waits until that transaction ends and then uses the latest values.
- For index records the search encounters, `SELECT ... FOR UPDATE` locks the rows and any associated index entries, the same as if you issued an `UPDATE` statement for those rows. Other transactions are blocked from updating those rows, from doing `SELECT ... LOCK IN SHARE MODE`, or from reading the data in certain transaction isolation levels. Consistent reads ignore any locks set on the records that exist in the read view. (Old versions of a record cannot be locked; they are reconstructed by applying [undo logs](#) on an in-memory copy of the record.)

These clauses are primarily useful when dealing with tree-structured or graph-structured data, either in a single table or split across multiple tables. You traverse edges or tree branches from one place to another, while reserving the right to come back and change any of these "pointer" values.

All locks set by `LOCK IN SHARE MODE` and `FOR UPDATE` queries are released when the transaction is committed or rolled back.



#### Note

Locking of rows for update using `SELECT FOR UPDATE` only applies when autocommit is disabled (either by beginning transaction with `START TRANSACTION` or by setting `autocommit` [464] to 0. If autocommit is enabled, the rows matching the specification are not locked.

#### Usage Examples

Suppose that you want to insert a new row into a table `child`, and make sure that the child row has a parent row in table `parent`. Your application code can ensure referential integrity throughout this sequence of operations.

First, use a consistent read to query the table `PARENT` and verify that the parent row exists. Can you safely insert the child row to table `CHILD`? No, because some other session could delete the parent row in the moment between your `SELECT` and your `INSERT`, without you being aware of it.



To avoid this potential issue, perform the `SELECT` using `LOCK IN SHARE MODE`:

```
SELECT * FROM parent WHERE NAME = 'Jones' LOCK IN SHARE MODE;
```

After the `LOCK IN SHARE MODE` query returns the parent 'Jones', you can safely add the child record to the `CHILD` table and commit the transaction. Any transaction that tries to read or write to the applicable row in the `PARENT` table waits until you are finished, that is, the data in all tables is in a consistent state.

For another example, consider an integer counter field in a table `CHILD_CODES`, used to assign a unique identifier to each child added to table `CHILD`. Do not use either consistent read or a shared mode read to read the present value of the counter, because two users of the database could see the same value for the counter, and a duplicate-key error occurs if two transactions attempt to add rows with the same identifier to the `CHILD` table.

Here, `LOCK IN SHARE MODE` is not a good solution because if two users read the counter at the same time, at least one of them ends up in deadlock when it attempts to update the counter.

Here are two ways to implement reading and incrementing the counter without interference from another transaction:

- First update the counter by incrementing it by 1, then read it and use the new value in the `CHILD` table. Any other transaction that tries to read the counter waits until your transaction commits. If another transaction is in the middle of this same sequence, your transaction waits until the other one commits.
- First perform a **locking read** of the counter using `FOR UPDATE`, and then increment the counter:

```
SELECT counter_field FROM child_codes FOR UPDATE;
UPDATE child_codes SET counter_field = counter_field + 1;
```

A `SELECT ... FOR UPDATE` reads the latest available data, setting exclusive locks on each row it reads. Thus, it sets the same locks a searched SQL `UPDATE` would set on the rows.

The preceding description is merely an example of how `SELECT ... FOR UPDATE` works. In MySQL, the specific task of generating a unique identifier actually can be accomplished using only a single access to the table:

```
UPDATE child_codes SET counter_field = LAST_INSERT_ID(counter_field + 1);
SELECT LAST_INSERT_ID();
```

The `SELECT` statement merely retrieves the identifier information (specific to the current connection). It does not access any table.

### 14.2.3.5. InnoDB Record, Gap, and Next-Key Locks

InnoDB has several types of record-level locks:

- Record lock: This is a lock on an index record.
- Gap lock: This is a lock on a gap between index records, or a lock on the gap before the first or after the last index record.
- Next-key lock: This is a combination of a record lock on the index record and a gap lock on the gap before the index record.

Record locks always lock index records, even if a table is defined with no indexes. For such cases, InnoDB creates a hidden clustered index and uses this index for record locking. See [Section 14.2.3.12.2, “Clustered and Secondary Indexes”](#).

By default, InnoDB operates in `REPEATABLE READ [1469]` transaction isolation level and with the `innodb_locks_unsafe_for_binlog [1729]` system variable disabled. In this case, InnoDB uses next-key locks for searches and index scans, which prevents phantom rows (see [Section 14.2.3.6, “Avoiding the Phantom Problem Using Next-Key Locking”](#)).

Next-key locking combines index-row locking with gap locking. InnoDB performs row-level locking in such a way that when it searches or scans a table index, it sets shared or exclusive locks on the index records it encounters. Thus, the row-level locks are actually index-record locks. In addition, a next-key lock on an index record also affects the “gap” before that index record. That is, a next-key lock is an index-record lock plus a gap lock on the gap preceding the index record. If one session has a shared or exclusive lock on record *R* in an index, another session cannot insert a new index record in the gap immediately before *R* in the index order.

Suppose that an index contains the values 10, 11, 13, and 20. The possible next-key locks for this index cover the following intervals, where ( or ) denote exclusion of the interval endpoint and [ or ] denote inclusion of the endpoint:

```
(negative infinity, 10]
(10, 11]
(11, 13]
(13, 20]
(20, positive infinity)
```

For the last interval, the next-key lock locks the gap above the largest value in the index and the “supremum” pseudo-record having a value higher than any value actually in the index. The supremum is not a real index record, so, in effect, this next-key lock locks only the gap following the largest index value.

The preceding example shows that a gap might span a single index value, multiple index values, or even be empty.

Gap locking is not needed for statements that lock rows using a unique index to search for a unique row. (This does not include the case that the search condition includes only some columns of a multiple-column unique index; in that case, gap locking does occur.) For example, if the *id* column has a unique index, the following statement uses only an index-record lock for the row having *id* value 100 and it does not matter whether other sessions insert rows in the preceding gap:

```
SELECT * FROM child WHERE id = 100;
```

If *id* is not indexed or has a nonunique index, the statement does lock the preceding gap.

A type of gap lock called an insertion intention gap lock is set by `INSERT` operations prior to row insertion. This lock signals the intent to insert in such a way that multiple transactions inserting into the same index gap need not wait for each other if they are not inserting at the same position within the gap. Suppose that there are index records with values of 4 and 7. Separate transactions that attempt to insert values of 5 and 6 each lock the gap between 4 and 7 with insert intention locks prior to obtaining the exclusive lock on the inserted row, but do not block each other because the rows are nonconflicting.

Gap locking can be disabled explicitly. This occurs if you change the transaction isolation level to `READ COMMITTED` [1469] or enable the `innodb_locks_unsafe_for_binlog` [1729] system variable (which is now deprecated). Under these circumstances, gap locking is disabled for searches and index scans and is used only for foreign-key constraint checking and duplicate-key checking.

There are also other effects of using the `READ COMMITTED` [1469] isolation level or enabling `innodb_locks_unsafe_for_binlog` [1729]: Record locks for nonmatching rows are released after MySQL has evaluated the `WHERE` condition. For `UPDATE` statements, InnoDB does a “semi-consistent” read, such that it returns the latest committed version to MySQL so that MySQL can determine whether the row matches the `WHERE` condition of the `UPDATE`.

#### 14.2.3.6. Avoiding the Phantom Problem Using Next-Key Locking

The so-called *phantom* problem occurs within a transaction when the same query produces different sets of rows at different times. For example, if a `SELECT` is executed twice, but returns a row the second time that was not returned the first time, the row is a “phantom” row.

Suppose that there is an index on the `id` column of the `child` table and that you want to read and lock all rows from the table having an identifier value larger than 100, with the intention of updating some column in the selected rows later:

```
SELECT * FROM child WHERE id > 100 FOR UPDATE;
```

The query scans the index starting from the first record where `id` is bigger than 100. Let the table contain rows having `id` values of 90 and 102. If the locks set on the index records in the scanned range do not lock out inserts made in the gaps (in this case, the gap between 90 and 102), another session can insert a new row into the table with an `id` of 101. If you were to execute the same `SELECT` within the same transaction, you would see a new row with an `id` of 101 (a “phantom”) in the result set returned by the query. If we regard a set of rows as a data item, the new phantom child would violate the isolation principle of transactions that a transaction should be able to run so that the data it has read does not change during the transaction.

To prevent phantoms, InnoDB uses an algorithm called *next-key locking* that combines index-row locking with gap locking. InnoDB performs row-level locking in such a way that when it searches or scans a table index, it sets shared or exclusive locks on the index records it encounters. Thus, the row-level locks are actually index-record locks. In addition, a next-key lock on an index record also affects the “gap” before that index record. That is, a next-key lock is an index-record lock plus a gap lock on the gap preceding the index record. If one session has a shared or exclusive lock on record `R` in an index, another session cannot insert a new index record in the gap immediately before `R` in the index order.

When InnoDB scans an index, it can also lock the gap after the last record in the index. Just that happens in the preceding example: To prevent any insert into the table where `id` would be bigger than 100, the locks set by InnoDB include a lock on the gap following `id` value 102.

You can use next-key locking to implement a uniqueness check in your application: If you read your data in share mode and do not see a duplicate for a row you are going to insert, then you can safely insert your row and know that the next-key lock set on the successor of your row during the read prevents anyone meanwhile inserting a duplicate for your row. Thus, the next-key locking enables you to “lock” the nonexistence of something in your table.

Gap locking can be disabled as discussed in [Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#). This may cause phantom problems because other sessions can insert new rows into the gaps when gap locking is disabled.

### 14.2.3.7. Locks Set by Different SQL Statements in InnoDB

A [locking read](#), an `UPDATE`, or a `DELETE` generally set record locks on every index record that is scanned in the processing of the SQL statement. It does not matter whether there are `WHERE` conditions in the statement that would exclude the row. InnoDB does not remember the exact `WHERE` condition, but only knows which index ranges were scanned. The locks are normally [next-key locks](#) that also block inserts into the “gap” immediately before the record. However, [gap locking](#) can be disabled explicitly, which causes next-key locking not to be used. For more information, see [Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#). The transaction isolation level also can affect which locks are set; see [Section 13.3.6, “SET TRANSACTION Syntax”](#).

If a secondary index is used in a search and index record locks to be set are exclusive, InnoDB also retrieves the corresponding clustered index records and sets locks on them.

Differences between shared and exclusive locks are described in [Section 14.2.3.2, “InnoDB Lock Modes”](#).

If you have no indexes suitable for your statement and MySQL must scan the entire table to process the statement, every row of the table becomes locked, which in turn blocks all inserts by other users to the table. It is important to create good indexes so that your queries do not unnecessarily scan many rows.

For `SELECT ... FOR UPDATE` or `SELECT ... LOCK IN SHARE MODE`, locks are acquired for scanned rows, and expected to be released for rows that do not qualify for inclusion in the result set

(for example, if they do not meet the criteria given in the [WHERE](#) clause). However, in some cases, rows might not be unlocked immediately because the relationship between a result row and its original source is lost during query execution. For example, in a [UNION](#), scanned (and locked) rows from a table might be inserted into a temporary table before evaluation whether they qualify for the result set. In this circumstance, the relationship of the rows in the temporary table to the rows in the original table is lost and the latter rows are not unlocked until the end of query execution.

InnoDB sets specific types of locks as follows.

- [SELECT ... FROM](#) is a consistent read, reading a snapshot of the database and setting no locks unless the transaction isolation level is set to [SERIALIZABLE \[1470\]](#). For [SERIALIZABLE \[1470\]](#) level, the search sets shared next-key locks on the index records it encounters.
- [SELECT ... FROM ... LOCK IN SHARE MODE](#) sets shared next-key locks on all index records the search encounters.
- For index records the search encounters, [SELECT ... FROM ... FOR UPDATE](#) blocks other sessions from doing [SELECT ... FROM ... LOCK IN SHARE MODE](#) or from reading in certain transaction isolation levels. Consistent reads will ignore any locks set on the records that exist in the read view.
- [UPDATE ... WHERE ...](#) sets an exclusive next-key lock on every record the search encounters.
- [DELETE FROM ... WHERE ...](#) sets an exclusive next-key lock on every record the search encounters.
- [INSERT](#) sets an exclusive lock on the inserted row. This lock is an index-record lock, not a next-key lock (that is, there is no gap lock) and does not prevent other sessions from inserting into the gap before the inserted row.

Prior to inserting the row, a type of gap lock called an insertion intention gap lock is set. This lock signals the intent to insert in such a way that multiple transactions inserting into the same index gap need not wait for each other if they are not inserting at the same position within the gap. Suppose that there are index records with values of 4 and 7. Separate transactions that attempt to insert values of 5 and 6 each lock the gap between 4 and 7 with insert intention locks prior to obtaining the exclusive lock on the inserted row, but do not block each other because the rows are nonconflicting.

If a duplicate-key error occurs, a shared lock on the duplicate index record is set. This use of a shared lock can result in deadlock should there be multiple sessions trying to insert the same row if another session already has an exclusive lock. This can occur if another session deletes the row. Suppose that an InnoDB table `t1` has the following structure:

```
CREATE TABLE t1 (i INT, PRIMARY KEY (i)) ENGINE = InnoDB;
```

Now suppose that three sessions perform the following operations in order:

Session 1:

```
START TRANSACTION;
INSERT INTO t1 VALUES(1);
```

Session 2:

```
START TRANSACTION;
INSERT INTO t1 VALUES(1);
```

Session 3:

```
START TRANSACTION;
INSERT INTO t1 VALUES(1);
```

Session 1:

```
ROLLBACK;
```

The first operation by session 1 acquires an exclusive lock for the row. The operations by sessions 2 and 3 both result in a duplicate-key error and they both request a shared lock for the row. When session 1 rolls back, it releases its exclusive lock on the row and the queued shared lock requests for sessions 2 and 3 are granted. At this point, sessions 2 and 3 deadlock: Neither can acquire an exclusive lock for the row because of the shared lock held by the other.

A similar situation occurs if the table already contains a row with key value 1 and three sessions perform the following operations in order:

Session 1:

```
START TRANSACTION;
DELETE FROM t1 WHERE i = 1;
```

Session 2:

```
START TRANSACTION;
INSERT INTO t1 VALUES(1);
```

Session 3:

```
START TRANSACTION;
INSERT INTO t1 VALUES(1);
```

Session 1:

```
COMMIT;
```

The first operation by session 1 acquires an exclusive lock for the row. The operations by sessions 2 and 3 both result in a duplicate-key error and they both request a shared lock for the row. When session 1 commits, it releases its exclusive lock on the row and the queued shared lock requests for sessions 2 and 3 are granted. At this point, sessions 2 and 3 deadlock: Neither can acquire an exclusive lock for the row because of the shared lock held by the other.

- `INSERT ... ON DUPLICATE KEY UPDATE` differs from a simple `INSERT` in that an exclusive next-key lock rather than a shared lock is placed on the row to be updated when a duplicate-key error occurs.
- `REPLACE` is done like an `INSERT` if there is no collision on a unique key. Otherwise, an exclusive next-key lock is placed on the row to be replaced.
- `INSERT INTO T SELECT ... FROM S WHERE ...` sets an exclusive index record without a gap lock on each row inserted into `T`. If the transaction isolation level is `READ COMMITTED [1469]` or `innodb_locks_unsafe_for_binlog [1729]` is enabled, and the transaction isolation level is not `SERIALIZABLE [1470]`, InnoDB does the search on `S` as a consistent read (no locks). Otherwise, InnoDB sets shared next-key locks on rows from `S`. InnoDB has to set locks in the latter case: In roll-forward recovery from a backup, every SQL statement must be executed in exactly the same way it was done originally.

`CREATE TABLE ... SELECT ...` performs the `SELECT` with shared next-key locks or as a consistent read, as for `INSERT ... SELECT`.

When a `SELECT` is used in the constructs `REPLACE INTO t SELECT ... FROM s WHERE ...` or `UPDATE t ... WHERE col IN (SELECT ... FROM s ...)`, InnoDB sets shared next-key locks on rows from table `s`.

- While initializing a previously specified `AUTO_INCREMENT` column on a table, InnoDB sets an exclusive lock on the end of the index associated with the `AUTO_INCREMENT` column. In accessing the auto-increment counter, InnoDB uses a specific `AUTO-INC` table lock mode where the lock lasts only to the end of the current SQL statement, not to the end of the entire transaction. Other sessions cannot insert into the table while the `AUTO-INC` table lock is held; see [Section 14.2.3.1, “The InnoDB Transaction Model and Locking”](#).

InnoDB fetches the value of a previously initialized `AUTO_INCREMENT` column without setting any locks.

- If a `FOREIGN KEY` constraint is defined on a table, any insert, update, or delete that requires the constraint condition to be checked sets shared record-level locks on the records that it looks at to check the constraint. InnoDB also sets these locks in the case where the constraint fails.
- `LOCK TABLES` sets table locks, but it is the higher MySQL layer above the InnoDB layer that sets these locks. InnoDB is aware of table locks if `innodb_table_locks = 1` (the default) and `autocommit = 0` [464], and the MySQL layer above InnoDB knows about row-level locks.

Otherwise, InnoDB's automatic deadlock detection cannot detect deadlocks where such table locks are involved. Also, because in this case the higher MySQL layer does not know about row-level locks, it is possible to get a table lock on a table where another session currently has row-level locks. However, this does not endanger transaction integrity, as discussed in [Section 14.2.3.9, “Deadlock Detection and Rollback”](#). See also [Section 14.2.7, “Limits on InnoDB Tables”](#).

#### 14.2.3.8. Implicit Transaction Commit and Rollback

By default, MySQL starts the session for each new connection with autocommit mode enabled, so MySQL does a commit after each SQL statement if that statement did not return an error. If a statement returns an error, the commit or rollback behavior depends on the error. See [Section 14.2.3.14, “InnoDB Error Handling”](#).

If a session that has autocommit disabled ends without explicitly committing the final transaction, MySQL rolls back that transaction.

Some statements implicitly end a transaction, as if you had done a `COMMIT` before executing the statement. For details, see [Section 13.3.3, “Statements That Cause an Implicit Commit”](#).

#### 14.2.3.9. Deadlock Detection and Rollback

InnoDB automatically detects transaction [deadlocks](#) and rolls back a transaction or transactions to break the deadlock. InnoDB tries to pick small transactions to roll back, where the size of a transaction is determined by the number of rows inserted, updated, or deleted.

InnoDB is aware of table locks if `innodb_table_locks = 1` (the default) and `autocommit = 0` [464], and the MySQL layer above it knows about row-level locks. Otherwise, InnoDB cannot detect deadlocks where a table lock set by a MySQL `LOCK TABLES` statement or a lock set by a storage engine other than InnoDB is involved. Resolve these situations by setting the value of the `innodb_lock_wait_timeout` [1729] system variable.

When InnoDB performs a complete rollback of a transaction, all locks set by the transaction are released. However, if just a single SQL statement is rolled back as a result of an error, some of the locks set by the statement may be preserved. This happens because InnoDB stores row locks in a format such that it cannot know afterward which lock was set by which statement.

If a `SELECT` calls a stored function in a transaction, and a statement within the function fails, that statement rolls back. Furthermore, if `ROLLBACK` is executed after that, the entire transaction rolls back.

For techniques to organize database operations to avoid deadlocks, see [Section 14.2.3.10, “How to Cope with Deadlocks”](#).

#### 14.2.3.10. How to Cope with Deadlocks

This section builds on the conceptual information about deadlocks in [Section 14.2.3.9, “Deadlock Detection and Rollback”](#). It explains how to organize database operations to minimize deadlocks and the subsequent error handling required in applications.

[Deadlocks](#) are a classic problem in transactional databases, but they are not dangerous unless they are so frequent that you cannot run certain transactions at all. Normally, you must write your



applications so that they are always prepared to re-issue a transaction if it gets rolled back because of a deadlock.

InnoDB uses automatic row-level locking. You can get deadlocks even in the case of transactions that just insert or delete a single row. That is because these operations are not really “atomic”; they automatically set locks on the (possibly several) index records of the row inserted or deleted.

You can cope with deadlocks and reduce the likelihood of their occurrence with the following techniques:

- At any time, issue the `SHOW ENGINE INNODB STATUS` command to determine the cause of the most recent deadlock. That can help you to tune your application to avoid deadlocks.
- If frequent deadlock warnings cause concern, collect more extensive debugging information by restarting the server with the `innodb_print_all_deadlocks [1741]` configuration option enabled. Information about each deadlock, not just the latest one, is recorded in the MySQL [error log](#). Remove this option and restart the server again once the debugging is finished.
- Always be prepared to re-issue a transaction if it fails due to deadlock. Deadlocks are not dangerous. Just try again.
- Commit your transactions immediately after making a set of related changes. Small transactions are less prone to collision. In particular, do not leave an interactive `mysql` session open for a long time with an uncommitted transaction.
- If you use [locking reads](#) (`SELECT ... FOR UPDATE` or `SELECT ... LOCK IN SHARE MODE`), try using a lower isolation level such as `READ COMMITTED [1469]`.
- When modifying multiple tables within a transaction, or different sets of rows in the same table, do those operations in a consistent order each time. Then transactions form well-defined queues and do not deadlock. For example, organize database operations into functions within your application, or call stored routines, rather than coding multiple similar sequences of `INSERT`, `UPDATE`, and `DELETE` statements in different places.
- Add well-chosen indexes to your tables. Then your queries need to scan fewer index records and consequently set fewer locks. Use `EXPLAIN SELECT` to determine which indexes the MySQL server regards as the most appropriate for your queries.
- Use less locking. If you can afford to permit a `SELECT` to return data from an old snapshot, do not add the clause `FOR UPDATE` or `LOCK IN SHARE MODE` to it. Using the `READ COMMITTED [1469]` isolation level is good here, because each consistent read within the same transaction reads from its own fresh snapshot.
- If nothing else helps, serialize your transactions with table-level locks. The correct way to use `LOCK TABLES` with transactional tables, such as InnoDB tables, is to begin a transaction with `SET autocommit = 0` (not `START TRANSACTION`) followed by `LOCK TABLES`, and to not call `UNLOCK TABLES` until you commit the transaction explicitly. For example, if you need to write to table `t1` and read from table `t2`, you can do this:

```
SET autocommit=0;
LOCK TABLES t1 WRITE, t2 READ, ...;... do something with tables t1 and t2 here ...
COMMIT;
UNLOCK TABLES;
```

Table-level locks prevent concurrent updates to the table, avoiding deadlocks at the expense of less responsiveness for a busy system.

- Another way to serialize transactions is to create an auxiliary “semaphore” table that contains just a single row. Have each transaction update that row before accessing other tables. In that way, all transactions happen in a serial fashion. Note that the InnoDB instant deadlock detection algorithm also works in this case, because the serializing lock is a row-level lock. With MySQL table-level locks, the timeout method must be used to resolve deadlocks.



### 14.2.3.11. InnoDB Multi-Versioning

InnoDB is a [multi-versioned storage engine](#): it keeps information about old versions of changed rows, to support transactional features such as concurrency and [rollback](#). This information is stored in the tablespace in a data structure called a [rollback segment](#) (after an analogous data structure in Oracle). InnoDB uses the information in the rollback segment to perform the undo operations needed in a transaction rollback. It also uses the information to build earlier versions of a row for a [consistent read](#).

#### Internal Details of Multi-Versioning

Internally, InnoDB adds three fields to each row stored in the database. A 6-byte `DB_TRX_ID` field indicates the transaction identifier for the last transaction that inserted or updated the row. Also, a deletion is treated internally as an update where a special bit in the row is set to mark it as deleted. Each row also contains a 7-byte `DB_ROLL_PTR` field called the roll pointer. The roll pointer points to an undo log record written to the rollback segment. If the row was updated, the undo log record contains the information necessary to rebuild the content of the row before it was updated. A 6-byte `DB_ROW_ID` field contains a row ID that increases monotonically as new rows are inserted. If InnoDB generates a clustered index automatically, the index contains row ID values. Otherwise, the `DB_ROW_ID` column does not appear in any index.

Undo logs in the rollback segment are divided into insert and update undo logs. Insert undo logs are needed only in transaction rollback and can be discarded as soon as the transaction commits. Update undo logs are used also in consistent reads, but they can be discarded only after there is no transaction present for which InnoDB has assigned a snapshot that in a consistent read could need the information in the update undo log to build an earlier version of a database row.

#### Guidelines for Managing Rollback Segments

Commit your transactions regularly, including those transactions that issue only consistent reads. Otherwise, InnoDB cannot discard data from the update undo logs, and the rollback segment may grow too big, filling up your tablespace.

The physical size of an undo log record in the rollback segment is typically smaller than the corresponding inserted or updated row. You can use this information to calculate the space needed for your rollback segment.

In the InnoDB multi-versioning scheme, a row is not physically removed from the database immediately when you delete it with an SQL statement. InnoDB only physically removes the corresponding row and its index records when it discards the update undo log record written for the deletion. This removal operation is called a [purge](#), and it is quite fast, usually taking the same order of time as the SQL statement that did the deletion.

If you insert and delete rows in smallish batches at about the same rate in the table, the purge thread can start to lag behind and the table can grow bigger and bigger because of all the “dead” rows, making everything disk-bound and very slow. In such a case, throttle new row operations, and allocate more resources to the purge thread by tuning the `innodb_max_purge_lag` [1735] system variable. See [Section 14.2.6, “InnoDB Startup Options and System Variables”](#) for more information.

### 14.2.3.12. InnoDB Table and Index Structures

This section describes how InnoDB tables, indexes, and their associated metadata is represented at the physical level. This information is primarily useful for performance tuning and troubleshooting.

#### 14.2.3.12.1. Role of the `.frm` File for InnoDB Tables

MySQL stores its data dictionary information for tables in [.frm files](#) in database directories. Unlike other MySQL storage engines, InnoDB also encodes information about the table in its own internal data dictionary inside the tablespace. When MySQL drops a table or a database, it deletes one or more `.frm` files as well as the corresponding entries inside the InnoDB data dictionary. You cannot move InnoDB tables between databases simply by moving the `.frm` files.

#### 14.2.3.12.2. Clustered and Secondary Indexes

Every InnoDB table has a special index called the [clustered index](#) where the data for the rows is stored. Typically, the clustered index is synonymous with the [primary key](#). To get the best performance from queries, inserts, and other database operations, you must understand how InnoDB uses the clustered index to optimize the most common lookup and DML operations for each table.

- When you define a [PRIMARY KEY](#) on your table, InnoDB uses it as the clustered index. Define a primary key for each table that you create. If there is no logical unique and non-null column or set of columns, add a new [auto-increment](#) column, whose values are filled in automatically.
- If you do not define a [PRIMARY KEY](#) for your table, MySQL locates the first [UNIQUE](#) index where all the key columns are [NOT NULL](#) and InnoDB uses it as the clustered index.
- If the table has no [PRIMARY KEY](#) or suitable [UNIQUE](#) index, InnoDB internally generates a hidden clustered index on a synthetic column containing row ID values. The rows are ordered by the ID that InnoDB assigns to the rows in such a table. The row ID is a 6-byte field that increases monotonically as new rows are inserted. Thus, the rows ordered by the row ID are physically in insertion order.

### How the Clustered Index Speeds Up Queries

Accessing a row through the clustered index is fast because the index search leads directly to the page with all the row data. If a table is large, the clustered index architecture often saves a disk I/O operation when compared to storage organizations that store row data using a different page from the index record. (For example, [MyISAM](#) uses one file for data rows and another for index records.)

### How Secondary Indexes Relate to the Clustered Index

All indexes other than the clustered index are known as [secondary indexes](#). In InnoDB, each record in a secondary index contains the primary key columns for the row, as well as the columns specified for the secondary index. InnoDB uses this primary key value to search for the row in the clustered index.

If the primary key is long, the secondary indexes use more space, so it is advantageous to have a short primary key.

For coding guidelines to take advantage of InnoDB clustered and secondary indexes, see [Section 8.3.2, “Using Primary Keys”](#) [Section 8.3, “Optimization and Indexes”](#) [Section 8.5, “Optimizing for InnoDB Tables”](#) [Section 8.3.2, “Using Primary Keys”](#).

#### 14.2.3.12.3. FULLTEXT Indexes

A special kind of index, the [FULLTEXT](#) index, helps InnoDB deal with queries and DML operations involving text-based columns and the words they contain. These indexes are physically represented as entire InnoDB tables, which are acted upon by SQL keywords such as the [FULLTEXT](#) clause of the [CREATE INDEX](#) statement, the [MATCH\(\) . . . AGAINST \[1229\]](#) syntax in a [SELECT](#) statement, and the [OPTIMIZE TABLE](#) statement. For usage information, see [Section 12.9, “Full-Text Search Functions”](#).

You can examine [FULLTEXT](#) indexes by querying tables in the [INFORMATION\\_SCHEMA](#) database. You can see basic index information for [FULLTEXT](#) indexes by querying [INNODB\\_SYS\\_INDEXES](#). Although InnoDB [FULLTEXT](#) indexes are represented by tables, which show up in [INNODB\\_SYS\\_TABLES](#) queries, the way to monitor the special text-processing aspects of a [FULLTEXT](#) index is to query the tables [INNODB\\_FT\\_CONFIG](#), [INNODB\\_FT\\_INDEX\\_TABLE](#), [INNODB\\_FT\\_INDEX\\_CACHE](#), [INNODB\\_FT\\_DEFAULT\\_STOPWORD](#), [INNODB\\_FT\\_DELETED](#), and [INNODB\\_FT\\_BEING\\_DELETED](#).

InnoDB [FULLTEXT](#) indexes are updated by the [OPTIMIZE TABLE](#) command, using a special mode controlled by the configuration options [innodb\\_ft\\_num\\_word\\_optimize \[1725\]](#) and [innodb\\_optimize\\_fulltext\\_only \[1740\]](#).

#### 14.2.3.12.4. Physical Structure of an InnoDB Index

All InnoDB indexes are [B-trees](#) where the index records are stored in the leaf pages of the tree. The default size of an index page is 16KB. When new records are inserted, InnoDB tries to leave 1/16 of the page free for future insertions and updates of the index records.

If index records are inserted in a sequential order (ascending or descending), the resulting index pages are about 15/16 full. If records are inserted in a random order, the pages are from 1/2 to 15/16 full. If the **fill factor** of an index page drops below 1/2, InnoDB tries to contract the index tree to free the page.



#### Note

You can specify the **page size** for all InnoDB tablespaces in a MySQL instance by setting the `innodb_page_size [1740]` configuration option before creating the instance. Once the page size for a MySQL instance is set, you cannot change it. Supported sizes are 16KB, 8KB, and 4KB, corresponding to the option values `16k`, `8k`, and `4k`.

A MySQL instance using a particular InnoDB page size cannot use data files or log files from an instance that uses a different page size.

### 14.2.3.12.5. Insert Buffering

Database applications often insert new rows in the ascending order of the primary key. In this case, due to the layout of the clustered index in the same order as the primary key, insertions into an InnoDB table do not require random reads from a disk.

On the other hand, secondary indexes are usually nonunique, and insertions into secondary indexes happen in a relatively random order. In the same way, deletes and updates can affect data pages that are not adjacent in secondary indexes. This would cause a lot of random disk I/O operations without a special mechanism used in InnoDB.

When an index record is inserted, marked for deletion, or deleted from a nonunique secondary index, InnoDB checks whether the secondary index page is in the **buffer pool**. If that is the case, InnoDB applies the change directly to the index page. If the index page is not found in the buffer pool, InnoDB records the change in a special structure known as the **insert buffer**. The insert buffer is kept small so that it fits entirely in the buffer pool, and changes can be applied very quickly. This process is known as **change buffering**. (Formerly, it applied only to inserts and was called insert buffering. The data structure is still called the insert buffer.)

#### Disk I/O for Flushing the Insert Buffer

Periodically, the insert buffer is merged into the secondary index trees in the database. Often, it is possible to merge several changes into the same page of the index tree, saving disk I/O operations. It has been measured that the insert buffer can speed up insertions into a table up to 15 times.

The insert buffer merging may continue to happen *after* the transaction has been committed. In fact, it may continue to happen after a server shutdown and restart (see [Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”](#)).

Insert buffer merging may take many hours when many secondary indexes must be updated and many rows have been inserted. During this time, disk I/O will be increased, which can cause significant slowdown on disk-bound queries. Another significant background I/O operation is the **purge thread** (see [Section 14.2.3.11, “InnoDB Multi-Versioning”](#)).

### 14.2.3.12.6. Adaptive Hash Indexes

The feature known as the **adaptive hash index** (AHI) lets InnoDB perform more like an in-memory database on systems with appropriate combinations of workload and ample memory for the **buffer pool**, without sacrificing any transactional features or reliability. This feature is enabled by the `innodb_adaptive_hash_index [1701]` option, or turned off by the `--skip-innodb_adaptive_hash_index` at server startup.

Based on the observed pattern of searches, MySQL builds a hash index using a prefix of the index key. The prefix of the key can be any length, and it may be that only some of the values in the B-tree appear in the hash index. Hash indexes are built on demand for those pages of the index that are often accessed.

If a table fits almost entirely in main memory, a hash index can speed up queries by enabling direct lookup of any element, turning the index value into a sort of pointer. InnoDB has a mechanism that monitors index searches. If InnoDB notices that queries could benefit from building a hash index, it does so automatically.

With some workloads, the speedup from hash index lookups greatly outweighs the extra work to monitor index lookups and maintain the hash index structure. Sometimes, the read/write lock that guards access to the adaptive hash index can become a source of contention under heavy workloads, such as multiple concurrent joins. Queries with LIKE operators and % wildcards also tend not to benefit from the AHI. For workloads where the adaptive hash index is not needed, turning it off reduces unnecessary performance overhead. Because it is difficult to predict in advance whether this feature is appropriate for a particular system, consider running benchmarks with it both enabled and disabled, using a realistic workload. The architectural changes in MySQL 5.6 and higher make more workloads suitable for disabling the adaptive hash index than in earlier releases, although it is still enabled by default.

The hash index is always built based on an existing B-tree index on the table. InnoDB can build a hash index on a prefix of any length of the key defined for the B-tree, depending on the pattern of searches that InnoDB observes for the B-tree index. A hash index can be partial, covering only those pages of the index that are often accessed.

You can monitor the use of the adaptive hash index and the contention for its use in the SEMAPHORES section of the output of the SHOW ENGINE INNODB STATUS command. If you see many threads waiting on an RW-latch created in btr0sea.c, then it might be useful to disable adaptive hash indexing.

For more information about the performance characteristics of hash indexes, see Section 8.3.8, “Comparison of B-Tree and Hash Indexes”.

#### 14.2.3.12.7. Physical Row Structure

The physical row structure for an InnoDB table depends on the row format specified when the table was created. By default, InnoDB uses the Antelope file format and its COMPACT row format. The REDUNDANT format is available to retain compatibility with older versions of MySQL. When you enable the innodb\_file\_per\_table [1718] setting, you can also make use of the newer Barracuda file format, with its DYNAMIC and COMPRESSED row formats, as explained in Section 5.4.8, “How InnoDB Stores Variable-Length Columns” and Section 5.4.6, “Working with InnoDB Compressed Tables”.

To check the row format of an InnoDB table, use SHOW TABLE STATUS.

The COMPACT row format decreases row storage space by about 20% at the cost of increasing CPU use for some operations. If your workload is a typical one that is limited by cache hit rates and disk speed, COMPACT format is likely to be faster. If the workload is a rare case that is limited by CPU speed, COMPACT format might be slower.

Rows in InnoDB tables that use REDUNDANT row format have the following characteristics:

- Each index record contains a 6-byte header. The header is used to link together consecutive records, and also in row-level locking.
- Records in the clustered index contain fields for all user-defined columns. In addition, there is a 6-byte transaction ID field and a 7-byte roll pointer field.
- If no primary key was defined for a table, each clustered index record also contains a 6-byte row ID field.
- Each secondary index record also contains all the primary key fields defined for the clustered index key that are not in the secondary index.
- A record contains a pointer to each field of the record. If the total length of the fields in a record is less than 128 bytes, the pointer is one byte; otherwise, two bytes. The array of these pointers is called the record directory. The area where these pointers point is called the data part of the record.

- Internally, InnoDB stores fixed-length character columns such as `CHAR(10)` in a fixed-length format. InnoDB does not truncate trailing spaces from `VARCHAR` columns.
- An SQL `NULL` value reserves one or two bytes in the record directory. Besides that, an SQL `NULL` value reserves zero bytes in the data part of the record if stored in a variable length column. In a fixed-length column, it reserves the fixed length of the column in the data part of the record. Reserving the fixed space for `NULL` values enables an update of the column from `NULL` to a non-`NULL` value to be done in place without causing fragmentation of the index page.

Rows in InnoDB tables that use `COMPACT` row format have the following characteristics:

- Each index record contains a 5-byte header that may be preceded by a variable-length header. The header is used to link together consecutive records, and also in row-level locking.
- The variable-length part of the record header contains a bit vector for indicating `NULL` columns. If the number of columns in the index that can be `NULL` is  $N$ , the bit vector occupies `CEILING(N/8)` bytes. (For example, if there are anywhere from 9 to 15 columns that can be `NULL`, the bit vector uses two bytes.) Columns that are `NULL` do not occupy space other than the bit in this vector. The variable-length part of the header also contains the lengths of variable-length columns. Each length takes one or two bytes, depending on the maximum length of the column. If all columns in the index are `NOT NULL` and have a fixed length, the record header has no variable-length part.
- For each non-`NULL` variable-length field, the record header contains the length of the column in one or two bytes. Two bytes will only be needed if part of the column is stored externally in overflow pages or the maximum length exceeds 255 bytes and the actual length exceeds 127 bytes. For an externally stored column, the 2-byte length indicates the length of the internally stored part plus the 20-byte pointer to the externally stored part. The internal part is 768 bytes, so the length is 768+20. The 20-byte pointer stores the true length of the column.
- The record header is followed by the data contents of the non-`NULL` columns.
- Records in the clustered index contain fields for all user-defined columns. In addition, there is a 6-byte transaction ID field and a 7-byte roll pointer field.
- If no primary key was defined for a table, each clustered index record also contains a 6-byte row ID field.
- Each secondary index record also contains all the primary key fields defined for the clustered index key that are not in the secondary index. If any of these primary key fields are variable length, the record header for each secondary index will have a variable-length part to record their lengths, even if the secondary index is defined on fixed-length columns.
- Internally, InnoDB stores fixed-length, fixed-width character columns such as `CHAR(10)` in a fixed-length format. InnoDB does not truncate trailing spaces from `VARCHAR` columns.
- Internally, InnoDB attempts to store UTF-8 `CHAR(N)` columns in  $N$  bytes by trimming trailing spaces. (With `REDUNDANT` row format, such columns occupy  $3 \times N$  bytes.) Reserving the minimum space  $N$  in many cases enables column updates to be done in place without causing fragmentation of the index page.

### 14.2.3.13. The InnoDB Recovery Process

InnoDB crash recovery consists of several steps:

The first step, applying the redo log, is performed during initialization, before accepting any connections. If all changes were flushed from the buffer pool to the tablespaces (`ibdata*` and `*.ibd` files) at the time of the shutdown or crash, the redo log application can be skipped. If the redo log files are missing at startup, InnoDB skips the redo log application.

The remaining steps after redo log application do not depend on the redo log (other than for logging the writes) and are performed in parallel with normal processing. These include:



- **Rolling back incomplete transactions:** Any transactions that were active at the time of crash or **fast shutdown**.
- **Insert buffer merge:** Applying changes from the insert buffer (part of the **system tablespace**) to leaf pages of secondary indexes, as the index pages are read to the buffer pool.
- **Purge:** Deleting delete-marked records that are no longer visible for any active transaction.

Of these, only rollback of incomplete transactions is special to crash recovery. The insert buffer merge and the purge are performed during normal processing.

In most situations, even if the MySQL server was killed unexpectedly in the middle of heavy activity, the recovery process happens automatically and no action is needed from the DBA. If a hardware failure or severe system error corrupted InnoDB data, MySQL might refuse to start. In that case, see [Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”](#) for the steps to troubleshoot such an issue.

### 14.2.3.14. InnoDB Error Handling

Error handling in InnoDB is not always the same as specified in the SQL standard. According to the standard, any error during an SQL statement should cause rollback of that statement. InnoDB sometimes rolls back only part of the statement, or the whole transaction. The following items describe how InnoDB performs error handling:

- If you run out of file space in a **tablespace**, a MySQL **Table is full** error occurs and InnoDB rolls back the SQL statement.
- A transaction **deadlock** causes InnoDB to **roll back** the entire **transaction**. Retry the whole transaction when this happens.

A lock wait timeout causes InnoDB to roll back only the single statement that was waiting for the lock and encountered the timeout. (To have the entire transaction roll back, start the server with the `--innodb_rollback_on_timeout [1744]` option.) Retry the statement if using the current behavior, or the entire transaction if using `--innodb_rollback_on_timeout [1744]`.

Both deadlocks and lock wait timeouts are normal on busy servers and it is necessary for applications to be aware that they may happen and handle them by retrying. You can make them less likely by doing as little work as possible between the first change to data during a transaction and the commit, so the locks are held for the shortest possible time and for the smallest possible number of rows. Sometimes splitting work between different transactions may be practical and helpful.

When a transaction rollback occurs due to a deadlock or lock wait timeout, it cancels the effect of the statements within the transaction. But if the start-transaction statement was `START TRANSACTION` or `BEGIN` statement, rollback does not cancel that statement. Further SQL statements become part of the transaction until the occurrence of `COMMIT`, `ROLLBACK`, or some SQL statement that causes an implicit commit.

- A duplicate-key error rolls back the SQL statement, if you have not specified the `IGNORE` option in your statement.
- A **row too long error** rolls back the SQL statement.
- Other errors are mostly detected by the MySQL layer of code (above the InnoDB storage engine level), and they roll back the corresponding SQL statement. Locks are not released in a rollback of a single SQL statement.

During implicit rollbacks, as well as during the execution of an explicit `ROLLBACK` SQL statement, `SHOW PROCESSLIST` displays **Rolling back** in the **State** column for the relevant connection.

#### 14.2.3.14.1. InnoDB Error Codes

The following is a nonexhaustive list of common InnoDB-specific errors that you may encounter, with information about why each occurs and how to resolve the problem.

- 1005 (ER\_CANT\_CREATE\_TABLE)

Cannot create table. If the error message refers to error 150, table creation failed because a [foreign key constraint](#) was not correctly formed. If the error message refers to error -1, table creation probably failed because the table includes a column name that matched the name of an internal InnoDB table.

- 1016 (ER\_CANT\_OPEN\_FILE)

Cannot find the InnoDB table from the InnoDB data files, although the .frm file for the table exists. See [Section 14.2.4.7, “Troubleshooting InnoDB Data Dictionary Operations”](#).

- 1114 (ER\_RECORD\_FILE\_FULL)

InnoDB has run out of free space in the tablespace. Reconfigure the tablespace to add a new data file.

- 1205 (ER\_LOCK\_WAIT\_TIMEOUT)

Lock wait timeout expired. The statement that waited too long was [rolled back](#) (not the entire [transaction](#)). You can increase the value of the `innodb_lock_wait_timeout` [1729] configuration option if SQL statements should wait longer for other transactions to complete, or decrease it if too many long-running transactions are causing [locking](#) problems and reducing [concurrency](#) on a busy system.

- 1206 (ER\_LOCK\_TABLE\_FULL)

The total number of locks exceeds the amount of memory InnoDB devotes to managing locks. To avoid this error, increase the value of `innodb_buffer_pool_size` [1708]. Within an individual application, a workaround may be to break a large operation into smaller pieces. For example, if the error occurs for a large `INSERT`, perform several smaller `INSERT` operations.

- 1213 (ER\_LOCK\_DEADLOCK)

The [transaction](#) encountered a [deadlock](#) and was automatically [rolled back](#) so that your application could take corrective action. To recover from this error, run all the operations in this transaction again. A deadlock occurs when requests for locks arrive in inconsistent order between transactions. The transaction that was rolled back released all its locks, and the other transaction can now get all the locks it requested. Thus when you re-run the transaction that was rolled back, it might have to wait for other transactions to complete, but typically the deadlock does not recur. If you encounter frequent deadlocks, make the sequence of locking operations (`LOCK TABLES`, `SELECT ... FOR UPDATE`, and so on) consistent between the different transactions or applications that experience the issue. See [Section 14.2.3.10, “How to Cope with Deadlocks”](#) for details.

- 1216 (ER\_NO\_REFERENCED\_ROW)

You are trying to add a row but there is no parent row, and a [foreign key constraint](#) fails. Add the parent row first.

- 1217 (ER\_ROW\_IS\_REFERENCED)

You are trying to delete a parent row that has children, and a [foreign key constraint](#) fails. Delete the children first.

#### 14.2.3.14.2. Operating System Error Codes

To print the meaning of an operating system error number, use the `perror` program that comes with the MySQL distribution.



- **Linux System Error Codes**

The following table provides a list of some common Linux system error codes. For a more complete list, see [Linux source code](#).

Number	Macro	Description
1	EPERM	Operation not permitted
2	ENOENT	No such file or directory
3	ESRCH	No such process
4	EINTR	Interrupted system call
5	EIO	I/O error
6	ENXIO	No such device or address
7	E2BIG	Arg list too long
8	ENOEXEC	Exec format error
9	EBADF	Bad file number
10	ECHILD	No child processes
11	EAGAIN	Try again
12	ENOMEM	Out of memory
13	EACCES	Permission denied
14	EFAULT	Bad address
15	ENOTBLK	Block device required
16	EBUSY	Device or resource busy
17	EEXIST	File exists
18	EXDEV	Cross-device link
19	ENODEV	No such device
20	ENOTDIR	Not a directory
21	EISDIR	Is a directory
22	EINVAL	Invalid argument
23	ENFILE	File table overflow
24	EMFILE	Too many open files
25	ENOTTY	Inappropriate ioctl for device
26	ETXTBSY	Text file busy
27	EFBIG	File too large
28	ENOSPC	No space left on device
29	ESPIPE	File descriptor does not allow seeking
30	EROFS	Read-only file system
31	EMLINK	Too many links

- **Windows System Error Codes**

The following table provides a list of some common Windows system error codes. For a complete list, see the [Microsoft Web site](#).

Number	Macro	Description
1	ERROR_INVALID_FUNCTION	Incorrect function.
2	ERROR_FILE_NOT_FOUND	The system cannot find the file specified.

Number	Macro	Description
3	ERROR_PATH_NOT_FOUND	The system cannot find the path specified.
4	ERROR_TOO_MANY_OPEN_FILES	The system cannot open the file.
5	ERROR_ACCESS_DENIED	Access is denied.
6	ERROR_INVALID_HANDLE	The handle is invalid.
7	ERROR_ARENA_TRASHED	The storage control blocks were destroyed.
8	ERROR_NOT_ENOUGH_MEMORY	Not enough storage is available to process this command.
9	ERROR_INVALID_BLOCK	The storage control block address is invalid.
10	ERROR_BAD_ENVIRONMENT	The environment is incorrect.
11	ERROR_BAD_FORMAT	An attempt was made to load a program with an incorrect format.
12	ERROR_INVALID_ACCESS	The access code is invalid.
13	ERROR_INVALID_DATA	The data is invalid.
14	ERROR_OUTOFMEMORY	Not enough storage is available to complete this operation.
15	ERROR_INVALID_DRIVE	The system cannot find the drive specified.
16	ERROR_CURRENT_DIRECTORY	The directory cannot be removed.
17	ERROR_NOT_SAME_DRIVE	The system cannot move the file to a different disk drive.
18	ERROR_NO_MORE_FILES	There are no more files.
19	ERROR_WRITE_PROTECT	The media is write protected.
20	ERROR_BAD_UNIT	The system cannot find the device specified.
21	ERROR_NOT_READY	The device is not ready.
22	ERROR_BAD_COMMAND	The device does not recognize the command.
23	ERROR_CRC	Data error (cyclic redundancy check).
24	ERROR_BAD_LENGTH	The program issued a command but the command length is incorrect.
25	ERROR_SEEK	The drive cannot locate a specific area or track on the disk.
26	ERROR_NOT_DOS_DISK	The specified disk or diskette cannot be accessed.
27	ERROR_SECTOR_NOT_FOUND	The drive cannot find the sector requested.
28	ERROR_OUT_OF_PAPER	The printer is out of paper.
29	ERROR_WRITE_FAULT	The system cannot write to the specified device.
30	ERROR_READ_FAULT	The system cannot read from the specified device.
31	ERROR_GEN_FAILURE	A device attached to the system is not functioning.
32	ERROR_SHARING_VIOLATION	The process cannot access the file because it is being used by another process.
33	ERROR_LOCK_VIOLATION	The process cannot access the file because another process has locked a portion of the file.
34	ERROR_WRONG_DISK	The wrong diskette is in the drive. Insert %2 (Volume Serial Number: %3) into drive %1.
36	ERROR_SHARING_BUFFER_EXCEEDED	Too many files opened for sharing.
38	ERROR_HANDLE_EOF	Reached the end of the file.
39	ERROR_HANDLE_DISK_FULL	The disk is full.
87	ERROR_INVALID_PARAMETER	The parameter is incorrect.
112	ERROR_DISK_FULL	The disk is full.
123	ERROR_INVALID_NAME	The file name, directory name, or volume label syntax is incorrect.
1450	ERROR_NO_SYSTEM_RESOURCES	Insufficient system resources exist to complete the requested service.

## 14.2.4. InnoDB Performance Tuning and Troubleshooting

This section explains the actions to take for InnoDB issues. Because much of the InnoDB focus is on performance, many issues are performance-related. You can use the tips here to track down problems after they occur, or do advance planning to avoid the problems entirely.

### 14.2.4.1. InnoDB Performance Tuning Tips

With InnoDB becoming the default storage engine in MySQL 5.5 and higher, the tips and guidelines for InnoDB tables are now part of the main optimization chapter. See [Section 8.5, “Optimizing for InnoDB Tables”](#).

### 14.2.4.2. InnoDB Performance and Scalability Enhancements

This section summarizes the major InnoDB features and enhancements for performance and scalability. This information is useful to any DBA or developer who is concerned with performance and scalability. Although some of the enhancements do not require any action on your part, knowing this information can still help you diagnose performance issues more quickly and modernize systems and applications that rely on older, inefficient behavior.

#### 14.2.4.2.1. Overview of InnoDB Performance

InnoDB has always been highly efficient, and includes several unique architectural elements to assure high performance and scalability. The latest InnoDB storage engine includes new features that take advantage of advances in operating systems and hardware platforms, such as multi-core processors and improved memory allocation systems. In addition, new configuration options let you better control some InnoDB internal subsystems to achieve the best performance with your workload.

Starting with MySQL 5.5 and InnoDB 1.1, the built-in InnoDB storage engine within MySQL is upgraded to the full feature set and performance of the former InnoDB Plugin. This change makes these performance and scalability enhancements available to a much wider audience than before, and eliminates the separate installation step of the InnoDB Plugin. After learning about the InnoDB performance features in this section, continue with [Chapter 8, Optimization](#) to learn the best practices for overall MySQL performance, and [Section 8.5, “Optimizing for InnoDB Tables”](#) in particular for InnoDB tips and guidelines.

#### 14.2.4.2.2. Compression Enhancements for OLTP Workloads

Traditionally, the InnoDB [compression](#) feature was recommended primarily for read-only or read-mostly [workloads](#), such as in a [data warehouse](#) configuration. The rise of [SSD](#) storage devices, which are fast but relatively small and expensive, makes compression attractive also for [OLTP](#) workloads: high-traffic, interactive web sites can reduce their storage requirements and their I/O operations per second ([IOPS](#)) by using compressed tables with applications that do frequent [INSERT](#), [UPDATE](#), and [DELETE](#) operations.

New configuration options in MySQL 5.6 let you adjust the way compression works for a particular MySQL instance, with an emphasis on performance and scalability for write-intensive operations:

- [innodb\\_compression\\_level \[1713\]](#) lets you turn the degree of compression up or down. A higher value lets you fit more data onto a storage device, at the expense of more CPU overhead during compression. A lower value lets you reduce CPU overhead when storage space is not critical, or you expect the data is not especially compressible.
- [innodb\\_compression\\_failure\\_threshold\\_pct \[1713\]](#) specifies a cutoff point for [compression failures](#) during updates to a compressed table. When this threshold is passed, MySQL begins to leave additional free space within each new compressed page, dynamically adjusting the amount of free space up to the percentage of page size specified by [innodb\\_compression\\_pad\\_pct\\_max \[1713\]](#)
- [innodb\\_compression\\_pad\\_pct\\_max \[1713\]](#) lets you adjust the maximum amount of space reserved within each [page](#) to record changes to compressed rows, without needing to compress the

entire page again. The higher the value, the more changes can be recorded without recompressing the page. MySQL uses a variable amount of free space for the pages within each compressed table, only when a designated percentage of compression operations “fail” at runtime, requiring an expensive operation to split the compressed page.

Because working with compressed data sometimes involves keeping both compressed and uncompressed versions of a page in memory at the same time, when using compression with an OLTP-style workload, be prepared to increase the value of the `innodb_buffer_pool_size` [1708] configuration option.

For more information on MySQL data compression, see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#). For the performance aspects, especially see the section [Section 5.4.6.3, “Tuning Compression for InnoDB Tables”](#).

#### 14.2.4.2.3. Optimizations for Read-Only Transactions

When a transaction is known to be read-only, InnoDB can avoid the overhead associated with setting up the [transaction ID](#) (`TRX_ID` field). The transaction ID is only needed for a [transaction](#) that might perform write operations or [locking reads](#) such as `SELECT ... FOR UPDATE`. Eliminating these unnecessary transaction IDs reduces the size of internal data structures that are consulted each time a query or DML statement constructs a [read view](#).

Currently, InnoDB detects the read-only nature of the transaction and applies this optimization when any of the following conditions are met:

- The transaction is started with the `START TRANSACTION READ ONLY` statement. In this case, attempting to make any changes to the database (for InnoDB, MyISAM, or other types of tables) causes an error, and the transaction continues in read-only state:

```
ERROR 1792 (25006): Cannot execute statement in a READ ONLY transaction.
```

You can still make changes to session-specific temporary tables in a read-only transaction, or issue locking queries for them, because those changes and locks are not visible to any other transaction.

- The `autocommit` setting is turned on, so that the transaction is guaranteed to be a single statement, and the single statement making up the transaction is a “non-locking” `SELECT` statement. That is, a `SELECT` that does not use a `FOR UPDATE` or `LOCK IN SHARED MODE` clause.

Thus, for a read-intensive application such as a report generator, you can tune a sequence of InnoDB queries by grouping them inside `START TRANSACTION READ ONLY` and `COMMIT`, or by turning on the `autocommit` setting before running the `SELECT` statements, or simply by avoiding any DML statements interspersed with the queries.



#### Note

Because these optimized transactions are kept out of certain internal InnoDB data structures, they are not listed in `SHOW ENGINE INNODB STATUS` output.

#### 14.2.4.2.4. Separate Tablespaces for InnoDB Undo Logs

This feature optionally moves the [InnoDB undo log](#) out of the [system tablespace](#) into one or more separate [tablespaces](#). The I/O patterns for the undo log make these new tablespaces good candidates to move to [SSD](#) storage, while keeping the system tablespace on hard disk storage. Users cannot drop the separate tablespaces created to hold InnoDB undo logs, or the individual [segments](#) inside those tablespaces.

Because these files handle I/O operations formerly done inside the system tablespace, we broaden the definition of system tablespace to include these new files.

The undo logs are also known as the [rollback segments](#).

This feature involves the following new or renamed configuration options:

- `innodb_undo_tablespaces` [1753].
- `innodb_undo_directory` [1752].
- `innodb_rollback_segments` becomes `innodb_undo_logs` [1752]. The old name is still available for compatibility.

## Usage Notes

To use this feature, follow these steps:

1. Decide on a path on a fast storage device to hold the undo logs. You will specify that path as the argument to the `innodb_undo_directory` [1752] option in your MySQL configuration file or startup script.
2. Decide on a non-zero starting value for the `innodb_undo_logs` [1752] option. You can start with a relatively low value and increase it over time to examine the effect on performance.
3. Decide on a non-zero value for the `innodb_undo_tablespaces` [1753] option. The multiple undo logs specified by the `innodb_undo_logs` value are divided between this many separate tablespaces (represented by `.ibd` files). This value is fixed for the life of the MySQL instance, so if you are uncertain about the optimal value, estimate on the high side.
4. Set up an entirely new MySQL instance for testing, using the values you chose in the configuration file or in your MySQL startup script. Use a realistic workload with data volume similar to your production servers.
5. Benchmark the performance of I/O intensive workloads.
6. Periodically increase the value of `innodb_undo_logs` and re-do the performance tests. Find the value where you stop experiencing gains in I/O performance.
7. Deploy a new production instance using the ideal settings for these options. Set it up as a [slave](#) server in a [replication](#) configuration, or transfer data from an earlier production instance.

## Performance and Scalability Considerations

Keeping the undo logs in separate files allows the MySQL team to implement I/O and memory optimizations related to this transactional data. For example, because the undo data is written to disk and then rarely used (only in case of crash recovery), it does not need to be kept in the filesystem memory cache, in turn allowing a higher percentage of system memory to be devoted to the [InnoDB buffer pool](#).

The typical SSD best practice of keeping the [InnoDB](#) system tablespace on a hard drive and moving the per-table tablespaces to SSD, is assisted by moving the undo information into separate tablespace files.

## Internals

The physical tablespace files are named `undoN`, where `N` is the space ID, including leading zeros.

Currently, MySQL instances containing separate undo tablespaces cannot be downgraded to earlier releases such as MySQL 5.5 or 5.1.

### 14.2.4.2.5. Faster Extension for InnoDB Data Files

The benefits of the [InnoDB file-per-table](#) setting come with the tradeoff that each `.ibd` file is extended as the data inside the table grows. This I/O operation can be a bottleneck for busy systems with many [InnoDB](#) tables. When all [InnoDB](#) tables are stored inside the [system tablespace](#), this extension operation happens less frequently, as space freed by `DELETE` or `TRUNCATE` operations within one table can be reused by another table.

MySQL 5.6 improves the concurrency of the extension operation, so that multiple `.ibd` files can be extended simultaneously, and this operation does not block read or write operations performed by other threads.

#### 14.2.4.2.6. Non-Recursive Deadlock Detection

The code that detects [deadlocks](#) in [InnoDB transactions](#) has been modified to use a fixed-size work area rather than a recursive algorithm. The resulting detection operation is faster as a result. You do not need to do anything to take advantage of this enhancement.

Under both the old and new detection mechanisms, you might encounter a `search too deep` error that is not a true deadlock, but requires you to re-try the transaction the same way as with a deadlock.

#### 14.2.4.2.7. Fast CRC32 Checksum Algorithm

You can enable the configuration option `innodb_checksum_algorithm=crc32` [\[1710\]](#) configuration setting to change the `checksum` algorithm to a faster one that scans the block 32 bits at a time rather than 8 bits at a time. When the CRC32 algorithm is enabled, data blocks that are written to disk by [InnoDB](#) contain different values in their checksum fields than before. This process could be gradual, with a mix of old and new checksum values within the same table or database.

For maximum downward compatibility, this setting is off by default:

- Current versions of [MySQL Enterprise Backup](#) (up to 3.8.0) do not support backing up tablespaces that use `crc32` checksums.
- `.ibd` files containing `crc32` checksums could cause problems downgrading to MySQL versions prior to 5.6.3. MySQL 5.6.3 and up recognizes either the new or old checksum values for the block as correct when reading the block from disk, ensuring that data blocks are compatible during upgrade and downgrade regardless of the algorithm setting. If data written with new checksum values is processed by an level of MySQL earlier than 5.6.3, it could be reported as corrupted.

When you set up a new MySQL instance, and can be sure that all the [InnoDB](#) data is created using the CRC32 checksum algorithm, you can use the setting `innodb_checksum_algorithm=strict_crc32` [\[1710\]](#), which can be faster than the `crc32` setting because it does not do the extra checksum calculations to support both old and new values.

The `innodb_checksum_algorithm` [\[1710\]](#) option has other values that allow it to replace the `innodb_checksums` [\[1711\]](#) option. `innodb_checksum_algorithm=none` is the same as `innodb_checksums=OFF`. `innodb_checksum_algorithm=innodb` is the same as `innodb_checksums=ON`. To avoid conflicts, remove references to `innodb_checksums` from your configuration file and MySQL startup scripts. The new option also accepts values `strict_none` and `strict_innodb`, again offering better performance in situations where all [InnoDB](#) data in an instance is created with the same checksum algorithm.

The following table illustrates the difference between the `none`, `innodb`, and `crc32` option values, and their `strict_` counterparts. `none`, `innodb`, and `crc32` write the specified type checksum value into each data block, but for compatibility accept any of the other checksum values when verifying a block during a read operation. The `strict_` form of each parameter only recognizes one kind of checksum, which makes verification faster but requires that all [InnoDB](#) data files in an instance be created under the identical `innodb_checksum_algorithm` value.

**Table 14.3. Allowed Settings for `innodb_checksum_algorithm`**

Value	Generated checksum (when writing)	Allowed checksums (when reading)
<code>none</code>	A constant number.	Any of the checksums generated by <code>none</code> , <code>innodb</code> , or <code>crc32</code> .
<code>innodb</code>	A checksum calculated in software, using the original algorithm from <a href="#">InnoDB</a> .	Any of the checksums generated by <code>none</code> , <code>innodb</code> , or <code>crc32</code> .



Value	Generated checksum (when writing)	Allowed checksums (when reading)
<code>crc32</code>	A checksum calculated using the <code>crc32</code> algorithm, possibly done with a hardware assist.	Any of the checksums generated by <code>none</code> , <code>innodb</code> , or <code>crc32</code> .
<code>strict_none</code>	A constant number	Only the checksum generated by <code>none</code> .
<code>strict_innodb</code>	A checksum calculated in software, using the original algorithm from InnoDB.	Only the checksum generated by <code>innodb</code> .
<code>strict_crc32</code>	A checksum calculated using the <code>crc32</code> algorithm, possibly done with a hardware assist.	Only the checksum generated by <code>crc32</code> .

#### 14.2.4.2.8. Faster Restart by Preloading the InnoDB Buffer Pool

After you restart a busy server, there is typically a [warmup](#) period with steadily increasing throughput, as disk pages that were in the [InnoDB buffer pool](#) are brought back into memory as the same data is queried, updated, and so on. Once the buffer pool holds a similar set of pages as before the restart, many operations are performed in memory rather than involving disk I/O, and throughput stabilizes at a high level.

This feature shortens the warmup period by immediately reloading disk pages that were in the buffer pool before the restart, rather than waiting for DML operations to access the corresponding rows. The I/O requests can be performed in large batches, making the overall I/O faster. The page loading happens in the background, and does not delay the database startup.

In addition to saving the buffer pool state at shutdown and restoring it at startup, you can also save or restore the state at any time. For example, you might save the state of the buffer pool after reaching a stable throughput under a steady workload. You might restore the previous buffer pool state after running reports or maintenance jobs that bring data pages into the buffer pool that are only needed during the time period for those operations, or after some other period with a non-typical workload.

Although the buffer pool itself could be many gigabytes in size, the data that InnoDB saves on disk to restore the buffer pool is tiny by comparison: just the tablespace and page IDs necessary to locate the appropriate pages on disk. This information is derived from the `information_schema` table `innodb_buffer_page_lru`.

Because the data is cached in and aged out of the buffer pool the same as with regular database operations, there is no problem if the disk pages were updated recently, or if a DML operation involves data that has not yet been loaded. The loading mechanism skips any requested pages that no longer exist.

This feature involves the configuration variables:

- [innodb\\_buffer\\_pool\\_dump\\_now](#) [1706]
- [innodb\\_buffer\\_pool\\_load\\_now](#) [1708]
- [innodb\\_buffer\\_pool\\_dump\\_at\\_shutdown](#) [1705]
- [innodb\\_buffer\\_pool\\_load\\_at\\_startup](#) [1707]
- [innodb\\_buffer\\_pool\\_filename](#) [1706]
- [innodb\\_buffer\\_pool\\_load\\_abort](#) [1707]

and the status variables:

- [Innodb\\_buffer\\_pool\\_dump\\_status](#) [590]
- [Innodb\\_buffer\\_pool\\_load\\_status](#) [590]



To save the current state of the [InnoDB](#) buffer pool, issue the statement:

```
SET innodb_buffer_pool_dump_now=ON;
```

The underlying mechanism involves a background thread that is dispatched to perform the dump and load operations.

By default, the buffer pool state is saved in a file `ib_buffer_pool` in the [InnoDB](#) data directory.

Disk pages from compressed tables are loaded into the buffer pool in their compressed form. Uncompression happens as usual when the page contents are accessed in the course of DML operations. Because decompression is a CPU-intensive process, it is more efficient for concurrency to perform that operation in one of the connection threads rather than the single thread that performs the buffer pool restore operation.

#### Example 14.1. Examples of Dumping and Restoring the InnoDB Buffer Pool

Trigger a dump of the buffer pool manually:

```
SET innodb_buffer_pool_dump_now=ON;
```

Specify that a dump should be taken at shutdown:

```
SET innodb_buffer_pool_dump_at_shutdown=ON;
```

Specify that a dump should be loaded at startup:

```
SET innodb_buffer_pool_load_at_startup=ON;
```

Trigger a load of the buffer pool manually:

```
SET innodb_buffer_pool_load_now=ON;
```

Specify which filename to use for storing the dump to and loading the dump from:

```
SET innodb_buffer_pool_filename='filename';
```

Display progress of dump:

```
SHOW STATUS LIKE 'innodb_buffer_pool_dump_status';
```

or:

```
SELECT variable_value FROM information_schema.global_status WHERE
variable_name = 'INNODB_BUFFER_POOL_DUMP_STATUS';
```

Outputs any of: not started, Dumping buffer pool 5/7, page 237/2873, Finished at 110505 12:18:02

Display progress of load:

```
SHOW STATUS LIKE 'innodb_buffer_pool_load_status';
```

or:

```
SELECT variable_value FROM information_schema.global_status WHERE
variable_name = 'INNODB_BUFFER_POOL_LOAD_STATUS';
```

Outputs any of: not started, Loaded 123/22301 pages, Finished at 110505 12:23:24

Abort a buffer pool load:

```
SET innodb_buffer_pool_load_abort=ON;
```

#### 14.2.4.2.9. Improvements to Buffer Pool Flushing

The new configuration options `innodb_flush_neighbors` [1721] and `innodb_lru_scan_depth` [1734] let you fine-tune certain aspects of the [flushing](#) process for the [InnoDB buffer pool](#). These options primarily help write-intensive [workloads](#). With heavy [DML](#) activity,

flushing can fall behind if it is not aggressive enough, resulting in excessive memory use in the buffer pool; or, disk writes due to flushing can saturate your I/O capacity if that mechanism is too aggressive. The ideal settings depend on your workload, data access patterns, and storage configuration (for example, whether data is stored on [HDD](#) or [SSD](#) devices).

For systems with constant heavy [workloads](#), or workloads that fluctuate widely, several new configuration options let you fine-tune the [flushing](#) behavior for InnoDB tables: [innodb\\_adaptive\\_flushing\\_lwm](#) [1701], [innodb\\_max\\_dirty\\_pages\\_pct\\_lwm](#) [1735], [innodb\\_io\\_capacity\\_max](#) [1728], and [innodb\\_flushing\\_avg\\_loops](#) [1722]. These options feed into an improved formula used by the [innodb\\_adaptive\\_flushing](#) [1700] option.

The existing [innodb\\_adaptive\\_flushing](#) [1700], [innodb\\_io\\_capacity](#) [1727] and [innodb\\_max\\_dirty\\_pages\\_pct](#) [1734] options work as before, except that they are limited or extended by other options: [innodb\\_adaptive\\_flushing\\_lwm](#) [1701], [innodb\\_io\\_capacity\\_max](#) [1728] and [innodb\\_max\\_dirty\\_pages\\_pct\\_lwm](#) [1735]:

- The [InnoDB adaptive flushing](#) mechanism is not appropriate in all cases. It gives the most benefit when the [redo log](#) is in danger of filling up. The [innodb\\_adaptive\\_flushing\\_lwm](#) [1701] option specifies a percentage of redo log capacity; when that threshold is crossed, InnoDB turns on adaptive flushing even if not specified by the [innodb\\_adaptive\\_flushing](#) [1700] option.
- If flushing activity falls far behind, InnoDB can flush more aggressively than specified by [innodb\\_io\\_capacity](#) [1727]. [innodb\\_io\\_capacity\\_max](#) [1728] represents an upper limit on the I/O capacity used in such emergency situations, so that the spike in I/O does not consume all the capacity of the server.
- When the [innodb\\_max\\_dirty\\_pages\\_pct](#) [1734] threshold is crossed, InnoDB can begin aggressively flushing pages to disk. The [innodb\\_max\\_dirty\\_pages\\_pct\\_lwm](#) [1735] option specifies a higher value at which InnoDB begins gently flushing pages, ideally preventing the percentage of dirty pages from reaching [innodb\\_max\\_dirty\\_pages\\_pct](#) [1734]. A value of [innodb\\_max\\_dirty\\_pages\\_pct\\_lwm=0](#) [1735] disables this “preflushing” behavior.

All of these options are most applicable for servers running heavy workloads for long periods of time, when there is rarely enough idle time to catch up with changes waiting to be written to disk. The [innodb\\_flushing\\_avg\\_loops](#) [1722] lets you distinguish between a server that is running at full capacity 24x7 and one that experiences periodic spikes in workload. For a server with a consistently high workload, keep this value high so that the adaptive flushing algorithm responds immediately to changes in the I/O rate. For a server that experiences peaks and troughs in its workload, keep this value low so that InnoDB does not overreact to sudden spikes in [DML](#) activity.

#### 14.2.4.2.10. Persistent Optimizer Statistics for InnoDB Tables

[Plan stability](#) is a desirable goal for your biggest and most important queries. InnoDB has always computed statistics for each InnoDB table to help the optimizer find the most efficient [query execution plan](#). Now you can make these statistics persistent, so that the index usage and join order for a particular query is less likely to change.

This feature is on by default, enabled by the configuration option [innodb\\_stats\\_persistent](#) [1747].

You control how much sampling is done to collect the statistics by setting the configuration options [innodb\\_stats\\_persistent\\_sample\\_pages](#) [1748] and [innodb\\_stats\\_transient\\_sample\\_pages](#) [1748].

The configuration option [innodb\\_stats\\_auto\\_recalc](#) [1746] determines whether the statistics are calculated automatically whenever a table undergoes substantial changes (to more than 10% of the rows). If that setting is disabled, ensure the accuracy of optimizer statistics by issuing the [ANALYZE TABLE](#) statement for each applicable table after creating an index or making substantial changes to indexed columns. You might run this statement in your setup scripts after representative data has been loaded into the table, and run it periodically after DML operations significantly change the contents of indexed columns, or on a schedule at times of low activity.



### Caution

To ensure statistics are gathered when a new index is created, either enable the `innodb_stats_auto_recalc` [1746] option, or run `ANALYZE TABLE` after creating each new index when the persistent statistics mode is enabled.

You can also set `innodb_stats_persistent` [1747], `innodb_stats_auto_recalc`, and `innodb_stats_sample_pages` options at the session level before creating a table, or use the `STATS_PERSISTENT`, `STATS_AUTO_RECALC`, and `STATS_SAMPLE_PAGES` clauses on the `CREATE TABLE` and `ALTER TABLE` statements, to override the system-wide setting and configure persistent statistics for individual tables.

Formerly, these statistics were cleared on each server restart and after some other operations, and recomputed when the table was next accessed. The statistics are computed using a random sampling technique that could produce different estimates the next time, leading to different choices in the execution plan and thus variations in query performance.

To revert to the previous method of collecting statistics that are periodically erased, run the command `ALTER TABLE tbl_name STATS_PERSISTENT=0`.

The persistent statistics feature relies on the internally managed tables in the `mysql` database, named `innodb_table_stats` and `innodb_index_stats`. These tables are set up automatically in all install, upgrade, and build-from-source procedures.

If you manually update the statistics in the tables during troubleshooting or tuning, issue the command `FLUSH TABLE tbl_name` to make MySQL reload the updated statistics.

#### 14.2.4.2.11. Faster Locking for Improved Scalability

In MySQL and InnoDB, multiple threads of execution access shared data structures. InnoDB synchronizes these accesses with its own implementation of [mutexes](#) and [read/write locks](#). InnoDB has historically protected the internal state of a read/write lock with an InnoDB mutex. On Unix and Linux platforms, the internal state of an InnoDB mutex is protected by a [Pthreads](#) mutex, as in IEEE Std 1003.1c (POSIX.1c).

On many platforms, there is a more efficient way to implement mutexes and read/write locks. [Atomic](#) operations can often be used to synchronize the actions of multiple threads more efficiently than Pthreads. Each operation to acquire or release a lock can be done in fewer CPU instructions, and thus result in less wasted time when threads are contending for access to shared data structures. This in turn means greater scalability on multi-core platforms.

InnoDB implements mutexes and read/write locks with the [built-in functions provided by the GNU Compiler Collection \(GCC\) for atomic memory access](#) instead of using the Pthreads approach previously used. More specifically, an InnoDB that is compiled with GCC version 4.1.2 or later uses the atomic builtins instead of a `pthread_mutex_t` to implement InnoDB mutexes and read/write locks.

On 32-bit Microsoft Windows, InnoDB has implemented mutexes (but not read/write locks) with hand-written assembler instructions. Beginning with Microsoft Windows 2000, functions for [Interlocked Variable Access](#) are available that are similar to the built-in functions provided by GCC. On Windows 2000 and higher, InnoDB makes use of the Interlocked functions. Unlike the old hand-written assembler code, the new implementation supports read/write locks and 64-bit platforms.

Solaris 10 introduced library functions for [atomic operations](#), and InnoDB uses these functions by default. When MySQL is compiled on Solaris 10 with a compiler that does not support the [built-in functions provided by the GNU Compiler Collection \(GCC\) for atomic memory access](#), InnoDB uses the library functions.

This change improves the scalability of InnoDB on multi-core systems. This feature is enabled out-of-the-box on the platforms where it is supported. You do not have to set any parameter or option to take advantage of the improved performance. On platforms where the GCC, Windows, or Solaris

functions for atomic memory access are not available, InnoDB uses the traditional Pthreads method of implementing mutexes and read/write locks.

When MySQL starts, InnoDB writes a message to the log file indicating whether atomic memory access is used for mutexes, for mutexes and read/write locks, or neither. If suitable tools are used to build InnoDB and the target CPU supports the atomic operations required, InnoDB uses the built-in functions for mutexing. If, in addition, the compare-and-swap operation can be used on thread identifiers (`pthread_t`), then InnoDB uses the instructions for read-write locks as well.

Note: If you are building from source, ensure that the build process properly takes advantage of your platform capabilities.

For more information about the performance implications of locking, see [Section 8.10, “Optimizing Locking Operations”](#).

#### 14.2.4.2.12. Using Operating System Memory Allocators

When InnoDB was developed, the memory allocators supplied with operating systems and run-time libraries were often lacking in performance and scalability. At that time, there were no memory allocator libraries tuned for multi-core CPUs. Therefore, InnoDB implemented its own memory allocator in the `mem` subsystem. This allocator is guarded by a single mutex, which may become a [bottleneck](#). InnoDB also implements a wrapper interface around the system allocator (`malloc` and `free`) that is likewise guarded by a single mutex.

Today, as multi-core systems have become more widely available, and as operating systems have matured, significant improvements have been made in the memory allocators provided with operating systems. New memory allocators perform better and are more scalable than they were in the past. The leading high-performance memory allocators include [Hoard](#), [libumem](#), [mtmalloc](#), [ptmalloc](#), [tbbmalloc](#), and [TCMalloc](#). Most workloads, especially those where memory is frequently allocated and released (such as multi-table joins), benefit from using a more highly tuned memory allocator as opposed to the internal, InnoDB-specific memory allocator.

You can control whether InnoDB uses its own memory allocator or an allocator of the operating system, by setting the value of the system configuration parameter `innodb_use_sys_malloc` [\[1754\]](#) in the MySQL option file (`my.cnf` or `my.ini`). If set to `ON` or `1` (the default), InnoDB uses the `malloc` and `free` functions of the underlying system rather than manage memory pools itself. This parameter is not dynamic, and takes effect only when the system is started. To continue to use the InnoDB memory allocator, set `innodb_use_sys_malloc` [\[1754\]](#) to `0`.



#### Note

When the InnoDB memory allocator is disabled, InnoDB ignores the value of the parameter `innodb_additional_mem_pool_size` [\[1702\]](#). The InnoDB memory allocator uses an additional memory pool for satisfying allocation requests without having to fall back to the system memory allocator. When the InnoDB memory allocator is disabled, all such allocation requests are fulfilled by the system memory allocator.

On Unix-like systems that use dynamic linking, replacing the memory allocator may be as easy as making the environment variable `LD_PRELOAD` or `LD_LIBRARY_PATH` point to the dynamic library that implements the allocator. On other systems, some relinking may be necessary. Please refer to the documentation of the memory allocator library of your choice.

Since InnoDB cannot track all memory use when the system memory allocator is used (`innodb_use_sys_malloc` [\[1754\]](#) is `ON`), the section “BUFFER POOL AND MEMORY” in the output of the `SHOW ENGINE INNODB STATUS` command only includes the buffer pool statistics in the “Total memory allocated”. Any memory allocated using the `mem` subsystem or using `ut_malloc` is excluded.

For more information about the performance implications of InnoDB memory usage, see [Section 8.9, “Buffering and Caching”](#).

#### 14.2.4.2.13. Controlling InnoDB Change Buffering

When `INSERT`, `UPDATE`, and `DELETE` operations are done to a table, often the values of indexed columns (particularly the values of secondary keys) are not in sorted order, requiring substantial I/O to bring secondary indexes up to date. InnoDB has an `insert buffer` that caches changes to secondary index entries when the relevant `page` is not in the `buffer pool`, thus avoiding I/O operations by not reading in the page from the disk. The buffered changes are merged when the page is loaded to the buffer pool, and the updated page is later flushed to disk using the normal mechanism. The InnoDB main thread merges buffered changes when the server is nearly idle, and during a `slow shutdown`.

Because it can result in fewer disk reads and writes, this feature is most valuable for workloads that are I/O-bound, for example applications with a high volume of DML operations such as bulk inserts.

However, the insert buffer occupies a part of the buffer pool, reducing the memory available to cache data pages. If the working set almost fits in the buffer pool, or if your tables have relatively few secondary indexes, it may be useful to disable insert buffering. If the working set entirely fits in the buffer pool, insert buffering does not impose any extra overhead, because it only applies to pages that are not in the buffer pool.

You can control the extent to which InnoDB performs insert buffering with the system configuration parameter `innodb_change_buffering` [1709]. You can turn on and off buffering for inserts, delete operations (when index records are initially marked for deletion) and purge operations (when index records are physically deleted). An update operation is represented as a combination of an insert and a delete. In MySQL 5.5 and higher, the default value is changed from `inserts` to `all`.

The allowed values of `innodb_change_buffering` [1709] are:

- `all`

The default value: buffer inserts, delete-marking operations, and purges.

- `none`

Do not buffer any operations.

- `inserts`

Buffer insert operations.

- `deletes`

Buffer delete-marking operations.

- `changes`

Buffer both inserts and delete-marking.

- `purges`

Buffer the physical deletion operations that happen in the background.

You can set the value of this parameter in the MySQL option file (`my.cnf` or `my.ini`) or change it dynamically with the `SET GLOBAL` command, which requires the `SUPER` privilege. Changing the setting affects the buffering of new operations; the merging of already buffered entries is not affected.

For more information about speeding up `INSERT`, `UPDATE`, and `DELETE` statements, see [Section 8.2.2, “Optimizing DML Statements”](#).

#### 14.2.4.2.14. Controlling Adaptive Hash Indexing

If a table fits almost entirely in main memory, the fastest way to perform queries on it is to use [hash indexes](#) rather than [B-tree](#) lookups. MySQL monitors searches on each index defined for an InnoDB table. If it notices that certain index values are being accessed frequently, it automatically builds an in-memory hash table for that index. See [Section 14.2.3.12.6, “Adaptive Hash Indexes”](#) for background information and usage guidelines for the [adaptive hash index](#) feature and the [innodb\\_adaptive\\_hash\\_index \[1701\]](#) configuration option.

#### 14.2.4.2.15. Changes Regarding Thread Concurrency

InnoDB uses operating system [threads](#) to process requests from user transactions. (Transactions may issue many requests to InnoDB before they commit or roll back.) On modern operating systems and servers with multi-core processors, where context switching is efficient, most workloads run well without any limit on the number of concurrent threads. Scalability improvements in MySQL 5.5 and up reduce the need to limit the number of concurrently executing threads inside InnoDB.

In situations where it is helpful to minimize context switching between threads, InnoDB can use a number of techniques to limit the number of concurrently executing operating system threads (and thus the number of requests that are processed at any one time). When InnoDB receives a new request from a user session, if the number of threads concurrently executing is at a pre-defined limit, the new request sleeps for a short time before it tries again. A request that cannot be rescheduled after the sleep is put in a first-in/first-out queue and eventually is processed. Threads waiting for locks are not counted in the number of concurrently executing threads.

You can limit the number of concurrent threads by setting the configuration parameter [innodb\\_thread\\_concurrency \[1751\]](#). Once the number of executing threads reaches this limit, additional threads sleep for a number of microseconds, set by the configuration parameter [innodb\\_thread\\_sleep\\_delay \[1752\]](#), before being placed into the queue.

Previously, it required experimentation to find the optimal value for [innodb\\_thread\\_sleep\\_delay](#), and the optimal value could change depending on the workload. In MySQL 5.6.3 and higher, you can set the configuration option [innodb\\_adaptive\\_max\\_sleep\\_delay \[1702\]](#) to the highest value you would allow for [innodb\\_thread\\_sleep\\_delay](#), and InnoDB automatically adjusts [innodb\\_thread\\_sleep\\_delay](#) up or down depending on the current thread-scheduling activity. This dynamic adjustment helps the thread scheduling mechanism to work smoothly during times when the system is lightly loaded and when it is operating near full capacity.

The default value for [innodb\\_thread\\_concurrency \[1751\]](#) and the implied default limit on the number of concurrent threads has been changed in various releases of MySQL and InnoDB. Currently, the default value of [innodb\\_thread\\_concurrency \[1751\]](#) is 0, so that by default there is no limit on the number of concurrently executing threads, as shown in [Table 14.4, “Changes to innodb\\_thread\\_concurrency”](#).

**Table 14.4. Changes to [innodb\\_thread\\_concurrency](#)**

InnoDB Version	MySQL Version	Default value	Default limit of concurrent threads	Value to allow unlimited threads
Built-in	Earlier than 5.1.11	20	No limit	20 or higher
Built-in	5.1.11 and newer	8	8	0
InnoDB before 1.0.3	(corresponding to Plugin)	8	8	0
InnoDB 1.0.3 and newer	(corresponding to Plugin)	0	No limit	0

Note that InnoDB causes threads to sleep only when the number of concurrent threads is limited. When there is no limit on the number of threads, all contend equally to be scheduled. That is, if [innodb\\_thread\\_concurrency \[1751\]](#) is 0, the value of [innodb\\_thread\\_sleep\\_delay \[1752\]](#) is ignored.



When there is a limit on the number of threads, InnoDB reduces context switching overhead by permitting multiple requests made during the execution of a single SQL statement to enter InnoDB without observing the limit set by `innodb_thread_concurrency` [1751]. Since an SQL statement (such as a join) may comprise multiple row operations within InnoDB, InnoDB assigns “tickets” that allow a thread to be scheduled repeatedly with minimal overhead.

When a new SQL statement starts, a thread has no tickets, and it must observe `innodb_thread_concurrency` [1751]. Once the thread is entitled to enter InnoDB, it is assigned a number of tickets that it can use for subsequently entering InnoDB. If the tickets run out, `innodb_thread_concurrency` [1751] is observed again and further tickets are assigned. The number of tickets to assign is specified by the global option `innodb_concurrency_tickets` [1714], which is 500 by default. A thread that is waiting for a lock is given one ticket once the lock becomes available.

The correct values of these variables depend on your environment and workload. Try a range of different values to determine what value works for your applications. Before limiting the number of concurrently executing threads, review configuration options that may improve the performance of InnoDB on multi-core and multi-processor computers, such as `innodb_use_sys_malloc` and `innodb_adaptive_hash_index`.

For general performance information about MySQL thread handling, see [Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”](#).

#### 14.2.4.2.16. Changes in the Read-Ahead Algorithm

A [read-ahead](#) request is an I/O request to prefetch multiple pages in the [buffer pool](#) asynchronously, in anticipation that these pages will be needed soon. The requests bring in all the pages in one [extent](#). InnoDB uses two read-ahead algorithms to improve I/O performance:

**Linear** read-ahead is a technique that predicts what pages might be needed soon based on pages in the buffer pool being accessed sequentially. You control when InnoDB performs a read-ahead operation by adjusting the number of sequential page accesses required to trigger an asynchronous read request, using the configuration parameter `innodb_read_ahead_threshold` [1742]. Before this parameter was added, InnoDB would only calculate whether to issue an asynchronous prefetch request for the entire next extent when it read in the last page of the current extent.

The new configuration parameter `innodb_read_ahead_threshold` [1742] controls how sensitive InnoDB is in detecting patterns of sequential page access. If the number of pages read sequentially from an extent is greater than or equal to `innodb_read_ahead_threshold` [1742], InnoDB initiates an asynchronous read-ahead operation of the entire following extent. It can be set to any value from 0-64. The default value is 56. The higher the value, the more strict the access pattern check. For example, if you set the value to 48, InnoDB triggers a linear read-ahead request only when 48 pages in the current extent have been accessed sequentially. If the value is 8, InnoDB would trigger an asynchronous read-ahead even if as few as 8 pages in the extent were accessed sequentially. You can set the value of this parameter in the MySQL [configuration file](#), or change it dynamically with the `SET GLOBAL` command, which requires the `SUPER` privilege.

**Random** read-ahead is a technique that predicts when pages might be needed soon based on pages already in the buffer pool, regardless of the order in which those pages were read. If 13 consecutive pages from the same extent are found in the buffer pool, InnoDB asynchronously issues a request to prefetch the remaining pages of the extent. This feature was initially turned off in MySQL 5.5. It is available once again starting in MySQL 5.1.59 and 5.5.16 and higher, turned off by default. To enable this feature, set the configuration variable `innodb_random_read_ahead` [1742].

The `SHOW ENGINE INNODB STATUS` command displays statistics to help you evaluate the effectiveness of the read-ahead algorithm. With the return of random read-ahead in MySQL 5.6, the `SHOW ENGINE INNODB STATUS` command once again includes `Innodb_buffer_pool_read_ahead_rnd`. `Innodb_buffer_pool_read_ahead` keeps its current name. (In earlier releases, it was listed as `Innodb_buffer_pool_read_ahead_seq`.) See [Section 14.2.5.11, “More Read-Ahead Statistics”](#) for more information.



For more information about I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#) and [Section 8.11.3, “Optimizing Disk I/O”](#).

#### 14.2.4.2.17. Multiple Background InnoDB I/O Threads

InnoDB uses background [threads](#) to service various types of I/O requests. You can configure the number of background threads that service read and write I/O on data pages, using the configuration parameters [innodb\\_read\\_io\\_threads \[1743\]](#) and [innodb\\_write\\_io\\_threads \[1754\]](#). These parameters signify the number of background threads used for read and write requests respectively. They are effective on all supported platforms. You can set the value of these parameters in the MySQL option file (`my.cnf` or `my.ini`); you cannot change them dynamically. The default value for these parameters is 4 and the permissible values range from 1-64.

The purpose of this change is to make InnoDB more scalable on high end systems. Each background thread can handle up to 256 pending I/O requests. A major source of background I/O is the [read-ahead](#) requests. InnoDB tries to balance the load of incoming requests in such way that most of the background threads share work equally. InnoDB also attempts to allocate read requests from the same extent to the same thread to increase the chances of coalescing the requests together. If you have a high end I/O subsystem and you see more than  $64 \times$  [innodb\\_read\\_io\\_threads \[1743\]](#) pending read requests in `SHOW ENGINE INNODB STATUS`, you might gain by increasing the value of [innodb\\_read\\_io\\_threads \[1743\]](#).

For more information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

#### 14.2.4.2.18. Asynchronous I/O on Linux

Starting in InnoDB 1.1 with MySQL 5.5, the [asynchronous I/O](#) capability that InnoDB has had on Windows systems is now available on Linux systems. (Other Unix-like systems continue to use synchronous I/O calls.) This feature improves the scalability of heavily I/O-bound systems, which typically show many pending reads/writes in the output of the command `SHOW ENGINE INNODB STATUS\G`.

Running with a large number of [InnoDB I/O threads](#), and especially running multiple such instances on the same server machine, can exceed capacity limits on Linux systems. In this case, you can fix the error:

```
EAGAIN: The specified maxevents exceeds the user's limit of available events.
```

by writing a higher limit to `/proc/sys/fs/aio-max-nr`.

In general, if a problem with the asynchronous I/O subsystem in the OS prevents InnoDB from starting, set the option [innodb\\_use\\_native\\_aio=0 \[1753\]](#) in the configuration file. This new configuration option applies to Linux systems only, and cannot be changed once the server is running.

For more information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

#### 14.2.4.2.19. Group Commit

InnoDB, like any other [ACID](#)-compliant database engine, flushes the [redo log](#) of a transaction before it is committed. Historically, InnoDB used [group commit](#) functionality to group multiple such flush requests together to avoid one flush for each commit. With group commit, InnoDB issues a single write to the log file to perform the commit action for multiple user transactions that commit at about the same time, significantly improving throughput.

Group commit in InnoDB worked until MySQL 4.x, and works once again with MySQL 5.1 with the InnoDB Plugin, and MySQL 5.5 and higher. The introduction of support for the distributed transactions and Two Phase Commit (2PC) in MySQL 5.0 interfered with the InnoDB group commit functionality. This issue is now resolved.

The group commit functionality inside InnoDB works with the Two Phase Commit protocol in MySQL. Re-enabling of the group commit functionality fully ensures that the ordering of commit in the

MySQL binlog and the InnoDB logfile is the same as it was before. It means it is **totally safe to use the MySQL Enterprise Backup product with InnoDB 1.0.4** (that is, the InnoDB Plugin with MySQL 5.1) and above. When the binlog is enabled, you typically also set the configuration option `sync_binlog=0`, because group commit for the binary log is only supported if it is set to 0.

Group commit is transparent; you do not need to do anything to take advantage of this significant performance improvement.

For more information about performance of `COMMIT` and other transactional operations, see [Section 8.5.2, “Optimizing InnoDB Transaction Management”](#).

#### 14.2.4.2.20. Controlling the InnoDB Master Thread I/O Rate

The `master thread` in InnoDB is a thread that performs various tasks in the background. Most of these tasks are I/O related, such as flushing dirty pages from the buffer pool or writing changes from the insert buffer to the appropriate secondary indexes. The master thread attempts to perform these tasks in a way that does not adversely affect the normal working of the server. It tries to estimate the free I/O bandwidth available and tune its activities to take advantage of this free capacity. Historically, InnoDB has used a hard coded value of 100 IOPs (input/output operations per second) as the total I/O capacity of the server.

The parameter `innodb_io_capacity [1727]` indicates the overall I/O capacity available to InnoDB, per buffer pool instance. This parameter should be set to approximately the number of I/O operations that the system can perform per second. The value depends on your system configuration. When `innodb_io_capacity [1727]` is set, the master threads estimates the I/O bandwidth available for background tasks based on the set value. Setting the value to `100` reverts to the old behavior.

You can set the value of `innodb_io_capacity [1727]` to any number 100 or greater. The default value is `200`, reflecting that the performance of typical modern I/O devices is higher than in the early days of MySQL. Typically, values around the previous default of 100 are appropriate for consumer-level storage devices, such as hard drives up to 7200 RPMs. Faster hard drives, RAID configurations, and SSDs benefit from higher values.

You can set the value of this parameter in the MySQL option file (`my.cnf` or `my.ini`) or change it dynamically with the `SET GLOBAL` command, which requires the `SUPER` privilege.

Formerly, the InnoDB master thread also performed any needed `purge` operations. In MySQL 5.6.5 and higher, those I/O operations are moved to other background threads, whose number is controlled by the `innodb_purge_threads [1742]` configuration option.

For more information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

#### 14.2.4.2.21. Controlling the Flushing Rate of Dirty Pages from the InnoDB Buffer Pool

InnoDB performs certain tasks in the background, including `flushing` of `dirty pages` (those pages that have been changed but are not yet written to the database files) from the `buffer pool`, a task performed by the `master thread`. Currently, InnoDB aggressively flushes buffer pool pages if the percentage of dirty pages in the buffer pool exceeds `innodb_max_dirty_pages_pct [1734]`.

InnoDB uses a new algorithm to estimate the required rate of flushing, based on the speed of redo log generation and the current rate of flushing. The intent is to smooth overall performance by ensuring that buffer flush activity keeps up with the need to keep the buffer pool “clean”. Automatically adjusting the rate of flushing can help to avoid sudden dips in throughput, when excessive buffer pool flushing limits the I/O capacity available for ordinary read and write activity.

InnoDB uses its log files in a circular fashion. Before reusing a portion of a log file, InnoDB flushes to disk all dirty buffer pool pages whose redo entries are contained in that portion of the log file, a process known as a `sharp checkpoint`. If a workload is write-intensive, it generates a lot of redo information, all written to the log file. If all available space in the log files is used up, a sharp checkpoint occurs, causing a temporary reduction in throughput. This situation can happen even though `innodb_max_dirty_pages_pct [1734]` is not reached.

InnoDB uses a heuristic-based algorithm to avoid such a scenario, by measuring the number of dirty pages in the buffer pool and the rate at which redo is being generated. Based on these numbers, InnoDB decides how many dirty pages to flush from the buffer pool each second. This self-adapting algorithm is able to deal with sudden changes in the workload.

Internal benchmarking has also shown that this algorithm not only maintains throughput over time, but can also improve overall throughput significantly.

Because adaptive flushing is a new feature that can significantly affect the I/O pattern of a workload, a new configuration parameter lets you turn off this feature. The default value of the boolean parameter `innodb_adaptive_flushing` [1700] is `TRUE`, enabling the new algorithm. You can set the value of this parameter in the MySQL option file (`my.cnf` or `my.ini`) or change it dynamically with the `SET GLOBAL` command, which requires the `SUPER` privilege.

For more information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

#### 14.2.4.2.22. Using the PAUSE Instruction in InnoDB Spin Loops

Synchronization inside InnoDB frequently involves the use of `spin` loops: while waiting, InnoDB executes a tight loop of instructions repeatedly to avoid having the InnoDB `process` and `threads` be rescheduled by the operating system. If the spin loops are executed too quickly, system resources are wasted, imposing a performance penalty on transaction throughput. Most modern processors implement the `PAUSE` instruction for use in spin loops, so the processor can be more efficient.

InnoDB uses a `PAUSE` instruction in its spin loops on all platforms where such an instruction is available. This technique increases overall performance with CPU-bound workloads, and has the added benefit of minimizing power consumption during the execution of the spin loops.

You do not have to do anything to take advantage of this performance improvement.

For performance considerations for InnoDB locking operations, see [Section 8.10, “Optimizing Locking Operations”](#).

#### 14.2.4.2.23. Control of Spin Lock Polling

Many InnoDB `mutexes` and `rw-locks` are reserved for a short time. On a multi-core system, it can be more efficient for a thread to continuously check if it can acquire a mutex or rw-lock for a while before sleeping. If the mutex or rw-lock becomes available during this polling period, the thread can continue immediately, in the same time slice. However, too-frequent polling by multiple threads of a shared object can cause “cache ping pong”, different processors invalidating portions of each others' cache. InnoDB minimizes this issue by waiting a random time between subsequent polls. The delay is implemented as a busy loop.

You can control the maximum delay between testing a mutex or rw-lock using the parameter `innodb_spin_wait_delay` [1745]. The duration of the delay loop depends on the C compiler and the target processor. (In the 100MHz Pentium era, the unit of delay was one microsecond.) On a system where all processor cores share a fast cache memory, you might reduce the maximum delay or disable the busy loop altogether by setting `innodb_spin_wait_delay=0`. On a system with multiple processor chips, the effect of cache invalidation can be more significant and you might increase the maximum delay.

The default value of `innodb_spin_wait_delay` [1745] is `6`. The spin wait delay is a dynamic global parameter that you can specify in the MySQL option file (`my.cnf` or `my.ini`) or change at runtime with the command `SET GLOBAL innodb_spin_wait_delay=delay`, where `delay` is the desired maximum delay. Changing the setting requires the `SUPER` privilege.

For performance considerations for InnoDB locking operations, see [Section 8.10, “Optimizing Locking Operations”](#).

#### 14.2.4.2.24. Making the Buffer Pool Scan Resistant

Rather than using a strictly LRU algorithm, InnoDB uses a technique to minimize the amount of data that is brought into the [buffer pool](#) and never accessed again. The goal is to make sure that frequently accessed (“hot”) pages remain in the buffer pool, even as [read-ahead](#) and [full table scans](#) bring in new blocks that might or might not be accessed afterward.

Newly read blocks are inserted into the middle of the list representing the buffer pool. of the LRU list. All newly read pages are inserted at a location that by default is  $3/8$  from the tail of the LRU list. The pages are moved to the front of the list (the most-recently used end) when they are accessed in the buffer pool for the first time. Thus pages that are never accessed never make it to the front portion of the LRU list, and “age out” sooner than with a strict LRU approach. This arrangement divides the LRU list into two segments, where the pages downstream of the insertion point are considered “old” and are desirable victims for LRU eviction.

For an explanation of the inner workings of the InnoDB buffer pool and the specifics of its LRU replacement algorithm, see [Section 8.9.1, “The InnoDB Buffer Pool”](#).

You can control the insertion point in the LRU list, and choose whether InnoDB applies the same optimization to blocks brought into the buffer pool by table or index scans. The configuration parameter [innodb\\_old\\_blocks\\_pct \[1738\]](#) controls the percentage of “old” blocks in the LRU list. The default value of [innodb\\_old\\_blocks\\_pct \[1738\]](#) is 37, corresponding to the original fixed ratio of  $3/8$ . The value range is 5 (new pages in the buffer pool age out very quickly) to 95 (only 5% of the buffer pool is reserved for hot pages, making the algorithm close to the familiar LRU strategy).

The optimization that keeps the buffer pool from being churned by read-ahead can avoid similar problems due to table or index scans. In these scans, a data page is typically accessed a few times in quick succession and is never touched again. The configuration parameter [innodb\\_old\\_blocks\\_time \[1738\]](#) specifies the time window (in milliseconds) after the first access to a page during which it can be accessed without being moved to the front (most-recently used end) of the LRU list. The default value of [innodb\\_old\\_blocks\\_time \[1738\]](#) is 0, corresponding to the original behavior of moving a page to the most-recently used end of the buffer pool list when it is first accessed in the buffer pool. Increasing this value makes more and more blocks likely to age out faster from the buffer pool.

Both [innodb\\_old\\_blocks\\_pct \[1738\]](#) and [innodb\\_old\\_blocks\\_time \[1738\]](#) are dynamic, global and can be specified in the MySQL option file ([my.cnf](#) or [my.ini](#)) or changed at runtime with the [SET GLOBAL](#) command. Changing the setting requires the [SUPER](#) privilege.

To help you gauge the effect of setting these parameters, the [SHOW ENGINE INNODB STATUS](#) command reports additional statistics. The [BUFFER POOL AND MEMORY](#) section looks like:

```
Total memory allocated 1107296256; in additional pool allocated 0
Dictionary memory allocated 80360
Buffer pool size 65535
Free buffers 0
Database pages 63920
Old database pages 23600
Modified db pages 34969
Pending reads 32
Pending writes: LRU 0, flush list 0, single page 0
Pages made young 414946, not young 2930673
1274.75 youngs/s, 16521.90 non-youngs/s
Pages read 486005, created 3178, written 160585
2132.37 reads/s, 3.40 creates/s, 323.74 writes/s
Buffer pool hit rate 950 / 1000, young-making rate 30 / 1000 not 392 / 1000
Pages read ahead 1510.10/s, evicted without access 0.00/s
LRU len: 63920, unzip_LRU len: 0
I/O sum[43690]:cur[221], unzip sum[0]:cur[0]
```

- [Old database pages](#) is the number of pages in the “old” segment of the LRU list.
- [Pages made young](#) and [not young](#) is the total number of “old” pages that have been made young or not respectively.

- `young/s` and `non-young/s` is the rate at which page accesses to the “old” pages have resulted in making such pages young or otherwise respectively since the last invocation of the command.
- `young-making rate` and `not` provides the same rate but in terms of overall buffer pool accesses instead of accesses just to the “old” pages.

Because the effects of these parameters can vary widely based on your hardware configuration, your data, and the details of your workload, always benchmark to verify the effectiveness before changing these settings in any performance-critical or production environment.

In mixed workloads where most of the activity is OLTP type with periodic batch reporting queries which result in large scans, setting the value of `innodb_old_blocks_time` [1738] during the batch runs can help keep the working set of the normal workload in the buffer pool.

When scanning large tables that cannot fit entirely in the buffer pool, setting `innodb_old_blocks_pct` [1738] to a small value keeps the data that is only read once from consuming a significant portion of the buffer pool. For example, setting `innodb_old_blocks_pct=5` restricts this data that is only read once to 5% of the buffer pool.

When scanning small tables that do fit into memory, there is less overhead for moving pages around within the buffer pool, so you can leave `innodb_old_blocks_pct` [1738] at its default value, or even higher, such as `innodb_old_blocks_pct=50`.

The effect of the `innodb_old_blocks_time` [1738] parameter is harder to predict than the `innodb_old_blocks_pct` [1738] parameter, is relatively small, and varies more with the workload. To arrive at an optimal value, conduct your own benchmarks if the performance improvement from adjusting `innodb_old_blocks_pct` [1738] is not sufficient.

For more information about the InnoDB buffer pool, see [Section 8.9.1, “The InnoDB Buffer Pool”](#).

#### 14.2.4.2.25. Improvements to Crash Recovery Performance

A number of optimizations speed up certain steps of the [recovery](#) that happens on the next startup after a crash. In particular, scanning the [redo log](#) and applying the redo log are faster than in MySQL 5.1 and earlier, due to improved algorithms for memory management. You do not need to take any actions to take advantage of this performance enhancement. If you kept the size of your redo log files artificially low because recovery took a long time, you can consider increasing the file size.

For more information about InnoDB recovery, see [Section 14.2.3.13, “The InnoDB Recovery Process”](#).

#### 14.2.4.2.26. Integration with the MySQL Performance Schema

Starting with InnoDB 1.1 with MySQL 5.5, you can profile certain internal InnoDB operations using the MySQL [Performance Schema feature](#). This type of tuning is primarily for expert users, those who push the limits of MySQL performance, read the MySQL source code, and evaluate optimization strategies to overcome performance bottlenecks. DBAs can also use this feature for capacity planning, to see whether their typical workload encounters any performance bottlenecks with a particular combination of CPU, RAM, and disk storage; and if so, to judge whether performance can be improved by increasing the capacity of some part of the system.

To use this feature to examine InnoDB performance:

- You must be running MySQL 5.5 or higher. You must build the database server from source, enabling the Performance Schema feature by building with the `--with-perfschema` option. Since the Performance Schema feature introduces some performance overhead, you should use it on a test or development system rather than on a production system.
- You must be running InnoDB 1.1 or higher.
- You must be generally familiar with how to use the [Performance Schema feature](#), for example to query tables in the `performance_schema` database.



- Examine the following kinds of InnoDB objects by querying the appropriate `performance_schema` tables. The items associated with InnoDB all contain the substring `innodb` in the `EVENT_NAME` column.

For the definitions of the `*_instances` tables, see [Section 20.9.2, “Performance Schema Instance Tables”](#). For the definitions of the `*_summary_*` tables, see [Section 20.9.8, “Performance Schema Summary Tables”](#). For the definition of the `thread` table, see [Section 20.9.9, “Performance Schema Miscellaneous Tables”](#). For the definition of the `*_current_*` and `*_history_*` tables, see [Section 20.9.3, “Performance Schema Wait Event Tables”](#).

- **Mutexes** in the `mutex_instances` table. (Mutexes and RW-locks related to the InnoDB buffer pool are not included in this coverage; the same applies to the output of the `SHOW ENGINE INNODB MUTEX` command.)
- **RW-locks** in the `rwlock_instances` table.
- RW-locks in the `rwlock_instances` table.
- File I/O operations in the `file_instances`, `file_summary_by_event_name`, and `file_summary_by_instance` tables.
- **Threads** in the `PROCESSLIST` table.
- During performance testing, examine the performance data in the `events_waits_current` and `events_waits_history_long` tables. If you are interested especially in InnoDB-related objects, use the clause `WHERE EVENT_NAME LIKE '%innodb%'` to see just those entries; otherwise, examine the performance statistics for the overall MySQL server.
- You must be running MySQL 5.5, with the Performance Schema enabled by building with the `--with-perfschema` build option.

For more information about the MySQL Performance Schema, see [Chapter 20, MySQL Performance Schema](#).

#### 14.2.4.2.27. Improvements to Performance from Multiple Buffer Pools

This performance enhancement is primarily useful for people with a large **buffer pool** size, typically in the multi-gigabyte range. To take advantage of this speedup, you must set the new `innodb_buffer_pool_instances` [1706] configuration option, and you might also adjust the `innodb_buffer_pool_size` [1708] value.

When the InnoDB buffer pool is large, many data requests can be satisfied by retrieving from memory. You might encounter bottlenecks from multiple threads trying to access the buffer pool at once. Starting in InnoDB 1.1 and MySQL 5.5, you can enable multiple buffer pools to minimize this contention. Each page that is stored in or read from the buffer pool is assigned to one of the buffer pools randomly, using a hashing function. Each buffer pool manages its own free lists, flush lists, LRUs, and all other data structures connected to a buffer pool, and is protected by its own buffer pool mutex.

To enable this feature, set the `innodb_buffer_pool_instances` configuration option to a value greater than 1 (the default) up to 64 (the maximum). This option takes effect only when you set the `innodb_buffer_pool_size` to a size of 1 gigabyte or more. The total size you specify is divided among all the buffer pools. For best efficiency, specify a combination of `innodb_buffer_pool_instances` [1706] and `innodb_buffer_pool_size` [1708] so that each buffer pool instance is at least 1 gigabyte.

For more information about the InnoDB buffer pool, see [Section 8.9.1, “The InnoDB Buffer Pool”](#).

#### 14.2.4.2.28. Better Scalability with Multiple Rollback Segments

Starting in InnoDB 1.1 with MySQL 5.5, the limit on concurrent **transactions** is greatly expanded, removing a bottleneck with the InnoDB **rollback segment** that affected high-capacity systems. The limit

applies to concurrent transactions that change any data; read-only transactions do not count against that maximum.

The single rollback segment is now divided into 128 segments, each of which can support up to 1023 transactions that perform writes, for a total of approximately 128K concurrent transactions. The original transaction limit was 1023.

Each transaction is assigned to one of the rollback segments, and remains tied to that rollback segment for the duration. This enhancement improves both scalability (higher number of concurrent transactions) and performance (less contention when different transactions access the rollback segments).

To take advantage of this feature, you do not need to create any new database or tables, or reconfigure anything. You must do a [slow shutdown](#) before upgrading from MySQL 5.1 or earlier, or some time afterward. InnoDB makes the required changes inside the [system tablespace](#) automatically, the first time you restart after performing a slow shutdown.

If your workload was not constrained by the original limit of 1023 concurrent transactions, you can reduce the number of rollback segments used within a MySQL instance or within a session by setting the configuration option `innodb_rollback_segments` [1744].

For more information about performance of InnoDB under high transactional load, see [Section 8.5.2, “Optimizing InnoDB Transaction Management”](#).

#### 14.2.4.2.29. Better Scalability with Improved Purge Scheduling

The [purge](#) operations (a type of garbage collection) that InnoDB performs automatically is now done in one or more separate threads, rather than as part of the [master thread](#). This change improves scalability, because the main database operations run independently from maintenance work happening in the background.

To control this feature, increase the value of the configuration option `innodb_purge_threads=n`. If DML action is concentrated on a single table or a few tables, keep the setting low so that the threads do not contend with each other for access to the busy tables. If DML operations are spread across many tables, increase the setting. Its maximum is 32.

There is another related configuration option, `innodb_purge_batch_size` with a default of 20 and maximum of 5000. This option is mainly intended for experimentation and tuning of purge operations, and should not be interesting to typical users.

For more information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

#### 14.2.4.2.30. Improved Log Sys Mutex

This is another performance improvement that comes for free, with no user action or configuration needed. The details here are intended for performance experts who delve into the InnoDB source code, or interpret reports with keywords such as “mutex” and “log\_sys”.

The [mutex](#) known as the log sys mutex has historically done double duty, controlling access to internal data structures related to log records and the [LSN](#), as well as pages in the [buffer pool](#) that are changed when a [mini-transaction](#) is committed. Starting in InnoDB 1.1 with MySQL 5.5, these two kinds of operations are protected by separate mutexes, with a new `log_buf` mutex controlling writes to buffer pool pages due to mini-transactions.

For performance considerations for InnoDB locking operations, see [Section 8.10, “Optimizing Locking Operations”](#).

#### 14.2.4.2.31. Separate Flush List Mutex

Starting with InnoDB 1.1 with MySQL 5.5, concurrent access to the [buffer pool](#) is faster. Operations involving the [flush list](#), a data structure related to the buffer pool, are now controlled by a separate



[mutex](#) and do not block access to the buffer pool. You do not need to configure anything to take advantage of this speedup; it is fully automatic.

For more information about the InnoDB buffer pool, see [Section 8.9.1, “The InnoDB Buffer Pool”](#).

#### 14.2.4.2.32. Kernel Mutex Split

The [mutex](#) controlling concurrent access to the InnoDB kernel is now divided into separate mutexes and [rw-locks](#) to reduce contention. You do not need to configure anything to take advantage of this speedup; it is fully automatic.

#### 14.2.4.2.33. InnoDB Configurable Data Dictionary Cache

To ease the memory load on systems with huge numbers of tables, InnoDB now frees up the memory associated with an opened table, using an LRU algorithm to select tables that have gone the longest without being accessed. To reserve more memory to hold metadata for open InnoDB tables, increase the value of the [table\\_definition\\_cache \[550\]](#) configuration option. InnoDB treats this value as a “soft limit”. The actual number of tables with cached metadata could be higher, because metadata for InnoDB system tables, and parent and child tables in foreign key relationships, is never evicted from memory.

#### 14.2.4.2.34. Improved Tablespace Management

Several new features extend the file-per-table mode enabled by the [innodb\\_file\\_per\\_table \[1718\]](#) configuration option, allowing more flexibility in how the `.ibd` files are placed, exported, and restored. We characterize this as a performance enhancement because it solves the common customer request to put data from different tables onto different storage devices, for best price/performance depending on the access patterns of the data. For example, tables with high levels of random reads and writes might be placed on an [SSD](#) device, while less-often-accessed data or data processed with large batches of sequential I/O might be placed on an [HDD](#) device. See [Section 5.4.1, “Managing InnoDB Tablespaces”](#) for details.

#### 14.2.4.2.35. memcached Plugin for InnoDB

The [memcached](#) daemon is frequently used as an in-memory caching layer in front of a MySQL database server. Now MySQL allows direct access to InnoDB tables using the familiar [memcached](#) protocol and client libraries. Instead of formulating queries in SQL, you can perform simple get, set, and increment operations that avoid the performance overhead of SQL parsing and constructing a query optimization plan. You can also access the underlying InnoDB tables through SQL to load data, generate reports, or perform multi-step transactional computations.

This technique allows the data to be stored in MySQL for reliability and consistency, while coding application logic that uses the database as a fast key-value store.

This feature combines the best of both worlds:

- Data that is written using the [memcached](#) protocol is transparently written to an InnoDB table, without going through the MySQL SQL layer. You can control the frequency of writes to achieve higher raw performance when updating non-critical data.
- Data that is requested data through the [memcached](#) protocol is transparently queried from an InnoDB table, without going through the MySQL SQL layer.
- Subsequent requests for the same data will be served from the InnoDB buffer pool. The buffer pool handles the in-memory caching. You can tune the performance of data-intensive operations using the familiar InnoDB configuration options.
- InnoDB can handle composing and decomposing multiple column values into a single [memcached](#) item value, reducing the amount of string parsing and concatenation required in your application. For example, you might store a string value `2|4|6|8` in the [memcached](#) cache, and InnoDB splits that value based on a separator character, then stores the result into four numeric columns.

For details on using this NoSQL-style interface to MySQL, see [Section 14.2.9, “InnoDB Integration with memcached”](#). For additional background on `memcached` and considerations for writing applications for its API, see [Section 15.6, “Using MySQL with memcached”](#).

#### 14.2.4.2.36. Online DDL

This feature is a continuation of the “Fast Index Creation” feature introduced in [Fast Index Creation in the InnoDB Storage Engine](#). Now you can perform other kinds of DDL operations on InnoDB tables [online](#): that is, with minimal delay for operations on that table, without rebuilding the entire table, or both. This enhancement improves responsiveness and availability in busy production environments, where making a table unavailable for minutes or hours whenever its column definitions change is not practical.

For full details, see [Section 5.5, “Online DDL for InnoDB Tables”](#).

The DDL operations enhanced by this feature are these variations on the `ALTER TABLE` statement:

- Create [secondary indexes](#): `CREATE INDEX name ON table (col_list)` or `ALTER TABLE table ADD INDEX name (col_list)`. (Creating a primary key or a `FULLTEXT` index still requires locking the table.)

Drop [secondary indexes](#): `DROP INDEX name ON table;` or `ALTER TABLE table DROP INDEX name`

Creating and dropping secondary indexes on InnoDB tables has avoided the table-copying behavior since the days of MySQL 5.1 with the InnoDB Plugin. Now, the table remains available for read and write operations while the index is being created or dropped. The `CREATE TABLE` or `DROP TABLE` statement only finishes after all transactions that are modifying the table are completed, so that the initial state of the index reflects the most recent contents of the table.

Previously, modifying the table while an index was being created or dropped typically resulted in a [deadlock](#) that cancelled the insert, update, or delete statement on the table.

- Changing the [auto-increment](#) value for a column: `ALTER TABLE table AUTO_INCREMENT=next_value;`

Especially in a distributed system using replication or sharding, you sometimes reset the auto-increment counter for a table to a specific value. The next row inserted into the table uses the specified value for its auto-increment column. You might also use this technique in a data warehousing environment where you periodically empty all the tables and reload them, and you can restart the auto-increment sequence from 1.

- Adding or dropping a [foreign key constraint](#):

```
ALTER TABLE tbl1 ADD CONSTRAINT fk_name FOREIGN KEY index (col1) REFERENCES tbl2(col2) referential_action;
ALTER TABLE tbl DROP FOREIGN KEY fk_name;
```

Dropping a foreign key can be performed online with the `foreign_key_checks [483]` option enabled or disabled. Creating a foreign key online requires `foreign_key_checks [483]` to be disabled.

If you do not know the names of the foreign key constraints on a particular table, issue the following statement and find the constraint name in the `CONSTRAINT` clause for each foreign key:

```
show create table table\G
```

Or, query the `information_schema.table_constraints` table and use the `constraint_name` and `constraint_type` columns to identify the foreign key names.

As a consequence of this enhancement, you can now also drop a foreign key and its associated index in a single statement, which previously required separate statements in a strict order:

```
ALTER TABLE table DROP FOREIGN KEY constraint, DROP INDEX index;
```

- Renaming a column: `ALTER TABLE tbl CHANGE old_col_name new_col_name datatype`

When you keep the same data type and only change the column name, this operation can always be performed online. As part of this enhancement, you can now rename a column that is part of a foreign key constraint, which was not allowed before.

- Some other `ALTER TABLE` operations are non-blocking, and are faster than before because the table-copying operation is optimized, even though a table copy is still required:
  - Changing the `ROW_FORMAT` or `KEY_BLOCK_SIZE` properties for a table.
  - Changing the nullable status for a column.
  - Adding, dropping, or reordering columns.



#### Note

As your database schema evolves with new columns, data types, constraints, indexes, and so on, keep your `CREATE TABLE` statements up to date with the latest table definitions. Even with the performance improvements of online DDL, it is more efficient to create stable database structures at the beginning, rather than creating part of the schema and then issuing `ALTER TABLE` statements afterward.

The main exception to this guideline is for secondary indexes on tables with large numbers of rows. It is typically most efficient to create the table with all details specified except the secondary indexes, load the data, then create the secondary indexes.

Whatever sequence of `CREATE TABLE`, `CREATE INDEX`, `ALTER TABLE`, and similar statements went into putting a table together, you can capture the SQL needed to reconstruct the current form of the table by issuing the statement `SHOW CREATE TABLE table\G` (uppercase `\G` required for tidy formatting). This output shows clauses such as numeric precision, `NOT NULL`, and `CHARACTER SET` that are sometimes added behind the scenes, and you might otherwise leave out when cloning the table on a new system or setting up foreign key columns with identical type.

### 14.2.4.3. InnoDB INFORMATION\_SCHEMA tables

The `INFORMATION_SCHEMA` is a MySQL feature that helps you monitor server activity to diagnose capacity and performance issues. Several InnoDB-related `INFORMATION_SCHEMA` tables (`INNODB_CMP`, `INNODB_CMP_RESET`, `INNODB_CMPMEM`, `INNODB_CMPMEM_RESET`, `INNODB_TRX`, `INNODB_LOCKS` and `INNODB_LOCK_WAITS`) contain live information about compressed InnoDB tables, the compressed InnoDB buffer pool, all transactions currently executing inside InnoDB, the locks that transactions hold and those that are blocking transactions waiting for access to a resource (a table or row).

This section describes the InnoDB-related Information Schema tables and shows some examples of their use.

#### 14.2.4.3.1. Information Schema Tables about Compression

Two new pairs of Information Schema tables can give you some insight into how well compression is working overall. One pair of tables contains information about the number of compression operations and the amount of time spent performing compression. Another pair of tables contains information on the way memory is allocated for compression.

##### 14.2.4.3.1.1. `INNODB_CMP` and `INNODB_CMP_RESET`

The tables `INNODB_CMP` and `INNODB_CMP_RESET` contain status information on the operations related to compressed tables, which are covered in [Section 5.4.6, “Working with InnoDB Compressed Tables”](#). The compressed [page size](#) is in the column `PAGE_SIZE`.

These two tables have identical contents, but reading from `INNODB_CMP_RESET` resets the statistics on compression and uncompression operations. For example, if you archive the output of `INNODB_CMP_RESET` every 60 minutes, you see the statistics for each hourly period. If you monitor the output of `INNODB_CMP` (making sure never to read `INNODB_CMP_RESET`), you see the cumulated statistics since InnoDB was started.

For the table definition, see [Table 19.1, “Columns of INNODB\\_CMP and INNODB\\_CMP\\_RESET”](#).

#### 14.2.4.3.1.2. INNODB\_CMPMEM and INNODB\_CMPMEM\_RESET

The tables `INNODB_CMPMEM` and `INNODB_CMPMEM_RESET` contain status information on the compressed pages that reside in the buffer pool. Please consult [Section 5.4.6, “Working with InnoDB Compressed Tables”](#) for further information on compressed tables and the use of the buffer pool. The tables `INNODB_CMP` and `INNODB_CMP_RESET` should provide more useful statistics on compression.

#### Internal Details

InnoDB uses a [buddy allocator](#) system to manage memory allocated to [pages of various sizes](#), from 1KB to 16KB. Each row of the two tables described here corresponds to a single page size.

These two tables have identical contents, but reading from `INNODB_CMPMEM_RESET` resets the statistics on relocation operations. For example, if every 60 minutes you archived the output of `INNODB_CMPMEM_RESET`, it would show the hourly statistics. If you never read `INNODB_CMPMEM_RESET` and monitored the output of `INNODB_CMPMEM` instead, it would show the cumulated statistics since InnoDB was started.

For the table definition, see [Table 19.3, “Columns of INNODB\\_CMPMEM and INNODB\\_CMPMEM\\_RESET”](#).

#### 14.2.4.3.1.3. Using the Compression Information Schema Tables

##### Example 14.2. Using the Compression Information Schema Tables

The following is sample output from a database that contains compressed tables (see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#), `INNODB_CMP`, `INNODB_CMP_PER_INDEX`, and `INNODB_CMPMEM`).

The following table shows the contents of `INFORMATION_SCHEMA.INNODB_CMP` under a light [workload](#). The only compressed page size that the buffer pool contains is 8K. Compressing or uncompressing pages has consumed less than a second since the time the statistics were reset, because the columns `COMPRESS_TIME` and `UNCOMPRESS_TIME` are zero.

page size	compress ops	compress ops ok	compress time	uncompress ops	uncompress time
1024	0	0	0	0	0
2048	0	0	0	0	0
4096	0	0	0	0	0
8192	1048	921	0	61	0
16384	0	0	0	0	0

According to `INNODB_CMPMEM`, there are 6169 compressed 8KB pages in the [buffer pool](#). The only other allocated block size is 64 bytes. The smallest `PAGE_SIZE` in `INNODB_CMPMEM` is used for block descriptors of those compressed pages for which no uncompressed page exists in the buffer pool. We see that there are 5910 such pages. Indirectly, we see that 259 (6169-5910) compressed pages also exist in the buffer pool in uncompressed form.

The following table shows the contents of `INFORMATION_SCHEMA.INNODB_CMPMEM` under a light workload. Some memory is unusable due to fragmentation of the memory allocator for compressed pages:  $\text{SUM}(\text{PAGE\_SIZE} * \text{PAGES\_FREE}) = 6784$ . This is because small memory allocation requests are fulfilled by splitting bigger blocks, starting from the 16K blocks that are allocated from the main buffer pool, using the buddy allocation system. The fragmentation is this low because some allocated blocks have been relocated (copied) to form bigger adjacent free blocks. This copying of  $\text{SUM}(\text{PAGE\_SIZE} * \text{RELOCATION\_OPS})$  bytes has consumed less than a second ( $\text{SUM}(\text{RELOCATION\_TIME}) = 0$ ).

page size	pages used	pages free	relocation ops	relocation time
64	5910	0	2436	0
128	0	1	0	0
256	0	0	0	0
512	0	1	0	0
1024	0	0	0	0
2048	0	1	0	0
4096	0	1	0	0
8192	6169	0	5	0
16384	0	0	0	0

#### 14.2.4.3.2. Information Schema Tables about Transactions

Three InnoDB-related Information Schema tables make it easy to monitor transactions and diagnose possible locking problems. The three tables are `INNODB_TRX`, `INNODB_LOCKS` and `INNODB_LOCK_WAITS`.

##### 14.2.4.3.2.1. `INNODB_TRX`

Contains information about every transaction currently executing inside InnoDB, including whether the transaction is waiting for a lock, when the transaction started, and the particular SQL statement the transaction is executing.

For the table definition, see [Table 19.4, “INNODB\\_TRX Columns”](#).

##### 14.2.4.3.2.2. `INNODB_LOCKS`

Each transaction in InnoDB that is waiting for another transaction to release a lock (`INNODB_TRX.TRX_STATE = 'LOCK WAIT'`) is blocked by exactly one “blocking lock request”. That blocking lock request is for a row or table lock held by another transaction in an incompatible mode. The waiting or blocked transaction cannot proceed until the other transaction commits or rolls back, thereby releasing the requested lock. For every blocked transaction, `INNODB_LOCKS` contains one row that describes each lock the transaction has requested, and for which it is waiting. `INNODB_LOCKS` also contains one row for each lock that is blocking another transaction, whatever the state of the transaction that holds the lock (`'RUNNING'`, `'LOCK WAIT'`, `'ROLLING BACK'` or `'COMMITTING'`). The lock that is blocking a transaction is always held in a mode (read vs. write, shared vs. exclusive) incompatible with the mode of requested lock.

For the table definition, see [Table 19.5, “INNODB\\_LOCKS Columns”](#).

##### 14.2.4.3.2.3. `INNODB_LOCK_WAITS`

Using this table, you can tell which transactions are waiting for a given lock, or for which lock a given transaction is waiting. This table contains one or more rows for each *blocked* transaction, indicating the lock it has requested and the lock(s) that is (are) blocking that request. The `REQUESTED_LOCK_ID` refers to the lock that a transaction is requesting, and the `BLOCKING_LOCK_ID` refers to the lock (held by another transaction) that is preventing the first transaction from proceeding. For any given blocked

transaction, all rows in `INNODB_LOCK_WAITS` have the same value for `REQUESTED_LOCK_ID` and different values for `BLOCKING_LOCK_ID`.

For the table definition, see [Table 19.6, “INNODB\\_LOCK\\_WAITS Columns”](#).

#### 14.2.4.3.2.4. Using the Transaction Information Schema Tables

##### Example 14.3. Identifying Blocking Transactions

It is sometimes helpful to be able to identify which transaction is blocking another. You can use the Information Schema tables to find out which transaction is waiting for another, and which resource is being requested.

Suppose you have the following scenario, with three users running concurrently. Each user (or session) corresponds to a MySQL thread, and executes one transaction after another. Consider the state of the system when these users have issued the following commands, but none has yet committed its transaction:

- User A:

```
BEGIN;
SELECT a FROM t FOR UPDATE;
SELECT SLEEP(100);
```

- User B:

```
SELECT b FROM t FOR UPDATE;
```

- User C:

```
SELECT c FROM t FOR UPDATE;
```

In this scenario, you can use this query to see who is waiting for whom:

```
SELECT r.trx_id waiting_trx_id,
       r.trx_mysql_thread_id waiting_thread,
       r.trx_query waiting_query,
       b.trx_id blocking_trx_id,
       b.trx_mysql_thread_id blocking_thread,
       b.trx_query blocking_query
FROM   information_schema.innodb_lock_waits w
INNER JOIN information_schema.innodb_trx b ON
        b.trx_id = w.blocking_trx_id
INNER JOIN information_schema.innodb_trx r ON
        r.trx_id = w.requesting_trx_id;
```

waiting trx id	waiting thread	waiting query	blocking trx id	blocking thread	blocking query
A4	6	SELECT b FROM t FOR UPDATE	A3	5	SELECT SLEEP(100)
A5	7	SELECT c FROM t FOR UPDATE	A3	5	SELECT SLEEP(100)
A5	7	SELECT c FROM t FOR UPDATE	A4	6	SELECT b FROM t FOR UPDATE

In the above result, you can identify users by the “waiting query” or “blocking query”. As you can see:

- User B (trx id 'A4', thread 6) and User C (trx id 'A5', thread 7) are both waiting for User A (trx id 'A3', thread 5).
- User C is waiting for User B as well as User A.

You can see the underlying data in the tables `INNODB_TRX`, `INNODB_LOCKS`, and `INNODB_LOCK_WAITS`.



The following table shows some sample contents of `INFORMATION_SCHEMA.INNODB_TRX`.

trx id	trx state	trx started	trx requested lock id	trx wait started	trx weight	trx mysql thread id	trx query
A3	RUNNING	2008-01-15 16:44:54	NULL	NULL	2	5	SELECT SLEEP(100)
A4	LOCK WAIT	2008-01-15 16:45:09	A4:1:3:2	2008-01-15 16:45:09	2	6	SELECT b FROM t FOR UPDATE
A5	LOCK WAIT	2008-01-15 16:45:14	A5:1:3:2	2008-01-15 16:45:14	2	7	SELECT c FROM t FOR UPDATE

The following table shows some sample contents of `INFORMATION_SCHEMA.INNODB_LOCKS`.

lock id	lock trx id	lock mode	lock type	lock table	lock index	lock space	lock page	lock rec	lock data
A3:1:3:2	A3	X	RECORD	`test`.`t`	`PRIMARY`	1	3	2	0x0200
A4:1:3:2	A4	X	RECORD	`test`.`t`	`PRIMARY`	1	3	2	0x0200
A5:1:3:2	A5	X	RECORD	`test`.`t`	`PRIMARY`	1	3	2	0x0200

The following table shows some sample contents of `INFORMATION_SCHEMA.INNODB_LOCK_WAITS`.

requesting trx id	requested lock id	blocking trx id	blocking lock id
A4	A4:1:3:2	A3	A3:1:3:2
A5	A5:1:3:2	A3	A3:1:3:2
A5	A5:1:3:2	A4	A4:1:3:2

#### Example 14.4. More Complex Example of Transaction Data in Information Schema Tables

Sometimes you would like to correlate the internal InnoDB locking information with session-level information maintained by MySQL. For example, you might like to know, for a given InnoDB transaction ID, the corresponding MySQL session ID and name of the user that may be holding a lock, and thus blocking another transaction.

The following output from the `INFORMATION_SCHEMA` tables is taken from a somewhat loaded system.

As can be seen in the following tables, there are several transactions running.

The following `INNODB_LOCKS` and `INNODB_LOCK_WAITS` tables shows that:

- Transaction `77F` (executing an `INSERT`) is waiting for transactions `77E`, `77D` and `77B` to commit.
- Transaction `77E` (executing an `INSERT`) is waiting for transactions `77D` and `77B` to commit.
- Transaction `77D` (executing an `INSERT`) is waiting for transaction `77B` to commit.
- Transaction `77B` (executing an `INSERT`) is waiting for transaction `77A` to commit.
- Transaction `77A` is running, currently executing `SELECT`.
- Transaction `E56` (executing an `INSERT`) is waiting for transaction `E55` to commit.
- Transaction `E55` (executing an `INSERT`) is waiting for transaction `19C` to commit.
- Transaction `19C` is running, currently executing an `INSERT`.

Note that there may be an inconsistency between queries shown in the two tables `INNODB_TRX.TRX_QUERY` and `PROCESSLIST.INFO`. The current transaction ID for a thread, and the query being executed in that transaction, may be different in these two tables for any given thread. See [Section 14.2.4.3.4.3, “Possible Inconsistency with `PROCESSLIST`”](#) for an explanation.



The following table shows the contents of `INFORMATION_SCHEMA.PROCESSLIST` in a system running a heavy [workload](#).

ID	USER	HOST	DB	COMMAND	TIME	STATE	INFO
384	root	localhost	test	Query	10	update	insert into t2 values ...
257	root	localhost	test	Query	3	update	insert into t2 values ...
130	root	localhost	test	Query	0	update	insert into t2 values ...
61	root	localhost	test	Query	1	update	insert into t2 values ...
8	root	localhost	test	Query	1	update	insert into t2 values ...
4	root	localhost	test	Query	0	preparing	SELECT * FROM processlist
2	root	localhost	test	Sleep	566		NULL

The following table shows the contents of `INFORMATION_SCHEMA.INNODB_TRX` in a system running a heavy [workload](#).

trx id	trx state	trx started	trx requested lock id	trx wait started	trx weight	trx mysql thread id	trx query
77F	LOCK WAIT	2008-01-15 13:10:16	77F:806	2008-01-15 13:10:16	1	876	insert into t09 (D, B, C) values ...
77E	LOCK WAIT	2008-01-15 13:10:16	77E:806	2008-01-15 13:10:16	1	875	insert into t09 (D, B, C) values ...
77D	LOCK WAIT	2008-01-15 13:10:16	77D:806	2008-01-15 13:10:16	1	874	insert into t09 (D, B, C) values ...
77B	LOCK WAIT	2008-01-15 13:10:16	77B:733:12:1	2008-01-15 13:10:16	4	873	insert into t09 (D, B, C) values ...
77A	RUNNING	2008-01-15 13:10:16	NULL	NULL	4	872	select b, c from t09 where ...
E56	LOCK WAIT	2008-01-15 13:10:06	E56:743:6:2	2008-01-15 13:10:06	5	384	insert into t2 values ...
E55	LOCK WAIT	2008-01-15 13:10:06	E55:743:38:2	2008-01-15 13:10:13	965	257	insert into t2 values ...
19C	RUNNING	2008-01-15 13:09:10	NULL	NULL	2900	130	insert into t2 values ...
E15	RUNNING	2008-01-15 13:08:59	NULL	NULL	5395	61	insert into t2 values ...
51D	RUNNING	2008-01-15 13:08:47	NULL	NULL	9807	8	insert into t2 values ...

The following table shows the contents of `INFORMATION_SCHEMA.INNODB_LOCK_WAITS` in a system running a heavy [workload](#).

requesting trx id	requested lock id	blocking trx id	blocking lock id
77F	77F:806	77E	77E:806
77F	77F:806	77D	77D:806
77F	77F:806	77B	77B:806
77E	77E:806	77D	77D:806
77E	77E:806	77B	77B:806
77D	77D:806	77B	77B:806
77B	77B:733:12:1	77A	77A:733:12:1
E56	E56:743:6:2	E55	E55:743:6:2
E55	E55:743:38:2	19C	19C:743:38:2

The following table shows the contents of `INFORMATION_SCHEMA.INNODB_LOCKS` in a system running a heavy [workload](#).

lock id	lock trx id	lock mode	lock type	lock table	lock index	lock space	lock page	lock rec	lock data
77F:806	77F	AUTO_INC	TABLE	`test` .`t09`	NULL	NULL	NULL	NULL	NULL
77E:806	77E	AUTO_INC	TABLE	`test` .`t09`	NULL	NULL	NULL	NULL	NULL
77D:806	77D	AUTO_INC	TABLE	`test` .`t09`	NULL	NULL	NULL	NULL	NULL
77B:806	77B	AUTO_INC	TABLE	`test` .`t09`	NULL	NULL	NULL	NULL	NULL
77B:733:12:1	77B	X	RECORD	`test` .`t09`	`PRIMARY`	733	12	1	supremum pseudo- record
77A:733:12:1	77A	X	RECORD	`test` .`t09`	`PRIMARY`	733	12	1	supremum pseudo- record
E56:743:6:2	E56	S	RECORD	`test` .`t2`	`PRIMARY`	743	6	2	0, 0
E55:743:6:2	E55	X	RECORD	`test` .`t2`	`PRIMARY`	743	6	2	0, 0
E55:743:38:2	E55	S	RECORD	`test` .`t2`	`PRIMARY`	743	38	2	1922, 1922
19C:743:38:2	19C	X	RECORD	`test` .`t2`	`PRIMARY`	743	38	2	1922, 1922

#### 14.2.4.3.3. Information Schema Tables about Full-Text Search

A set of related `INFORMATION_SCHEMA` tables contains information about `FULLTEXT` search indexes on InnoDB tables:

- `INNODB_FT_CONFIG`
- `INNODB_FT_INDEX_TABLE`
- `INNODB_FT_INDEX_CACHE`
- `INNODB_FT_DEFAULT_STOPWORD`

- `INNODB_FT_DELETED`
- `INNODB_FT_BEING_DELETED`

#### 14.2.4.3.4. Special Locking Considerations for InnoDB `INFORMATION_SCHEMA` Tables

##### 14.2.4.3.4.1. Understanding InnoDB Locking

When a transaction updates a row in a table, or locks it with `SELECT FOR UPDATE`, InnoDB establishes a list or queue of locks on that row. Similarly, InnoDB maintains a list of locks on a table for table-level locks transactions hold. If a second transaction wants to update a row or lock a table already locked by a prior transaction in an incompatible mode, InnoDB adds a lock request for the row to the corresponding queue. For a lock to be acquired by a transaction, all incompatible lock requests previously entered into the lock queue for that row or table must be removed (the transactions holding or requesting those locks either commit or roll back).

A transaction may have any number of lock requests for different rows or tables. At any given time, a transaction may be requesting a lock that is held by another transaction, in which case it is blocked by that other transaction. The requesting transaction must wait for the transaction that holds the blocking lock to commit or rollback. If a transaction is not waiting for a lock, it is in the `'RUNNING'` state. If a transaction is waiting for a lock, it is in the `'LOCK WAIT'` state.

The table `INNODB_LOCKS` holds one or more row for each `'LOCK WAIT'` transaction, indicating the lock request(s) that is (are) preventing its progress. This table also contains one row describing each lock in a queue of locks pending for a given row or table. The table `INNODB_LOCK_WAITS` shows which locks already held by a transaction are blocking locks requested by other transactions.

##### 14.2.4.3.4.2. Granularity of `INFORMATION_SCHEMA` Data

The data exposed by the transaction and locking tables represent a glimpse into fast-changing data. This is not like other (user) tables, where the data only changes when application-initiated updates occur. The underlying data is internal system-managed data, and can change very quickly.

For performance reasons, and to minimize the chance of misleading `JOINS` between the `INFORMATION_SCHEMA` tables, InnoDB collects the required transaction and locking information into an intermediate buffer whenever a `SELECT` on any of the tables is issued. This buffer is refreshed only if more than 0.1 seconds has elapsed since the last time the buffer was read. The data needed to fill the three tables is fetched atomically and consistently and is saved in this global internal buffer, forming a point-in-time “snapshot”. If multiple table accesses occur within 0.1 seconds (as they almost certainly do when MySQL processes a join among these tables), then the same snapshot is used to satisfy the query.

A correct result is returned when you `JOIN` any of these tables together in a single query, because the data for the three tables comes from the same snapshot. Because the buffer is not refreshed with every query of any of these tables, if you issue separate queries against these tables within a tenth of a second, the results are the same from query to query. On the other hand, two separate queries of the same or different tables issued more than a tenth of a second apart may see different results, since the data come from different snapshots.

Because InnoDB must temporarily stall while the transaction and locking data is collected, too frequent queries of these tables can negatively impact performance as seen by other users.

As these tables contain sensitive information (at least `INNODB_LOCKS.LOCK_DATA` and `INNODB_TRX.TRX_QUERY`), for security reasons, only the users with the `PROCESS` privilege are allowed to `SELECT` from them.

##### 14.2.4.3.4.3. Possible Inconsistency with `PROCESSLIST`

As just described, while the transaction and locking data is correct and consistent when these `INFORMATION_SCHEMA` tables are populated, the underlying data changes so fast that similar glimpses at other, similarly fast-changing data, may not be in sync. Thus, you should be careful

in comparing the data in the InnoDB transaction and locking tables with that in the [MySQL table](#) `PROCESSLIST`. The data from the `PROCESSLIST` table does not come from the same snapshot as the data about locking and transactions. Even if you issue a single `SELECT` (Joining `INNODB_TRX` and `PROCESSLIST`, for example), the content of those tables is generally not consistent. `INNODB_TRX` may reference rows that are not present in `PROCESSLIST` or the currently executing SQL query of a transaction, shown in `INNODB_TRX.TRX_QUERY` may be different from the one in `PROCESSLIST.INFO`. The query in `INNODB_TRX` is always consistent with the rest of `INNODB_TRX`, `INNODB_LOCKS` and `INNODB_LOCK_WAITS` when the data comes from the same snapshot.

#### 14.2.4.4. `SHOW ENGINE INNODB STATUS` and the InnoDB Monitors

InnoDB Monitors provide information about the InnoDB internal state. This information is useful for performance tuning. Each Monitor can be enabled by creating a table with a special name, which causes InnoDB to write Monitor output periodically. Also, output for the standard InnoDB Monitor is available on demand through the `SHOW ENGINE INNODB STATUS` SQL statement.

There are several types of InnoDB Monitors:

- The standard InnoDB Monitor displays the following types of information:
  - Table and record locks held by each active transaction.
  - Lock waits of a transaction.
  - Semaphore waits of threads.
  - Pending file I/O requests.
  - Buffer pool statistics.
  - Purge and insert buffer merge activity of the main InnoDB thread.

For a discussion of InnoDB lock modes, see [Section 14.2.3.2, “InnoDB Lock Modes”](#).

To enable the standard InnoDB Monitor for periodic output, create a table named `innodb_monitor`. To obtain Monitor output on demand, use the `SHOW ENGINE INNODB STATUS` SQL statement to fetch the output to your client program. If you are using the `mysql` interactive client, the output is more readable if you replace the usual semicolon statement terminator with `\G`:

```
mysql> SHOW ENGINE INNODB STATUS\G
```

- The InnoDB Lock Monitor is like the standard Monitor but also provides extensive lock information. To enable this Monitor for periodic output, create a table named `innodb_lock_monitor`.
- The InnoDB Tablespace Monitor prints a list of file segments in the shared tablespace and validates the tablespace allocation data structures. To enable this Monitor for periodic output, create a table named `innodb_tablespace_monitor`.
- The InnoDB Table Monitor prints the contents of the InnoDB internal data dictionary. To enable this Monitor for periodic output, create a table named `innodb_table_monitor`.



#### Note

As of MySQL 5.6.3, the `innodb_table_monitor` table is deprecated and will be removed in a future MySQL release. Similar information can be obtained from `INNODB INFORMATION_SCHEMA` tables. See [Section 19.30, “INFORMATION\\_SCHEMA Tables for InnoDB”](#).

To enable an InnoDB Monitor for periodic output, use a `CREATE TABLE` statement to create the table associated with the Monitor. For example, to enable the standard InnoDB Monitor, create the `innodb_monitor` table:

```
CREATE TABLE innodb_monitor (a INT) ENGINE=INNODB;
```

To stop the Monitor, drop the table:

```
DROP TABLE innodb_monitor;
```

The `CREATE TABLE` syntax is just a way to pass a command to the InnoDB engine through MySQL's SQL parser: The only things that matter are the table name `innodb_monitor` and that it be an InnoDB table. The structure of the table is not relevant at all for the InnoDB Monitor. If you shut down the server, the Monitor does not restart automatically when you restart the server. Drop the Monitor table and issue a new `CREATE TABLE` statement to start the Monitor. (This syntax may change in a future release.)

The `PROCESS` [755] privilege is required to start or stop the InnoDB Monitor tables.

When you enable InnoDB Monitors for periodic output, InnoDB writes their output to the `mysqld` server standard error output (`stderr`). In this case, no output is sent to clients. When switched on, InnoDB Monitors print data about every 15 seconds. Server output usually is directed to the error log (see Section 5.2.2, “The Error Log”). This data is useful in performance tuning. On Windows, start the server from a command prompt in a console window with the `--console` [417] option if you want to direct the output to the window rather than to the error log.

InnoDB sends diagnostic output to `stderr` or to files rather than to `stdout` or fixed-size memory buffers, to avoid potential buffer overflows. As a side effect, the output of `SHOW ENGINE INNODB STATUS` is written to a status file in the MySQL data directory every fifteen seconds. The name of the file is `innodb_status.pid`, where `pid` is the server process ID. InnoDB removes the file for a normal shutdown. If abnormal shutdowns have occurred, instances of these status files may be present and must be removed manually. Before removing them, you might want to examine them to see whether they contain useful information about the cause of abnormal shutdowns. The `innodb_status.pid` file is created only if the configuration option `innodb-status-file=1` [1698] is set.

InnoDB Monitors should be enabled only when you actually want to see Monitor information because output generation does result in some performance decrement. Also, if you enable monitor output by creating the associated table, your error log may become quite large if you forget to remove the table later.

For additional information about InnoDB monitors, see:

- Mark Leith: [InnoDB Table and Tablespace Monitors](#)

Each monitor begins with a header containing a timestamp and the monitor name. For example:

```
=====
090407 12:06:19 INNODB TABLESPACE MONITOR OUTPUT
=====
```

The header for the standard Monitor (`INNODB MONITOR OUTPUT`) is also used for the Lock Monitor because the latter produces the same output with the addition of extra lock information.

The following sections describe the output for each Monitor.

#### 14.2.4.4.1. InnoDB Standard Monitor and Lock Monitor Output

The Lock Monitor is the same as the standard Monitor except that it includes additional lock information. Enabling either monitor for periodic output by creating the associated InnoDB table turns on the same output stream, but the stream includes the extra information if the Lock Monitor is enabled. For example, if you create the `innodb_monitor` and `innodb_lock_monitor` tables, that turns on a single output stream. The stream includes extra lock information until you disable the Lock Monitor by removing the `innodb_lock_monitor` table.

## Example InnoDB Monitor output:

```

mysql> SHOW ENGINE INNODB STATUS\G
***** 1. row *****
Status:
=====
030709 13:00:59 INNODB MONITOR OUTPUT
=====
Per second averages calculated from the last 18 seconds
-----
BACKGROUND THREAD
-----
srv_master_thread loops: 53 1_second, 44 sleeps, 5 10_second, 7 background,
  7 flush
srv_master_thread log flush and writes: 48
-----
SEMAPHORES
-----
OS WAIT ARRAY INFO: reservation count 413452, signal count 378357
--Thread 32782 has waited at btr0sea.c line 1477 for 0.00 seconds the
semaphore: X-lock on RW-latch at 41a28668 created in file btr0sea.c line 135
a writer (thread id 32782) has reserved it in mode wait exclusive
number of readers 1, waiters flag 1
Last time read locked in file btr0sea.c line 731
Last time write locked in file btr0sea.c line 1347
Mutex spin waits 0, rounds 0, OS waits 0
RW-shared spins 2, rounds 60, OS waits 2
RW-excl spins 0, rounds 0, OS waits 0
Spin rounds per wait: 0.00 mutex, 20.00 RW-shared, 0.00 RW-excl
-----
LATEST FOREIGN KEY ERROR
-----
030709 13:00:59 Transaction:
TRANSACTION 0 290328284, ACTIVE 0 sec, process no 3195
inserting
15 lock struct(s), heap size 2496, undo log entries 9
MySQL thread id 25, query id 4668733 localhost heikki update
insert into ibtest11a (D, B, C) values (5, 'khDk' , 'khDk')
Foreign key constraint fails for table test/ibtest11a:
'
  CONSTRAINT `0_219242` FOREIGN KEY (`A`, `D`) REFERENCES `ibtest11b` (`A`,
  `D`) ON DELETE CASCADE ON UPDATE CASCADE
Trying to add in child table, in index PRIMARY tuple:
  0: len 4; hex 80000101; asc ....; 1: len 4; hex 80000005; asc ....; 2:
  len 4; hex 6b68446b; asc khDk; 3: len 6; hex 0000114e0edc; asc ...N...; 4:
  len 7; hex 00000000c3e0a7; asc .....; 5: len 4; hex 6b68446b; asc khDk;
But in parent table test/ibtest11b, in index PRIMARY,
the closest match we can find is record:
RECORD: info bits 0 0: len 4; hex 8000015b; asc ...[; 1: len 4; hex
80000005; asc ....; 2: len 3; hex 6b6864; asc khD; 3: len 6; hex
0000111ef3eb; asc .....; 4: len 7; hex 800001001e0084; asc .....; 5:
len 3; hex 6b6864; asc khD;
-----
LATEST DETECTED DEADLOCK
-----
030709 12:59:58
*** (1) TRANSACTION:
TRANSACTION 0 290252780, ACTIVE 1 sec, process no 3185
inserting
LOCK WAIT 3 lock struct(s), heap size 320, undo log entries 146
MySQL thread id 21, query id 4553379 localhost heikki update
INSERT INTO alex1 VALUES(86, 86, 794, 'aA35818', 'bb', 'c79166', 'd4766t',
'e187358f', 'g84586', 'h794', date_format('2001-04-03 12:54:22', '%Y-%m-%d
%H:%i'), 7
*** (1) WAITING FOR THIS LOCK TO BE GRANTED:
RECORD LOCKS space id 0 page no 48310 n bits 568 table test/alex1 index
symbole trx id 0 290252780 lock mode S waiting
Record lock, heap no 324 RECORD: info bits 0 0: len 7; hex 61613335383138;
asc aa35818; 1:
*** (2) TRANSACTION:
TRANSACTION 0 290251546, ACTIVE 2 sec, process no 3190
inserting

```





```

I/O thread 0 state: waiting for i/o request (insert buffer thread)
I/O thread 1 state: waiting for i/o request (log thread)
I/O thread 2 state: waiting for i/o request (read thread)
I/O thread 3 state: waiting for i/o request (write thread)
Pending normal aio reads: 0, aio writes: 0,
  ibuf aio reads: 0, log i/o's: 0, sync i/o's: 0
Pending flushes (fsync) log: 0; buffer pool: 0
151671 OS file reads, 94747 OS file writes, 8750 OS fsyncs
25.44 reads/s, 18494 avg bytes/read, 17.55 writes/s, 2.33 fsyncs/s
-----
INSERT BUFFER AND ADAPTIVE HASH INDEX
-----
Ibuf for space 0: size 1, free list len 19, seg size 21,
85004 inserts, 85004 merged recs, 26669 merges
Hash table size 207619, used cells 14461, node heap has 16 buffer(s)
1877.67 hash searches/s, 5121.10 non-hash searches/s
---
LOG
---
Log sequence number 18 1212842764
Log flushed up to 18 1212665295
Last checkpoint at 18 1135877290
0 pending log writes, 0 pending chkp writes
4341 log i/o's done, 1.22 log i/o's/second
-----
BUFFER POOL AND MEMORY
-----
Total memory allocated 84966343; in additional pool allocated 1402624
Buffer pool size 3200
Free buffers 110
Database pages 3074
Modified db pages 2674
Pending reads 0
Pending writes: LRU 0, flush list 0, single page 0
Pages read 171380, created 51968, written 194688
28.72 reads/s, 20.72 creates/s, 47.55 writes/s
Buffer pool hit rate 999 / 1000
-----
ROW OPERATIONS
-----
0 queries inside InnoDB, 0 queries in queue
Main thread process no. 3004, id 7176, state: purging
Number of rows inserted 3738558, updated 127415, deleted 33707, read 755779
1586.13 inserts/s, 50.89 updates/s, 28.44 deletes/s, 107.88 reads/s
-----
END OF INNODB MONITOR OUTPUT
=====

```

InnoDB Monitor output is limited to 1MB when produced using the `SHOW ENGINE INNODB STATUS` statement. This limit does not apply to output written to the server's error output.

Some notes on the output sections:

#### BACKGROUND THREAD

The `srv_master_thread` lines shows work done by the main background thread.

#### SEMAPHORES

This section reports threads waiting for a semaphore and statistics on how many times threads have needed a spin or a wait on a mutex or a rw-lock semaphore. A large number of threads waiting for semaphores may be a result of disk I/O, or contention problems inside InnoDB. Contention can be due to heavy parallelism of queries or problems in operating system thread scheduling. Setting the `innodb_thread_concurrency [1751]` system variable smaller than the default value might help in such situations. The `Spin rounds per wait` line shows the number of spinlock rounds per OS wait for a mutex.

#### LATEST FOREIGN KEY ERROR

This section provides information about the most recent foreign key constraint error. It is not present if no such error has occurred. The contents include the statement that failed as well as information about the constraint that failed and the referenced and referencing tables.

#### LATEST DETECTED DEADLOCK

This section provides information about the most recent deadlock. It is not present if no deadlock has occurred. The contents show which transactions are involved, the statement each was attempting to execute, the locks they have and need, and which transaction InnoDB decided to roll back to break the deadlock. The lock modes reported in this section are explained in [Section 14.2.3.2, “InnoDB Lock Modes”](#).

#### TRANSACTIONS

If this section reports lock waits, your applications might have lock contention. The output can also help to trace the reasons for transaction deadlocks.

#### FILE I/O

This section provides information about threads that InnoDB uses to perform various types of I/O. The first few of these are dedicated to general InnoDB processing. The contents also display information for pending I/O operations and statistics for I/O performance.

The number of these threads are controlled by the `innodb_read_io_threads` [1743] and `innodb_write_io_threads` [1754] parameters. See [Section 14.2.6, “InnoDB Startup Options and System Variables”](#).

#### INSERT BUFFER AND ADAPTIVE HASH INDEX

This section shows the status of the InnoDB insert buffer and adaptive hash index. (See [Section 14.2.3.12.5, “Insert Buffering”](#), and [Section 14.2.3.12.6, “Adaptive Hash Indexes”](#).) The contents include the number of operations performed for each, plus statistics for hash index performance.

#### LOG

This section displays information about the InnoDB log. The contents include the current log sequence number, how far the log has been flushed to disk, and the position at which InnoDB last took a checkpoint. (See [Section 5.3.3, “InnoDB Checkpoints”](#).) The section also displays information about pending writes and write performance statistics.

#### BUFFER POOL AND MEMORY

This section gives you statistics on pages read and written. You can calculate from these numbers how many data file I/O operations your queries currently are doing.

For additional information about the operation of the buffer pool, see [Section 8.9.1, “The InnoDB Buffer Pool”](#).

#### ROW OPERATIONS

This section shows what the main thread is doing, including the number and performance rate for each type of row operation.

### 14.2.4.4.2. InnoDB Tablespace Monitor Output

The InnoDB Tablespace Monitor prints information about the file segments in the shared tablespace and validates the tablespace allocation data structures. If you use individual tablespaces by enabling `innodb_file_per_table` [1718], the Tablespace Monitor does not describe those tablespaces.

Example InnoDB Tablespace Monitor output:

```

=====
090408 21:28:09 INNODB TABLESPACE MONITOR OUTPUT
=====
FILE SPACE INFO: id 0
size 13440, free limit 3136, free extents 28
not full frag extents 2: used pages 78, full frag extents 3
first seg id not used 0 23845
SEGMENT id 0 1 space 0; page 2; res 96 used 46; full ext 0
fragm pages 32; free extents 0; not full extents 1: pages 14
SEGMENT id 0 2 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
SEGMENT id 0 3 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
...
SEGMENT id 0 15 space 0; page 2; res 160 used 160; full ext 2
fragm pages 32; free extents 0; not full extents 0: pages 0
SEGMENT id 0 488 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
SEGMENT id 0 17 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
...
SEGMENT id 0 171 space 0; page 2; res 592 used 481; full ext 7
fragm pages 16; free extents 0; not full extents 2: pages 17
SEGMENT id 0 172 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
SEGMENT id 0 173 space 0; page 2; res 96 used 44; full ext 0
fragm pages 32; free extents 0; not full extents 1: pages 12
...
SEGMENT id 0 601 space 0; page 2; res 1 used 1; full ext 0
fragm pages 1; free extents 0; not full extents 0: pages 0
NUMBER of file segments: 73
Validating tablespace
Validation ok
-----
END OF INNODB TABLESPACE MONITOR OUTPUT
=====

```

The Tablespace Monitor output includes information about the shared tablespace as a whole, followed by a list containing a breakdown for each segment within the tablespace.

In this example using the default [page size](#), the tablespace consists of database pages that are 16KB each. The pages are grouped into [extents](#) of size 1MB (64 consecutive pages).

The initial part of the output that displays overall tablespace information has this format:

```

FILE SPACE INFO: id 0
size 13440, free limit 3136, free extents 28
not full frag extents 2: used pages 78, full frag extents 3
first seg id not used 0 23845

```

Overall tablespace information includes these values:

- **id**: The tablespace ID. A value of 0 refers to the shared tablespace.
- **size**: The current tablespace size in pages.
- **free limit**: The minimum page number for which the free list has not been initialized. Pages at or above this limit are free.
- **free extents**: The number of free extents.
- **not full frag extents, used pages**: The number of fragment extents that are not completely filled, and the number of pages in those extents that have been allocated.
- **full frag extents**: The number of completely full fragment extents.
- **first seg id not used**: The first unused segment ID.

Individual segment information has this format:

```
SEGMENT id 0 15 space 0; page 2; res 160 used 160; full ext 2
fragm pages 32; free extents 0; not full extents 0: pages 0
```

Segment information includes these values:

**id**: The segment ID.

**space**, **page**: The tablespace number and page within the tablespace where the segment “inode” is located. A tablespace number of 0 indicates the shared tablespace. InnoDB uses inodes to keep track of segments in the tablespace. The other fields displayed for a segment (**id**, **res**, and so forth) are derived from information in the inode.

**res**: The number of pages allocated (reserved) for the segment.

**used**: The number of allocated pages in use by the segment.

**full ext**: The number of extents allocated for the segment that are completely used.

**fragm pages**: The number of initial pages that have been allocated to the segment.

**free extents**: The number of extents allocated for the segment that are completely unused.

**not full extents**: The number of extents allocated for the segment that are partially used.

**pages**: The number of pages used within the not-full extents.

When a segment grows, it starts as a single page, and InnoDB allocates the first pages for it one at a time, up to 32 pages (this is the **fragm pages** value). After that, InnoDB allocates complete extents. InnoDB can add up to 4 extents at a time to a large segment to ensure good sequentiality of data.

For the example segment shown earlier, it has 32 fragment pages, plus 2 full extents (64 pages each), for a total of 160 pages used out of 160 pages allocated. The following segment has 32 fragment pages and one partially full extent using 14 pages for a total of 46 pages used out of 96 pages allocated:

```
SEGMENT id 0 1 space 0; page 2; res 96 used 46; full ext 0
fragm pages 32; free extents 0; not full extents 1: pages 14
```

It is possible for a segment that has extents allocated to it to have a **fragm pages** value less than 32 if some of the individual pages have been deallocated subsequent to extent allocation.

#### 14.2.4.4.3. InnoDB Table Monitor Output

The InnoDB Table Monitor prints the contents of the InnoDB internal data dictionary.

The output contains one section per table. The **SYS\_FOREIGN** and **SYS\_FOREIGN\_COLS** sections are for internal data dictionary tables that maintain information about foreign keys. There are also sections for the Table Monitor table and each user-created InnoDB table. Suppose that the following two tables have been created in the **test** database:

```
CREATE TABLE parent
(
  par_id    INT NOT NULL,
  fname    CHAR(20),
  lname    CHAR(20),
  PRIMARY KEY (par_id),
  UNIQUE INDEX (lname, fname)
) ENGINE = INNODB;

CREATE TABLE child
(
  par_id    INT NOT NULL,
  child_id  INT NOT NULL,
  name     VARCHAR(40),
  birth    DATE,
  weight   DECIMAL(10,2),
```

```

misc_info  VARCHAR(255),
last_update  TIMESTAMP,
PRIMARY KEY (par_id, child_id),
INDEX (name),
FOREIGN KEY (par_id) REFERENCES parent (par_id)
  ON DELETE CASCADE
  ON UPDATE CASCADE
) ENGINE = INNODB;

```

Then the Table Monitor output will look something like this (reformatted slightly):

```

=====
090420 12:09:32 INNODB TABLE MONITOR OUTPUT
=====
-----
TABLE: name SYS_FOREIGN, id 0 11, columns 7, indexes 3, appr.rows 1
  COLUMNS: ID: DATA_VARCHAR DATA_ENGLISH len 0;
            FOR_NAME: DATA_VARCHAR DATA_ENGLISH len 0;
            REF_NAME: DATA_VARCHAR DATA_ENGLISH len 0;
            N_COLS: DATA_INT len 4;
            DB_ROW_ID: DATA_SYS prtype 256 len 6;
            DB_TRX_ID: DATA_SYS prtype 257 len 6;
  INDEX: name ID_IND, id 0 11, fields 1/6, uniq 1, type 3
        root page 46, appr.key vals 1, leaf pages 1, size pages 1
        FIELDS:  ID DB_TRX_ID DB_ROLL_PTR FOR_NAME REF_NAME N_COLS
  INDEX: name FOR_IND, id 0 12, fields 1/2, uniq 2, type 0
        root page 47, appr.key vals 1, leaf pages 1, size pages 1
        FIELDS:  FOR_NAME ID
  INDEX: name REF_IND, id 0 13, fields 1/2, uniq 2, type 0
        root page 48, appr.key vals 1, leaf pages 1, size pages 1
        FIELDS:  REF_NAME ID
-----
TABLE: name SYS_FOREIGN_COLS, id 0 12, columns 7, indexes 1, appr.rows 1
  COLUMNS: ID: DATA_VARCHAR DATA_ENGLISH len 0;
            POS: DATA_INT len 4;
            FOR_COL_NAME: DATA_VARCHAR DATA_ENGLISH len 0;
            REF_COL_NAME: DATA_VARCHAR DATA_ENGLISH len 0;
            DB_ROW_ID: DATA_SYS prtype 256 len 6;
            DB_TRX_ID: DATA_SYS prtype 257 len 6;
  INDEX: name ID_IND, id 0 14, fields 2/6, uniq 2, type 3
        root page 49, appr.key vals 1, leaf pages 1, size pages 1
        FIELDS:  ID POS DB_TRX_ID DB_ROLL_PTR FOR_COL_NAME REF_COL_NAME
-----
TABLE: name test/child, id 0 14, columns 10, indexes 2, appr.rows 201
  COLUMNS: par_id: DATA_INT DATA_BINARY_TYPE DATA_NOT_NULL len 4;
            child_id: DATA_INT DATA_BINARY_TYPE DATA_NOT_NULL len 4;
            name: DATA_VARCHAR prtype 524303 len 40;
            birth: DATA_INT DATA_BINARY_TYPE len 3;
            weight: DATA_FIXBINARY DATA_BINARY_TYPE len 5;
            misc_info: DATA_VARCHAR prtype 524303 len 255;
            last_update: DATA_INT DATA_UNSIGNED DATA_BINARY_TYPE DATA_NOT_NULL len 4;
            DB_ROW_ID: DATA_SYS prtype 256 len 6;
            DB_TRX_ID: DATA_SYS prtype 257 len 6;
  INDEX: name PRIMARY, id 0 17, fields 2/9, uniq 2, type 3
        root page 52, appr.key vals 201, leaf pages 5, size pages 6
        FIELDS:  par_id child_id DB_TRX_ID DB_ROLL_PTR name birth weight misc_info last_update
  INDEX: name name, id 0 18, fields 1/3, uniq 3, type 0
        root page 53, appr.key vals 210, leaf pages 1, size pages 1
        FIELDS:  name par_id child_id
  FOREIGN KEY CONSTRAINT test/child_ibfk_1: test/child ( par_id )
        REFERENCES test/parent ( par_id )
-----
TABLE: name test/innodb_table_monitor, id 0 15, columns 4, indexes 1, appr.rows 0
  COLUMNS: i: DATA_INT DATA_BINARY_TYPE len 4;
            DB_ROW_ID: DATA_SYS prtype 256 len 6;
            DB_TRX_ID: DATA_SYS prtype 257 len 6;
  INDEX: name GEN_CLUST_INDEX, id 0 19, fields 0/4, uniq 1, type 1
        root page 193, appr.key vals 0, leaf pages 1, size pages 1
        FIELDS:  DB_ROW_ID DB_TRX_ID DB_ROLL_PTR i
-----
TABLE: name test/parent, id 0 13, columns 6, indexes 2, appr.rows 299
  COLUMNS: par_id: DATA_INT DATA_BINARY_TYPE DATA_NOT_NULL len 4;

```

```

      fname: DATA_CHAR prtype 524542 len 20;
      lname: DATA_CHAR prtype 524542 len 20;
      DB_ROW_ID: DATA_SYS prtype 256 len 6;
      DB_TRX_ID: DATA_SYS prtype 257 len 6;
INDEX: name PRIMARY, id 0 15, fields 1/5, uniq 1, type 3
      root page 50, appr.key vals 299, leaf pages 2, size pages 3
      FIELDS:  par_id DB_TRX_ID DB_ROLL_PTR fname lname
INDEX: name lname, id 0 16, fields 2/3, uniq 2, type 2
      root page 51, appr.key vals 300, leaf pages 1, size pages 1
      FIELDS:  lname fname par_id
FOREIGN KEY CONSTRAINT test/child_ibfk_1: test/child ( par_id )
      REFERENCES test/parent ( par_id )
-----
END OF INNODB TABLE MONITOR OUTPUT
=====

```

For each table, Table Monitor output contains a section that displays general information about the table and specific information about its columns, indexes, and foreign keys.

The general information for each table includes the table name (in *db\_name/tbl\_name* format except for internal tables), its ID, the number of columns and indexes, and an approximate row count.

The **COLUMNS** part of a table section lists each column in the table. Information for each column indicates its name and data type characteristics. Some internal columns are added by InnoDB, such as **DB\_ROW\_ID** (row ID), **DB\_TRX\_ID** (transaction ID), and **DB\_ROLL\_PTR** (a pointer to the rollback/undo data).

- **DATA\_xxx**: These symbols indicate the data type. There may be multiple **DATA\_xxx** symbols for a given column.
- **prtype**: The column's "precise" type. This field includes information such as the column data type, character set code, nullability, signedness, and whether it is a binary string. This field is described in the `innobase/include/data0type.h` source file.
- **len**: The column length in bytes.

Each **INDEX** part of the table section provides the name and characteristics of one table index:

- **name**: The index name. If the name is **PRIMARY**, the index is a primary key. If the name is **GEN\_CLUST\_INDEX**, the index is the clustered index that is created automatically if the table definition doesn't include a primary key or non-NULL unique index. See [Section 14.2.3.12.2, "Clustered and Secondary Indexes"](#).
- **id**: The index ID.
- **fields**: The number of fields in the index, as a value in *m/n* format:
  - **m** is the number of user-defined columns; that is, the number of columns you would see in the index definition in a `CREATE TABLE` statement.
  - **n** is the total number of index columns, including those added internally. For the clustered index, the total includes the other columns in the table definition, plus any columns added internally. For a secondary index, the total includes the columns from the primary key that are not part of the secondary index.
- **uniq**: The number of leading fields that are enough to determine index values uniquely.
- **type**: The index type. This is a bit field. For example, 1 indicates a clustered index and 2 indicates a unique index, so a clustered index (which always contains unique values), will have a **type** value of 3. An index with a **type** value of 0 is neither clustered nor unique. The flag values are defined in the `innobase/include/dict0mem.h` source file.
- **root page**: The index root page number.
- **appr. key vals**: The approximate index cardinality.

- `leaf pages`: The approximate number of leaf pages in the index.
- `size pages`: The approximate total number of pages in the index.
- `FIELDS`: The names of the fields in the index. For a clustered index that was generated automatically, the field list begins with the internal `DB_ROW_ID` (row ID) field. `DB_TRX_ID` and `DB_ROLL_PTR` are always added internally to the clustered index, following the fields that comprise the primary key. For a secondary index, the final fields are those from the primary key that are not part of the secondary index.

The end of the table section lists the `FOREIGN KEY` definitions that apply to the table. This information appears whether the table is a referencing or referenced table.

#### 14.2.4.5. InnoDB General Troubleshooting

The following general guidelines apply to troubleshooting InnoDB problems:

- When an operation fails or you suspect a bug, look at the MySQL server error log (see [Section 5.2.2, “The Error Log”](#)).
- If the failure is related to a `deadlock`, run with the `innodb_print_all_deadlocks [1741]` option enabled so that details about each InnoDB deadlock are printed to the MySQL server error log.
- Issues relating to the InnoDB data dictionary include failed `CREATE TABLE` statements (orphaned table files), inability to open `.InnoDB` files, and `system cannot find the path specified` errors. For information about these sorts of problems and errors, see [Section 14.2.4.7, “Troubleshooting InnoDB Data Dictionary Operations”](#).
- When troubleshooting, it is usually best to run the MySQL server from the command prompt, rather than through `mysqld_safe` or as a Windows service. You can then see what `mysqld` prints to the console, and so have a better grasp of what is going on. On Windows, start `mysqld` with the `--console [417]` option to direct the output to the console window.
- Use the InnoDB Monitors to obtain information about a problem (see [Section 14.2.4.4, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”](#)). If the problem is performance-related, or your server appears to be hung, use the standard Monitor to print information about the internal state of InnoDB. If the problem is with locks, use the Lock Monitor. If the problem is in creation of tables or other data dictionary operations, use the Table Monitor to print the contents of the InnoDB internal data dictionary. To see tablespace information use the Tablespace Monitor.
- If you suspect that a table is corrupt, run `CHECK TABLE` on that table.

##### 14.2.4.5.1. Troubleshooting InnoDB I/O Problems

The troubleshooting steps for InnoDB I/O problems depend on when the problem occurs: during startup of the MySQL server, or during normal operations when a DML or DDL statement fails due to problems at the file system level.

#### Initialization Problems

If something goes wrong when InnoDB attempts to initialize its tablespace or its log files, delete all files created by InnoDB: all `ibdata` files and all `ib_logfile` files. If you already created some InnoDB tables, also delete the corresponding `.frm` files for these tables, and any `.ibd` files if you are using multiple tablespaces, from the MySQL database directories. Then try the InnoDB database creation again. For easiest troubleshooting, start the MySQL server from a command prompt so that you see what is happening.

#### Runtime Problems

If InnoDB prints an operating system error during a file operation, usually the problem has one of the following solutions:



- Make sure the [InnoDB](#) data file directory and the [InnoDB](#) log directory exist.
- Make sure `mysqld` has access rights to create files in those directories.
- Make sure `mysqld` can read the proper `my.cnf` or `my.ini` option file, so that it starts with the options that you specified.
- Make sure the disk is not full and you are not exceeding any disk quota.
- Make sure that the names you specify for subdirectories and data files do not clash.
- Doublecheck the syntax of the `innodb_data_home_dir` [1715] and `innodb_data_file_path` [1714] values. In particular, any `MAX` value in the `innodb_data_file_path` [1714] option is a hard limit, and exceeding that limit causes a fatal error.

#### 14.2.4.6. Starting [InnoDB](#) on a Corrupted Database

To investigate database page corruption, you might dump your tables from the database with `SELECT ... INTO OUTFILE`. Usually, most of the data obtained in this way is intact. Serious corruption might cause `SELECT * FROM tbl_name` statements or [InnoDB](#) background operations to crash or assert, or even cause [InnoDB](#) roll-forward recovery to crash. In such cases, use the `innodb_force_recovery` [1722] option to force the [InnoDB](#) storage engine to start up while preventing background operations from running, so that you can dump your tables. For example, you can add the following line to the `[mysqld]` section of your option file before restarting the server:

```
[mysqld]
innodb_force_recovery = 4
```

`innodb_force_recovery` [1722] is 0 by default (normal startup without forced recovery). The permissible nonzero values for `innodb_force_recovery` [1722] follow. A larger number includes all precautions of smaller numbers. If you can dump your tables with an option value of at most 4, then you are relatively safe that only some data on corrupt individual pages is lost. A value of 6 is more drastic because database pages are left in an obsolete state, which in turn may introduce more corruption into B-trees and other database structures.

- 1 (`SRV_FORCE_IGNORE_CORRUPT`)

Lets the server run even if it detects a corrupt [page](#). Tries to make `SELECT * FROM tbl_name` jump over corrupt index records and pages, which helps in dumping tables.

- 2 (`SRV_FORCE_NO_BACKGROUND`)

Prevents the [master thread](#) and any [purge threads](#) from running. If a crash would occur during the [purge](#) operation, this recovery value prevents it.

- 3 (`SRV_FORCE_NO_TRX_UNDO`)

Does not run transaction [rollbacks](#) after [crash recovery](#).

- 4 (`SRV_FORCE_NO_IBUF_MERGE`)

Prevents [insert buffer](#) merge operations. If they would cause a crash, does not do them. Does not calculate table [statistics](#).

- 5 (`SRV_FORCE_NO_UNDO_LOG_SCAN`)

Does not look at [undo logs](#) when starting the database: [InnoDB](#) treats even incomplete transactions as committed.

- 6 (`SRV_FORCE_NO_LOG_REDO`)

Does not do the [redo log](#) roll-forward in connection with recovery.

With this value, you might not be able to do queries other than a basic `SELECT * FROM t`, with no `WHERE`, `ORDER BY`, or other clauses. More complex queries could encounter corrupted data structures and fail.

If corruption within the table data prevents you from dumping the entire table contents, a query with an `ORDER BY primary_key DESC` clause might be able to dump the portion of the table after the corrupted part.

*The database must not otherwise be used with any nonzero value of `innodb_force_recovery` [1722].* As a safety measure, InnoDB prevents `INSERT`, `UPDATE`, or `DELETE` operations when `innodb_force_recovery` [1722] is greater than 0.

You can `SELECT` from tables to dump them, or `DROP` or `CREATE` tables even if forced recovery is used. If you know that a given table is causing a crash on rollback, you can drop it. You can also use this to stop a runaway rollback caused by a failing mass import or `ALTER TABLE`: kill the `mysqld` process and set `innodb_force_recovery` [1722] to 3 to bring the database up without the rollback, then `DROP` the table that is causing the runaway rollback.

#### 14.2.4.7. Troubleshooting InnoDB Data Dictionary Operations

Information about table definitions is stored both in the `.frm` files, and in the InnoDB data dictionary. If you move `.frm` files around, or if the server crashes in the middle of a data dictionary operation, these sources of information can become inconsistent.

##### Problem with CREATE TABLE

A symptom of an out-of-sync data dictionary is that a `CREATE TABLE` statement fails. If this occurs, look in the server's error log. If the log says that the table already exists inside the InnoDB internal data dictionary, you have an orphaned table inside the InnoDB tablespace files that has no corresponding `.frm` file. The error message looks like this:

```
InnoDB: Error: table test/parent already exists in InnoDB internal
InnoDB: data dictionary. Have you deleted the .frm file
InnoDB: and not used DROP TABLE? Have you used DROP DATABASE
InnoDB: for InnoDB tables in MySQL version <= 3.23.43?
InnoDB: See the Restrictions section of the InnoDB manual.
InnoDB: You can drop the orphaned table inside InnoDB by
InnoDB: creating an InnoDB table with the same name in another
InnoDB: database and moving the .frm file to the current database.
InnoDB: Then MySQL thinks the table exists, and DROP TABLE will
InnoDB: succeed.
```

You can drop the orphaned table by following the instructions given in the error message. If you are still unable to use `DROP TABLE` successfully, the problem may be due to name completion in the `mysql` client. To work around this problem, start the `mysql` client with the `--skip-auto-rehash` [248] option and try `DROP TABLE` again. (With name completion on, `mysql` tries to construct a list of table names, which fails when a problem such as just described exists.)

##### Problem Opening Table

Another symptom of an out-of-sync data dictionary is that MySQL prints an error that it cannot open a `.InnoDB` file:

```
ERROR 1016: Can't open file: 'child2.InnoDB'. (errno: 1)
```

In the error log you can find a message like this:

```
InnoDB: Cannot find table test/child2 from the internal data dictionary
InnoDB: of InnoDB though the .frm file for the table exists. Maybe you
InnoDB: have deleted and recreated InnoDB data files but have forgotten
InnoDB: to delete the corresponding .frm files of InnoDB tables?
```

This means that there is an orphaned `.frm` file without a corresponding table inside InnoDB. You can drop the orphaned `.frm` file by deleting it manually.

## Problem with Temporary Table

If MySQL crashes in the middle of an `ALTER TABLE` operation, you may end up with an orphaned temporary table inside the InnoDB tablespace. Using the Table Monitor, you can see listed a table with a name that begins with `#sql-`. You can perform SQL statements on tables whose name contains the character “#” if you enclose the name within backticks. Thus, you can drop such an orphaned table like any other orphaned table using the method described earlier. To copy or rename a file in the Unix shell, you need to put the file name in double quotation marks if the file name contains “#”.

## Problem with Missing Tablespace

With `innodb_file_per_table` [1718] enabled, the following message might occur if the `.frm` or `.ibd` files (or both) are missing:

```
InnoDB: in InnoDB data dictionary has tablespace id N,
InnoDB: but tablespace with that id or name does not exist. Have
InnoDB: you deleted or moved .ibd files?
InnoDB: This may also be a table created with CREATE TEMPORARY TABLE
InnoDB: whose .ibd and .frm files MySQL automatically removed, but the
InnoDB: table still exists in the InnoDB internal data dictionary.
```

If this occurs, try the following procedure to resolve the problem:

1. Create a matching `.frm` file in some other database directory and copy it to the database directory where the orphan table is located.
2. Issue `DROP TABLE` for the original table. That should successfully drop the table and InnoDB should print a warning to the error log that the `.ibd` file was missing.

## 14.2.5. InnoDB Features for Flexibility, Ease of Use and Reliability

This section describes several recently added InnoDB features that offer new flexibility and improve ease of use, reliability and performance. The Barracuda file format improves efficiency for storing large variable-length columns, and enables table [compression](#). Configuration options that once were unchangeable after startup, are now flexible and can be changed dynamically. Some improvements are automatic, such as faster and more efficient `TRUNCATE TABLE`. Others allow you the flexibility to control InnoDB behavior; for example, you can control whether certain problems cause errors or just warnings. And informational messages and error reporting continue to be made more user-friendly.

### 14.2.5.1. Support for Read-Only Media

You can now query InnoDB tables where the MySQL data directory is on read-only media, by enabling the `--innodb-read-only` [1743] configuration option at server startup.

#### How to Enable

To prepare an instance for read-only operation, make sure all the necessary information is [flushed](#) to the data files before storing it on the read-only medium. Run the server with change buffering disabled (`innodb_change_buffering=0` [1709]) and do a [slow shutdown](#).

To enable read-only mode for an entire MySQL instance, specify the following configuration options at server startup:

- `--innodb-read-only=1` [1743]
- If the instance is on read-only media such as a DVD or CD, or the `/var` directory is not writeable by all: `--pid-file=path_on_writeable_media` and `--event-scheduler=disabled`

#### Usage Scenarios

This mode of operation is appropriate in situations such as:

- Distributing a MySQL application, or a set of MySQL data, on a read-only storage medium such as a DVD or CD.

- Multiple MySQL instances querying the same data directory simultaneously, typically in a data warehousing configuration. You might use this technique to avoid [bottlenecks](#) that can occur with a heavily loaded MySQL instance, or you might use different configuration options for the various instances to tune each one for particular kinds of queries.
- Querying data that has been put into a read-only state for security or data integrity reasons, such as archived backup data.



### Note

This feature is mainly intended for flexibility in distribution and deployment, rather than raw performance based on the read-only aspect. See [Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”](#) for ways to tune the performance of read-only queries, which do not require making the entire server read-only.

## How It Works

When the server is run in read-only mode through the `--innodb-read-only [1743]` option, certain InnoDB features and components are reduced or turned off entirely:

- No [change buffering](#) is done, in particular no merges from the change buffer. To make sure the change buffer is empty when you prepare the instance for read-only operation, disable change buffering (`innodb_change_buffering=0 [1709]`) and do a [slow shutdown](#) first.
- There is no [crash recovery](#) phase at startup. The instance must have performed a [slow shutdown](#) before being put into the read-only state.
- Because the [redo log](#) is not used in read-only operation, you can set `innodb_log_file_size [1732]` to the smallest size possible (1 MB) before making the instance read-only.
- All background threads other than I/O read threads are turned off. As a consequence, a read-only instance cannot encounter any [deadlocks](#).
- Information about deadlocks, monitor output, and so on is not written to temporary files. As a consequence, `SHOW ENGINE INNODB STATUS` does not produce any output.
- If the MySQL server is started with `--innodb-read-only [1743]` but the data directory is still on writeable media, the root user can still perform [DCL](#) operations such as `GRANT` and `REVOKE`.
- Changes to configuration option settings that would normally change the the behavior of write operations, have no effect when the server is in read-only mode.
- The [MVCC](#) processing to enforce [isolation levels](#) is turned off. All queries read the latest version of a record, because update and deletes are not possible.
- The [undo log](#) is not used. Disable any settings for the `innodb_undo_tablespaces [1753]` and `innodb_undo_directory [1752]` configuration options.

### 14.2.5.2. Larger Size Limit for Redo Log Files

Formerly, the combined size of the InnoDB [redo log](#) files was limited to 4 gigabytes. Starting in MySQL 5.6.3, this size limit is raised to 512GB.

You do not need any special upgrade process or file format to take advantage of this feature. The bytes that record the extra size information were already reserved and set to zero in the InnoDB [system tablespace](#).

If you develop applications that interact with the InnoDB [logical sequence number \(LSN\)](#) value, change your code to use guaranteed 64-bit variables to store and compare LSN values, rather than 32-bit variables.

### 14.2.5.3. 2-Byte Collation IDs for InnoDB Tables

InnoDB tables can now use collation IDs greater than 255. Currently, the collation IDs in this range are all user-defined. For example, the following InnoDB table can now be created, where formerly the collation ID of 359 was beyond the range supported by InnoDB.

```
sql> show collation like 'ucs2_vn_ci';
+-----+-----+-----+-----+-----+-----+
| Collation | Charset | Id | Default | Compiled | Sortlen |
+-----+-----+-----+-----+-----+-----+
| ucs2_vn_ci | ucs2    | 359 |         |          | 8       |
+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)

mysql> create table two_byte_collation (c1 char(1) character set ucs2 collate ucs2_vn_ci)
-> engine = InnoDB;
Query OK, 0 rows affected (0.16 sec)
```

### 14.2.5.4. The Barracuda File Format

InnoDB has started using named file formats to improve compatibility in upgrade and downgrade situations, or heterogeneous systems running different levels of MySQL. Many important InnoDB features, such as table compression and the `DYNAMIC` row format for more efficient BLOB storage, require creating tables in the `Barracuda` file format. The original file format, which previously didn't have a name, is known now as `Antelope`.

To create new tables that take advantage of the Barracuda features, enable that file format using the configuration parameter `innodb_file_format` [1716]. The value of this parameter determines whether a newly created table or index can use compression or the new `DYNAMIC` row format.

To preclude the use of new features that would make your database inaccessible to the built-in InnoDB in MySQL 5.1 and prior releases, omit `innodb_file_format` [1716] or set it to `Antelope`.

You can set the value of `innodb_file_format` [1716] on the command line when you start `mysqld`, or in the option file `my.cnf` (Unix operating systems) or `my.ini` (Windows). You can also change it dynamically with the `SET GLOBAL` statement.

For more information about managing file formats, see [Section 5.4.7, “InnoDB File-Format Management”](#).

### 14.2.5.5. Dynamic Control of System Configuration Parameters

In MySQL 5.5 and higher, you can change certain system configuration parameters without shutting down and restarting the server, as was necessary in MySQL 5.1 and lower. This increases uptime, and makes it easier to test and prototype new SQL and application code. The following sections explain these parameters.

#### 14.2.5.5.1. Dynamically Changing `innodb_file_per_table`

Since MySQL version 4.1, InnoDB has provided two alternatives for how tables are stored on disk. You can create a new table and its indexes in the shared `system tablespace`, physically stored in the `ibdata files`. Or, you can store a new table and its indexes in a separate tablespace (a `.ibd file`). The storage layout for each InnoDB table is determined by the the configuration parameter `innodb_file_per_table` [1718] at the time the table is created.

In MySQL 5.5 and higher, the configuration parameter `innodb_file_per_table` [1718] is dynamic, and can be set `ON` or `OFF` using the `SET GLOBAL`. Previously, the only way to set this parameter was in the MySQL `configuration file` (`my.cnf` or `my.ini`), and changing it required shutting down and restarting the server.

The default setting is `OFF`, so new tables and indexes are created in the system tablespace. Dynamically changing the value of this parameter requires the `SUPER` privilege and immediately affects the operation of all connections.

Tables created when `innodb_file_per_table` [1718] is enabled can use the [Barracuda](#) file format, and `TRUNCATE` returns the disk space for those tables to the operating system. The Barracuda file format in turn enables features such as table compression and the `DYNAMIC` row format. Tables created when `innodb_file_per_table` [1718] is off cannot use these features. To take advantage of those features for an existing table, you can turn on the file-per-table setting and run `ALTER TABLE t ENGINE=INNODB` for that table.

When you redefine the primary key for an InnoDB table, the table is re-created using the current settings for `innodb_file_per_table` [1718] and `innodb_file_format` [1716]. This behavior does not apply when adding or dropping InnoDB secondary indexes, as explained in [Fast Index Creation in the InnoDB Storage Engine](#). When a secondary index is created without rebuilding the table, the index is stored in the same file as the table data, regardless of the current `innodb_file_per_table` [1718] setting.

#### 14.2.5.5.2. Dynamically Changing `innodb_stats_on_metadata`

In MySQL 5.5 and higher, you can change the setting of `innodb_stats_on_metadata` [1747] dynamically at runtime, to control whether or not InnoDB performs statistics gathering when metadata statements are executed. To change the setting, issue the statement `SET GLOBAL innodb_stats_on_metadata=mode`, where *mode* is either `ON` or `OFF` (or `1` or `0`). Changing this setting requires the `SUPER` privilege and immediately affects the operation of all connections.

This setting is related to the feature described in [Section 14.2.5.8, “Controlling Optimizer Statistics Estimation”](#).

#### 14.2.5.5.3. Dynamically Changing `innodb_lock_wait_timeout`

The length of time a transaction waits for a resource, before giving up and rolling back the statement, is determined by the value of the configuration parameter `innodb_lock_wait_timeout` [1729]. (In MySQL 5.0.12 and earlier, the entire transaction was rolled back, not just the statement.) Your application can try the statement again (usually after waiting for a while), or roll back the entire transaction and restart.

The error returned when the timeout period is exceeded is:

```
ERROR HY000: Lock wait timeout exceeded; try restarting transaction
```

In MySQL 5.5 and higher, the configuration parameter `innodb_lock_wait_timeout` [1729] can be set at runtime with the `SET GLOBAL` or `SET SESSION` statement. Changing the `GLOBAL` setting requires the `SUPER` privilege and affects the operation of all clients that subsequently connect. Any client can change the `SESSION` setting for `innodb_lock_wait_timeout` [1729], which affects only that client.

In MySQL 5.1 and earlier, the only way to set this parameter was in the MySQL [configuration file](#) (`my.cnf` or `my.ini`), and changing it required shutting down and restarting the server.

#### 14.2.5.5.4. Dynamically Changing `innodb_adaptive_hash_index`

As described in [Section 14.2.4.2.14, “Controlling Adaptive Hash Indexing”](#), it may be desirable, depending on your workload, to dynamically enable or disable the [adaptive hash indexing](#) scheme InnoDB uses to improve query performance.

The configuration option `innodb_adaptive_hash_index` [1701] lets you disable the adaptive hash index. It is enabled by default. You can modify this parameter through the `SET GLOBAL` statement, without restarting the server. Changing the setting requires the `SUPER` privilege.

Disabling the adaptive hash index empties the hash table immediately. Normal operations can continue while the hash table is emptied, and executing queries that were using the hash table access the index B-trees directly instead. When the adaptive hash index is re-enabled, the hash table is populated again during normal operation.

#### 14.2.5.6. `TRUNCATE TABLE` Reclaims Space



When you `truncate` a table that is stored in a `.ibd` file of its own (because `innodb_file_per_table` [1718] was enabled when the table was created), and if the table is not referenced in a `FOREIGN KEY` constraint, the table is dropped and re-created in a new `.ibd` file. This operation is much faster than deleting the rows one by one. The operating system can reuse the disk space, in contrast to tables within the InnoDB `system tablespace`, where only InnoDB can use the space after they are truncated. `Physical backups` can also be smaller, without big blocks of unused space in the middle of the system tablespace.

MySQL 5.1 and earlier would re-use the existing `.ibd` file, thus releasing the space only to InnoDB for storage management, but not to the operating system. Note that when the table is truncated, the count of rows affected by the `TRUNCATE TABLE` statement is an arbitrary number.



#### Note

If there is a `foreign key constraint` between two columns in the same table, that table can still be truncated using this fast technique.

If there are foreign key constraints between the table being truncated and other tables, the truncate operation fails. This is a change to the previous behavior, which would transform the `TRUNCATE` operation to a `DELETE` operation that removed all the rows and triggered `ON DELETE` operations on `child tables`.

### 14.2.5.7. InnoDB Strict Mode

To guard against ignored typos and syntax errors in SQL, or other unintended consequences of various combinations of operational modes and SQL statements, InnoDB provides a `strict mode` of operations. In this mode, InnoDB raises error conditions in certain cases, rather than issuing a warning and processing the specified statement (perhaps with unintended behavior). This is analogous to `sql_mode` in MySQL, which controls what SQL syntax MySQL accepts, and determines whether it silently ignores errors, or validates input syntax and data values. Since InnoDB strict mode is relatively new, some statements that execute without errors with earlier versions of MySQL might generate errors unless you disable strict mode.

The setting of InnoDB strict mode affects the handling of syntax errors on the `CREATE TABLE`, `ALTER TABLE` and `CREATE INDEX` statements. The strict mode also enables a record size check, so that an `INSERT` or `UPDATE` never fails due to the record being too large for the selected page size.

Oracle recommends enabling `innodb_strict_mode` [1749] when using the `ROW_FORMAT` and `KEY_BLOCK_SIZE` clauses on `CREATE TABLE`, `ALTER TABLE`, and `CREATE INDEX` statements. Without strict mode, InnoDB ignores conflicting clauses and creates the table or index, with only a warning in the message log. The resulting table might have different behavior than you intended, such as having no compression when you tried to create a compressed table. When InnoDB strict mode is on, such problems generate an immediate error and the table or index is not created, avoiding a troubleshooting session later.

InnoDB strict mode is set with the configuration parameter `innodb_strict_mode` [1749], which can be specified as `ON` or `OFF`. You can set the value on the command line when you start `mysqld`, or in the `configuration file` `my.cnf` or `my.ini`. You can also enable or disable InnoDB strict mode at runtime with the statement `SET [GLOBAL|SESSION] innodb_strict_mode=mode`, where `mode` is either `ON` or `OFF`. Changing the `GLOBAL` setting requires the `SUPER` privilege and affects the operation of all clients that subsequently connect. Any client can change the `SESSION` setting for `innodb_strict_mode` [1749], and the setting affects only that client.

### 14.2.5.8. Controlling Optimizer Statistics Estimation

The MySQL query optimizer uses estimated `statistics` about key distributions to choose the indexes for an execution plan, based on the relative `selectivity` of the index. Certain operations cause InnoDB to sample random pages from each index on a table to estimate the `cardinality` of the index. (This technique is known as `random dives`.) These operations include the `ANALYZE TABLE` statement, the `SHOW TABLE STATUS` statement, and accessing the table for the first time after a restart.



To give you control over the quality of the statistics estimate (and thus better information for the query optimizer), you can now change the number of sampled pages using the parameter `innodb_stats_sample_pages` [1748]. Previously, the number of sampled pages was always 8, which could be insufficient to produce an accurate estimate, leading to poor index choices by the query optimizer. This technique is especially important for large tables and tables used in `joins`. Unnecessary `full table scans` for such tables can be a substantial performance issue. See [Section 8.2.1.4, “How to Avoid Full Table Scans”](#) for tips on tuning such queries.

You can set the global parameter `innodb_stats_sample_pages` [1748], at runtime. The default value for this parameter is 8, preserving the same behavior as in past releases.



#### Note

The value of `innodb_stats_sample_pages` [1748] affects the index sampling for *all* InnoDB tables and indexes. There are the following potentially significant impacts when you change the index sample size:

- Small values like 1 or 2 can result in very inaccurate estimates of cardinality.
- Increasing the `innodb_stats_sample_pages` [1748] value might require more disk reads. Values much larger than 8 (say, 100), can cause a big slowdown in the time it takes to open a table or execute `SHOW TABLE STATUS`.
- The optimizer might choose very different query plans based on different estimates of index selectivity.

To disable the cardinality estimation for metadata statements such as `SHOW TABLE STATUS`, execute the statement `SET GLOBAL innodb_stats_on_metadata=OFF` (or 0). The ability to set this option dynamically is also relatively new.

All InnoDB tables are opened, and the statistics are re-estimated for all associated indexes, when the `mysql` client starts if the auto-rehash setting is set on (the default). To improve the start up time of the `mysql` client, you can turn auto-rehash off. The auto-rehash feature enables automatic name completion of database, table, and column names for interactive users.

Whatever value of `innodb_stats_sample_pages` [1748] works best for a system, set the option and leave it at that value. Choose a value that results in reasonably accurate estimates for all tables in your database without requiring excessive I/O. Because the statistics are automatically recalculated at various times other than on execution of `ANALYZE TABLE`, it does not make sense to increase the index sample size, run `ANALYZE TABLE`, then decrease sample size again. The more accurate statistics calculated by `ANALYZE` running with a high value of `innodb_stats_sample_pages` [1748] can be wiped away later.

Although it is not possible to specify the sample size on a per-table basis, smaller tables generally require fewer index samples than larger tables do. If your database has many large tables, consider using a higher value for `innodb_stats_sample_pages` [1748] than if you have mostly smaller tables.

### 14.2.5.9. Better Error Handling when Dropping Indexes

For optimal performance with DML statements, InnoDB requires an index to exist on `foreign key` columns, so that `UPDATE` and `DELETE` operations on a `parent table` can easily check whether corresponding rows exist in the `child table`. MySQL creates or drops such indexes automatically when needed, as a side-effect of `CREATE TABLE`, `CREATE INDEX`, and `ALTER TABLE` statements.

When you drop an index, InnoDB checks whether the index is not used for checking a foreign key constraint. It is still OK to drop the index if there is another index that can be used to enforce the same constraint. InnoDB prevents you from dropping the last index that can enforce a particular referential constraint.

The message that reports this error condition is:

```
ERROR 1553 (HY000): Cannot drop index 'fooIdx':
needed in a foreign key constraint
```

This message is friendlier than the earlier message it replaces:

```
ERROR 1025 (HY000): Error on rename of './db2/#sql-18eb_3'
to './db2/foo'(errno: 150)
```

A similar change in error reporting applies to an attempt to drop the primary key index. For tables without an explicit `PRIMARY KEY`, InnoDB creates an implicit `clustered index` using the first columns of the table that are declared `UNIQUE` and `NOT NULL`. When you drop such an index, InnoDB automatically copies the table and rebuilds the index using a different `UNIQUE NOT NULL` group of columns or a system-generated key. Since this operation changes the primary key, it uses the slow method of copying the table and re-creating the index, rather than the Fast Index Creation technique from [Section 5.5.6, "Implementation Details of Online DDL"](#).

Previously, an attempt to drop an implicit clustered index (the first `UNIQUE NOT NULL` index) failed if the table did not contain a `PRIMARY KEY`:

```
ERROR 42000: This table type requires a primary key
```

#### 14.2.5.10. More Compact Output of `SHOW ENGINE INNODB MUTEX`

The statement `SHOW ENGINE INNODB MUTEX` displays information about InnoDB `mutexes` and `rw-locks`. Although this information is useful for tuning on multi-core systems, the amount of output can be overwhelming on systems with a big `buffer pool`. There is one mutex and one rw-lock in each 16K buffer pool block, and there are 65,536 blocks per gigabyte. It is unlikely that a single block mutex or rw-lock from the buffer pool could become a performance bottleneck.

`SHOW ENGINE INNODB MUTEX` now skips the mutexes and rw-locks of buffer pool blocks. It also does not list any mutexes or rw-locks that have never been waited on (`os_waits=0`). Thus, `SHOW ENGINE INNODB MUTEX` only displays information about mutexes and rw-locks outside of the buffer pool that have caused at least one OS-level `wait`.

#### 14.2.5.11. More Read-Ahead Statistics

As described in [Section 14.2.4.2.16, "Changes in the Read-Ahead Algorithm"](#), a read-ahead request is an asynchronous I/O request issued in anticipation that a page will be used in the near future. Knowing how many pages are read through this read-ahead mechanism, and how many of them are evicted from the buffer pool without ever being accessed, can be useful to help fine-tune the parameter `innodb_read_ahead_threshold` [1742].

`SHOW ENGINE INNODB STATUS` output displays the global status variables `Innodb_buffer_pool_read_ahead` [591] and `Innodb_buffer_pool_read_ahead_evicted` [591]. These variables indicate the number of pages brought into the `buffer pool` by read-ahead requests, and the number of such pages `evicted` from the buffer pool without ever being accessed respectively. These counters provide global values since the last server restart.

`SHOW ENGINE INNODB STATUS` also shows the rate at which the read-ahead pages are read in and the rate at which such pages are evicted without being accessed. The per-second averages are based on the statistics collected since the last invocation of `SHOW ENGINE INNODB STATUS` and are displayed in the `BUFFER POOL AND MEMORY` section of the output.

### 14.2.6. InnoDB Startup Options and System Variables

This section describes the InnoDB-related command options and system variables. System variables that are true or false can be enabled at server startup by naming them, or disabled by using a `--skip-` prefix. For example, to enable or disable InnoDB checksums, you can use `--innodb_checksums` [1711] or `--skip-innodb_checksums` [1711] on the command line, or `innodb_checksums` [1711] or `skip-innodb_checksums` in an option file. System variables

that take a numeric value can be specified as `--var_name=value` on the command line or as `var_name=value` in option files. For more information on specifying options and system variables, see [Section 4.2.3, “Specifying Program Options”](#). Many of the system variables can be changed at runtime (see [Section 5.1.5.2, “Dynamic System Variables”](#)).

Certain options control the locations and layout of the InnoDB data files. [Section 14.2.1.2, “Configuring InnoDB”](#) explains how to use these options. Many other options, that you might not use initially, help to tune InnoDB performance characteristics based on machine capacity and your database workload. The performance-related options are explained in [Section 14.2.4, “InnoDB Performance Tuning and Troubleshooting”](#) and [Section 14.2.4.2, “InnoDB Performance and Scalability Enhancements”](#).

**Table 14.5. InnoDB Option/Variable Reference**

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">daemon_memcached</a>	<a href="#">Yes enable_binlog [1698]</a>	<a href="#">Yes [1698]</a>	Yes		Global	No
<a href="#">daemon_memcached</a>	<a href="#">Yes engine_library_path [1698]</a>	<a href="#">Yes [1698]</a>	Yes		Global	No
<a href="#">daemon_memcached</a>	<a href="#">Yes engine_library_path [1699]</a>	<a href="#">Yes [1699]</a>	Yes		Global	No
<a href="#">daemon_memcached</a>	<a href="#">Yes option [1699]</a>	<a href="#">Yes [1699]</a>	Yes		Global	No
<a href="#">daemon_memcached</a>	<a href="#">Yes r_batch_size [1699]</a>	<a href="#">Yes [1699]</a>	Yes		Global	No
<a href="#">daemon_memcached</a>	<a href="#">Yes w_batch_size [1700]</a>	<a href="#">Yes [1700]</a>	Yes		Global	No
<a href="#">foreign_key_checks</a> [483]			Yes		Both	Yes
<a href="#">have_innodb</a> [487]			Yes		Global	No
<a href="#">ignore-builtin-innodb</a> [1697]	Yes	Yes			Global	No
- Variable: <a href="#">ignore_builtin_innodb</a>			Yes		Global	No
<a href="#">innodb</a> [1698]	Yes	Yes				
<a href="#">innodb_adaptive_flushing</a> [1700]	<a href="#">Yes [1700]</a>	<a href="#">Yes [1700]</a>	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_lwm</a> [1701]	<a href="#">Yes [1701]</a>	<a href="#">Yes [1701]</a>	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_index</a> [1702]	<a href="#">Yes [1702]</a>	<a href="#">Yes [1702]</a>	Yes		Global	Yes
<a href="#">innodb_adaptive_flushing_sleep_delay</a> [1702]	<a href="#">Yes [1702]</a>	<a href="#">Yes [1702]</a>	Yes		Global	Yes
<a href="#">innodb_additional_mem_pool_size</a> [1702]	<a href="#">Yes [1702]</a>	<a href="#">Yes [1702]</a>	Yes		Global	No
<a href="#">innodb_api_bk_commit_interval</a> [1702]	<a href="#">Yes [1702]</a>	<a href="#">Yes [1702]</a>	Yes		Global	Yes
<a href="#">innodb_api_disable_rowlock</a> [1703]	<a href="#">Yes [1703]</a>	<a href="#">Yes [1703]</a>	Yes		Global	No
<a href="#">innodb_api_enable_binlog</a> [1703]	<a href="#">Yes [1703]</a>	<a href="#">Yes [1703]</a>	Yes		Global	No
<a href="#">innodb_api_enable_binlog</a> [1703]	<a href="#">Yes [1703]</a>	<a href="#">Yes [1703]</a>	Yes		Global	No
<a href="#">innodb_api_trx_level</a> [1704]	<a href="#">Yes [1704]</a>	<a href="#">Yes [1704]</a>	Yes		Global	Yes
<a href="#">innodb_autoextend_increment</a> [1705]	<a href="#">Yes [1705]</a>	<a href="#">Yes [1705]</a>	Yes		Global	Yes
<a href="#">innodb_autoinc_lock_mode</a> [1705]	<a href="#">Yes [1705]</a>	<a href="#">Yes [1705]</a>	Yes		Global	No
<a href="#">Innodb_available_undo_logs</a> [590]				Yes	Global	No
<a href="#">Innodb_buffer_pool_bytes_data</a> [590]				Yes	Global	No
<a href="#">Innodb_buffer_pool_bytes_dirty</a> [590]				Yes	Global	No
<a href="#">innodb_buffer_pool_dump_at_shutdown</a> [1705]	<a href="#">Yes [1705]</a>	<a href="#">Yes [1705]</a>	Yes		Global	Yes
<a href="#">innodb_buffer_pool_dump_now</a> [1706]	<a href="#">Yes [1706]</a>	<a href="#">Yes [1706]</a>	Yes		Global	Yes
<a href="#">Innodb_buffer_pool_dump_status</a> [590]				Yes	Global	No
<a href="#">innodb_buffer_pool_filename</a> [1706]	<a href="#">Yes [1706]</a>	<a href="#">Yes [1706]</a>	Yes		Global	Yes
<a href="#">innodb_buffer_pool_instances</a> [1706]	<a href="#">Yes [1706]</a>	<a href="#">Yes [1706]</a>	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
innodb_buffer_pool_read_abort [1707]	Yes	Yes	Yes		Global	Yes
innodb_buffer_pool_read_at_startup [1707]	Yes	Yes	Yes		Global	No
innodb_buffer_pool_read_now [1707]	Yes	Yes	Yes		Global	Yes
Innodb_buffer_pool_load_status [590]				Yes	Global	No
Innodb_buffer_pool_pages_data [590]				Yes	Global	No
Innodb_buffer_pool_pages_dirty [590]				Yes	Global	No
Innodb_buffer_pool_pages_flushed [591]				Yes	Global	No
Innodb_buffer_pool_pages_free [591]				Yes	Global	No
Innodb_buffer_pool_pages_latched [591]				Yes	Global	No
Innodb_buffer_pool_pages_misc [591]				Yes	Global	No
Innodb_buffer_pool_pages_total [591]				Yes	Global	No
Innodb_buffer_pool_read_ahead [591]				Yes	Global	No
Innodb_buffer_pool_read_ahead_evicted [591]				Yes	Global	No
Innodb_buffer_pool_read_requests [591]				Yes	Global	No
Innodb_buffer_pool_reads [591]				Yes	Global	No
innodb_buffer_pool_size [1708]	Yes	Yes	Yes		Global	No
Innodb_buffer_pool_wait_free [591]				Yes	Global	No
Innodb_buffer_pool_write_requests [591]				Yes	Global	No
innodb_change_buffer_max_size [1709]	Yes	Yes	Yes		Global	Yes
innodb_change_buffering [1709]	Yes	Yes	Yes		Global	Yes
innodb_checksum_algorithm [1710]	Yes	Yes	Yes		Global	Yes
innodb_checksums [1711]	Yes	Yes	Yes		Global	No
innodb_cmp_per_index_enabled [1712]	Yes	Yes	Yes		Global	Yes
innodb_commit_concurrency [1713]	Yes	Yes	Yes		Global	Yes
innodb_compression_failure_threshold_pct [1713]	Yes	Yes	Yes		Global	Yes
innodb_compression_level [1713]	Yes	Yes	Yes		Global	Yes
innodb_compression_pad_pct_max [1713]	Yes	Yes	Yes		Global	Yes
innodb_concurrency_tickets [1714]	Yes	Yes	Yes		Global	Yes
innodb_data_file_path [1714]	Yes	Yes	Yes		Global	No
Innodb_data_fsyncs [591]				Yes	Global	No
innodb_data_home_dir [1715]	Yes	Yes	Yes		Global	No
Innodb_data_pending_fsyncs [591]				Yes	Global	No
Innodb_data_pending_reads [592]				Yes	Global	No
Innodb_data_pending_writes [592]				Yes	Global	No
Innodb_data_read [592]				Yes	Global	No
Innodb_data_reads [592]				Yes	Global	No
Innodb_data_writes [592]				Yes	Global	No
Innodb_data_written [592]				Yes	Global	No
Innodb_dblwr_pages_written [592]				Yes	Global	No
Innodb_dblwr_writes [592]				Yes	Global	No
innodb_disable_socket_cache [1715]	Yes	Yes	Yes		Global	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
innodb_doublewrite	Yes [1715]	Yes	Yes		Global	No
innodb_fast_shutdown	Yes [1716]	Yes	Yes		Global	Yes
innodb_file_format	Yes [1716]	Yes	Yes		Global	Yes
innodb_file_format_check	Yes [1717]	Yes	Yes		Global	No
innodb_file_format_max	Yes [1717]	Yes	Yes		Global	Yes
innodb_file_per_table	Yes [1718]	Yes	Yes		Global	Yes
innodb_flush_log_at_timeout	[1718]		Yes		Global	Yes
innodb_flush_log_at_trx_commit	Yes [1718]	Yes	Yes		Global	Yes
innodb_flush_method	Yes [1719]	Yes	Yes		Global	No
innodb_flush_neighbors	Yes [1721]	Yes	Yes		Global	Yes
innodb_flushing_avg_loops	Yes [1722]	Yes	Yes		Global	Yes
innodb_force_load_corrupted	Yes [1722]	Yes	Yes		Global	No
innodb_force_recovery	Yes [1722]	Yes	Yes		Global	No
innodb_ft_aux_table	Yes [1723]	Yes	Yes		Global	Yes
innodb_ft_cache_size	Yes [1724]	Yes	Yes		Global	No
innodb_ft_enable_synonym	Yes [1724]	Yes	Yes		Global	Yes
innodb_ft_max_token_size	Yes [1724]	Yes	Yes		Global	No
innodb_ft_min_token_size	Yes [1725]	Yes	Yes		Global	No
innodb_ft_num_word_optimize	Yes [1725]	Yes	Yes		Global	Yes
innodb_ft_server_stopword_table	Yes [1725]	Yes	Yes		Global	Yes
innodb_ft_sort_pll_degree	Yes [1726]	Yes	Yes		Global	No
innodb_ft_user_stopword_table	Yes [1726]	Yes	Yes		Both	Yes
InnoDB_have_atomic_builtins	[592]			Yes	Global	No
innodb_io_capacity	Yes [1727]	Yes	Yes		Global	Yes
innodb_io_capacity_max	Yes [1728]	Yes	Yes		Global	Yes
innodb_large_prefix	Yes [1728]	Yes	Yes		Global	Yes
innodb_lock_wait_timeout	Yes [1729]	Yes	Yes		Both	Yes
innodb_locks_unsafe_for_binlog	Yes [1729]	Yes	Yes		Global	No
innodb_log_buffer_size	Yes [1732]	Yes	Yes		Global	No
innodb_log_compressed_pages	Yes [1732]	Yes	Yes		Global	Yes
innodb_log_file_size	Yes [1732]	Yes	Yes		Global	No
innodb_log_files_in_group	Yes [1733]	Yes	Yes		Global	No
innodb_log_group_home_dir	Yes [1733]	Yes	Yes		Global	No
InnoDB_log_waits	[592]			Yes	Global	No
InnoDB_log_write_requests	[592]			Yes	Global	No
InnoDB_log_writes	[592]			Yes	Global	No
innodb_lru_scan_depth	Yes [1734]	Yes	Yes		Global	Yes
innodb_max_dirty_pages_pct	Yes [1734]	Yes	Yes		Global	Yes
innodb_max_dirty_pages_pct_lwm	Yes [1735]	Yes	Yes		Global	Yes
innodb_max_purge_age	Yes [1735]	Yes	Yes		Global	Yes
innodb_max_purge_age_delay	Yes [1735]	Yes	Yes		Global	Yes

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
innodb_mirrored_log_groups [1736]	Yes	Yes	Yes		Global	No
innodb_monitor_disable [1736]	Yes	Yes	Yes		Global	Yes
innodb_monitor_enable [1737]	Yes	Yes	Yes		Global	Yes
innodb_monitor_reset [1737]	Yes	Yes	Yes		Global	Yes
innodb_monitor_reset_all [1737]	Yes	Yes	Yes		Global	Yes
Innodb_num_open_files [592]				Yes	Global	No
innodb_old_blocks_remove [1738]	Yes	Yes	Yes		Global	Yes
innodb_old_blocks_size [1738]	Yes	Yes	Yes		Global	Yes
innodb_online_alter_log_max_size [1739]	Yes	Yes	Yes		Global	Yes
innodb_open_files [539]	Yes	Yes	Yes		Global	No
innodb_optimize_fulltext_only [1740]	Yes	Yes	Yes		Global	Yes
Innodb_os_log_fsyncs [592]				Yes	Global	No
Innodb_os_log_pending_fsyncs [592]				Yes	Global	No
Innodb_os_log_pending_writes [593]				Yes	Global	No
Innodb_os_log_written [593]				Yes	Global	No
innodb_page_size [540]	Yes	Yes	Yes		Global	No
Innodb_page_size [593]				Yes	Global	No
Innodb_pages_created [593]				Yes	Global	No
Innodb_pages_read [593]				Yes	Global	No
Innodb_pages_written [593]				Yes	Global	No
innodb_print_all_deadlocks [1741]	Yes	Yes	Yes		Global	Yes
innodb_purge_batch_size [1741]	Yes	Yes	Yes		Global	Yes
innodb_purge_threads [1742]	Yes	Yes	Yes		Global	No
innodb_random_read_ahead [1742]	Yes	Yes	Yes		Global	Yes
innodb_read_ahead_threshold [1743]	Yes	Yes	Yes		Global	Yes
innodb_read_io_threads [1743]	Yes	Yes	Yes		Global	No
innodb_read_only [1743]	Yes	Yes	Yes		Global	No
innodb_replication_delay [1744]	Yes	Yes	Yes		Global	Yes
innodb_rollback_on_timeout [1744]	Yes	Yes	Yes		Global	No
innodb_rollback_segments [1744]	Yes	Yes	Yes		Global	Yes
Innodb_row_lock_current_waits [593]				Yes	Global	No
Innodb_row_lock_time [593]				Yes	Global	No
Innodb_row_lock_time_avg [593]				Yes	Global	No
Innodb_row_lock_time_max [593]				Yes	Global	No
Innodb_row_lock_waits [593]				Yes	Global	No
Innodb_rows_deleted [593]				Yes	Global	No
Innodb_rows_inserted [593]				Yes	Global	No
Innodb_rows_read [593]				Yes	Global	No
Innodb_rows_updated [593]				Yes	Global	No
innodb_sort_buffer_size [1745]	Yes	Yes	Yes		Global	Yes
innodb_spin_wait_delay [1745]	Yes	Yes	Yes		Global	Yes



Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">innodb_stats_auto</a>	<a href="#">Yes</a> <a href="#">calc [1746]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_method</a>	<a href="#">Yes</a> <a href="#">[1746]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_on_metadata</a>	<a href="#">Yes</a> <a href="#">data [1747]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_persistent</a>	<a href="#">Yes</a> <a href="#">[1747]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_persistent_sample_pages</a>	<a href="#">Yes</a> <a href="#">pages [1748]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_sample_pages</a>	<a href="#">Yes</a> <a href="#">pages [1748]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_stats_transient_sample_pages</a>	<a href="#">Yes</a> <a href="#">pages [1748]</a>	Yes	Yes		Global	Yes
<a href="#">innodb-status-file</a> [1698]	Yes	Yes				
<a href="#">innodb_strict_mode</a>	<a href="#">Yes</a> <a href="#">[1749]</a>	Yes	Yes		Both	Yes
<a href="#">innodb_support_xa</a>	<a href="#">Yes</a> <a href="#">[1749]</a>	Yes	Yes		Both	Yes
<a href="#">innodb_sync_array_size</a>	<a href="#">Yes</a> <a href="#">[1750]</a>	Yes	Yes		Global	No
<a href="#">innodb_sync_spin_ops</a>	<a href="#">Yes</a> <a href="#">[1750]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_table_locks</a>	<a href="#">Yes</a> <a href="#">[1751]</a>	Yes	Yes		Both	Yes
<a href="#">innodb_thread_concurrency</a>	<a href="#">Yes</a> <a href="#">[1751]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_thread_sleep_delay</a>	<a href="#">Yes</a> <a href="#">[1752]</a>	Yes	Yes		Global	Yes
<a href="#">InnoDB_truncated_status_writes</a> [593]				Yes	Global	No
<a href="#">innodb_undo_directory</a>	<a href="#">Yes</a> <a href="#">[1752]</a>	Yes	Yes		Global	No
<a href="#">innodb_undo_logs</a>	<a href="#">Yes</a> <a href="#">[1752]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_undo_tablespaces</a>	<a href="#">Yes</a> <a href="#">[1753]</a>	Yes	Yes		Global	Yes
<a href="#">innodb_use_native_aio</a>	<a href="#">Yes</a> <a href="#">[1753]</a>	Yes	Yes		Global	No
<a href="#">innodb_use_sys_malloc</a>	<a href="#">Yes</a> <a href="#">[1754]</a>	Yes	Yes		Global	No
<a href="#">innodb_version</a> [1754]			Yes		Global	No
<a href="#">innodb_write_io_threads</a>	<a href="#">Yes</a> <a href="#">[1754]</a>	Yes	Yes		Global	No
<a href="#">timed_mutexes</a> [557]	Yes	Yes	Yes		Global	Yes
<a href="#">unique_checks</a> [561]			Yes		Both	Yes

## InnoDB Command Options

- `--ignore-builtin-innodb`

<b>Command-Line Format</b>	<code>--ignore-builtin-innodb</code>
<b>Option-File Format</b>	<code>ignore-builtin-innodb</code>
<b>Option Sets Variable</b>	Yes, <code>ignore_builtin_innodb</code> [1700]
<b>Variable Name</b>	<code>ignore-builtin-innodb</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Deprecated</b>	5.2.22
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>

In MySQL 5.1, this option caused the server to behave as if the built-in InnoDB were not present, which enabled InnoDB Plugin to be used instead. In MySQL 5.6, InnoDB is the default storage



engine and `InnoDB Plugin` is not used, so this option has no effect. As of MySQL 5.6.5, it is ignored.

- `--innodb[=value]`

Controls loading of the `InnoDB` storage engine, if the server was compiled with `InnoDB` support. This option has a tristate format, with possible values of `OFF`, `ON`, or `FORCE`. See [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

To disable `InnoDB`, use `--innodb=OFF` [1698] or `--skip-innodb` [1698]. In this case, because the default storage engine is `InnoDB`, the server will not start unless you also use `--default-storage-engine` [474] and `--default-tmp-storage-engine` [474] to set the default to some other engine for both permanent and `TEMPORARY` tables.

- `--innodb-status-file`

<b>Command-Line Format</b>	<code>--innodb-status-file</code>	
<b>Option-File Format</b>	<code>innodb-status-file</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Controls whether `InnoDB` creates a file named `innodb_status.pid` in the MySQL data directory. If enabled, `InnoDB` periodically writes the output of `SHOW ENGINE INNODB STATUS` to this file.

By default, the file is not created. To create it, start `mysqld` with the `--innodb-status-file=1` option. The file is deleted during normal shutdown.

- `--skip-innodb` [1698]

Disable the `InnoDB` storage engine. See the description of `--innodb` [1698].

## InnoDB System Variables

- `daemon_memcached_enable_binlog`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon_memcached_enable_binlog=#</code>	
<b>Option-File Format</b>	<code>daemon_memcached_enable_binlog</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_enable_binlog</code> [1698]	
<b>Variable Name</b>	<code>daemon_memcached_enable_binlog</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>false</code>

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `daemon_memcached_engine_lib_name`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon<sup>1698</sup>_memcached_engine_lib_name=library</code>	

<b>Option-File Format</b>	<code>daemon_memcached_engine_lib_name</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_engine_lib_name</code> [1698]	
<b>Variable Name</b>	<code>daemon_memcached_engine_lib_name</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>innodb_engine.so</code>

Specifies the shared library that implements the `InnoDB memcached` plugin.

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `daemon_memcached_engine_lib_path`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon_memcached_engine_lib_path=directory</code>	
<b>Option-File Format</b>	<code>daemon_memcached_engine_lib_path</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_engine_lib_path</code> [1699]	
<b>Variable Name</b>	<code>daemon_memcached_engine_lib_path</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	

The path of the directory containing the shared library that implements the `InnoDB memcached` plugin.

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `daemon_memcached_option`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon_memcached_option=options</code>	
<b>Option-File Format</b>	<code>daemon_memcached_option</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_option</code> [1699]	
<b>Variable Name</b>	<code>daemon_memcached_option</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	

Space-separated options that are passed to the underlying `memcached` daemon on startup.

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `daemon_memcached_r_batch_size`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon_memcached_r_batch_size=#</code>	
<b>Option-File Format</b>	<code>daemon_memcached_r_batch_size</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_r_batch_size</code> [1699]	
<b>Variable Name</b>	<code>daemon_memcached_r_batch_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>

Specifies how many `memcached` read operations (`get`) to perform before doing a `COMMIT` to start a new transaction. Counterpart of `daemon_memcached_w_batch_size` [1700].

This value is set to 1 by default, so that any changes made to the table through SQL statements are immediately visible to the `memcached` operations. You might increase it to reduce the overhead from frequent commits on a system where the underlying table is only being accessed through the `memcached` interface. If you set the value too large, the amount of undo or redo data could impose some storage overhead, as with any long-running transaction.

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `daemon_memcached_w_batch_size`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--daemon_memcached_w_batch_size=#</code>	
<b>Option-File Format</b>	<code>daemon_memcached_w_batch_size</code>	
<b>Option Sets Variable</b>	Yes, <code>daemon_memcached_w_batch_size</code> [1700]	
<b>Variable Name</b>	<code>daemon_memcached_w_batch_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>

Specifies how many `memcached` write operations, such as `add`, `set`, or `incr`, to perform before doing a `COMMIT` to start a new transaction. Counterpart of `daemon_memcached_r_batch_size` [1699].

This value is set to 1 by default, on the assumption that any data being stored is important to preserve in case of an outage and should immediately be committed. When storing non-critical data, you might increase this value to reduce the overhead from frequent commits; but then the last  $N-1$  uncommitted write operations could be lost in case of a crash.

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `ignore_builtin_innodb`

See the description of `--ignore-builtin-innodb` [1697] under “InnoDB Command Options” earlier in this section.

- `innodb_adaptive_flushing`

<b>Command-Line Format</b>	<code>--innodb_adaptive_flushing=#</code>	
<b>Option-File Format</b>	<code>innodb_adaptive_flushing</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_adaptive_flushing</code> [1700]	
<b>Variable Name</b>	<code>innodb_adaptive_flushing</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Specifies whether to dynamically adjust the rate of flushing [dirty pages](#) in the [InnoDB buffer pool](#) based on the workload. Adjusting the flush rate dynamically is intended to avoid bursts of I/O activity. This setting is enabled by default. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_adaptive_flushing_lwm`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_adaptive_flushing_lwm=#</code>	
<b>Option-File Format</b>	<code>innodb_adaptive_flushing_lwm</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_adaptive_flushing_lwm</code> [1701]	
<b>Variable Name</b>	<code>innodb_adaptive_flushing_lwm</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10</code>
	<b>Range</b>	<code>0 .. 70</code>

Low water mark representing percentage of [redo log](#) capacity at which [adaptive flushing](#) is enabled.

- `innodb_adaptive_hash_index`

<b>Command-Line Format</b>	<code>--innodb_adaptive_hash_index=#</code>	
<b>Option-File Format</b>	<code>innodb_adaptive_hash_index</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_adaptive_hash_index</code> [1701]	
<b>Variable Name</b>	<code>innodb_adaptive_hash_index</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Whether the [InnoDB adaptive hash index](#) is enabled or disabled. The adaptive hash index feature is useful for some workloads, and not for others; conduct benchmarks with it both enabled and disabled, using realistic workloads. See [Section 14.2.3.12.6, “Adaptive Hash Indexes”](#) for details.

This variable is enabled by default. Use `--skip-innodb_adaptive_hash_index` at server startup to disable it.

- `innodb_adaptive_max_sleep_delay` [1702]

<b>Version Introduced</b>	5.6.3
<b>Command-Line Format</b>	<code>--innodb_adaptive_max_sleep_delay=#</code>
<b>Option-File Format</b>	<code>innodb_adaptive_max_sleep_delay</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_adaptive_max_sleep_delay</code> [1702]
<b>Variable Name</b>	<code>innodb_adaptive_max_sleep_delay</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>150000</code>
	<b>Range</b> <code>0 .. 1000000</code>

Allows InnoDB to automatically adjust the value of `innodb_thread_sleep_delay` [1752] up or down according to the current workload. Any non-zero value enables automated, dynamic adjustment of the `innodb_thread_sleep_delay` value, up to the maximum value specified in the `innodb_adaptive_max_sleep_delay` option. The value represents the number of microseconds. This option can be useful in busy systems, with greater than 16 InnoDB threads. (In practice, it is most valuable for MySQL systems with hundreds or thousands of simultaneous connections.)

- `innodb_additional_mem_pool_size`

<b>Version Deprecated</b>	5.6.3
<b>Command-Line Format</b>	<code>--innodb_additional_mem_pool_size=#</code>
<b>Option-File Format</b>	<code>innodb_additional_mem_pool_size</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_additional_mem_pool_size</code> [1702]
<b>Variable Name</b>	<code>innodb_additional_mem_pool_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
<b>Deprecated</b>	5.6.3
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>8388608</code>
	<b>Range</b> <code>2097152 .. 4294967295</code>

The size in bytes of a memory pool InnoDB uses to store **data dictionary** information and other internal data structures. The more tables you have in your application, the more memory you allocate here. If InnoDB runs out of memory in this pool, it starts to allocate memory from the operating system and writes warning messages to the MySQL error log. The default value is 8MB.

This variable relates to the InnoDB internal memory allocator, which is unused if `innodb_use_sys_malloc` [1754] is enabled. As of MySQL 5.6.3, `innodb_additional_mem_pool_size` is deprecated and will be removed in a future MySQL release.

- `innodb_api_bk_commit_interval`

<b>Version Introduced</b>	5.6.7	
<b>Command-Line Format</b>	<code>--innodb_api_bk_commit_interval=#</code>	
<b>Option-File Format</b>	<code>innodb_api_bk_commit_interval</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_api_bk_commit_interval</code> [1702]	
<b>Variable Name</b>	<code>innodb_api_bk_commit_interval</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>5</code>
	<b>Range</b>	<code>1 .. 1073741824</code>

Specifies how often to auto-commit idle connections that use the `InnoDB memcached` interface. See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `innodb_api_disable_rowlock`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_api_disable_rowlock=#</code>	
<b>Option-File Format</b>	<code>innodb_api_disable_rowlock</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_api_disable_rowlock</code> [1703]	
<b>Variable Name</b>	<code>innodb_api_disable_rowlock</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `innodb_api_enable_binlog`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_api_enable_binlog=#</code>	
<b>Option-File Format</b>	<code>innodb_api_enable_binlog</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_api_enable_binlog</code> [1703]	
<b>Variable Name</b>	<code>innodb_api_enable_binlog</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Lets you use the `InnoDB memcached` plugin with the MySQL `binary log`. See [Section 14.2.9, “InnoDB Integration with memcached”](#) for usage details for this option.

- `innodb_api_enable_md1`

<b>Version Introduced</b>	5.6.6
<b>Command-Line Format</b>	<code>--innodb_api_enable_md1=#</code>
<b>Option-File Format</b>	<code>innodb_api_enable_md1</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_api_enable_md1</code> [1703]
<b>Variable Name</b>	<code>innodb_api_enable_md1</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>OFF</code>

Locks the table used by the InnoDB `memcached` plugin, so that it cannot be dropped or altered by DDL through the SQL interface. See Section 14.2.9, “InnoDB Integration with memcached” for usage details for this option.

- `innodb_api_trx_level`

<b>Version Introduced</b>	5.6.6
<b>Command-Line Format</b>	<code>--innodb_api_trx_level=#</code>
<b>Option-File Format</b>	<code>innodb_api_trx_level</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_api_trx_level</code> [1704]
<b>Variable Name</b>	<code>innodb_api_trx_level</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>0</code>

Lets you control the transaction [isolation level](#) on queries processed by the `memcached` interface. See Section 14.2.9, “InnoDB Integration with memcached” for usage details for this option. The constants corresponding to the familiar names are:

- `0 = READ UNCOMMITTED` [1470]
- `1 = READ COMMITTED` [1469]
- `2 = REPEATABLE READ` [1469]
- `3 = SERIALIZABLE` [1470]
- `innodb_autoextend_increment`

<b>Command-Line Format</b>	<code>--innodb_autoextend_increment=#</code>
<b>Option-File Format</b>	<code>innodb_autoextend_increment</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_autoextend_increment</code> [1704]
<b>Variable Name</b>	<code>innodb_autoextend_increment</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>



	<b>Type</b>	numeric
	<b>Default</b>	8
	<b>Range</b>	1 .. 1000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	64
	<b>Range</b>	1 .. 1000

The increment size (in MB) for extending the size of an auto-extend [InnoDB system tablespace](#) file when it becomes full. The default value is 64 as of MySQL 5.6.6, 8 before that. This variable does not affect the per-table tablespace files that are created if you use [innodb\\_file\\_per\\_table=1 \[1718\]](#). Those files are auto-extending regardless of the value of [innodb\\_autoextend\\_increment \[1704\]](#). The initial extensions are by small amounts, after which extensions occur in increments of 4MB.

- [innodb\\_autoinc\\_lock\\_mode](#)

<b>Command-Line Format</b>	<code>--innodb_autoinc_lock_mode=#</code>	
<b>Option-File Format</b>	<code>innodb_autoinc_lock_mode</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_autoinc_lock_mode [1705]</a>	
<b>Variable Name</b>	<code>innodb_autoinc_lock_mode</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Valid Values</b>	0 1 2

The [lock mode](#) to use for generating [auto-increment](#) values. The permissible values are 0, 1, or 2, for “traditional”, “consecutive”, or “interleaved” lock mode, respectively. [Section 5.4.4, “AUTO\\_INCREMENT Handling in InnoDB”](#), describes the characteristics of these modes.

This variable has a default of 1 (“consecutive” lock mode).

- [innodb\\_buffer\\_pool\\_dump\\_at\\_shutdown](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_dump_at_shutdown=#</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_dump_at_shutdown</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_dump_at_shutdown [1705]</a>	
<b>Variable Name</b>	<code>innodb_buffer_pool_dump_at_shutdown</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean

	<b>Default</b>	OFF
--	----------------	-----

Specifies whether to record the pages cached in the InnoDB [buffer pool](#) when the MySQL server is shut down, to shorten the [warmup](#) process at the next restart. Typically used in combination with [innodb\\_buffer\\_pool\\_load\\_at\\_startup](#) [1707].

- [innodb\\_buffer\\_pool\\_dump\\_now](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_dump_now=#</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_dump_now</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_dump_now</a> [1706]	
<b>Variable Name</b>	<code>innodb_buffer_pool_dump_now</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	OFF

Immediately records the pages cached in the InnoDB [buffer pool](#). Typically used in combination with [innodb\\_buffer\\_pool\\_load\\_now](#) [1708].

- [innodb\\_buffer\\_pool\\_filename](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_filename=file</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_filename</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_filename</a> [1706]	
<b>Variable Name</b>	<code>innodb_buffer_pool_filename</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>ib_buffer_pool</code>

Specifies the file that holds the list of page numbers produced by [innodb\\_buffer\\_pool\\_dump\\_at\\_shutdown](#) [1705] or [innodb\\_buffer\\_pool\\_dump\\_now](#) [1706].

- [innodb\\_buffer\\_pool\\_instances](#)

<b>Command-Line Format</b>	<code>--innodb_buffer_pool_instances=#</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_instances</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_instances</a> [1706]	
<b>Variable Name</b>	<code>innodb_buffer_pool_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>

	<b>Default</b>	1
	<b>Range</b>	1 .. 64
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)
	<b>Range</b>	1 .. 64

The number of regions that the InnoDB buffer pool is divided into. For systems with buffer pools in the multi-gigabyte range, dividing the buffer pool into separate instances can improve concurrency, by reducing contention as different threads read and write to cached pages. Each page that is stored in or read from the buffer pool is assigned to one of the buffer pool instances randomly, using a hashing function. Each buffer pool manages its own free lists, flush lists, LRU, and all other data structures connected to a buffer pool, and is protected by its own buffer pool mutex.

This option takes effect only when you set the `innodb_buffer_pool_size` to a size of 1 gigabyte or more. The total size you specify is divided among all the buffer pools. For best efficiency, specify a combination of `innodb_buffer_pool_instances` [1706] and `innodb_buffer_pool_size` [1708] so that each buffer pool instance is at least 1 gigabyte.

Before MySQL 5.6.6, the default is 1. As of MySQL 5.6.6, the default is 8, except on 32-bit Windows systems, where the default depends on the value of `innodb_buffer_pool_size` [1708]:

- If `innodb_buffer_pool_size` [1708] is greater than 1.3GB, the default for `innodb_buffer_pool_instances` [1706] is `innodb_buffer_pool_size` [1708]/128MB, with individual memory allocation requests for each chunk. 1.3GB was chosen as the boundary at which there is significant risk for 32-bit Windows to be unable to allocate the contiguous address space needed for a single buffer pool.
- Otherwise, the default is 1.
- `innodb_buffer_pool_load_abort`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_load_abort=#</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_load_abort</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_buffer_pool_load_abort</code> [1707]	
<b>Variable Name</b>	<code>innodb_buffer_pool_load_abort</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Interrupts the process of restoring InnoDB buffer pool contents triggered by `innodb_buffer_pool_load_at_startup` [1707] or `innodb_buffer_pool_load_now` [1708].

- `innodb_buffer_pool_load_at_startup`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_load_at_startup=#</code>	
<b>Option-File Format</b>	<code>innodb_buffer_pool_load_at_startup</code>	

<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_load_at_startup</a> [1707]	
<b>Variable Name</b>	<a href="#">innodb_buffer_pool_load_at_startup</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">boolean</a>
	<b>Default</b>	OFF

Specifies that, on MySQL server startup, the InnoDB [buffer pool](#) is automatically [warmed up](#) by loading the same pages it held at an earlier time. Typically used in combination with [innodb\\_buffer\\_pool\\_dump\\_at\\_shutdown](#) [1705].

- [innodb\\_buffer\\_pool\\_load\\_now](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_buffer_pool_load_now=#</code>	
<b>Option-File Format</b>	<a href="#">innodb_buffer_pool_load_now</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_load_now</a> [1708]	
<b>Variable Name</b>	<a href="#">innodb_buffer_pool_load_now</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">boolean</a>
	<b>Default</b>	OFF

Immediately [warms up](#) the InnoDB [buffer pool](#) by loading a set of data pages, without waiting for a server restart. Can be useful to bring cache memory back to a known state during benchmarking, or to ready the MySQL server to resume its normal workload after running queries for reports or maintenance.

- [innodb\\_buffer\\_pool\\_size](#)

<b>Command-Line Format</b>	<code>--innodb_buffer_pool_size=#</code>	
<b>Option-File Format</b>	<a href="#">innodb_buffer_pool_size</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_buffer_pool_size</a> [1708]	
<b>Variable Name</b>	<a href="#">innodb_buffer_pool_size</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<a href="#">numeric</a>
	<b>Default</b>	134217728
	<b>Range</b>	1048576 .. 2**32-1

The size in bytes of the [buffer pool](#), the memory area where InnoDB caches table and index data. The default value is 128MB. The maximum value depends on the CPU architecture; the maximum is 4294967295 ( $2^{32}-1$ ) on 32-bit systems and 18446744073709551615 ( $2^{64}-1$ ) on 64-bit systems. On 32-bit systems, the CPU architecture and operating system may impose a lower practical

maximum size than the stated maximum. When the size of the buffer pool is greater than 1GB, setting `innodb_buffer_pool_instances` [1706] to a value greater than 1 can improve the scalability on a busy server.

The larger you set this value, the less disk I/O is needed to access the same data in tables more than once. On a dedicated database server, you might set this to up to 80% of the machine physical memory size. Be prepared to scale back this value if these other issues occur:

- Competition for physical memory might cause paging in the operating system.
- InnoDB reserves additional memory for buffers and control structures, so that the total allocated space is approximately 10% greater than the specified size.
- The address space must be contiguous, which can be an issue on Windows systems with DLLs that load at specific addresses.
- The time to initialize the buffer pool is roughly proportional to its size. On large installations, this initialization time might be significant. For example, on a modern Linux x86\_64 server, initialization of a 10GB buffer pool takes approximately 6 seconds. See Section 8.9.1, “The InnoDB Buffer Pool”.
- `innodb_change_buffer_max_size`

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_change_buffer_max_size=#</code>	
<b>Option-File Format</b>	<code>innodb_change_buffer_max_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_change_buffer_max_size</code> [1709]	
<b>Variable Name</b>	<code>innodb_change_buffer_max_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>25</code>
	<b>Range</b>	<code>0 .. 50</code>

Maximum size for the InnoDB `change buffer`, as a percentage of the total size of the `buffer pool`. You might increase this value for a MySQL server with heavy insert, update, and delete activity, or decrease it for a MySQL server with unchanging data used for reporting. For general I/O tuning advice, see Section 8.5.7, “Optimizing InnoDB Disk I/O”.

- `innodb_change_buffering`

<b>Command-Line Format</b>	<code>--innodb_change_buffering=#</code>	
<b>Option-File Format</b>	<code>innodb_change_buffering</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_change_buffering</code> [1709]	
<b>Variable Name</b>	<code>innodb_change_buffering</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>all</code>

	<b>Valid Values</b>	inserts deletes purges changes all none
--	---------------------	--

Whether InnoDB performs [change buffering](#), an optimization that delays write operations to secondary indexes so that the I/O operations can be performed sequentially. The permitted values are `inserts` (buffer insert operations), `deletes` (buffer delete operations; strictly speaking, the writes that mark index records for later deletion during a purge operation), `changes` (buffer insert and delete-marking operations), `purges` (buffer [purge](#) operations, the writes when deleted index entries are finally garbage-collected), `all` (buffer insert, delete-marking, and purge operations) and `none` (do not buffer any operations). The default is `all`. For details, see [Section 14.2.4.2.13, “Controlling InnoDB Change Buffering”](#). For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_checksum_algorithm` [\[1710\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_checksum_algorithm=#</code>	
<b>Option-File Format</b>	<code>innodb_checksum_algorithm</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_checksum_algorithm</code> <a href="#">[1710]</a>	
<b>Variable Name</b>	<code>innodb_checksum_algorithm</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>innodb</code>
	<b>Valid Values</b>	innodb crc32 none strict_innodb strict_crc32 strict_none
	<b>Permitted Values (&gt;= 5.6.6, &lt;= 5.6.6)</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>crc32</code>
	<b>Valid Values</b>	innodb crc32 none strict_innodb

		<code>strict_crc32</code>
		<code>strict_none</code>
	<b>Permitted Values (&gt;= 5.6.7)</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>innodb</code>
	<b>Valid Values</b>	<code>innodb</code> <code>crc32</code> <code>none</code> <code>strict_innodb</code> <code>strict_crc32</code> <code>strict_none</code>

Specifies how to generate and verify the [checksum](#) stored in each disk block of each [InnoDB tablespace](#). Replaces the `innodb_checksums [1711]` option.

The value `innodb` is backward-compatible with all versions of MySQL. The value `crc32` uses an algorithm that is faster to compute the checksum for every modified block, and to check the checksums for each disk read. The value `none` writes a constant value in the checksum field rather than computing a value based on the block data. The blocks in a tablespace can use a mix of old, new, and no checksum values, being updated gradually as the data is modified; once any blocks in a tablespace are modified to use the `crc32` algorithm, the associated tables cannot be read by earlier versions of MySQL.

The default value was changed from `innodb` to `crc32` in MySQL 5.6.6, but switched back to `innodb` in 5.6.7 for improved compatibility of [InnoDB data files](#) during a downgrade to an earlier MySQL version, and for use of [MySQL Enterprise Backup](#) for backups.

The `strict_*` forms work the same as `innodb`, `crc32`, and `none`, except that [InnoDB](#) halts if it encounters a mix of checksum values in the same tablespace. You can only use these options in a completely new instance, to set up all tablespaces for the first time. The `strict_*` settings are somewhat faster, because they do not need to compute both new and old checksum values to accept both during disk reads.

For usage information, including a matrix of valid combinations of checksum values during read and write operations, see [Section 14.2.4.2.7, “Fast CRC32 Checksum Algorithm”](#).

- [innodb\\_checksums](#)

<b>Command-Line Format</b>	<code>--innodb_checksums</code>	
<b>Option-File Format</b>	<code>innodb_checksums</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_checksums [1711]</code>	
<b>Variable Name</b>	<code>innodb_checksums</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>



InnoDB can use [checksum](#) validation on all tablespace pages read from the disk to ensure extra fault tolerance against hardware faults or corrupted data files. This validation is enabled by default. Under specialized circumstances (such as when running benchmarks) this extra safety feature can be disabled with `--skip-innodb-checksums`. You can specify the method of calculating the checksum with `innodb_checksum_algorithm` [\[1710\]](#).

In MySQL 5.6.3 and higher, this option is deprecated, replaced by `innodb_checksum_algorithm` [\[1710\]](#). `innodb_checksum_algorithm=innodb` is the same as `innodb_checksums=ON` (the default). `innodb_checksum_algorithm=none` is the same as `innodb_checksums=OFF`. Remove any `innodb_checksums` options from your configuration files and startup scripts, to avoid conflicts with `innodb_checksum_algorithm`: `innodb_checksums=OFF` would automatically set `innodb_checksum_algorithm=none`; `innodb_checksums=ON` would be ignored and overridden by any other setting for `innodb_checksum_algorithm`.

- `innodb_cmp_per_index_enabled`

<b>Version Introduced</b>	5.6.7	
<b>Command-Line Format</b>	<code>--innodb_cmp_per_index_enabled=#</code>	
<b>Option-File Format</b>	<code>innodb_cmp_per_index_enabled</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_cmp_per_index_enabled</code> <a href="#">[1712]</a>	
<b>Variable Name</b>	<code>innodb_cmp_per_index_enabled</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>
	<b>Valid Values</b>	<code>OFF</code> <code>ON</code>

Enables per-index compression-related statistics in the `INFORMATION_SCHEMA.INNODB_CMP_PER_INDEX` table. Because these statistics can be expensive to gather, only enable this option on development, test, or slave instances during performance tuning related to [InnoDB compressed](#) tables.

- `innodb_commit_concurrency`

<b>Command-Line Format</b>	<code>--innodb_commit_concurrency=#</code>	
<b>Option-File Format</b>	<code>innodb_commit_concurrency</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_commit_concurrency</code> <a href="#">[1712]</a>	
<b>Variable Name</b>	<code>innodb_commit_concurrency</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1000</code>

The number of [threads](#) that can [commit](#) at the same time. A value of 0 (the default) permits any number of [transactions](#) to commit simultaneously.

The value of `innodb_commit_concurrency` cannot be changed at runtime from zero to nonzero or vice versa. The value can be changed from one nonzero value to another.

- `innodb_compression_failure_threshold_pct`

<b>Version Introduced</b>	5.6.7	
<b>Command-Line Format</b>	<code>--innodb_compression_failure_threshold_pct=#</code>	
<b>Option-File Format</b>	<code>innodb_compression_failure_threshold_pct</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_compression_failure_threshold_pct</code> [1713]	
<b>Variable Name</b>	<code>innodb_compression_failure_threshold_pct</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>5</code>
	<b>Range</b>	<code>0 .. 100</code>

Sets the cutoff point at which MySQL begins adding padding within `compressed` pages to avoid expensive `compression failures`. A value of zero disables the mechanism that monitors compression efficiency and dynamically adjusts the padding amount.

- `innodb_compression_level`

<b>Version Introduced</b>	5.6.7	
<b>Command-Line Format</b>	<code>--innodb_compression_level=#</code>	
<b>Option-File Format</b>	<code>innodb_compression_level</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_compression_level</code> [1713]	
<b>Variable Name</b>	<code>innodb_compression_level</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>6</code>
	<b>Range</b>	<code>0 .. 9</code>

Specifies the level of zlib compression to use for InnoDB `compressed` tables and indexes.

- `innodb_compression_pad_pct_max`

<b>Version Introduced</b>	5.6.7	
<b>Command-Line Format</b>	<code>--innodb_compression_pad_pct_max=#</code>	
<b>Option-File Format</b>	<code>innodb_compression_pad_pct_max</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_compression_pad_pct_max</code> [1713]	
<b>Variable Name</b>	<code>innodb_compression_pad_pct_max</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	

<b>Type</b>	numeric
<b>Default</b>	50
<b>Range</b>	0 .. 75

Specifies the maximum percentage that can be reserved as free space within each compressed page, allowing room to reorganize the data and modification log within the page when a compressed table or index is updated and the data might be recompressed. Only applies when `innodb_compression_failure_threshold_pct` [1713] is set to a non-zero value, and the rate of compression failures passes the cutoff point.

- `innodb_concurrency_tickets`

<b>Command-Line Format</b>	<code>--innodb_concurrency_tickets=#</code>
<b>Option-File Format</b>	<code>innodb_concurrency_tickets</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_concurrency_tickets</code> [1714]
<b>Variable Name</b>	<code>innodb_concurrency_tickets</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Type</b>	numeric
<b>Default</b>	500
<b>Range</b>	1 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Type</b>	numeric
<b>Default</b>	5000
<b>Range</b>	1 .. 4294967295

Determines the number of threads that can enter InnoDB concurrently. A thread is placed in a queue when it tries to enter InnoDB if the number of threads has already reached the concurrency limit. When a thread is permitted to enter InnoDB, it is given a number of “free tickets” equal to the value of `innodb_concurrency_tickets` [1714], and the thread can enter and leave InnoDB freely until it has used up its tickets. After that point, the thread again becomes subject to the concurrency check (and possible queuing) the next time it tries to enter InnoDB. The default value is 5000 as of MySQL 5.6.6, 500 before that.

- `innodb_data_file_path`

<b>Command-Line Format</b>	<code>--innodb_data_file_path=name</code>
<b>Option-File Format</b>	<code>innodb_data_file_path</code>
<b>Variable Name</b>	<code>innodb_data_file_path</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&lt;= 5.6.6)</b>
<b>Type</b>	string
<b>Default</b>	<code>ibdata1:10M:autoextend</code>
	<b>Permitted Values (&gt;= 5.6.7)</b>
<b>Type</b>	string
<b>Default</b>	<code>ibdata1:12M:autoextend</code>

The paths to individual [InnoDB data files](#) and their sizes. The full directory path to each data file is formed by concatenating [innodb\\_data\\_home\\_dir \[1715\]](#) to each path specified here. The file sizes are specified in KB, MB, or GB (1024MB) by appending [K](#), [M](#), or [G](#) to the size value. The sum of the sizes of the files must be at least slightly larger than 10MB. If you do not specify [innodb\\_data\\_file\\_path](#), the default behavior is to create a single auto-extending data file, slightly larger than 12MB, named `ibdata1`. The size limit of individual files is determined by your operating system. You can set the file size to more than 4GB on those operating systems that support big files. You can also [use raw disk partitions as data files](#). For detailed information on configuring [InnoDB tablespace](#) files, see [Section 14.2.1.2, “Configuring InnoDB”](#).

- [innodb\\_data\\_home\\_dir](#)

<b>Command-Line Format</b>	<code>--innodb_data_home_dir=path</code>	
<b>Option-File Format</b>	<code>innodb_data_home_dir</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_data_home_dir [1715]</a>	
<b>Variable Name</b>	<code>innodb_data_home_dir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The common part of the directory path for all [InnoDB data files](#) in the [system tablespace](#). This setting does not affect the location [offile-per-table](#) tablespaces when [innodb\\_file\\_per\\_table \[1718\]](#) is enabled. The default value is the MySQL data directory. If you specify the value as an empty string, you can use absolute file paths in [innodb\\_data\\_file\\_path \[1714\]](#).

- [innodb\\_disable\\_sort\\_file\\_cache](#)

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_disable_sort_file_cache=#</code>	
<b>Option-File Format</b>	<code>innodb_disable_sort_file_cache</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_disable_sort_file_cache [1715]</a>	
<b>Variable Name</b>	<code>innodb_disable_sort_file_cache</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

If enabled, this variable disables the operating system file system cache for merge-sort temporary files. The effect is to open such files with the equivalent of `O_DIRECT`. This variable was added in MySQL 5.6.4.

- [innodb\\_doublewrite](#)

<b>Command-Line Format</b>	<code>--innodb_doublewrite</code>	
<b>Option-File Format</b>	<code>innodb_doublewrite</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_doublewrite [1715]</a>	
<b>Variable Name</b>	<code>innodb_doublewrite</code>	

<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean

If this variable is enabled (the default), InnoDB stores all data twice, first to the [doublewrite buffer](#), then to the actual [data files](#). This variable can be turned off with `--skip-innodb_doublewrite` for benchmarks or cases when top performance is needed rather than concern for data integrity or possible failures.

- `innodb_fast_shutdown`

<b>Command-Line Format</b>	<code>--innodb_fast_shutdown[=#]</code>	
<b>Option-File Format</b>	<code>innodb_fast_shutdown</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_fast_shutdown</code> <a href="#">[1716]</a>	
<b>Variable Name</b>	<code>innodb_fast_shutdown</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Valid Values</b>	0 1 2

The [InnoDB shutdown](#) mode. If the value is 0, InnoDB does a [slow shutdown](#), a full [purge](#) and an insert buffer merge before shutting down. If the value is 1 (the default), InnoDB skips these operations at shutdown, a process known as a [fast shutdown](#). If the value is 2, InnoDB flushes its logs and shuts down cold, as if MySQL had crashed; no committed transactions are lost, but the [crash recovery](#) operation makes the next startup take longer.

The slow shutdown can take minutes, or even hours in extreme cases where substantial amounts of data are still buffered. Use the slow shutdown technique before upgrading or downgrading between MySQL major releases, so that all data files are fully prepared in case the upgrade process updates the file format.

Use `innodb_fast_shutdown=2` in emergency or troubleshooting situations, to get the absolute fastest shutdown if data is at risk of corruption.

- `innodb_file_format`

<b>Command-Line Format</b>	<code>--innodb_file_format=#</code>	
<b>Option-File Format</b>	<code>innodb_file_format</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_file_format</code> <a href="#">[1716]</a>	
<b>Variable Name</b>	<code>innodb_file_format</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

	<b>Default</b>	Antelope
	<b>Valid Values</b>	Antelope Barracuda

The [file format](#) to use for new InnoDB tables. Currently, [Antelope](#) and [Barracuda](#) are supported. This applies only for tables that have their own [tablespace](#), so for it to have an effect, [innodb\\_file\\_per\\_table](#) [1718] must be enabled. The [Barracuda](#) file format is required for certain InnoDB features such as [table compression](#).

- [innodb\\_file\\_format\\_check](#)

<b>Command-Line Format</b>	--innodb_file_format_check=#	
<b>Option-File Format</b>	innodb_file_format_check	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_file_format_check</a> [1717]	
<b>Variable Name</b>	<a href="#">innodb_file_format_check</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	ON

This variable can be set to 1 or 0 at server startup to enable or disable whether InnoDB checks the [file format](#) tag in the [system tablespace](#) (for example, [Antelope](#) or [Barracuda](#)). If the tag is checked and is higher than that supported by the current version of InnoDB, an error occurs and InnoDB does not start. If the tag is not higher, InnoDB sets the value of [innodb\\_file\\_format\\_max](#) [1717] to the file format tag.



**Note**

Despite the default value sometimes being displayed as [ON](#) or [OFF](#), always use the numeric values 1 or 0 to turn this option on or off in your configuration file or command line.

- [innodb\\_file\\_format\\_max](#)

<b>Command-Line Format</b>	--innodb_file_format_max=#	
<b>Option-File Format</b>	<a href="#">innodb_file_format_max</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_file_format_max</a> [1717]	
<b>Variable Name</b>	<a href="#">innodb_file_format_max</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string
	<b>Default</b>	Antelope
	<b>Valid Values</b>	Antelope Barracuda

At server startup, InnoDB sets the value of this variable to the [file format](#) tag in the [system tablespace](#) (for example, [Antelope](#) or [Barracuda](#)). If the server creates or opens a table with a “higher” file format, it sets the value of [innodb\\_file\\_format\\_max](#) [1717] to that format.

- `innodb_file_per_table`

<b>Command-Line Format</b>	<code>--innodb_file_per_table</code>	
<b>Option-File Format</b>	<code>innodb_file_per_table</code>	
<b>Variable Name</b>	<code>innodb_file_per_table</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&gt;= 5.6.0, &lt;= 5.6.5)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

When `innodb_file_per_table` is enabled (the default in 5.6.6 and higher), InnoDB stores the data and indexes for each newly created table in a separate `.ibd` file, rather than in the system tablespace. The storage for these InnoDB tables is reclaimed when such tables are dropped or truncated. This setting enables several other InnoDB features, such as table [compression](#). See [Section 5.4.1, “Managing InnoDB Tablespaces”](#) for details about such features.

When `innodb_file_per_table` is disabled, InnoDB stores the data for all tables and indexes in the `ibdata` files that make up the [system tablespace](#). This setting reduces the performance overhead of filesystem operations for operations such as `DROP TABLE` or `TRUNCATE TABLE`. It is most appropriate for a server environment where entire storage devices are devoted to MySQL data. Because the system tablespace never shrinks, and is shared across all databases in an [instance](#), avoid loading huge amounts of temporary data on a space-constrained system when `innodb_file_per_table=OFF`. Set up a separate instance in such cases, so that you can drop the entire instance to reclaim the space.

By default, `innodb_file_per_table` is enabled as of MySQL 5.6.6, disabled before that. Consider disabling it if backward compatibility with MySQL 5.5 or 5.1 is a concern. This will prevent `ALTER TABLE` from moving InnoDB tables from the system tablespace to individual `.ibd` files.

- `innodb_flush_log_at_timeout`

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>innodb_flush_log_at_timeout</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>
	<b>Range</b>	<code>0 .. 2700</code>

Write and flush the logs every *N* seconds. This setting has an effect only when `innodb_flush_log_at_trx_commit` [\[1718\]](#) has a value of 2.

This variable was added in MySQL 5.6.6.

- `innodb_flush_log_at_trx_commit`

<b>Command-Line Format</b>	<code>--innodb_flush_log_at_trx_commit[=#]</code>
----------------------------	---



<b>Option-File Format</b>	<code>innodb_flush_log_at_trx_commit</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_flush_log_at_trx_commit</code> [1718]	
<b>Variable Name</b>	<code>innodb_flush_log_at_trx_commit</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>1</code>
	<b>Valid Values</b>	<code>0</code> <code>1</code> <code>2</code>

Controls the balance between strict [ACID](#) compliance for [commit](#) operations, and higher performance that is possible when commit-related I/O operations are rearranged and done in batches. You can achieve better performance by changing the default value, but then you can lose up to one second worth of [transactions](#) in a crash.

- The default value of 1 is required for full ACID compliance. With this value, the [log buffer](#) is written out to the [log file](#) at each transaction commit and the [flush](#) to disk operation is performed on the log file.
- With a value of 0, any `mysqld` process crash can erase the last second of transactions. The log buffer is written out to the log file once per second and the flush to disk operation is performed on the log file, but no writes are done at a transaction commit.
- With a value of 2, only an operating system crash or a power outage can erase the last second of transactions. The log buffer is written out to the file at each commit, but the flush to disk operation is not performed on it. Before MySQL 5.6.6, the flushing on the log file takes place once per second. Note that the once-per-second flushing is not 100% guaranteed to happen every second, due to process scheduling issues. As of MySQL 5.6.6, flushing frequency is controlled by `innodb_flush_log_at_timeout` [1718] instead.
- InnoDB's [crash recovery](#) works regardless of the value. Transactions are either applied entirely or erased entirely.

For the greatest possible durability and consistency in a replication setup using InnoDB with transactions, use `innodb_flush_log_at_trx_commit=1` and `sync_binlog=1` in your master server `my.cnf` file.



### Caution

Many operating systems and some disk hardware fool the flush-to-disk operation. They may tell `mysqld` that the flush has taken place, even though it has not. Then the durability of transactions is not guaranteed even with the setting 1, and in the worst case a power outage can even corrupt InnoDB data. Using a battery-backed disk cache in the SCSI disk controller or in the disk itself speeds up file flushes, and makes the operation safer. You can also try using the Unix command `hdparm` to disable the caching of disk writes in hardware caches, or use some other command specific to the hardware vendor.

- `innodb_flush_method`

<b>Command-Line Format</b>	<code>--innodb_flush_method=name</code>
----------------------------	---

<b>Option-File Format</b>	innodb_flush_method	
<b>Option Sets Variable</b>	Yes, innodb_flush_method [1719]	
<b>Variable Name</b>	innodb_flush_method	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.6)</b>	
	<b>Type</b> (solaris)	enumeration
	<b>Default</b>	fdatasync
	<b>Valid Values</b>	O_DSYNC O_DIRECT
	<b>Permitted Values (&lt;= 5.6.6)</b>	
	<b>Type</b> (linux)	enumeration
	<b>Default</b>	fdatasync
	<b>Valid Values</b>	O_DSYNC O_DIRECT
	<b>Permitted Values (&lt;= 5.6.6)</b>	
	<b>Type</b> (hpux)	enumeration
	<b>Default</b>	fdatasync
	<b>Valid Values</b>	O_DSYNC O_DIRECT
	<b>Permitted Values (&gt;= 5.6.7)</b>	
	<b>Type</b> (linux)	enumeration
	<b>Default</b>	fdatasync
	<b>Valid Values</b>	fdatasync O_DSYNC O_DIRECT O_DIRECT_NO_FSYNC
	<b>Permitted Values (&gt;= 5.6.7)</b>	
	<b>Type</b> (hpux)	enumeration
	<b>Default</b>	fdatasync
	<b>Valid Values</b>	fdatasync O_DSYNC O_DIRECT O_DIRECT_NO_FSYNC
	<b>Permitted Values (&gt;= 5.6.7)</b>	

<b>Type</b> (solaris)	enumeration
<b>Default</b>	<code>fdatasync</code>
<b>Valid Values</b>	<code>fdatasync</code>  <code>O_DSYNC</code>  <code>O_DIRECT</code>  <code>O_DIRECT_NO_FSYNC</code>

Controls the system calls used to flush data to the InnoDB data files and log files, which can influence I/O throughput. This variable is relevant only for Unix and Linux systems. On Windows systems, the flush method is always `async_unbuffered` and cannot be changed.

By default, InnoDB uses the `fsync()` system call to flush both the data and log files. If `innodb_flush_method` [1719] option is set to `O_DSYNC`, InnoDB uses `O_SYNC` to open and flush the log files, and `fsync()` to flush the data files. If `O_DIRECT` is specified (available on some GNU/Linux versions, FreeBSD, and Solaris), InnoDB uses `O_DIRECT` (or `directio()` on Solaris) to open the data files, and uses `fsync()` to flush both the data and log files. Note that InnoDB uses `fsync()` instead of `fdatasync()`, and it does not use `O_DSYNC` by default because there have been problems with it on many varieties of Unix.

An alternative setting is `O_DIRECT_NO_FSYNC`: it uses the `O_DIRECT` flag during flushing I/O, but skips the `fsync()` system call afterwards. This setting is suitable for some types of filesystems but not others. For example, it is not suitable for XFS. If you are not sure whether the filesystem you use requires an `fsync()`, for example to preserve all file metadata, use `O_DIRECT` instead.

Depending on hardware configuration, setting `innodb_flush_method` to `O_DIRECT` or `O_DIRECT_NO_FSYNC` can have either a positive or negative effect on performance. Benchmark your particular configuration to decide which setting to use, or whether to keep the default. Examine the `InnoDB_data_fsyncs` [591] status variable to see the overall number of `fsync()` calls done with each setting. The mix of read and write operations in your workload can also affect which setting performs better for you. For example, on a system with a hardware RAID controller and battery-backed write cache, `O_DIRECT` can help to avoid double buffering between the InnoDB buffer pool and the operating system's filesystem cache. On some systems where InnoDB data and log files are located on a SAN, the default value or `O_DSYNC` might be faster for a read-heavy workload with mostly `SELECT` statements. Always test this parameter with the same type of hardware and workload that reflects your production environment. For general I/O tuning advice, see [Section 8.5.7, "Optimizing InnoDB Disk I/O"](#).

Formerly, a value of `fdatasync` also specified the default behavior. This value was removed, due to confusion that a value of `fdatasync` caused `fsync()` system calls rather than `fdatasync()` for flushing. To obtain the default value now, do not set any value for `innodb_flush_method` [1719] at startup.

- `innodb_flush_neighbors`

<b>Version Introduced</b>	5.6.3
<b>Command-Line Format</b>	<code>--innodb_flush_neighbors</code>
<b>Option-File Format</b>	<code>innodb_flush_neighbors</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_flush_neighbors</code> [1721]
<b>Variable Name</b>	<code>innodb_flush_neighbors</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes

	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	ON	

Specifies whether [flushing](#) a page from the InnoDB [buffer pool](#) also flushes other [dirty pages](#) in the same [extent](#). When the table data is stored on a traditional [HDD](#) storage device, flushing such [neighbor pages](#) in one operation reduces I/O overhead (primarily for disk seek operations) compared to flushing individual pages at different times. For table data stored on [SSD](#), seek time is not a significant factor and you can turn this setting off to spread out the write operations. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- [innodb\\_flushing\\_avg\\_loops](#)

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	--innodb_flushing_avg_loops=#	
<b>Option-File Format</b>	<a href="#">innodb_flushing_avg_loops</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_flushing_avg_loops</a> [1722]	
<b>Variable Name</b>	<a href="#">innodb_flushing_avg_loops</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	30	
<b>Range</b>	1 .. 1000	

Number of iterations for which InnoDB keeps the previously calculated snapshot of the flushing state, controlling how quickly [adaptive flushing](#) responds to changing [workloads](#). Increasing the value makes the the rate of [flush](#) operations change smoothly and gradually as the workload changes. Decreasing the value makes adaptive flushing adjust quickly to workload changes, which can cause spikes in flushing activity if the workload increases and decreases suddenly.

- [innodb\\_force\\_load\\_corrupted](#) [1722]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	--innodb_force_load_corrupted	
<b>Option-File Format</b>	<a href="#">innodb_force_load_corrupted</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_force_load_corrupted</a> [1722]	
<b>Variable Name</b>	<a href="#">innodb_force_load_corrupted</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	OFF	

Lets InnoDB load tables at startup that are marked as corrupted. Use only during troubleshooting, to recover data that is otherwise inaccessible. When troubleshooting is complete, turn this setting back off and restart the server.

- [innodb\\_force\\_recovery](#)

<b>Command-Line Format</b>	<code>--innodb_force_recovery=#</code>	
<b>Option-File Format</b>	<code>innodb_force_recovery</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_force_recovery</code> [1722]	
<b>Variable Name</b>	<code>innodb_force_recovery</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	0
	<b>Valid Values</b>	0 1 2 3 4 5 6

The [crash recovery](#) mode, typically only changed in serious troubleshooting situations. Possible values are from 0 to 6. The meanings of these values are described in [Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”](#).



### Warning

Only set this variable greater than 0 in an emergency situation, to dump your tables from a corrupt database. As a safety measure, InnoDB prevents any changes to its data when this variable is greater than 0. This restriction also prohibits some queries that use `WHERE` or `ORDER BY` clauses, because high values can prevent queries from using indexes, to guard against possible corrupt index data.

- [innodb\\_ft\\_aux\\_table](#) [1723]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_aux_table=db_name/table_name</code>	
<b>Option-File Format</b>	<code>innodb_ft_aux_table</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_aux_table</code> [1723]	
<b>Variable Name</b>	<code>innodb_ft_aux_table</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

Specifies the qualified name of an InnoDB table containing a `FULLTEXT` index. After you set this variable to a name in the format `db_name/table_name` the `INFORMATION_SCHEMA` tables `INNODB_FT_INDEX_TABLE`, `INNODB_FT_INDEX_CACHE`, `INNODB_FT_CONFIG`,

`INNODB_FT_DELETED`, and `INNODB_FT_BEING_DELETED` then reflect information about the search index for the specified table.

- `innodb_ft_cache_size` [1724]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_cache_size=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_cache_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_cache_size</code> [1724]	
<b>Variable Name</b>	<code>innodb_ft_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&gt;= 5.6.4, &lt;= 5.6.9)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>32000000</code>
	<b>Permitted Values (&gt;= 5.6.10)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>8000000</code>

Size of the cache that holds a parsed document in memory while creating an InnoDB `FULLTEXT` index.

- `innodb_ft_enable_stopword` [1724]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_enable_stopword=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_enable_stopword</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_enable_stopword</code> [1724]	
<b>Variable Name</b>	<code>innodb_ft_enable_stopword</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Specifies that a set of `stopwords` is associated with an InnoDB `FULLTEXT` index at the time the index is created. If the `innodb_ft_user_stopword_table` [1726] option is set, the stopwords are taken from that table. Else, if the `innodb_ft_server_stopword_table` [1725] option is set, the stopwords are taken from that table. Otherwise, a built-in set of default stopwords is used.

- `innodb_ft_max_token_size` [1724]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_max_token_size=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_max_token_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_max_token_size</code> [1724]	
<b>Variable Name</b>	<code>innodb_ft_max_token_size</code>	
<b>Variable Scope</b>	Global	

<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	84
	<b>Range</b>	10 .. 252

Maximum length of words that are stored in an InnoDB `FULLTEXT` index. Setting a limit on this value reduces the size of the index, thus speeding up queries, by omitting long keywords or arbitrary collections of letters that are not real words and are not likely to be search terms.

- [innodb\\_ft\\_min\\_token\\_size \[1725\]](#)

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_min_token_size=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_min_token_size</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_ft_min_token_size [1725]</a>	
<b>Variable Name</b>	<code>innodb_ft_min_token_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	3
	<b>Range</b>	0 .. 16

Minimum length of words that are stored in an InnoDB `FULLTEXT` index. Increasing this value reduces the size of the index, thus speeding up queries, by omitting common word that are unlikely to be significant in a search context, such as the English words “a” and “to”. For content using a CJK (Chinese, Japanese, Korean) character set, specify a value of 1.

- [innodb\\_ft\\_num\\_word\\_optimize \[1725\]](#)

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_num_word_optimize=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_num_word_optimize</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_ft_num_word_optimize [1725]</a>	
<b>Variable Name</b>	<code>innodb_ft_num_word_optimize</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	2000

Number of words to process during each `OPTIMIZE TABLE` operation on an InnoDB `FULLTEXT` index. Because a bulk insert or update operation to a table containing a full-text search index could require substantial index maintenance to incorporate all changes, you might do a series of `OPTIMIZE TABLE` statements, each picking up where the last left off.

- [innodb\\_ft\\_server\\_stopword\\_table \[1725\]](#)

<b>Version Introduced</b>	5.6.4
---------------------------	-------



<b>Command-Line Format</b>	<code>--innodb_ft_server_stopword_table=db_name/table_name</code>	
<b>Option-File Format</b>	<code>innodb_ft_server_stopword_table</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_server_stopword_table</code> [1725]	
<b>Variable Name</b>	<code>innodb_ft_server_stopword_table</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>NULL</code>

Name of the table containing a list of words to ignore when creating an InnoDB `FULLTEXT` index, in the format `db_name/table_name`.



#### Note

The stopword table must be an InnoDB table, containing a single `VARCHAR` column named `VALUE`. The stopword table must exist before you specify its name in the configuration option value.

- `innodb_ft_sort_pll_degree` [1726]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_sort_pll_degree=#</code>	
<b>Option-File Format</b>	<code>innodb_ft_sort_pll_degree</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_sort_pll_degree</code> [1726]	
<b>Variable Name</b>	<code>innodb_ft_sort_pll_degree</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>2</code>
	<b>Range</b>	<code>1 .. 32</code>

Number of threads used in parallel to index and tokenize text in an InnoDB `FULLTEXT` index, when building a `search index` for a large table.

- `innodb_ft_user_stopword_table` [1726]

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_ft_user_stopword_table=db_name/table_name</code>	
<b>Option-File Format</b>	<code>innodb_ft_user_stopword_table</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_ft_user_stopword_table</code> [1726]	
<b>Variable Name</b>	<code>innodb_ft_user_stopword_table</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>NULL</code>

Name of the table containing a list of words to ignore when creating an InnoDB `FULLTEXT` index, in the format `db_name/table_name`.



### Note

The stopword table must be an InnoDB table, containing a single `VARCHAR` column named `VALUE`. The stopword table must exist before you specify its name in the configuration option value.

- `innodb_io_capacity`

<b>Command-Line Format</b>	<code>--innodb_io_capacity=#</code>
<b>Option-File Format</b>	<code>innodb_io_capacity</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_io_capacity</code> [1727]
<b>Variable Name</b>	<code>innodb_io_capacity</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 200
	<b>Range</b> 100 .. $2^{32}-1$
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 200
	<b>Range</b> 100 .. $2^{64}-1$

The `innodb_io_capacity` [1727] parameter sets an upper limit, per buffer pool instance, on the I/O activity performed by the InnoDB background tasks, such as `flushing` pages from the `buffer pool` and merging data from the `insert buffer`. The default value is 200. For busy systems capable of higher I/O rates, you can set a higher value at server startup, to help the server handle the background maintenance work associated with a high rate of row changes.



### Note

Specifying a server startup value for `innodb_io_capacity` [1727] will render `innodb_io_capacity_max` [1728] inoperable.

For systems with individual 5400 RPM or 7200 RPM drives, you might lower the value to the former default of 100.

This parameter should be set to approximately the number of I/O operations that the system can perform per second. Ideally, keep this setting as low as practical, but not so low that these background activities fall behind. If the value is too high, data is removed from the buffer pool and insert buffer too quickly to provide significant benefit from the caching.

The value represents an estimated proportion of the I/O operations per second (IOPS) available to older-generation disk drives that could perform about 100 IOPS. The current default of 200 reflects that modern storage devices are capable of much higher I/O rates.

In general, you can increase the value as a function of the number of drives used for InnoDB I/O, particularly fast drives capable of high numbers of IOPS. For example, systems that use multiple disks or solid-state disks for InnoDB are likely to benefit from the ability to control this parameter.

Although you can specify a very high number, in practice such large values have little if any benefit; for example, a value of one million would be considered very high.

You can set the `innodb_io_capacity` value to any number 100 or greater, and the default value is 200. You can set the value of this parameter in the MySQL option file (`my.cnf` or `my.ini`) or change it dynamically with the `SET GLOBAL` command, which requires the `SUPER` privilege.

See [Section 14.2.4.2.20, “Controlling the InnoDB Master Thread I/O Rate”](#) for more guidelines about this option. For general information about InnoDB I/O performance, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_io_capacity_max`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_io_capacity_max=#</code>	
<b>Option-File Format</b>	<code>innodb_io_capacity_max</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_io_capacity_max [1728]</a>	
<b>Variable Name</b>	<code>innodb_io_capacity_max</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>see formula in description</code>
	<b>Min Value</b>	<code>2000</code>
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>
	<b>Max Value</b>	<code>18446744073709547520</code>

The limit up to which InnoDB is allowed to extend the `innodb_io_capacity [1727]` setting, per buffer pool instance, in case of emergency. Its default value is twice the default value of `innodb_io_capacity [1727]`, with a lower limit of 2000. It is inoperative if you have specified any value for `innodb_io_capacity [1727]` at server startup.

For a brief period during MySQL 5.6 development, this variable was known as `innodb_max_io_capacity`.

- `innodb_large_prefix`

<b>Version Introduced</b>	5.6.3
<b>Command-Line Format</b>	<code>--innodb_large_prefix</code>
<b>Option-File Format</b>	<code>innodb_large_prefix</code>
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_large_prefix [1728]</a>
<b>Variable Name</b>	<code>innodb_large_prefix</code>

<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	boolean
<b>Default</b>	OFF

Enable this option to allow [index key prefixes](#) longer than 767 bytes (up to 3072 bytes), for [InnoDB](#) tables that use the [DYNAMIC](#) and [COMPRESSED](#) row formats. (Creating such tables also requires the option values `innodb_file_format=barracuda` [1716] and `innodb_file_per_table=true` [1718].) See [Section 14.2.7, “Limits on InnoDB Tables”](#) for the relevant maximums associated with index key prefixes under various settings.

For tables using the [REDUNDANT](#) and [COMPACT](#) row formats, this option does not affect the allowed key prefix length. It does introduce a new error possibility. When this setting is enabled, attempting to create an index prefix with a key length greater than 3072 for a [REDUNDANT](#) or [COMPACT](#) table causes an error `ER_INDEX_COLUMN_TOO_LONG` (1727).

- `innodb_lock_wait_timeout`

<b>Command-Line Format</b>	<code>--innodb_lock_wait_timeout=#</code>
<b>Option-File Format</b>	<code>innodb_lock_wait_timeout</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_lock_wait_timeout</code> [1729]
<b>Variable Name</b>	<code>innodb_lock_wait_timeout</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	50
<b>Range</b>	1 .. 1073741824

The timeout in seconds an [InnoDB transaction](#) waits for a [row lock](#) before giving up. The default value is 50 seconds. A transaction that tries to access a row that is locked by another [InnoDB](#) transaction waits at most this many seconds for write access to the row before issuing the following error:

```
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting transaction
```

When a lock wait timeout occurs, the current statement is [rolled back](#) (not the entire transaction). To have the entire transaction roll back, start the server with the `--innodb_rollback_on_timeout` [1744] option. See also [Section 14.2.3.14, “InnoDB Error Handling”](#).

You might decrease this value for highly interactive applications or [OLTP](#) systems, to display user feedback quickly or put the update into a queue for processing later. You might increase this value for long-running back-end operations, such as a transform step in a data warehouse that waits for other large insert or update operations to finish.

`innodb_lock_wait_timeout` applies to [InnoDB](#) row locks only. A MySQL [table lock](#) does not happen inside [InnoDB](#) and this timeout does not apply to waits for table locks.

The lock wait timeout value does not apply to [deadlocks](#), because [InnoDB](#) detects them immediately and rolls back one of the deadlocked transactions.

- `innodb_locks_unsafe_for_binlog`

<b>Version Deprecated</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_locks_unsafe_for_binlog</code>	
<b>Option-File Format</b>	<code>innodb_locks_unsafe_for_binlog</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_locks_unsafe_for_binlog</code> [1729]	
<b>Variable Name</b>	<code>innodb_locks_unsafe_for_binlog</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Deprecated</b>	5.6.3	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

This variable affects how InnoDB uses [gap locking](#) for searches and index scans. As of MySQL 5.6.3, `innodb_locks_unsafe_for_binlog` is deprecated and will be removed in a future MySQL release.

Normally, InnoDB uses an algorithm called next-key locking that combines index-row locking with [gap locking](#). InnoDB performs row-level locking in such a way that when it searches or scans a table index, it sets shared or exclusive locks on the index records it encounters. Thus, the row-level locks are actually index-record locks. In addition, a next-key lock on an index record also affects the “gap” before that index record. That is, a next-key lock is an index-record lock plus a gap lock on the gap preceding the index record. If one session has a shared or exclusive lock on record *R* in an index, another session cannot insert a new index record in the gap immediately before *R* in the index order. See [Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#).

By default, the value of `innodb_locks_unsafe_for_binlog` is 0 (disabled), which means that gap locking is enabled: InnoDB uses next-key locks for searches and index scans. To enable the variable, set it to 1. This causes gap locking to be disabled: InnoDB uses only index-record locks for searches and index scans.

Enabling `innodb_locks_unsafe_for_binlog` does not disable the use of gap locking for foreign-key constraint checking or duplicate-key checking.

The effect of enabling `innodb_locks_unsafe_for_binlog` is similar to but not identical to setting the transaction isolation level to `READ COMMITTED` [1469]:

- Enabling `innodb_locks_unsafe_for_binlog` [1729] is a global setting and affects all sessions, whereas the isolation level can be set globally for all sessions, or individually per session.
- `innodb_locks_unsafe_for_binlog` [1729] can be set only at server startup, whereas the isolation level can be set at startup or changed at runtime.

`READ COMMITTED` [1469] therefore offers finer and more flexible control than `innodb_locks_unsafe_for_binlog` [1729]. For additional details about the effect of isolation level on gap locking, see [Section 13.3.6, “SET TRANSACTION Syntax”](#).

Enabling `innodb_locks_unsafe_for_binlog` may cause phantom problems because other sessions can insert new rows into the gaps when gap locking is disabled. Suppose that there is an index on the `id` column of the `child` table and that you want to read and lock all rows from the table having an identifier value larger than 100, with the intention of updating some column in the selected rows later:

```
SELECT * FROM child WHERE id > 100 FOR UPDATE;
```

The query scans the index starting from the first record where `id` is greater than 100. If the locks set on the index records in that range do not lock out inserts made in the gaps, another session can insert a new row into the table. Consequently, if you were to execute the same `SELECT` again within the same transaction, you would see a new row in the result set returned by the query. This also means that if new items are added to the database, InnoDB does not guarantee serializability. Therefore, if `innodb_locks_unsafe_for_binlog` is enabled, InnoDB guarantees at most an isolation level of `READ COMMITTED` [1469]. (Conflict serializability is still guaranteed.) For additional information about phantoms, see [Section 14.2.3.6, “Avoiding the Phantom Problem Using Next-Key Locking”](#).

Enabling `innodb_locks_unsafe_for_binlog` has additional effects:

- For `UPDATE` or `DELETE` statements, InnoDB holds locks only for rows that it updates or deletes. Record locks for nonmatching rows are released after MySQL has evaluated the `WHERE` condition. This greatly reduces the probability of deadlocks, but they can still happen.
- For `UPDATE` statements, if a row is already locked, InnoDB performs a “semi-consistent” read, returning the latest committed version to MySQL so that MySQL can determine whether the row matches the `WHERE` condition of the `UPDATE`. If the row matches (must be updated), MySQL reads the row again and this time InnoDB either locks it or waits for a lock on it.

Consider the following example, beginning with this table:

```
CREATE TABLE t (a INT NOT NULL, b INT) ENGINE = InnoDB;
INSERT INTO t VALUES (1,2),(2,3),(3,2),(4,3),(5,2);
COMMIT;
```

In this case, table has no indexes, so searches and index scans use the hidden clustered index for record locking (see [Section 14.2.3.12.2, “Clustered and Secondary Indexes”](#)).

Suppose that one client performs an `UPDATE` using these statements:

```
SET autocommit = 0;
UPDATE t SET b = 5 WHERE b = 3;
```

Suppose also that a second client performs an `UPDATE` by executing these statements following those of the first client:

```
SET autocommit = 0;
UPDATE t SET b = 4 WHERE b = 2;
```

As InnoDB executes each `UPDATE`, it first acquires an exclusive lock for each row, and then determines whether to modify it. If InnoDB does not modify the row and `innodb_locks_unsafe_for_binlog` is enabled, it releases the lock. Otherwise, InnoDB retains the lock until the end of the transaction. This affects transaction processing as follows.

If `innodb_locks_unsafe_for_binlog` is disabled, the first `UPDATE` acquires x-locks and does not release any of them:

```
x-lock(1,2); retain x-lock
x-lock(2,3); update(2,3) to (2,5); retain x-lock
x-lock(3,2); retain x-lock
x-lock(4,3); update(4,3) to (4,5); retain x-lock
x-lock(5,2); retain x-lock
```

The second `UPDATE` blocks as soon as it tries to acquire any locks (because first update has retained locks on all rows), and does not proceed until the first `UPDATE` commits or rolls back:

```
x-lock(1,2); block and wait for first UPDATE to commit or roll back
```

If `innodb_locks_unsafe_for_binlog` is enabled, the first `UPDATE` acquires x-locks and releases those for rows that it does not modify:

```
x-lock(1,2); unlock(1,2)
x-lock(2,3); update(2,3) to (2,5); retain x-lock
x-lock(3,2); unlock(3,2)
x-lock(4,3); update(4,3) to (4,5); retain x-lock
x-lock(5,2); unlock(5,2)
```

For the second `UPDATE`, InnoDB does a “semi-consistent” read, returning the latest committed version of each row to MySQL so that MySQL can determine whether the row matches the `WHERE` condition of the `UPDATE`:

```
x-lock(1,2); update(1,2) to (1,4); retain x-lock
x-lock(2,3); unlock(2,3)
x-lock(3,2); update(3,2) to (3,4); retain x-lock
x-lock(4,3); unlock(4,3)
x-lock(5,2); update(5,2) to (5,4); retain x-lock
```

- `innodb_log_buffer_size`

<b>Command-Line Format</b>	<code>--innodb_log_buffer_size=#</code>	
<b>Option-File Format</b>	<code>innodb_log_buffer_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_log_buffer_size</code> [1732]	
<b>Variable Name</b>	<code>innodb_log_buffer_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>8388608</code>
	<b>Range</b>	<code>262144 .. 4294967295</code>

The size in bytes of the buffer that InnoDB uses to write to the `log files` on disk. The default value is 8MB. A large `log buffer` enables large `transactions` to run without a need to write the log to disk before the transactions `commit`. Thus, if you have transactions that update, insert, or delete many rows, making the log buffer larger saves disk I/O. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_log_compressed_pages`

<b>Version Introduced</b>	5.6.11	
<b>Command-Line Format</b>	<code>--innodb_log_compressed_pages=#</code>	
<b>Option-File Format</b>	<code>innodb_log_compressed_pages</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_log_compressed_pages</code> [1732]	
<b>Variable Name</b>	<code>innodb_log_compressed_pages</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Specifies whether images of `re-compressed pages` are stored in `InnoDB redo logs`.

This variable was added in MySQL 5.6.11.

- `innodb_log_file_size`



<b>Command-Line Format</b>	<code>--innodb_log_file_size=#</code>	
<b>Option-File Format</b>	<code>innodb_log_file_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_log_file_size</code> [1732]	
<b>Variable Name</b>	<code>innodb_log_file_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&gt;= 5.6.3, &lt;= 5.6.7)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	5242880
	<b>Range</b>	1048576 .. 512GB / <code>innodb_log_files_in_group</code>
	<b>Permitted Values (&gt;= 5.6.8)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	50331648
	<b>Range</b>	1048576 .. 512GB / <code>innodb_log_files_in_group</code>

The size in bytes of each [log file](#) in a [log group](#). The combined size of log files (`innodb_log_file_size * innodb_log_files_in_group` [1733]) can be up to 512GB. The default value is 48MB. Sensible values range from 1MB to 1/*N*-th of the size of the [buffer pool](#), where *N* is the number of log files in the group. The larger the value, the less checkpoint flush activity is needed in the buffer pool, saving disk I/O. Larger log files also make [crash recovery](#) slower, although improvements to recovery performance in MySQL 5.5 and higher make the log file size less of a consideration. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_log_files_in_group`

<b>Command-Line Format</b>	<code>--innodb_log_files_in_group=#</code>	
<b>Option-File Format</b>	<code>innodb_log_files_in_group</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_log_files_in_group</code> [1733]	
<b>Variable Name</b>	<code>innodb_log_files_in_group</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	2
	<b>Range</b>	2 .. 100

The number of [log files](#) in the [log group](#). InnoDB writes to the files in a circular fashion. The default (and recommended) value is 2. The location of these files is specified by `innodb_log_group_home_dir` [1733]. The combined size of log files (`innodb_log_file_size` [1732] \* `innodb_log_files_in_group`) can be up to 512GB.

- `innodb_log_group_home_dir`

<b>Command-Line Format</b>	<code>--innodb_log_group_home_dir=path</code>	
<b>Option-File Format</b>	<code>innodb_log_group_home_dir</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_log_group_home_dir</code> [1733]	
<b>Variable Name</b>	<code>innodb_log_group_home_dir</code>	

<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name

The directory path to the InnoDB redo log files, whose number is specified by `innodb_log_files_in_group` [1733]. If you do not specify any InnoDB log variables, the default is to create two files named `ib_logfile0` and `ib_logfile1` in the MySQL data directory. Their size is given by the size of the `innodb_log_file_size` [1732] system variable.

- `innodb_lru_scan_depth`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_lru_scan_depth=#</code>	
<b>Option-File Format</b>	<code>innodb_lru_scan_depth</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_lru_scan_depth</code> [1734]	
<b>Variable Name</b>	<code>innodb_lru_scan_depth</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	numeric
	<b>Default</b>	1024
	<b>Range</b>	100 .. 2**32-1
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	1024
	<b>Range</b>	100 .. 2**64-1

A parameter that influences the algorithms and heuristics for the flush operation for the InnoDB buffer pool. Primarily of interest to performance experts tuning I/O-intensive workloads. It specifies, per buffer pool instance, how far down the buffer pool LRU list the `page_cleaner` thread scans looking for dirty pages to flush. This is a background operation performed once a second. If you have spare I/O capacity under a typical workload, increase the value. If a write-intensive workload saturates your I/O capacity, decrease the value, especially if you have a large buffer pool. For general I/O tuning advice, see Section 8.5.7, “Optimizing InnoDB Disk I/O”.

- `innodb_max_dirty_pages_pct`

<b>Command-Line Format</b>	<code>--innodb_max_dirty_pages_pct=#</code>
<b>Option-File Format</b>	<code>innodb_max_dirty_pages_pct</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_max_dirty_pages_pct</code> [1734]
<b>Variable Name</b>	<code>innodb_max_dirty_pages_pct</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes

	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	75	
<b>Range</b>	0 .. 99	

InnoDB tries to flush data from the buffer pool so that the percentage of dirty pages does not exceed this value. Specify an integer in the range from 0 to 99. The default value is 75. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_max_dirty_pages_pct_lwm`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_max_dirty_pages_pct_lwm=#</code>	
<b>Option-File Format</b>	<code>innodb_max_dirty_pages_pct_lwm</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_max_dirty_pages_pct_lwm</code> <a href="#">[1735]</a>	
<b>Variable Name</b>	<code>innodb_max_dirty_pages_pct_lwm</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	0	
<b>Range</b>	0 .. 99	

Low water mark representing percentage of dirty pages where preflushing is enabled to control the dirty page ratio. The default of 0 disables the preflushing behavior entirely.

- `innodb_max_purge_lag`

<b>Command-Line Format</b>	<code>--innodb_max_purge_lag=#</code>	
<b>Option-File Format</b>	<code>innodb_max_purge_lag</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_max_purge_lag</code> <a href="#">[1735]</a>	
<b>Variable Name</b>	<code>innodb_max_purge_lag</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	0	
<b>Range</b>	0 .. 4294967295	

This variable controls how to delay `INSERT`, `UPDATE`, and `DELETE` operations when purge operations are lagging (see [Section 14.2.3.11, “InnoDB Multi-Versioning”](#)). The default value is 0 (no delays).

The InnoDB transaction system maintains a list of transactions that have index records delete-marked by `UPDATE` or `DELETE` operations. The length of this list represents the `purge_lag` value. When `purge_lag` exceeds `innodb_max_purge_lag` [\[1735\]](#), each `INSERT`, `UPDATE`, and `DELETE` operation is delayed.

The amount of the delay is  $((\text{purge\_lag}/\text{innodb\_max\_purge\_lag}) * 10) - 5$  milliseconds. To prevent excessive delays in extreme situations where `purge_lag` becomes huge, you can put a cap on

the amount of delay by setting the `innodb_max_purge_lag_delay` [1736] configuration option. The delay is computed in the beginning of a purge batch, every ten seconds. The operations are not delayed if purge cannot run because of an old `consistent read` view that could see the rows to be purged.

A typical setting for a problematic workload might be 1 million, assuming that transactions are small, only 100 bytes in size, and it is permissible to have 100MB of unpurged InnoDB table rows.

The lag value is displayed as the history list length in the `TRANSACTIONS` section of InnoDB Monitor output. For example, if the output includes the following lines, the lag value is 20:

```
-----
TRANSACTIONS
-----
Trx id counter 0 290328385
Purge done for trx's n:o < 0 290315608 undo n:o < 0 17
History list length 20
```

For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_max_purge_lag_delay`

<b>Version Introduced</b>	5.6.5
<b>Command-Line Format</b>	<code>--innodb_max_purge_lag_delay=#</code>
<b>Option-File Format</b>	<code>innodb_max_purge_lag_delay</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_max_purge_lag_delay</code> [1736]
<b>Variable Name</b>	<code>innodb_max_purge_lag_delay</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>0</code>
	<b>Min Value</b> <code>0</code>

Specifies the maximum delay in milliseconds for the delay imposed by the `innodb_max_purge_lag` [1735] configuration option. Any non-zero value represents an upper limit on the delay period computed from the formula based on the value of `innodb_max_purge_lag`. The default of zero means that there is no upper limit imposed on the delay interval.

For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_mirrored_log_groups`

Has no effect. This variable is deprecated as of MySQL 5.6.11 and will be removed in a future MySQL release.

- `innodb_monitor_disable`

<b>Version Introduced</b>	5.6.2
<b>Command-Line Format</b>	<code>--innodb_monitor_disable=[counter module pattern all]</code>
<b>Option-File Format</b>	<code>innodb_monitor_disable</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_monitor_disable</code> [1736]
<b>Variable Name</b>	<code>innodb_monitor_disable</code>

<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Turns off one or more [counters](#) in the `INFORMATION_SCHEMA.INNODB_METRICS` table. For usage information, see [Section 19.30.17, “The INFORMATION\\_SCHEMA INNODB\\_METRICS Table”](#).

- `innodb_monitor_enable`

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_monitor_enable=name</code>	
<b>Option-File Format</b>	<code>innodb_monitor_enable</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_monitor_enable</code> <a href="#">[1737]</a>	
<b>Variable Name</b>	<code>innodb_monitor_enable</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Turns on one or more [counters](#) in the `INFORMATION_SCHEMA.INNODB_METRICS` table. For usage information, see [Section 19.30.17, “The INFORMATION\\_SCHEMA INNODB\\_METRICS Table”](#).

- `innodb_monitor_reset`

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_monitor_reset=[counter module pattern all]</code>	
<b>Option-File Format</b>	<code>innodb_monitor_reset</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_monitor_reset</code> <a href="#">[1737]</a>	
<b>Variable Name</b>	<code>innodb_monitor_reset</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Resets to zero the count value for one or more [counters](#) in the `INFORMATION_SCHEMA.INNODB_METRICS` table. For usage information, see [Section 19.30.17, “The INFORMATION\\_SCHEMA INNODB\\_METRICS Table”](#).

- `innodb_monitor_reset_all`

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_monitor_reset_all=[counter module pattern all]</code>	
<b>Option-File Format</b>	<code>innodb_monitor_reset_all</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_monitor_reset_all</code> <a href="#">[1737]</a>	
<b>Variable Name</b>	<code>innodb_monitor_reset_all</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	

	<b>Permitted Values</b>	
<b>Type</b>	string	

Resets all values (minimum, maximum, and so on) for one or more [counters](#) in the `INFORMATION_SCHEMA.INNODB_METRICS` table. For usage information, see [Section 19.30.17, “The INFORMATION\\_SCHEMA INNODB\\_METRICS Table”](#).

- `innodb_old_blocks_pct`

<b>Command-Line Format</b>	<code>--innodb_old_blocks_pct=#</code>	
<b>Option-File Format</b>	<code>innodb_old_blocks_pct</code>	
<b>Variable Name</b>	<code>innodb_old_blocks_pct</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	37	
<b>Range</b>	5 .. 95	

Specifies the approximate percentage of the [InnoDB buffer pool](#) used for the old block [sublist](#). The range of values is 5 to 95. The default value is 37 (that is, 3/8 of the pool). Often used in combination with `innodb_old_blocks_time` [1738]. See [Section 8.9.1, “The InnoDB Buffer Pool”](#) for information about buffer pool management, such as the [LRU](#) algorithm and [eviction](#) policies.

- `innodb_old_blocks_time`

<b>Command-Line Format</b>	<code>--innodb_old_blocks_time=#</code>	
<b>Option-File Format</b>	<code>innodb_old_blocks_time</code>	
<b>Variable Name</b>	<code>innodb_old_blocks_time</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
<b>Type</b>	numeric	
<b>Default</b>	0	
<b>Range</b>	0 .. 2**32-1	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
<b>Type</b>	numeric	
<b>Default</b>	1000	
<b>Range</b>	0 .. 2**32-1	

Non-zero values protect against the [buffer pool](#) being filled up by data that is referenced only for a brief period, such as during a [full table scan](#). Increasing this value offers more protection against full table scans interfering with data cached in the buffer pool.

Specifies how long in milliseconds (ms) a block inserted into the old [sublist](#) must stay there after its first access before it can be moved to the new sublist. If the value is 0, a block inserted into the old sublist moves immediately to the new sublist the first time it is accessed, no matter how soon after insertion the access occurs. If the value is greater than 0, blocks remain in the old sublist until an access occurs at least that many ms after the first access. For example, a value of 1000 causes

blocks to stay in the old sublist for 1 second after the first access before they become eligible to move to the new sublist.

The default value is 1000 as of MySQL 5.6.6, 0 before that.

This variable is often used in combination with `innodb_old_blocks_pct` [1738]. See [Section 8.9.1, “The InnoDB Buffer Pool”](#) for information about buffer pool management, such as the LRU algorithm and [eviction](#) policies.

- `innodb_online_alter_log_max_size`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_online_alter_log_max_size=#</code>	
<b>Option-File Format</b>	<code>innodb_online_alter_log_max_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_online_alter_log_max_size</code> [1739]	
<b>Variable Name</b>	<code>innodb_online_alter_log_max_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>134217728</code>
	<b>Range</b>	<code>65536 .. 2**64-1</code>

Specifies an upper limit on the size of the temporary log files used during [online DDL](#) operations for [InnoDB](#) tables. There is one such log file for each index being created or table being altered. This log file stores data inserted, updated, or deleted in the table during the DDL operation. The temporary log file is extended when needed by the value of `innodb_sort_buffer_size` [1745], up to the maximum specified by `innodb_online_alter_log_max_size`. If any temporary log file exceeds the upper size limit, the `ALTER TABLE` operation fails and all uncommitted concurrent DML operations are rolled back. Thus, a large value for this option allows more DML to happen during an online DDL operation, but also causes a longer period at the end of the DDL operation when the table is locked to apply the data from the log.

- `innodb_open_files`

<b>Command-Line Format</b>	<code>--innodb_open_files=#</code>	
<b>Option-File Format</b>	<code>innodb_open_files</code>	
<b>Variable Name</b>	<code>innodb_open_files</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>300</code>
	<b>Range</b>	<code>10 .. 4294967295</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>
	<b>Range</b>	<code>10 .. 4294967295</code>



This variable is relevant only if you use multiple [InnoDB tablespaces](#). It specifies the maximum number of [.ibd files](#) that MySQL can keep open at one time. The minimum value is 10. As of MySQL 5.6.6, the default value is 300 if [innodb\\_file\\_per\\_table \[1718\]](#) is not enabled, and the higher of 300 and [table\\_open\\_cache \[551\]](#) otherwise. Before 5.6.6, the default value is 300.

The file descriptors used for [.ibd files](#) are for [InnoDB tables](#) only. They are independent of those specified by the [--open-files-limit \[434\]](#) server option, and do not affect the operation of the table cache. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- [innodb\\_optimize\\_fulltext\\_only](#)

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_optimize_fulltext_only=#</code>	
<b>Option-File Format</b>	<code>innodb_optimize_fulltext_only</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_optimize_fulltext_only [1740]</a>	
<b>Variable Name</b>	<code>innodb_optimize_fulltext_only</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Changes the way the [OPTIMIZE TABLE](#) statement operates on [InnoDB tables](#). Intended to be enabled temporarily, during maintenance operations for [InnoDB tables](#) with [FULLTEXT](#) indexes.

By default, [OPTIMIZE TABLE](#) reorganizes the data in the [clustered index](#) of the table. When this option is enabled, [OPTIMIZE TABLE](#) skips this reorganization of the table data, and instead processes the newly added, deleted, and updated token data for a [FULLTEXT](#) index. See [Section 14.2.3.12.3, “FULLTEXT Indexes”](#) for more information about [FULLTEXT](#) indexes for [InnoDB tables](#).

- [innodb\\_page\\_size](#)

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_page_size=#k</code>	
<b>Option-File Format</b>	<code>innodb_page_size</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_page_size [1740]</a>	
<b>Variable Name</b>	<code>innodb_page_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>16384</code>
	<b>Valid Values</b>	<code>4k</code> <code>8k</code> <code>16k</code> <code>4096</code>

	8192
	16384

Specifies the [page size](#) for all [InnoDB tablespaces](#) in a MySQL [instance](#). This value is set when the instance is created and remains constant afterwards. You can specify page size using the values [16k](#) (the default), [8k](#), or [4k](#).

The default, with the largest page size, is appropriate for a wide range of [workloads](#), particularly for queries involving table scans and DML operations involving bulk updates. Smaller page sizes might be more efficient for [OLTP](#) workloads involving many small writes, where contention can be an issue when a single page contains many rows. Smaller pages might also be efficient with [SSD](#) storage devices, which typically use small block sizes. Keeping the [InnoDB](#) page size close to the storage device block size minimizes the amount of unchanged data that is rewritten to disk. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- [innodb\\_print\\_all\\_deadlocks](#)

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_print_all_deadlocks=#</code>	
<b>Option-File Format</b>	<code>innodb_print_all_deadlocks</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_print_all_deadlocks</a> [1741]	
<b>Variable Name</b>	<code>innodb_print_all_deadlocks</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

When this option is enabled, information about all [deadlocks](#) in [InnoDB](#) user transactions is recorded in the `mysqld` [error log](#). Otherwise, you see information about only the last deadlock, using the `SHOW ENGINE INNODB STATUS` command. An occasional [InnoDB](#) deadlock is not necessarily an issue, because [InnoDB](#) detects the condition immediately, and rolls back one of the transactions automatically. You might use this option to troubleshoot why deadlocks are happening if an application does not have appropriate error-handling logic to detect the rollback and retry its operation. A large number of deadlocks might indicate the need to restructure transactions that issue [DML](#) or `SELECT ... FOR UPDATE` statements for multiple tables, so that each transaction accesses the tables in the same order, thus avoiding the deadlock condition.

- [innodb\\_purge\\_batch\\_size](#)

<b>Command-Line Format</b>	<code>--innodb_purge_batch_size=#</code>	
<b>Option-File Format</b>	<code>innodb_purge_batch_size</code>	
<b>Variable Name</b>	<code>innodb_purge_batch_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&gt;= 5.6.3)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>300</code>
	<b>Range</b>	<code>1 .. 5000</code>

The granularity of changes, expressed in units of [redo log](#) records, that trigger a [purge](#) operation, flushing the changed [buffer pool](#) blocks to disk. This option is intended for tuning performance in combination with the setting `innodb_purge_threads=n`, and typical users do not need to modify it.

- [innodb\\_purge\\_threads](#)

<b>Command-Line Format</b>	<code>--innodb_purge_threads=#</code>
<b>Option-File Format</b>	<code>innodb_purge_threads</code>
<b>Variable Name</b>	<code>innodb_purge_threads</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&gt;= 5.6.2)</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>0</code>
	<b>Range</b> <code>0 .. 32</code>
	<b>Permitted Values (&gt;= 5.6.5)</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>1</code>
	<b>Range</b> <code>1 .. 32</code>

The number of background threads devoted to the InnoDB [purge](#) operation. The new default and minimum value of 1 in MySQL 5.6.5 signifies that the purge operation is always performed by background threads, never as part of the [master thread](#). Non-zero values runs the purge operation in one or more background threads, which can reduce internal contention within InnoDB, improving scalability. Increasing the value to greater than 1 creates that many separate purge threads, which can improve efficiency on systems where [DML](#) operations are performed on multiple tables. The maximum is 32.

- [innodb\\_random\\_read\\_ahead](#)

<b>Version Introduced</b>	5.6.3
<b>Command-Line Format</b>	<code>--innodb_random_read_ahead=#</code>
<b>Option-File Format</b>	<code>innodb_random_read_ahead</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_random_read_ahead</code> [1742]
<b>Variable Name</b>	<code>innodb_random_read_ahead</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>OFF</code>

Enables the random [read-ahead](#) technique for optimizing InnoDB I/O. This is a setting that was originally on by default, then was removed in MySQL 5.5, and now is available but turned off by default. See [Section 14.2.4.2.16, “Changes in the Read-Ahead Algorithm”](#) for details about the performance considerations for the different types of read-ahead requests. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- [innodb\\_read\\_ahead\\_threshold](#)

<b>Command-Line Format</b>	<code>--innodb_read_ahead_threshold=#</code>	
<b>Option-File Format</b>	<code>innodb_read_ahead_threshold</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_read_ahead_threshold</code> [1742]	
<b>Variable Name</b>	<code>innodb_read_ahead_threshold</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>56</code>
	<b>Range</b>	<code>0 .. 64</code>

Controls the sensitivity of linear [read-ahead](#) that InnoDB uses to prefetch pages into the [buffer pool](#). If InnoDB reads at least `innodb_read_ahead_threshold` pages sequentially from an [extent](#) (64 pages), it initiates an asynchronous read for the entire following extent. The permissible range of values is 0 to 64. The default is 56: InnoDB must read at least 56 pages sequentially from an extent to initiate an asynchronous read for the following extent. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_read_io_threads`

<b>Command-Line Format</b>	<code>--innodb_read_io_threads=#</code>	
<b>Option-File Format</b>	<code>innodb_read_io_threads</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_read_io_threads</code> [1743]	
<b>Variable Name</b>	<code>innodb_read_io_threads</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4</code>
	<b>Range</b>	<code>1 .. 64</code>

The number of I/O threads for read operations in InnoDB. The default value is 4. Its counterpart for write threads is `innodb_write_io_threads` [1754]. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).



#### Note

On Linux systems, running multiple MySQL servers (typically more than 12) with default settings for `innodb_read_io_threads`, `innodb_write_io_threads` [1754], and the Linux `aio-max-nr` setting can exceed system limits. Ideally, increase the `aio-max-nr` setting; as a workaround, you might reduce the settings for one or both of the MySQL configuration options.

- `innodb_read_only`

<b>Version Introduced</b>	<code>5.6.7</code>
<b>Command-Line Format</b>	<code>--innodb_read_only=#</code>
<b>Option-File Format</b>	<code>innodb_read_only</code>

<b>Option Sets Variable</b>	Yes, <code>innodb_read_only</code> [1743]	
<b>Variable Name</b>	<code>innodb_read_only</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Starts the server in read-only mode. For distributing database applications or data sets on read-only media. Can also be used in data warehouses to share the same data directory between multiple instances. See [Section 14.2.5.1, “Support for Read-Only Media”](#) for usage instructions.

- `innodb_replication_delay`

<b>Command-Line Format</b>	<code>--innodb_replication_delay=#</code>	
<b>Option-File Format</b>	<code>innodb_replication_delay</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_replication_delay</code> [1744]	
<b>Variable Name</b>	<code>innodb_replication_delay</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 4294967295</code>

The replication thread delay (in ms) on a slave server if `innodb_thread_concurrency` [1751] is reached.

- `innodb_rollback_on_timeout`

<b>Command-Line Format</b>	<code>--innodb_rollback_on_timeout</code>	
<b>Option-File Format</b>	<code>innodb_rollback_on_timeout</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_rollback_on_timeout</code> [1744]	
<b>Variable Name</b>	<code>innodb_rollback_on_timeout</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

In MySQL 5.6, InnoDB **rolls back** only the last statement on a transaction timeout by default. If `--innodb_rollback_on_timeout` [1744] is specified, a transaction timeout causes InnoDB to abort and roll back the entire transaction (the same behavior as in MySQL 4.1).

- `innodb_rollback_segments`

<b>Version Introduced</b>	5.6.2
<b>Command-Line Format</b>	<code>--innodb_rollback_segments=#</code>
<b>Option-File Format</b>	<code>innodb_rollback_segments</code>

<b>Option Sets Variable</b>	Yes, <code>innodb_rollback_segments</code> [1744]	
<b>Variable Name</b>	<code>innodb_rollback_segments</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>128</code>
	<b>Range</b>	<code>1 .. 128</code>

Defines how many of the [rollback segments](#) in the system tablespace that InnoDB uses within a [transaction](#). You might reduce this value from its default of 128 if a smaller number of rollback segments performs better for your workload. For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

- `innodb_sort_buffer_size`

<b>Version Introduced</b>	5.6.4	
<b>Command-Line Format</b>	<code>--innodb_sort_buffer_size=#</code>	
<b>Option-File Format</b>	<code>innodb_sort_buffer_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_sort_buffer_size</code> [1745]	
<b>Variable Name</b>	<code>innodb_sort_buffer_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&gt;= 5.6.4)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1048576</code>
	<b>Range</b>	<code>524288 .. 67108864</code>
	<b>Permitted Values (&gt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1048576</code>
	<b>Range</b>	<code>65536 .. 67108864</code>

Specifies the sizes of several buffers used for sorting data during creation of an [InnoDB](#) index. Before this setting was made configurable, the size was hardcoded to 1MB, and that value remains the default. This sort area is only used for merge sorts during index creation, not during later index maintenance operations. During an `ALTER TABLE` or `CREATE TABLE` statement that creates an index, 3 buffers are allocated, each with a size defined by this option. These buffers are deallocated when the index creation completes.

The value of this option also controls the amount by which the temporary log file is extended, to record concurrent DML during [online DDL](#) operations.

- `innodb_spin_wait_delay`

<b>Command-Line Format</b>	<code>--innodb_spin_wait_delay=#</code>
<b>Option-File Format</b>	<code>innodb_spin_wait_delay</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_spin_wait_delay</code> [1745]
<b>Variable Name</b>	<code>innodb_spin_wait_delay</code>
<b>Variable Scope</b>	Global

<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	6
	<b>Range</b>	0 .. 4294967295

The maximum delay between polls for a [spin](#) lock. The low-level implementation of this mechanism varies depending on the combination of hardware and operating system, so the delay does not correspond to a fixed time interval. The default value is 6.

- [innodb\\_stats\\_auto\\_recalc](#)

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	--innodb_stats_auto_recalc=#	
<b>Option-File Format</b>	<a href="#">innodb_stats_auto_recalc</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_stats_auto_recalc</a> [1746]	
<b>Variable Name</b>	<a href="#">innodb_stats_auto_recalc</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	ON

Causes [InnoDB](#) to automatically recalculate [persistent statistics](#) after the data in a table is changed substantially. The threshold value is currently 10% of the rows in the table. This setting applies to tables created when the [innodb\\_stats\\_persistent](#) [1747] option is enabled, or where the clause `STATS_PERSISTENT=1` is enabled by a `CREATE TABLE` or `ALTER TABLE` statement. The amount of data sampled to produce the statistics is controlled by the [innodb\\_stats\\_persistent\\_sample\\_pages](#) [1748] configuration option.

- [innodb\\_stats\\_method](#)

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	--innodb_stats_method=name	
<b>Option-File Format</b>	<a href="#">innodb_stats_method</a>	
<b>Option Sets Variable</b>	Yes, <a href="#">innodb_stats_method</a> [1746]	
<b>Variable Name</b>	<a href="#">innodb_stats_method</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	enumeration
	<b>Default</b>	nulls_equal
	<b>Valid Values</b>	<a href="#">nulls_equal</a> <a href="#">nulls_unequal</a> <a href="#">nulls_ignored</a>

How the server treats `NULL` values when collecting [statistics](#) about the distribution of index values for [InnoDB](#) tables. This variable has three possible values, `nulls_equal`, `nulls_unequal`, and



`nulls_ignored`. For `nulls_equal`, all `NULL` index values are considered equal and form a single value group that has a size equal to the number of `NULL` values. For `nulls_unequal`, `NULL` values are considered unequal, and each `NULL` forms a distinct value group of size 1. For `nulls_ignored`, `NULL` values are ignored.

The method that is used for generating table statistics influences how the optimizer chooses indexes for query execution, as described in [Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”](#).

- `innodb_stats_on_metadata`

<b>Command-Line Format</b>	<code>--innodb_stats_on_metadata</code>	
<b>Option-File Format</b>	<code>innodb_stats_on_metadata</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_stats_on_metadata</code> <a href="#">[1747]</a>	
<b>Variable Name</b>	<code>innodb_stats_on_metadata</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

When this variable is enabled, InnoDB updates [statistics](#) during metadata statements such as `SHOW TABLE STATUS` or `SHOW INDEX`, or when accessing the `INFORMATION_SCHEMA` tables `TABLES` or `STATISTICS`. (These updates are similar to what happens for `ANALYZE TABLE`.) When disabled, InnoDB does not update statistics during these operations. Leaving this setting disabled can improve access speed for schemas that have a large number of tables or indexes. It can also improve the stability of [execution plans](#) for queries that involve InnoDB tables.

This variable is disabled by default as of MySQL 5.6.6, enabled before that.

- `innodb_stats_persistent`

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--innodb_stats_persistent=setting</code>	
<b>Option-File Format</b>	<code>innodb_stats_persistent</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_stats_persistent</code> <a href="#">[1747]</a>	
<b>Variable Name</b>	<code>innodb_stats_persistent</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>
	<b>Valid Values</b>	<code>OFF</code> <code>ON</code> <code>0</code> <code>1</code>

	default
--	---------

Specifies whether the InnoDB index statistics produced by the `ANALYZE TABLE` command are stored on disk, remaining consistent until a subsequent `ANALYZE TABLE`. Otherwise, the statistics are recalculated more frequently, such as at each server restart, which can lead to variations in [query execution plans](#). This setting is stored with each table when the table is created. You can specify or change it through SQL with the `STATS_PERSISTENT` clause of the `CREATE TABLE` and `ALTER TABLE` commands.

- `innodb_stats_persistent_sample_pages`

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--innodb_stats_persistent_sample_pages=#</code>	
<b>Option-File Format</b>	<code>innodb_stats_persistent_sample_pages</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_stats_persistent_sample_pages</code> [1748]	
<b>Variable Name</b>	<code>innodb_stats_persistent_sample_pages</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>20</code>

The number of index [pages](#) to sample when estimating [cardinality](#) and other [statistics](#) for an indexed column, such as those calculated by `ANALYZE TABLE`. Increasing the value improves the accuracy of index statistics, which can improve the [query execution plan](#), at the expense of increased I/O during the execution of `ANALYZE TABLE` for an InnoDB table.

This option only applies when the `innodb_stats_persistent` [1747] setting is turned on for a table; when that option is turned off for a table, the `innodb_stats_transient_sample_pages` [1748] setting applies instead.

- `innodb_stats_sample_pages`

<b>Version Deprecated</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_stats_sample_pages=#</code>	
<b>Option-File Format</b>	<code>innodb_stats_sample_pages</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_stats_sample_pages</code> [1748]	
<b>Variable Name</b>	<code>innodb_stats_sample_pages</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Deprecated</b>	5.6.3	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>8</code>
	<b>Range</b>	<code>1 .. 2**64-1</code>

Deprecated, use `innodb_stats_transient_sample_pages` [1748] instead.

- `innodb_stats_transient_sample_pages`

<b>Version Introduced</b>	5.6.2
<b>Command-Line Format</b>	<code>--innodb_stats_transient_sample_pages=#</code>
<b>Option-File Format</b>	<code>innodb_stats_transient_sample_pages</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_stats_transient_sample_pages</code> [1748]
<b>Variable Name</b>	<code>innodb_stats_transient_sample_pages</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>8</code>

The number of index [pages](#) to sample when estimating [cardinality](#) and other [statistics](#) for an indexed column, such as those calculated by `ANALYZE TABLE`. The default value is 8. Increasing the value improves the accuracy of index statistics, which can improve the [query execution plan](#), at the expense of increased I/O when opening an InnoDB table or recalculating statistics.

This option only applies when the `innodb_stats_persistent` [1747] setting is turned off for a table; when this option is turned on for a table, the `innodb_stats_persistent_sample_pages` [1748] setting applies instead. Takes the place of the `innodb_stats_sample_pages` option. For more information, see [Section 14.2.5, “InnoDB Features for Flexibility, Ease of Use and Reliability”](#).

- `innodb_strict_mode`

<b>Command-Line Format</b>	<code>--innodb_strict_mode=#</code>
<b>Option-File Format</b>	<code>innodb_strict_mode</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_strict_mode</code> [1749]
<b>Variable Name</b>	<code>innodb_strict_mode</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>OFF</code>

Whether InnoDB returns errors rather than warnings for certain conditions. This is analogous to strict SQL mode. The default value is `OFF`. See [Section 14.2.5.7, “InnoDB Strict Mode”](#) for a list of the conditions that are affected.

- `innodb_support_xa`

<b>Command-Line Format</b>	<code>--innodb_support_xa</code>
<b>Option-File Format</b>	<code>innodb_support_xa</code>
<b>Option Sets Variable</b>	Yes, <code>innodb_support_xa</code> [1749]
<b>Variable Name</b>	<code>innodb_support_xa</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>

	<b>Default</b>	TRUE
--	----------------	------

Enables InnoDB support for two-phase commit in XA transactions, causing an extra disk flush for transaction preparation. This setting is the default. The XA mechanism is used internally and is essential for any server that has its binary log turned on and is accepting changes to its data from more than one thread. If you turn it off, transactions can be written to the binary log in a different order from the one in which the live database is committing them. This can produce different data when the binary log is replayed in disaster recovery or on a replication slave. Do not turn it off on a replication master server unless you have an unusual setup where only one thread is able to change data.

For a server that is accepting data changes from only one thread, it is safe and recommended to turn off this option to improve performance for InnoDB tables. For example, you can turn it off on replication slaves where only the replication SQL thread is changing data.

You can also turn off this option if you do not need it for safe binary logging or replication, and you also do not use an external XA transaction manager.

- `innodb_sync_array_size`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_sync_array_size=#</code>	
<b>Option-File Format</b>	<code>innodb_sync_array_size</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_sync_array_size</code> [1750]	
<b>Variable Name</b>	<code>innodb_sync_array_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1
	<b>Range</b>	1 .. 1024

Splits an internal data structure used to coordinate threads, for higher concurrency in workloads with large numbers of waiting threads. This setting must be configured when the MySQL instance is starting up, and cannot be changed afterward. Increasing this option value is recommended for workloads that frequently produce a large number of waiting threads, typically greater than 768.

- `innodb_sync_spin_loops`

<b>Command-Line Format</b>	<code>--innodb_sync_spin_loops=#</code>	
<b>Option-File Format</b>	<code>innodb_sync_spin_loops</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_sync_spin_loops</code> [1750]	
<b>Variable Name</b>	<code>innodb_sync_spin_loops</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	30
	<b>Range</b>	0 .. 4294967295

The number of times a thread waits for an InnoDB mutex to be freed before the thread is suspended. The default value is 30.

- `innodb_table_locks`

<b>Command-Line Format</b>	<code>--innodb_table_locks</code>	
<b>Option-File Format</b>	<code>innodb_table_locks</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_table_locks</code> [1751]	
<b>Variable Name</b>	<code>innodb_table_locks</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>TRUE</code>

If `autocommit = 0` [464], InnoDB honors `LOCK TABLES`; MySQL does not return from `LOCK TABLES ... WRITE` until all other threads have released all their locks to the table. The default value of `innodb_table_locks` [1751] is 1, which means that `LOCK TABLES` causes InnoDB to lock a table internally if `autocommit = 0` [464].

In MySQL 5.6, `innodb_table_locks = 0` [1751] has no effect for tables locked explicitly with `LOCK TABLES ... WRITE`. It does have an effect for tables locked for read or write by `LOCK TABLES ... WRITE` implicitly (for example, through triggers) or by `LOCK TABLES ... READ`.

- `innodb_thread_concurrency`

<b>Command-Line Format</b>	<code>--innodb_thread_concurrency=#</code>	
<b>Option-File Format</b>	<code>innodb_thread_concurrency</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_thread_concurrency</code> [1751]	
<b>Variable Name</b>	<code>innodb_thread_concurrency</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1000</code>

InnoDB tries to keep the number of operating system threads concurrently inside InnoDB less than or equal to the limit given by this variable. Once the number of threads reaches this limit, additional threads are placed into a wait state within a FIFO queue for execution. Threads waiting for locks are not counted in the number of concurrently executing threads.

The correct value for this variable is dependent on environment and workload. Try a range of different values to determine what value works for your applications. A recommended value is 2 times the number of CPUs plus the number of disks.

The range of this variable is 0 to 1000. A value of 0 (the default) is interpreted as infinite concurrency (no concurrency checking). Disabling thread concurrency checking enables InnoDB to create as many threads as it needs. A value of 0 also disables the `queries inside InnoDB` and `queries in queue counters` in the `ROW OPERATIONS` section of `SHOW ENGINE INNODB STATUS` output.

- `innodb_thread_sleep_delay`

<b>Command-Line Format</b>	<code>--innodb_thread_sleep_delay=#</code>	
<b>Option-File Format</b>	<code>innodb_thread_sleep_delay</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_thread_sleep_delay</code> [1752]	
<b>Variable Name</b>	<code>innodb_thread_sleep_delay</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10000</code>

How long InnoDB threads sleep before joining the InnoDB queue, in microseconds. The default value is 10,000. A value of 0 disables sleep.

- `innodb_undo_directory`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_undo_directory=name</code>	
<b>Option-File Format</b>	<code>innodb_undo_directory</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_undo_directory</code> [1752]	
<b>Variable Name</b>	<code>innodb_undo_directory</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>.</code>

The relative or absolute directory path where InnoDB creates separate tablespaces for the undo logs. Typically used to place those logs on a different storage device. Used in conjunction with `innodb_undo_logs` [1752] and `innodb_undo_tablespaces` [1753], which determine the disk layout of the undo logs outside the `system tablespace`. Its default value of `.` represents the same directory where InnoDB creates its other log files by default.

- `innodb_undo_logs`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_undo_logs=#</code>	
<b>Option-File Format</b>	<code>innodb_undo_logs</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_undo_logs</code> [1752]	
<b>Variable Name</b>	<code>innodb_undo_logs</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>128</code>
	<b>Range</b>	<code>0 .. 128</code>

Defines how many of the [rollback segments](#) in the [system tablespace](#) that InnoDB uses within a [transaction](#). This setting is appropriate for tuning performance if you observe mutex contention related to the undo logs. Replaces the `innodb_rollback_segments` setting. For the total number of available undo logs, rather than the number of active ones, see the `InnoDB_available_undo_logs` [590] status variable.

Although you can increase or decrease how many rollback segments are used within a transaction, the number of rollback segments physically present in the system never decreases. Thus you might start with a low value for this parameter and gradually increase it, to avoid allocating rollback segments that are not needed later.

- `innodb_undo_tablespaces`

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_undo_tablespaces=#</code>	
<b>Option-File Format</b>	<code>innodb_undo_tablespaces</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_undo_tablespaces</code> [1753]	
<b>Variable Name</b>	<code>innodb_undo_tablespaces</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>126</code>
	<b>Min Value</b>	<code>0</code>

The number of [tablespace](#) files that the [undo logs](#) are divided between, when you use a non-zero `innodb_undo_logs` [1752] setting. By default, all the undo logs are part of the [system tablespace](#). Because the undo logs can become large during long-running transactions, splitting the undo logs between multiple tablespaces reduces the maximum size of any one tablespace. The tablespace files are created in the location defined by `innodb_undo_directory` [1752], with names of the form `innodbN`, where *N* is a sequential series of integers, including leading zeros.

- `innodb_use_native_aio`

<b>Command-Line Format</b>	<code>--innodb_use_native_aio=#</code>	
<b>Option-File Format</b>	<code>innodb_use_native_aio</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_use_native_aio</code> [1753]	
<b>Variable Name</b>	<code>innodb_use_native_aio</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Specifies whether to use the Linux asynchronous I/O subsystem. This variable applies to Linux systems only, and cannot be changed while the server is running.

Normally, you do not need to touch this option, because it is enabled by default. If a problem with the asynchronous I/O subsystem in the OS prevents InnoDB from starting, start the server with this variable disabled (use `innodb_use_native_aio=0` [1753] in the option file). This option could



also be turned off automatically during startup, if InnoDB detects a potential problem such as a combination of `tmpdir` location, `tmpfs` filesystem, and Linux kernel that does not support AIO on `tmpfs`.

- `innodb_use_sys_malloc`

<b>Version Deprecated</b>	5.6.3	
<b>Command-Line Format</b>	<code>--innodb_use_sys_malloc=#</code>	
<b>Option-File Format</b>	<code>innodb_use_sys_malloc</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_use_sys_malloc</code> [1754]	
<b>Variable Name</b>	<code>innodb_use_sys_malloc</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
<b>Deprecated</b>	5.6.3	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>ON</code>

Whether InnoDB uses the operating system memory allocator (`ON`) or its own (`OFF`). The default value is `ON`.

As of MySQL 5.6.3, `innodb_use_sys_malloc` [1754] is deprecated and will be removed in a future MySQL release.

- `innodb_version`

The InnoDB version number. Starting in 5.6.11, the separate numbering for InnoDB is discontinued and this value is the same as for the `version` [563] variable.

- `innodb_write_io_threads`

<b>Command-Line Format</b>	<code>--innodb_write_io_threads=#</code>	
<b>Option-File Format</b>	<code>innodb_write_io_threads</code>	
<b>Option Sets Variable</b>	Yes, <code>innodb_write_io_threads</code> [1754]	
<b>Variable Name</b>	<code>innodb_write_io_threads</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>4</code>
	<b>Range</b>	<code>1 .. 64</code>

The number of I/O threads for write operations in InnoDB. The default value is 4. Its counterpart for read threads is `innodb_read_io_threads` [1743]. For general I/O tuning advice, see Section 8.5.7, “Optimizing InnoDB Disk I/O”.



#### Note

On Linux systems, running multiple MySQL servers (typically more than 12) with default settings for `innodb_read_io_threads` [1743], `innodb_write_io_threads`, and the Linux `aio-max-nr` setting can exceed system limits. Ideally, increase the `aio-max-nr` setting; as a

workaround, you might reduce the settings for one or both of the MySQL configuration options.

- `sync_binlog`

<b>Command-Line Format</b>	<code>--sync-binlog=#</code>
<b>Option-File Format</b>	<code>sync_binlog</code>
<b>Option Sets Variable</b>	Yes, <code>sync_binlog</code> [2017]
<b>Variable Name</b>	<code>sync_binlog</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> numeric
	<b>Default</b> 0
	<b>Range</b> 0 .. 4294967295
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> numeric
	<b>Default</b> 0
	<b>Range</b> 0 .. 18446744073709547520

If the value of this variable is greater than 0, the MySQL server synchronizes its binary log to disk (using `fdatasync()`) after every `sync_binlog` [2017] writes to the binary log. There is one write to the binary log per statement if autocommit is enabled, and one write per transaction otherwise. The default value of `sync_binlog` [2017] is 0, which does no synchronizing to disk. A value of 1 is the safest choice, because in the event of a crash you lose at most one statement or transaction from the binary log. However, it is also the slowest choice (unless the disk has a battery-backed cache, which makes synchronization very fast). For general I/O tuning advice, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

## 14.2.6.1. List of Parameters Recently Changed in InnoDB

### 14.2.6.1.1. New Parameters

Throughout the course of development, InnoDB 1.1 and its predecessor the InnoDB Plugin introduced new configuration parameters.

The following InnoDB-related configuration parameters were added between MySQL 5.5 and 5.6:

```
innodb_adaptive_max_sleep_delay [1702],
innodb_buffer_pool_dump_at_shutdown [1705], innodb_buffer_pool_dump_now [1706],
innodb_buffer_pool_filename [1706], innodb_buffer_pool_load_abort [1707],
innodb_buffer_pool_load_at_startup [1707], innodb_buffer_pool_load_now [1708],
innodb_change_buffer_max_size [1709], innodb_checksum_algorithm [1710],
innodb_disable_sort_file_cache, innodb_flush_neighbors [1721],
innodb_force_load_corrupted [1722], innodb_ft_aux_table [1723],
innodb_ft_cache_size [1724], innodb_ft_enable_diag_print,
innodb_ft_enable_stopword [1724], innodb_ft_max_token_size [1724],
innodb_ft_min_token_size [1725], innodb_ft_num_word_optimize [1725],
innodb_ft_server_stopword_table [1725], innodb_ft_sort_pll_degree [1726],
innodb_ft_user_stopword_table [1726], innodb_large_prefix [1728],
```

`innodb_lru_scan_depth` [1734], `innodb_monitor_disable` [1736],  
`innodb_monitor_enable` [1737], `innodb_monitor_reset` [1737],  
`innodb_monitor_reset_all` [1737], `innodb_optimize_fulltext_only` [1740],  
`innodb_page_size` [1740], `innodb_print_all_deadlocks` [1741],  
`innodb_random_read_ahead` [1742], `innodb_sort_buffer_size` [1745],  
`innodb_stats_persistent_sample_pages` [1748],  
`innodb_stats_transient_sample_pages` [1748], `innodb_sync_array_size` [1750],  
`innodb_undo_directory` [1752], `innodb_undo_logs` [1752],  
`innodb_undo_tablespace` [1753]

#### 14.2.6.1.2. Parameters with New Defaults

For better out-of-the-box performance, the following InnoDB configuration parameters have new default values since MySQL 5.1:

**Table 14.6. InnoDB Parameters with New Defaults**

Name	Old Default	New Default
<code>innodb_additional_mem_pool_size</code>	1MB	8MB
<code>innodb_buffer_pool_size</code>	8MB	128MB
<code>innodb_change_buffering</code>	<code>inserts</code>	<code>all</code>
<code>innodb_file_format_check</code>	ON	1
<code>innodb_log_buffer_size</code>	1MB	8MB
<code>innodb_max_dirty_pages_pct</code>	90	75
<code>innodb_sync_spin_loops</code>	20	30
<code>innodb_thread_concurrency</code>	8	0

#### 14.2.7. Limits on InnoDB Tables



##### Warning

Do *not* convert MySQL system tables in the `mysql` database from `MyISAM` to `InnoDB` tables! This is an unsupported operation. If you do this, MySQL does not restart until you restore the old system tables from a backup or re-generate them with the `mysql_install_db` script.



##### Warning

It is not a good idea to configure `InnoDB` to use data files or log files on NFS volumes. Otherwise, the files might be locked by other processes and become unavailable for use by MySQL.

#### Maximums and Minimums

- A table can contain a maximum of 1020 columns. (Raised in MySQL 5.6.9 from the earlier limit of 1000.)
- A table can contain a maximum of 64 `secondary indexes`.
- By default, an index key for a single-column index can be up to 767 bytes. The same length limit applies to any index key prefix. See [Section 13.1.11, “CREATE INDEX Syntax”](#). For example, you might hit this limit with a `column prefix` index of more than 255 characters on a `TEXT` or `VARCHAR` column, assuming a UTF-8 character set and the maximum of 3 bytes for each character. When the `innodb_large_prefix` [1728] configuration option is enabled, this length limit is raised to 3072 bytes, for `InnoDB` tables that use the `DYNAMIC` and `COMPRESSED` row formats.

When you attempt to specify an index prefix length longer than allowed, the length is silently reduced to the maximum length for a nonunique index. For a unique index, exceeding the index

prefix limit produces an error. To avoid such errors for replication configurations, avoid setting the `innodb_large_prefix` [1728] option on the master if it cannot also be set on the slaves, and the slaves have unique indexes that could be affected by this limit.

This configuration option changes the error handling for some combinations of row format and prefix length longer than the maximum allowed. See `innodb_large_prefix` [1728] for details.

- The InnoDB internal maximum key length is 3500 bytes, but MySQL itself restricts this to 3072 bytes. This limit applies to the length of the combined index key in a multi-column index.
- If you reduce the InnoDB page size to 8KB or 4KB by specifying the `innodb_page_size` [1740] option when creating the MySQL instance, the maximum length of the index key is lowered proportionally, based on the limit of 3072 bytes for a 16KB page size. That is, the maximum index key length is 1536 bytes when the page size is 8KB, and 768 bytes when the page size is 4KB.
- The maximum row length, except for variable-length columns (`VARBINARY`, `VARCHAR`, `BLOB` and `TEXT`), is slightly less than half of a database page. That is, the maximum row length is about 8000 bytes for the default page size of 16KB; if you reduce the page size by specifying the `innodb_page_size` [1740] option when creating the MySQL instance, the maximum row length is 4000 bytes for 8KB pages and 2000 bytes for 4KB pages. `LONGBLOB` and `LONGTEXT` columns must be less than 4GB, and the total row length, including `BLOB` and `TEXT` columns, must be less than 4GB.

If a row is less than half a page long, all of it is stored locally within the page. If it exceeds half a page, variable-length columns are chosen for external off-page storage until the row fits within half a page, as described in [Section 5.3.2, “File Space Management”](#).

- Although InnoDB supports row sizes larger than 65,535 bytes internally, MySQL itself imposes a row-size limit of 65,535 for the combined size of all columns:

```
mysql> CREATE TABLE t (a VARCHAR(8000), b VARCHAR(10000),
-> c VARCHAR(10000), d VARCHAR(10000), e VARCHAR(10000),
-> f VARCHAR(10000), g VARCHAR(10000)) ENGINE=InnoDB;
ERROR 1118 (42000): Row size too large. The maximum row size for the
used table type, not counting BLOBs, is 65535. You have to change some
columns to TEXT or BLOBs
```

See [Section E.10.4, “Limits on Table Column Count and Row Size”](#).

- On some older operating systems, files must be less than 2GB. This is not a limitation of InnoDB itself, but if you require a large tablespace, you will need to configure it using several smaller data files rather than one or a file large data files.
- The combined size of the InnoDB log files can be up to 512GB.
- The minimum tablespace size is slightly larger than 10MB. The maximum tablespace size is four billion database pages (64TB). This is also the maximum size for a table.
- The default database page size in InnoDB is 16KB, or you can lower the page size to 8KB or 4KB by specifying the `innodb_page_size` [1740] option when creating the MySQL instance.



#### Note

Increasing the page size is not a supported operation: there is no guarantee that InnoDB will function normally with a page size greater than 16KB. Problems compiling or running InnoDB may occur. In particular, `ROW_FORMAT=COMPRESSED` in the Barracuda file format assumes that the page size is at most 16KB and uses 14-bit pointers.

A MySQL instance using a particular InnoDB page size cannot use data files or log files from an instance that uses a different page size. This limitation

could affect restore or downgrade operations using data from MySQL 5.6, which does support page sizes other than 16KB.

## Index Types

- InnoDB tables support `FULLTEXT` indexes, starting in MySQL 5.6.4. See [Section 14.2.3.12.3, “FULLTEXT Indexes”](#) for details.
- InnoDB tables support spatial data types, but not indexes on them.

## Restrictions on InnoDB Tables

- `ANALYZE TABLE` determines index cardinality (as displayed in the `Cardinality` column of `SHOW INDEX` output) by doing [random dives](#) to each of the index trees and updating index cardinality estimates accordingly. Because these are only estimates, repeated runs of `ANALYZE TABLE` could produce different numbers. This makes `ANALYZE TABLE` fast on InnoDB tables but not 100% accurate because it does not take all rows into account.

You can make the [statistics](#) collected by `ANALYZE TABLE` more precise and more stable by turning on the `innodb_stats_persistent` [\[1747\]](#) configuration option, as explained in [Section 14.2.4.2.10, “Persistent Optimizer Statistics for InnoDB Tables”](#). When that setting is enabled, it is important to run `ANALYZE TABLE` after major changes to indexed column data, because the statistics are not recalculated periodically (such as after a server restart) as they traditionally have been.

You can change the number of random dives by modifying the `innodb_stats_persistent_sample_pages` [\[1748\]](#) system variable (if the persistent statistics setting is turned on), or the `innodb_stats_transient_sample_pages` [\[1748\]](#) system variable (if the persistent statistics setting is turned off). For more information, see [Section 14.2.5, “InnoDB Features for Flexibility, Ease of Use and Reliability”](#).

MySQL uses index cardinality estimates only in join optimization. If some join is not optimized in the right way, you can try using `ANALYZE TABLE`. In the few cases that `ANALYZE TABLE` does not produce values good enough for your particular tables, you can use `FORCE INDEX` with your queries to force the use of a particular index, or set the `max_seeks_for_key` [\[507\]](#) system variable to ensure that MySQL prefers index lookups over table scans. See [Section 5.1.4, “Server System Variables”](#), and [Section C.5.6, “Optimizer-Related Issues”](#).

- `SHOW TABLE STATUS` does not give accurate statistics on InnoDB tables, except for the physical size reserved by the table. The row count is only a rough estimate used in SQL optimization.
- InnoDB does not keep an internal count of rows in a table because concurrent transactions might “see” different numbers of rows at the same time. To process a `SELECT COUNT(*) FROM t` statement, InnoDB scans an index of the table, which takes some time if the index is not entirely in the buffer pool. If your table does not change often, using the MySQL query cache is a good solution. To get a fast count, you have to use a counter table you create yourself and let your application update it according to the inserts and deletes it does. If an approximate row count is sufficient, `SHOW TABLE STATUS` can be used. See [Section 14.2.4.1, “InnoDB Performance Tuning Tips”](#).
- On Windows, InnoDB always stores database and table names internally in lowercase. To move databases in a binary format from Unix to Windows or from Windows to Unix, create all databases and tables using lowercase names.
- An `AUTO_INCREMENT` column `ai_col` must be defined as part of an index such that it is possible to perform the equivalent of an indexed `SELECT MAX(ai_col)` lookup on the table to obtain the maximum column value. Typically, this is achieved by making the column the first column of some table index.
- While initializing a previously specified `AUTO_INCREMENT` column on a table, InnoDB sets an exclusive lock on the end of the index associated with the `AUTO_INCREMENT` column. While

accessing the auto-increment counter, `InnoDB` uses a specific `AUTO-INC` table lock mode where the lock lasts only to the end of the current SQL statement, not to the end of the entire transaction. Other clients cannot insert into the table while the `AUTO-INC` table lock is held. See [Section 5.4.4, “AUTO\\_INCREMENT Handling in InnoDB”](#).

- When you restart the MySQL server, `InnoDB` may reuse an old value that was generated for an `AUTO_INCREMENT` column but never stored (that is, a value that was generated during an old transaction that was rolled back).
- When an `AUTO_INCREMENT` integer column runs out of values, a subsequent `INSERT` operation returns a duplicate-key error. This is general MySQL behavior, similar to how `MyISAM` works.
- `DELETE FROM tbl_name` does not regenerate the table but instead deletes all rows, one by one.
- Currently, cascaded foreign key actions do not activate triggers.
- You cannot create a table with a column name that matches the name of an internal `InnoDB` column (including `DB_ROW_ID`, `DB_TRX_ID`, `DB_ROLL_PTR`, and `DB_MIX_ID`). The server reports error 1005 and refers to error -1 in the error message. This restriction applies only to use of the names in uppercase.

## Locking and Transactions

- `LOCK TABLES` acquires two locks on each table if `innodb_table_locks=1` (the default). In addition to a table lock on the MySQL layer, it also acquires an `InnoDB` table lock. Versions of MySQL before 4.1.2 did not acquire `InnoDB` table locks; the old behavior can be selected by setting `innodb_table_locks=0`. If no `InnoDB` table lock is acquired, `LOCK TABLES` completes even if some records of the tables are being locked by other transactions.

In MySQL 5.6, `innodb_table_locks=0` [\[1751\]](#) has no effect for tables locked explicitly with `LOCK TABLES ... WRITE`. It does have an effect for tables locked for read or write by `LOCK TABLES ... WRITE` implicitly (for example, through triggers) or by `LOCK TABLES ... READ`.

- All `InnoDB` locks held by a transaction are released when the transaction is committed or aborted. Thus, it does not make much sense to invoke `LOCK TABLES` on `InnoDB` tables in `autocommit=1` [\[464\]](#) mode because the acquired `InnoDB` table locks would be released immediately.
- You cannot lock additional tables in the middle of a transaction because `LOCK TABLES` performs an implicit `COMMIT` and `UNLOCK TABLES`.
- The limit of 1023 concurrent data-modifying transactions has been raised in MySQL 5.5 and above. The limit is now  $128 * 1023$  concurrent transactions that generate undo records. You can remove any workarounds that require changing the proper structure of your transactions, such as committing more frequently.

### 14.2.8. MySQL and the ACID Model

The `ACID` model is a set of database design principles that emphasize aspects of reliability that are important for business data and mission-critical applications. MySQL includes components such as the `InnoDB` storage engine that adhere closely to the ACID model, so that data is not corrupted and results are not distorted by exceptional conditions such as software crashes and hardware malfunctions. When you rely on ACID-compliant features, you do not need to reinvent the wheel of consistency checking and crash recovery mechanisms. In cases where you have additional software safeguards, ultra-reliable hardware, or an application that can tolerate a small amount of data loss or inconsistency, you can adjust MySQL settings to trade some of the ACID reliability for greater performance or throughput.

The following sections discuss how MySQL features, in particular the `InnoDB` storage engine, interact with the categories of the ACID model:

- **A**: atomicity.
- **C**: consistency.
- **I**: isolation.
- **D**: durability.

## Atomicity

The **atomicity** aspect of the ACID model mainly involves [InnoDB transactions](#). Related MySQL features include:

- Autocommit setting.
- `COMMIT` statement.
- `ROLLBACK` statement.
- Operational data from the `INFORMATION_SCHEMA` tables.

## Consistency

The **consistency** aspect of the ACID model mainly involves internal [InnoDB](#) processing to protect data from crashes. Related MySQL features include:

- [InnoDB doublewrite buffer](#).
- [InnoDB crash recovery](#).

## Isolation

The **isolation** aspect of the ACID model mainly involves [InnoDB transactions](#), in particular the [isolation level](#) that applies to each transaction. Related MySQL features include:

- `Autocommit` setting.
- `SET ISOLATION LEVEL` statement.
- The low-level details of [InnoDB locking](#). During performance tuning, you see these details through `INFORMATION_SCHEMA` tables.

## Durability

The **durability** aspect of the ACID model involves MySQL software features interacting with your particular hardware configuration. Because of the many possibilities depending on the capabilities of your CPU, network, and storage devices, this aspect is the most complicated to provide concrete guidelines for. (And those guidelines might take the form of buy “new hardware”.) Related MySQL features include:

- [InnoDB doublewrite buffer](#), turned on and off by the `innodb_doublewrite` [1715] configuration option.
- Configuration option `innodb_flush_log_at_trx_commit` [1718].
- Configuration option `sync_binlog` [2017].
- Configuration option `innodb_file_per_table` [1718].
- Write buffer in a storage device, such as a disk drive, SSD, or RAID array.
- Battery-backed cache in a storage device.
- The operating system used to run MySQL, in particular its support for the `fsync()` system call.



- Uninterruptible power supply (UPS) protecting the electrical power to all computer servers and storage devices that run MySQL servers and store MySQL data.
- Your backup strategy, such as frequency and types of backups, and backup retention periods.
- For distributed or hosted data applications, the particular characteristics of the data centers where the hardware for the MySQL servers is located, and network connections between the data centers.

## 14.2.9. InnoDB Integration with memcached

The ever-increasing performance demands of web-based services have generated significant interest in simple data access methods that maximize performance. These techniques are broadly classified under the name “NoSQL”: to increase performance and throughput, they take away the overhead of parsing an SQL statement, constructing an execution plan, and dealing with strongly typed data values split into multiple fields.

MySQL 5.6 includes a NoSQL interface, using an integrated `memcached` daemon that can automatically store data and retrieve it from `InnoDB` tables, turning the MySQL server into a fast “key-value store” for single-row insert, update, or delete operations. You can still also access the same tables through SQL for convenience, complex queries, bulk operations, application compatibility, and other strengths of traditional database software.

With this NoSQL interface, you use the familiar `memcached` API to speed up database operations, letting `InnoDB` handle memory caching using its `buffer pool` mechanism. Data modified through `memcached` operations such as `ADD`, `SET`, `INCR` are stored to disk, using the familiar `InnoDB` mechanisms such as `change buffering`, the `doublewrite buffer`, and `crash recovery`. The combination of `memcached` simplicity and `InnoDB` durability provides users with the best of both worlds, as explained in [Section 14.2.9.1, “Benefits of the InnoDB / memcached Combination”](#). For architectural details about how the components fit together, see [Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”](#).

### 14.2.9.1. Benefits of the InnoDB / memcached Combination

This section outlines advantages and usage scenarios for the `memcached` interface to `InnoDB` tables introduced in [Section 14.2.9, “InnoDB Integration with memcached”](#). The combination of `InnoDB` tables and `memcached` offers advantages over using either by themselves:

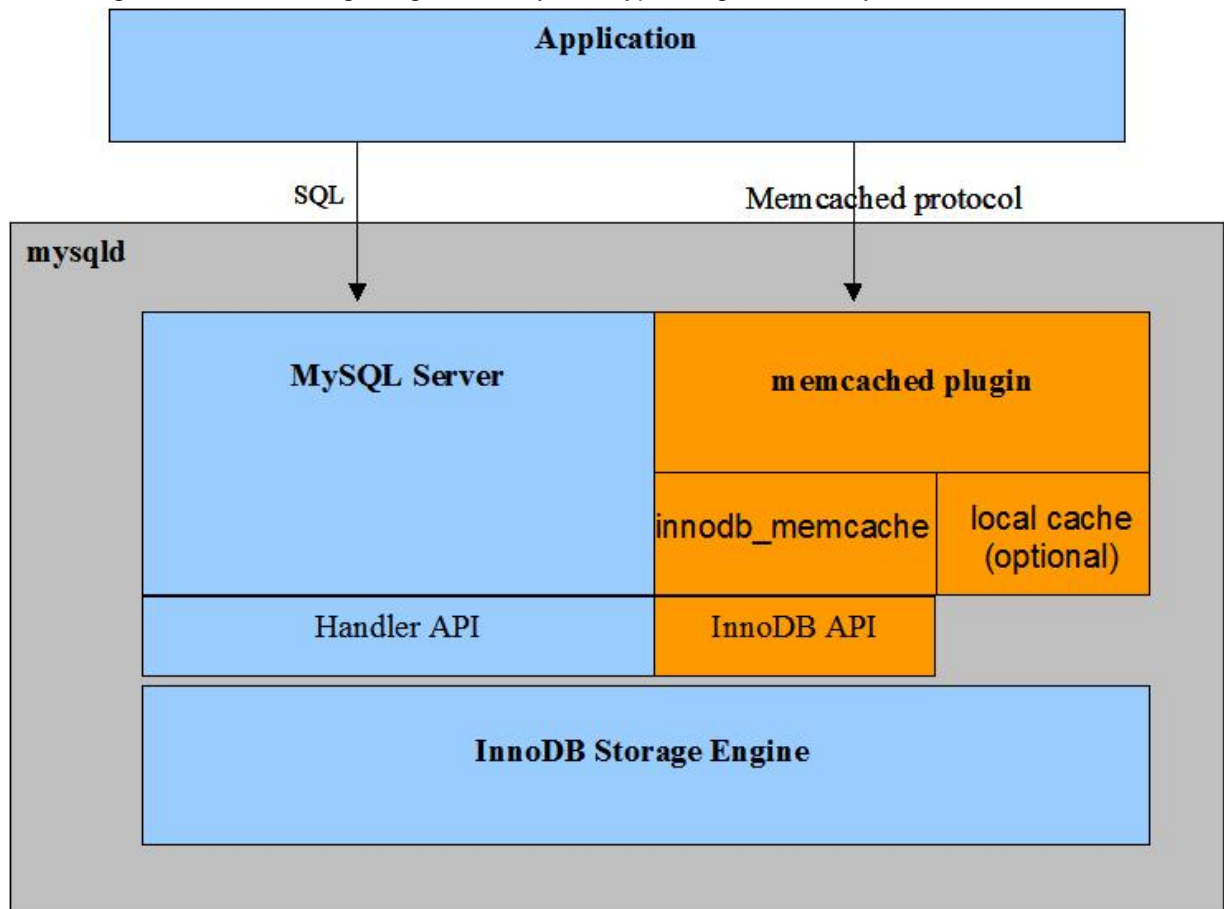
- Raw performance for simple lookups. Direct access to the `InnoDB` storage engine avoids the parsing and planning overhead of SQL. Running `memcached` in the same process space as the MySQL server avoids the network overhead of passing requests back and forth.
- Data is stored in a MySQL database to protect against crashes, outages, and corruption.
- The transfer between memory and disk is handled automatically, simplifying application logic.
- Data can be unstructured or structured, depending on the type of application. You can make an all-new table for the data, or map the NoSQL-style processing to one or more existing tables.
- You can still access the underlying table through SQL, for reporting, analysis, ad hoc queries, bulk loading, set operations such as union and intersection, and other operations well suited to the expressiveness and flexibility of SQL.
- You can ensure high availability of the NoSQL data by using this feature on a `master server` in combination with MySQL replication.
- The integration of `memcached` with MySQL provides a painless way to make the in-memory data persistent, so you can use it for more significant kinds of data. You can put more `add`, `incr`, and similar write operations into your application, without worrying that the data could disappear at any moment. You can bounce the `memcached` server without losing updates made to the cached data. To guard against unexpected outages, you can take advantage of `InnoDB` crash recovery, replication, and backup procedures.

- The way [InnoDB](#) does fast [primary key](#) lookups is a natural fit for [memcached](#) single-item queries. The direct, low-level database access path used by the [memcached](#) plugin is much more efficient for key-value lookups than equivalent SQL queries.
- The serialization features of [memcached](#), which can turn complex data structures, binary files, or even code blocks into storable strings, offer a simple way to get such objects into a database.
- Because you can access the underlying data through SQL, you can produce reports, search or update across multiple keys, and call functions such as `AVG()` and `MAX()` on the [memcached](#) data. All of these operations are expensive or complicated with the standalone [memcached](#).
- You do not need to manually load data into [memcached](#) at startup. As particular keys are requested by an application, the values are retrieved from the database automatically, and cached in memory using the [InnoDB buffer pool](#).
- Because [memcached](#) consumes relatively little CPU, and its memory footprint is easy to control, it can run comfortably alongside a MySQL instance on the same system.
- Because data consistency is enforced through the usual mechanism as with regular [InnoDB](#) tables, you do not have to worry about stale [memcached](#) data or fallback logic to query the database in the case of a missing key.

### 14.2.9.2. Architecture of InnoDB and memcached Integration

This section describes how the [memcached](#) daemon is integrated into the MySQL Server, to help understand how this approach compares with other techniques that combine NoSQL components or interfaces with a MySQL back-end.

When integrated with MySQL Server, [memcached](#) is implemented as a MySQL plugin daemon, accessing the [InnoDB](#) storage engine directly and bypassing the SQL layer:



Features provided in the current release:

- `memcached` as a daemon plugin of `mysqld`: both `mysqld` and `memcached` run in the same process space, with very low latency access to data.
- Direct access to `InnoDB` tables, bypassing the SQL parser, the optimizer, and even the Handler API layer.
- Standard `memcached` protocols, both the text-based protocol and the binary protocol. The `InnoDB + memcached` combination passes all 55 compatibility tests from the `memcapable` command.
- Multi-column support: you can map multiple columns into the “value” part of the key/value store, with column values delimited by a user-specified separator character.
- By default, you use the `memcached` protocol to read and write data directly to `InnoDB`, and let MySQL manage the in-memory caching through the `InnoDB buffer pool`. The default settings represent the combination of high reliability with the fewest surprises for database applications. For example, the default settings avoid uncommitted data on the database side, or stale data returned for `memcached get` requests.
- You can control how often data is passed back and forth between `InnoDB` and `memcached` operations through the `innodb_api_bk_commit_interval` [1702], `daemon_memcached_r_batch_size` [1699], and `daemon_memcached_w_batch_size` [1700] configuration options. Both the batch size options default to a value of 1 for maximum reliability.
- You can specify any `memcached` configuration options through the MySQL configuration variable `daemon_memcached_option` [1699]. For example, you might change the port that `memcached` listens on, reduce the maximum number of simultaneous connections, change the maximum memory size for a key/value pair, or enable debugging messages for the error log.
- A configuration option `innodb_api_trx_level` [1704] lets you control the transaction `isolation level` on queries processed by the `memcached` interface. Although `memcached` has concept of `transactions`, you might use this property to control how soon `memcached` sees changes caused by SQL statements, if you issue `DML` statements on the same table that `memcached` interfaces with. By default, it is set to `READ UNCOMMITTED` [1470].
- Another configuration option is `innodb_api_enable_md1` [1703]. “MDL” stands for “metadata locking”. This basically locks the table from the MySQL level, so that the mapped table cannot be dropped or altered by `DDL` through the SQL interface. Without the lock, the table can be dropped from MySQL layer, but will be kept in the `InnoDB` storage until `memcached` or any other user stops using it.

### Differences Between Using `memcached` Standalone or with `InnoDB`

MySQL users might already be familiar with using `memcached` along with MySQL, as described in [Section 15.6, “Using MySQL with `memcached`”](#). This section describes the similarities and differences between the information in that section, and when using the `InnoDB` integration features of the `memcached` that is built into MySQL. The link at the start of each item goes to the associated information about the traditional `memcached` server.

- **Installation:** Because the `memcached` library comes with the MySQL server, installation and setup are straightforward. You run a SQL script to set up a table for `memcached` to use, issue a one-time `install plugin` statement to enable `memcached`, and add to the MySQL configuration file or startup script any desired `memcached` options, for example to use a different port. You might still install the regular `memcached` distribution to get the additional utilities such as `memcp`, `memcat`, and `memcapable`.
- **Deployment:** It is typical to run large numbers of low-capacity `memcached` servers. Because the `InnoDB + memcached` combination has a 1:1 ratio between database and `memcached` servers, the typical deployment involves a smaller number of moderate or high-powered servers, machines that were already running MySQL. The benefit of this server configuration is more for improving the efficiency of each individual database server than in tapping into unused memory or distributing

lookups across large numbers of servers. In the default configuration, very little memory is used for `memcached`, and the in-memory lookups are served from the `InnoDB buffer pool`, which automatically caches the most recently used and most frequently used data. As in a traditional MySQL server instance, keep the value of the `innodb_buffer_pool_size` [1708] configuration option as high as practical (without causing paging at the OS level), so that as much of the workload as possible is done in memory.

- **Expiry:** By default (that is, with the caching policy `innodb_only`), the latest data from the `InnoDB` table is always returned, so the expiry options have no practical effect. If you change the caching policy to `caching` or `cache-only`, the expiry options work as usual, but requested data might be stale if it was updated in the underlying table before it expires from the memory cache.
- **Namespaces:** `memcached` is like a single giant directory, where to keep files from conflicting with each other you might give them elaborate names with prefixes and suffixes. The integrated `InnoDB / memcached` server lets you use these same naming conventions for keys, with one addition. Key names of the format `@@table_id.key.table_id` are decoded to reference a specific a table, using mapping data from the `innodb_memcache.containers` table. The `key` is looked up in or written to the specified table.

The `@@` notation only works for individual calls to the `get`, `add`, and `set` functions, not the others such as `incr` or `delete`. To designate the default table for all subsequent `memcached` operations within a session, perform a `get` request using the `@@` notation and a table ID, but without the key portion. For example:

```
get @@table_x
```

Subsequent `get`, `set`, `incr`, `delete` and other operations use the table designated by `table_x` in the `innodb_memcache.containers.name` column.

- **Hashing and distribution:** The default configuration, with the caching policy `innodb_only`, is suitable for the traditional deployment configuration where all data is available on all servers, such as a set of replication slave servers.

If you physically divide the data, as in a sharded configuration, you can split the data across several machines running the `InnoDB` and `memcached` combined server, and use the traditional `memcached` hashing mechanism to route requests to a particular machine. On the MySQL side, typically you would let all the data be inserted by `add` requests to `memcached` so the appropriate values were stored in the database on the appropriate server.

These types of deployment best practices are still being codified.

- **Memory usage:** By default (with the caching policy `innodb_only`), the `memcached` protocol passes information back and forth with `InnoDB` tables, and the fixed-size `InnoDB` buffer pool handles the in-memory lookups rather than `memcached` memory usage growing and shrinking. Relatively little memory is used on the `memcached` side.

If you switch the caching policy to `caching` or `cache-only`, the normal rules of `memcached` memory usage apply. Memory for the `memcached` data values is allocated in terms of “slabs”. You can control the slab size and maximum memory used for `memcached`.

Either way, you can monitor and troubleshoot the integrated `memcached` daemon using the familiar `statistics` system, accessed through the standard protocol, for example over a `telnet` session. Because extra utilities are not included with the integrated daemon, to use the `memcached-tool script`, install a full `memcached` distribution.

- **Thread usage:** MySQL threads and `memcached` threads must co-exist on the same server, so any limits imposed on threads by the operating system apply to this total number.
- **Log usage:** Because the `memcached` daemon is run alongside the MySQL server and writes to `stderr`, the `-v`, `-vv`, and `-vvv` options for logging write their output to the MySQL `error log`.

- **memcached operations:** All the familiar operations such as `get`, `set`, `add`, and `delete` are available. Serialization (that is, the exact string format to represent complex data structures) depends on the language interface. Still to confirm if there are any special nuances for these operations:

```
delete(key, [, time]): Deletes the key and its associated item from the cache.
    If you supply a time, then adding another item with the specified key is blocked
    for the specified period.
flush_all: Invalidates (or expires) all the current items in the cache.
    Technically they still exist (they are not deleted), but they are silently destroyed
    the next time you try to access them.
```

- **Using memcached as a MySQL front end:** That is what the [InnoDB](#) integration with [memcached](#) is all about. Putting these components together improves the performance of your application. Making [InnoDB](#) handle data transfers between memory and disk simplifies the logic of your application.
- **Utilities:** The MySQL server includes the `libmemcached` library but not the additional command-line utilities. To get the commands such as `memcp`, `memcat`, and `memcapable` commands, install a full [memcached](#) distribution. When `memrm` and `memflush` remove items from the cache, they are also removed from the underlying [InnoDB](#) table.
- **Programming interfaces:** You can access the MySQL server through the [InnoDB](#) and [memcached](#) combination using the same language as always: [C](#) and [C++](#), [Java](#), [Perl](#), [Python](#), [PHP](#), and [Ruby](#). Specify the server hostname and port as with any other [memcached](#) server. By default, the integrated [memcached](#) server listens on the same port as usual, `11211`. You can use both the [text and binary protocols](#). You can customize the [behavior](#) of the [memcached](#) functions at runtime. Serialization (that is, the exact string format to represent complex data structures) depends on the language interface.
- **Frequently asked questions:** MySQL has had an extensive [memcached](#) FAQ for several releases. In MySQL 5.6, the answers are largely the same, except that using [InnoDB](#) tables as a storage medium for [memcached](#) data means that you can use this combination for more write-intensive applications than before, rather than as a read-only cache.

For a more detailed look at the workings of this feature, see [Section 14.2.9.7, “Internals of the InnoDB memcached Plugin”](#).

### 14.2.9.3. Getting Started with InnoDB Memcached Plugin

This section describes the steps to activate the [InnoDB / memcached](#) integration on a MySQL Server. Because the [memcached](#) daemon is tightly integrated with the MySQL Server to avoid network traffic and minimize latency, you perform this process on each MySQL instance that uses this feature.



#### Note

Before setting up the [memcached](#) interface for any data, consult [Section 14.2.9.4, “Security Considerations for the InnoDB memcached Plugin”](#) to understand the security procedures needed to prevent any unauthorized access.

#### 14.2.9.3.1. Prerequisites for the InnoDB memcached Plugin

Before you set up the plugin itself and the internal tables it uses, verify that your server has the required prerequisite software.

##### Platform Support

Currently, the [memcached](#) Daemon Plugin is only supported on Linux, Solaris, and OS X platforms.

##### Software Prerequisites

You must have [libevent](#) installed, since it is required by [memcached](#). The way to get this library is different if you use the MySQL installer or build from source, as described in the following sections.

## Using a MySQL Installation Package

When you use a MySQL installer, the `libevent` library is not included. Use the particular method for your operating system to download and install `libevent` 1.4.3 or later: for example, depending on the operating system, you might use the command `apt-get`, `yum`, or `port install`.

The libraries for `memcached` and the InnoDB plugin for `memcached` are put into the right place by the MySQL installer. For typical operation, the files `lib/plugin/libmemcached.so` and `lib/plugin/innodb_engine.so` are used.

## Building from Source

For a brief introduction on the setup steps, see the file `README-innodb_memcached` in the source distribution. This is a more detailed explanation of that procedure.

If you have the source code release, then there is a `libevent` 1.4.3 included in the package.

1. Go to the directory `plugin/innodb_memcached/libevent`.
2. Issue the following commands to build and install the `libevent` library:

```
autoconf
./configure
make
make install
```

When building the MySQL Server from source, once `libevent` is installed, build the MySQL server as usual. No special build steps are required.

The source code for the InnoDB-`memcached` plugin is in the `plugin/innodb_memcached` directory. As part of the server build, it generates two shared libraries:

- `libmemcached.so`: the `memcached` daemon plugin to MySQL.
- `innodb_engine.so`: an InnoDB API plugin to `memcached`.

Put these two shared libraries in the MySQL plugin directory. To find the MySQL plugin directory, issue the following command:

```
mysql> select @@plugin_dir;
+-----+
| @@plugin_dir |
+-----+
| /Users/cirrus/sandbox-setup/5.6.10/lib/plugin/ |
+-----+
1 row in set (0.00 sec)
```

## Setting Operating System Limits

The `memcached` daemon can sometimes cause the MySQL server to exceed the OS limit on the number of open files. You might need to run the `ulimit` command to increase the limit, and then start the MySQL server from that same shell. See [Section 14.2.9.8, “Troubleshooting the InnoDB memcached Plugin”](#) for the steps to resolve this issue.

### 14.2.9.3.2. Installing and Configuring the InnoDB memcached Plugin

#### Setting Up Required Tables

To configure the `memcached` plugin so it can interact with InnoDB tables, run the configuration script `scripts/innodb_memcached_config.sql` to install the necessary tables used behind the scenes:

```
mysql: source MYSQL_HOME/share/innodb_memcached_config.sql
```



This is a one-time operation. The tables remain in place if you later disable and re-enable the `memcached` support. For information about the layout and purpose of these tables, see [Section 14.2.9.7, “Internals of the InnoDB memcached Plugin”](#).

### Installing the Daemon Plugin

To activate the daemon plugin, use the `install plugin` statement, just as when installing any other MySQL plugin:

```
mysql> install plugin daemon_memcached soname "libmemcached.so";
```

Once the plugin is installed this way, it is automatically activated each time the MySQL server is booted or restarted.

### Disabling the Daemon Plugin

When making major changes to the plugin configuration, you might need to turn off the plugin. To do so, issue the following statement:

```
mysql> uninstall plugin daemon_memcached;
```

To re-enable it, issue the preceding `install plugin` statement again. All the previous configuration settings, internal tables, and data are preserved when the plugin is bounced this way.

### Specifying `memcached` Configuration Options

If you have any `memcached` specific configuration parameters, specify them on the `mysqld` command line or enter them in the MySQL configuration file, encoded in the argument to the `daemon_memcached_option` [1699] MySQL configuration option. The `memcached` configuration options take effect when the plugin is installed, which you do each time the MySQL server is started.

For example, to make `memcached` listen on port 11222 instead of the default port 11211, add `-p11222` to the MySQL configuration option `daemon_memcached_option` [1699]:

```
mysqld ... --loose-daemon_memcached_option="-p11222"
```

You can add other `memcached` command line options to the `daemon_memcached_option` [1699] string. The other configuration options are:

- `daemon_memcached_engine_lib_name` [1698] (default `innodb_engine.so`)
- `daemon_memcached_engine_lib_path` [1699] (default `NULL`, representing the plugin directory).
- `daemon_memcached_r_batch_size` [1699], batch commit size for read operations (`get`). It specifies after how many `memcached` read operations the system automatically does a `commit`. By default, this is set to 1 so that every `get` request can access the very latest committed data in the `InnoDB` table, whether the data was updated through `memcached` or by SQL. When its value is greater than 1, the counter for read operations is incremented once for every `get` call. The `flush_all` call resets both the read and write counters.
- `daemon_memcached_w_batch_size` [1700], batch commit for any write operations (`set`, `replace`, `append`, `prepend`, `incr`, `decr`, and so on) By default, this is set as 1, so that no uncommitted data is lost in case of an outage, and any SQL queries on the underlying table can access the very latest data. When its value is greater than 1, the counter for write operations is incremented once for every `add`, `set`, `incr`, `decr`, and `delete` call. The `flush_all` call resets both the read and write counters.

By default, you do not need to change anything with the first two configuration options. Those options allow you to load any other storage engine for `memcached` (such as the `NDB memcached` engine).



Again, please note that you will have these configuration parameter in your MySQL configure file or MySQL boot command line. They take effect when you load the `memcached` plugin.

## Summary

Now you have everything set up. You can directly interact with InnoDB tables through the `memcached` interface. To verify that the feature is working properly, see [Section 14.2.9.3.3, “Verifying the InnoDB and memcached Setup”](#).

### 14.2.9.3.3. Verifying the InnoDB and memcached Setup

Now that everything is set up, you can experiment with the InnoDB and `memcached` combination:

Here is an example using the Unix, Linux, or OS X command shell:

```
# Point memcached-related commands at the memcached attached to the mysqld process.
export MEMCACHED_SERVERS=127.0.0.1:11211
# Store the contents of a modestly sized text file in memcached, with the data passed
# to MySQL and stored in a table. The key is the basename of the file, 'mime.types'.
memcp /etc/apache2/mime.types
# Retrieve the data we just stored, from the memory cache.
memcat mime.types
```

Here is an example using `telnet` to send `memcached` commands and receive results through the ASCII protocol:

```
telnet 127.0.0.1 11211
set all 10 0 9
123456789
STORED
get all
VALUE all 0 9
123456789
END
quit
```

To prove that all the same data has been stored in MySQL, connect to the MySQL server and issue:

```
mysql> select * from test.demo_test;
```

Now, shut down the MySQL server, which also shuts off the integrated `memcached` server. Further attempts to access the `memcached` data now fail with a connection error. Normally, the `memcached` data would disappear at this point, and you would write application logic to load the data back into memory when `memcached` was restarted. But the MySQL / `memcached` integration automates this process:

- Restart the MySQL server.
- Run the `install plugin` statement to start the `daemon_memcached` plugin again.
- Now any `memcat` commands or `get` operations once again return the key/value pairs you stored in the earlier `memcached` session. When a key is requested and the associated value is not already in the memory cache, it is automatically queried from the MySQL table, by default `test.demo_test`.

### 14.2.9.4. Security Considerations for the InnoDB memcached Plugin



#### Caution

Consult this section before deploying the `InnoDB memcached` plugin on any production servers, or even test servers if the MySQL instance contains any sensitive information.

Because `memcached` does not use an authentication mechanism by default, and the optional SASL authentication is not as strong as traditional DBMS security measures, make sure to keep only non-

sensitive data in the MySQL instance using the [InnoDB memcached](#) plugin, and wall off any servers using this configuration from potential intruders. Do not allow [memcached](#) access to such servers from the Internet, only from within a firewalled intranet, ideally from a subnet whose membership you can restrict.

#### 14.2.9.4.1. Password-Protecting the [memcached](#) Interface through SASL

SASL support gives you the capability to protect your MySQL database from unauthenticated access through [memcached](#) clients. This section explains the steps to enable this option. The steps to enable such support are almost identical to those you would do to enable SASL for a traditional [memcached](#) server.

##### Background Info:

SASL stands for “Simple Authentication and Security Layer”, a standard for adding authentication support to connection-based protocols. [memcached](#) added SASL support starting in its 1.4.3 release.

For the InnoDB + [memcached](#) combination, the table that stores the [memcached](#) data must be registered in the [container](#) system table. And [memcached](#) clients can only access such a registered table. Even though the DBA can add access restrictions on a table that is registered with the [memcached](#) plugin, they have no control over who can access it through [memcached](#) applications. This is why we provide a means (through SASL) to control who can access [InnoDB](#) tables associated with the [memcached](#) plugin.

The following section shows how to build, enable, and test an SASL-enabled [InnoDB memcached](#) plugin.

##### Steps to Build and Enable SASL in InnoDB Memcached Plugin:

By default, SASL-enabled [InnoDB memcached](#) is not included in the release package, since it relies on building [memcached](#) with SASL libraries. To enable this feature, download the MySQL source and rebuild the [InnoDB memcached](#) plugin after downloading the SASL libraries:

1. First, get the SASL development and utility libraries. For example, on Ubuntu, you can get these libraries through:

```
sudo apt-get -f install libsasl2-2 sasl2-bin libsasl2-2 libsasl2-dev libsasl2-modules
```

2. Then build the [InnoDB memcached](#) plugin (shared libraries) with SASL capability, by adding `ENABLE_MEMCACHED_SASL=1` to the `cmake` options. In addition, [memcached](#) provides a simple plaintext password support, which is easier to use for testing. To enable this, set the option `ENABLE_MEMCACHED_SASL_PWDB=1`.

Overall, you will add following three options to the `cmake`:

```
cmake ... -DWITH_INNO_DB_MEMCACHED=1
-DENABLE_MEMCACHED_SASL=1 -DENABLE_MEMCACHED_SASL_PWDB=1
```

3. The third step is to install the [InnoDB memcached](#) plugin as before, as explained in [Section 14.2.9.3, “Getting Started with InnoDB Memcached Plugin”](#).
4. As previously mentioned, [memcached](#) provides a simple plaintext password support through SASL, which will be used for this demo.

- a. Create a user named `testname` and its password as `testpasswd` in a file:

```
echo "testname:testpasswd:::::::::" >/home/jy/memcached-sasl-db
```

- b. Let [memcached](#) know about it by setting the environment variable `MEMCACHED_SASL_PWDB`:

```
export MEMCACHED_SASL_PWDB=/home/jy/memcached-sasl-db
```

- c. Also tell `memcached` that it is a plaintext password:

```
echo "mech_list: plain" > /home/jy/work2/msasl/clients/memcached.conf
export SASL_CONF_PATH=/home/jy/work2/msasl/clients/memcached.conf
```

5. Then reboot the server, and add a `daemon_memcached_option` [1699] option `-S` to enable SASL:

```
mysqld ... --daemon_memcached_option="-S"
```

6. Now the setup is complete. To test it, you might need an SASL-enabled client, such as this [SASL-enabled libmemcached](#).

```
memcp --servers=localhost:11211 --binary --username=testname
--password=testpasswd myfile.txt

memcat --servers=localhost:11211 --binary --username=testname
--password=testpasswd myfile.txt
```

Without appropriate user name or password, the above operation is rejected with the error message `memcache error AUTHENTICATION FAILURE`. Otherwise, the operation succeed. You can also examine the plaintext password set in the `memcached-sasl-db` file to verify it.

There are other methods to test SASL authentication with `memcached`. But the one described above is the most straightforward.

#### 14.2.9.5. Writing Applications for the InnoDB memcached Interface

Typically, writing an application for the [InnoDB memcached](#) interface involves some degree of rewriting or adapting existing code that uses MySQL or the `memcached` API:

- Instead of many `memcached` servers running on low-powered machines, you have the same number of `memcached` servers as MySQL servers, running on relatively high-powered machines with substantial disk storage and memory. You might reuse some existing code that works with the `memcached` API, but some adaptation is likely needed due to the different server configuration.
- The data stored through this interface all goes into `VARCHAR`, `TEXT`, or `BLOB` columns, and must be converted to do numeric operations. You can do the conversion on the application side, or by using the `CAST()` function in queries.
- Coming from a database background, you might be used to general-purpose SQL tables with many columns. The tables accessed by the `memcached` code likely have only a few or even just a single column holding data values.
- You might adapt parts of your application that do single-row queries, inserts, updates, or deletes, to squeeze more performance out of critical sections of code. Both [queries](#) (read) and [DML](#) (write) operations can be substantially faster when performed through the `memcached` interface. The speedup for writes is typically greater than the speedup for reads, so you might focus on adapting the code that performs logging or records interactive choices on a web site.

The following sections explore these aspects in more detail.

##### 14.2.9.5.1. Adapting an Existing MySQL Schema for a `memcached` Application

Consider these aspects of `memcached` applications when adapting an existing MySQL schema or application to use the `memcached` interface:

- `memcached` keys cannot contain spaces or newlines, because those characters are used as separators in the ASCII protocol. If you are using lookup values that contain spaces, transform or hash them into values without spaces before using them as keys in calls to `add()`, `set()`, `get()` and so on. Although theoretically those characters are allowed in keys in programs that use the

binary protocol, you should always restrict the characters used in keys to ensure compatibility with a broad range of clients.

- If you have a short numeric [primary key](#) column in an [InnoDB](#) table, you can use that as the unique lookup key for [memcached](#) by converting the integer to a string value. If the [memcached](#) server is being used for more than one application, or with more than one [InnoDB](#) table, consider modifying the name to make sure it is unique. For example, you might prepend the table name, or the database name and the table name, before the numeric value.
- You cannot use a partitioned table for data queried or stored through the [memcached](#) interface.
- The [memcached](#) protocol passes numeric values around as strings. To store numeric values in the underlying [InnoDB](#) table, for example to implement counters that can be used in SQL functions such as [SUM\(\)](#) or [AVG\(\)](#):
  - Use [VARCHAR](#) columns with enough characters to hold all the digits of the largest expected number (and additional characters if appropriate for the negative sign, decimal point, or both).
  - In any query that performs arithmetic using the column values, use the [CAST\(\)](#) function to convert from string to integer or other numeric type. For example:

```
-- Alphabetic entries are returned as zero.
select cast(c2 as unsigned integer) from demo_test;
-- Since there could be numeric values of 0, can't disqualify them.
-- Test the string values to find the ones that are integers, and average only those.
select avg(cast(c2 as unsigned integer)) from demo_test
  where c2 between '0' and '9999999999';
-- Views let you hide the complexity of queries. The results are already converted;
-- no need to repeat conversion functions and WHERE clauses each time.
create view numbers as select c1 key, cast(c2 as unsigned integer) val
  from demo_test where c2 between '0' and '9999999999';
select sum(val) from numbers;
```

Note that any alphabetic values in the result set are turned into 0 by the call to [CAST\(\)](#). When using functions such as [AVG\(\)](#) that depend on the number of rows in the result set, include [WHERE](#) clauses to filter out any non-numeric values.

- If the [InnoDB](#) column you use as a key can be longer than 250 bytes, hash it to a value that is less than 250 bytes.
- To use an existing table with the [memcached](#) interface, define an entry for it in the [innodb\\_memcache.containers](#) table. To make that the table the default for all requests relayed through [memcached](#), specify the value [default](#) in the [name](#) column, then restart the MySQL server to make that change take effect. If you are using multiple tables for different classes of [memcached](#) data, set up multiple entries in the [innodb\\_memcache.containers](#) table with [name](#) values of your choosing, then issue a [memcached](#) request of the form [get @@name](#) or [set @@name](#) within the application to switch the table used for subsequent requests through the [memcached](#) API.

For an example of using a table other than the predefined [test.demo\\_test](#) table, see [Example 14.5, “Specifying the Table and Column Mapping for an InnoDB + memcached Application”](#). For the required layout and meaning of the columns in such a table, see [Section 14.2.9.7, “Internals of the InnoDB memcached Plugin”](#).

- To use multiple MySQL column values with [memcached](#) key/value pairs, in the [innodb\\_memcache.containers](#) entry associated with the MySQL table, specify in the [value\\_columns](#) field several column names separated by comma, semicolon, space, or pipe characters; for example, [col1,col2,col3](#) or [col1|col2|col3](#).

Concatenate the column values into a single string using the pipe character as a separator, before passing that string to [memcached add](#) or [set](#) calls. The string is unpacked automatically into the various columns. Each [get](#) call returns a single string with the column values, also delimited by the pipe separator character. you unpack those values using the appropriate syntax depending on your application language.

### Example 14.5. Specifying the Table and Column Mapping for an InnoDB + memcached Application

Here is an example showing how to use your own table for a MySQL application going through the [InnoDB memcached](#) plugin for data manipulation.

First, we set up a table to hold some country data: the population, area in metric units, and 'R' or 'L' indicating if people drive on the right or on the left.

```
use test;

CREATE TABLE `multicol` (
  `country` varchar(128) NOT NULL DEFAULT '',
  `population` varchar(10) DEFAULT NULL,
  `area_sq_km` varchar(9) DEFAULT NULL,
  `drive_side` varchar(1) DEFAULT NULL,
  `c3` int(11) DEFAULT NULL,
  `c4` bigint(20) unsigned DEFAULT NULL,
  `c5` int(11) DEFAULT NULL,
  PRIMARY KEY (`country`),
) ENGINE=InnoDB DEFAULT CHARSET=latin1;
```

Now we make a descriptor for this table so that the [InnoDB memcached](#) plugin knows how to access it:

- The sample entry in the [CONTAINERS](#) table has a `name` column 'aaa'; we set up another identifier 'bbb'. If we made a single master table for all [memcached](#) applications to use, we would make the ID 'default' and skip the @@ requests to switch tables.
- We specify the `test.multicol` table. The schema name is stored in one column and the table name is stored in another column.
- The key column will be our unique `country` value. That column was specified as the primary key when we created the table above, so we also specify the index name 'PRIMARY' here.
- Rather than a single column to hold a composite data value, we will divide the data among three table columns, so we specify a comma-separated list of those columns that will be used when storing or retrieving values.
- And for the flags, expire, and CAS values, we specify corresponding columns based on the settings from the sample table `demo.test`. These values are typically not significant in applications using the [InnoDB memcached](#) plugin, because MySQL keeps the data synchronized and there is no need to worry about data expiring or being stale.

```
insert into innodb_memcache.containers
  (name,db_schema,db_table,key_columns,value_columns,flags,cas_column,
  expire_time_column,unique_idx_name_on_key)
values
  ('bbb','test','multicol','country','population,area_sq_km,drive_side',
  'c3','c4','c5','PRIMARY');

commit;
```

Here is a sample Python program showing how we would access this table from a program:

- No database authorization is needed, since all data manipulation is done through the [memcached](#) interface. All we need to know is the port number the [memcached](#) daemon is listening to on the local system.
- We load sample values for a few arbitrary countries. (Area and population figures from Wikipedia.)
- To make the program use the `multicol` table, we call the `switch_table()` function that does a dummy `GET` or `SET` request using @@ notation. The name in the request is `bbb`, which is the value we stored in `innodb_memcache.containers.name`. (In a real application, we would use a more

descriptive name. This example just illustrates that you specify a table identifier, not the table name, with the `GET @@...` request.

- The utility functions to insert and query the data demonstrate how we might turn a Python data structure into pipe-separated values for sending to MySQL with `ADD` or `SET` requests, and unpack the pipe-separated values returned by `GET` requests. This extra processing is only required when mapping the single `memcached` value to multiple MySQL table columns.

```
import sys, os
import memcache

def connect_to_memcached():
    memc = memcache.Client(['127.0.0.1:11211'], debug=0);
    print "Connected to memcached."
    return memc

def banner(message):
    print
    print "=" * len(message)
    print message
    print "=" * len(message)

country_data = [
    ("Canada", "34820000", "9984670", "R"),
    ("USA", "314242000", "9826675", "R"),
    ("Ireland", "6399152", "84421", "L"),
    ("UK", "62262000", "243610", "L"),
    ("Mexico", "113910608", "1972550", "R"),
    ("Denmark", "5543453", "43094", "R"),
    ("Norway", "5002942", "385252", "R"),
    ("UAE", "8264070", "83600", "R"),
    ("India", "1210193422", "3287263", "L"),
    ("China", "1347350000", "9640821", "R"),
]

def switch_table(memc, table):
    key = "@@" + table
    print "Switching default table to '" + table + "' by issuing GET for '" + key + "'."
    result = memc.get(key)

def insert_country_data(memc):
    banner("Inserting initial data via memcached interface")
    for item in country_data:
        country = item[0]
        population = item[1]
        area = item[2]
        drive_side = item[3]

        key = country
        value = "|".join([population, area, drive_side])
        print "Key = " + key
        print "Value = " + value

        if memc.add(key, value):
            print "Added new key, value pair."
        else:
            print "Updating value for existing key."
            memc.set(key, value)

def query_country_data(memc):
    banner("Retrieving data for all keys (country names)")
    for item in country_data:
        key = item[0]
        result = memc.get(key)
        print "Here is the result retrieved from the database for key " + key + ":"
        print result
        (m_population, m_area, m_drive_side) = result.split("|")
        print "Unpacked population value: " + m_population
        print "Unpacked area value      : " + m_area
        print "Unpacked drive side value: " + m_drive_side
```

```
if __name__ == '__main__':

    memc = connect_to_memcached()
    switch_table(memc, "bbb")
    insert_country_data(memc)
    query_country_data(memc)

    sys.exit(0)
```

Here are some SQL queries to illustrate the state of the MySQL data after the script is run, and show how you could access the same data directly through SQL, or from an application written in any language using the appropriate [MySQL Connector or API](#).

The table descriptor 'bbb' is in place, allowing us to switch to the `multicol` table by issuing a `memcached` request `GET @bbb:`

```
mysql: use innodb_memcache;
Database changed

mysql: select * from containers;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| name | db_schema | db_table | key_columns | value_columns | flags | cas_column | expi |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| aaa  | test      | demo_test | c1          | c2          | c3    | c4        | c5   |
| bbb  | test      | multicol  | country     | population,area_sq_km,drive_side | c3    | c4        | c5   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.01 sec)
```

After running the script, the data is in the `multicol` table, available for traditional MySQL [queries](#) or [DML](#) statements:

```
mysql: use test;
Database changed

mysql: select * from multicol;
+-----+-----+-----+-----+-----+-----+-----+
| country | population | area_sq_km | drive_side | c3 | c4 | c5 |
+-----+-----+-----+-----+-----+-----+-----+
| Canada  | 34820000  | 9984670    | R          | 0  | 11 | 0  |
| China   | 1347350000 | 9640821    | R          | 0  | 20 | 0  |
| Denmark | 5543453   | 43094      | R          | 0  | 16 | 0  |
| India   | 1210193422 | 3287263    | L          | 0  | 19 | 0  |
| Ireland | 6399152   | 84421      | L          | 0  | 13 | 0  |
| Mexico  | 113910608 | 1972550    | R          | 0  | 15 | 0  |
| Norway  | 5002942   | 385252     | R          | 0  | 17 | 0  |
| UAE     | 8264070   | 83600      | R          | 0  | 18 | 0  |
| UK      | 62262000  | 243610     | L          | 0  | 14 | 0  |
| USA     | 314242000 | 9826675    | R          | 0  | 12 | 0  |
+-----+-----+-----+-----+-----+-----+-----+
10 rows in set (0.00 sec)

mysql: desc multicol;
+-----+-----+-----+-----+-----+-----+
| Field | Type | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| country | varchar(128) | NO | PRI | | |
| population | varchar(10) | YES | | NULL | |
| area_sq_km | varchar(9) | YES | | NULL | |
| drive_side | varchar(1) | YES | | NULL | |
| c3 | int(11) | YES | | NULL | |
| c4 | bigint(20) unsigned | YES | | NULL | |
| c5 | int(11) | YES | | NULL | |
+-----+-----+-----+-----+-----+-----+
7 rows in set (0.01 sec)
```

Allow sufficient size to hold all necessary digits, decimal points, sign characters, leading zeros, and so on when defining the length for columns that will be treated as numbers. Too-long values in a string column such as a `VARCHAR` are truncated by removing some characters, which might produce a nonsensical numeric value.



We can produce reports through SQL queries, doing calculations and tests across any columns, not just the `country` key column. (Because these examples use data from only a few countries, the numbers are for illustration purposes only.) Here, we find the average population of countries where people drive on the right, and the average size of countries whose names start with “U”:

```
mysql: select avg(population) from multicol where drive_side = 'R';
+-----+
| avg(population) |
+-----+
| 261304724.7142857 |
+-----+
1 row in set (0.00 sec)

mysql: select sum(area_sq_km) from multicol where country like 'U%';
+-----+
| sum(area_sq_km) |
+-----+
| 10153885 |
+-----+
1 row in set (0.00 sec)
```

Because the `population` and `area_sq_km` columns store character data rather than strongly typed numeric data, functions such as `avg()` and `sum()` work by converting each value to a number first. This approach *does not work* for operators such as `<` or `>`: for example, when comparing character-based values, `9 > 1000`, which is not you expect from a clause such as `ORDER BY population DESC`. For the most accurate type treatment, perform queries against views that cast numeric columns to the appropriate types. This technique lets you issue very simple `SELECT *` queries from your database applications, while ensuring that all casting, filtering, and ordering is correct. Here, we make a view that can be queried to find the top 3 countries in descending order of population, with the results always reflecting the latest data from the `multicol` table, and with the population and area figures always treated as numbers:

```
mysql: create view populous_countries as
select
  country,
  cast(population as unsigned integer) population,
  cast(area_sq_km as unsigned integer) area_sq_km,
  drive_side from multicol
order by cast(population as unsigned integer) desc
limit 3;
Query OK, 0 rows affected (0.01 sec)

mysql: select * from populous_countries;
+-----+-----+-----+-----+
| country | population | area_sq_km | drive_side |
+-----+-----+-----+-----+
| China   | 1347350000 | 9640821    | R          |
| India   | 1210193422 | 3287263    | L          |
| USA     | 314242000  | 9826675    | R          |
+-----+-----+-----+-----+
3 rows in set (0.00 sec)

mysql: desc populous_countries;
+-----+-----+-----+-----+-----+-----+
| Field      | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| country    | varchar(128)        | NO   |     |          |       |
| population | bigint(10) unsigned | YES  |     | NULL    |       |
| area_sq_km | int(9) unsigned     | YES  |     | NULL    |       |
| drive_side | varchar(1)          | YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.02 sec)
```

#### 14.2.9.5.2. Adapting an Existing `memcached` Application for the Integrated `memcached` Daemon

Consider these aspects of MySQL and `InnoDB` tables when adapting an existing `memcached` application to use the MySQL integration:

- If you have key values longer than a few bytes, you might find it more efficient to use a numeric auto-increment column for the [primary key](#) in the [InnoDB](#) table, and create a unique [secondary index](#) on the column holding the [memcached](#) key values. This is because [InnoDB](#) performs best for large-scale insertions if the primary key values are added in sorted order (as they are with auto-increment values), and the primary key values are duplicated in each secondary index, which can take up unnecessary space when the primary key is a long string value.
- If you store several different classes of information in [memcached](#), you might set up a separate [InnoDB](#) table for each kind of data. Define additional table identifiers in the [innodb\\_memcache.containers](#) table, and use the notation [@@table\\_id.key](#) to store or retrieve items from different tables. Physically dividing the items lets you tune the characteristics of each table for best space utilization, performance, and reliability. For example, you might enable [compression](#) for a table that holds blog posts, but not for one that holds thumbnail images. You might back up one table more frequently than another because it holds critical data. You might create additional [secondary indexes](#) on tables that are frequently used to generate reports through SQL.
- Preferably, set up a stable set of table definitions for use with the [memcached](#) interface and leave them in place permanently. Changes to the [containers](#) table take effect the next time that table is queried. The entries in that table are processed at startup, and are consulted whenever an unrecognized table ID is requested by the [@@](#) notation. Thus, new entries are visible as soon as you try to use the associated table ID, but changes to existing entries require a server restart before they take effect.
- When you use the default caching policy [innodb\\_only](#), your calls to [add\(\)](#), [set\(\)](#), [incr\(\)](#), and so on can succeed but still trigger debugging messages such as [while expecting 'STORED', got unexpected response 'NOT\\_STORED'](#). This is because in the [innodb\\_only](#) configuration, new and updated values are sent directly to the [InnoDB](#) table without being saved in the memory cache.

#### 14.2.9.5.3. Tuning Performance of the InnoDB memcached Plugin

Because using [InnoDB](#) in combination with [memcached](#) involves writing all data to disk, whether immediately or sometime later, understand that raw performance is expected to be somewhat lower than using [memcached](#) by itself. Focus your tuning goals for the [InnoDB memcached](#) plugin on achieving higher performance than equivalent SQL operations.

Benchmarks suggest that both queries and [DML](#) operations (inserts, updates, and deletes) are faster going through the [memcached](#) interface than with traditional SQL. DML operations typically see a larger speedup. Thus, the types of applications you might adapt to use the [memcached](#) interface first are those that are write-intensive. You might also use MySQL as a data store for types of write-intensive applications that formerly used some fast, lightweight mechanism where reliability was not a priority.

#### Adapting SQL Queries

The types of queries that are most suited to the simple [GET](#) request style are those with a single clause, or a set of [AND](#) conditions, in the [WHERE](#) clause:

```
SQL:
select col from tbl where key = 'key_value';

memcached:
GET key_value

SQL:
select col from tbl where col1 = val1 and col2 = val2 and col3 = val3;

memcached:
# Since you must always know these 3 values to look up the key,
# combine them into a unique string and use that as the key
# for all ADD, SET, and GET operations.
key_value = val1 + ":" + val2 + ":" + val3
GET key_value
```

```

SQL:
select 'key exists!' from tbl
  where exists (select coll from tbl where key = 'key_value') limit 1;

memcached:
# Test for existence of key by asking for its value and checking if the call succeeds,
# ignoring the value itself. For existence checking, you typically only store a very
# short value such as "1".
GET key_value

```

## Taking Advantage of System Memory

For best performance, deploy the [InnoDB memcached](#) plugin on machines that are configured like typical database servers: in particular, with the majority of system RAM devoted to the [InnoDB buffer pool](#) through the [innodb\\_buffer\\_pool\\_size \[1708\]](#) configuration option. For systems with multi-gigabyte buffer pools, consider raising the value of the [innodb\\_buffer\\_pool\\_instances \[1706\]](#) configuration option for maximum throughput when most operations involve data already cached in memory.

## Reducing Redundant I/O

[InnoDB](#) has a number of settings that let you choose the balance between high reliability in case of a crash, and the amount of I/O overhead during high write workloads. For example, consider setting the configuration options [innodb\\_doublewrite=0 \[1715\]](#) and [innodb\\_flush\\_log\\_at\\_trx\\_commit=2 \[1718\]](#). Measure the performance with different settings for the [innodb\\_flush\\_method \[1719\]](#) option. If the [binary log](#) is not turned on for the server, use the setting [innodb\\_support\\_xa=0 \[1749\]](#).

For other ways to reduce or tune I/O for table operations, see [Section 8.5.7, “Optimizing InnoDB Disk I/O”](#).

## Reducing Transactional Overhead

The default value of 1 for the configuration options [daemon\\_memcached\\_r\\_batch\\_size \[1699\]](#) and [daemon\\_memcached\\_w\\_batch\\_size \[1700\]](#) is intended for maximum reliability of results and safety of stored or updated data.

Depending on the type of application, you might increase one or both of these settings to reduce the overhead of frequent [commit](#) operations. On a busy system, you might increase [daemon\\_memcached\\_r\\_batch\\_size \[1699\]](#), knowing that changes to the data made through SQL might not become visible to [memcached](#) immediately (that is, until  $N$  more [get](#) operations were processed). When processing data where every write operation must be reliably stored, you would leave [daemon\\_memcached\\_w\\_batch\\_size \[1700\]](#) set to 1. You might increase it when processing large numbers of updates intended to only be used for statistical analysis, where it is not critical if the last  $N$  updates are lost in case of a crash.

For example, imagine a system that monitors traffic crossing a busy bridge, recording approximately 100,000 vehicles each day. If the application simply counts different types of vehicles to analyze traffic patterns, it might change [daemon\\_memcached\\_w\\_batch\\_size \[1700\]](#) from 1 to 100, reducing the I/O overhead for commit operations by 99%. In case of an unexpected outage, only a maximum of 100 records could be lost, which might be an acceptable margin of error. If instead the application was doing automated toll collection for each car, it would keep [daemon\\_memcached\\_w\\_batch\\_size \[1700\]](#) set to 1 to ensure that every toll record was immediately saved to disk.

Because of the way [InnoDB](#) organizes the [memcached](#) key values on disk, if you have a large number of keys to create, it can be faster to sort all the data items by the key value in your application and [add](#) them in sorted order, rather than creating them in arbitrary order.

The [memslap](#) command, which is part of the regular [memcached](#) distribution but not included with the MySQL server, can be useful for benchmarking different configurations. It can also be used to

generate sample key/value pairs that you can use in your own benchmarking. See [Section 15.6.3.3.6, “libmemcached Command-Line Utilities”](#) for details.

#### 14.2.9.5.4. Controlling Transactional Behavior of the InnoDB memcached Plugin

Unlike with the traditional [memcached](#), with the [InnoDB + memcached](#) combination you can control how “durable” are the data values produced through calls to [add](#), [set](#), [incr](#), and so on. Because MySQL places a high priority on durability and consistency of data, by default all data written through the [memcached](#) interface is always stored to disk, and calls to [get](#) always return the most recent value from disk. Although this default setting does not give the highest possible raw performance, it is still very fast compared to the traditional SQL interface for [InnoDB](#) tables.

As you gain experience with this feature, you can make the decision to relax the durability settings for non-critical classes of data, at the risk of possibly losing some updated values in case of an outage, or returning data that is slightly out-of-date.

##### Frequency of Commits

One tradeoff between durability and raw performance is how frequently new and changed data is [committed](#). If the data is critical, you want it to be committed immediately so that it is safe in case of any crash or outage. If the data is less critical, such as counters that would be reset after a crash, or debugging or logging data where you could afford to lose a few seconds worth, you might prefer the higher raw throughput that comes with less frequent commits.

When a [memcached](#) operation causes an insert, update, or delete in the underlying [InnoDB](#) table, that change might be committed to the underlying table instantly (if [daemon\\_memcached\\_w\\_batch\\_size=1 \[1700\]](#)) or some time later (if that configuration option value is greater than 1). In either case, the change cannot be rolled back. If you increase the value of [daemon\\_memcached\\_w\\_batch\\_size=1 \[1700\]](#) to avoid high I/O overhead during busy times, commits could become very infrequent when the workload decreases. As a safety measure, a background thread automatically commits changes made through the [memcached](#) API at regular intervals. The interval is controlled by the [innodb\\_api\\_bk\\_commit\\_interval \[1702\]](#) configuration option, and by default is 5 seconds.

When a [memcached](#) operation causes an insert or update in the underlying [InnoDB](#) table, the changed data is immediately visible to other [memcached](#) requests because the new value remains in the memory cache, even if it is not committed yet on the MySQL side.

##### Transaction Isolation

When a [memcached](#) operation such as [get](#) or [incr](#) causes a query or DML operation in the underlying [InnoDB](#) table, you can control whether it sees the very latest data written to the table, only data that has been committed, or other variations of transaction [isolation level](#). You control this feature through the [innodb\\_api\\_trx\\_level \[1704\]](#) configuration option. The numeric values specified with this option correspond to the familiar isolation level names such as [REPEATABLE READ \[1469\]](#). See the description of the [innodb\\_api\\_trx\\_level \[1704\]](#) option for the full list.

The stricter the isolation level, the more certain you can be that the data you retrieve will not be rolled back or changed suddenly so that a subsequent query sees a different value. But that strictness comes with greater [locking](#) overhead that can cause waits. For a NoSQL-style application that does not use long-running transactions, you can typically stay with the default isolation level or switch to a less strict one.

##### Allowing or Disallowing DDL

By default, you can perform [DDL](#) operations such as [ALTER TABLE](#) on the tables being used by the [InnoDB memcached](#) plugin. To avoid potential slowdowns when these tables are being used for high-throughput applications, you can disable DDL operations on these tables by turning on the [innodb\\_api\\_enable\\_mdl \[1703\]](#) configuration option at startup. This option is less appropriate

when you are accessing the same underlying tables through both the `memcached` interface and SQL, because it blocks `CREATE INDEX` statements on the tables, which could be important for configuring the system to run reporting queries.

### Data Stored on Disk, in Memory, or Both

Table `innodb_memcache.cache_policies` specifies whether to store data written through the `memcached` on disk (`innodb_only`, the default); to store the data in memory only, as in the traditional `memcached` (`cache-only`); or both (`caching`).

With the `caching` setting, if `memcached` cannot find a key in memory, it searches for the value in an `InnoDB` table. Values returned from `get` calls under the `caching` setting could be out-of-date, if they were updated on disk in the `InnoDB` table but not yet expired from the memory cache.

The caching policy can be set independently for `get`, `set` (including `incr` and `decr`), `delete`, and `flush` operations. For example:

- You might allow `get` and `set` operations to query or update a table and the `memcached` memory cache at the same time (through the `caching` setting), while making `delete`, `flush`, or both operate only on the in-memory copy (through the `cache_only` setting). That way, deleting or flushing an item just expires it from the cache, and the latest value is returned from the `InnoDB` table the next time the item is requested.

```
mysql> desc innodb_memcache.cache_policies;
+-----+-----+-----+-----+-----+-----+
| Field          | Type                               | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| policy_name    | varchar(40)                        | NO   | PRI | NULL    |      |
| get_policy     | enum('innodb_only','cache_only','caching','disabled') | NO   |     | NULL    |      |
| set_policy     | enum('innodb_only','cache_only','caching','disabled') | NO   |     | NULL    |      |
| delete_policy  | enum('innodb_only','cache_only','caching','disabled') | NO   |     | NULL    |      |
| flush_policy   | enum('innodb_only','cache_only','caching','disabled') | NO   |     | NULL    |      |
+-----+-----+-----+-----+-----+-----+

mysql> select * from innodb_memcache.cache_policies;
+-----+-----+-----+-----+-----+
| policy_name | get_policy | set_policy | delete_policy | flush_policy |
+-----+-----+-----+-----+-----+
| cache_policy | innodb_only | innodb_only | innodb_only | innodb_only |
+-----+-----+-----+-----+-----+

mysql> update innodb_memcache.cache_policies set set_policy = 'caching'
-> where policy_name = 'cache_policy';
```

The `cache_policies` values are only read at startup, and are tightly integrated with the operation of the `memcached` plugin. After changing any of the values in this table, uninstall the plugin and reinstall it:

```
mysql> uninstall plugin daemon_memcached;
Query OK, 0 rows affected (2.00 sec)
mysql> install plugin daemon_memcached soname "libmemcached.so";
Query OK, 0 rows affected (0.00 sec)
```

#### 14.2.9.5.5. Adapting DML Statements to `memcached` Operations

Benchmarks suggest that the `InnoDB memcached` plugin speeds up `DML` operations (inserts, updates, and deletes) more than it speeds up queries. You might focus your initial development efforts on write-intensive applications that are I/O-bound, and look for opportunities to use MySQL for new kinds of write-intensive applications.

- `INSERT INTO t1 (key,val) VALUES (some_key,some_value);`  
`SELECT val FROM t1 WHERE key = some_key;`  
`UPDATE t1 SET val = new_value WHERE key = some_key;`

```
UPDATE t1 SET val = val + x WHERE key = some_key;
DELETE FROM t1 WHERE key = some_key;
```

Single-row DML statements are the most straightforward kinds of statements to turn into `memcached` operations: `INSERT` becomes `add`, `UPDATE` becomes `set`, `incr` or `decr`, and `DELETE` becomes `delete`. When issued through the `memcached` interface, these operations are guaranteed to affect only 1 row because `key` is unique within the table.

In the preceding SQL examples, `t1` refers to the table currently being used by the `InnoDB memcached` plugin based on the configuration settings in the `innodb_memcache.containers` table, `key` refers to the column listed under `key_columns`, and `val` refers to the column listed under `value_columns`.

- ```
TRUNCATE TABLE t1;
DELETE FROM t1;
```

Corresponds to the `flush_all` operation, when `t1` is configured as the table for `memcached` operations as in the previous step. Removes all the rows in the table.

#### 14.2.9.5.6. Performing DML and DDL Statements on the Underlying InnoDB Table

You can access the InnoDB table (by default, `test.demo_test`) through the standard SQL interfaces. However, there are some restrictions:

- When query a table through SQL that is also being accessed through the `memcached` interface, remember that `memcached` operations can be configured to be committed periodically rather than after every write operation. This behavior is controlled by the `daemon_memcached_w_batch_size [1700]` option. If this option is set to a value greater than 1, use `READ UNCOMMITTED [1470]` queries to find the just-inserted rows:

```
mysql> set session TRANSACTION ISOLATION LEVEL read uncommitted;
Query OK, 0 rows affected (0.00 sec)

mysql> select * from demo_test;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| cx   | cy   | cl   | cz   | c2   | ca   | CB   | c3   | cu   | c4   | C5   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| NULL | NULL | all  | NULL | 123456789 | NULL | NULL | 10  | NULL | 3   | NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

- To modify a table through SQL that is also being accessed through the `memcached` interface, remember that `memcached` operations can be configured to be start a new transaction periodically rather than for every read operation. This behavior is controlled by the `daemon_memcached_r_batch_size [1699]` option. If this option is set to a value greater than 1, ...
- The `InnoDB` table is locked IS (shared intention) or IX (exclusive intentional) for all operations in a transaction. If you increase `daemon_memcached_r_batch_size [1699]` and `daemon_memcached_w_batch_size [1700]` substantially from their default value of 1, the table is most likely intentionally locked between each operation, preventing you from running `DDL` statements on the table.

#### 14.2.9.6. Using the InnoDB memcached Plugin with Replication

Because the `InnoDB memcached` daemon plugin supports the MySQL `binary log`, any updates made on a `master server` through the `memcached` interface can be replicated for backup, balancing intensive read workloads, and high availability. All `memcached` commands are supported for binlogging.

You do not need to set up the `InnoDB memcached` plugin on the `slave servers`. In this configuration, the primary advantage is increased write throughput on the master. The speed of the replication mechanism is not affected.



The following sections show how to use the binlog capability, to use the [InnoDB memcached](#) plugin along with MySQL replication. It assumes you have already done the basic setup described in [Section 14.2.9.3, “Getting Started with InnoDB Memcached Plugin”](#).

### Enable InnoDB Memcached Binary Log with `innodb_api_enable_binlog`:

- To use the [InnoDB memcached](#) plugin with the MySQL [binary log](#), enable the `innodb_api_enable_binlog` [1703] configuration option on the [master server](#). This option can only be set at server boot time. You must also enable the MySQL binary log on the master server with the `--log-bin` [2008] option. You can add these options to your server configuration file such as `my.cnf`, or on the `mysqld` command line.

```
mysqld ... --log-bin --innodb_api_enable_binlog=1
```

- Then configure your master and slave server, as described in [Section 16.1.1, “How to Set Up Replication”](#).
- Use `mysqldump` to create a master data snapshot, and sync it to the slave server.

```
master shell: mysqldump --all-databases --lock-all-tables > dbdump.db
slave shell: mysql < dbdump.db
```

- On the master server, issue `show master status` to obtain the Master Binary Log Coordinates:

```
mysql> show master status;
```

- On the slave server, use a `change master to` statement to set up a slave server with the above coordinates:

```
mysql> CHANGE MASTER TO
  MASTER_HOST='localhost',
  MASTER_USER='root',
  MASTER_PASSWORD='',
  MASTER_PORT = 13000,
  MASTER_LOG_FILE='0.000001',
  MASTER_LOG_POS=114;
```

- Then start the slave:

```
mysql> start slave;
```

If the error log prints output similar to the following, the slave is ready for replication:

```
111002 18:47:15 [Note] Slave I/O thread: connected to master 'root@localhost:13000',
replication started in log '0.000001' at position 114
```

### Test with the [memcached telnet](#) interface:

To test the server with the above replication setup, we use the [memcached telnet](#) interface, and also query the master and slave servers using SQL to verify the results.

In our configuration setup SQL, one example table `demo_test` is created in the `test` database for use by [memcached](#). We will use this default table for the demonstrations:

- Use `set` to insert a record, key `test1`, value `t1`, and flag `10`:

```
telnet 127.0.0.1 11211
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.
set test1 10 0 2
t1
STORED
```



In the master server, you can see that the row is inserted. `c1` maps to the key, `c2` maps to the value, `c3` is the flag, `c4` is the `cas` value, and `c5` is the expiration.

```
mysql> select * from test.demo_test;
```

| c1    | c2 | c3 | c4 | c5 |
|-------|----|----|----|----|
| test1 | t1 | 10 | 2  | 0  |

```
1 row in set (0.00 sec)
```

In the slave server, you will see the same record is inserted by replication:

```
mysql> select * from test.demo_test;
```

| c1    | c2 | c3 | c4 | c5 |
|-------|----|----|----|----|
| test1 | t1 | 10 | 2  | 0  |

```
1 row in set (0.00 sec)
```

- Use `set` command to update the key `test1` to a new value `new`:

```
Connected to 127.0.0.1.
Escape character is '^]'.
set test1 10 0 3
new
STORED
```

From the slave server, the update is replicated (notice the `cas` value also updated):

```
mysql> select * from test.demo_test;
```

| c1    | c2  | c3 | c4 | c5 |
|-------|-----|----|----|----|
| test1 | new | 10 | 3  | 0  |

```
1 row in set (0.00 sec)
```

- Delete the record with a `delete` command:

```
Connected to 127.0.0.1.
Escape character is '^]'.
delete test1
DELETED
```

When the delete is replicated to the slave, the record on the slave is also deleted:

```
mysql> select * from test.demo_test;
Empty set (0.00 sec)
```

- Truncate the table with the `flush_all` command.

First, insert two records by telnetting to the master server:

```
Connected to 127.0.0.1.
Escape character is '^]'.
set test2 10 0 5
again
STORED
set test3 10 0 6
again1
```

```
STORED
```

In the slave server, confirm these two records are replicated:

```
mysql> select * from test.demo_test;
```

| c1    | c2     | c3 | c4 | c5 |
|-------|--------|----|----|----|
| test2 | again  | 10 | 5  | 0  |
| test3 | again1 | 10 | 6  | 0  |

```
2 rows in set (0.00 sec)
```

Call `flush_all` in the telnet interface to truncate the table:

```
Connected to 127.0.0.1.
Escape character is '^]'.
flush_all
OK
```

Then check that the truncation operation is replicated on the slave server:

```
mysql> select * from test.demo_test;
Empty set (0.00 sec)
```

All `memcached` commands are supported in terms of replication.

### Notes for the InnoDB Memcached Binlog:

Binlog Format:

- Most `memcached` operations are mapped to **DML** statements (analogous to insert, delete, update). Since there is no actual SQL statement being processed by the MySQL server, all `memcached` commands (except for `flush_all`) use Row-Based Replication (RBR) logging. This is independent of any server `binlog_format` [2010] setting.
- The `memcached flush_all` command is mapped to the `TRUNCATE TABLE` command. Since **DDL** commands can only use statement-based logging, this `flush_all` command is replicated by sending a `TRUNCATE TABLE` statement.

Transactions:

- The concept of **transactions** has not typically been part of `memcached` applications. We use `daemon_memcached_r_batch_size` [1699] and `daemon_memcached_w_batch_size` [1700] to control the read and write transaction batch size for performance considerations. These settings do not affect replication: each SQL operation on the underlying table is replicated right after successful completion.
- The default value of `daemon_memcached_w_batch_size` [1700] is 1, so each `memcached` write operation is committed immediately. This default setting incurs a certain amount of performance overhead, to avoid any inconsistency in the data visible on the master and slave servers. The replicated records will always be available immediately on the slave server. If you set `daemon_memcached_w_batch_size` greater than 1, records inserted or updated through the `memcached` interface are not immediately visible on the master server; to view these records on the master server before they are committed, issue `set transaction isolation level read uncommitted`.

### 14.2.9.7. Internals of the InnoDB memcached Plugin

This section explains the details of the underlying tables used by the `InnoDB / memcached` plugin.

The configuration script installs 3 tables needed by the InnoDB `memcached`. These tables are created in a dedicated database `innodb_memcache`:

```
mysql> use innodb_memcache;
Database changed
mysql> show tables;
+-----+
| Tables_in_innodb_memcache |
+-----+
| cache_policies            |
| config_options           |
| containers                |
+-----+
3 rows in set (0.01 sec)
```

The tables are:

`containers` - This table is the most important table for the `memcached` daemon. It describes the table or tables used to store the `memcached` values. You must make changes to this table to start using the `memcached` interface with one or more of your own tables, rather than just experimenting with the `test.demo_test` table.

The mapping is done through specifying corresponding column values in the table:

```
mysql> desc containers;
+-----+-----+-----+-----+-----+-----+
| Field          | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
name	varchar(50)	NO	PRI	NULL	
db_schema	varchar(250)	NO		NULL	
db_table	varchar(250)	NO		NULL	
key_columns	varchar(250)	NO		NULL	
value_columns	varchar(250)	YES		NULL	
flags	varchar(250)	NO		0	
cas_column	varchar(250)	YES		NULL	
expire_time_column	varchar(250)	YES		NULL	
unique_idx_name_on_key	varchar(250)	NO		NULL	
+-----+-----+-----+-----+-----+-----+
9 rows in set (0.02 sec)
```

- `db_schema` and `db_table` columns specify the database and table name for storing the `memcached` value.
- `key_columns` specifies the single column name used as the lookup key for `memcached` operations.
- `value_columns` describes the columns (one or more) used as values for `memcached` operations. To specify multiple columns, separate them with pipe characters (such as `col1|col2|col3` and so on).
- `unique_idx_name_on_key` is the name of the index on the key column. It must be a unique index. It can be the [primary key](#) or a [secondary index](#). Preferably, make the key column the primary key of the `InnoDB` table. Doing so saves a lookup step over using a secondary index for this column. You cannot make a [covering index](#) for `memcached` lookups; `InnoDB` returns an error if you try to define a composite secondary index over both the key and value columns.

The above 5 column values (table name, key column, value column and index) must be supplied. Otherwise, the setup will fail.

Although the following values are optional, they are needed for full compliance with the `memcached` protocol:

- `flags` specifies the columns used as flags (a user-defined numeric value that is stored and retrieved along with the main value) for `memcached`. It is also used as the column specifier for some operations (such as `incr`, `prepend`) if `memcached` value is mapped to multiple columns. So the

operation would be done on the specified column. For example, if you have mapped a value to 3 columns, and only want the increment operation performed on one of these columns, you can use `flags` to specify which column will be used for these operations.

- `cas_column` and `exp_column` are used specifically to store the `cas` (compare-and-swap) and `exp` (expiry) value of `memcached`. Those values are related to the way `memcached` hashes requests to different servers and caches data in memory. Because the `InnoDB memcached` plugin is so tightly integrated with a single `memcached` daemon, and the in-memory caching mechanism is handled by MySQL and the `buffer pool`, these columns are rarely needed in this type of deployment.

Table `cache_policies` specifies whether to use `InnoDB` as the data store of `memcached` (`innodb_only`), or to use the traditional `memcached` engine as the backstore (`cache-only`), or both (`caching`). In the last case, if `memcached` cannot find a key in memory, it searches for the value in an `InnoDB` table.

Table `config_options` stores `memcached`-related settings that are appropriate to change at runtime, through SQL. Currently, MySQL supports the following configuration options through this table:

`separator`: The separator used to separate values of a long string into smaller values for multiple columns values. By default, this is the `|` character. For example, if you defined `col1`, `col2` as value columns, And you define `|` as separator, you could issue the following command in `memcached` to insert values into `col1` and `col2` respectively:

```
set keyx 10 0 19
valuecolx|valuecoly
```

So `valuecol1x` is stored in `col1` and `valuecoly` is stored in `col2`.

`table_map_delimiter`: The character separating the schema name and the table name when you use the `@@` notation in a key name to access a key in a specific table. For example, `@@t1.some_key` and `@@t2.some_key` have the same key value, but are stored in different tables and so do not conflict.

### Example Tables

Finally, the configuration script creates a table `demo_test` in the `test` database as an example. It also allows the Daemon Memcached to work immediately, without creating any additional tables.

The entries in the `container` table define which column is used for what purpose as described above:

```
mysql> select * from innodb_memcache.containers;
+-----+-----+-----+-----+-----+-----+-----+-----+
| name | db_schema | db_table | key_columns | value_columns | flags | cas_column | expire_time_column |
+-----+-----+-----+-----+-----+-----+-----+-----+
| aaa  | test      | demo_test | c1          | c2          | c3    | c4        | c5          |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)

mysql> desc test.demo_test;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
c1	varchar(32)	NO	PRI		
c2	varchar(1024)	YES		NULL	
c3	int(11)	YES		NULL	
c4	bigint(20) unsigned	YES		NULL	
c5	int(11)	YES		NULL	
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.01 sec)
```

When no table ID is requested through the `@@` notation in the key name:

- If a row has a `name` value of `default`, the corresponding table is used by the `memcached` plugin. Thus, when you make your first entry in `innodb_memcache.containers` to move beyond the `demo_test` table, use a `name` value of `default`.

- If there is no `innodb_memcache.containers.name` value of `default`, the row with the first `name` value in alphabetical order is used.

### 14.2.9.8. Troubleshooting the InnoDB memcached Plugin

The following list shows some potential issues you might encounter using the `InnoDB memcached` daemon, and solutions or workarounds where available:

- If you see this error in your MySQL error log, the server might fail to start:

```
failed to set rlimit for open files. Try running as root or requesting
smaller maxconns value.
```

The error message is actually from the `memcached` daemon. One solution is to raise the OS limit for the number of open files. The command varies depending on the operating system. For example, here are the commands to check and increase the limit on several operating systems:

```
# Linux
$ ulimit -n
1024
ulimit -n 4096
$ ulimit -n
4096

# OS X Lion (10.6)
$ ulimit -n
256
ulimit -n 4096
$ ulimit -n
4096
```

The other solution is to reduce the number of concurrent connections available for the `memcached` daemon, using the `-c` option which defaults to 1024. Encode that `memcached` option using the MySQL option `daemon_memcached_option` [1699] inside the MySQL configuration file:

```
[mysqld]
...
loose-daemon_memcached_option='-c 64'
```

- To troubleshoot problems where the `memcached` daemon is unable to store data in or retrieve data from the `InnoDB` table, specify the `memcached` option `-vvv` through the MySQL configuration option `daemon_memcached_option` [1699]. Examine the MySQL error log for debug output related to `memcached` operations.
- If the column specified to hold the `memcached` item values is the wrong data type, such as a numeric type instead of a string type, attempts to store key/value pairs will fail with no specific error code or message.
- If the `daemon_memcached` plugin causes any issues with starting the MySQL server, disable it during troubleshooting by adding this line under the `[mysqld]` group in your MySQL configuration file:

```
daemon_memcached=OFF
```

For example, if you run the `install plugin` command before running the `scripts/innodb_memcached_config.sql` script to set up the necessary database and tables, the server might crash and be unable to start. Or, if you set up an incorrect entry in the `innodb_memcache.containers` table, the server might be unable to start.

To permanently turn off the `memcached` plugin for a MySQL instance, issue the following command:

```
mysql> uninstall plugin daemon_memcached;
```

- If you run more than one instance of MySQL on the same machine, with the `memcached` daemon plugin enabled in each, make sure to specify a unique `memcached` port for each one using the `daemon_memcached_option` [1699] configuration option.
- You might find that a SQL statement cannot find an expected table, or there is no data in the table, but `memcached` API calls still work and retrieve the expected data. This can happen if you do not set up the entry in the `innodb_memcache.containers` table, or do not switch to that table by issuing a `GET` or `SET` request with the key `@@table_id`, or make a change to an existing entry in `innodb_memcache.containers` without restarting the MySQL server afterward. The free-form storage mechanism is flexible enough that your requests to store or retrieve a multi-column value like `col1|col2|col3` will usually still work, even if the daemon is using the `test.demo_test` table which stores all the data within a single column.
- As the length of the `memcached` key and values increase, you encounter size and length limits at different points:
  - When the key exceeds 250 bytes in size, `memcached` operations return an error. This is currently a fixed limit within `memcached`.
  - You might encounter `InnoDB`-related limits when the value exceeds 768 bytes in size, or 3072 bytes in size, or 1/2 of the size specified by `innodb_page_size` [1740]. These limits primarily apply if you intend to create an index on the value column to run report-generating queries on that column from SQL. See [Section 14.2.7, “Limits on InnoDB Tables”](#) for details.
  - The maximum size for the combination of the key and the value is 1 MB.
- If you share configuration files across MySQL servers with different versions, using the latest configuration options for the `memcached` plugin could cause startup errors for older MySQL versions. To avoid compatibility problems, use the `loose` forms of these option names, for example `loose-daemon_memcached_option=-c 64` instead of `daemon_memcached_option=-c 64`.

### 14.3. The **MyISAM** Storage Engine

**MyISAM** is based on the older (and no longer available) **ISAM** storage engine but has many useful extensions.

**Table 14.7. **MyISAM** Storage Engine Features**

|                                                    |       |                                        |                  |                                              |       |
|----------------------------------------------------|-------|----------------------------------------|------------------|----------------------------------------------|-------|
| <b>Storage limits</b>                              | 256TB | <b>Transactions</b>                    | No               | <b>Locking granularity</b>                   | Table |
| <b>MVCC</b>                                        | No    | <b>Geospatial data type support</b>    | Yes              | <b>Geospatial indexing support</b>           | Yes   |
| <b>B-tree indexes</b>                              | Yes   | <b>T-tree indexes</b>                  | No               | <b>Hash indexes</b>                          | No    |
| <b>Full-text search indexes</b>                    | Yes   | <b>Clustered indexes</b>               | No               | <b>Data caches</b>                           | No    |
| <b>Index caches</b>                                | Yes   | <b>Compressed data</b>                 | Yes <sup>a</sup> | <b>Encrypted data<sup>b</sup></b>            | Yes   |
| <b>Cluster database support</b>                    | No    | <b>Replication support<sup>c</sup></b> | Yes              | <b>Foreign key support</b>                   | No    |
| <b>Backup / point-in-time recovery<sup>d</sup></b> | Yes   | <b>Query cache support</b>             | Yes              | <b>Update statistics for data dictionary</b> | Yes   |

<sup>a</sup>Compressed **MyISAM** tables are supported only when using the compressed row format. Tables using the compressed row format with **MyISAM** are read only.

<sup>b</sup>Implemented in the server (via encryption functions), rather than in the storage engine.

<sup>c</sup>Implemented in the server, rather than in the storage engine.

<sup>d</sup>Implemented in the server, rather than in the storage engine.

Each **MyISAM** table is stored on disk in three files. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table format. The data file has an `.MYD` (**MYData**) extension. The index file has an `.MYI` (**MYIndex**) extension.

To specify explicitly that you want a `MyISAM` table, indicate that with an `ENGINE` table option:

```
CREATE TABLE t (i INT) ENGINE = MYISAM;
```

In MySQL 5.6, it is normally necessary to use `ENGINE` to specify the `MyISAM` storage engine because `InnoDB` is the default engine.

You can check or repair `MyISAM` tables with the `mysqlcheck` client or `myisamchk` utility. You can also compress `MyISAM` tables with `myisampack` to take up much less space. See [Section 4.5.3, “mysqlcheck — A Table Maintenance Program”](#), [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#), and [Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).

`MyISAM` tables have the following characteristics:

- All data values are stored with the low byte first. This makes the data machine and operating system independent. The only requirements for binary portability are that the machine uses two's-complement signed integers and IEEE floating-point format. These requirements are widely used among mainstream machines. Binary compatibility might not be applicable to embedded systems, which sometimes have peculiar processors.

There is no significant speed penalty for storing data low byte first; the bytes in a table row normally are unaligned and it takes little more processing to read an unaligned byte in order than in reverse order. Also, the code in the server that fetches column values is not time critical compared to other code.

- All numeric key values are stored with the high byte first to permit better index compression.
- Large files (up to 63-bit file length) are supported on file systems and operating systems that support large files.
- There is a limit of  $(2^{32})^2$  (1.844E+19) rows in a `MyISAM` table.
- The maximum number of indexes per `MyISAM` table is 64.

The maximum number of columns per index is 16.

- The maximum key length is 1000 bytes. This can also be changed by changing the source and recompiling. For the case of a key longer than 250 bytes, a larger key block size than the default of 1024 bytes is used.
- When rows are inserted in sorted order (as when you are using an `AUTO_INCREMENT` column), the index tree is split so that the high node only contains one key. This improves space utilization in the index tree.
- Internal handling of one `AUTO_INCREMENT` column per table is supported. `MyISAM` automatically updates this column for `INSERT` and `UPDATE` operations. This makes `AUTO_INCREMENT` columns faster (at least 10%). Values at the top of the sequence are not reused after being deleted. (When an `AUTO_INCREMENT` column is defined as the last column of a multiple-column index, reuse of values deleted from the top of a sequence does occur.) The `AUTO_INCREMENT` value can be reset with `ALTER TABLE` or `myisamchk`.
- Dynamic-sized rows are much less fragmented when mixing deletes with updates and inserts. This is done by automatically combining adjacent deleted blocks and by extending blocks if the next block is deleted.
- `MyISAM` supports concurrent inserts: If a table has no free blocks in the middle of the data file, you can `INSERT` new rows into it at the same time that other threads are reading from the table. A free block can occur as a result of deleting rows or an update of a dynamic length row with more data than its current contents. When all free blocks are used up (filled in), future inserts become concurrent again. See [Section 8.10.3, “Concurrent Inserts”](#).



- You can put the data file and index file in different directories on different physical devices to get more speed with the `DATA DIRECTORY` and `INDEX DIRECTORY` table options to `CREATE TABLE`. See [Section 13.1.14, “CREATE TABLE Syntax”](#).
- `BLOB` and `TEXT` columns can be indexed.
- `NULL` values are permitted in indexed columns. This takes 0 to 1 bytes per key.
- Each character column can have a different character set. See [Section 10.1, “Character Set Support”](#).
- There is a flag in the `MyISAM` index file that indicates whether the table was closed correctly. If `mysqld` is started with the `--myisam-recover-options` [\[432\]](#) option, `MyISAM` tables are automatically checked when opened, and are repaired if the table wasn't closed properly.
- `myisamchk` marks tables as checked if you run it with the `--update-state` [\[321\]](#) option. `myisamchk --fast` checks only those tables that don't have this mark.
- `myisamchk --analyze` stores statistics for portions of keys, as well as for entire keys.
- `myisampack` can pack `BLOB` and `VARCHAR` columns.

`MyISAM` also supports the following features:

- Support for a true `VARCHAR` type; a `VARCHAR` column starts with a length stored in one or two bytes.
- Tables with `VARCHAR` columns may have fixed or dynamic row length.
- The sum of the lengths of the `VARCHAR` and `CHAR` columns in a table may be up to 64KB.
- Arbitrary length `UNIQUE` constraints.

## Additional Resources

- A forum dedicated to the `MyISAM` storage engine is available at <http://forums.mysql.com/list.php?21>.

### 14.3.1. `MyISAM` Startup Options

The following options to `mysqld` can be used to change the behavior of `MyISAM` tables. For additional information, see [Section 5.1.3, “Server Command Options”](#).

**Table 14.8. `MyISAM` Option/Variable Reference**

| Name                                                        | Cmd-Line                  | Option file | System Var | Status Var | Var Scope | Dynamic |
|-------------------------------------------------------------|---------------------------|-------------|------------|------------|-----------|---------|
| <code>bulk_insert_buffer</code>                             | Yes <a href="#">[466]</a> | Yes         | Yes        |            | Both      | Yes     |
| <code>concurrent_insert</code>                              | Yes <a href="#">[466]</a> | Yes         | Yes        |            | Global    | Yes     |
| <code>delay-key-write</code> <a href="#">[421]</a>          | Yes                       | Yes         |            |            | Global    | Yes     |
| - Variable:<br><code>delay_key_write</code>                 |                           |             | Yes        |            | Global    | Yes     |
| <code>have_rtree_keys</code> <a href="#">[487]</a>          |                           |             | Yes        |            | Global    | No      |
| <code>key_buffer_size</code> <a href="#">[492]</a>          | Yes                       | Yes         | Yes        |            | Global    | Yes     |
| <code>log-isam</code> <a href="#">[427]</a>                 | Yes                       | Yes         |            |            |           |         |
| <code>myisam-block-size</code> <a href="#">[432]</a>        | Yes                       | Yes         |            |            |           |         |
| <code>myisam_data_pointer_size</code> <a href="#">[511]</a> | Yes                       | Yes         | Yes        |            | Global    | Yes     |
| <code>myisam_max_sort_length</code> <a href="#">[511]</a>   | Yes                       | Yes         | Yes        |            | Global    | Yes     |

| Name                                               | Cmd-Line | Option file | System Var | Status Var | Var Scope | Dynamic |
|----------------------------------------------------|----------|-------------|------------|------------|-----------|---------|
| <a href="#">myisam_mmap_size [512]</a>             | Yes      | Yes         | Yes        |            | Global    | No      |
| <a href="#">myisam-recover-options [432]</a>       | Yes      | Yes         |            |            |           |         |
| - Variable:<br><code>myisam_recover_options</code> |          |             |            |            |           |         |
| <a href="#">myisam_recover_options [512]</a>       |          |             | Yes        |            | Global    | No      |
| <a href="#">myisam_repair_threads [512]</a>        | Yes      | Yes         | Yes        |            | Both      | Yes     |
| <a href="#">myisam_sort_buffer_size [513]</a>      | Yes      | Yes         | Yes        |            | Both      | Yes     |
| <a href="#">myisam_stats_method [513]</a>          | Yes      | Yes         | Yes        |            | Both      | Yes     |
| <a href="#">myisam_use_mmap [514]</a>              | Yes      | Yes         | Yes        |            | Global    | Yes     |
| <a href="#">skip-concurrent-insert [439]</a>       | Yes      | Yes         |            |            |           |         |
| - Variable:<br><code>concurrent_insert</code>      |          |             |            |            |           |         |
| <a href="#">tmp_table_size [558]</a>               | Yes      | Yes         | Yes        |            | Both      | Yes     |

- `--myisam-recover-options=mode [432]`  
Set the mode for automatic recovery of crashed [MyISAM](#) tables.
- `--delay-key-write=ALL [421]`  
Don't flush key buffers between writes for any [MyISAM](#) table.



#### Note

If you do this, you should not access [MyISAM](#) tables from another program (such as from another MySQL server or with `myisamchk`) when the tables are in use. Doing so risks index corruption. Using `--external-locking [422]` does not eliminate this risk.

The following system variables affect the behavior of [MyISAM](#) tables. For additional information, see [Section 5.1.4, "Server System Variables"](#).

- `bulk_insert_buffer_size [466]`  
The size of the tree cache used in bulk insert optimization.



#### Note

This is a limit *per thread*!

- `myisam_max_sort_file_size [511]`  
The maximum size of the temporary file that MySQL is permitted to use while re-creating a [MyISAM](#) index (during `REPAIR TABLE`, `ALTER TABLE`, or `LOAD DATA INFILE`). If the file size would be larger than this value, the index is created using the key cache instead, which is slower. The value is given in bytes.
- `myisam_sort_buffer_size [513]`  
Set the size of the buffer used when recovering tables.

Automatic recovery is activated if you start `mysqld` with the `--myisam-recover-options [432]` option. In this case, when the server opens a [MyISAM](#) table, it checks whether the table is marked as

crashed or whether the open count variable for the table is not 0 and you are running the server with external locking disabled. If either of these conditions is true, the following happens:

- The server checks the table for errors.
- If the server finds an error, it tries to do a fast table repair (with sorting and without re-creating the data file).
- If the repair fails because of an error in the data file (for example, a duplicate-key error), the server tries again, this time re-creating the data file.
- If the repair still fails, the server tries once more with the old repair option method (write row by row without sorting). This method should be able to repair any type of error and has low disk space requirements.

If the recovery wouldn't be able to recover all rows from previously completed statements and you didn't specify `FORCE` in the value of the `--myisam-recover-options` [432] option, automatic repair aborts with an error message in the error log:

```
Error: Couldn't repair table: test.g00pages
```

If you specify `FORCE`, a warning like this is written instead:

```
Warning: Found 344 of 354 rows when repairing ./test/g00pages
```

Note that if the automatic recovery value includes `BACKUP`, the recovery process creates files with names of the form `tbl_name-datetime.BAK`. You should have a `cron` script that automatically moves these files from the database directories to backup media.

### 14.3.2. Space Needed for Keys

`MyISAM` tables use B-tree indexes. You can roughly calculate the size for the index file as  $(key\_length+4)/0.67$ , summed over all keys. This is for the worst case when all keys are inserted in sorted order and the table doesn't have any compressed keys.

String indexes are space compressed. If the first index part is a string, it is also prefix compressed. Space compression makes the index file smaller than the worst-case figure if a string column has a lot of trailing space or is a `VARCHAR` column that is not always used to the full length. Prefix compression is used on keys that start with a string. Prefix compression helps if there are many strings with an identical prefix.

In `MyISAM` tables, you can also prefix compress numbers by specifying the `PACK_KEYS=1` table option when you create the table. Numbers are stored with the high byte first, so this helps when you have many integer keys that have an identical prefix.

### 14.3.3. `MyISAM` Table Storage Formats

`MyISAM` supports three different storage formats. Two of them, fixed and dynamic format, are chosen automatically depending on the type of columns you are using. The third, compressed format, can be created only with the `myisampack` utility (see Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only `MyISAM` Tables”).

When you use `CREATE TABLE` or `ALTER TABLE` for a table that has no `BLOB` or `TEXT` columns, you can force the table format to `FIXED` or `DYNAMIC` with the `ROW_FORMAT` table option.

See Section 13.1.14, “`CREATE TABLE Syntax`”, for information about `ROW_FORMAT`.

You can decompress (unpack) compressed `MyISAM` tables using `myisamchk --unpack` [323]; see Section 4.6.3, “`myisamchk` — `MyISAM` Table-Maintenance Utility”, for more information.

### 14.3.3.1. Static (Fixed-Length) Table Characteristics

Static format is the default for MyISAM tables. It is used when the table contains no variable-length columns (`VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT`). Each row is stored using a fixed number of bytes.

Of the three MyISAM storage formats, static format is the simplest and most secure (least subject to corruption). It is also the fastest of the on-disk formats due to the ease with which rows in the data file can be found on disk: To look up a row based on a row number in the index, multiply the row number by the row length to calculate the row position. Also, when scanning a table, it is very easy to read a constant number of rows with each disk read operation.

The security is evidenced if your computer crashes while the MySQL server is writing to a fixed-format MyISAM file. In this case, `myisamchk` can easily determine where each row starts and ends, so it can usually reclaim all rows except the partially written one. Note that MyISAM table indexes can always be reconstructed based on the data rows.



#### Note

Fixed-length row format is only available for tables without `BLOB` or `TEXT` columns. Creating a table with these columns with an explicit `ROW_FORMAT` clause will not raise an error or warning; the format specification will be ignored.

Static-format tables have these characteristics:

- `CHAR` and `VARCHAR` columns are space-padded to the specified column width, although the column type is not altered. `BINARY` and `VARBINARY` columns are padded with `0x00` bytes to the column width.
- Very quick.
- Easy to cache.
- Easy to reconstruct after a crash, because rows are located in fixed positions.
- Reorganization is unnecessary unless you delete a huge number of rows and want to return free disk space to the operating system. To do this, use `OPTIMIZE TABLE` or `myisamchk -r`.
- Usually require more disk space than dynamic-format tables.

### 14.3.3.2. Dynamic Table Characteristics

Dynamic storage format is used if a MyISAM table contains any variable-length columns (`VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT`), or if the table was created with the `ROW_FORMAT=DYNAMIC` table option.

Dynamic format is a little more complex than static format because each row has a header that indicates how long it is. A row can become fragmented (stored in noncontiguous pieces) when it is made longer as a result of an update.

You can use `OPTIMIZE TABLE` or `myisamchk -r` to defragment a table. If you have fixed-length columns that you access or change frequently in a table that also contains some variable-length columns, it might be a good idea to move the variable-length columns to other tables just to avoid fragmentation.

Dynamic-format tables have these characteristics:

- All string columns are dynamic except those with a length less than four.
- Each row is preceded by a bitmap that indicates which columns contain the empty string (for string columns) or zero (for numeric columns). Note that this does not include columns that contain `NULL` values. If a string column has a length of zero after trailing space removal, or a numeric column has

a value of zero, it is marked in the bitmap and not saved to disk. Nonempty strings are saved as a length byte plus the string contents.

- Much less disk space usually is required than for fixed-length tables.
- Each row uses only as much space as is required. However, if a row becomes larger, it is split into as many pieces as are required, resulting in row fragmentation. For example, if you update a row with information that extends the row length, the row becomes fragmented. In this case, you may have to run `OPTIMIZE TABLE` or `myisamchk -r` from time to time to improve performance. Use `myisamchk -ei` to obtain table statistics.
- More difficult than static-format tables to reconstruct after a crash, because rows may be fragmented into many pieces and links (fragments) may be missing.
- The expected row length for dynamic-sized rows is calculated using the following expression:

```
3
+ (number of columns + 7) / 8
+ (number of char columns)
+ (packed size of numeric columns)
+ (length of strings)
+ (number of NULL columns + 7) / 8
```

There is a penalty of 6 bytes for each link. A dynamic row is linked whenever an update causes an enlargement of the row. Each new link is at least 20 bytes, so the next enlargement probably goes in the same link. If not, another link is created. You can find the number of links using `myisamchk -ed`. All links may be removed with `OPTIMIZE TABLE` or `myisamchk -r`.

### 14.3.3.3. Compressed Table Characteristics

Compressed storage format is a read-only format that is generated with the `myisampack` tool. Compressed tables can be uncompressed with `myisamchk`.

Compressed tables have the following characteristics:

- Compressed tables take very little disk space. This minimizes disk usage, which is helpful when using slow disks (such as CD-ROMs).
- Each row is compressed separately, so there is very little access overhead. The header for a row takes up one to three bytes depending on the biggest row in the table. Each column is compressed differently. There is usually a different Huffman tree for each column. Some of the compression types are:
  - Suffix space compression.
  - Prefix space compression.
  - Numbers with a value of zero are stored using one bit.
  - If values in an integer column have a small range, the column is stored using the smallest possible type. For example, a `BIGINT` column (eight bytes) can be stored as a `TINYINT` column (one byte) if all its values are in the range from `-128` to `127`.
  - If a column has only a small set of possible values, the data type is converted to `ENUM`.
  - A column may use any combination of the preceding compression types.
- Can be used for fixed-length or dynamic-length rows.



#### Note

While a compressed table is read only, and you cannot therefore update or add rows in the table, DDL (Data Definition Language) operations are still valid. For

example, you may still use `DROP` to drop the table, and `TRUNCATE TABLE` to empty the table.

### 14.3.4. MyISAM Table Problems

The file format that MySQL uses to store data has been extensively tested, but there are always circumstances that may cause database tables to become corrupted. The following discussion describes how this can happen and how to handle it.

#### 14.3.4.1. Corrupted MyISAM Tables

Even though the MyISAM table format is very reliable (all changes to a table made by an SQL statement are written before the statement returns), you can still get corrupted tables if any of the following events occur:

- The `mysqld` process is killed in the middle of a write.
- An unexpected computer shutdown occurs (for example, the computer is turned off).
- Hardware failures.
- You are using an external program (such as `myisamchk`) to modify a table that is being modified by the server at the same time.
- A software bug in the MySQL or MyISAM code.

Typical symptoms of a corrupt table are:

- You get the following error while selecting data from the table:

```
Incorrect key file for table: '...'. Try to repair it
```

- Queries don't find rows in the table or return incomplete results.

You can check the health of a MyISAM table using the `CHECK TABLE` statement, and repair a corrupted MyISAM table with `REPAIR TABLE`. When `mysqld` is not running, you can also check or repair a table with the `myisamchk` command. See [Section 13.7.2.2, “CHECK TABLE Syntax”](#), [Section 13.7.2.5, “REPAIR TABLE Syntax”](#), and [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

If your tables become corrupted frequently, you should try to determine why this is happening. The most important thing to know is whether the table became corrupted as a result of a server crash. You can verify this easily by looking for a recent `restarted mysqld` message in the error log. If there is such a message, it is likely that table corruption is a result of the server dying. Otherwise, corruption may have occurred during normal operation. This is a bug. You should try to create a reproducible test case that demonstrates the problem. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#), and [MySQL Internals: Porting to Other Systems](#).

#### 14.3.4.2. Problems from Tables Not Being Closed Properly

Each MyISAM index file (`.MYI` file) has a counter in the header that can be used to check whether a table has been closed properly. If you get the following warning from `CHECK TABLE` or `myisamchk`, it means that this counter has gone out of sync:

```
clients are using or haven't closed the table properly
```

This warning doesn't necessarily mean that the table is corrupted, but you should at least check the table.

The counter works as follows:

- The first time a table is updated in MySQL, a counter in the header of the index files is incremented.
- The counter is not changed during further updates.
- When the last instance of a table is closed (because a `FLUSH TABLES` operation was performed or because there is no room in the table cache), the counter is decremented if the table has been updated at any point.
- When you repair the table or check the table and it is found to be okay, the counter is reset to zero.
- To avoid problems with interaction with other processes that might check the table, the counter is not decremented on close if it was zero.

In other words, the counter can become incorrect only under these conditions:

- A `MyISAM` table is copied without first issuing `LOCK TABLES` and `FLUSH TABLES`.
- MySQL has crashed between an update and the final close. (Note that the table may still be okay, because MySQL always issues writes for everything between each statement.)
- A table was modified by `myisamchk --recover` or `myisamchk --update-state` at the same time that it was in use by `mysqld`.
- Multiple `mysqld` servers are using the table and one server performed a `REPAIR TABLE` or `CHECK TABLE` on the table while it was in use by another server. In this setup, it is safe to use `CHECK TABLE`, although you might get the warning from other servers. However, `REPAIR TABLE` should be avoided because when one server replaces the data file with a new one, this is not known to the other servers.

In general, it is a bad idea to share a data directory among multiple servers. See [Section 5.6, “Running Multiple MySQL Instances on One Machine”](#), for additional discussion.

## 14.4. The MEMORY Storage Engine

The `MEMORY` storage engine (formerly known as `HEAP`) creates special-purpose tables with contents that are stored in memory. Because the data is vulnerable to crashes, hardware issues, or power outages, only use these tables as temporary work areas or read-only caches for data pulled from other tables.

**Table 14.9. MEMORY Storage Engine Features**

|                                                    |     |                                        |     |                                              |       |
|----------------------------------------------------|-----|----------------------------------------|-----|----------------------------------------------|-------|
| <b>Storage limits</b>                              | RAM | <b>Transactions</b>                    | No  | <b>Locking granularity</b>                   | Table |
| <b>MVCC</b>                                        | No  | <b>Geospatial data type support</b>    | No  | <b>Geospatial indexing support</b>           | No    |
| <b>B-tree indexes</b>                              | Yes | <b>T-tree indexes</b>                  | No  | <b>Hash indexes</b>                          | Yes   |
| <b>Full-text search indexes</b>                    | No  | <b>Clustered indexes</b>               | No  | <b>Data caches</b>                           | N/A   |
| <b>Index caches</b>                                | N/A | <b>Compressed data</b>                 | No  | <b>Encrypted data<sup>a</sup></b>            | Yes   |
| <b>Cluster database support</b>                    | No  | <b>Replication support<sup>b</sup></b> | Yes | <b>Foreign key support</b>                   | No    |
| <b>Backup / point-in-time recovery<sup>c</sup></b> | Yes | <b>Query cache support</b>             | Yes | <b>Update statistics for data dictionary</b> | Yes   |

<sup>a</sup>Implemented in the server (via encryption functions), rather than in the storage engine.

<sup>b</sup>Implemented in the server, rather than in the storage engine.

<sup>c</sup>Implemented in the server, rather than in the storage engine.

**When to Use MEMORY or MySQL Cluster.** Developers looking to deploy applications that use the `MEMORY` storage engine for important, highly available, or frequently updated data should consider



whether MySQL Cluster is a better choice. A typical use case for the `MEMORY` engine involves these characteristics:

- Operations involving transient, non-critical data such as session management or caching. When the MySQL server halts or restarts, the data in `MEMORY` tables is lost.
- In-memory storage for fast access and low latency. Data volume can fit entirely in memory without causing the operating system to swap out virtual memory pages.
- A read-only or read-mostly data access pattern (limited updates).

MySQL Cluster offers the same features as the `MEMORY` engine with higher performance levels, and provides additional features not available with `MEMORY`:

- Row-level locking and multiple-thread operation for low contention between clients.
- Scalability even with statement mixes that include writes.
- Optional disk-backed operation for data durability.
- Shared-nothing architecture and multiple-host operation with no single point of failure, enabling 99.999% availability.
- Automatic data distribution across nodes; application developers need not craft custom sharding or partitioning solutions.
- Support for variable-length data types (including `BLOB` and `TEXT`) not supported by `MEMORY`.

For a white paper with more detailed comparison of the `MEMORY` storage engine and MySQL Cluster, see [Scaling Web Services with MySQL Cluster: An Alternative to the MySQL Memory Storage Engine](#). This white paper includes a performance study of the two technologies and a step-by-step guide describing how existing `MEMORY` users can migrate to MySQL Cluster.

## Performance Characteristics

`MEMORY` performance is constrained by contention resulting from single-thread execution and table lock overhead when processing updates. This limits scalability when load increases, particularly for statement mixes that include writes.

Despite the in-memory processing for `MEMORY` tables, they are not necessarily faster than `InnoDB` tables on a busy server, for general-purpose queries, or under a read/write workload. In particular, the table locking involved with performing updates can slow down concurrent usage of `MEMORY` tables from multiple sessions.

Depending on the kinds of queries performed on a `MEMORY` table, you might create indexes as either the default hash data structure (for looking up single values based on a unique key), or a general-purpose B-tree data structure (for all kinds of queries involving equality, inequality, or range operators such as less than or greater than). The following sections illustrate the syntax for creating both kinds of indexes. A common performance issue is using the default hash indexes in workloads where B-tree indexes are more efficient.

## Physical Characteristics of `MEMORY` Tables

The `MEMORY` storage engine associates each table with one disk file, which stores the table definition (not the data). The file name begins with the table name and has an extension of `.frm`.

`MEMORY` tables have the following characteristics:

- Space for `MEMORY` tables is allocated in small blocks. Tables use 100% dynamic hashing for inserts. No overflow area or extra key space is needed. No extra space is needed for free lists. Deleted rows

are put in a linked list and are reused when you insert new data into the table. MEMORY tables also have none of the problems commonly associated with deletes plus inserts in hashed tables.

- MEMORY tables use a fixed-length row-storage format. Variable-length types such as VARCHAR are stored using a fixed length.
- MEMORY tables cannot contain BLOB or TEXT columns.
- MEMORY includes support for AUTO\_INCREMENT columns.
- Non-TEMPORARY MEMORY tables are shared among all clients, just like any other non-TEMPORARY table.

## DDL Operations for MEMORY Tables

To create a MEMORY table, specify the clause ENGINE=MEMORY on the CREATE TABLE statement.

```
CREATE TABLE t (i INT) ENGINE = MEMORY;
```

As indicated by the engine name, MEMORY tables are stored in memory. They use hash indexes by default, which makes them very fast for single-value lookups, and very useful for creating temporary tables. However, when the server shuts down, all rows stored in MEMORY tables are lost. The tables themselves continue to exist because their definitions are stored in .frm files on disk, but they are empty when the server restarts.

This example shows how you might create, use, and remove a MEMORY table:

```
mysql> CREATE TABLE test ENGINE=MEMORY
->     SELECT ip,SUM(downloads) AS down
->     FROM log_table GROUP BY ip;
mysql> SELECT COUNT(ip),AVG(down) FROM test;
mysql> DROP TABLE test;
```

The maximum size of MEMORY tables is limited by the `max_heap_table_size` [504] system variable, which has a default value of 16MB. To enforce different size limits for MEMORY tables, change the value of this variable. The value in effect for CREATE TABLE, or a subsequent ALTER TABLE or TRUNCATE TABLE, is the value used for the life of the table. A server restart also sets the maximum size of existing MEMORY tables to the global `max_heap_table_size` [504] value. You can set the size for individual tables as described later in this section.

## Indexes

The MEMORY storage engine supports both HASH and BTREE indexes. You can specify one or the other for a given index by adding a USING clause as shown here:

```
CREATE TABLE lookup
(id INT, INDEX USING HASH (id))
ENGINE = MEMORY;
CREATE TABLE lookup
(id INT, INDEX USING BTREE (id))
ENGINE = MEMORY;
```

For general characteristics of B-tree and hash indexes, see [Section 8.3.1, “How MySQL Uses Indexes”](#).

MEMORY tables can have up to 64 indexes per table, 16 columns per index and a maximum key length of 3072 bytes.

If a MEMORY table hash index has a high degree of key duplication (many index entries containing the same value), updates to the table that affect key values and all deletes are significantly slower. The degree of this slowdown is proportional to the degree of duplication (or, inversely proportional to the index cardinality). You can use a BTREE index to avoid this problem.

`MEMORY` tables can have nonunique keys. (This is an uncommon feature for implementations of hash indexes.)

Columns that are indexed can contain `NULL` values.

## User-Created and Temporary Tables

`MEMORY` table contents are stored in memory, which is a property that `MEMORY` tables share with internal temporary tables that the server creates on the fly while processing queries. However, the two types of tables differ in that `MEMORY` tables are not subject to storage conversion, whereas internal temporary tables are:

- If an internal temporary table becomes too large, the server automatically converts it to on-disk storage, as described in [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#).
- User-created `MEMORY` tables are never converted to disk tables.

## Loading Data

To populate a `MEMORY` table when the MySQL server starts, you can use the `--init-file` [\[424\]](#) option. For example, you can put statements such as `INSERT INTO ... SELECT` or `LOAD DATA INFILE` into this file to load the table from a persistent data source. See [Section 5.1.3, “Server Command Options”](#), and [Section 13.2.6, “LOAD DATA INFILE Syntax”](#).

For loading data into `MEMORY` tables accessed by other sessions concurrently, `MEMORY` supports `INSERT DELAYED`. See [Section 13.2.5.2, “INSERT DELAYED Syntax”](#).

## `MEMORY` Tables and Replication

A server's `MEMORY` tables become empty when it is shut down and restarted. If the server is a replication master, its slaves are not aware that these tables have become empty, so you see out-of-date content if you select data from the tables on the slaves. To synchronize master and slave `MEMORY` tables, when a `MEMORY` table is used on a master for the first time since it was started, a `DELETE` statement is written to the master's binary log, to empty the table on the slaves also. The slave still has outdated data in the table during the interval between the master's restart and its first use of the table. To avoid this interval when a direct query to the slave could return stale data, use the `--init-file` [\[424\]](#) option to populate the `MEMORY` table on the master at startup.

## Managing Memory Use

The server needs sufficient memory to maintain all `MEMORY` tables that are in use at the same time.

Memory is not reclaimed if you delete individual rows from a `MEMORY` table. Memory is reclaimed only when the entire table is deleted. Memory that was previously used for deleted rows is re-used for new rows within the same table. To free all the memory used by a `MEMORY` table when you no longer require its contents, execute `DELETE` or `TRUNCATE TABLE` to remove all rows, or remove the table altogether using `DROP TABLE`. To free up the memory used by deleted rows, use `ALTER TABLE ENGINE=MEMORY` to force a table rebuild.

The memory needed for one row in a `MEMORY` table is calculated using the following expression:

```
SUM_OVER_ALL_BTREE_KEYS(max_length_of_key + sizeof(char*) * 4)
+ SUM_OVER_ALL_HASH_KEYS(sizeof(char*) * 2)
+ ALIGN(length_of_row+1, sizeof(char*))
```

`ALIGN()` represents a round-up factor to cause the row length to be an exact multiple of the `char` pointer size. `sizeof(char*)` is 4 on 32-bit machines and 8 on 64-bit machines.

As mentioned earlier, the `max_heap_table_size` [\[504\]](#) system variable sets the limit on the maximum size of `MEMORY` tables. To control the maximum size for individual tables,

set the session value of this variable before creating each table. (Do not change the global `max_heap_table_size` [504] value unless you intend the value to be used for `MEMORY` tables created by all clients.) The following example creates two `MEMORY` tables, with a maximum size of 1MB and 2MB, respectively:

```
mysql> SET max_heap_table_size = 1024*1024;
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t1 (id INT, UNIQUE(id)) ENGINE = MEMORY;
Query OK, 0 rows affected (0.01 sec)

mysql> SET max_heap_table_size = 1024*1024*2;
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE t2 (id INT, UNIQUE(id)) ENGINE = MEMORY;
Query OK, 0 rows affected (0.00 sec)
```

Both tables revert to the server's global `max_heap_table_size` [504] value if the server restarts.

You can also specify a `MAX_ROWS` table option in `CREATE TABLE` statements for `MEMORY` tables to provide a hint about the number of rows you plan to store in them. This does not enable the table to grow beyond the `max_heap_table_size` [504] value, which still acts as a constraint on maximum table size. For maximum flexibility in being able to use `MAX_ROWS`, set `max_heap_table_size` [504] at least as high as the value to which you want each `MEMORY` table to be able to grow.

## Additional Resources

A forum dedicated to the `MEMORY` storage engine is available at <http://forums.mysql.com/list.php?92>.

## 14.5. The CSV Storage Engine

The `CSV` storage engine stores data in text files using comma-separated values format.

The `CSV` storage engine is always compiled into the MySQL server.

To examine the source for the `CSV` engine, look in the `storage/csv` directory of a MySQL source distribution.

When you create a `CSV` table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. The storage engine also creates a data file. Its name begins with the table name and has a `.CSV` extension. The data file is a plain text file. When you store data into the table, the storage engine saves it into the data file in comma-separated values format.

```
mysql> CREATE TABLE test (i INT NOT NULL, c CHAR(10) NOT NULL)
-> ENGINE = CSV;
Query OK, 0 rows affected (0.12 sec)

mysql> INSERT INTO test VALUES(1,'record one'),(2,'record two');
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM test;
+-----+-----+
| i    | c          |
+-----+-----+
| 1    | record one |
| 2    | record two |
+-----+-----+
2 rows in set (0.00 sec)
```

Creating a `CSV` table also creates a corresponding Metafile that stores the state of the table and the number of rows that exist in the table. The name of this file is the same as the name of the table with the extension `CSM`.

If you examine the `test.CSV` file in the database directory created by executing the preceding statements, its contents should look like this:

```
"1","record one"
"2","record two"
```

This format can be read, and even written, by spreadsheet applications such as Microsoft Excel or StarOffice Calc.

### 14.5.1. Repairing and Checking CSV Tables

The CSV storage engines supports the `CHECK` and `REPAIR` statements to verify and if possible repair a damaged CSV table.

When running the `CHECK` statement, the CSV file will be checked for validity by looking for the correct field separators, escaped fields (matching or missing quotation marks), the correct number of fields compared to the table definition and the existence of a corresponding CSV metafile. The first invalid row discovered will report an error. Checking a valid table produces output like that shown below:

```
mysql> check table csvtest;
+-----+-----+-----+-----+
| Table      | Op    | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.csvtest | check | status   | OK       |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

A check on a corrupted table returns a fault:

```
mysql> check table csvtest;
+-----+-----+-----+-----+
| Table      | Op    | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.csvtest | check | error    | Corrupt  |
+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

If the check fails, the table is marked as crashed (corrupt). Once a table has been marked as corrupt, it is automatically repaired when you next run `CHECK` or execute a `SELECT` statement. The corresponding corrupt status and new status will be displayed when running `CHECK`:

```
mysql> check table csvtest;
+-----+-----+-----+-----+
| Table      | Op    | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.csvtest | check | warning  | Table is marked as crashed |
| test.csvtest | check | status   | OK       |
+-----+-----+-----+-----+
2 rows in set (0.08 sec)
```

To repair a table you can use `REPAIR`, this copies as many valid rows from the existing CSV data as possible, and then replaces the existing CSV file with the recovered rows. Any rows beyond the corrupted data are lost.

```
mysql> repair table csvtest;
+-----+-----+-----+-----+
| Table      | Op    | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.csvtest | repair | status   | OK       |
+-----+-----+-----+-----+
1 row in set (0.02 sec)
```



#### Warning

Note that during repair, only the rows from the CSV file up to the first damaged row are copied to the new table. All other rows from the first damaged row to the end of the table are removed, even valid rows.

## 14.5.2. CSV Limitations

The [CSV](#) storage engine does not support indexing.

Partitioning is not supported for tables using the [CSV](#) storage engine.

All tables that you create using the [CSV](#) storage engine must have the [NOT NULL](#) attribute on all columns. However, for backward compatibility, you can continue to use tables with nullable columns that were created in previous MySQL releases. (Bug #32050)

## 14.6. The [ARCHIVE](#) Storage Engine

The [ARCHIVE](#) storage engine produces special-purpose tables that store large amounts of unindexed data in a very small footprint.

**Table 14.10. [ARCHIVE](#) Storage Engine Features**

|                                                    |      |                                        |     |                                              |       |
|----------------------------------------------------|------|----------------------------------------|-----|----------------------------------------------|-------|
| <b>Storage limits</b>                              | None | <b>Transactions</b>                    | No  | <b>Locking granularity</b>                   | Table |
| <b>MVCC</b>                                        | No   | <b>Geospatial data type support</b>    | Yes | <b>Geospatial indexing support</b>           | No    |
| <b>B-tree indexes</b>                              | No   | <b>T-tree indexes</b>                  | No  | <b>Hash indexes</b>                          | No    |
| <b>Full-text search indexes</b>                    | No   | <b>Clustered indexes</b>               | No  | <b>Data caches</b>                           | No    |
| <b>Index caches</b>                                | No   | <b>Compressed data</b>                 | Yes | <b>Encrypted data<sup>a</sup></b>            | Yes   |
| <b>Cluster database support</b>                    | No   | <b>Replication support<sup>b</sup></b> | Yes | <b>Foreign key support</b>                   | No    |
| <b>Backup / point-in-time recovery<sup>c</sup></b> | Yes  | <b>Query cache support</b>             | Yes | <b>Update statistics for data dictionary</b> | Yes   |

<sup>a</sup>Implemented in the server (via encryption functions), rather than in the storage engine.

<sup>b</sup>Implemented in the server, rather than in the storage engine.

<sup>c</sup>Implemented in the server, rather than in the storage engine.

The [ARCHIVE](#) storage engine is included in MySQL binary distributions. To enable this storage engine if you build MySQL from source, invoke `CMake` with the `-DWITH_ARCHIVE_STORAGE_ENGINE` option.

To examine the source for the [ARCHIVE](#) engine, look in the `storage/archive` directory of a MySQL source distribution.

You can check whether the [ARCHIVE](#) storage engine is available with the `SHOW ENGINES` statement.

When you create an [ARCHIVE](#) table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. The storage engine creates other files, all having names beginning with the table name. The data file has an extension of `.ARZ`. An `.ARN` file may appear during optimization operations.

The [ARCHIVE](#) engine supports `INSERT` and `SELECT`, but not `DELETE`, `REPLACE`, or `UPDATE`. It does support `ORDER BY` operations, `BLOB` columns, and basically all but spatial data types (see [Section 12.18.4.1, “MySQL Spatial Data Types”](#)). The [ARCHIVE](#) engine uses row-level locking.

The [ARCHIVE](#) engine supports the `AUTO_INCREMENT` column attribute. The `AUTO_INCREMENT` column can have either a unique or nonunique index. Attempting to create an index on any other column results in an error. The [ARCHIVE](#) engine also supports the `AUTO_INCREMENT` table option in `CREATE TABLE` and `ALTER TABLE` statements to specify the initial sequence value for a new table or reset the sequence value for an existing table, respectively.

The [ARCHIVE](#) engine ignores `BLOB` columns if they are not requested and scans past them while reading.

**Storage:** Rows are compressed as they are inserted. The `ARCHIVE` engine uses `zlib` lossless data compression (see <http://www.zlib.net/>). You can use `OPTIMIZE TABLE` to analyze the table and pack it into a smaller format (for a reason to use `OPTIMIZE TABLE`, see later in this section). The engine also supports `CHECK TABLE`. There are several types of insertions that are used:

- An `INSERT` statement just pushes rows into a compression buffer, and that buffer flushes as necessary. The insertion into the buffer is protected by a lock. A `SELECT` forces a flush to occur, unless the only insertions that have come in were `INSERT DELAYED` (those flush as necessary). See [Section 13.2.5.2, “INSERT DELAYED Syntax”](#).
- A bulk insert is visible only after it completes, unless other inserts occur at the same time, in which case it can be seen partially. A `SELECT` never causes a flush of a bulk insert unless a normal insert occurs while it is loading.

**Retrieval:** On retrieval, rows are uncompressed on demand; there is no row cache. A `SELECT` operation performs a complete table scan: When a `SELECT` occurs, it finds out how many rows are currently available and reads that number of rows. `SELECT` is performed as a consistent read. Note that lots of `SELECT` statements during insertion can deteriorate the compression, unless only bulk or delayed inserts are used. To achieve better compression, you can use `OPTIMIZE TABLE` or `REPAIR TABLE`. The number of rows in `ARCHIVE` tables reported by `SHOW TABLE STATUS` is always accurate. See [Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#), [Section 13.7.2.5, “REPAIR TABLE Syntax”](#), and [Section 13.7.5.37, “SHOW TABLE STATUS Syntax”](#).

## Additional Resources

- A forum dedicated to the `ARCHIVE` storage engine is available at <http://forums.mysql.com/list.php?112>.

## 14.7. The `BLACKHOLE` Storage Engine

The `BLACKHOLE` storage engine acts as a “black hole” that accepts data but throws it away and does not store it. Retrievals always return an empty result:

```
mysql> CREATE TABLE test(i INT, c CHAR(10)) ENGINE = BLACKHOLE;
Query OK, 0 rows affected (0.03 sec)

mysql> INSERT INTO test VALUES(1,'record one'),(2,'record two');
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM test;
Empty set (0.00 sec)
```

To enable the `BLACKHOLE` storage engine if you build MySQL from source, invoke `CMake` with the `-DWITH_BLACKHOLE_STORAGE_ENGINE` option.

To examine the source for the `BLACKHOLE` engine, look in the `sql` directory of a MySQL source distribution.

When you create a `BLACKHOLE` table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. There are no other files associated with the table.


The `BLACKHOLE` storage engine supports all kinds of indexes. That is, you can include index declarations in the table definition.

You can check whether the `BLACKHOLE` storage engine is available with the `SHOW ENGINES` statement.

Inserts into a `BLACKHOLE` table do not store any data, but if the binary log is enabled, the SQL statements are logged (and replicated to slave servers). This can be useful as a repeater or filter mechanism. Suppose that your application requires slave-side filtering rules, but transferring all



binary log data to the slave first results in too much traffic. In such a case, it is possible to set up on the master host a “dummy” slave process whose default storage engine is `BLACKHOLE`, depicted as follows:



The master writes to its binary log. The “dummy” `mysqld` process acts as a slave, applying the desired combination of `replicate-do-*` and `replicate-ignore-*` rules, and writes a new, filtered binary log of its own. (See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).) This filtered log is provided to the slave.

The dummy process does not actually store any data, so there is little processing overhead incurred by running the additional `mysqld` process on the replication master host. This type of setup can be repeated with additional replication slaves.

`INSERT` triggers for `BLACKHOLE` tables work as expected. However, because the `BLACKHOLE` table does not actually store any data, `UPDATE` and `DELETE` triggers are not activated: The `FOR EACH ROW` clause in the trigger definition does not apply because there are no rows.

Other possible uses for the `BLACKHOLE` storage engine include:

- Verification of dump file syntax.
- Measurement of the overhead from binary logging, by comparing performance using `BLACKHOLE` with and without binary logging enabled.
- `BLACKHOLE` is essentially a “no-op” storage engine, so it could be used for finding performance bottlenecks not related to the storage engine itself.

The `BLACKHOLE` engine is transaction-aware, in the sense that committed transactions are written to the binary log and rolled-back transactions are not.

### Blackhole Engine and Auto Increment Columns

The Blackhole engine is a no-op engine. Any operations performed on a table using Blackhole will have no effect. This should be born in mind when considering the behavior of primary key columns that auto increment. The engine will not automatically increment field values, and does not retain auto increment field state. This has important implications in replication.

Consider the following replication scenario where all three of the following conditions apply:

1. On a master server there is a blackhole table with an auto increment field that is a primary key.
2. On a slave the same table exists but using the MyISAM engine.

3. Inserts are performed into the master's table without explicitly setting the auto increment value in the `INSERT` statement itself or through using a `SET INSERT_ID` statement.

In this scenario replication will fail with a duplicate entry error on the primary key column.

In statement based replication, the value of `INSERT_ID` in the context event will always be the same. Replication will therefore fail due to trying insert a row with a duplicate value for a primary key column.

In row based replication, the value that the engine returns for the row always be the same for each insert. This will result in the slave attempting to replay two insert log entries using the same value for the primary key column, and so replication will fail.

### Column Filtering

When using row-based replication, (`binlog_format=ROW` [2010]), a slave where the last columns are missing from a table is supported, as described in the section [Section 16.4.1.8, “Replication with Differing Table Definitions on Master and Slave”](#).

This filtering works on the slave side, that is, the columns are copied to the slave before they are filtered out. There are at least two cases where it is not desirable to copy the columns to the slave:

1. If the data is confidential, so the slave server should not have access to it.
2. If the master has many slaves, filtering before sending to the slaves may reduce network traffic.

Master column filtering can be achieved using the `BLACKHOLE` engine. This is carried out in a way similar to how master table filtering is achieved - by using the `BLACKHOLE` engine and the `--replicate-do-table` [1978] or `--replicate-ignore-table` [1978] option.

The setup for the master is:

```
CREATE TABLE t1 (public_col_1, ..., public_col_N,
                 secret_col_1, ..., secret_col_M) ENGINE=MyISAM;
```

The setup for the trusted slave is:

```
CREATE TABLE t1 (public_col_1, ..., public_col_N) ENGINE=BLACKHOLE;
```

The setup for the untrusted slave is:

```
CREATE TABLE t1 (public_col_1, ..., public_col_N) ENGINE=MyISAM;
```

## 14.8. The `MERGE` Storage Engine

The `MERGE` storage engine, also known as the `MRG_MyISAM` engine, is a collection of identical `MyISAM` tables that can be used as one. “Identical” means that all tables have identical column and index information. You cannot merge `MyISAM` tables in which the columns are listed in a different order, do not have exactly the same columns, or have the indexes in different order. However, any or all of the `MyISAM` tables can be compressed with `myisampack`. See [Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#). Differences in table options such as `AVG_ROW_LENGTH`, `MAX_ROWS`, or `PACK_KEYS` do not matter.

An alternative to a `MERGE` table is a partitioned table, which stores partitions of a single table in separate files. Partitioning enables some operations to be performed more efficiently and is not limited to the `MyISAM` storage engine. For more information, see [Chapter 17, Partitioning](#).

When you create a `MERGE` table, MySQL creates two files on disk. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table format, and an `.MRG` file contains the names of the underlying `MyISAM` tables that should be used as one. The tables do not have to be in the same database as the `MERGE` table.

You can use `SELECT`, `DELETE`, `UPDATE`, and `INSERT` on `MERGE` tables. You must have [SELECT](#) [755], [DELETE](#) [754], and [UPDATE](#) [756] privileges on the `MyISAM` tables that you map to a `MERGE` table.

**Note**

The use of MERGE tables entails the following security issue: If a user has access to MyISAM table *t*, that user can create a MERGE table *m* that accesses *t*. However, if the user's privileges on *t* are subsequently revoked, the user can continue to access *t* by doing so through *m*.

Use of DROP TABLE with a MERGE table drops only the MERGE specification. The underlying tables are not affected.

To create a MERGE table, you must specify a UNION=(*list-of-tables*) option that indicates which MyISAM tables to use. You can optionally specify an INSERT\_METHOD option to control how inserts into the MERGE table take place. Use a value of FIRST or LAST to cause inserts to be made in the first or last underlying table, respectively. If you specify no INSERT\_METHOD option or if you specify it with a value of NO, inserts into the MERGE table are not permitted and attempts to do so result in an error.

The following example shows how to create a MERGE table:

```
mysql> CREATE TABLE t1 (
->   a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   message CHAR(20)) ENGINE=MyISAM;
mysql> CREATE TABLE t2 (
->   a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   message CHAR(20)) ENGINE=MyISAM;
mysql> INSERT INTO t1 (message) VALUES ('Testing'),('table'),('t1');
mysql> INSERT INTO t2 (message) VALUES ('Testing'),('table'),('t2');
mysql> CREATE TABLE total (
->   a INT NOT NULL AUTO_INCREMENT,
->   message CHAR(20), INDEX(a))
->   ENGINE=MERGE UNION=(t1,t2) INSERT_METHOD=LAST;
```

Note that column *a* is indexed as a PRIMARY KEY in the underlying MyISAM tables, but not in the MERGE table. There it is indexed but not as a PRIMARY KEY because a MERGE table cannot enforce uniqueness over the set of underlying tables. (Similarly, a column with a UNIQUE index in the underlying tables should be indexed in the MERGE table but not as a UNIQUE index.)

After creating the MERGE table, you can use it to issue queries that operate on the group of tables as a whole:

```
mysql> SELECT * FROM total;
+----+-----+
| a | message |
+----+-----+
1	Testing
2	table
3	t1
1	Testing
2	table
3	t2
+----+-----+
```

To remap a MERGE table to a different collection of MyISAM tables, you can use one of the following methods:

- DROP the MERGE table and re-create it.
- Use ALTER TABLE *tbl\_name* UNION=(...) to change the list of underlying tables.

It is also possible to use ALTER TABLE ... UNION=( ) (that is, with an empty UNION clause) to remove all of the underlying tables.

The underlying table definitions and indexes must conform closely to the definition of the MERGE table. Conformance is checked when a table that is part of a MERGE table is opened, not when the MERGE table is created. If any table fails the conformance checks, the operation that triggered the opening of the table fails. This means that changes to the definitions of tables within a MERGE may cause a failure when the MERGE table is accessed. The conformance checks applied to each table are:

- The underlying table and the `MERGE` table must have the same number of columns.
- The column order in the underlying table and the `MERGE` table must match.
- Additionally, the specification for each corresponding column in the parent `MERGE` table and the underlying tables are compared and must satisfy these checks:
  - The column type in the underlying table and the `MERGE` table must be equal.
  - The column length in the underlying table and the `MERGE` table must be equal.
  - The column of the underlying table and the `MERGE` table can be `NULL`.
- The underlying table must have at least as many indexes as the `MERGE` table. The underlying table may have more indexes than the `MERGE` table, but cannot have fewer.

**Note**

A known issue exists where indexes on the same columns must be in identical order, in both the `MERGE` table and the underlying `MyISAM` table. See Bug #33653.

Each index must satisfy these checks:

- The index type of the underlying table and the `MERGE` table must be the same.
- The number of index parts (that is, multiple columns within a compound index) in the index definition for the underlying table and the `MERGE` table must be the same.
- For each index part:
  - Index part lengths must be equal.
  - Index part types must be equal.
  - Index part languages must be equal.
  - Check whether index parts can be `NULL`.

If a `MERGE` table cannot be opened or used because of a problem with an underlying table, `CHECK TABLE` displays information about which table caused the problem.

## Additional Resources

- A forum dedicated to the `MERGE` storage engine is available at <http://forums.mysql.com/list.php?93>.

### 14.8.1. `MERGE` Table Advantages and Disadvantages

`MERGE` tables can help you solve the following problems:

- Easily manage a set of log tables. For example, you can put data from different months into separate tables, compress some of them with `myisampack`, and then create a `MERGE` table to use them as one.
- Obtain more speed. You can split a large read-only table based on some criteria, and then put individual tables on different disks. A `MERGE` table structured this way could be much faster than using a single large table.
- Perform more efficient searches. If you know exactly what you are looking for, you can search in just one of the underlying tables for some queries and use a `MERGE` table for others. You can even have many different `MERGE` tables that use overlapping sets of tables.

- Perform more efficient repairs. It is easier to repair individual smaller tables that are mapped to a [MERGE](#) table than to repair a single large table.
- Instantly map many tables as one. A [MERGE](#) table need not maintain an index of its own because it uses the indexes of the individual tables. As a result, [MERGE](#) table collections are very fast to create or remap. (You must still specify the index definitions when you create a [MERGE](#) table, even though no indexes are created.)
- If you have a set of tables from which you create a large table on demand, you can instead create a [MERGE](#) table from them on demand. This is much faster and saves a lot of disk space.
- Exceed the file size limit for the operating system. Each [MyISAM](#) table is bound by this limit, but a collection of [MyISAM](#) tables is not.
- You can create an alias or synonym for a [MyISAM](#) table by defining a [MERGE](#) table that maps to that single table. There should be no really notable performance impact from doing this (only a couple of indirect calls and `memcpy()` calls for each read).

The disadvantages of [MERGE](#) tables are:

- You can use only identical [MyISAM](#) tables for a [MERGE](#) table.
- Some [MyISAM](#) features are unavailable in [MERGE](#) tables. For example, you cannot create [FULLTEXT](#) indexes on [MERGE](#) tables. (You can create [FULLTEXT](#) indexes on the underlying [MyISAM](#) tables, but you cannot search the [MERGE](#) table with a full-text search.)
- If the [MERGE](#) table is nontemporary, all underlying [MyISAM](#) tables must be nontemporary. If the [MERGE](#) table is temporary, the [MyISAM](#) tables can be any mix of temporary and nontemporary.
- [MERGE](#) tables use more file descriptors than [MyISAM](#) tables. If 10 clients are using a [MERGE](#) table that maps to 10 tables, the server uses  $(10 \times 10) + 10$  file descriptors. (10 data file descriptors for each of the 10 clients, and 10 index file descriptors shared among the clients.)
- Index reads are slower. When you read an index, the [MERGE](#) storage engine needs to issue a read on all underlying tables to check which one most closely matches a given index value. To read the next index value, the [MERGE](#) storage engine needs to search the read buffers to find the next value. Only when one index buffer is used up does the storage engine need to read the next index block. This makes [MERGE](#) indexes much slower on [eq\\_ref \[898\]](#) searches, but not much slower on [ref \[898\]](#) searches. For more information about [eq\\_ref \[898\]](#) and [ref \[898\]](#), see [Section 13.8.2, “EXPLAIN Syntax”](#).

## 14.8.2. MERGE Table Problems

The following are known problems with [MERGE](#) tables:

- In versions of MySQL Server prior to 5.1.23, it was possible to create temporary merge tables with nontemporary child [MyISAM](#) tables.

From versions 5.1.23, [MERGE](#) children were locked through the parent table. If the parent was temporary, it was not locked and so the children were not locked either. Parallel use of the [MyISAM](#) tables corrupted them.

- If you use [ALTER TABLE](#) to change a [MERGE](#) table to another storage engine, the mapping to the underlying tables is lost. Instead, the rows from the underlying [MyISAM](#) tables are copied into the altered table, which then uses the specified storage engine.
- The [INSERT\\_METHOD](#) table option for a [MERGE](#) table indicates which underlying [MyISAM](#) table to use for inserts into the [MERGE](#) table. However, use of the [AUTO\\_INCREMENT](#) table option for that [MyISAM](#) table has no effect for inserts into the [MERGE](#) table until at least one row has been inserted directly into the [MyISAM](#) table.

- A `MERGE` table cannot maintain uniqueness constraints over the entire table. When you perform an `INSERT`, the data goes into the first or last `MyISAM` table (as determined by the `INSERT_METHOD` option). MySQL ensures that unique key values remain unique within that `MyISAM` table, but not over all the underlying tables in the collection.
- Because the `MERGE` engine cannot enforce uniqueness over the set of underlying tables, `REPLACE` does not work as expected. The two key facts are:
  - `REPLACE` can detect unique key violations only in the underlying table to which it is going to write (which is determined by the `INSERT_METHOD` option). This differs from violations in the `MERGE` table itself.
  - If `REPLACE` detects a unique key violation, it will change only the corresponding row in the underlying table it is writing to; that is, the first or last table, as determined by the `INSERT_METHOD` option.

Similar considerations apply for `INSERT ... ON DUPLICATE KEY UPDATE`.

- `MERGE` tables do not support partitioning. That is, you cannot partition a `MERGE` table, nor can any of a `MERGE` table's underlying `MyISAM` tables be partitioned.
- You should not use `ANALYZE TABLE`, `REPAIR TABLE`, `OPTIMIZE TABLE`, `ALTER TABLE`, `DROP TABLE`, `DELETE` without a `WHERE` clause, or `TRUNCATE TABLE` on any of the tables that are mapped into an open `MERGE` table. If you do so, the `MERGE` table may still refer to the original table and yield unexpected results. To work around this problem, ensure that no `MERGE` tables remain open by issuing a `FLUSH TABLES` statement prior to performing any of the named operations.

The unexpected results include the possibility that the operation on the `MERGE` table will report table corruption. If this occurs after one of the named operations on the underlying `MyISAM` tables, the corruption message is spurious. To deal with this, issue a `FLUSH TABLES` statement after modifying the `MyISAM` tables.

- `DROP TABLE` on a table that is in use by a `MERGE` table does not work on Windows because the `MERGE` storage engine's table mapping is hidden from the upper layer of MySQL. Windows does not permit open files to be deleted, so you first must flush all `MERGE` tables (with `FLUSH TABLES`) or drop the `MERGE` table before dropping the table.
- The definition of the `MyISAM` tables and the `MERGE` table are checked when the tables are accessed (for example, as part of a `SELECT` or `INSERT` statement). The checks ensure that the definitions of the tables and the parent `MERGE` table definition match by comparing column order, types, sizes and associated indexes. If there is a difference between the tables, an error is returned and the statement fails. Because these checks take place when the tables are opened, any changes to the definition of a single table, including column changes, column ordering, and engine alterations will cause the statement to fail.
- The order of indexes in the `MERGE` table and its underlying tables should be the same. If you use `ALTER TABLE` to add a `UNIQUE` index to a table used in a `MERGE` table, and then use `ALTER TABLE` to add a nonunique index on the `MERGE` table, the index ordering is different for the tables if there was already a nonunique index in the underlying table. (This happens because `ALTER TABLE` puts `UNIQUE` indexes before nonunique indexes to facilitate rapid detection of duplicate keys.) Consequently, queries on tables with such indexes may return unexpected results.
- If you encounter an error message similar to `ERROR 1017 (HY000): Can't find file: 'tbl_name.MRG' (errno: 2)`, it generally indicates that some of the underlying tables do not use the `MyISAM` storage engine. Confirm that all of these tables are `MyISAM`.
- The maximum number of rows in a `MERGE` table is  $2^{64}$  (~1.844E+19; the same as for a `MyISAM` table). It is not possible to merge multiple `MyISAM` tables into a single `MERGE` table that would have more than this number of rows.
- The `MERGE` storage engine does not support `INSERT DELAYED` statements.



- Use of underlying **MyISAM** tables of differing row formats with a parent **MERGE** table is currently known to fail. See Bug #32364.
- You cannot change the union list of a nontemporary **MERGE** table when **LOCK TABLES** is in effect. The following does *not* work:

```
CREATE TABLE m1 ... ENGINE=MRG_MYISAM ...;
LOCK TABLES t1 WRITE, t2 WRITE, m1 WRITE;
ALTER TABLE m1 ... UNION=(t1,t2) ...;
```

However, you can do this with a temporary **MERGE** table.

- You cannot create a **MERGE** table with **CREATE ... SELECT**, neither as a temporary **MERGE** table, nor as a nontemporary **MERGE** table. For example:

```
CREATE TABLE m1 ... ENGINE=MRG_MYISAM ... SELECT ...;
```

Attempts to do this result in an error: *tbl\_name* is not **BASE TABLE**.

- In some cases, differing **PACK\_KEYS** table option values among the **MERGE** and underlying tables cause unexpected results if the underlying tables contain **CHAR** or **BINARY** columns. As a workaround, use **ALTER TABLE** to ensure that all involved tables have the same **PACK\_KEYS** value. (Bug #50646)

## 14.9. The **FEDERATED** Storage Engine

The **FEDERATED** storage engine lets you access data from a remote MySQL database without using replication or cluster technology. Querying a local **FEDERATED** table automatically pulls the data from the remote (federated) tables. No data is stored on the local tables.

To include the **FEDERATED** storage engine if you build MySQL from source, invoke **CMake** with the **-DWITH\_FEDERATED\_STORAGE\_ENGINE** option.

The **FEDERATED** storage engine is not enabled by default in the running server; to enable **FEDERATED**, you must start the MySQL server binary using the **--federated** option.

To examine the source for the **FEDERATED** engine, look in the **storage/federated** directory of a MySQL source distribution.

### 14.9.1. **FEDERATED** Storage Engine Overview

When you create a table using one of the standard storage engines (such as **MyISAM**, **CSV** or **InnoDB**), the table consists of the table definition and the associated data. When you create a **FEDERATED** table, the table definition is the same, but the physical storage of the data is handled on a remote server.

A **FEDERATED** table consists of two elements:

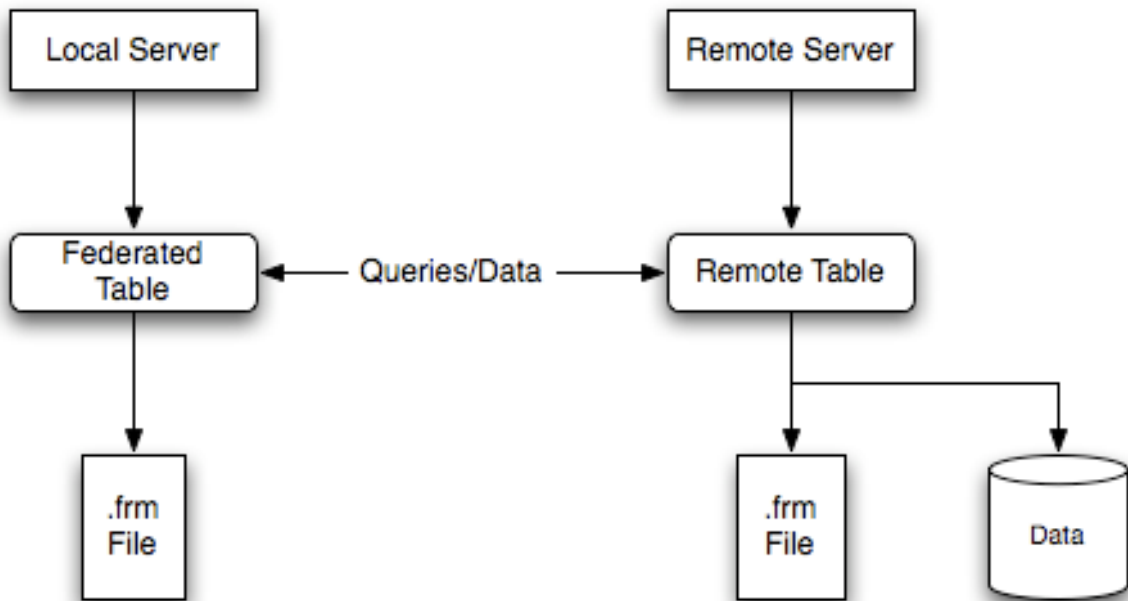
- A *remote server* with a database table, which in turn consists of the table definition (stored in the **.frm** file) and the associated table. The table type of the remote table may be any type supported by the remote **mysqld** server, including **MyISAM** or **InnoDB**.
- A *local server* with a database table, where the table definition matches that of the corresponding table on the remote server. The table definition is stored within the **.frm** file. However, there is no data file on the local server. Instead, the table definition includes a connection string that points to the remote table.

When executing queries and statements on a **FEDERATED** table on the local server, the operations that would normally insert, update or delete information from a local data file are instead sent to the remote server for execution, where they update the data file on the remote server or return matching rows from the remote server.



The basic structure of a `FEDERATED` table setup is shown in Figure 14.1, “`FEDERATED` Table Structure”.

**Figure 14.1. `FEDERATED` Table Structure**



When a client issues an SQL statement that refers to a `FEDERATED` table, the flow of information between the local server (where the SQL statement is executed) and the remote server (where the data is physically stored) is as follows:

1. The storage engine looks through each column that the `FEDERATED` table has and constructs an appropriate SQL statement that refers to the remote table.
2. The statement is sent to the remote server using the MySQL client API.
3. The remote server processes the statement and the local server retrieves any result that the statement produces (an affected-rows count or a result set).
4. If the statement produces a result set, each column is converted to internal storage engine format that the `FEDERATED` engine expects and can use to display the result to the client that issued the original statement.

The local server communicates with the remote server using MySQL client C API functions. It invokes `mysql_real_query()` to send the statement. To read a result set, it uses `mysql_store_result()` and fetches rows one at a time using `mysql_fetch_row()`.

## 14.9.2. How to Create `FEDERATED` Tables

To create a `FEDERATED` table you should follow these steps:

1. Create the table on the remote server. Alternatively, make a note of the table definition of an existing table, perhaps using the `SHOW CREATE TABLE` statement.
2. Create the table on the local server with an identical table definition, but adding the connection information that links the local table to the remote table.

For example, you could create the following table on the remote server:

```
CREATE TABLE test_table (
  id      INT(20) NOT NULL AUTO_INCREMENT,
```

```

name  VARCHAR(32) NOT NULL DEFAULT '',
other INT(20) NOT NULL DEFAULT '0',
PRIMARY KEY (id),
INDEX name (name),
INDEX other_key (other)
)
ENGINE=MyISAM
DEFAULT CHARSET=latin1;

```

To create the local table that will be federated to the remote table, there are two options available. You can either create the local table and specify the connection string (containing the server name, login, password) to be used to connect to the remote table using the **CONNECTION**, or you can use an existing connection that you have previously created using the **CREATE SERVER** statement.



**Important**

When you create the local table it *must* have an identical field definition to the remote table.



**Note**

You can improve the performance of a **FEDERATED** table by adding indexes to the table on the host. The optimization will occur because the query sent to the remote server will include the contents of the **WHERE** clause and will be sent to the remote server and subsequently executed locally. This reduces the network traffic that would otherwise request the entire table from the server for local processing.

**14.9.2.1. Creating a **FEDERATED** Table Using **CONNECTION****

To use the first method, you must specify the **CONNECTION** string after the engine type in a **CREATE TABLE** statement. For example:

```

CREATE TABLE federated_table (
  id      INT(20) NOT NULL AUTO_INCREMENT,
  name   VARCHAR(32) NOT NULL DEFAULT '',
  other  INT(20) NOT NULL DEFAULT '0',
  PRIMARY KEY (id),
  INDEX name (name),
  INDEX other_key (other)
)
ENGINE=FEDERATED
DEFAULT CHARSET=latin1
CONNECTION='mysql://fed_user@remote_host:9306/federated/test_table';

```



**Note**

**CONNECTION** replaces the **COMMENT** used in some previous versions of MySQL.

The **CONNECTION** string contains the information required to connect to the remote server containing the table that will be used to physically store the data. The connection string specifies the server name, login credentials, port number and database/table information. In the example, the remote table is on the server `remote_host`, using port 9306. The name and port number should match the host name (or IP address) and port number of the remote MySQL server instance you want to use as your remote table.

The format of the connection string is as follows:

```

scheme://user_name[:password]@host_name[:port_num]/db_name/tbl_name

```

Where:

- *scheme*: A recognized connection protocol. Only `mysql` is supported as the *scheme* value at this point.

- *user\_name*: The user name for the connection. This user must have been created on the remote server, and must have suitable privileges to perform the required actions (**SELECT**, **INSERT**, **UPDATE**, and so forth) on the remote table.
- *password*: (Optional) The corresponding password for *user\_name*.
- *host\_name*: The host name or IP address of the remote server.
- *port\_num*: (Optional) The port number for the remote server. The default is 3306.
- *db\_name*: The name of the database holding the remote table.
- *tbl\_name*: The name of the remote table. The name of the local and the remote table do not have to match.

Sample connection strings:

```
CONNECTION='mysql://username:password@hostname:port/database/tablename'
CONNECTION='mysql://username@hostname/database/tablename'
CONNECTION='mysql://username:password@hostname/database/tablename'
```

### 14.9.2.2. Creating a **FEDERATED** Table Using **CREATE SERVER**

If you are creating a number of **FEDERATED** tables on the same server, or if you want to simplify the process of creating **FEDERATED** tables, you can use the **CREATE SERVER** statement to define the server connection parameters, just as you would with the **CONNECTION** string.

The format of the **CREATE SERVER** statement is:

```
CREATE SERVER server_name
FOREIGN DATA WRAPPER wrapper_name
OPTIONS (option [, option] ...)
```

The *server\_name* is used in the connection string when creating a new **FEDERATED** table.

For example, to create a server connection identical to the **CONNECTION** string:

```
CONNECTION='mysql://fed_user@remote_host:9306/federated/test_table';
```

You would use the following statement:

```
CREATE SERVER fedlink
FOREIGN DATA WRAPPER mysql
OPTIONS (USER 'fed_user', HOST 'remote_host', PORT 9306, DATABASE 'federated');
```

To create a **FEDERATED** table that uses this connection, you still use the **CONNECTION** keyword, but specify the name you used in the **CREATE SERVER** statement.

```
CREATE TABLE test_table (
  id      INT(20) NOT NULL AUTO_INCREMENT,
  name    VARCHAR(32) NOT NULL DEFAULT '',
  other   INT(20) NOT NULL DEFAULT '0',
  PRIMARY KEY (id),
  INDEX name (name),
  INDEX other_key (other)
)
ENGINE=FEDERATED
DEFAULT CHARSET=latin1
CONNECTION='fedlink/test_table';
```

The connection name in this example contains the name of the connection (*fedlink*) and the name of the table (*test\_table*) to link to, separated by a slash. If you specify only the connection name without a table name, the table name of the local table is used instead.

For more information on **CREATE SERVER**, see [Section 13.1.13, “CREATE SERVER Syntax”](#).

The `CREATE SERVER` statement accepts the same arguments as the `CONNECTION` string. The `CREATE SERVER` statement updates the rows in the `mysql.servers` table. See the following table for information on the correspondence between parameters in a connection string, options in the `CREATE SERVER` statement, and the columns in the `mysql.servers` table. For reference, the format of the `CONNECTION` string is as follows:

```
scheme://user_name[:password]@host_name[:port_num]/db_name/tbl_name
```

| Description       | CONNECTION string | CREATE SERVER option      | mysql.servers column  |
|-------------------|-------------------|---------------------------|-----------------------|
| Connection scheme | <i>scheme</i>     | <code>wrapper_name</code> | <code>Wrapper</code>  |
| Remote user       | <i>user_name</i>  | <code>USER</code>         | <code>Username</code> |
| Remote password   | <i>password</i>   | <code>PASSWORD</code>     | <code>Password</code> |
| Remote host       | <i>host_name</i>  | <code>HOST</code>         | <code>Host</code>     |
| Remote port       | <i>port_num</i>   | <code>PORT</code>         | <code>Port</code>     |
| Remote database   | <i>db_name</i>    | <code>DATABASE</code>     | <code>Db</code>       |

### 14.9.3. FEDERATED Storage Engine Notes and Tips

You should be aware of the following points when using the `FEDERATED` storage engine:

- `FEDERATED` tables may be replicated to other slaves, but you must ensure that the slave servers are able to use the user/password combination that is defined in the `CONNECTION` string (or the row in the `mysql.servers` table) to connect to the remote server.

The following items indicate features that the `FEDERATED` storage engine does and does not support:

- The remote server must be a MySQL server.
- The remote table that a `FEDERATED` table points to *must* exist before you try to access the table through the `FEDERATED` table.
- It is possible for one `FEDERATED` table to point to another, but you must be careful not to create a loop.
- A `FEDERATED` table does not support indexes per se. Because access to the table is handled remotely, it is the remote table that supports the indexes. Care should be taken when creating a `FEDERATED` table since the index definition from an equivalent `MyISAM` or other table may not be supported. For example, creating a `FEDERATED` table with an index prefix on `VARCHAR`, `TEXT` or `BLOB` columns will fail. The following definition in `MyISAM` is valid:

```
CREATE TABLE `T1` (`A` VARCHAR(100), UNIQUE KEY(`A` (30))) ENGINE=MYISAM;
```

The key prefix in this example is incompatible with the `FEDERATED` engine, and the equivalent statement will fail:

```
CREATE TABLE `T1` (`A` VARCHAR(100), UNIQUE KEY(`A` (30))) ENGINE=FEDERATED
CONNECTION='MYSQL://127.0.0.1:3306/TEST/T1';
```

If possible, you should try to separate the column and index definition when creating tables on both the remote server and the local server to avoid these index issues.

- Internally, the implementation uses `SELECT`, `INSERT`, `UPDATE`, and `DELETE`, but not `HANDLER`.
- The `FEDERATED` storage engine supports `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `TRUNCATE TABLE`, and indexes. It does not support `ALTER TABLE`, or any Data Definition Language statements that directly affect the structure of the table, other than `DROP TABLE`. The current implementation does not use prepared statements.

- **FEDERATED** accepts `INSERT ... ON DUPLICATE KEY UPDATE` statements, but if a duplicate-key violation occurs, the statement fails with an error.
- Performance on a **FEDERATED** table when performing bulk inserts (for example, on a `INSERT INTO ... SELECT ...` statement) is slower than with other table types because each selected row is treated as an individual `INSERT` statement on the **FEDERATED** table.
- Transactions are not supported.
- **FEDERATED** performs bulk-insert handling such that multiple rows are sent to the remote table in a batch. This provides a performance improvement and enables the remote table to perform improvement. Also, if the remote table is transactional, it enables the remote storage engine to perform statement rollback properly should an error occur. This capability has the following limitations:
  - The size of the insert cannot exceed the maximum packet size between servers. If the insert exceeds this size, it is broken into multiple packets and the rollback problem can occur.
  - Bulk-insert handling does not occur for `INSERT ... ON DUPLICATE KEY UPDATE`.
- There is no way for the **FEDERATED** engine to know if the remote table has changed. The reason for this is that this table must work like a data file that would never be written to by anything other than the database system. The integrity of the data in the local table could be breached if there was any change to the remote database.
- When using a **CONNECTION** string, you cannot use an '@' character in the password. You can get round this limitation by using the `CREATE SERVER` statement to create a server connection.
- The `insert_id` [490] and `timestamp` [558] options are not propagated to the data provider.
- Any `DROP TABLE` statement issued against a **FEDERATED** table drops only the local table, not the remote table.
- **FEDERATED** tables do not work with the query cache.
- User-defined partitioning is not supported for **FEDERATED** tables.

#### 14.9.4. **FEDERATED** Storage Engine Resources

The following additional resources are available for the **FEDERATED** storage engine:

- A forum dedicated to the **FEDERATED** storage engine is available at <http://forums.mysql.com/list.php?105>.

### 14.10. The **EXAMPLE** Storage Engine

The **EXAMPLE** storage engine is a stub engine that does nothing. Its purpose is to serve as an example in the MySQL source code that illustrates how to begin writing new storage engines. As such, it is primarily of interest to developers.

To enable the **EXAMPLE** storage engine if you build MySQL from source, invoke `CMake` with the `-DWITH_EXAMPLE_STORAGE_ENGINE` option.

To examine the source for the **EXAMPLE** engine, look in the `storage/example` directory of a MySQL source distribution.

When you create an **EXAMPLE** table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. No other files are created. No data can be stored into the table. Retrievals return an empty result.

```
mysql> CREATE TABLE test (i INT) ENGINE = EXAMPLE;
```

```
Query OK, 0 rows affected (0.78 sec)

mysql> INSERT INTO test VALUES(1),(2),(3);
ERROR 1031 (HY000): Table storage engine for 'test' doesn't »
      have this option

mysql> SELECT * FROM test;
Empty set (0.31 sec)
```

The `EXAMPLE` storage engine does not support indexing.

## 14.11. Other Storage Engines

Other storage engines may be available from third parties and community members that have used the Custom Storage Engine interface.

Third party engines are not supported by MySQL. For further information, documentation, installation guides, bug reporting or for any help or assistance with these engines, please contact the developer of the engine directly.

For more information on developing a customer storage engine that can be used with the Pluggable Storage Engine Architecture, see [MySQL Internals: Writing a Custom Storage Engine](#).

## 14.12. Overview of MySQL Storage Engine Architecture

The MySQL pluggable storage engine architecture enables a database professional to select a specialized storage engine for a particular application need while being completely shielded from the need to manage any specific application coding requirements. The MySQL server architecture isolates the application programmer and DBA from all of the low-level implementation details at the storage level, providing a consistent and easy application model and API. Thus, although there are different capabilities across different storage engines, the application is shielded from these differences.

The pluggable storage engine architecture provides a standard set of management and support services that are common among all underlying storage engines. The storage engines themselves are the components of the database server that actually perform actions on the underlying data that is maintained at the physical server level.

This efficient and modular architecture provides huge benefits for those wishing to specifically target a particular application need—such as data warehousing, transaction processing, or high availability situations—while enjoying the advantage of utilizing a set of interfaces and services that are independent of any one storage engine.

The application programmer and DBA interact with the MySQL database through Connector APIs and service layers that are above the storage engines. If application changes bring about requirements that demand the underlying storage engine change, or that one or more storage engines be added to support new needs, no significant coding or process changes are required to make things work. The MySQL server architecture shields the application from the underlying complexity of the storage engine by presenting a consistent and easy-to-use API that applies across storage engines.

### 14.12.1. Pluggable Storage Engine Architecture

MySQL Server uses a pluggable storage engine architecture that enables storage engines to be loaded into and unloaded from a running MySQL server.

#### Plugging in a Storage Engine

Before a storage engine can be used, the storage engine plugin shared library must be loaded into MySQL using the `INSTALL PLUGIN` statement. For example, if the `EXAMPLE` engine plugin is named `example` and the shared library is named `ha_example.so`, you load it with the following statement:

```
mysql> INSTALL PLUGIN example SONAME 'ha_example.so';
```

To install a pluggable storage engine, the plugin file must be located in the MySQL plugin directory, and the user issuing the `INSTALL PLUGIN` statement must have `INSERT [755]` privilege for the `mysql.plugin` table.

The shared library must be located in the MySQL server plugin directory, the location of which is given by the `plugin_dir [525]` system variable.

### Unplugging a Storage Engine

To unplug a storage engine, use the `UNINSTALL PLUGIN` statement:

```
mysql> UNINSTALL PLUGIN example;
```

If you unplug a storage engine that is needed by existing tables, those tables become inaccessible, but will still be present on disk (where applicable). Ensure that there are no tables using a storage engine before you unplug the storage engine.

## 14.12.2. The Common Database Server Layer

A MySQL pluggable storage engine is the component in the MySQL database server that is responsible for performing the actual data I/O operations for a database as well as enabling and enforcing certain feature sets that target a specific application need. A major benefit of using specific storage engines is that you are only delivered the features needed for a particular application, and therefore you have less system overhead in the database, with the end result being more efficient and higher database performance. This is one of the reasons that MySQL has always been known to have such high performance, matching or beating proprietary monolithic databases in industry standard benchmarks.

From a technical perspective, what are some of the unique supporting infrastructure components that are in a storage engine? Some of the key feature differentiations include:

- *Concurrency*: Some applications have more granular lock requirements (such as row-level locks) than others. Choosing the right locking strategy can reduce overhead and therefore improve overall performance. This area also includes support for capabilities such as multi-version concurrency control or “snapshot” read.
- *Transaction Support*: Not every application needs transactions, but for those that do, there are very well defined requirements such as ACID compliance and more.
- *Referential Integrity*: The need to have the server enforce relational database referential integrity through DDL defined foreign keys.
- *Physical Storage*: This involves everything from the overall page size for tables and indexes as well as the format used for storing data to physical disk.
- *Index Support*: Different application scenarios tend to benefit from different index strategies. Each storage engine generally has its own indexing methods, although some (such as B-tree indexes) are common to nearly all engines.
- *Memory Caches*: Different applications respond better to some memory caching strategies than others, so although some memory caches are common to all storage engines (such as those used for user connections or MySQL's high-speed Query Cache), others are uniquely defined only when a particular storage engine is put in play.
- *Performance Aids*: This includes multiple I/O threads for parallel operations, thread concurrency, database checkpointing, bulk insert handling, and more.
- *Miscellaneous Target Features*: This may include support for geospatial operations, security restrictions for certain data manipulation operations, and other similar features.

Each set of the pluggable storage engine infrastructure components are designed to offer a selective set of benefits for a particular application. Conversely, avoiding a set of component features helps



reduce unnecessary overhead. It stands to reason that understanding a particular application's set of requirements and selecting the proper MySQL storage engine can have a dramatic impact on overall system efficiency and performance.



---

# Chapter 15. High Availability and Scalability

## Table of Contents

|                                                                         |      |
|-------------------------------------------------------------------------|------|
| 15.1. Oracle VM Template for MySQL Enterprise Edition .....             | 1822 |
| 15.2. Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux ..... | 1822 |
| 15.3. Overview of MySQL with Windows Failover Clustering .....          | 1825 |
| 15.4. Using MySQL within an Amazon EC2 Instance .....                   | 1827 |
| 15.4.1. Setting Up MySQL on an EC2 AMI .....                            | 1828 |
| 15.4.2. EC2 Instance Limitations .....                                  | 1829 |
| 15.4.3. Deploying a MySQL Database Using EC2 .....                      | 1829 |
| 15.5. Using ZFS Replication .....                                       | 1832 |
| 15.5.1. Using ZFS for File System Replication .....                     | 1833 |
| 15.5.2. Configuring MySQL for ZFS Replication .....                     | 1834 |
| 15.5.3. Handling MySQL Recovery with ZFS .....                          | 1835 |
| 15.6. Using MySQL with <a href="#">memcached</a> .....                  | 1835 |
| 15.6.1. Installing <a href="#">memcached</a> .....                      | 1837 |
| 15.6.2. Using <a href="#">memcached</a> .....                           | 1838 |
| 15.6.3. Developing a <a href="#">memcached</a> Application .....        | 1856 |
| 15.6.4. Getting <a href="#">memcached</a> Statistics .....              | 1880 |
| 15.6.5. <a href="#">memcached</a> FAQ .....                             | 1888 |
| 15.7. MySQL Proxy .....                                                 | 1891 |
| 15.7.1. MySQL Proxy Supported Platforms .....                           | 1892 |
| 15.7.2. Installing MySQL Proxy .....                                    | 1892 |
| 15.7.3. MySQL Proxy Command Options .....                               | 1895 |
| 15.7.4. MySQL Proxy Scripting .....                                     | 1904 |
| 15.7.5. Using MySQL Proxy .....                                         | 1918 |
| 15.7.6. MySQL Proxy FAQ .....                                           | 1923 |

Data is the currency of today's web, mobile, social, enterprise and cloud applications. Ensuring data is always available is a top priority for any organization. Minutes of downtime can result in significant loss of revenue and reputation.

There is no “one size fits all” approach to delivering High Availability (HA). Unique application attributes, business requirements, operational capabilities and legacy infrastructure can all influence HA technology selection. And technology is only one element in delivering HA: people and processes are just as critical as the technology itself.

MySQL is deployed into many applications demanding availability and scalability. **Availability** refers to the ability to cope with, and if necessary recover from, failures on the host, including failures of MySQL, the operating system, or the hardware and maintenance activity that may otherwise cause downtime. **Scalability** refers to the ability to spread both the database and the load of your application queries across multiple MySQL servers.

Because each application has different operational and availability requirements, MySQL offers a range of certified and supported solutions, delivering the appropriate levels of High Availability (HA) and scalability to meet service level requirements. Such solutions extend from replication, through virtualization and geographically redundant, multi-data center solutions delivering 99.999% uptime.

Selecting the right high availability solution for an application largely depends on:

- The level of availability required.
- The type of application being deployed.
- Accepted best practices within your own environment.

The primary solutions supported by MySQL include:

- MySQL Replication. Learn more: [Chapter 16, Replication](#)
- MySQL Cluster. Learn more: [MySQL Cluster NDB 7.2](#)
- Oracle VM Template for MySQL. Learn more: [Section 15.1, “Oracle VM Template for MySQL Enterprise Edition”](#).
- MySQL with DRBD with Corosync and Pacemaker. Learn more: [Section 15.2, “Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux”](#).
- MySQL with Windows Failover Clustering. Learn more: [Section 15.3, “Overview of MySQL with Windows Failover Clustering”](#).
- MySQL with Solaris Cluster. [Learn more about Solaris Cluster](#).

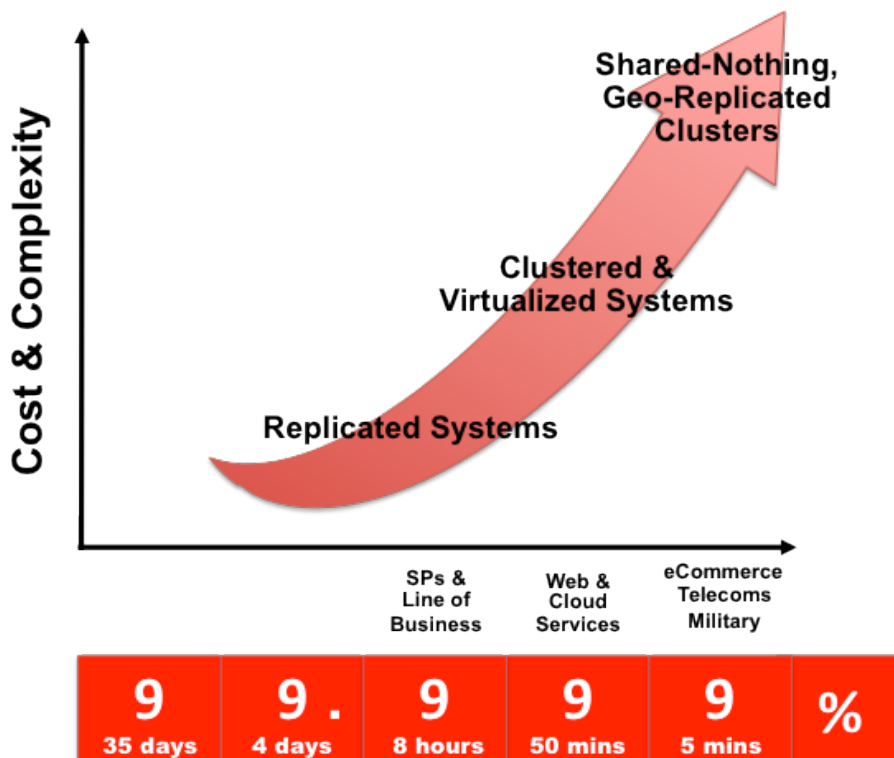
Further options are available using third-party solutions.

Each architecture used to achieve highly available database services is differentiated by the levels of uptime it offers. These architectures can be grouped into three main categories:

- Data Replication.
- Clustered & Virtualized Systems.
- Shared-Nothing, Geographically-Replicated Clusters.

As illustrated in the following figure, each of these architectures offers progressively higher levels of uptime, which must be balanced against potentially greater levels of cost and complexity that each can incur. Simply deploying a high availability architecture is not a guarantee of actually delivering HA. In fact, a poorly implemented and maintained shared-nothing cluster could easily deliver lower levels of availability than a simple data replication solution.

**Figure 15.1. Tradeoffs: Cost and Complexity versus Availability**



The following table compares the HA and Scalability capabilities of the various MySQL solutions:

**Table 15.1. Feature Comparison of MySQL HA Solutions**

| Requirement                                  | MySQL Replication                    | Windows Server | DRBD                                                                 | Oracle VM Template                                                   | MySQL Cluster                          |
|----------------------------------------------|--------------------------------------|----------------|----------------------------------------------------------------------|----------------------------------------------------------------------|----------------------------------------|
| <b>Availability</b>                          |                                      |                |                                                                      |                                                                      |                                        |
| Platform Support                             | All Supported by MySQL Server        | Linux          | Linux                                                                | Oracle Linux                                                         | All Supported by MySQL Cluster         |
| Automated IP Failover                        | No                                   |                | Yes                                                                  | Yes                                                                  | Depends on Connector and Configuration |
| Automated Database Failover                  | No                                   |                | Yes                                                                  | Yes                                                                  | Yes                                    |
| Automatic Data Resynchronization             | No                                   |                | Yes                                                                  | N/A - Shared Storage                                                 | Yes                                    |
| Typical Failover Time                        | User / Script Dependent              |                | Configuration Dependent, 60 seconds and Above                        | Configuration Dependent, 60 seconds and Above                        | 1 Second and Less                      |
| Synchronous Replication                      | No, Asynchronous and Semisynchronous |                | Yes                                                                  | N/A - Shared Storage                                                 | Yes                                    |
| Shared Storage                               | No, Distributed                      |                | No, Distributed                                                      | Yes                                                                  | No, Distributed                        |
| Geographic redundancy support                | Yes                                  |                | Yes, via MySQL Replication                                           | Yes, via MySQL Replication                                           | Yes, via MySQL Replication             |
| Update Schema On-Line                        | No                                   |                | No                                                                   | No                                                                   | Yes                                    |
| <b>Scalability</b>                           |                                      |                |                                                                      |                                                                      |                                        |
| Number of Nodes                              | One Master, Multiple Slaves          |                | One Active (primary), one Passive (secondary) Node                   | One Active (primary), one Passive (secondary) Node                   | 255                                    |
| Built-in Load Balancing                      | Reads, via MySQL Replication         |                | Reads, via MySQL Replication                                         | Reads, via MySQL Replication & During Failover                       | Yes, Reads and Writes                  |
| Supports Read-Intensive Workloads            | Yes                                  |                | Yes                                                                  | Yes                                                                  | Yes                                    |
| Supports Write-Intensive Workloads           | Yes, via Application-Level Sharding  |                | Yes, via Application-Level Sharding to Multiple Active/Passive Pairs | Yes, via Application-Level Sharding to Multiple Active/Passive Pairs | Yes, via Auto-Sharding                 |
| Scale On-Line (add nodes, repartition, etc.) | No                                   |                | No                                                                   | No                                                                   | Yes                                    |

## 15.1. Oracle VM Template for MySQL Enterprise Edition

Virtualization is a key technology to enable data center efficiency and high availability while providing the foundation for cloud computing. Integrating MySQL Enterprise Edition with Oracle Linux, the Oracle VM Template is the fastest, easiest, and most reliable way to provision virtualized MySQL instances, enabling users to meet the explosive demand for highly available services.

The Oracle VM Template enables rapid deployment and eliminates manual configuration efforts. It provides a pre-installed and pre-configured virtualized MySQL 5.5 Enterprise Edition software image running on Oracle Linux and Oracle VM, certified for production use. The MySQL software image has undergone extensive integration and quality assurance testing as part of the development process.

In addition to rapid provisioning, MySQL users also benefit from the integrated high availability features of Oracle VM which are designed to enable organizations to meet stringent SLA (Service Level Agreement) demands through a combination of:

- **Automatic recovery from failures**, with Oracle VM automatically restarting failed instances on available servers in the server pool after outages of the physical server, VM or MySQL database.
- **Live Migration**, enabling operations staff to move running instances of MySQL to alternative hosts within a server pool during maintenance operations.

Instructions for the creation, deployment and use of the Oracle VM Template for MySQL Enterprise Edition are available from:

- The Oracle VM Template for MySQL Enterprise Edition whitepaper: [http://www.mysql.com/why-mysql/white-papers/mysql\\_wp\\_oracle-vm-template-for-mee.php](http://www.mysql.com/why-mysql/white-papers/mysql_wp_oracle-vm-template-for-mee.php).
- The README file accompanying the download of the Template.

To download the Oracle VM Template for MySQL Enterprise, go to <http://edelivery.oracle.com/oraclevm> and follow these instructions:

- Complete your registration information (Name, Company Name, Email Address and Country) and click on the download agreement.
- Select "Oracle VM Templates" from the "Select a Product Pack" pull-down menu and click "Go".
- Select MySQL Enterprise from the list of Oracle VM Templates.
- Download and unzip the files and refer to the README for further instructions.

## 15.2. Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux

DRBD (Distributed Replication Block Device) is one of the leading solutions for MySQL HA (High Availability). When combined with Pacemaker and Corosync, users have:

- An end-to-end, integrated stack of mature and proven open source technologies, fully supported by Oracle (as part of MySQL Enterprise Edition).
- Automatic failover and recovery for service continuity.
- Mirroring, via synchronous replication, to ensure failover between nodes without the risk of losing committed transactions.
- Building of HA clusters from commodity hardware, without the requirement for shared-storage.

The following figure illustrates the stack that can be used to deliver a level of High Availability for the MySQL service.

At the lowest level, 2 hosts are required in order to provide physical redundancy; if using a virtual environment, those 2 hosts should be on different physical machines. It is an important feature that no shared storage is required. At any point in time, the services will be active on one host and in standby mode on the other.

Pacemaker and Corosync combine to provide the clustering layer that sits between the services and the underlying hosts and operating systems. Pacemaker is responsible for starting and stopping services, ensuring that they are running on exactly one host, thus delivering high availability and avoiding data corruption. Corosync provides the underlying messaging infrastructure between the nodes that enables Pacemaker to do its job; it also handles the nodes membership within the cluster and informs Pacemaker of any changes.



**Figure 15.2. MySQL, DRBD, Pacemaker, and Corosync Stack**

The core Pacemaker process does not have built-in knowledge of the specific services to be managed; instead, it uses agents that provide a wrapper for the service-specific actions. For example, in this solution we use agents for Virtual IP Addresses, MySQL and DRBD: these are all existing agents and come packaged with Pacemaker.

The essential services managed by Pacemaker in this configuration are DRBD, MySQL and the Virtual IP Address that applications use to connect to the active MySQL service.

DRBD synchronizes data at the block device (typically a spinning or solid state disk) – transparent to the application, database and even the file system. DRBD requires the use of a journaling file system such as `ext3` or `ext4`. For this solution, it acts in an active-standby mode: at any point in time, the directories being managed by DRBD are accessible for reads and writes on exactly one of the two hosts and inaccessible (even for reads) on the other. Any changes made on the active host are synchronously replicated to the standby host by DRBD.

Download the following guide for detailed instructions on installing, configuring, provisioning and testing the complete MySQL and DRBD stack, including:

- MySQL Database.
- DRBD kernel module and userland utilities.
- Pacemaker and Corosync cluster messaging and management processes.
- Oracle Linux operating system.

Download the guide at: <http://www.mysql.com/why-mysql/white-papers/mysql-high-availability-drbd-configuration-deployment-guide/>.

## Support for DRBD

The complete DRBD stack for MySQL has been certified by Oracle, and commercial support is available as part of [MySQL Enterprise Edition](#) and [Oracle Linux Premier Support](#), providing a single point of contact for the entire stack, whether issues relate to the operating system, DRBD, clustering software or MySQL.


## 15.3. Overview of MySQL with Windows Failover Clustering

Microsoft Windows is consistently ranked as the top development platform for MySQL, based on surveys of the MySQL user community.

MySQL Enterprise Edition is certified and supported with Windows Server 2008 R2 Failover Clustering (WSFC), enabling organizations to safely deploy business-critical applications demanding high levels of availability using Microsoft's native Windows clustering services.

The following figure illustrates the integration of MySQL with Windows Server Failover Clustering to provide a highly available service:

Figure 15.3. Typical MySQL HA Configuration with Windows Server Failover Clustering



In this architecture, MySQL is deployed in an Active / Passive configuration. Failures of either MySQL or the underlying server are automatically detected and the MySQL instance is restarted on the Passive node. Applications accessing the database, as well as any MySQL replication slaves, can automatically reconnect to the new MySQL process using the same Virtual IP address once MySQL recovery has completed and it starts accepting connections.

MySQL with Windows Failover Clustering requires at least 2 servers within the cluster together with shared storage (for example, FC-AL SAN or iSCSI disks).

The MySQL binaries and data files are stored in the shared storage and Windows Failover Clustering ensures that only one of the cluster nodes will access those files at any point in time.

Clients connect to the MySQL service through a Virtual IP Address (VIP). In the event of failover they experience a brief loss of connection, but otherwise do not need to be aware that the failover has happened, other than to handle the failure of any transactions that were active when the failover occurred.

You can learn more about configuring MySQL with Windows Server Failover Clustering from the whitepaper posted here: [http://www.mysql.com/why-mysql/white-papers/mysql\\_wp\\_windows\\_failover\\_clustering.php](http://www.mysql.com/why-mysql/white-papers/mysql_wp_windows_failover_clustering.php)

For background and usage information about Windows Server Failover Clustering, see these pages on the Microsoft Technet site:

- [Failover Clustering](#)
- [Failover Clusters in Windows Server 2008 R2](#)

## 15.4. Using MySQL within an Amazon EC2 Instance

The Amazon Elastic Compute Cloud (EC2) service provides virtual servers that you can build and deploy to run a variety of different applications and services, including MySQL. The EC2 service is based around the Xen framework, supporting x86, Linux based, platforms with individual instances of a virtual machine referred to as an Amazon Machine Image (AMI). You have complete (root) access to the AMI instance that you create, enabling you to configure and install your AMI in any way you choose.

To use EC2, you create an AMI based on the configuration and applications that you intend to use, and upload the AMI to the Amazon Simple Storage Service (S3). From the S3 resource, you can deploy one or more copies of the AMI to run as an instance within the EC2 environment. The EC2 environment provides management and control of the instance and contextual information about the instance while it is running.

Because you can create and control the AMI, the configuration, and the applications, you can deploy and create any environment you choose. This includes a basic MySQL server in addition to more extensive replication, HA and scalability scenarios that enable you to take advantage of the EC2 environment, and the ability to deploy additional instances as the demand for your MySQL services and applications grow.

To aid the deployment and distribution of work, three different Amazon EC2 instances are available, small (identified as `m1.small`), large (`m1.large`) and extra large (`m1.xlarge`). The different types provide different levels of computing power measured in EC2 computer units (ECU). A summary of the different instance configurations is shown in the following table.

| EC2 Attribute   | Small  | Large  | Extra Large |
|-----------------|--------|--------|-------------|
| Platform        | 32-bit | 64-bit | 64-bit      |
| CPU cores       | 1      | 2      | 4           |
| ECUs            | 1      | 4      | 8           |
| RAM             | 1.7GB  | 7.5GB  | 15GB        |
| Storage         | 150GB  | 840GB  | 1680GB      |
| I/O Performance | Medium | High   | High        |

The typical model for deploying and using MySQL within the EC2 environment is to create a basic AMI that you can use to hold your database data and application. Once the basic environment for your database and application has been created you can then choose to deploy the AMI to a suitable instance. Here the flexibility of having an AMI that can be re-deployed from the small to the large or extra large EC2 instance makes it easy to upgrade the hardware environment without rebuilding your application or database stack.

To get started with MySQL on EC2, including information on how to set up and install MySQL within an EC2 installation and how to port and migrate your data to the running instance, see [Section 15.4.1, "Setting Up MySQL on an EC2 AMI"](#).

For tips and advice on how to create a scalable EC2 environment using MySQL, including guides on setting up replication, see [Section 15.4.3, “Deploying a MySQL Database Using EC2”](#).

### 15.4.1. Setting Up MySQL on an EC2 AMI

There are many different ways of setting up an EC2 AMI with MySQL, including using any of the pre-configured AMIs supplied by Amazon.

The default *Getting Started* AMI provided by Amazon uses Fedora Core 4, and you can install MySQL by using `yum`:

```
shell> yum install mysql
```

This installs both the MySQL server and the Perl DBD::mysql driver for the Perl DBI API.

Alternatively, you can use one of the AMIs that include MySQL within the standard installation.

Finally, you can also install a standard version of MySQL downloaded from the MySQL Web site. The installation process and instructions are identical to any other installation of MySQL on Linux. See [Chapter 2, \*Installing and Upgrading MySQL\*](#).

The standard configuration for MySQL places the data files in the default location, `/var/lib/mysql`. The default data directory on an EC2 instance is `/mnt` (although on the large and extra large instance you can alter this configuration). You must edit `/etc/my.cnf` to set the `datadir` [418] option to point to the larger storage area.



#### Important

The first time you use the main storage location within an EC2 instance it needs to be initialized. The initialization process starts automatically the first time you write to the device. You can start using the device right away, but the write performance of the new device is significantly lower on the initial writes until the initialization process has finished.

To avoid this problem when setting up a new instance, you should start the initialization process before populating your MySQL database. One way to do this is to use `dd` to write to the file system:

```
root-shell> dd if=/dev/zero of=initialize bs=1024M count=50
```

The preceding creates a 50GB on the file system and starts the initialization process. Delete the file once the process has finished.

The initialization process can be time-consuming. On the small instance, initialization takes between two and three hours. For the large and extra large drives, the initialization can be 10 or 20 hours, respectively.

In addition to configuring the correct storage location for your MySQL data files, also consider setting the following other settings in your instance before you save the instance configuration for deployment:

- Set the MySQL server ID, so that when you use it for replication, the ID information is set correctly.
- Enabling binary logging, so that replication can be initialized without starting and stopping the server.
- Set the caching and memory parameters for your storage engines. There are no limitations or restrictions on what storage engines you use in your EC2 environment. Choose a configuration, possibly using one of the standard configurations provided with MySQL appropriate for the instance on which you expect to deploy. The large and extra large instances have RAM that can be dedicated to caching. Be aware that if you choose to install `memcached` on the servers as part of your application stack you must ensure there is enough memory for both MySQL and `memcached`.

Once you have configured your AMI with MySQL and the rest of your application stack, save the AMI so that you can deploy and reuse the instance.

Once you have your application stack configured in an AMI, populating your MySQL database with data should be performed by creating a dump of your database using `mysqldump`, transferring the dump to the EC2 instance, and then reloading the information into the EC2 instance database.

Before using your instance with your application in a production situation, be aware of the limitations of the EC2 instance environment. See [Section 15.4.2, “EC2 Instance Limitations”](#). To begin using your MySQL AMI, consult the notes on deployment. See [Section 15.4.3, “Deploying a MySQL Database Using EC2”](#).

## 15.4.2. EC2 Instance Limitations

Be aware of the following limitations of the EC2 instances before deploying your applications. Although these shouldn't affect your ability to deploy within the Amazon EC2 environment, they may alter the way you setup and configure your environment to support your application.

- Data stored within instances is not persistent. If you create an instance and populate the instance with data, then the data only remains in place while the machine is running, and does not survive a reboot. If you shut down the instance, any data it contained is lost.

To ensure that you do not lose information, take regular backups using `mysqldump`. If the data being stored is critical, consider using replication to keep a “live” backup of your data in the event of a failure. When creating a backup, write the data to the Amazon S3 service to avoid the transfer charges applied when copying data offsite.

- EC2 instances are not persistent. If the hardware on which an instance is running fails, the instance is shut down. This can lead to loss of data or service.

However, if you use EBS, you can attach an EBS storage volume to an EC2 instance, and that EBS volume is persistent. Like a disk, an EBS volume can fail, but it is possible to create point-in-time snapshots of the volume. Snapshots are persisted to Amazon S3 and can be used to restore data in the event of volume failure.

- To replicate your EC2 instances to a non-EC2 environment, be aware of the transfer costs to and from the EC2 service. Data transfer between different EC2 instances is free, so using replication within the EC2 environment does not incur additional charges.
- Certain HA features are either not directly supported, or have limiting factors or problems that could reduce their utility. For example, using DRBD or MySQL Cluster might not work. The default storage configuration is also not redundant. You can use software-based RAID to improve redundancy, but this implies a further performance hit.

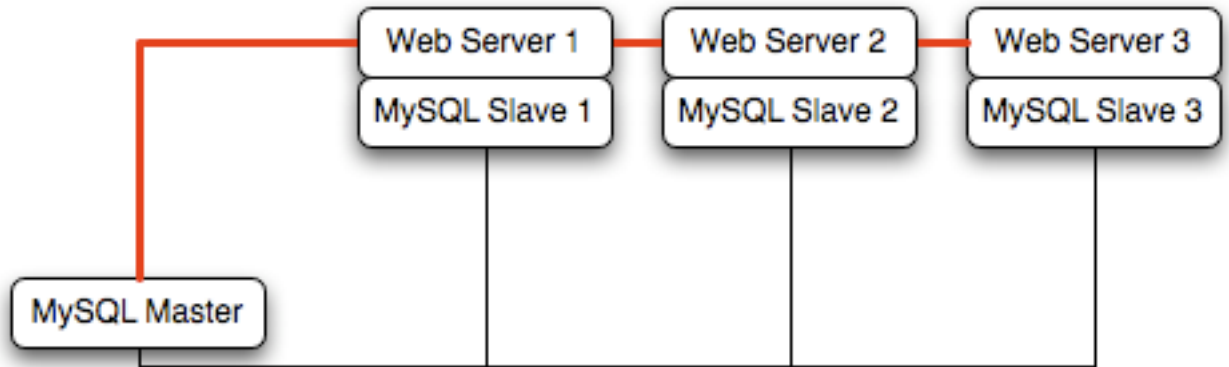
## 15.4.3. Deploying a MySQL Database Using EC2

Because you cannot guarantee the uptime and availability of your EC2 instances, when deploying MySQL within the EC2 environment, use an approach that enables you to easily distribute work among your EC2 instances. There are a number of ways of doing this. Using sharding techniques, where you split the application across multiple servers dedicating specific blocks of your dataset and users to different servers is an effective way of doing this. As a general rule, it is easier to create more EC2 instances to support more users than to upgrade the instance to a larger machine.

The EC2 architecture works best when you treat the EC2 instances as temporary, cache-based solutions, rather than as a long-term, high availability solution. In addition to using multiple machines, take advantage of other services, such as `memcached` to provide additional caching for your application to help reduce the load on the MySQL server so that it can concentrate on writes. On the large and extra large instances within EC2, the RAM available can provide a large memory cache for data.

Most types of scale-out topology that you would use with your own hardware can be used and applied within the EC2 environment. However, use the limitations and advice already given to ensure that any potential failures do not lose you any data. Also, because the relative power of each EC2 instance is so low, be prepared to alter your application to use sharding and add further EC2 instances to improve the performance of your application.

For example, take the typical scale-out environment shown following, where a single master replicates to one or more slaves (three in this example), with a web server running on each replication slave.



You can reproduce this structure completely within the EC2 environment, using an EC2 instance for the master, and one instance for each of the web and MySQL slave servers.




### Note

Within the EC2 environment, internal (private) IP addresses used by the EC2 instances are constant. Always use these internal addresses and names when communicating between instances. Only use public IP addresses when communicating with the outside world - for example, when publicizing your application.

To ensure reliability of your database, add at least one replication slave dedicated to providing an active backup and storage to the Amazon S3 facility. You can see an example of this in the following topology.

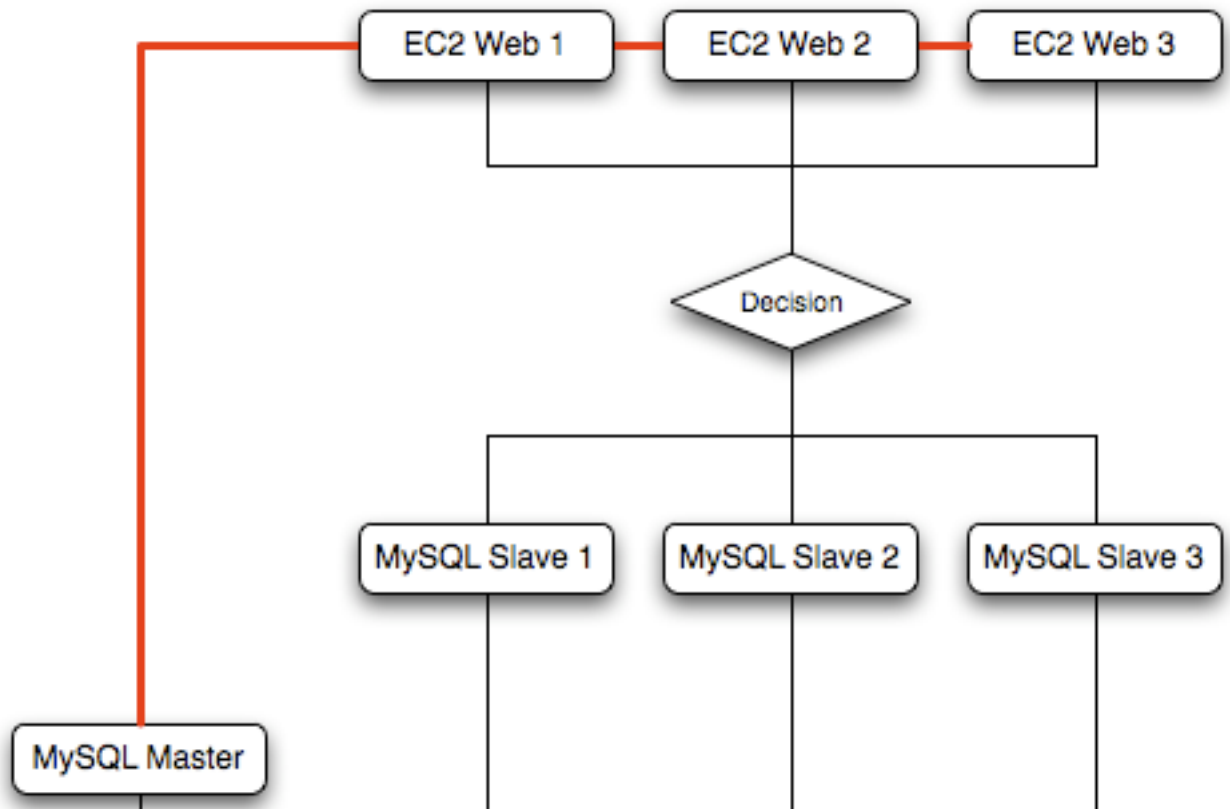




**Using [memcached](#)** within your EC2 instances should provide better performance. The large and extra large instances have a significant amount of RAM. To use [memcached](#) in your application, when loading information from the database, first check whether the item exists in the cache. If the data you are looking for exists in the cache, use it. If not, reload the data from the database and populate the cache.

**Sharding** divides up data in your entire database by allocating individual machines or machine groups to provide a unique set of data according to an appropriate group. For example, you might put all users with a surname ending in the letters A-D onto a single server. When a user connects to the application and their surname is known, queries can be redirected to the appropriate MySQL server.

When using sharding with EC2, separate the web server and MySQL server into separate EC2 instances, and then apply the sharding decision logic into your application. Once you know which MySQL server you should be using for accessing the data you then distribute queries to the appropriate server. You can see a sample of this in the following illustration.



### Warning

With sharding and EC2, be careful that the potential for failure of an instance does not affect your application. If the EC2 instance that provides the MySQL server for a particular shard fails, then all of the data on that shard becomes unavailable.

## 15.5. Using ZFS Replication

To support high availability environments, providing an instant copy of the information on both the currently active machine and the hot backup is a critical part of the HA solution. There are many solutions to this problem, including [Chapter 16, Replication](#) and [Section 15.2, “Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux”](#).

The ZFS file system provides functionality to create a snapshot of the file system contents, transfer the snapshot to another machine, and extract the snapshot to recreate the file system. You can create a snapshot at any time, and you can create as many snapshots as you like. By continually creating, transferring, and restoring snapshots, you can provide synchronization between one or more machines in a fashion similar to DRBD.

The following example shows a simple Solaris system running with a single ZFS pool, mounted at `/scratchpool`:

| Filesystem      | size | used | avail | capacity | Mounted on        |
|-----------------|------|------|-------|----------|-------------------|
| /dev/dsk/c0d0s0 | 4.6G | 3.7G | 886M  | 82%      | /                 |
| /devices        | 0K   | 0K   | 0K    | 0%       | /devices          |
| ctfs            | 0K   | 0K   | 0K    | 0%       | /system/contract  |
| proc            | 0K   | 0K   | 0K    | 0%       | /proc             |
| mnttab          | 0K   | 0K   | 0K    | 0%       | /etc/mnttab       |
| swap            | 1.4G | 892K | 1.4G  | 1%       | /etc/svc/volatile |
| objfs           | 0K   | 0K   | 0K    | 0%       | /system/object    |

```

/usr/lib/libc/libc_hwcapi.so.1
4.6G 3.7G 886M 82% /lib/libc.so.1
fd
0K 0K 0K 0% /dev/fd
swap
1.4G 40K 1.4G 1% /tmp
swap
1.4G 28K 1.4G 1% /var/run
/dev/dsk/c0d0s7
26G 913M 25G 4% /export/home
scratchpool
16G 24K 16G 1% /scratchpool

```

The MySQL data is stored in a directory on `/scratchpool`. To help demonstrate some of the basic replication functionality, there are also other items stored in `/scratchpool` as well:

```

total 17
drwxr-xr-x 31 root bin 50 Jul 21 07:32 DTT/
drwxr-xr-x 4 root bin 5 Jul 21 07:32 SUNWmlib/
drwxr-xr-x 14 root sys 16 Nov 5 09:56 SUNWspro/
drwxrwxrwx 19 1000 1000 40 Nov 6 19:16 emacs-22.1/

```

To create a snapshot of the file system, you use `zfs snapshot`, specifying the pool and the snapshot name:

```
root-shell> zfs snapshot scratchpool@snap1
```

To list the snapshots already taken:

```

root-shell> zfs list -t snapshot
NAME          USED  AVAIL  REFER  MOUNTPOINT
scratchpool@snap1  0    -    24.5K  -
scratchpool@snap2  0    -    24.5K  -

```

The snapshots themselves are stored within the file system metadata, and the space required to keep them varies as time goes on because of the way the snapshots are created. The initial creation of a snapshot is very quick, because instead of taking an entire copy of the data and metadata required to hold the entire snapshot, ZFS records only the point in time and metadata of when the snapshot was created.

As more changes to the original file system are made, the size of the snapshot increases because more space is required to keep the record of the old blocks. If you create lots of snapshots, say one per day, and then delete the snapshots from earlier in the week, the size of the newer snapshots might also increase, as the changes that make up the newer state have to be included in the more recent snapshots, rather than being spread over the seven snapshots that make up the week.

You cannot directly back up the snapshots because they exist within the file system metadata rather than as regular files. To get the snapshot into a format that you can copy to another file system, tape, and so on, you use the `zfs send` command to create a stream version of the snapshot.

For example, to write the snapshot out to a file:

```
root-shell> zfs send scratchpool@snap1 >/backup/scratchpool-snap1
```

Or tape:

```
root-shell> zfs send scratchpool@snap1 >/dev/rmt/0
```

You can also write out the incremental changes between two snapshots using `zfs send`:

```
root-shell> zfs send scratchpool@snap1 scratchpool@snap2 >/backup/scratchpool-changes
```

To recover a snapshot, you use `zfs recv`, which applies the snapshot information either to a new file system, or to an existing one.

## 15.5.1. Using ZFS for File System Replication

Because `zfs send` and `zfs recv` use streams to exchange data, you can use them to replicate information from one system to another by combining `zfs send`, `ssh`, and `zfs recv`.

For example, to copy a snapshot of the `scratchpool` file system to a new file system called `slavepool` on a new server, you would use the following command. This sequence combines the snapshot of `scratchpool`, the transmission to the slave machine (using `ssh` with login credentials), and the recovery of the snapshot on the slave using `zfs recv`:

```
root-shell> zfs send scratchpool@snap1 |ssh id@host pfexec zfs recv -F slavepool
```

The first part of the pipeline, `zfs send scratchpool@snap1`, streams the snapshot. The `ssh` command, and the command that it executes on the other server, `pfexec zfs recv -F slavepool`, receives the streamed snapshot data and writes it to `slavepool`. In this instance, I've specified the `-F` option which forces the snapshot data to be applied, and is therefore destructive. This is fine, as I'm creating the first version of my replicated file system.

On the slave machine, the replicated file system contains the exact same content:

```
root-shell> ls -al /slavepool/
total 23
drwxr-xr-x  6 root  root      7 Nov  8 09:13 ./
drwxr-xr-x 29 root  root     34 Nov  9 07:06 ../
drwxr-xr-x 31 root  bin     50 Jul 21 07:32 DTT/
drwxr-xr-x  4 root  bin     5 Jul 21 07:32 SUNWmlib/
drwxr-xr-x 14 root  sys    16 Nov  5 09:56 SUNWspr0/
drwxrwxrwx 19 1000 1000   40 Nov  6 19:16 emacs-22.1/
```

Once a snapshot has been created, to synchronize the file system again, you create a new snapshot and then use the incremental snapshot feature of `zfs send` to send the changes between the two snapshots to the slave machine again:

```
root-shell> zfs send -i scratchpool@snapshot1 scratchpool@snapshot2 |ssh id@host pfexec zfs recv slavepool
```

This operation only succeeds if the file system on the slave machine has not been modified at all. You cannot apply the incremental changes to a destination file system that has changed. In the example above, the `ls` command would cause problems by changing the metadata, such as the last access time for files or directories.

To prevent changes on the slave file system, set the file system on the slave to be read-only:

```
root-shell> zfs set readonly=on slavepool
```

Setting `readonly` means that you cannot change the file system on the slave by normal means, including the file system metadata. Operations that would normally update metadata (like our `ls`) silently perform their function without attempting to update the file system state.

In essence, the slave file system is nothing but a static copy of the original file system. However, even when configured to be read-only, a file system can have snapshots applied to it. With the file system set to read only, re-run the initial copy:

```
root-shell> zfs send scratchpool@snap1 |ssh id@host pfexec zfs recv -F slavepool
```

Now you can make changes to the original file system and replicate them to the slave.

### 15.5.2. Configuring MySQL for ZFS Replication

Configuring MySQL on the source file system is a case of creating the data on the file system that you intend to replicate. The configuration file in the example below has been updated to use `scratchpool/mysql-data` as the data directory, and now you can initialize the tables:

```
root-shell> mysql_install_db --defaults-file=/etc/mysql/5.5/my.cnf --user=mysql
```

To synchronize the initial information, perform a new snapshot and then send an incremental snapshot to the slave using `zfs send`:

```
root-shell> zfs snapshot scratchpool@snap2
```

```
root-shell> zfs send -i scratchpool@snap1 scratchpool@snap2|ssh id@host pfexec zfs recv slavepool
```

Doublecheck that the slave has the data by looking at the MySQL data directory on the [slavepool](#):

```
root-shell> ls -al /slavepool/mysql-data/
```

Now you can start up MySQL, create some data, and then replicate the changes using `zfs send/ zfs recv` to the slave to synchronize the changes.

The rate at which you perform the synchronization depends on your application and environment. The limitation is the speed required to perform the snapshot and then to send the changes over the network.

To automate the process, create a script that performs the snapshot, send, and receive operation, and use `cron` to synchronize the changes at set times or intervals.

### 15.5.3. Handling MySQL Recovery with ZFS

When using ZFS replication to provide a constant copy of your data, ensure that you can recover your tables, either manually or automatically, in the event of a failure of the original system.

In the event of a failure, follow this sequence:

1. Stop the script on the master, if it is still up and running.
2. Set the slave file system to be read/write:

```
root-shell> zfs set readonly=off slavepool
```

3. Start up `mysqld` on the slave. If you are using [InnoDB](#), you get auto-recovery, if it is needed, to make sure the table data is correct, as shown here when I started up from our mid-INSERT snapshot:

```
InnoDB: The log sequence number in ibdata files does not match
InnoDB: the log sequence number in the ib_logfiles!
081109 15:59:59 InnoDB: Database was not shut down normally!
InnoDB: Starting crash recovery.
InnoDB: Reading tablespace information from the .ibd files...
InnoDB: Restoring possible half-written data pages from the doublewrite
InnoDB: buffer...
081109 16:00:03 InnoDB: Started; log sequence number 0 1142807951
081109 16:00:03 [Note] /slavepool/mysql-5.0.67-solaris10-i386/bin/mysqld: ready for connections.
Version: '5.0.67' socket: '/tmp/mysql.sock' port: 3306 MySQL Community Server (GPL)
```

Use [InnoDB](#) tables and a regular synchronization schedule to reduce the risk for significant data loss. On MyISAM tables, you might need to run `REPAIR TABLE`, and you might even have lost some information.

## 15.6. Using MySQL with [memcached](#)

[memcached](#) is a simple, highly scalable key-based cache that stores data and objects wherever dedicated or spare RAM is available for quick access by applications, without going through layers of parsing or disk I/O. To use, you run the `memcached` command on one or more hosts and then use the shared cache to store objects. For more usage instructions, see [Section 15.6.2, “Using memcached”](#)

Benefits of using [memcached](#) include:


- Because all information is stored in RAM, the access speed is faster than loading the information each time from disk.
- Because the “value” portion of the key-value pair does not have any data type restrictions, you can cache data such as complex structures, documents, images, or a mixture of such things.

- If you use the in-memory cache to hold transient information, or as a read-only cache for information also stored in a database, the failure of any memcached server is not critical. For persistent data, you can fall back to an alternative lookup method using database queries, and reload the data into RAM on a different server.

The typical usage environment is to modify your application so that information is read from the cache provided by memcached. If the information is not in memcached, then the data is loaded from the MySQL database and written into the cache so that future requests for the same object benefit from the cached data.

For a typical deployment layout, see [Figure 15.4, “memcached Architecture Overview”](#).

**Figure 15.4. memcached Architecture Overview**



In the example structure, any of the clients can contact one of the memcached servers to request a given key. Each client is configured to talk to all of the servers shown in the illustration. Within the client, when the request is made to store the information, the key used to reference the data is hashed and this hash is then used to select one of the memcached servers. The selection of the memcached server takes place on the client before the server is contacted, keeping the process lightweight.

The same algorithm is used again when a client requests the same key. The same key generates the same hash, and the same memcached server is selected as the source for the data. Using this method, the cached data is spread among all of the memcached servers, and the cached information is accessible from any client. The result is a distributed, memory-based, cache that can return information, particularly complex data and structures, much faster than natively reading the information from the database.

The data held within a traditional memcached server is never stored on disk (only in RAM, which means there is no persistence of data), and the RAM cache is always populated from the backing store (a MySQL database). If a memcached server fails, the data can always be recovered from the MySQL database.

## memcached Integration with MySQL Storage Engines

In April 2011, MySQL announced the preview of a new memcached interface for the InnoDB and MySQL Cluster storage engines.

Using the memcached API, web services can directly access the InnoDB and MySQL Cluster storage engines without transformations to SQL, ensuring low latency and high throughput for read/write queries. Operations such as SQL parsing are eliminated and more of the server's hardware resources (CPU, memory and I/O) are dedicated to servicing the query within the storage engine itself. The memcached data can be persisted to disk while still cached in memory for fast retrieval.

These are targeted to be incorporated into future MySQL 5.6 Milestone and MySQL Cluster Development Releases.

You can learn more about these interfaces from this Dev Zone article: <http://dev.mysql.com/tech-resources/articles/nosql-to-mysql-with-memcached.html>.

## 15.6.1. Installing `memcached`

You can build and install `memcached` from the source code directly, or you can use an existing operating system package or installation.

### Installing `memcached` from a Binary Distribution

To install `memcached` on a Red Hat, or Fedora host, use `yum`:

```
root-shell> yum install memcached
```



#### Note

On CentOS, you may be able to obtain a suitable RPM from another source, or use the source tarball.

To install `memcached` on a Debian or Ubuntu host, use `apt-get`:

```
root-shell> apt-get install memcached
```

To install `memcached` on a Gentoo host, use `emerge`:

```
root-shell> emerge install memcached
```

### Building `memcached` from Source

On other Unix-based platforms, including Solaris, AIX, HP-UX and Mac OS X, and Linux distributions not mentioned already, you must install from source. For Linux, make sure you have a 2.6-based kernel, which includes the improved `epoll` interface. For all platforms, ensure that you have `libevent` 1.1 or higher installed. You can obtain `libevent` from [libevent web page](#).

You can obtain the source for `memcached` from [memcached Web site](#).

To build `memcached`, follow these steps:

1. Extract the `memcached` source package:

```
shell> gunzip -c memcached-1.2.5.tar.gz | tar xf -
```

2. Change to the `memcached-1.2.5` directory:

```
shell> cd memcached-1.2.5
```

3. Run `configure`

```
shell> ./configure
```

Some additional options you might specify to the `configure`:

- `--prefix`

To specify a different installation directory, use the `--prefix` option:

```
shell> ./configure --prefix=/opt
```

The default is to use the `/usr/local` directory.

- `--with-libevent`



If you have installed `libevent` and `configure` cannot find the library, use the `--with-libevent` option to specify the location of the installed library.

- `--enable-64bit`

To build a 64-bit version of `memcached` (which enables you to use a single instance with a large RAM allocation), use `--enable-64bit`.

- `--enable-threads`

To enable multi-threading support in `memcached`, which improves the response times on servers with a heavy load, use `--enable-threads`. You must have support for the POSIX threads within your operating system to enable thread support. For more information on the threading support, see [Section 15.6.2.7, “memcached Thread Support”](#).

- `--enable-dtrace`

`memcached` includes a range of DTrace threads that can be used to monitor and benchmark a `memcached` instance. For more information, see [Section 15.6.2.5, “Using memcached and DTrace”](#).

4. Run `make` to build `memcached`:

```
shell> make
```

5. Run `make install` to install `memcached`:

```
shell> make install
```

## 15.6.2. Using memcached

To start using `memcached`, start the `memcached` service on one or more servers. Running `memcached` sets up the server, allocates the memory and starts listening for connections from clients.



### Note

You do not need to be a privileged user (`root`) to run `memcached` except to listen on one of the privileged TCP/IP ports (below 1024). You must, however, use a user that has not had their memory limits restricted using `setrlimit` or similar.

To start the server, run `memcached` as a nonprivileged (that is, non-`root`) user:

```
shell> memcached
```

By default, `memcached` uses the following settings:

- Memory allocation of 64MB
- Listens for connections on all network interfaces, using port 11211
- Supports a maximum of 1024 simultaneous connections

Typically, you would specify the full combination of options that you want when starting `memcached`, and normally provide a startup script to handle the initialization of `memcached`. For example, the following line starts `memcached` with a maximum of 1024MB RAM for the cache, listening on port 11211 on the IP address 192.168.0.110, running as a background daemon:

```
shell> memcached -d -m 1024 -p 11211 -l 192.168.0.110
```

To ensure that `memcached` is started up on boot, check the init script and configuration parameters.

memcached supports the following options:

- `-u user`

If you start memcached as `root`, use the `-u` option to specify the user for executing memcached:

```
shell> memcached -u memcache
```

- `-m memory`

Set the amount of memory allocated to memcached for object storage. Default is 64MB.

To increase the amount of memory allocated for the cache, use the `-m` option to specify the amount of RAM to be allocated (in megabytes). The more RAM you allocate, the more data you can store and therefore the more effective your cache is.



### Warning

Do not specify a memory allocation larger than your available RAM. If you specify too large a value, then some RAM allocated for memcached uses swap space, and not physical RAM. This may lead to delays when storing and retrieving values, because data is swapped to disk, instead of storing the data directly in RAM.

You can use the output of the `vmstat` command to get the free memory, as shown in `free` column:

```
shell> vmstat
kthr      memory          page        disk          faults        cpu
r  b  w    swap  free  re  mf  pi  po  fr  de  sr  sl  s2  --  --    in   sy   cs  us  sy  id
0  0  0  5170504 3450392 2   7   2  0  0  0  4  0  0  0  0   296  54  199  0   0 100
```

For example, to allocate 3GB of RAM:

```
shell> memcached -m 3072
```

On 32-bit x86 systems where you are using PAE to access memory above the 4GB limit, you cannot allocate RAM beyond the maximum process size. You can get around this by running multiple instances of memcached, each listening on a different port:

```
shell> memcached -m 1024 -p11211
shell> memcached -m 1024 -p11212
shell> memcached -m 1024 -p11213
```



### Note

On all systems, particularly 32-bit, ensure that you leave enough room for both memcached application in addition to the memory setting. For example, if you have a dedicated memcached host with 4GB of RAM, do not set the memory size above 3500MB. Failure to do this may cause either a crash or severe performance issues.

- `-l interface`

Specify a network interface/address to listen for connections. The default is to listen on all available address (`INADDR_ANY`).

```
shell> memcached -l 192.168.0.110
```

Support for IPv6 address support was added in memcached 1.2.5.

- `-p port`

Specify the TCP port to use for connections. Default is 18080.

```
shell> memcached -p 18080
```

- `-U port`

Specify the UDP port to use for connections. Default is 11211, 0 switches UDP off.

```
shell> memcached -U 18080
```

- `-s socket`

Specify a Unix socket to listen on.

If you are running `memcached` on the same server as the clients, you can disable the network interface and use a local UNIX socket using the `-s` option:

```
shell> memcached -s /tmp/memcached
```

Using a UNIX socket automatically disables network support, and saves network ports (allowing more ports to be used by your web server or other process).

- `-a mask`

Specify the access mask to be used for the Unix socket, in octal. Default is 0700.

- `-c connections`

Specify the maximum number of simultaneous connections to the `memcached` service. The default is 1024.

```
shell> memcached -c 2048
```

Use this option, either to reduce the number of connections (to prevent overloading `memcached` service) or to increase the number to make more effective use of the server running `memcached` server.

- `-t threads`

Specify the number of threads to use when processing incoming requests.

By default, `memcached` is configured to use 4 concurrent threads. The threading improves the performance of storing and retrieving data in the cache, using a locking system to prevent different threads overwriting or updating the same values. To increase or decrease the number of threads, use the `-t` option:

```
shell> memcached -t 8
```

- `-d`

Run `memcached` as a daemon (background) process:

```
shell> memcached -d
```

- `-r`

Maximize the size of the core file limit. In the event of a failure, this attempts to dump the entire memory space to disk as a core file, up to any limits imposed by `setrlimit`.

- `-M`

Return an error to the client when the memory has been exhausted. This replaces the normal behavior of removing older items from the cache to make way for new items.

- `-k`

Lock down all paged memory. This reserves the memory before use, instead of allocating new slabs of memory as new items are stored in the cache.

**Note**

There is a user-level limit on how much memory you can lock. Trying to allocate more than the available memory fails. You can set the limit for the user you started the daemon with (not for the `-u user` user) within the shell by using `ulimit -S -l NUM_KB`

- `-v`

Verbose mode. Prints errors and warnings while executing the main event loop.

- `-vv`

Very verbose mode. In addition to information printed by `-v`, also prints each client command and the response.

- `-vvv`

Extremely verbose mode. In addition to information printed by `-vv`, also show the internal state transitions.

- `-h`

Print the help message and exit.

- `-i`

Print the `memcached` and `libevent` license.

- `-I mem`

Specify the maximum size permitted for storing an object within the `memcached` instance. The size supports a unit postfix (`k` for kilobytes, `m` for megabytes). For example, to increase the maximum supported object size to 32MB:

```
shell> memcached -I 32m
```

The maximum object size you can specify is 128MB, the default remains at 1MB.

This option was added in 1.4.2.

- `-b`

Set the backlog queue limit. The backlog queue configures how many network connections can be waiting to be processed by `memcached`. Increasing this limit may reduce errors received by the client that it is not able to connect to the `memcached` instance, but does not improve the performance of the server. The default is 1024.

- `-P pidfile`

Save the process ID of the `memcached` instance into `file`.

- `-f`

Set the chunk size growth factor. When allocating new memory chunks, the allocated size of new chunks is determined by multiplying the default slab size by this factor.

To see the effects of this option without extensive testing, use the `-vv` command-line option to show the calculated slab sizes. For more information, see [Section 15.6.2.8, “memcached Logs”](#).

- `-n bytes`

The minimum space allocated for the key+value+flags information. The default is 48 bytes.

- `-L`

On systems that support large memory pages, enables large memory page use. Using large memory pages enables `memcached` to allocate the item cache in one large chunk, which can improve the performance by reducing the number misses when accessing memory.

- `-C`

Disable the use of compare and swap (CAS) operations.

This option was added in `memcached` 1.3.x.

- `-D char`

Set the default character to be used as a delimiter between the key prefixes and IDs. This is used for the per-prefix statistics reporting (see [Section 15.6.4, “Getting `memcached` Statistics”](#)). The default is the colon (:). If this option is used, statistics collection is turned on automatically. If not used, you can enable stats collection by sending the `stats detail on` command to the server.

This option was added in `memcached` 1.3.x.

- `-R num`

Sets the maximum number of requests per event process. The default is 20.

- `-B protocol`

Set the binding protocol, that is, the default `memcached` protocol support for client connections. Options are `ascii`, `binary` or `auto`. Automatic (`auto`) is the default.

This option was added in `memcached` 1.4.0.

### 15.6.2.1. `memcached` Deployment

When using `memcached` you can use a number of different potential deployment strategies and topologies. The exact strategy to use depends on your application and environment. When developing a system for deploying `memcached` within your system, keep in mind the following points:

- `memcached` is only a caching mechanism. It shouldn't be used to store information that you cannot otherwise afford to lose and then load from a different location.
- There is no security built into the `memcached` protocol. At a minimum, make sure that the servers running `memcached` are only accessible from inside your network, and that the network ports being used are blocked (using a firewall or similar). If the information on the `memcached` servers that is being stored is any sensitive, then encrypt the information before storing it in `memcached`.
- `memcached` does not provide any sort of failover. Because there is no communication between different `memcached` instances. If an instance fails, your application must be capable of removing it from the list, reloading the data and then writing data to another `memcached` instance.
- Latency between the clients and the `memcached` can be a problem if you are using different physical machines for these tasks. If you find that the latency is a problem, move the `memcached` instances to be on the clients.
- Key length is determined by the `memcached` server. The default maximum key size is 250 bytes.
- Try to use at least two `memcached` instances, especially for multiple clients, to avoid having a single point of failure. Ideally, create as many `memcached` nodes as possible. When adding and removing

---

`memcached` instances from a pool, the hashing and distribution of key/value pairs may be affected. For information on how to avoid problems, see [Section 15.6.2.4, “memcached Hashing/Distribution Types”](#).

### 15.6.2.2. Using Namespaces

The `memcached` cache is a very simple massive key/value storage system, and as such there is no way of compartmentalizing data automatically into different sections. For example, if you are storing information by the unique ID returned from a MySQL database, then storing the data from two different tables could run into issues because the same ID might be valid in both tables.

Some interfaces provide an automated mechanism for creating *namespaces* when storing information into the cache. In practice, these namespaces are merely a prefix before a given ID that is applied every time a value is stored or retrieve from the cache.

You can implement the same basic principle by using keys that describe the object and the unique identifier within the key that you supply when the object is stored. For example, when storing user data, prefix the ID of the user with `user:` or `user-`.



#### Note

Using namespaces or prefixes only controls the keys stored/retrieved. There is no security within `memcached`, and therefore no way to enforce that a particular client only accesses keys with a particular namespace. Namespaces are only useful as a method of identifying data and preventing corruption of key/value pairs.

### 15.6.2.3. Data Expiry

There are two types of data expiry within a `memcached` instance. The first type is applied at the point when you store a new key/value pair into the `memcached` instance. If there is not enough space within a suitable slab to store the value, then an existing least recently used (LRU) object is removed (evicted) from the cache to make room for the new item.

The LRU algorithm ensures that the object that is removed is one that is either no longer in active use or that was used so long ago that its data is potentially out of date or of little value. However, in a system where the memory allocated to `memcached` is smaller than the number of regularly used objects required in the cache, a lot of expired items could be removed from the cache even though they are in active use. You use the statistics mechanism to get a better idea of the level of evictions (expired objects). For more information, see [Section 15.6.4, “Getting memcached Statistics”](#).

You can change this eviction behavior by setting the `-M` command-line option when starting `memcached`. This option forces an error to be returned when the memory has been exhausted, instead of automatically evicting older data.

The second type of expiry system is an explicit mechanism that you can set when a key/value pair is inserted into the cache, or when deleting an item from the cache. Using an expiration time can be a useful way of ensuring that the data in the cache is up to date and in line with your application needs and requirements.

A typical scenario for explicitly setting the expiry time might include caching session data for a user when accessing a Web site. `memcached` uses a lazy expiry mechanism where the explicit expiry time that has been set is compared with the current time when the object is requested. Only objects that have not expired are returned.

You can also set the expiry time when explicitly deleting an object from the cache. In this case, the expiry time is really a timeout and indicates the period when any attempts to set the value for a given key are rejected.

### 15.6.2.4. memcached Hashing/Distribution Types

The `memcached` client interface supports a number of different distribution algorithms that are used in multi-server configurations to determine which host should be used when setting or getting data from a given `memcached` instance. When you get or set a value, a hash is constructed from the supplied key and then used to select a host from the list of configured servers. Because the hashing mechanism uses the supplied key as the basis for the hash, the same server is selected during both set and get operations.

You can think of this process as follows. Given an array of servers (a, b, and c), the client uses a hashing algorithm that returns an integer based on the key being stored or retrieved. The resulting value is then used to select a server from the list of servers configured in the client. Most standard client hashing within `memcache` clients uses a simple modulus calculation on the value against the number of configured `memcached` servers. You can summarize the process in pseudocode as:

```
@memcservers = ['a.memc', 'b.memc', 'c.memc'];
$value = hash($key);
$chosen = $value % length(@memcservers);
```

Replacing the above with values:

```
@memcservers = ['a.memc', 'b.memc', 'c.memc'];
$value = hash('myid');
$chosen = 7009 % 3;
```

In the above example, the client hashing algorithm chooses the server at index 1 ( $7009 \% 3 = 1$ ), and store or retrieve the key and value with that server.




#### Note

This selection and hashing process is handled automatically by the `memcached` client you are using; you need only provide the list of `memcached` servers to use.

You can see a graphical representation of this below in [Figure 15.5, “memcached Hash Selection”](#).

**Figure 15.5. memcached Hash Selection**



The same hashing and selection process takes place during any operation on the specified key within the `memcached` client.

Using this method provides a number of advantages:

- The hashing and selection of the server to contact is handled entirely within the client. This eliminates the need to perform network communication to determine the right machine to contact.
- Because the determination of the `memcached` server occurs entirely within the client, the server can be selected automatically regardless of the operation being executed (set, get, increment, etc.).



- Because the determination is handled within the client, the hashing algorithm returns the same value for a given key; values are not affected or reset by differences in the server environment.
- Selection is very fast. The hashing algorithm on the key value is quick and the resulting selection of the server is from a simple array of available machines.
- Using client-side hashing simplifies the distribution of data over each `memcached` server. Natural distribution of the values returned by the hashing algorithm means that keys are automatically spread over the available servers.

Providing that the list of servers configured within the client remains the same, the same stored key returns the same value, and therefore selects the same server.

However, if you do not use the same hashing mechanism then the same data may be recorded on different servers by different interfaces, both wasting space on your `memcached` and leading to potential differences in the information.



#### Note

One way to use a multi-interface compatible hashing mechanism is to use the `libmemcached` library and the associated interfaces. Because the interfaces for the different languages (including C, Ruby, Perl and Python) use the same client library interface, they always generate the same hash code from the ID.

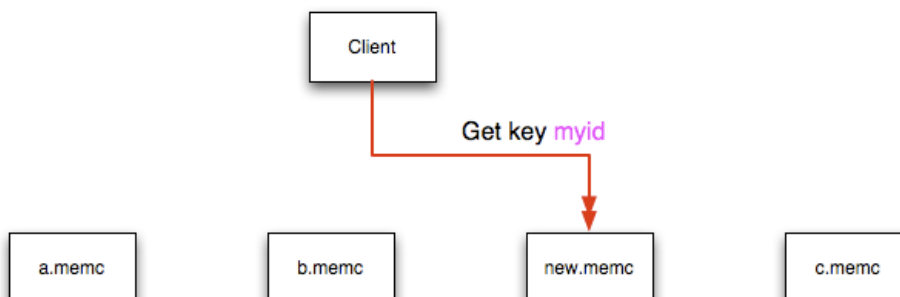
The problem with client-side selection of the server is that the list of the servers (including their sequential order) *must* remain consistent on each client using the `memcached` servers, and the servers must be available. If you try to perform an operation on a key when:

- A new `memcached` instance has been added to the list of available instances
- A `memcached` instance has been removed from the list of available instances
- The order of the `memcached` instances has changed

When the hashing algorithm is used on the given key, but with a different list of servers, the hash calculation may choose a different server from the list.

If a new `memcached` instance is added into the list of servers, as `new.memc` is in the example below, then a GET operation using the same key, `myid`, can result in a cache-miss. This is because the same value is computed from the key, which selects the same index from the array of servers, but index 2 now points to the new server, not the server `c.memc` where the data was originally stored. This would result in a cache miss, even though the key exists within the cache on another `memcached` instance.

**Figure 15.6.** `memcached` Hash Selection with New `memcached` instance



This means that servers `c.memc` and `new.memc` both contain the information for key `myid`, but the information stored against the key in each server may be different in each instance. A more significant problem is a much higher number of cache-misses when retrieving data, as the addition of a new

---

server changes the distribution of keys, and this in turn requires rebuilding the cached data on the `memcached` instances, causing an increase in database reads.

The same effect can occur if you actively manage the list of servers configured in your clients, adding and removing the configured `memcached` instances as each instance is identified as being available. For example, removing a `memcached` instance when the client notices that the instance can no longer be contacted can cause the server selection to fail as described here.

To prevent this causing significant problems and invalidating your cache, you can select the hashing algorithm used to select the server. There are two common types of hashing algorithm, *consistent* and *modula*.

With *consistent* hashing algorithms, the same key when applied to a list of servers always uses the same server to store or retrieve the keys, even if the list of configured servers changes. This means that you can add and remove servers from the configure list and always use the same server for a given key. There are two types of consistent hashing algorithms available, Ketama and Wheel. Both types are supported by `libmemcached`, and implementations are available for PHP and Java.

Any consistent hashing algorithm has some limitations. When you add servers to an existing list of configured servers, keys are distributed to the new servers as part of the normal distribution. When you remove servers from the list, the keys are re-allocated to another server within the list, meaning that the cache needs to be re-populated with the information. Also, a consistent hashing algorithm does not resolve the issue where you want consistent selection of a server across multiple clients, but where each client contains a different list of servers. The consistency is enforced only within a single client.

With a *modula* hashing algorithm, the client selects a server by first computing the hash and then choosing a server from the list of configured servers. As the list of servers changes, so the server selected when using a modula hashing algorithm also changes. The result is the behavior described above; changes to the list of servers mean that different servers are selected when retrieving data, leading to cache misses and increase in database load as the cache is re-seeded with information.

If you use only a single `memcached` instance for each client, or your list of `memcached` servers configured for a client never changes, then the selection of a hashing algorithm is irrelevant, as it has no noticeable effect.

If you change your servers regularly, or you use a common set of servers that are shared among a large number of clients, then using a consistent hashing algorithm should help to ensure that your cache data is not duplicated and the data is evenly distributed.

### 15.6.2.5. Using `memcached` and DTrace

`memcached` includes a number of different DTrace probes that can be used to monitor the operation of the server. The probes included can monitor individual connections, slab allocations, and modifications to the hash table when a key/value pair is added, updated, or removed.

For more information on DTrace and writing DTrace scripts, read the [DTrace User Guide](#).

Support for DTrace probes was added to `memcached` 1.2.6 includes a number of DTrace probes that can be used to help monitor your application. DTrace is supported on Solaris 10, OpenSolaris, Mac OS X 10.5 and FreeBSD. To enable the DTrace probes in `memcached`, build from source and use the `--enable-dtrace` option. For more information, see [Section 15.6.1, "Installing `memcached`"](#).

The probes supported by `memcached` are:

- `conn-allocate(connid)`

Fired when a connection object is allocated from the connection pool.

- `connid`: The connection ID.
- `conn-release(connid)`

Fired when a connection object is released back to the connection pool.

Arguments:

- `connid`: The connection ID.
- `conn-create(ptr)`

Fired when a new connection object is being created (that is, there are no free connection objects in the connection pool).

Arguments:

- `ptr`: A pointer to the connection. object
- `conn-destroy(ptr)`

Fired when a connection object is being destroyed.

Arguments:

- `ptr`: A pointer to the connection object.
- `conn-dispatch(connid, threadid)`

Fired when a connection is dispatched from the main or connection-management thread to a worker thread.

Arguments:

- `connid`: The connection ID.
- `threadid`: The thread ID.
- `slabs-allocate(size, slabclass, slabsize, ptr)`

Allocate memory from the slab allocator.

Arguments:

- `size`: The requested size.
- `slabclass`: The allocation is fulfilled in this class.
- `slabsize`: The size of each item in this class.
- `ptr`: A pointer to allocated memory.
- `slabs-allocate-failed(size, slabclass)`

Failed to allocate memory (out of memory).

Arguments:

- `size`: The requested size.
- `slabclass`: The class that failed to fulfill the request.
- `slabs-slabclass-allocate(slabclass)`

Fired when a slab class needs more space.

Arguments:

- `slabclass`: The class that needs more memory.

- `slabs-slabclass-allocate-failed(slabclass)`

Failed to allocate memory (out of memory).

Arguments:

- `slabclass`: The class that failed to grab more memory.

- `slabs-free(size, slabclass, ptr)`

Release memory.

Arguments:

- `size`: The amount of memory to release, in bytes.
- `slabclass`: The class the memory belongs to.
- `ptr`: A pointer to the memory to release.

- `assoc-find(key, depth)`

Fired when we have searched the hash table for a named key. These two elements provide an insight into how well the hash function operates. Traversals are a sign of a less optimal function, wasting CPU capacity.

Arguments:

- `key`: The key searched for.
- `depth`: The depth in the list of hash table.

- `assoc-insert(key, nokeys)`

Fired when a new item has been inserted.

Arguments:

- `key`: The key just inserted.
- `nokeys`: The total number of keys currently being stored, including the key for which insert was called.

- `assoc-delete(key, nokeys)`

Fired when a new item has been removed.

Arguments:

- `key`: The key just deleted.
- `nokeys`: The total number of keys currently being stored, excluding the key for which delete was called.

- `item-link(key, size)`

Fired when an item is being linked in the cache.

Arguments:

- `key`: The items key.

- `size`: The size of the data.

- `item-unlink(key, size)`

Fired when an item is being deleted.

Arguments:

- `key`: The items key.
- `size`: The size of the data.

- `item-remove(key, size)`

Fired when the refcount for an item is reduced.

Arguments:

- `key`: The item's key.
- `size`: The size of the data.

- `item-update(key, size)`

Fired when the "last referenced" time is updated.

Arguments:

- `key`: The item's key.
- `size`: The size of the data.

- `item-replace(oldkey, oldsize, newkey, newsize)`

Fired when an item is being replaced with another item.

Arguments:

- `oldkey`: The key of the item to replace.
- `oldsize`: The size of the old item.
- `newkey`: The key of the new item.
- `newsize`: The size of the new item.

- `process-command-start(connid, request, size)`

Fired when the processing of a command starts.

Arguments:

- `connid`: The connection ID.
- `request`: The incoming request.
- `size`: The size of the request.

- `process-command-end(connid, response, size)`

Fired when the processing of a command is done.

Arguments:

- `connid`: The connection ID.
- `response`: The response to send back to the client.
- `size`: The size of the response.
- `command-get(connid, key, size)`  
Fired for a `get` command.  
Arguments:
  - `connid`: The connection ID.
  - `key`: The requested key.
  - `size`: The size of the key's data (or -1 if not found).
- `command-gets(connid, key, size, casid)`  
Fired for a `gets` command.  
Arguments:
  - `connid`: The connection ID.
  - `key`: The requested key.
  - `size`: The size of the key's data (or -1 if not found).
  - `casid`: The casid for the item.
- `command-add(connid, key, size)`  
Fired for a `add` command.  
Arguments:
  - `connid`: The connection ID.
  - `key`: The requested key.
  - `size`: The new size of the key's data (or -1 if not found).
- `command-set(connid, key, size)`  
Fired for a `set` command.  
Arguments:
  - `connid`: The connection ID.
  - `key`: The requested key.
  - `size`: The new size of the key's data (or -1 if not found).
- `command-replace(connid, key, size)`  
Fired for a `replace` command.  
Arguments:
  - `connid`: The connection ID.

- `key`: The requested key.
- `size`: The new size of the key's data (or -1 if not found).

- `command-prepend(connid, key, size)`

Fired for a `prepend` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.
- `size`: The new size of the key's data (or -1 if not found).

- `command-append(connid, key, size)`

Fired for a `append` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.
- `size`: The new size of the key's data (or -1 if not found).

- `command-cas(connid, key, size, casid)`

Fired for a `cas` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.
- `size`: The size of the key's data (or -1 if not found).
- `casid`: The cas ID requested.

- `command-incr(connid, key, val)`

Fired for `incr` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.
- `val`: The new value.

- `command-decr(connid, key, val)`

Fired for `decr` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.



- `val`: The new value.
- `command-delete(connid, key, exptime)`

Fired for a `delete` command.

Arguments:

- `connid`: The connection ID.
- `key`: The requested key.
- `exptime`: The expiry time.

### 15.6.2.6. Memory Allocation within memcached

When you first start `memcached`, the memory that you have configured is not automatically allocated. Instead, `memcached` only starts allocating and reserving physical memory once you start saving information into the cache.

When you start to store data into the cache, `memcached` does not allocate the memory for the data on an item by item basis. Instead, a slab allocation is used to optimize memory usage and prevent memory fragmentation when information expires from the cache.

With slab allocation, memory is reserved in blocks of 1MB. The slab is divided up into a number of blocks of equal size. When you try to store a value into the cache, `memcached` checks the size of the value that you are adding to the cache and determines which slab contains the right size allocation for the item. If a slab with the item size already exists, the item is written to the block within the slab.

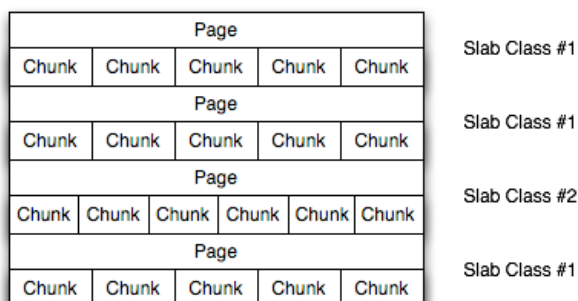
If the new item is bigger than the size of any existing blocks, then a new slab is created, divided up into blocks of a suitable size. If an existing slab with the right block size already exists, but there are no free blocks, a new slab is created. If you update an existing item with data that is larger than the existing block allocation for that key, then the key is re-allocated into a suitable slab.

For example, the default size for the smallest block is 88 bytes (40 bytes of value, and the default 48 bytes for the key and flag data). If the size of the first item you store into the cache is less than 40 bytes, then a slab with a block size of 88 bytes is created and the value stored.

If the size of the data that you intend to store is larger than this value, then the block size is increased by the chunk size factor until a block size large enough to hold the value is determined. The block size is always a function of the scale factor, rounded up to a block size which is exactly divisible into the chunk size.

For a sample of the structure, see [Figure 15.7, “Memory Allocation in memcached”](#).

**Figure 15.7. Memory Allocation in memcached**



The result is that you have multiple pages allocated within the range of memory allocated to `memcached`. Each page is 1MB in size (by default), and is split into a different number of chunks, according to the chunk size required to store the key/value pairs. Each instance has multiple pages

---

allocated, and a page is always created when a new item needs to be created requiring a chunk of a particular size. A slab may consist of multiple pages, and each page within a slab contains an equal number of chunks.

The chunk size of a new slab is determined by the base chunk size combined with the chunk size growth factor. For example, if the initial chunks are 104 bytes in size, and the default chunk size growth factor is used (1.25), then the next chunk size allocated would be the best power of 2 fit for  $104 \times 1.25$ , or 136 bytes.

Allocating the pages in this way ensures that memory does not get fragmented. However, depending on the distribution of the objects that you store, it may lead to an inefficient distribution of the slabs and chunks if you have significantly different sized items. For example, having a relatively small number of items within each chunk size may waste a lot of memory with just few chunks in each allocated page.

You can tune the growth factor to reduce this effect by using the `-f` command line option, which adapts the growth factor applied to make more effective use of the chunks and slabs allocated. For information on how to determine the current slab allocation statistics, see [Section 15.6.4.2, “memcached Slabs Statistics”](#).

If your operating system supports it, you can also start `memcached` with the `-L` command line option. This option preallocates all the memory during startup using large memory pages. This can improve performance by reducing the number of misses in the CPU memory cache.

### 15.6.2.7. `memcached` Thread Support

If you enable the thread implementation within when building `memcached` from source, then `memcached` uses multiple threads in addition to the `libevent` system to handle requests.

When enabled, the threading implementation operates as follows:

- Threading is handled by wrapping functions within the code to provide basic protection from updating the same global structures at the same time.
- Each thread uses its own instance of the `libevent` to help improve performance.
- TCP/IP connections are handled with a single thread listening on the TCP/IP socket. Each connection is then distributed to one of the active threads on a simple round-robin basis. Each connection then operates solely within this thread while the connection remains open.
- For UDP connections, all the threads listen to a single UDP socket for incoming requests. Threads that are not currently dealing with another request ignore the incoming packet. One of the remaining, nonbusy, threads reads the request and sends the response. This implementation can lead to increased CPU load as threads wake from sleep to potentially process the request.

Using threads can increase the performance on servers that have multiple CPU cores available, as the requests to update the hash table can be spread between the individual threads. To minimize overhead from the locking mechanism employed, experiment with different thread values to achieve the best performance based on the number and type of requests within your given workload.

### 15.6.2.8. `memcached` Logs

If you enable verbose mode, using the `-v`, `-vv`, or `-vvv` options, then the information output by `memcached` includes details of the operations being performed.

Without the verbose options, `memcached` normally produces no output during normal operating.

- **Output when using `-v`**

The lowest verbosity level shows you:

- Errors and warnings
- Transient errors

- Protocol and socket errors, including exhausting available connections
- Each registered client connection, including the socket descriptor number and the protocol used.

For example:

```
32: Client using the ascii protocol
33: Client using the ascii protocol
```

Note that the socket descriptor is only valid while the client remains connected. Non-persistent connections may not be effectively represented.

Examples of the error messages output at this level include:

```
<%d send buffer was %d, now %d
Can't listen for events on fd %d
Can't read from libevent pipe
Catastrophic: event fd doesn't match conn fd!
Couldn't build response
Couldn't realloc input buffer
Couldn't update event
Failed to build UDP headers
Failed to read, and not due to blocking
Too many open connections
Unexpected state %d
```

- **Output when using `-vv`**

When using the second level of verbosity, you get more detailed information about protocol operations, keys updated, chunk and network operations and details.

During the initial start-up of `memcached` with this level of verbosity, you are shown the sizes of the individual slab classes, the chunk sizes, and the number of entries per slab. These do not show the allocation of the slabs, just the slabs that would be created when data is added. You are also given information about the listen queues and buffers used to send information. A sample of the output generated for a TCP/IP based system with the default memory and growth factors is given below:

```
shell> memcached -vv
slab class 1: chunk size 80 per slab 13107
slab class 2: chunk size 104 per slab 10082
slab class 3: chunk size 136 per slab 7710
slab class 4: chunk size 176 per slab 5957
slab class 5: chunk size 224 per slab 4681
slab class 6: chunk size 280 per slab 3744
slab class 7: chunk size 352 per slab 2978
slab class 8: chunk size 440 per slab 2383
slab class 9: chunk size 552 per slab 1899
slab class 10: chunk size 696 per slab 1506
slab class 11: chunk size 872 per slab 1202
slab class 12: chunk size 1096 per slab 956
slab class 13: chunk size 1376 per slab 762
slab class 14: chunk size 1720 per slab 609
slab class 15: chunk size 2152 per slab 487
slab class 16: chunk size 2696 per slab 388
slab class 17: chunk size 3376 per slab 310
slab class 18: chunk size 4224 per slab 248
slab class 19: chunk size 5280 per slab 198
slab class 20: chunk size 6600 per slab 158
slab class 21: chunk size 8256 per slab 127
slab class 22: chunk size 10320 per slab 101
slab class 23: chunk size 12904 per slab 81
slab class 24: chunk size 16136 per slab 64
slab class 25: chunk size 20176 per slab 51
slab class 26: chunk size 25224 per slab 41
slab class 27: chunk size 31536 per slab 33
slab class 28: chunk size 39424 per slab 26
slab class 29: chunk size 49280 per slab 21
slab class 30: chunk size 61600 per slab 17
```

```

slab class 31: chunk size 77000 perslab 13
slab class 32: chunk size 96256 perslab 10
slab class 33: chunk size 120320 perslab 8
slab class 34: chunk size 150400 perslab 6
slab class 35: chunk size 188000 perslab 5
slab class 36: chunk size 235000 perslab 4
slab class 37: chunk size 293752 perslab 3
slab class 38: chunk size 367192 perslab 2
slab class 39: chunk size 458992 perslab 2
<26 server listening (auto-negotiate)
<29 server listening (auto-negotiate)
<30 send buffer was 57344, now 2097152
<31 send buffer was 57344, now 2097152
<30 server listening (udp)
<30 server listening (udp)
<31 server listening (udp)
<30 server listening (udp)
<30 server listening (udp)
<31 server listening (udp)
<31 server listening (udp)
<31 server listening (udp)

```

Using this verbosity level can be a useful way to check the effects of the growth factor used on slabs with different memory allocations, which in turn can be used to better tune the growth factor to suit the data you are storing in the cache. For example, if you set the growth factor to 4 (quadrupling the size of each slab):

```

shell> memcached -f 4 -m 1g -vv
slab class 1: chunk size 80 perslab 13107
slab class 2: chunk size 320 perslab 3276
slab class 3: chunk size 1280 perslab 819
slab class 4: chunk size 5120 perslab 204
slab class 5: chunk size 20480 perslab 51
slab class 6: chunk size 81920 perslab 12
slab class 7: chunk size 327680 perslab 3
...

```

During use of the cache, this verbosity level also prints out detailed information on the storage and recovery of keys and other information. An example of the output during a typical set/get and increment/decrement operation is shown below.

```

32: Client using the ascii protocol
<32 set my_key 0 0 10
>32 STORED
<32 set object_key 1 0 36
>32 STORED
<32 get my_key
>32 sending key my_key
>32 END
<32 get object_key
>32 sending key object_key
>32 END
<32 set key 0 0 6
>32 STORED
<32 incr key 1
>32 789544
<32 decr key 1
>32 789543
<32 incr key 2
>32 789545
<32 set my_key 0 0 10
>32 STORED
<32 set object_key 1 0 36
>32 STORED
<32 get my_key
>32 sending key my_key
>32 END
<32 get object_key
>32 sending key object_key1 1 36

```

```
>32 END
<32 set key 0 0 6
>32 STORED
<32 incr key 1
>32 789544
<32 decr key 1
>32 789543
<32 incr key 2
>32 789545
```

During client communication, for each line, the initial character shows the direction of flow of the information. The < for communication from the client to the `memcached` server and > for communication back to the client. The number is the numeric socket descriptor for the connection.

- **Output when using `-vvv`**

This level of verbosity includes the transitions of connections between different states in the event library while reading and writing content to/from the clients. It should be used to diagnose and identify issues in client communication. For example, you can use this information to determine if `memcached` is taking a long time to return information to the client, during the read of the client operation or before returning and completing the operation. An example of the typical sequence for a set operation is provided below:

```
<32 new auto-negotiating client connection
32: going from conn_new_cmd to conn_waiting
32: going from conn_waiting to conn_read
32: going from conn_read to conn_parse_cmd
32: Client using the ascii protocol
<32 set my_key 0 0 10
32: going from conn_parse_cmd to conn_nread
> NOT FOUND my_key
>32 STORED
32: going from conn_nread to conn_write
32: going from conn_write to conn_new_cmd
32: going from conn_new_cmd to conn_waiting
32: going from conn_waiting to conn_read
32: going from conn_read to conn_closing
<32 connection closed.
```

All of the verbosity levels in `memcached` are designed to be used during debugging or examination of issues. The quantity of information generated, particularly when using `-vvv`, is significant, particularly on a busy server. Also be aware that writing the error information out, especially to disk, may negate some of the performance gains you achieve by using `memcached`. Therefore, use in production or deployment environments is not recommended.

### 15.6.3. Developing a `memcached` Application

A number of language interfaces let applications store and retrieve information with `memcached` servers. You can write `memcached` applications in popular languages such as Perl, PHP, Python, Ruby, C, and Java.

Data stored into a `memcached` server is referred to by a single string (the key), with storage into the cache and retrieval from the cache using the key as the reference. The cache therefore operates like a large associative array or hash table. It is not possible to structure or otherwise organize the information stored in the cache. To emulate database notions such as multiple tables or composite key values, you must use encode the extra information into the strings used as keys. For example, to store or look up the address corresponding to a specific latitude and longitude, you might turn those two numeric values into a single comma-separated string to use as a key.

#### 15.6.3.1. Basic `memcached` Operations

The interface to `memcached` supports the following methods for storing and retrieving information in the cache, and these are consistent across all the different APIs, although the language specific mechanics might be different:

- `get(key)`: Retrieves information from the cache. Returns the value associated with the key if the specified key exists. Returns `NULL`, `nil`, `undefined`, or the closest equivalent in the corresponding language, if the specified key does not exist.
- `set(key, , value, [, expiry])`: Sets the item associated with a key in the cache to the specified value. Note that this either updates an existing item if the key already exists, or adds a new key/value pair if the key doesn't exist. If the expiry time is specified, then the item expires (and is deleted) when the expiry time is reached. The time is specified in seconds, and is taken as a relative time if the value is less than 30 days (30\*24\*60\*60), or an absolute time (epoch) if larger than this value.
- `add(key, , value, [, expiry])`: Adds the key and associated value to the cache, if the specified key does not already exist.
- `replace(key, , value, [, expiry])`: Replaces the item associated with the specified `key`, only if the key already exists. The new value is given by the `value` parameter.
- `delete(key, [, time])`: Deletes the `key` and its associated item from the cache. If you supply a `time`, then adding another item with the specified `key` is blocked for the specified period.
- `incr(key, [, value])`: Increments the item associated with the `key` by one or the optional `value`.
- `decr(key, [, value])`: Decrements the item associated with the `key` by one or the optional `value`.
- `flush_all`: Invalidates (or expires) all the current items in the cache. Technically they still exist (they are not deleted), but they are silently destroyed the next time you try to access them.

In all implementations, most or all of these functions are duplicated through the corresponding native language interface.

When practical, use `memcached` to store full items, rather than caching a single column value from the database. For example, when displaying a record about an object (invoice, user history, or blog post), load all the data for the associated entry from the database, and compile it into the internal structure that would normally be required by the application. Save the complete object in the cache.

Complex data structures cannot be stored directly. Most interfaces serialize the data for you, that is, put it in a textual form that can reconstruct the original pointers and nesting. Perl uses `Storable`, PHP uses `serialize`, Python uses `cPickle` (or `Pickle`) and Java uses the `Serializable` interface. In most cases, the serialization interface used is customizable. To share data stored in `memcached` instances between different language interfaces, consider using a common serialization solution such as JSON (Javascript Object Notation).

### 15.6.3.2. Using `memcached` as a MySQL Caching Layer

When using `memcached` to cache MySQL data, your application must retrieve data from the database and load the appropriate key-value pairs into the cache. Then, subsequent lookups can be done directly from the cache.

Because MySQL has its own in-memory caching mechanisms for queried data, such as the `InnoDB buffer pool` and the MySQL query cache, look for opportunities beyond loading individual column values or rows into the cache. Prefer to cache composite values, such as those retrieved from multiple tables through a join query, or result sets assembled from multiple rows.



#### Caution

Limit the information in the cache to non-sensitive data, because there is no security required to access or update the information within a `memcached` instance. Anybody with access to the machine has the ability to read, view

and potentially update the information. To keep the data secure, encrypt the information before caching it. To restrict the users capable of connecting to the server, either disable network access, or use IPTables or similar techniques to restrict access to the memcached ports to a select set of hosts.

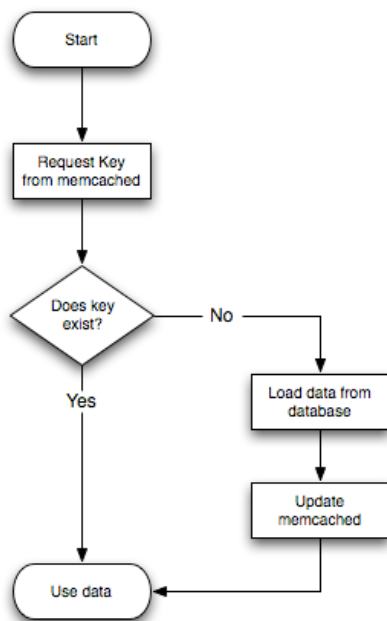
You can introduce memcached to an existing application, even if caching was not part of the original design. In many languages and environments the changes to the application will be just a few lines, first to attempt to read from the cache when loading data, fall back to the old method if the information is not cached, and to update the cache with information once the data has been read.

The general sequence for using memcached in any language as a caching solution for MySQL is as follows:

1. Request the item from the cache.
2. If the item exists, use the item data.
3. If the item does not exist, load the data from MySQL, and store the value into the cache. This means the value is available to the next client that requests it from the cache.

For a flow diagram of this sequence, see [Figure 15.8, “Typical memcached Application Flowchart”](#).

**Figure 15.8. Typical memcached Application Flowchart**



### Adapting Database Best Practices to memcached Applications

The most direct way to cache MySQL data is to use a 2-column table, where the first column is a **primary key**. Because of the uniqueness requirements for memcached keys, make sure your database schema makes appropriate use of primary keys and **unique constraints**.

If you combine multiple column values into a single memcached item value, choose data types to make it easy to parse the value back into its components, for example by using a separator character between numeric values.

The queries that map most easily to memcached lookups are those with a single **WHERE** clause, using an **=** or **IN** operator. For complicated **WHERE** clauses, or those using operators such as **<**, **>**, **BETWEEN**, or **LIKE**, memcached does not provide a simple or efficient way to scan through or filter the keys



or associated values, so typically you perform those operations as SQL queries on the underlying database.

### 15.6.3.3. Using `libmemcached` with C and C++

The `libmemcached` library provides both C and C++ interfaces to `memcached` and is also the basis for a number of different additional API implementations, including Perl, Python and Ruby. Understanding the core `libmemcached` functions can help when using these other interfaces.

The C library is the most comprehensive interface library for `memcached` and provides functions and operational systems not always exposed in interfaces not based on the `libmemcached` library.

The different functions can be divided up according to their basic operation. In addition to functions that interface to the core API, a number of utility functions provide extended functionality, such as appending and prepending data.

To build and install `libmemcached`, download the `libmemcached` package, run `configure`, and then build and install:

```
shell> tar xjf libmemcached-0.21.tar.gz
shell> cd libmemcached-0.21
shell> ./configure
shell> make
shell> make install
```

On many Linux operating systems, you can install the corresponding `libmemcached` package through the usual `yum`, `apt-get`, or similar commands.

To build an application that uses the library, first set the list of servers. Either directly manipulate the servers configured within the main `memcached_st` structure, or separately populate a list of servers, and then add this list to the `memcached_st` structure. The latter method is used in the following example. Once the server list has been set, you can call the functions to store or retrieve data. A simple application for setting a preset value to `localhost` is provided here:

```
#include <stdio.h>
#include <string.h>
#include <unistd.h>
#include <libmemcached/memcached.h>

int main(int argc, char *argv[])
{
    memcached_server_st *servers = NULL;
    memcached_st *memc;
    memcached_return rc;
    char *key= "keystring";
    char *value= "keyvalue";

    memcached_server_st *memcached_servers_parse (char *server_strings);
    memc= memcached_create(NULL);

    servers= memcached_server_list_append(servers, "localhost", 11211, &rc);
    rc= memcached_server_push(memc, servers);

    if (rc == MEMCACHED_SUCCESS)
        fprintf(stderr, "Added server successfully\n");
    else
        fprintf(stderr, "Couldn't add server: %s\n", memcached_strerror(memc, rc));

    rc= memcached_set(memc, key, strlen(key), value, strlen(value), (time_t)0, (uint32_t)0);

    if (rc == MEMCACHED_SUCCESS)
        fprintf(stderr, "Key stored successfully\n");
    else
        fprintf(stderr, "Couldn't store key: %s\n", memcached_strerror(memc, rc));

    return 0;
}
```

```
}

```

To test the success of an operation, use the return value, or populated result code, for a given function. The value is always set to `MEMCACHED_SUCCESS` if the operation succeeded. In the event of a failure, use the `memcached_strerror()` function to translate the result code into a printable string.

To build the application, specify the `memcached` library:

```
shell> gcc -o memc_basic memc_basic.c -lmemcached

```

Running the above sample application, after starting a `memcached` server, should return a success message:

```
shell> memc_basic
Added server successfully
Key stored successfully

```

### 15.6.3.3.1. `libmemcached` Base Functions

The base `libmemcached` functions let you create, destroy and clone the main `memcached_st` structure that is used to interface with the `memcached` servers. The main functions are defined below:

```
memcached_st *memcached_create (memcached_st *ptr);

```

Creates a new `memcached_st` structure for use with the other `libmemcached` API functions. You can supply an existing, static, `memcached_st` structure, or `NULL` to have a new structure allocated. Returns a pointer to the created structure, or `NULL` on failure.

```
void memcached_free (memcached_st *ptr);

```

Frees the structure and memory allocated to a previously created `memcached_st` structure.

```
memcached_st *memcached_clone(memcached_st *clone, memcached_st *source);

```

Clones an existing `memcached` structure from the specified `source`, copying the defaults and list of servers defined in the structure.

### 15.6.3.3.2. `libmemcached` Server Functions

The `libmemcached` API uses a list of servers, stored within the `memcached_server_st` structure, to act as the list of servers used by the rest of the functions. To use `memcached`, you first create the server list, and then apply the list of servers to a valid `libmemcached` object.

Because the list of servers, and the list of servers within an active `libmemcached` object can be manipulated separately, you can update and manage server lists while an active `libmemcached` interface is running.

The functions for manipulating the list of servers within a `memcached_st` structure are:

```
memcached_return
memcached_server_add (memcached_st *ptr,
                    char *hostname,
                    unsigned int port);

```

Adds a server, using the given `hostname` and `port` into the `memcached_st` structure given in `ptr`.

```
memcached_return
memcached_server_add_unix_socket (memcached_st *ptr,
                                char *socket);

```

Adds a Unix socket to the list of servers configured in the `memcached_st` structure.

```
unsigned int memcached_server_count (memcached_st *ptr);

```

Returns a count of the number of configured servers within the `memcached_st` structure.

```
memcached_server_st *
    memcached_server_list (memcached_st *ptr);
```

Returns an array of all the defined hosts within a `memcached_st` structure.

```
memcached_return
    memcached_server_push (memcached_st *ptr,
                          memcached_server_st *list);
```

Pushes an existing list of servers onto list of servers configured for a current `memcached_st` structure. This adds servers to the end of the existing list, and duplicates are not checked.

The `memcached_server_st` structure can be used to create a list of `memcached` servers which can then be applied individually to `memcached_st` structures.

```
memcached_server_st *
    memcached_server_list_append (memcached_server_st *ptr,
                                  char *hostname,
                                  unsigned int port,
                                  memcached_return *error);
```

Adds a server, with `hostname` and `port`, to the server list in `ptr`. The result code is handled by the `error` argument, which should point to an existing `memcached_return` variable. The function returns a pointer to the returned list.

```
unsigned int memcached_server_list_count (memcached_server_st *ptr);
```

Returns the number of the servers in the server list.

```
void memcached_server_list_free (memcached_server_st *ptr);
```

Frees the memory associated with a server list.

```
memcached_server_st *memcached_servers_parse (char *server_strings);
```

Parses a string containing a list of servers, where individual servers are separated by a comma, space, or both, and where individual servers are of the form `server[:port]`. The return value is a server list structure.

### 15.6.3.3.3. `libmemcached` Set Functions

The set-related functions within `libmemcached` provide the same functionality as the core functions supported by the `memcached` protocol. The full definition for the different functions is the same for all the base functions (`add`, `replace`, `prepend`, `append`). For example, the function definition for `memcached_set()` is:

```
memcached_return
    memcached_set (memcached_st *ptr,
                  const char *key,
                  size_t key_length,
                  const char *value,
                  size_t value_length,
                  time_t expiration,
                  uint32_t flags);
```

The `ptr` is the `memcached_st` structure. The `key` and `key_length` define the key name and length, and `value` and `value_length` the corresponding value and length. You can also set the expiration and optional flags. For more information, see [Section 15.6.3.3.5, “Controlling `libmemcached` Behaviors”](#).

This table outlines the remainder of the set-related `libmemcached` functions and the equivalent core functions supported by the `memcached` protocol.

| Libmemcached Function                                                                                                               | Equivalent Core Function                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>memcached_set(memc, key, key_length, value, value_length, expiration, flags)</code>                                           | Generic <code>set()</code> operation.                                                                                                               |
| <code>memcached_add(memc, key, key_length, value, value_length, expiration, flags)</code>                                           | Generic <code>add()</code> function.                                                                                                                |
| <code>memcached_replace(memc, key, key_length, value, value_length, expiration, flags)</code>                                       | Generic <code>replace()</code> .                                                                                                                    |
| <code>memcached_prepend(memc, key, key_length, value, value_length, expiration, flags)</code>                                       | Prepends the specified <code>value</code> before the current value of the specified <code>key</code> .                                              |
| <code>memcached_append(memc, key, key_length, value, value_length, expiration, flags)</code>                                        | Appends the specified <code>value</code> after the current value of the specified <code>key</code> .                                                |
| <code>memcached_cas(memc, key, key_length, value, value_length, expiration, flags, cas)</code>                                      | Overwrites the data for a given key as long as the corresponding <code>cas</code> value is still the same within the server.                        |
| <code>memcached_set_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code>     | Similar to the generic <code>set()</code> , but has the option of an additional master key that can be used to identify an individual server.       |
| <code>memcached_add_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code>     | Similar to the generic <code>add()</code> , but has the option of an additional master key that can be used to identify an individual server.       |
| <code>memcached_replace_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code> | Similar to the generic <code>replace()</code> , but has the option of an additional master key that can be used to identify an individual server.   |
| <code>memcached_prepend_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code> | Similar to the <code>memcached_prepend()</code> , but has the option of an additional master key that can be used to identify an individual server. |
| <code>memcached_append_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code>  | Similar to the <code>memcached_append()</code> , but has the option of an additional master key that can be used to identify an individual server.  |
| <code>memcached_cas_by_key(memc, master_key, master_key_length, key, key_length, value, value_length, expiration, flags)</code>     | Similar to the <code>memcached_cas()</code> , but has the option of an additional master key that can be used to identify an individual server.     |

The `by_key` methods add two further arguments that define the master key, to be used and applied during the hashing stage for selecting the servers. You can see this in the following definition:

```
memcached_return
memcached_set_by_key(memcached_st *ptr,
                    const char *master_key,
                    size_t master_key_length,
                    const char *key,
                    size_t key_length,
                    const char *value,
                    size_t value_length,
```

```
time_t expiration,
uint32_t flags);
```

All the functions return a value of type `memcached_return`, which you can compare against the `MEMCACHED_SUCCESS` constant.

#### 15.6.3.3.4. `libmemcached` Get Functions

The `libmemcached` functions provide both direct access to a single item, and a multiple-key request mechanism that provides much faster responses when fetching a large number of keys simultaneously.

The main get-style function, which is equivalent to the generic `get()` is `memcached_get()`. This function returns a string pointer, pointing to the value associated with the specified key.

```
char *memcached_get (memcached_st *ptr,
                    const char *key, size_t key_length,
                    size_t *value_length,
                    uint32_t *flags,
                    memcached_return *error);
```

A multi-key get, `memcached_mget()`, is also available. Using a multiple key get operation is much quicker to do in one block than retrieving the key values with individual calls to `memcached_get()`. To start the multi-key get, call `memcached_mget()`:

```
memcached_return
memcached_mget (memcached_st *ptr,
               char **keys, size_t *key_length,
               unsigned int number_of_keys);
```

The return value is the success of the operation. The `keys` parameter should be an array of strings containing the keys, and `key_length` an array containing the length of each corresponding key. `number_of_keys` is the number of keys supplied in the array.

To fetch the individual values, use `memcached_fetch()` to get each corresponding value.

```
char *memcached_fetch (memcached_st *ptr,
                      const char *key, size_t *key_length,
                      size_t *value_length,
                      uint32_t *flags,
                      memcached_return *error);
```

The function returns the key value, with the `key`, `key_length` and `value_length` parameters being populated with the corresponding key and length information. The function returns `NULL` when there are no more values to be returned. A full example, including the populating of the key data and the return of the information is provided here.

```
#include <stdio.h>
#include <sstring.h>
#include <unistd.h>
#include <libmemcached/memcached.h>

int main(int argc, char *argv[])
{
    memcached_server_st *servers = NULL;
    memcached_st *memc;
    memcached_return rc;
    char *keys[] = {"huey", "dewey", "louie"};
    size_t key_length[3];
    char *values[] = {"red", "blue", "green"};
    size_t value_length[3];
    unsigned int x;
    uint32_t flags;

    char return_key[MEMCACHED_MAX_KEY];
    size_t return_key_length;
    char *return_value;
    size_t return_value_length;
```

```

memc= memcached_create(NULL);

servers= memcached_server_list_append(servers, "localhost", 11211, &rc);
rc= memcached_server_push(memc, servers);

if (rc == MEMCACHED_SUCCESS)
    fprintf(stderr, "Added server successfully\n");
else
    fprintf(stderr, "Couldn't add server: %s\n", memcached_strerror(memc, rc));

for(x= 0; x < 3; x++)
    {
        key_length[x] = strlen(keys[x]);
        value_length[x] = strlen(values[x]);

        rc= memcached_set(memc, keys[x], key_length[x], values[x],
                           value_length[x], (time_t)0, (uint32_t)0);
        if (rc == MEMCACHED_SUCCESS)
            fprintf(stderr, "Key %s stored successfully\n", keys[x]);
        else
            fprintf(stderr, "Couldn't store key: %s\n", memcached_strerror(memc, rc));
    }

rc= memcached_mget(memc, keys, key_length, 3);

if (rc == MEMCACHED_SUCCESS)
    {
        while ((return_value= memcached_fetch(memc, return_key, &return_key_length,
   &return_value_length, &flags, &rc)) != NULL)
            {
                if (rc == MEMCACHED_SUCCESS)
                    {
                        fprintf(stderr, "Key %s returned %s\n", return_key, return_value);
                    }
            }
    }

return 0;
}

```

Running the above application produces the following output:

```

shell> memc_multi_fetch
Added server successfully
Key huey stored successfully
Key dewey stored successfully
Key louie stored successfully
Key huey returned red
Key dewey returned blue
Key louie returned green

```

### 15.6.3.3.5. Controlling libmemcached Behaviors

The behavior of `libmemcached` can be modified by setting one or more behavior flags. These can either be set globally, or they can be applied during the call to individual functions. Some behaviors also accept an additional setting, such as the hashing mechanism used when selecting servers.

To set global behaviors:

```

memcached_return
memcached_behavior_set (memcached_st *ptr,
                        memcached_behavior flag,
                        uint64_t data);

```

To get the current behavior setting:

```

uint64_t
memcached_behavior_get (memcached_st *ptr,
                        memcached_behavior flag);

```

The following table describes libmemcached behavior flags.

| Behavior                           | Description                                                                                                                                                                                                                                                                                                        |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MEMCACHED_BEHAVIOR_NO_BLOCK        | Caused libmemcached to use asynchronous I/O.                                                                                                                                                                                                                                                                       |
| MEMCACHED_BEHAVIOR_TCP_NODELAY     | Turns on no-delay for network sockets.                                                                                                                                                                                                                                                                             |
| MEMCACHED_BEHAVIOR_HASH            | Without a value, sets the default hashing algorithm for keys to use MD5. Other valid values include MEMCACHED_HASH_DEFAULT, MEMCACHED_HASH_MD5, MEMCACHED_HASH_CRC, MEMCACHED_HASH_FNV1_64, MEMCACHED_HASH_FNV1A_64, MEMCACHED_HASH_FNV1_32, and MEMCACHED_HASH_FNV1A_32.                                          |
| MEMCACHED_BEHAVIOR_DISTRIBUTION    | Changes the method of selecting the server used to store a given value. The default method is MEMCACHED_DISTRIBUTION_MODULA. You can enable consistent hashing by setting MEMCACHED_DISTRIBUTION_CONSISTENT. MEMCACHED_DISTRIBUTION_CONSISTENT is an alias for the value MEMCACHED_DISTRIBUTION_CONSISTENT_KETAMA. |
| MEMCACHED_BEHAVIOR_CACHE_LOOKUPS   | Cache the lookups made to the DNS service. This can improve the performance if you are using names instead of IP addresses for individual hosts.                                                                                                                                                                   |
| MEMCACHED_BEHAVIOR_SUPPORT_CAS     | Support CAS operations. By default, this is disabled because it imposes a performance penalty.                                                                                                                                                                                                                     |
| MEMCACHED_BEHAVIOR_KETAMA          | Sets the default distribution to MEMCACHED_DISTRIBUTION_CONSISTENT_KETAMA and the hash to MEMCACHED_HASH_MD5.                                                                                                                                                                                                      |
| MEMCACHED_BEHAVIOR_POLL_TIMEOUT    | Modify the timeout value used by poll(). Supply a signed int pointer for the timeout value.                                                                                                                                                                                                                        |
| MEMCACHED_BEHAVIOR_BUFFER_REQUESTS | Buffers IO requests instead of them being sent. A get operation, or closing the connection causes the data to be flushed.                                                                                                                                                                                          |
| MEMCACHED_BEHAVIOR_VERIFY_KEY      | Forces libmemcached to verify that a specified key is valid.                                                                                                                                                                                                                                                       |
| MEMCACHED_BEHAVIOR_SORT_HOSTS      | If set, hosts added to the list of configured hosts for a memcached_st structure are placed into the host list in sorted order. This breaks consistent hashing if that behavior has been enabled.                                                                                                                  |
| MEMCACHED_BEHAVIOR_CONNECT_TIMEOUT | In nonblocking mode this changes the value of the timeout during socket connection.                                                                                                                                                                                                                                |

### 15.6.3.3.6. libmemcached Command-Line Utilities

In addition to the main C library interface, libmemcached also includes a number of command-line utilities that can be useful when working with and debugging memcached applications.

All of the command-line tools accept a number of arguments, the most critical of which is servers, which specifies the list of servers to connect to when returning information.

The main tools are:

- memcat: Display the value for each ID given on the command line:

```
shell> memcat --servers=localhost hwkey
```



```
Hello world
```

- `memcp`: Copy the contents of a file into the cache, using the file name as the key:

```
shell> echo "Hello World" > hwkey
shell> memcp --servers=localhost hwkey
shell> memcat --servers=localhost hwkey
Hello world
```

- `memrm`: Remove an item from the cache:

```
shell> memcat --servers=localhost hwkey
Hello world
shell> memrm --servers=localhost hwkey
shell> memcat --servers=localhost hwkey
```

- `memslap`: Test the load on one or more `memcached` servers, simulating get/set and multiple client operations. For example, you can simulate the load of 100 clients performing get operations:

```
shell> memslap --servers=localhost --concurrency=100 --flush --test=get
memslap --servers=localhost --concurrency=100 --flush --test=get Threads connecting to servers 100
Took 13.571 seconds to read data
```

- `memflush`: Flush (empty) the contents of the `memcached` cache.

```
shell> memflush --servers=localhost
```

### 15.6.3.4. Using MySQL and `memcached` with Perl

The `Cache::Memcached` module provides a native interface to the Memcache protocol, and provides support for the core functions offered by `memcached`. Install the module using your operating system's package management system, or using `CPAN`:

```
root-shell> perl -MCPAN -e 'install Cache::Memcached'
```

To use `memcached` from Perl through the `Cache::Memcached` module, first create a new `Cache::Memcached` object that defines the list of servers and other parameters for the connection. The only argument is a hash containing the options for the cache interface. For example, to create a new instance that uses three `memcached` servers:

```
use Cache::Memcached;

my $cache = new Cache::Memcached {
    'servers' => [
        '192.168.0.100:11211',
        '192.168.0.101:11211',
        '192.168.0.102:11211',
    ],
};
```



#### Note

When using the `Cache::Memcached` interface with multiple servers, the API automatically performs certain operations across all the servers in the group. For example, getting statistical information through `Cache::Memcached` returns a hash that contains data on a host-by-host basis, as well as generalized statistics for all the servers in the group.

You can set additional properties on the cache object instance when it is created by specifying the option as part of the option hash. Alternatively, you can use a corresponding method on the instance:

- `servers` or method `set_servers()`: Specifies the list of the servers to be used. The servers list should be a reference to an array of servers, with each element as the address and port number combination (separated by a colon). You can also specify a local connection through a UNIX socket (for example `/tmp/sock/memcached`). To specify the server with a weight (indicating how much

more frequently the server should be used during hashing), specify an array reference with the `memcached` server instance and a weight number. Higher numbers give higher priority.

- `compress_threshold` or method `set_compress_threshold()`: Specifies the threshold when values are compressed. Values larger than the specified number are automatically compressed (using `zlib`) during storage and retrieval.
- `no_rehash` or method `set_norehash()`: Disables finding a new server if the original choice is unavailable.
- `readonly` or method `set_readonly()`: Disables writes to the `memcached` servers.

Once the `Cache::Memcached` object instance has been configured, you can use the `set()` and `get()` methods to store and retrieve information from the `memcached` servers. Objects stored in the cache are automatically serialized and deserialized using the `Storable` module.

The `Cache::Memcached` interface supports the following methods for storing/retrieving data, and relate to the generic methods as shown in the table.

| <code>Cache::Memcached</code> Function | Equivalent Generic Method                                                                                        |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <code>get()</code>                     | Generic <code>get()</code> .                                                                                     |
| <code>get_multi(keys)</code>           | Gets multiple <code>keys</code> from memcache using just one query. Returns a hash reference of key/value pairs. |
| <code>set()</code>                     | Generic <code>set()</code> .                                                                                     |
| <code>add()</code>                     | Generic <code>add()</code> .                                                                                     |
| <code>replace()</code>                 | Generic <code>replace()</code> .                                                                                 |
| <code>delete()</code>                  | Generic <code>delete()</code> .                                                                                  |
| <code>incr()</code>                    | Generic <code>incr()</code> .                                                                                    |
| <code>decr()</code>                    | Generic <code>decr()</code> .                                                                                    |

Below is a complete example for using `memcached` with Perl and the `Cache::Memcached` module:

```
#!/usr/bin/perl

use Cache::Memcached;
use DBI;
use Data::Dumper;

# Configure the memcached server

my $cache = new Cache::Memcached {
    'servers' => [
        'localhost:11211',
    ],
};

# Get the film name from the command line
# memcached keys must not contain spaces, so create
# a key name by replacing spaces with underscores

my $filmname = shift or die "Must specify the film name\n";
my $filmkey = $filmname;
$filmkey =~ s/ /_/;

# Load the data from the cache

my $filmdata = $cache->get($filmkey);

# If the data wasn't in the cache, then we load it from the database

if (!defined($filmdata))
```

```

{
    $filmdata = load_filmdata($filename);

    if (defined($filmdata))
    {
# Set the data into the cache, using the key

        if ($cache->set($filmkey,$filmdata))
            {
                print STDERR "Film data loaded from database and cached\n";
            }
        else
            {
                print STDERR "Couldn't store to cache\n";
            }
    }
    else
    {
        die "Couldn't find $filename\n";
    }
}
else
{
    print STDERR "Film data loaded from Memcached\n";
}

sub load_filmdata
{
    my ($filename) = @_ ;

    my $dsn = "DBI:mysql:database=sakila;host=localhost;port=3306";

    $dbh = DBI->connect($dsn, 'sakila','password');

    my ($filmbase) = $dbh->selectrow_hashref(sprintf('select * from film where title = %s',
  $dbh->quote($filename)));

    if (!defined($filename))
    {
        return (undef);
    }

    $filmbase->{stars} =
    $dbh->selectall_arrayref(sprintf('select concat(first_name," ",last_name) ' .
                                    'from film_actor left join (actor) ' .
                                    'on (film_actor.actor_id = actor.actor_id) ' .
                                    ' where film_id=%s',
                                    $dbh->quote($filmbase->{film_id})));

    return($filmbase);
}

```

The example uses the Sakila database, obtaining film data from the database and writing a composite record of the film and actors to [memcached](#). When calling it for a film does not exist, you get this result:

```

shell> memcached-sakila.pl "ROCK INSTINCT"
Film data loaded from database and cached

```

When accessing a film that has already been added to the cache:

```

shell> memcached-sakila.pl "ROCK INSTINCT"
Film data loaded from Memcached

```

### 15.6.3.5. Using MySQL and [memcached](#) with Python

The Python [memcache](#) module interfaces to [memcached](#) servers, and is written in pure Python (that is, without using one of the C APIs). You can download and install a copy from [Python Memcached](#).

To install, download the package and then run the Python installer:

```
python setup.py install
running install
running bdist_egg
running egg_info
creating python_memcached.egg-info
...
removing 'build/bdist.linux-x86_64/egg' (and everything under it)
Processing python_memcached-1.43-py2.4.egg
creating /usr/lib64/python2.4/site-packages/python_memcached-1.43-py2.4.egg
Extracting python_memcached-1.43-py2.4.egg to /usr/lib64/python2.4/site-packages
Adding python-memcached 1.43 to easy-install.pth file

Installed /usr/lib64/python2.4/site-packages/python_memcached-1.43-py2.4.egg
Processing dependencies for python-memcached==1.43
Finished processing dependencies for python-memcached==1.43
```

Once installed, the `memcache` module provides a class-based interface to your `memcached` servers. When you store Python data structures as `memcached` items, they are automatically serialized (turned into string values) using the Python `cPickle` or `pickle` modules.

To create a new `memcache` interface, import the `memcache` module and create a new instance of the `memcache.Client` class. For example, if the `memcached` daemon is running on localhost using the default port:

```
import memcache
memc = memcache.Client(['127.0.0.1:11211'])
```

The first argument is an array of strings containing the server and port number for each `memcached` instance to use. To enable debugging, set the optional `debug` parameter to 1.

By default, the hashing mechanism used to divide the items among multiple servers is `crc32`. To change the function used, set the value of `memcache.serverHashFunction` to the alternate function to use. For example:

```
from zlib import Adler32
memcache.serverHashFunction = Adler32
```

Once you have defined the servers to use within the `memcache` instance, the core functions provide the same functionality as in the generic interface specification. The following table provides a summary of the supported functions:

| Python <code>memcache</code> Function                      | Equivalent Generic Function                                                                                      |
|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <code>get()</code>                                         | Generic <code>get()</code> .                                                                                     |
| <code>get_multi(keys)</code>                               | Gets multiple values from the supplied array of <code>keys</code> . Returns a hash reference of key/value pairs. |
| <code>set()</code>                                         | Generic <code>set()</code> .                                                                                     |
| <code>set_multi(dict [, expiry [, key_prefix]])</code>     | Sets multiple key/value pairs from the supplied <code>dict</code> .                                              |
| <code>add()</code>                                         | Generic <code>add()</code> .                                                                                     |
| <code>replace()</code>                                     | Generic <code>replace()</code> .                                                                                 |
| <code>prepend(key, value [, expiry])</code>                | Prepends the supplied <code>value</code> to the value of the existing <code>key</code> .                         |
| <code>append(key, value [, expiry])</code>                 | Appends the supplied <code>value</code> to the value of the existing <code>key</code> .                          |
| <code>delete()</code>                                      | Generic <code>delete()</code> .                                                                                  |
| <code>delete_multi(keys [, expiry [, key_prefix]] )</code> | Deletes all the keys from the hash matching each string in the array <code>keys</code> .                         |

| Python <code>memcache</code> Function | Equivalent Generic Function   |
|---------------------------------------|-------------------------------|
| <code>incr()</code>                   | Generic <code>incr()</code> . |
| <code>decr()</code>                   | Generic <code>decr()</code> . |



### Note

Within the Python `memcache` module, all the `*_multi()` functions support an optional `key_prefix` parameter. If supplied, then the string is used as a prefix to all key lookups. For example, if you call:

```
memc.get_multi(['a','b'], key_prefix='users:')
```

The function retrieves the keys `users:a` and `users:b` from the servers.

Here is an example showing the storage and retrieval of information to a `memcache` instance, loading the raw data from MySQL:

```
import sys
import MySQLdb
import memcache

memc = memcache.Client(['127.0.0.1:11211'], debug=1);

try:
    conn = MySQLdb.connect (host = "localhost",
                            user = "sakila",
                            passwd = "password",
                            db = "sakila")
except MySQLdb.Error, e:
    print "Error %d: %s" % (e.args[0], e.args[1])
    sys.exit (1)

popularfilms = memc.get('top5films')

if not popularfilms:
    cursor = conn.cursor()
    cursor.execute('select film_id,title from film order by rental_rate desc limit 5')
    rows = cursor.fetchall()
    memc.set('top5films',rows,60)
    print "Updated memcached with MySQL data"
else:
    print "Loaded data from memcached"
    for row in popularfilms:
        print "%s, %s" % (row[0], row[1])
```

When executed for the first time, the data is loaded from the MySQL database and stored to the `memcached` server.

```
shell> python memc_python.py
Updated memcached with MySQL data
```

Because the data is automatically serialized using `cPickle/pickle`, when you load the data back from `memcached`, you can use the object directly. In the example above, the information stored to `memcached` is in the form of rows from a Python DB cursor. When accessing the information (within the 60 second expiry time), the data is loaded from `memcached` and dumped:

```
shell> python memc_python.py
Loaded data from memcached
2, ACE GOLDFINGER
7, AIRPLANE SIERRA
8, AIRPORT POLLOCK
10, ALADDIN CALENDAR
13, ALI FOREVER
```

The serialization and deserialization happens automatically. Because serialization of Python data may be incompatible with other interfaces and languages, you can change the serialization module used

during initialization. For example, you might use JSON format when you store complex data structures using a script written in one language, and access them in a script written in a different language.

### 15.6.3.6. Using MySQL and `memcached` with PHP

PHP provides support for the Memcache functions through a PECL extension. To enable the PHP `memcache` extensions, build PHP using the `--enable-memcache` option to `configure` when building from source.

If you are installing on a Red Hat-based server, you can install the `php-pecl-memcache` RPM:

```
root-shell> yum --install php-pecl-memcache
```

On Debian-based distributions, use the `php-memcache` package.

To set global runtime configuration options, specify the configuration option values within your `php.ini` file. The following table provides the name, default value, and a description for each global runtime configuration option.

| Configuration option                        | Default  | Description                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>memcache.allow_failover</code>        | 1        | Specifies whether another server in the list should be queried if the first server selected fails.                                                                                                                                                                                                                |
| <code>memcache.max_failover_attempts</code> | 20       | Specifies the number of servers to try before returning a failure.                                                                                                                                                                                                                                                |
| <code>memcache.chunk_size</code>            | 8192     | Defines the size of network chunks used to exchange data with the <code>memcached</code> server.                                                                                                                                                                                                                  |
| <code>memcache.default_port</code>          | 11211    | Defines the default port to use when communicating with the <code>memcached</code> servers.                                                                                                                                                                                                                       |
| <code>memcache.hash_strategy</code>         | standard | Specifies which hash strategy to use. Set to <code>consistent</code> to enable servers to be added or removed from the pool without causing the keys to be remapped to other servers. When set to <code>standard</code> , an older (modula) strategy is used that potentially uses different servers for storage. |
| <code>memcache.hash_function</code>         | crc32    | Specifies which function to use when mapping keys to servers. <code>crc32</code> uses the standard CRC32 hash. <code>fnv</code> uses the FNV-1a hashing algorithm.                                                                                                                                                |

To create a connection to a `memcached` server, create a new `Memcache` object and then specify the connection options. For example:

```
<?php
$cache = new Memcache;
$cache->connect('localhost',11211);
?>
```

This opens an immediate connection to the specified server.

To use multiple `memcached` servers, you need to add servers to the `memcache` object using `addServer()`:

```
bool Memcache::addServer ( string $host [, int $port [, bool $persistent
                        [, int $weight [, int $timeout [, int $retry_interval
                        [, bool $status [, callback $failure_callback
```

```
]]]]]]]] )
```

The server management mechanism within the `php-memcache` module is a critical part of the interface as it controls the main interface to the `memcached` instances and how the different instances are selected through the hashing mechanism.

To create a simple connection to two `memcached` instances:

```
<?php
$cache = new Memcache;
$cache->addServer('192.168.0.100',11211);
$cache->addServer('192.168.0.101',11211);
?>
```

In this scenario, the instance connection is not explicitly opened, but only opened when you try to store or retrieve a value. To enable persistent connections to `memcached` instances, set the `$persistent` argument to true. This is the default setting, and causes the connections to remain open.

To help control the distribution of keys to different instances, use the global `memcache.hash_strategy` setting. This sets the hashing mechanism used to select. You can also add another weight to each server, which effectively increases the number of times the instance entry appears in the instance list, therefore increasing the likelihood of the instance being chosen over other instances. To set the weight, set the value of the `$weight` argument to more than one.

The functions for setting and retrieving information are identical to the generic functional interface offered by `memcached`, as shown in this table:

| PECL <code>memcache</code> Function | Generic Function                 |
|-------------------------------------|----------------------------------|
| <code>get()</code>                  | Generic <code>get()</code> .     |
| <code>set()</code>                  | Generic <code>set()</code> .     |
| <code>add()</code>                  | Generic <code>add()</code> .     |
| <code>replace()</code>              | Generic <code>replace()</code> . |
| <code>delete()</code>               | Generic <code>delete()</code> .  |
| <code>increment()</code>            | Generic <code>incr()</code> .    |
| <code>decrement()</code>            | Generic <code>decr()</code> .    |

A full example of the PECL `memcache` interface is provided below. The code loads film data from the Sakila database when the user provides a film name. The data stored into the `memcached` instance is recorded as a `mysqli` result row, and the API automatically serializes the information for you.

```
<?php
$memc = new Memcache;
$memc->addServer('localhost', '11211');

if(empty($_POST['film'])) {
?>
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">
<head>
<meta http-equiv="Content-Type" content="text/html; charset=utf-8" />
<title>Simple Memcache Lookup</title>
</head>
<body>
<form method="post">
<p><b>Film</b>: <input type="text" size="20" name="film"></p>
<input type="submit">
</form>
<hr/>
<?php
```



```

} else {

    echo "Loading data...\n";

    $film    = htmlspecialchars($_POST['film'], ENT_QUOTES, 'UTF-8');
    $mfilms = $memc->get($film);

    if ($mfilms) {

        printf("<p>Film data for %s loaded from memcache</p>", $mfilms['title']);

        foreach (array_keys($mfilms) as $key) {
            printf("<p><b>%s</b>: %s</p>", $key, $mfilms[$key]);
        }

    } else {

        $mysqli = mysqli('localhost','sakila','password','sakila');

        if (mysqli_connect_error()) {
            printf("Database error: (%d) %s", mysqli_connect_errno(), mysqli_connect_error());
            exit;
        }

        $sql = sprintf('SELECT * FROM film WHERE title="%s"', $mysqli->real_escape_string($film));

        $result = $mysqli->query($sql);

        if (!$result) {
            printf("Database error: (%d) %s", $mysqli->errno, $mysqli->error);
            exit;
        }

        $row = $result->fetch_assoc();

        $memc->set($row['title'], $row);

        printf("<p>Loaded (%s) from MySQL</p>", htmlspecialchars($row['title'], ENT_QUOTES, 'UTF-8'));
    }
}
?>
</body>
</html>

```

With PHP, the connections to the `memcached` instances are kept open as long as the PHP and associated Apache instance remain running. When adding or removing servers from the list in a running instance (for example, when starting another script that mentions additional servers), the connections are shared, but the script only selects among the instances explicitly configured within the script.

To ensure that changes to the server list within a script do not cause problems, make sure to use the consistent hashing mechanism.

### 15.6.3.7. Using MySQL and `memcached` with Ruby

There are a number of different modules for interfacing to `memcached` within Ruby. The `Ruby-MemCache` client library provides a native interface to `memcached` that does not require any external libraries, such as `libmemcached`. You can obtain the installer package from <http://www.deveiate.org/projects/RMemCache>.

To install, extract the package and then run `install.rb`:

```
shell> install.rb
```

If you have RubyGems, you can install the `Ruby-MemCache` gem:

```
shell> gem install Ruby-MemCache
Bulk updating Gem source index for: http://gems.rubyforge.org
```

```
Install required dependency io-reactor? [Yn] y
Successfully installed Ruby-MemCache-0.0.1
Successfully installed io-reactor-0.05
Installing ri documentation for io-reactor-0.05...
Installing RDoc documentation for io-reactor-0.05...
```

To use a `memcached` instance from within Ruby, create a new instance of the `MemCache` object.

```
require 'memcache'
memc = MemCache::new '192.168.0.100:11211'
```

You can add a weight to each server to increase the likelihood of the server being selected during hashing by appending the weight count to the server host name/port string:

```
require 'memcache'
memc = MemCache::new '192.168.0.100:11211:3'
```

To add servers to an existing list, you can append them directly to the `MemCache` object:

```
memc += ["192.168.0.101:11211"]
```

To set data into the cache, you can just assign a value to a key within the new cache object, which works just like a standard Ruby hash object:

```
memc["key"] = "value"
```

Or to retrieve the value:

```
print memc["key"]
```

For more explicit actions, you can use the method interface, which mimics the main `memcached` API functions, as summarized in the following table:

Ruby <code>MemCache</code> Method	Equivalent <code>memcached</code> API Functions
<code>get()</code>	Generic <code>get()</code> .
<code>get_hash(keys)</code>	Get the values of multiple <code>keys</code> , returning the information as a hash of the keys and their values.
<code>set()</code>	Generic <code>set()</code> .
<code>set_many(pairs)</code>	Set the values of the keys and values in the hash <code>pairs</code> .
<code>add()</code>	Generic <code>add()</code> .
<code>replace()</code>	Generic <code>replace()</code> .
<code>delete()</code>	Generic <code>delete()</code> .
<code>incr()</code>	Generic <code>incr()</code> .
<code>decr()</code>	Generic <code>decr()</code> .

### 15.6.3.8. Using MySQL and `memcached` with Java

The `com.danga.MemCached` class within Java provides a native interface to `memcached` instances. You can obtain the client from <http://whalin.com/memcached/>. The Java class uses hashes that are compatible with `libmemcached`, so you can mix and match Java and `libmemcached` applications accessing the same `memcached` instances. The serialization between Java and other interfaces are not compatible. If this is a problem, use JSON or a similar nonbinary serialization format.

On most systems, you can download the package and use the `jar` directly.

To use the `com.danga.MemCached` interface, you create a `MemCachedClient` instance and then configure the list of servers by configuring the `SockIOPool`. Through the pool specification you set up

the server list, weighting, and the connection parameters to optimized the connections between your client and the `memcached` instances that you configure.

Generally, you can configure the `memcached` interface once within a single class, then use this interface throughout the rest of your application.

For example, to create a basic interface, first configure the `MemCachedClient` and base `SockIOPool` settings:

```
public class MyClass {
    protected static MemCachedClient mcc = new MemCachedClient();

    static {
        String[] servers =
            {
                "localhost:11211",
            };

        Integer[] weights = { 1 };

        SockIOPool pool = SockIOPool.getInstance();

        pool.setServers( servers );
        pool.setWeights( weights );
    }
}
```

In the above sample, the list of servers is configured by creating an array of the `memcached` instances to use. You can then configure individual weights for each server.

The remainder of the properties for the connection are optional, but you can set the connection numbers (initial connections, minimum connections, maximum connections, and the idle timeout) by setting the pool parameters:

```
pool.setInitConn( 5 );
pool.setMinConn( 5 );
pool.setMaxConn( 250 );
pool.setMaxIdle( 1000 * 60 * 60 * 6
```

Once the parameters have been configured, initialize the connection pool:

```
pool.initialize();
```

The pool, and the connection to your `memcached` instances should now be ready to use.

To set the hashing algorithm used to select the server used when storing a given key, use `pool.setHashingAlg()`:

```
pool.setHashingAlg( SockIOPool.NEW_COMPAT_HASH );
```

Valid values are `NEW_COMPAT_HASH`, `OLD_COMPAT_HASH` and `NATIVE_HASH` are also basic modular hashing algorithms. For a consistent hashing algorithm, use `CONSISTENT_HASH`. These constants are equivalent to the corresponding hash settings within `libmemcached`.

The following table outlines the Java `com.danga.MemCached` methods and the equivalent generic methods in the `memcached` interface specification.

Java <code>com.danga.MemCached</code> Method	Equivalent Generic Method
<code>get()</code>	Generic <code>get()</code> .
<code>getMulti(keys)</code>	Get the values of multiple <code>keys</code> , returning the information as Hash map using <code>java.lang.String</code> for the keys and

Java <code>com.danga.MemCached</code> Method	Equivalent Generic Method
	<code>java.lang.Object</code> for the corresponding values.
<code>set()</code>	Generic <code>set()</code> .
<code>add()</code>	Generic <code>add()</code> .
<code>replace()</code>	Generic <code>replace()</code> .
<code>delete()</code>	Generic <code>delete()</code> .
<code>incr()</code>	Generic <code>incr()</code> .
<code>decr()</code>	Generic <code>decr()</code> .

### 15.6.3.9. Using the `memcached` TCP Text Protocol

Communicating with a `memcached` server can be achieved through either the TCP or UDP protocols. When using the TCP protocol, you can use a simple text based interface for the exchange of information.

When communicating with `memcached`, you can connect to the server using the port configured for the server. You can open a connection with the server without requiring authorization or login. As soon as you have connected, you can start to send commands to the server. When you have finished, you can terminate the connection without sending any specific disconnection command. Clients are encouraged to keep their connections open to decrease latency and improve performance.

Data is sent to the `memcached` server in two forms:

- Text lines, which are used to send commands to the server, and receive responses from the server.
- Unstructured data, which is used to receive or send the value information for a given key. Data is returned to the client in exactly the format it was provided.

Both text lines (commands and responses) and unstructured data are always terminated with the string `\r\n`. Because the data being stored may contain this sequence, the length of the data (returned by the client before the unstructured data is transmitted) should be used to determine the end of the data.

Commands to the server are structured according to their operation:

- **Storage commands:** `set`, `add`, `replace`, `append`, `prepend`, `cas`

Storage commands to the server take the form:

```
command key [flags] [exptime] length [noreply]
```

Or when using compare and swap (`cas`):

```
cas key [flags] [exptime] length [casunique] [noreply]
```

Where:

- `command`: The command name.
  - `set`: Store value against key
  - `add`: Store this value against key if the key does not already exist
  - `replace`: Store this value against key if the key already exists
  - `append`: Append the supplied value to the end of the value for the specified key. The `flags` and `exptime` arguments should not be used.
  - `prepend`: Append value currently in the cache to the end of the supplied value for the specified key. The `flags` and `exptime` arguments should not be used.

- `cas`: Set the specified key to the supplied value, only if the supplied `casunique` matches. This is effectively the equivalent of change the information if nobody has updated it since I last fetched it.
- `key`: The key. All data is stored using a the specific key. The key cannot contain control characters or whitespace, and can be up to 250 characters in size.
- `flags`: The flags for the operation (as an integer). Flags in `memcached` are transparent. The `memcached` server ignores the contents of the flags. They can be used by the client to indicate any type of information. In `memcached` 1.2.0 and lower the value is a 16-bit integer value. In `memcached` 1.2.1 and higher the value is a 32-bit integer.
- `exptime`: The expiry time, or zero for no expiry.
- `length`: The length of the supplied value block in bytes, excluding the terminating `\r\n` characters.
- `casunique`: A unique 64-bit value of an existing entry. This is used to compare against the existing value. Use the value returned by the `gets` command when issuing `cas` updates.
- `noreply`: Tells the server not to reply to the command.

For example, to store the value `abcdef` into the key `xyzkey`, you would use:

```
set xyzkey 0 0 6\r\nabcdef\r\n
```

The return value from the server is one line, specifying the status or error information. For more information, see [Table 15.3, “memcached Protocol Responses”](#).

- **Retrieval commands:** `get`, `gets`

Retrieval commands take the form:

```
get key1 [key2 ... keyn]
gets key1 [key2 ... keyn]
```

You can supply multiple keys to the commands, with each requested key separated by whitespace.

The server responds with an information line of the form:

```
VALUE key flags bytes [casunique]
```

Where:

- `key`: The key name.
- `flags`: The value of the flag integer supplied to the `memcached` server when the value was stored.
- `bytes`: The size (excluding the terminating `\r\n` character sequence) of the stored value.
- `casunique`: The unique 64-bit integer that identifies the item.

The information line is immediately followed by the value data block. For example:

```
get xyzkey\r\n
VALUE xyzkey 0 6\r\n
abcdef\r\n
```

If you have requested multiple keys, an information line and data block is returned for each key found. If a requested key does not exist in the cache, no information is returned.

- **Delete commands:** `delete`

Deletion commands take the form:

```
delete key [time] [noreply]
```

Where:

- `key`: The key name.
- `time`: The time in seconds (or a specific Unix time) for which the client wishes the server to refuse `add` or `replace` commands on this key. All `add`, `replace`, `get`, and `gets` commands fail during this period. `set` operations succeed. After this period, the key is deleted permanently and all commands are accepted.

If not supplied, the value is assumed to be zero (delete immediately).

- `noreply`: Tells the server not to reply to the command.

Responses to the command are either `DELETED` to indicate that the key was successfully removed, or `NOT_FOUND` to indicate that the specified key could not be found.

- **Increment/Decrement:** `incr`, `decr`

The increment and decrement commands change the value of a key within the server without performing a separate `get/set` sequence. The operations assume that the currently stored value is a 64-bit integer. If the stored value is not a 64-bit integer, then the value is assumed to be zero before the increment or decrement operation is applied.

Increment and decrement commands take the form:

```
incr key value [noreply]
decr key value [noreply]
```

Where:

- `key`: The key name.
- `value`: An integer to be used as the increment or decrement value.
- `noreply`: Tells the server not to reply to the command.

The response is:

- `NOT_FOUND`: The specified key could not be located.
- `value`: The new value associated with the specified key.

Values are assumed to be unsigned. For `decr` operations, the value is never decremented below 0. For `incr` operations, the value wraps around the 64-bit maximum.

- **Statistics commands:** `stats`

The `stats` command provides detailed statistical information about the current status of the `memcached` instance and the data it is storing.

Statistics commands take the form:

```
STAT [name] [value]
```

Where:

- `name`: The optional name of the statistics to return. If not specified, the general statistics are returned.

- `value`: A specific value to be used when performing certain statistics operations.

The return value is a list of statistics data, formatted as follows:

```
STAT name value
```

The statistics are terminated with a single line, `END`.

For more information, see [Section 15.6.4, “Getting `memcached` Statistics”](#).

For reference, a list of the different commands supported and their formats is provided below.

**Table 15.2. `memcached` Command Reference**

Command	Command Formats
<code>set</code>	<code>set key flags exptime length, set key flags exptime length noreply</code>
<code>add</code>	<code>add key flags exptime length, add key flags exptime length noreply</code>
<code>replace</code>	<code>replace key flags exptime length, replace key flags exptime length noreply</code>
<code>append</code>	<code>append key length, append key length noreply</code>
<code>prepend</code>	<code>prepend key length, prepend key length noreply</code>
<code>cas</code>	<code>cas key flags exptime length casunique, cas key flags exptime length casunique noreply</code>
<code>get</code>	<code>get key1 [key2 ... keyn]</code>
<code>gets</code>	
<code>delete</code>	<code>delete key, delete key noreply, delete key expiry, delete key expiry noreply</code>
<code>incr</code>	<code>incr key, incr key noreply, incr key value, incr key value noreply</code>
<code>decr</code>	<code>decr key, decr key noreply, decr key value, decr key value noreply</code>
<code>stat</code>	<code>stat, stat name, stat name value</code>

When sending a command to the server, the response from the server is one of the settings in the following table. All response values from the server are terminated by `\r\n`:

**Table 15.3. `memcached` Protocol Responses**

String	Description
<code>STORED</code>	Value has successfully been stored.
<code>NOT_STORED</code>	The value was not stored, but not because of an error. For commands where you are adding a or updating a value if it exists (such as <code>add</code> and <code>replace</code> ), or where the item has already been set to be deleted.
<code>EXISTS</code>	When using a <code>cas</code> command, the item you are trying to store already exists and has been modified since you last checked it.
<code>NOT_FOUND</code>	The item you are trying to store, update or delete does not exist or has already been deleted.
<code>ERROR</code>	You submitted a nonexistent command name.
<code>CLIENT_ERROR errorstring</code>	There was an error in the input line, the detail is contained in <code>errorstring</code> .



String	Description
<code>SERVER_ERROR</code> <code>errorstring</code>	There was an error in the server that prevents it from returning the information. In extreme conditions, the server may disconnect the client after this error occurs.
<code>VALUE</code> <code>keys</code> <code>flags</code> <code>length</code>	The requested key has been found, and the stored <code>key</code> , <code>flags</code> and data block are returned, of the specified <code>length</code> .
<code>DELETED</code>	The requested key was deleted from the server.
<code>STAT</code> <code>name</code> <code>value</code>	A line of statistics data.
<code>END</code>	The end of the statistics data.

## 15.6.4. Getting `memcached` Statistics

The `memcached` system has a built-in statistics system that collects information about the data being stored into the cache, cache hit ratios, and detailed information on the memory usage and distribution of information through the slab allocation used to store individual items. Statistics are provided at both a basic level that provide the core statistics, and more specific statistics for specific areas of the `memcached` server.

This information can be useful to ensure that you are getting the correct level of cache and memory usage, and that your slab allocation and configuration properties are set at an optimal level.

The stats interface is available through the standard `memcached` protocol, so the reports can be accessed by using `telnet` to connect to the `memcached`. The supplied `memcached-tool` includes support for obtaining the [Section 15.6.4.2, “memcached Slabs Statistics”](#) and [Section 15.6.4.1, “memcached General Statistics”](#) information. For more information, see [Section 15.6.4.6, “Using memcached-tool”](#).

Alternatively, most of the language API interfaces provide a function for obtaining the statistics from the server.

For example, to get the basic stats using `telnet`:

```
shell> telnet localhost 11211
Trying ::1...
Connected to localhost.
Escape character is '^]'.
stats
STAT pid 23599
STAT uptime 675
STAT time 1211439587
STAT version 1.2.5
STAT pointer_size 32
STAT rusage_user 1.404992
STAT rusage_system 4.694685
STAT curr_items 32
STAT total_items 56361
STAT bytes 2642
STAT curr_connections 53
STAT total_connections 438
STAT connection_structures 55
STAT cmd_get 113482
STAT cmd_set 80519
STAT get_hits 78926
STAT get_misses 34556
STAT evictions 0
STAT bytes_read 6379783
STAT bytes_written 4860179
STAT limit_maxbytes 67108864
STAT threads 1
END
```

When using Perl and the `Cache::Memcached` module, the `stats()` function returns information about all the servers currently configured in the connection object, and total statistics for all the `memcached` servers as a whole.

For example, the following Perl script obtains the stats and dumps the hash reference that is returned:

```
use Cache::Memcached;
use Data::Dumper;

my $memc = new Cache::Memcached;
$memc->set_servers(\@ARGV);

print Dumper($memc->stats());
```

When executed on the same `memcached` as used in the `Telnet` example above we get a hash reference with the host by host and total statistics:

```
$VAR1 = {
  'hosts' => {
    'localhost:11211' => {
      'misc' => {
        'bytes' => '2421',
        'curr_connections' => '3',
        'connection_structures' => '56',
        'pointer_size' => '32',
        'time' => '1211440166',
        'total_items' => '410956',
        'cmd_set' => '588167',
        'bytes_written' => '35715151',
        'evictions' => '0',
        'curr_items' => '31',
        'pid' => '23599',
        'limit_maxbytes' => '67108864',
        'uptime' => '1254',
        'rusage_user' => '9.857805',
        'cmd_get' => '838451',
        'rusage_system' => '34.096988',
        'version' => '1.2.5',
        'get_hits' => '581511',
        'bytes_read' => '46665716',
        'threads' => '1',
        'total_connections' => '3104',
        'get_misses' => '256940'
      },
      'sizes' => {
        '128' => '16',
        '64' => '15'
      }
    }
  },
  'self' => {},
  'total' => {
    'cmd_get' => 838451,
    'bytes' => 2421,
    'get_hits' => 581511,
    'connection_structures' => 56,
    'bytes_read' => 46665716,
    'total_items' => 410956,
    'total_connections' => 3104,
    'cmd_set' => 588167,
    'bytes_written' => 35715151,
    'curr_items' => 31,
    'get_misses' => 256940
  }
};
```

The statistics are divided up into a number of distinct sections, and then can be requested by adding the type to the `stats` command. Each statistics output is covered in more detail in the following sections.

- General statistics, see [Section 15.6.4.1, “memcached General Statistics”](#).

- Slab statistics (`slabs`), see [Section 15.6.4.2, “memcached Slabs Statistics”](#).
- Item statistics (`items`), see [Section 15.6.4.3, “memcached Item Statistics”](#).
- Size statistics (`sizes`), see [Section 15.6.4.4, “memcached Size Statistics”](#).
- Detailed status (`detail`), see [Section 15.6.4.5, “memcached Detail Statistics”](#).

#### 15.6.4.1. `memcached` General Statistics

The output of the general statistics provides an overview of the performance and use of the `memcached` instance. The statistics returned by the command and their meaning is shown in the following table.

The following terms are used to define the value type for each statistics value:

- `32u`: 32-bit unsigned integer
- `64u`: 64-bit unsigned integer
- `32u:32u`: Two 32-bit unsigned integers separated by a colon
- `String`: Character string

Statistic	Data type	Description	Version
<code>pid</code>	<code>32u</code>	Process ID of the <code>memcached</code> instance.	
<code>uptime</code>	<code>32u</code>	Uptime (in seconds) for this <code>memcached</code> instance.	
<code>time</code>	<code>32u</code>	Current time (as epoch).	
<code>version</code>	string	Version string of this instance.	
<code>pointer_size</code>	string	Size of pointers for this host specified in bits (32 or 64).	
<code>rusage_user</code>	<code>32u:32u</code>	Total user time for this instance (seconds:microseconds).	
<code>rusage_system</code>	<code>32u:32u</code>	Total system time for this instance (seconds:microseconds).	
<code>curr_items</code>	<code>32u</code>	Current number of items stored by this instance.	
<code>total_items</code>	<code>32u</code>	Total number of items stored during the life of this instance.	
<code>bytes</code>	<code>64u</code>	Current number of bytes used by this server to store items.	
<code>curr_connections</code>	<code>32u</code>	Current number of open connections.	
<code>total_connections</code>	<code>32u</code>	Total number of connections opened since the server started running.	
<code>connection_structures</code>	<code>32u</code>	Number of connection structures allocated by the server.	
<code>cmd_get</code>	<code>64u</code>	Total number of retrieval requests ( <code>get</code> operations).	
<code>cmd_set</code>	<code>64u</code>	Total number of storage requests ( <code>set</code> operations).	
<code>get_hits</code>	<code>64u</code>	Number of keys that have been requested and found present.	
<code>get_misses</code>	<code>64u</code>	Number of items that have been requested and not found.	
<code>delete_hits</code>	<code>64u</code>	Number of keys that have been deleted and found present.	1.3.x
<code>delete_misses</code>	<code>64u</code>	Number of items that have been delete and not found.	1.3.x

Statistic	Data type	Description	Version
<code>incr_hits</code>	64u	Number of keys that have been incremented and found present.	1.3.x
<code>incr_misses</code>	64u	Number of items that have been incremented and not found.	1.3.x
<code>decr_hits</code>	64u	Number of keys that have been decremented and found present.	1.3.x
<code>decr_misses</code>	64u	Number of items that have been decremented and not found.	1.3.x
<code>cas_hits</code>	64u	Number of keys that have been compared and swapped and found present.	1.3.x
<code>cas_misses</code>	64u	Number of items that have been compared and swapped and not found.	1.3.x
<code>cas_badvalue</code>	64u	Number of keys that have been compared and swapped, but the comparison (original) value did not match the supplied value.	1.3.x
<code>evictions</code>	64u	Number of valid items removed from cache to free memory for new items.	
<code>bytes_read</code>	64u	Total number of bytes read by this server from network.	
<code>bytes_written</code>	64u	Total number of bytes sent by this server to network.	
<code>limit_maxbytes</code>	32u	Number of bytes this server is permitted to use for storage.	
<code>threads</code>	32u	Number of worker threads requested.	
<code>conn_yields</code>	64u	Number of yields for connections (related to the <code>-R</code> option).	1.4.0

The most useful statistics from those given here are the number of cache hits, misses, and evictions.

A large number of `get_misses` may just be an indication that the cache is still being populated with information. The number should, over time, decrease in comparison to the number of cache `get_hits`. If, however, you have a large number of cache misses compared to cache hits after an extended period of execution, it may be an indication that the size of the cache is too small and you either need to increase the total memory size, or increase the number of the `memcached` instances to improve the hit ratio.

A large number of `evictions` from the cache, particularly in comparison to the number of items stored is a sign that your cache is too small to hold the amount of information that you regularly want to keep cached. Instead of items being retained in the cache, items are being evicted to make way for new items keeping the turnover of items in the cache high, reducing the efficiency of the cache.

#### 15.6.4.2. `memcached` Slabs Statistics

To get the `slabs` statistics, use the `stats slabs` command, or the API equivalent.

The slab statistics provide you with information about the slabs that have created and allocated for storing information within the cache. You get information both on each individual slab-class and total statistics for the whole slab.

```
STAT 1:chunk_size 104
STAT 1:chunks_per_page 10082
STAT 1:total_pages 1
STAT 1:total_chunks 10082
STAT 1:used_chunks 10081
STAT 1:free_chunks 1
STAT 1:free_chunks_end 10079
```

## Getting memcached Statistics

```
STAT 9:chunk_size 696
STAT 9:chunks_per_page 1506
STAT 9:total_pages 63
STAT 9:total_chunks 94878
STAT 9:used_chunks 94878
STAT 9:free_chunks 0
STAT 9:free_chunks_end 0
STAT active_slabs 2
STAT total_malloced 67083616
END
```

Individual stats for each slab class are prefixed with the slab ID. A unique ID is given to each allocated slab from the smallest size up to the largest. The prefix number indicates the slab class number in relation to the calculated chunk from the specified growth factor. Hence in the example, 1 is the first chunk size and 9 is the 9th chunk allocated size.

The parameters returned for each chunk size and a description of each parameter are provided in the following table.

Statistic	Description	Version
<code>chunk_size</code>	Space allocated to each chunk within this slab class.	
<code>chunks_per_page</code>	Number of chunks within a single page for this slab class.	
<code>total_pages</code>	Number of pages allocated to this slab class.	
<code>total_chunks</code>	Number of chunks allocated to the slab class.	
<code>used_chunks</code>	Number of chunks allocated to an item..	
<code>free_chunks</code>	Number of chunks not yet allocated to items.	
<code>free_chunks_end</code>	Number of free chunks at the end of the last allocated page.	
<code>get_hits</code>	Number of get hits to this chunk	1.3.x
<code>cmd_set</code>	Number of set commands on this chunk	1.3.x
<code>delete_hits</code>	Number of delete hits to this chunk	1.3.x
<code>incr_hits</code>	Number of increment hits to this chunk	1.3.x
<code>decr_hits</code>	Number of decrement hits to this chunk	1.3.x
<code>cas_hits</code>	Number of CAS hits to this chunk	1.3.x
<code>cas_badval</code>	Number of CAS hits on this chunk where the existing value did not match	1.3.x
<code>mem_requested</code>	The true amount of memory of memory requested within this chunk	1.4.1

The following additional statistics cover the information for the entire server, rather than on a chunk by chunk basis:

Statistic	Description	Version
<code>active_slabs</code>	Total number of slab classes allocated.	
<code>total_malloced</code>	Total amount of memory allocated to slab pages.	

The key values in the slab statistics are the `chunk_size`, and the corresponding `total_chunks` and `used_chunks` parameters. These given an indication of the size usage of the chunks within the system. Remember that one key/value pair is placed into a chunk of a suitable size.

From these stats, you can get an idea of your size and chunk allocation and distribution. If you store many items with a number of largely different sizes, consider adjusting the chunk size growth factor to increase in larger steps to prevent chunk and memory wastage. A good indication of a bad growth factor is a high number of different slab classes, but with relatively few chunks actually in use within

each slab. Increasing the growth factor creates fewer slab classes and therefore makes better use of the allocated pages.

### 15.6.4.3. memcached Item Statistics

To get the `items` statistics, use the `stats items` command, or the API equivalent.

The `items` statistics give information about the individual items allocated within a given slab class.

```
STAT items:2:number 1
STAT items:2:age 452
STAT items:2:evicted 0
STAT items:2:evicted_nonzero 0
STAT items:2:evicted_time 2
STAT items:2:outofmemory 0
STAT items:2:tailrepairs 0
...
STAT items:27:number 1
STAT items:27:age 452
STAT items:27:evicted 0
STAT items:27:evicted_nonzero 0
STAT items:27:evicted_time 2
STAT items:27:outofmemory 0
STAT items:27:tailrepairs 0
```

The prefix number against each statistics relates to the corresponding chunk size, as returned by the `stats slabs` statistics. The result is a display of the number of items stored within each chunk within each slab size, and specific statistics about their age, eviction counts, and out of memory counts. A summary of the statistics is given in the following table.

Statistic	Description	
<code>number</code>	The number of items currently stored in this slab class.	
<code>age</code>	The age of the oldest item within the slab class, in seconds.	
<code>evicted</code>	The number of items evicted to make way for new entries.	
<code>evicted_time</code>	The time of the last evicted entry	
<code>evicted_nonzero</code>	The time of the last evicted non-zero entry	1.4.0
<code>outofmemory</code>	The number of items for this slab class that have triggered an out of memory error (only value when the <code>-M</code> command line option is in effect).	
<code>tailrepairs</code>	Number of times the entries for a particular ID need repairing	

Item level statistics can be used to determine how many items are stored within a given slab and their freshness and recycle rate. You can use this to help identify whether there are certain slab classes that are triggering a much larger number of evictions than others.

### 15.6.4.4. memcached Size Statistics

To get size statistics, use the `stats sizes` command, or the API equivalent.

The size statistics provide information about the sizes and number of items of each size within the cache. The information is returned as two columns, the first column is the size of the item (rounded up to the nearest 32 byte boundary), and the second column is the count of the number of items of that size within the cache:

```
96 35
128 38
160 807
192 804
224 410
```

```
256 222
288 83
320 39
352 53
384 33
416 64
448 51
480 30
512 54
544 39
576 10065
```



### Caution

Running this statistic locks up your cache as each item is read from the cache and its size calculated. On a large cache, this may take some time and prevent any set or get operations until the process completes.

The item size statistics are useful only to determine the sizes of the objects you are storing. Since the actual memory allocation is relevant only in terms of the chunk size and page size, the information is only useful during a careful debugging or diagnostic session.

#### 15.6.4.5. `memcached` Detail Statistics

For `memcached` 1.3.x and higher, you can enable and obtain detailed statistics about the `get`, `set`, and `del` operations on the individual keys stored in the cache, and determine whether the attempts hit (found) a particular key. These operations are only recorded while the detailed stats analysis is turned on.

To enable detailed statistics, you must send the `stats detail on` command to the `memcached` server:

```
$ telnet localhost 11211
Trying 127.0.0.1...
Connected to tiger.
Escape character is '^]'.stats detail on
OK
```

Individual statistics are recorded for every `get`, `set` and `del` operation on a key, including keys that are not currently stored in the server. For example, if an attempt is made to obtain the value of key `abckey` and it does not exist, the `get` operating on the specified key are recorded while detailed statistics are in effect, even if the key is not currently stored. The `hits`, that is, the number of `get` or `del` operations for a key that exists in the server are also counted.

To turn detailed statistics off, send the `stats detail off` command to the `memcached` server:

```
$ telnet localhost 11211
Trying 127.0.0.1...
Connected to tiger.
Escape character is '^]'.stats detail on
OK
```

To obtain the detailed statistics recorded during the process, send the `stats detail dump` command to the `memcached` server:

```
stats detail dump
PREFIX hykkey get 0 hit 0 set 1 del 0
PREFIX xyzkey get 0 hit 0 set 1 del 0
PREFIX yukkey get 1 hit 0 set 0 del 0
PREFIX abckey get 3 hit 3 set 1 del 0
END
```

You can use the detailed statistics information to determine whether your `memcached` clients are using a large number of keys that do not exist in the server by comparing the `hit` and `get` or `del` counts.



Because the information is recorded by key, you can also determine whether the failures or operations are clustered around specific keys.

#### 15.6.4.6. Using memcached-tool

The `memcached-tool`, located within the `scripts` directory within the `memcached` source directory. The tool provides convenient access to some reports and statistics from any `memcached` instance.

The basic format of the command is:

```
shell> ./memcached-tool hostname:port [command]
```

The default output produces a list of the slab allocations and usage. For example:

```
shell> memcached-tool localhost:11211 display
# Item_Size Max_age Pages Count Full? Evicted Evict_Time OOM
1 80B 93s 1 20 no 0 0 0
2 104B 93s 1 16 no 0 0 0
3 136B 1335s 1 28 no 0 0 0
4 176B 1335s 1 24 no 0 0 0
5 224B 1335s 1 32 no 0 0 0
6 280B 1335s 1 34 no 0 0 0
7 352B 1335s 1 36 no 0 0 0
8 440B 1335s 1 46 no 0 0 0
9 552B 1335s 1 58 no 0 0 0
10 696B 1335s 1 66 no 0 0 0
11 872B 1335s 1 89 no 0 0 0
12 1.1K 1335s 1 112 no 0 0 0
13 1.3K 1335s 1 145 no 0 0 0
14 1.7K 1335s 1 123 no 0 0 0
15 2.1K 1335s 1 198 no 0 0 0
16 2.6K 1335s 1 199 no 0 0 0
17 3.3K 1335s 1 229 no 0 0 0
18 4.1K 1335s 1 248 yes 36 2 0
19 5.2K 1335s 2 328 no 0 0 0
20 6.4K 1335s 2 316 yes 387 1 0
21 8.1K 1335s 3 381 yes 492 1 0
22 10.1K 1335s 3 303 yes 598 2 0
23 12.6K 1335s 5 405 yes 605 1 0
24 15.8K 1335s 6 384 yes 766 2 0
25 19.7K 1335s 7 357 yes 908 170 0
26 24.6K 1336s 7 287 yes 1012 1 0
27 30.8K 1336s 7 231 yes 1193 169 0
28 38.5K 1336s 4 104 yes 1323 169 0
29 48.1K 1336s 1 21 yes 1287 1 0
30 60.2K 1336s 1 17 yes 1093 169 0
31 75.2K 1337s 1 13 yes 713 168 0
32 94.0K 1337s 1 10 yes 278 168 0
33 117.5K 1336s 1 3 no 0 0 0
```

This output is the same if you specify the `command` as `display`:

```
shell> memcached-tool localhost:11211 display
# Item_Size Max_age Pages Count Full? Evicted Evict_Time OOM
1 80B 93s 1 20 no 0 0 0
2 104B 93s 1 16 no 0 0 0
...
```

The output shows a summarized version of the output from the `slabs` statistics. The columns provided in the output are shown below:

- `#`: The slab number
- `Item_Size`: The size of the slab
- `Max_age`: The age of the oldest item in the slab
- `Pages`: The number of pages allocated to the slab

- **Count**: The number of items in this slab
- **Full?**: Whether the slab is fully populated
- **Evicted**: The number of objects evicted from this slab
- **Evict\_Time**: The time (in seconds) since the last eviction
- **OOM**: The number of items that have triggered an out of memory error

You can also obtain a dump of the general statistics for the server using the `stats` command:

```
shell> memcached-tool localhost:11211 stats
#localhost:11211  Field      Value
    accepting_conns      1
        bytes           162
    bytes_read           485
    bytes_written        6820
    cas_badval           0
    cas_hits             0
    cas_misses           0
    cmd_flush            0
    cmd_get              4
    cmd_set              2
    conn_yields          0
connection_structures  11
    curr_connections     10
        curr_items      2
    decr_hits            0
    decr_misses         1
    delete_hits         0
    delete_misses      0
    evictions            0
    get_hits             4
    get_misses          0
    incr_hits            0
    incr_misses         2
    limit_maxbytes       67108864
listen_disabled_num    0
    pid                 12981
    pointer_size         32
    rusage_system        0.013911
    rusage_user          0.011876
    threads              4
    time                 1255518565
total_connections     20
    total_items          2
    uptime              880
    version              1.4.2
```

## 15.6.5. memcached FAQ

### Questions

- **16.6.5.1: [1889]** Can memcached be run on a Windows environment?
- **16.6.5.2: [1889]** What is the maximum size of an object you can store in memcached? Is that configurable?
- **16.6.5.3: [1889]** Is it true `memcached` will be much more effective with db-read-intensive applications than with db-write-intensive applications?
- **16.6.5.4: [1889]** Is there any overhead in not using persistent connections? If persistent is always recommended, what are the downsides (for example, locking up)?
- **16.6.5.5: [1890]** How is an event such as a crash of one of the `memcached` servers handled by the `memcached` client?

- [16.6.5.6: \[1890\]](#) What is a recommended hardware configuration for a memcached server?
- [16.6.5.7: \[1890\]](#) Is memcached more effective for video and audio as opposed to textual read/writes?
- [16.6.5.8: \[1890\]](#) Can memcached work with ASPX?
- [16.6.5.9: \[1890\]](#) How expensive is it to establish a memcache connection? Should those connections be pooled?
- [16.6.5.10: \[1890\]](#) How is the data handled when the memcached server is down?
- [16.6.5.11: \[1890\]](#) How are auto-increment columns in the MySQL database coordinated across multiple instances of memcached?
- [16.6.5.12: \[1891\]](#) Is compression available?
- [16.6.5.13: \[1891\]](#) Can we implement different types of memcached as different nodes in the same server, so can there be deterministic and non-deterministic in the same server?
- [16.6.5.14: \[1891\]](#) What are best practices for testing an implementation, to ensure that it improves performance, and to measure the impact of memcached configuration changes? And would you recommend keeping the configuration very simple to start?

## Questions and Answers

### 16.6.5.1: Can memcached be run on a Windows environment?

No. Currently memcached is available only on the Unix/Linux platform. There is an unofficial port available, see <http://www.codeplex.com/memcachedproviders>.

### 16.6.5.2: What is the maximum size of an object you can store in memcached? Is that configurable?

The default maximum object size is 1MB. In memcached 1.4.2 and later, you can change the maximum size of an object using the `-I` command line option.

For versions before this, to increase this size, you have to re-compile memcached. You can modify the value of the `POWER_BLOCK` within the `slabs.c` file within the source.

In memcached 1.4.2 and higher, you can configure the maximum supported object size by using the `-I` command-line option. For example, to increase the maximum object size to 5MB:

```
$ memcached -I 5m
```

If an object is larger than the maximum object size, you must manually split it. memcached is very simple: you give it a key and some data, it tries to cache it in RAM. If you try to store more than the default maximum size, the value is just truncated for speed reasons.

### 16.6.5.3: Is it true memcached will be much more effective with db-read-intensive applications than with db-write-intensive applications?

Yes. memcached plays no role in database writes, it is a method of caching data already read from the database in RAM.

### 16.6.5.4: Is there any overhead in not using persistent connections? If persistent is always recommended, what are the downsides (for example, locking up)?

If you don't use persistent connections when communicating with memcached, there will be a small increase in the latency of opening the connection each time. The effect is comparable to use nonpersistent connections with MySQL.

In general, the chance of locking or other issues with persistent connections is minimal, because there is very little locking within `memcached`. If there is a problem, eventually your request will time out and return no result, so your application will need to load from MySQL again.

**16.6.5.5: How is an event such as a crash of one of the `memcached` servers handled by the `memcached` client?**

There is no automatic handling of this. If your client fails to get a response from a server, code a fallback mechanism to load the data from the MySQL database.

The client APIs all provide the ability to add and remove `memcached` instances on the fly. If within your application you notice that `memcached` server is no longer responding, you can remove the server from the list of servers, and keys will automatically be redistributed to another `memcached` server in the list. If retaining the cache content on all your servers is important, make sure you use an API that supports a consistent hashing algorithm. For more information, see [Section 15.6.2.4, “memcached Hashing/Distribution Types”](#).

**16.6.5.6: What is a recommended hardware configuration for a `memcached` server?**

`memcached` has a very low processing overhead. All that is required is spare physical RAM capacity. A `memcached` server does not require a dedicated machine. If you have web, application, or database servers that have spare RAM capacity, then use them with `memcached`.

To build and deploy a dedicated `memcached` server, use a relatively low-power CPU, lots of RAM, and one or more Gigabit Ethernet interfaces.

**16.6.5.7: Is `memcached` more effective for video and audio as opposed to textual read/writes?**

`memcached` works equally well for all kinds of data. To `memcached`, any value you store is just a stream of data. Remember, though, that the maximum size of an object you can store in `memcached` is 1MB, but can be configured to be larger by using the `-I` option in `memcached` 1.4.2 and later, or by modifying the source in versions before 1.4.2. If you plan on using `memcached` with audio and video content, you will probably want to increase the maximum object size. Also remember that `memcached` is a solution for caching information for reading. It shouldn't be used for writes, except when updating the information in the cache.

**16.6.5.8: Can `memcached` work with ASPX?**

There are ports and interfaces for many languages and environments. ASPX relies on an underlying language such as C# or VisualBasic, and if you are using ASP.NET then there is a C# `memcached` library. For more information, see <https://sourceforge.net/projects/memcacheddotnet/>.

**16.6.5.9: How expensive is it to establish a memcache connection? Should those connections be pooled?**

Opening the connection is relatively inexpensive, because there is no security, authentication or other handshake taking place before you can start sending requests and getting results. Most APIs support a persistent connection to a `memcached` instance to reduce the latency. Connection pooling would depend on the API you are using, but if you are communicating directly over TCP/IP, then connection pooling would provide some small performance benefit.

**16.6.5.10: How is the data handled when the `memcached` server is down?**

The behavior is entirely application dependent. Most applications fall back to loading the data from the database (just as if they were updating the `memcached` information). If you are using multiple `memcached` servers, you might also remove a downed server from the list to prevent it from affecting performance. Otherwise, the client will still attempt to communicate with the `memcached` server that corresponds to the key you are trying to load.

**16.6.5.11: How are auto-increment columns in the MySQL database coordinated across multiple instances of `memcached`?**

They aren't. There is no relationship between MySQL and `memcached` unless your application (or, if you are using the MySQL UDFs for `memcached`, your database definition) creates one.

If you are storing information based on an auto-increment key into multiple instances of `memcached`, the information is only stored on one of the `memcached` instances anyway. The client uses the key value to determine which `memcached` instance to store the information. It doesn't store the same information across all the instances, as that would be a waste of cache memory.

#### **16.6.5.12: Is compression available?**

Yes. Most of the client APIs support some sort of compression, and some even allow you to specify the threshold at which a value is deemed appropriate for compression during storage.

#### **16.6.5.13: Can we implement different types of `memcached` as different nodes in the same server, so can there be deterministic and non-deterministic in the same server?**

Yes. You can run multiple instances of `memcached` on a single server, and in your client configuration you choose the list of servers you want to use.

#### **16.6.5.14: What are best practices for testing an implementation, to ensure that it improves performance, and to measure the impact of `memcached` configuration changes? And would you recommend keeping the configuration very simple to start?**

The best way to test the performance is to start up a `memcached` instance. First, modify your application so that it stores the data just before the data is about to be used or displayed into `memcached`. Since the APIs handle the serialization of the data, it should just be a one-line modification to your code. Then, modify the start of the process that would normally load that information from MySQL with the code that requests the data from `memcached`. If the data cannot be loaded from `memcached`, default to the MySQL process.

All of the changes required will probably amount to just a few lines of code. To get the best benefit, make sure you cache entire objects (for example, all the components of a web page, blog post, discussion thread, and so on), rather than using `memcached` as a simple cache of individual rows of MySQL tables.

Keeping the configuration simple at the start, or even over the long term, is easy with `memcached`. Once you have the basic structure up and running, often the only ongoing change is to add more servers into the list of servers used by your applications. You don't need to manage the `memcached` servers, and there is no complex configuration; just add more servers to the list and let the client API and the `memcached` servers make the decisions.

## **15.7. MySQL Proxy**

The MySQL Proxy is an application that communicates over the network using the MySQL network protocol and provides communication between one or more MySQL servers and one or more MySQL clients. Because MySQL Proxy uses the MySQL network protocol, it can be used without modification with any MySQL-compatible client that uses the protocol. This includes the `mysql` command-line client, any clients that uses the MySQL client libraries, and any connector that supports the MySQL network protocol.

In the most basic configuration, MySQL Proxy simply interposes itself between the server and clients, passing queries from the clients to the MySQL Server and returning the responses from the MySQL Server to the appropriate client. In more advanced configurations, the MySQL Proxy can also monitor and alter the communication between the client and the server. Query interception enables you to add profiling, and interception of the exchanges is scriptable using the Lua scripting language.

By intercepting the queries from the client, the proxy can insert additional queries into the list of queries sent to the server, and remove the additional results when they are returned by the server. Using this functionality you can return the results from the original query to the client while adding informational statements to each query, for example, to monitor their execution time or progress, and separately log the results.

The proxy enables you to perform additional monitoring, filtering, or manipulation of queries without requiring you to make any modifications to the client and without the client even being aware that it is communicating with anything but a genuine MySQL server.

This documentation covers MySQL Proxy 0.8.2. And MySQL Proxy contains third-party code. For license information on third-party code, see [Appendix A, Licenses for Third-Party Components](#).

**Warning**

MySQL Proxy is currently an Alpha release and should not be used within production environments.

**Important**

MySQL Proxy is compatible with MySQL 5.0 or later. Testing has not been performed with Version 4.1. Please provide feedback on your experiences using the [MySQL Proxy Forum](#).

For release notes detailing the changes in each release of MySQL Proxy, see [MySQL Proxy Release Notes](#).

## 15.7.1. MySQL Proxy Supported Platforms

MySQL Proxy is currently available as a precompiled binary for the following platforms:

- Linux (including Red Hat, Fedora, Debian, SuSE) and derivatives
- Mac OS X
- FreeBSD
- IBM AIX
- Sun Solaris
- Microsoft Windows (including Microsoft Windows XP, Microsoft Windows Vista, Microsoft Windows Server 2003, Microsoft Windows Server 2008)

**Note**

You must have the .NET Framework 1.1 or higher installed.

Other Unix/Linux platforms not listed should be compatible by using the source package and building MySQL Proxy locally.

System requirements for the MySQL Proxy application are the same as the main MySQL server. Currently MySQL Proxy is compatible only with MySQL 5.0.1 and later. MySQL Proxy is provided as a standalone, statically linked binary. You need not have MySQL or Lua installed.

## 15.7.2. Installing MySQL Proxy

You have three choices for installing MySQL Proxy:

- Precompiled binaries are available for a number of different platforms. See [Section 15.7.2.1, “Installing MySQL Proxy from a Binary Distribution”](#).
- You can install from the source code to build on an environment not supported by the binary distributions. See [Section 15.7.2.2, “Installing MySQL Proxy from a Source Distribution”](#).
- The latest version of the MySQL Proxy source code is available through a development repository is the best way to stay up to date with the latest fixes and revisions. See [Section 15.7.2.3, “Installing MySQL Proxy from the Bazaar Repository”](#).

### 15.7.2.1. Installing MySQL Proxy from a Binary Distribution

If you download a binary package, you must extract and copy the package contents to your desired installation directory. The package contains files required by MySQL Proxy, including additional Lua scripts and other components required for execution.

To install, unpack the archive into the desired directory, then modify your `PATH` environment variable so that you can use the `mysql-proxy` command directly:

```
shell> cd /usr/local
shell> tar zxf mysql-proxy-0.8.2-platform.tar.gz
shell> PATH=$PATH:/usr/local/mysql-proxy-0.8.2-platform/sbin
```

To update the path globally on a system, you might need administrator privileges to modify the appropriate `/etc/profile`, `/etc/bashrc`, or other system configuration file.

On Windows, you can update the `PATH` environment variable using this procedure:

1. On the Windows desktop, right-click the My Computer icon, and select Properties.
2. Next select the Advanced tab from the System Properties menu that appears, and click the Environment Variables button.
3. Under **System Variables**, select Path, then click the Edit button. The Edit System Variable dialogue should appear.

The Microsoft Visual C++ runtime libraries are a requirement for running MySQL Proxy as of version 0.8.2. Users that do not have these libraries must download and install the Microsoft Visual C++ 2008 Service Pack 1 Redistributable Package MFC Security Update. Use the following link to obtain the package:

<http://www.microsoft.com/download/en/details.aspx?id=26368>

### 15.7.2.2. Installing MySQL Proxy from a Source Distribution

You can download a source package and compile the MySQL Proxy yourself. To build from source, you must have the following prerequisite components installed:

- `libevent` 1.x or higher (1.3b or later is preferred).
- `lua` 5.1.x or higher.
- `glib2` 2.6.0 or higher.
- `pkg-config`.
- `libtool` 1.5 or higher.
- MySQL 5.0.x or higher developer files.



#### Note

On some operating systems, you might need to manually build the required components to get the latest version. If you have trouble compiling MySQL Proxy, consider using a binary distributions instead.

After verifying that the prerequisite components are installed, configure and build MySQL Proxy:

```
shell> tar zxf mysql-proxy-0.8.2.tar.gz
shell> cd mysql-proxy-0.8.2
shell> ./configure
shell> make
```

To test the build, use the `check` target to `make`:



```
shell> make check
```

The tests try to connect to `localhost` using the `root` user. To provide a password, set the `MYSQL_PASSWORD` environment variable:

```
shell> MYSQL_PASSWORD=root_pwd make check
```

You can install using the `install` target:

```
shell> make install
```

By default, `mysql-proxy` is installed into `/usr/local/sbin/mysql-proxy`. The Lua example scripts are installed into `/usr/local/share`.

### 15.7.2.3. Installing MySQL Proxy from the Bazaar Repository

The MySQL Proxy source is available through a public Bazaar repository and is the quickest way to get the latest releases and fixes.

A build from the Bazaar repository requires that the following prerequisite components be installed:

- Bazaar 1.10.0 or later.
- `libtool` 1.5 or higher.
- `autoconf` 2.56 or higher.
- `automake` 1.10 or higher.
- `libevent` 1.x or higher (1.3b or later is preferred).
- `lua` 5.1.x or higher.
- `glib2` 2.4.0 or higher.
- `pkg-config`.
- MySQL 5.0.x or higher developer files.

The `mysql-proxy` source is hosted on Launchpad. To check out a local copy of the Bazaar repository, use `bzr`:

```
shell> bzr branch lp:mysql-proxy
```

The preceding command downloads a complete version of the Bazaar repository for `mysql-proxy`. The main source files are located within the `trunk` subdirectory. The configuration scripts must be generated before you can configure and build `mysql-proxy`. The `autogen.sh` script generates the required configuration scripts for you:

```
shell> sh ./autogen.sh
```

The `autogen.sh` script creates the standard `configure` script, which you then use to configure and build with `make`:

```
shell> ./configure
shell> make
shell> make install
```

To create a standalone source distribution, identical to the source distribution available for download, use this command:

```
shell> make distcheck
```

The preceding command creates the file `mysql-proxy-0.8.2.tar.gz` (with the corresponding current version) within the current directory.

### 15.7.2.4. Setting Up MySQL Proxy as a Windows Service

The MySQL distribution on Windows includes the `mysql-proxy-svc.exe` command that enables a MySQL Proxy instance to be managed by the Windows service control manager. You can control the service, including automatically starting and stopping it during boot, reboot and shutdown, without separately running the MySQL Proxy application.

To set up a MySQL Proxy service, use the `sc` command to create a new service using the MySQL Proxy service command. Specify the MySQL Proxy options on the `sc` command line, and identify the service with a unique name. For example, to configure a new MySQL Proxy instance that will automatically start when your system boots, redirecting queries to the local MySQL server:

```
C:\> sc create "Proxy" DisplayName= "MySQL Proxy" start= "auto" »
      binPath= "C:\Program Files\MySQL\mysql-proxy-0.8.2\bin\mysql-proxy-svc.exe »
      --proxy-backend-addresses=127.0.0.1:3306"
```



#### Note

The space following the equal sign after each property is required; failure to include it results in an error.

The preceding command creates a new service called `Proxy`. You can start and stop the service using the `net start|stop` command with the service name. The service is not automatically started after it is created. To start the service:

```
C:\> net start proxy
The MySQL Proxy service is starting.
The MySQL Proxy service was started successfully.
```

You can specify additional command-line options to the `sc` command. You can also set up multiple MySQL Proxy services on the same machine (providing they are configured to listen on different ports and/or IP addresses).

You can delete a service that you have created:

```
C:\> sc delete proxy
```

For more information on creating services using `sc`, see [How to create a Windows service by using Sc.exe](#).

### 15.7.3. MySQL Proxy Command Options

To start MySQL Proxy, you can run it directly from the command line:

```
shell> mysql-proxy
```

For most situations, you specify at least the host name or address and the port number of the backend MySQL server to which the MySQL Proxy should pass queries.

You can specify options to `mysql-proxy` either on the command line, or by using a configuration file and the `--defaults-file` [1899] command-line option to specify the file location.

If you use a configuration file, format it as follows:

- Specify the options within a `[mysql-proxy]` configuration group. For example:

```
[mysql-proxy]
admin-address = host:port
```

- Specify all configuration options in the form of a configuration name and the value to set.
- For options that are a simple toggle on the command line (for example, `--proxy-skip-profiling` [1903]), use `true` or `false`. For example, the following is invalid:

```
[mysql-proxy]
proxy-skip-profiling
```

But this is valid:

```
[mysql-proxy]
proxy-skip-profiling = true
```

- Give the configuration file Unix permissions of 0660 (readable and writable by user and group, no access for others).

Failure to adhere to any of these requirements causes `mysql-proxy` to generate an error during startup.

The following tables list the supported configuration file and command-line options.

**Table 15.4. `mysql-proxy` Help Options**

Format	Option File	Description
<code>--help</code> [1897]		Show help options
<code>--help-admin</code> [1898]		Show admin module options
<code>--help-all</code> [1898]		Show all help options
<code>--help-proxy</code> [1898]		Show proxy module options

**Table 15.5. `mysql-proxy` Admin Options**

Format	Option File	Description
<code>--admin-address=host:port</code> [1898]	<code>admin-address=host:port</code> [1898]	The admin module listening host and port
<code>--admin-lua-script=file_name</code> [1898]	<code>admin-lua-script=file_name</code> [1898]	Script to execute by the admin module
<code>--admin-password=password</code> [1898]	<code>admin-password=password</code> [1898]	Authentication password for admin module
<code>--admin-username=user_name</code> [1898]	<code>admin-username=user_name</code> [1898]	Authentication user name for admin module
<code>--proxy-address=host:port</code> [1901]	<code>proxy-address=host:port</code> [1901]	The listening proxy server host and port

**Table 15.6. `mysql-proxy` Proxy Options**

Format	Option File	Description	Removed
<code>--no-proxy</code> [1901]	<code>no-proxy</code> [1901]	Do not start the proxy module	
<code>--proxy-backend-addresses=host:port</code> [1902]	<code>proxy-backend-addresses=host:port</code> [1902]	The MySQL server host and port	
<code>--proxy-fix-bug-25371</code> [1903]	<code>proxy-fix-bug-25371</code> [1903]	Enable the fix for Bug #25371 for older libmysql versions	0.8.1
<code>--proxy-lua-script=file_name</code> [1903]	<code>proxy-lua-script=file_name</code> [1903]	Filename for Lua script for proxy operations	
<code>--proxy-pool-no-change-user</code> [1903]	<code>proxy-pool-no-change-user</code> [1903]	Do not use the protocol <code>CHANGE_USER</code> command to reset the connection when coming from the connection pool	
<code>--proxy-read-only-backend-addresses=host:port</code> [1902]	<code>proxy-read-only-backend-addresses=host:port</code> [1902]	The MySQL server host and port (read only)	

Format	Option File	Description	Removed
<code>--proxy-skip-profiling [1903]</code>	<code>proxy-skip-profiling [1903]</code>	Disable query profiling	

**Table 15.7. `mysql-proxy` Applications Options**

Format	Option File	Description
<code>--basedir=dir_name [1899]</code>	<code>basedir=dir_name [1899]</code>	The base directory prefix for paths in the configuration
<code>--daemon [1899]</code>	<code>daemon [1899]</code>	Start in daemon mode
<code>--defaults-file=file_name [1899]</code>		The configuration file to use
<code>--event-threads=count [1899]</code>	<code>event-threads=count [1899]</code>	The number of event-handling threads
<code>--keepalive [1899]</code>	<code>keepalive [1899]</code>	Try to restart the proxy if a crash occurs
<code>--log-backtrace-on-crash [1900]</code>	<code>log-backtrace-on-crash [1900]</code>	Try to invoke the debugger and generate a backtrace on crash
<code>--log-file=file_name [1900]</code>	<code>log-file=file_name [1900]</code>	The file where error messages are logged
<code>--log-level=level [1900]</code>	<code>log-level=level [1900]</code>	The logging level
<code>--log-use-syslog [1900]</code>	<code>log-use-syslog [1900]</code>	Log errors to syslog
<code>--lua-cpath=dir_name [1900]</code>	<code>lua-cpath=dir_name [1900]</code>	Set the LUA_CPATH
<code>--lua-path=dir_name [1900]</code>	<code>lua-path=dir_name [1900]</code>	Set the LUA_PATH
<code>--max-open-files=count [1901]</code>	<code>max-open-files=count [1901]</code>	The maximum number of open files to support
<code>--pid-file=file_name [1903]</code>	<code>pid-file=file_name [1903]</code>	File in which to store the process ID
<code>--plugin-dir=dir_name [1901]</code>	<code>plugin-dir=dir_name [1901]</code>	Directory containing plugin files
<code>--plugins=plugin,... [1901]</code>	<code>plugins=plugin,... [1901]</code>	List of plugins to load
<code>--user=user_name [1904]</code>	<code>user=user_name [1904]</code>	The user to use when running mysql-proxy
<code>--version [1904]</code>		Show version information

Except as noted in the following details, all of the options can be used within the configuration file by supplying the option and the corresponding value. For example:

```
[mysql-proxy]
log-file = /var/log/mysql-proxy.log
log-level = message
```

- `--help [1897]`, `-h`

Command-Line Format	
	<code>--help</code>
	<code>-h</code>

Show available help options.

- `--help-admin` [1898]

<b>Command-Line Format</b>	<code>--help-admin</code>
----------------------------	---------------------------

Show options for the admin module.

- `--help-all` [1898]

<b>Command-Line Format</b>	<code>--help-all</code>
----------------------------	-------------------------

Show all help options.

- `--help-proxy` [1898]

<b>Command-Line Format</b>	<code>--help-proxy</code>
----------------------------	---------------------------

Show options for the proxy module.

- `--admin-address=host:port` [1898]

<b>Command-Line Format</b>	<code>--admin-address=host:port</code>	
<b>Option-File Format</b>	<code>admin-address=host:port</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>:4041</code>

The host name (or IP address) and port for the administration port. The default is `localhost:4041`.

- `--admin-lua-script=file_name` [1898]

<b>Command-Line Format</b>	<code>--admin-lua-script=file_name</code>	
<b>Option-File Format</b>	<code>admin-lua-script=file_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	

The script to use for the proxy administration module.

- `--admin-password=password` [1898]

<b>Command-Line Format</b>	<code>--admin-password=password</code>	
<b>Option-File Format</b>	<code>admin-password=password</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	

The password to use to authenticate users wanting to connect to the MySQL Proxy administration module. This module uses the MySQL protocol to request a user name and password for connections.

- `--admin-username=user_name` [1898]

<b>Command-Line Format</b>	<code>--admin-username=user_name</code>	
<b>Option-File Format</b>	<code>admin-username=user_name</code>	

	<b>Permitted Values</b>	
<b>Type</b>	string	
<b>Default</b>	root	

The user name to use to authenticate users wanting to connect to the MySQL Proxy administration module. This module uses the MySQL protocol to request a user name and password for connections. The default user name is `root`.

- `--basedir=dir_name` [1899]

<b>Command-Line Format</b>	<code>--basedir=dir_name</code>	
<b>Option-File Format</b>	<code>basedir=dir_name</code>	
	<b>Permitted Values</b>	
<b>Type</b>	directory name	

The base directory to use as a prefix for all other file name configuration options. The base name should be an absolute (not relative) directory. If you specify a relative directory, `mysql-proxy` generates an error during startup.

- `--daemon` [1899]

<b>Command-Line Format</b>	<code>--daemon</code>
<b>Option-File Format</b>	<code>daemon</code>

Starts the proxy in daemon mode.

- `--defaults-file=file_name` [1899]

<b>Command-Line Format</b>	<code>--defaults-file=file_name</code>
----------------------------	----------------------------------------

The file to read for configuration options. If not specified, MySQL Proxy takes options only from the command line.

- `--event-threads=count` [1899]

<b>Command-Line Format</b>	<code>--event-threads=count</code>	
<b>Option-File Format</b>	<code>event-threads=count</code>	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	1	

The number of event threads to reserve to handle incoming requests.

- `--keepalive` [1899]

<b>Command-Line Format</b>	<code>--keepalive</code>
<b>Option-File Format</b>	<code>keepalive</code>

Create a process surrounding the main `mysql-proxy` process that attempts to restart the main `mysql-proxy` process in the event of a crash or other failure.



**Note**

The `--keepalive` [1899] option is not available on Microsoft Windows. When running as a service, `mysql-proxy` automatically restarts.

- `--log-backtrace-on-crash` [1900]

<b>Command-Line Format</b>	<code>--log-backtrace-on-crash</code>
<b>Option-File Format</b>	<code>log-backtrace-on-crash</code>

Log a backtrace to the error log and try to initialize the debugger in the event of a failure.

- `--log-file=file_name` [1900]

<b>Command-Line Format</b>	<code>--log-file=file_name</code>
<b>Option-File Format</b>	<code>log-file=file_name</code>
	<b>Permitted Values</b>
	<b>Type</b>   <code>file name</code>

The file to use to record log information. If this option is not given, `mysql-proxy` logs to the standard error output.

- `--log-level=level` [1900]

<b>Command-Line Format</b>	<code>--log-level=level</code>
<b>Option-File Format</b>	<code>log-level=level</code>
	<b>Permitted Values</b>
	<b>Type</b>   <code>enumeration</code>
	<b>Valid Values</b>   <code>error</code> <code>warning</code> <code>info</code> <code>message</code> <code>debug</code>

The log level to use when outputting error messages. Messages with that level (or lower) are output. For example, `message` level also outputs message with `info`, `warning`, and `error` levels.

- `--log-use-syslog` [1900]

<b>Command-Line Format</b>	<code>--log-use-syslog</code>
<b>Option-File Format</b>	<code>log-use-syslog</code>

Log errors to the syslog (Unix/Linux only).

- `--lua-cpath=dir_name` [1900]

<b>Command-Line Format</b>	<code>--lua-cpath=dir_name</code>
<b>Option-File Format</b>	<code>lua-cpath=dir_name</code>
	<b>Permitted Values</b>
	<b>Type</b>   <code>directory name</code>

The `LUA_CPATH` to use when loading compiled modules or libraries for Lua scripts.

- `--lua-path=dir_name` [1900]

<b>Command-Line Format</b>	<code>--lua-path=dir_name</code>
----------------------------	----------------------------------



## MySQL Proxy Command Options

<b>Option-File Format</b>	<code>lua-path=dir_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The `LUA_CPATH` to use when loading modules for Lua.

- `--max-open-files=count` [1901]

<b>Command-Line Format</b>	<code>--max-open-files=count</code>	
<b>Option-File Format</b>	<code>max-open-files=count</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>

The maximum number of open files and sockets supported by the `mysql-proxy` process. Certain scripts might require a higher value.

- `--no-proxy` [1901]

<b>Command-Line Format</b>	<code>--no-proxy</code>	
<b>Option-File Format</b>	<code>no-proxy</code>	

Disable the proxy module.

- `--plugin-dir=dir_name` [1901]

<b>Command-Line Format</b>	<code>--plugin-dir=dir_name</code>	
<b>Option-File Format</b>	<code>plugin-dir=dir_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>directory name</code>

The directory to use when loading plugins for `mysql-proxy`.

- `--plugins=plugin` [1901]

<b>Command-Line Format</b>	<code>--plugins=plugin,...</code>	
<b>Option-File Format</b>	<code>plugins=plugin,...</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

Lloads a plugin.

When using this option on the command line, you can specify the option multiple times to specify multiple plugins. For example:

```
shell> mysql-proxy --plugins=proxy --plugins=admin
```

When using the option within the configuration file, you should separate multiple plugins by commas. The equivalent of the preceding example would be:

```
...
plugins=proxy,admin
```

- `--proxy-address=host:port` [1901], `-P host:port`  
1901

<b>Command-Line Format</b>	<code>--proxy-address=host:port</code>	
----------------------------	----------------------------------------	--

	<code>-P host:port</code>
<b>Option-File Format</b>	<code>proxy-address=host:port</code>
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>
<b>Default</b>	<code>:4040</code>

The listening host name (or IP address) and port of the proxy server. The default is `:4040` (all IPs on port 4040).

- `--proxy-read-only-backend-addresses=host:port [1902]`, `-r host:port`

<b>Command-Line Format</b>	<code>--proxy-read-only-backend-addresses=host:port</code>
	<code>-r host:port</code>
<b>Option-File Format</b>	<code>proxy-read-only-backend-addresses=host:port</code>
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>

The listening host name (or IP address) and port of the proxy server for read-only connections. The default is for this information not to be set.



**Note**

Setting this value only configures the servers within the corresponding internal structure (see `proxy.global.backends [1909]`). You can determine the backend type by checking the `type` field for each connection.

You should therefore only use this option in combination with a script designed to make use of the different backend types.

When using this option on the command line, you can specify the option and the server multiple times to specify multiple backends. For example:

```
shell> mysql-proxy --proxy-read-only-backend-addresses=192.168.0.1:3306 --proxy-read-only-backend-addresses=192.168.0.2:3306
```

When using the option within the configuration file, you should separate multiple servers by commas. The equivalent of the preceding example would be:

```
...
proxy-read-only-backend-addresses = 192.168.0.1:3306,192.168.0.2:3306
```

- `--proxy-backend-addresses=host:port [1902]`, `-b host:port`

<b>Command-Line Format</b>	<code>--proxy-backend-addresses=host:port</code>
	<code>-b host:port</code>
<b>Option-File Format</b>	<code>proxy-backend-addresses=host:port</code>
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>
<b>Default</b>	<code>127.0.0.1:3306</code>

The host name (or IP address) and port of the MySQL server to connect to. You can specify multiple backend servers by supplying multiple options. Clients are connected to each backend server in round-robin fashion. For example, if you specify two servers A and B, the first client connection will go to server A; the second client connection to server B and the third client connection to server A.

When using this option on the command line, you can specify the option and the server multiple times to specify multiple backends. For example:

```
shell> mysql-proxy --proxy-backend-addresses 192.168.0.1:3306 --proxy-backend-addresses 192.168.0.2:3306
```

When using the option within the configuration file, you should separate multiple servers by commas. The equivalent of the preceding example would be:

```
...
proxy-backend-addresses = 192.168.0.1:3306,192.168.0.2:3306
```

- `--proxy-pool-no-change-user` [1903]

<b>Command-Line Format</b>	<code>--proxy-pool-no-change-user</code>
<b>Option-File Format</b>	<code>proxy-pool-no-change-user</code>

Disable use of the MySQL protocol `CHANGE_USER` command when reusing a connection from the pool of connections specified by the `proxy-backend-addresses` list.

- `--proxy-skip-profiling` [1903]

<b>Command-Line Format</b>	<code>--proxy-skip-profiling</code>
<b>Option-File Format</b>	<code>proxy-skip-profiling</code>

Disable query profiling (statistics time tracking). The default is for tracking to be enabled.

- `--proxy-fix-bug-25371` [1903]

<b>Version Removed</b>	0.8.1
<b>Command-Line Format</b>	<code>--proxy-fix-bug-25371</code>
<b>Option-File Format</b>	<code>proxy-fix-bug-25371</code>

Enable a workaround for an issue when connecting to a MySQL server later than 5.1.12 when using a MySQL client library of any earlier version.

This option was removed in `mysql-proxy` 0.8.1. Now, `mysql-proxy` returns an error message at the protocol level if it sees a `COM_CHANGE_USER` being sent to a server that has a version from 5.1.14 to 5.1.17.

- `--proxy-lua-script=file_name` [1903], `-s file_name`

<b>Command-Line Format</b>	<code>--proxy-lua-script=file_name</code>	
	<code>-s file_name</code>	
<b>Option-File Format</b>	<code>proxy-lua-script=file_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The Lua script file to be loaded. Note that the script file is not physically loaded and parsed until a connection is made. Also note that the specified Lua script is reloaded for each connection; if the content of the Lua script changes while `mysql-proxy` is running, the updated content is automatically used when a new connection is made.

- `--pid-file=file_name` [1903]

<b>Command-Line Format</b>	<code>--pid-file=file_name</code>
<b>Option-File Format</b>	<code>pid-file=file_name</code>

	<b>Permitted Values</b>	
	<b>Type</b>	file name

The name of the file in which to store the process ID.

- `--user=user_name` [1904]

<b>Command-Line Format</b>	<code>--user=user_name</code>	
<b>Option-File Format</b>	<code>user=user_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Run `mysql-proxy` as the specified `user`.

- `--version` [1904], `-V`

<b>Command-Line Format</b>	<code>--version</code>	
	<code>-V</code>	

Show the version number.

The most common usage is as a simple proxy service (that is, without additional scripting). For basic proxy operation, you must specify at least one `proxy-backend-addresses` option to specify the MySQL server to connect to by default:

```
shell> mysql-proxy --proxy-backend-addresses=MySQL.example.com:3306
```

The default proxy port is `4040`, so you can connect to your MySQL server through the proxy by specifying the host name and port details:

```
shell> mysql --host=localhost --port=4040
```

If your server requires authentication information, this will be passed through natively without alteration by `mysql-proxy`, so you must also specify the required authentication information:

```
shell> mysql --host=localhost --port=4040 \  
--user=user_name --password=password
```


You can also connect to a read-only port (which filters out `UPDATE` and `INSERT` queries) by connecting to the read-only port. By default the host name is the default, and the port is `4042`, but you can alter the host/port information by using the `proxy-read-only-backend-addresses` [1902] command-line option.

For more detailed information on how to use these command-line options, and `mysql-proxy` in general in combination with Lua scripts, see [Section 15.7.5, "Using MySQL Proxy"](#).

## 15.7.4. MySQL Proxy Scripting

You can control how MySQL Proxy manipulates and works with the queries and results that are passed on to the MySQL server through the use of the embedded Lua scripting language. You can find out more about the Lua programming language from the [Lua Web site](#).

The following diagram shows an overview of the classes exposed by MySQL Proxy.



The primary interaction between MySQL Proxy and the server is provided by defining one or more functions through an Lua script. A number of functions are supported, according to different events and operations in the communication sequence between a client and one or more backend MySQL servers:

- `connect_server()`: This function is called each time a connection is made to MySQL Proxy from a client. You can use this function during load-balancing to intercept the original connection and decide which server the client should ultimately be attached to. If you do not define a special solution, a simple round-robin style distribution is used by default.
- `read_handshake()`: This function is called when the initial handshake information is returned by the server. You can capture the handshake information returned and provide additional checks before the authorization exchange takes place.
- `read_auth()`: This function is called when the authorization packet (user name, password, default database) are submitted by the client to the server for authentication.
- `read_auth_result()`: This function is called when the server returns an authorization packet to the client indicating whether the authorization succeeded.
- `read_query()`: This function is called each time a query is sent by the client to the server. You can use this to edit and manipulate the original query, including adding new queries before and after the original statement. You can also use this function to return information directly to the client, bypassing the server, which can be useful to filter unwanted queries or queries that exceed known limits.
- `read_query_result()`: This function is called each time a result is returned from the server, providing you have manually injected queries into the query queue. If you have not explicitly injected queries within the `read_query()` function, this function is not triggered. You can use this to edit the result set, or to remove or filter the result sets generated from additional queries you injected into the queue when using `read_query()`.

The following table lists MySQL proxy and server communication functions, the supplied information, and the direction of information flow when the function is triggered.

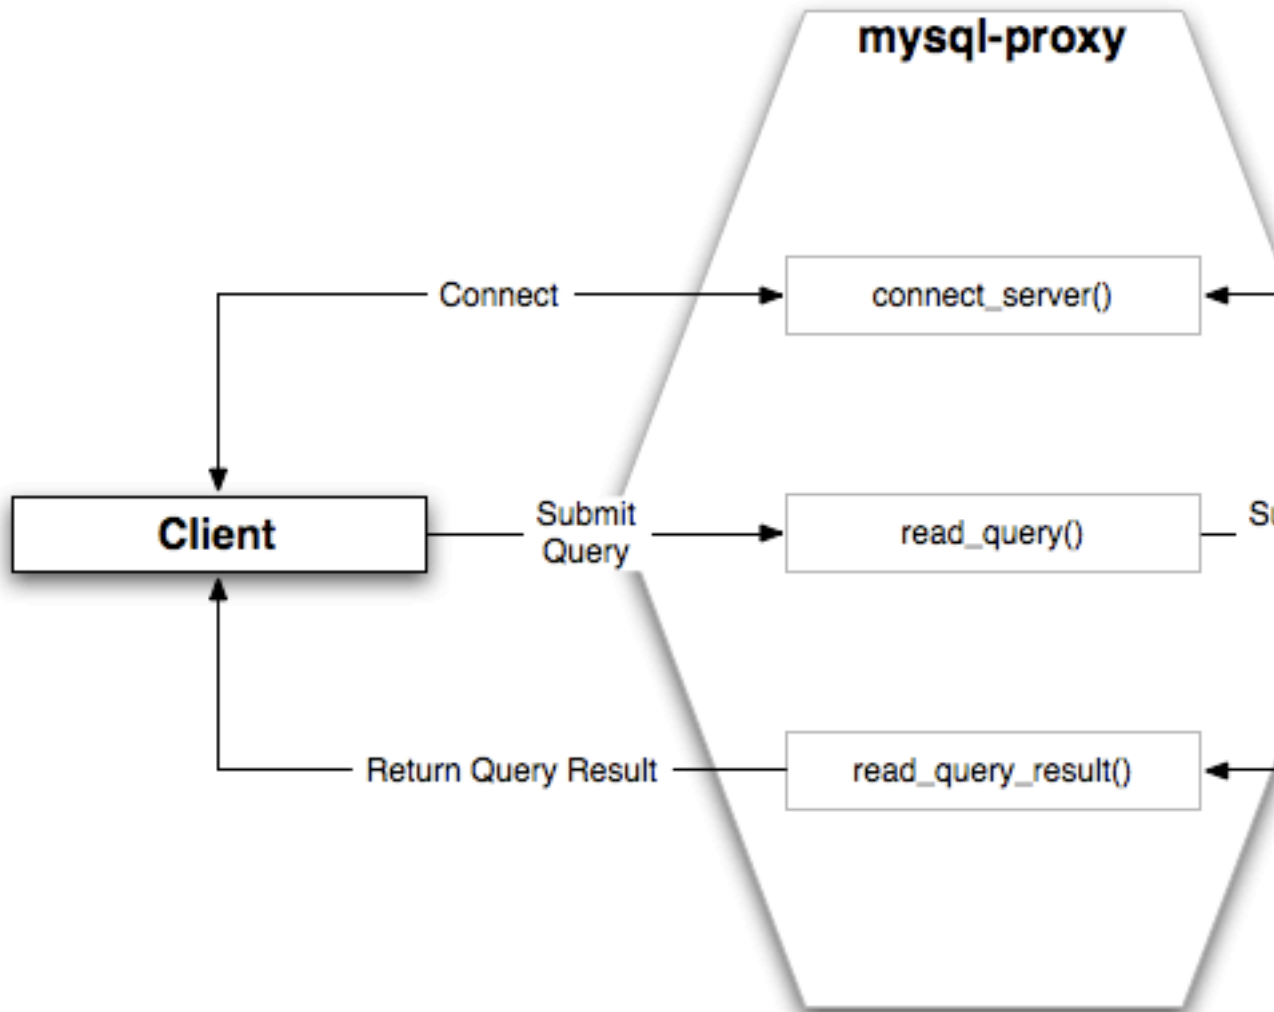
Function	Supplied Information	Direction
<code>connect_server()</code>	None	Client to Server
<code>read_handshake()</code>	None	Server to Client
<code>read_auth()</code>	None	Client to Server
<code>read_auth_result()</code>	None	Server to Client
<code>read_query()</code>	Query	Client to Server
<code>read_query_result()</code>	Query result	Server to Client

By default, all functions return a result that indicates whether the data should be passed on to the client or server (depending on the direction of the information being transferred). This return value can be overridden by explicitly returning a constant indicating that a particular response should be sent. For example, it is possible to construct result set information by hand within `read_query()` and to return the result set directly to the client without ever sending the original query to the server.

In addition to these functions, a number of built-in structures provide control over how MySQL Proxy forwards queries and returns the results by providing a simplified interface to elements such as the list of queries and the groups of result sets that are returned.

#### 15.7.4.1. Proxy Scripting Sequence During Query Injection

The following figure gives an example of how the proxy might be used when injecting queries into the query queue. Because the proxy sits between the client and MySQL server, what the proxy sends to the server, and the information that the proxy ultimately returns to the client, need not match or correlate. Once the client has connected to the proxy, the sequence shown in the following diagram occurs for each individual query sent by the client.



1. When the client submits one query to the proxy, the `read_query()` function within the proxy is triggered. The function adds the query to the query queue.
2. Once manipulation by `read_query()` has completed, the queries are submitted, sequentially, to the MySQL server.
3. The MySQL server returns the results from each query, one result set for each query submitted. The `read_query_result()` function is triggered for each result set, and each invocation can decide which result set to return to the client.

For example, you can queue additional queries into the global query queue to be processed by the server. This can be used to add statistical information by adding queries before and after the original query, changing the original query:

```
SELECT * FROM City;
```

Into a sequence of queries:

```
SELECT NOW();
SELECT * FROM City;
SELECT NOW();
```

You can also modify the original statement; for example, to add `EXPLAIN` to each statement executed to get information on how the statement was processed, again altering our original SQL statement into a number of statements:



```
SELECT * FROM City;
EXPLAIN SELECT * FROM City;
```

In both of these examples, the client would have received more result sets than expected. Regardless of how you manipulate the incoming query and the returned result, the number of queries returned by the proxy must match the number of original queries sent by the client.

You could adjust the client to handle the multiple result sets sent by the proxy, but in most cases you will want the existence of the proxy to remain transparent. To ensure that the number of queries and result sets match, you can use the MySQL Proxy `read_query_result()` to extract the additional result set information and return only the result set the client originally requested back to the client. You can achieve this by giving each query that you add to the query queue a unique ID, then filter out queries that do not match the original query ID when processing them with `read_query_result()`.

### 15.7.4.2. Internal Structures

There are a number of internal structures within the scripting element of MySQL Proxy. The primary structure is `proxy` and this provides an interface to the many common structures used throughout the script, such as connection lists and configured backend servers. Other structures, such as the incoming packet from the client and result sets are only available within the context of one of the scriptable functions.

The following table describes common attributes of the MySQL `proxy` scripting element.

Attribute	Description
<code>connection</code>	A structure containing the active client connections. For a list of attributes, see <code>proxy.connection</code> [1908].
<code>servers</code>	A structure containing the list of configured backend servers. For a list of attributes, see <code>proxy.global.backends</code> [1909].
<code>queries</code>	A structure containing the queue of queries that will be sent to the server during a single client query. For a list of attributes, see <code>proxy.queries</code> [1909].
<code>PROXY_VERSION</code>	The version number of MySQL Proxy, encoded in hex. You can use this to check that the version number supports a particular option from within the Lua script. Note that the value is encoded as a hex value, so to check the version is at least 0.5.1 you compare against <code>0x00501</code> .

#### `proxy.connection`

The `proxy.connection` object is read only, and provides information about the current connection, and is split into a `client` and `server` tables. This enables you to examine information about both the incoming client connections to the proxy (`client`), and to the backend servers (`server`).

The following table describes the client and server attributes of the `proxy.connection` object.

Attribute	Description
<code>client.default_db</code>	Default database requested by the client
<code>client.username</code>	User name used to authenticate
<code>client.scrambled_password</code>	The scrambled version of the password used to authenticate
<code>client.dst.name</code>	The combined <code>address:port</code> of the Proxy port used by this client (should match the <code>--proxy-address</code> [1901] configuration parameter)
<code>client.dst.address</code>	The IP address of the of the Proxy port used by this client
<code>client.dst.port</code>	The port number of the of the Proxy port used by this client
<code>client.src.name</code>	The combined <code>address:port</code> of the client (originating) TCP/IP endpoint

Attribute	Description
<code>client.src.address</code>	The IP address of the client (originating) TCP/IP port
<code>client.src.port</code>	The port of the client (originating) TCP/IP endpoint
<code>server.scramble_buffer</code>	The scramble buffer used to scramble the password
<code>server.mysql_version</code>	The MySQL version number of the server
<code>server.thread_id</code>	The ID of the thread handling the connection to the current server
<code>server.dst.name</code>	The combined <code>address:port</code> for the backend server for the current connection (i.e. the connection to the MySQL server)
<code>server.dst.address</code>	The address for the backend server
<code>server.dst.port</code>	The port for the backend server
<code>server.src.name</code>	The combined <code>address:port</code> for the TCP/IP endpoint used by the Proxy to connect to the backend server
<code>server.src.address</code>	The address of the endpoint for the proxy-side connection to the MySQL server
<code>server.src.port</code>	The port of the endpoint for the proxy-side connection to the MySQL server

### `proxy.global.backends`

The `proxy.global.backends` table is partially writable and contains an array of all the configured backend servers and the server metadata (IP address, status, etc.). You can determine the array index of the current connection using `proxy.connection["backend_ndx"]` which is the index into this table of the backend server being used by the active connection.

The attributes for each entry within the `proxy.global.backends` table are shown in the following table.

Attribute	Description
<code>dst.name</code>	The combined <code>address:port</code> of the backend server.
<code>dst.address</code>	The IP address of the backend server.
<code>dst.port</code>	The port of the backend server.
<code>connected_clients</code>	The number of clients currently connected.
<code>state</code>	The status of the backend server. See <a href="#">Backend State/Type Constants [1912]</a> .
<code>type</code>	The type of the backend server. You can use this to identify whether the backed was configured as a standard read/write backend, or a read-only backend. You can compare this value to the <code>proxy.BACKEND_TYPE_RW</code> and <code>proxy.BACKEND_TYPE_RO</code> .

### `proxy.queries`

The `proxy.queries` object is a queue representing the list of queries to be sent to the server. The queue is not populated automatically, but if you do not explicitly populate the queue, queries are passed on to the backend server verbatim. Also, if you do not populate the query queue by hand, the `read_query_result()` function is not triggered.

The following functions are supported for populating the `proxy.queries` object.

Function	Description
<code>append(id, packet, [options])</code>	Appends a query to the end of the query queue. The <code>id</code> is an integer identifier that you can use to recognize the query results when they are returned by the server. The packet should be a properly formatted

Function	Description
	query packet. The optional <code>options</code> should be a table containing the options specific to this packet.
<code>prepend(id, packet)</code>	Prepends a query to the query queue. The <code>id</code> is an identifier that you can use to recognize the query results when they are returned by the server. The packet should be a properly formatted query packet.
<code>reset()</code>	Empties the query queue.
<code>len()</code>	Returns the number of query packets in the queue.

For example, you could append a query packet to the `proxy.queries` queue by using the `append()`:

```
proxy.queries:append(1, packet)
```

The optional third argument to `append()` should contain the options for the packet. To have access to the result set through the `read_query_result()` function, set the `resultset_is_needed` flag to `true`:

```
proxy.queries:append( 1, packet, { resultset_is_needed = true } )
```

If that flag is `false` (the default), proxy will:

- Send the result set to the client as soon as it is received
- Reduce memory usage (because the result set is not stored internally for processing)
- Reduce latency of returning results to the client
- Pass data from server to client unaltered

The default mode is therefore quicker and useful if you only want to monitor the queries sent, and the basic statistics.

To perform any kind of manipulation on the returned data, you must set the flag to `true`, which will:

- Store the result set so that it can be processed.
- Enable modification of the result set before it is returned to the client.
- Enable you to discard the result set instead of returning it to the client.

### `proxy.response`

The `proxy.response` structure is used when you want to return your own MySQL response, instead of forwarding a packet that you have received a backend server. The structure holds the response type information, an optional error message, and the result set (rows/columns) to return.

The following table describes the attributes of the `proxy.response` structure.

Attribute	Description
<code>type</code>	The type of the response. The type must be either <code>MYSQLD_PACKET_OK</code> or <code>MYSQLD_PACKET_ERR</code> . If the <code>MYSQLD_PACKET_ERR</code> , you should set the value of the <code>mysql.response.errmsg</code> with a suitable error message.
<code>errmsg</code>	A string containing the error message that will be returned to the client.
<code>resultset</code>	A structure containing the result set information (columns and rows), identical to what would be returned when returning a results from a <code>SELECT</code> query.

When using `proxy.response` you either set `proxy.response.type` to `proxy.MYSQLD_PACKET_OK` and then build `resultset` to contain the results to return, or set

`proxy.response.type` to `proxy.MYSQLD_PACKET_ERR` and set the `proxy.response.errmsg` to a string with the error message. To send the completed result set or error message, you should return the `proxy.PROXY_SEND_RESULT` to trigger the return of the packet information.

An example of this can be seen in the `tutorial-resultset.lua` script within the MySQL Proxy package:

```
if string.lower(command) == "show" and string.lower(option) == "querycounter" then
    ---
    -- proxy.PROXY_SEND_RESULT requires
    --
    -- proxy.response.type to be either
    -- * proxy.MYSQLD_PACKET_OK or
    -- * proxy.MYSQLD_PACKET_ERR
    --
    -- for proxy.MYSQLD_PACKET_OK you need a resultset
    -- * fields
    -- * rows
    --
    -- for proxy.MYSQLD_PACKET_ERR
    -- * errmsg
    proxy.response.type = proxy.MYSQLD_PACKET_OK
    proxy.response.resultset = {
        fields = {
            { type = proxy.MYSQL_TYPE_LONG, name = "global_query_counter", },
            { type = proxy.MYSQL_TYPE_LONG, name = "query_counter", },
        },
        rows = {
            { proxy.global.query_counter, query_counter }
        }
    }

    -- we have our result, send it back
    return proxy.PROXY_SEND_RESULT
elseif string.lower(command) == "show" and string.lower(option) == "myerror" then
    proxy.response.type = proxy.MYSQLD_PACKET_ERR
    proxy.response.errmsg = "my first error"

    return proxy.PROXY_SEND_RESULT
```

### `proxy.response.resultset`

The `proxy.response.resultset` structure should be populated with the rows and columns of data to return. The structure contains the information about the entire result set, with the individual elements of the data shown in the following table.

The following table describes the attributes of the `proxy.response.resultset` structure.

Attribute	Description
<code>fields</code>	The definition of the columns being returned. This should be a dictionary structure with the <code>type</code> specifying the MySQL data type, and the <code>name</code> specifying the column name. Columns should be listed in the order of the column data that will be returned.
<code>flags</code>	A number of flags related to the result set. Valid flags include <code>auto_commit</code> (whether an automatic commit was triggered), <code>no_good_index_used</code> (the query executed without using an appropriate index), and <code>no_index_used</code> (the query executed without using any index).
<code>rows</code>	The actual row data. The information should be returned as an array of arrays. Each inner array should contain the column data, with the outer array making up the entire result set.
<code>warning_count</code>	The number of warnings for this result set.
<code>affected_rows</code>	The number of rows affected by the original statement.

Attribute	Description
<code>insert_id</code>	The last insert ID for an auto-incremented column in a table.
<code>query_status</code>	The status of the query operation. You can use the <code>MYSQLD_PACKET_OK</code> or <code>MYSQLD_PACKET_ERR</code> constants to populate this parameter.

For an example showing how to use this structure, see [proxy.response \[1910\]](#).

### Proxy Return State Constants

The following constants are used internally by the proxy to specify the response to send to the client or server. All constants are exposed as values within the main `proxy` table.

Constant	Description
<code>PROXY_SEND_QUERY</code>	Causes the proxy to send the current contents of the queries queue to the server.
<code>PROXY_SEND_RESULT</code>	Causes the proxy to send a result set back to the client.
<code>PROXY_IGNORE_RESULT</code>	Causes the proxy to drop the result set (nothing is returned to the client).

As constants, these entities are available without qualification in the Lua scripts. For example, at the end of the `read_query_result()` you might return `PROXY_IGNORE_RESULT`:

```
return proxy.PROXY_IGNORE_RESULT
```

### Packet State Constants

The following states describe the status of a network packet. These items are entries within the main `proxy` table.

Constant	Description
<code>MYSQLD_PACKET_OK</code>	The packet is OK
<code>MYSQLD_PACKET_ERR</code>	The packet contains error information
<code>MYSQLD_PACKET_RAW</code>	The packet contains raw data

### Backend State/Type Constants

The following constants are used either to define the status or type of the backend MySQL server to which the proxy is connected. These items are entries within the main `proxy` table.

Constant	Description
<code>BACKEND_STATE_UNKNOWN</code>	The current status is unknown
<code>BACKEND_STATE_UP</code>	The backend is known to be up (available)
<code>BACKEND_STATE_DOWN</code>	The backend is known to be down (unavailable)
<code>BACKEND_TYPE_UNKNOWN</code>	Backend type is unknown
<code>BACKEND_TYPE_RW</code>	Backend is available for read/write
<code>BACKEND_TYPE_RO</code>	Backend is available only for read-only use

### Server Command Constants

The values described in the table below are used in the packets exchanged between the client and server to identify the information in the rest of the packet. These items are entries within the main `proxy` table. The packet type is defined as the first character in the sent packet. For example, when intercepting packets from the client to edit or monitor a query, you would check that the first byte of the packet was of type `proxy.COM_QUERY`.

Constant	Description
COM_SLEEP	Sleep
COM_QUIT	Quit
COM_INIT_DB	Initialize database
COM_QUERY	Query
COM_FIELD_LIST	Field List
COM_CREATE_DB	Create database
COM_DROP_DB	Drop database
COM_REFRESH	Refresh
COM_SHUTDOWN	Shutdown
COM_STATISTICS	Statistics
COM_PROCESS_INFO	Process List
COM_CONNECT	Connect
COM_PROCESS_KILL	Kill
COM_DEBUG	Debug
COM_PING	Ping
COM_TIME	Time
COM_DELAYED_INSERT	Delayed insert
COM_CHANGE_USER	Change user
COM_BINLOG_DUMP	Binlog dump
COM_TABLE_DUMP	Table dump
COM_CONNECT_OUT	Connect out
COM_REGISTER_SLAVE	Register slave
COM_STMT_PREPARE	Prepare server-side statement
COM_STMT_EXECUTE	Execute server-side statement
COM_STMT_SEND_LONG_DATA	Long data
COM_STMT_CLOSE	Close server-side statement
COM_STMT_RESET	Reset statement
COM_SET_OPTION	Set option
COM_STMT_FETCH	Fetch statement
COM_DAEMON	Daemon (MySQL 5.1 only)
COM_ERROR	Error

### MySQL Type Constants

These constants are used to identify the field types in the query result data returned to clients from the result of a query. These items are entries within the main `proxy` table.

Constant	Field Type
MYSQL_TYPE_DECIMAL	Decimal
MYSQL_TYPE_NEWDECIMAL	Decimal (MySQL 5.0 or later)
MYSQL_TYPE_TINY	Tiny
MYSQL_TYPE_SHORT	Short
MYSQL_TYPE_LONG	Long

Constant	Field Type
<code>MYSQL_TYPE_FLOAT</code>	Float
<code>MYSQL_TYPE_DOUBLE</code>	Double
<code>MYSQL_TYPE_NULL</code>	Null
<code>MYSQL_TYPE_TIMESTAMP</code>	Timestamp
<code>MYSQL_TYPE_LONGLONG</code>	Long long
<code>MYSQL_TYPE_INT24</code>	Integer
<code>MYSQL_TYPE_DATE</code>	Date
<code>MYSQL_TYPE_TIME</code>	Time
<code>MYSQL_TYPE_DATETIME</code>	Datetime
<code>MYSQL_TYPE_YEAR</code>	Year
<code>MYSQL_TYPE_NEWDATE</code>	Date (MySQL 5.0 or later)
<code>MYSQL_TYPE_ENUM</code>	Enumeration
<code>MYSQL_TYPE_SET</code>	Set
<code>MYSQL_TYPE_TINY_BLOB</code>	Tiny Blob
<code>MYSQL_TYPE_MEDIUM_BLOB</code>	Medium Blob
<code>MYSQL_TYPE_LONG_BLOB</code>	Long Blob
<code>MYSQL_TYPE_BLOB</code>	Blob
<code>MYSQL_TYPE_VAR_STRING</code>	Varstring
<code>MYSQL_TYPE_STRING</code>	String
<code>MYSQL_TYPE_TINY</code>	Tiny (compatible with <code>MYSQL_TYPE_CHAR</code> )
<code>MYSQL_TYPE_ENUM</code>	Enumeration (compatible with <code>MYSQL_TYPE_INTERVAL</code> )
<code>MYSQL_TYPE_GEOMETRY</code>	Geometry
<code>MYSQL_TYPE_BIT</code>	Bit

#### 15.7.4.3. Capturing a Connection with `connect_server()`

When the proxy accepts a connection from a MySQL client, the `connect_server()` function is called.

There are no arguments to the function, but you can use and if necessary manipulate the information in the `proxy.connection` table, which is unique to each client session.

For example, if you have multiple backend servers, you can specify which server that connection should use by setting the value of `proxy.connection.backend_ndx` to a valid server number. The following code chooses between two servers based on whether the current time in minutes is odd or even:

```
function connect_server()
    print("--> a client really wants to talk to a server")
    if (tonumber(os.date("%M")) % 2 == 0) then
        proxy.connection.backend_ndx = 2
        print("Choosing backend 2")
    else
        proxy.connection.backend_ndx = 1
        print("Choosing backend 1")
    end
    print("Using " .. proxy.global.backends[proxy.connection.backend_ndx].dst.name)
end
```

This example also displays the IP address/port combination by accessing the information from the internal `proxy.global.backends` table.



#### 15.7.4.4. Examining the Handshake with `read_handshake()`

Handshake information is sent by the server to the client after the initial connection (through `connect_server()`) has been made. The handshake information contains details about the MySQL version, the ID of the thread that will handle the connection information, and the IP address of the client and server. This information is exposed through the `proxy.connection` structure.

- `proxy.connection.server.mysql_version`: The version of the MySQL server.
- `proxy.connection.server.thread_id`: The thread ID.
- `proxy.connection.server.scramble_buffer`: The password scramble buffer.
- `proxy.connection.server.dst.name`: The IP address of the server.
- `proxy.connection.client.src.name`: The IP address of the client.

For example, you can print out the handshake data and refuse clients by IP address with the following function:

```
function read_handshake()
    print("<-- let's send him some information about us")
    print("    mysql-version: " .. proxy.connection.server.mysql_version)
    print("    thread-id      : " .. proxy.connection.server.thread_id)
    print("    scramble-buf   : " .. string.format("%q", proxy.connection.server.scramble_buffer))
    print("    server-addr    : " .. proxy.connection.server.dst.name)
    print("    client-addr    : " .. proxy.connection.client.dst.name)

    if not proxy.connection.client.src.name:match("^127.0.0.1:") then
        proxy.response.type = proxy.MYSQLD_PACKET_ERR
        proxy.response.errmsg = "only local connects are allowed"

        print("we don't like this client");

        return proxy.PROXY_SEND_RESULT
    end
end
```

Note that you must return an error packet to the client by using `proxy.PROXY_SEND_RESULT`.

#### 15.7.4.5. Examining the Authentication Credentials with `read_auth()`

The `read_auth()` function is triggered when an authentication handshake is initiated by the client. In the execution sequence, `read_auth()` occurs immediately after `read_handshake()`, so the server selection has already been made, but the connection and authorization information has not yet been provided to the backend server.

You can obtain the authentication information by examining the `proxy.connection.client` structure. For more information, see [proxy.connection \[1908\]](#).

For example, you can print the user name and password supplied during authorization using:

```
function read_auth()
    print("    username      : " .. proxy.connection.client.username)
    print("    password     : " .. string.format("%q", proxy.connection.client.scrambled_password))
end
```

You can interrupt the authentication process within this function and return an error packet back to the client by constructing a new packet and returning `proxy.PROXY_SEND_RESULT`:

```
proxy.response.type = proxy.MYSQLD_PACKET_ERR
proxy.response.errmsg = "Logins are not allowed"
return proxy.PROXY_SEND_RESULT
```

#### 15.7.4.6. Accessing Authentication Information with `read_auth_result()`

The return packet from the server during authentication is captured by `read_auth_result()`. The only argument to this function is the authentication packet returned by the server. As the packet is a

raw MySQL network protocol packet, you must access the first byte to identify the packet type and contents. The `MYSQLD_PACKET_ERR` and `MYSQLD_PACKET_OK` constants can be used to identify whether the authentication was successful:

```
function read_auth_result(auth)
    local state = auth.packet:byte()

    if state == proxy.MYSQLD_PACKET_OK then
        print("<-- auth ok");
    elseif state == proxy.MYSQLD_PACKET_ERR then
        print("<-- auth failed");
    else
        print("<-- auth ... don't know: " .. string.format("%q", auth.packet));
    end
end
```

If a long-password capable client tries to authenticate to a server that supports long passwords, but the user password provided is actually short, `read_auth_result()` will be called twice. The first time, `auth.packet:byte()` will equal 254, indicating that the client should try again using the old password protocol. The second time `read_auth_result()` is called, `auth.packet:byte()` will indicate whether the authentication actually succeeded.

#### 15.7.4.7. Manipulating Queries with `read_query()`

The `read_query()` function is called once for each query submitted by the client and accepts a single argument, the query packet that was provided. To access the content of the packet, you must parse the packet contents manually.

For example, you can intercept a query packet and print out the contents using the following function definition:

```
function read_query( packet )
    if packet:byte() == proxy.COM_QUERY then
        print("we got a normal query: " .. packet:sub(2))
    end
end
```

This example checks the first byte of the packet to determine the type. If the type is `COM_QUERY` (see [Server Command Constants \[1912\]](#)), we extract the query from the packet and print it. The structure of the packet type supplied is important. In the case of a `COM_QUERY` packet, the remaining contents of the packet are the text of the query string. In this example, no changes have been made to the query or the list of queries that will ultimately be sent to the MySQL server.

To modify a query, or add new queries, you must populate the query queue (`proxy.queries`), then execute the queries that you have placed into the queue. If you do not modify the original query or the queue, the query received from the client is sent to the MySQL server verbatim.

When adding queries to the queue, you should follow these guidelines:

- The packets inserted into the queue must be valid query packets. For each packet, you must set the initial byte to the packet type. If you are appending a query, you can append the query statement to the rest of the packet.
- Once you add a query to the queue, the queue is used as the source for queries sent to the server. If you add a query to the queue to add more information, you must also add the original query to the queue or it will not be executed.
- Once the queue has been populated, you must set the return value from `read_query()` to indicate whether the query queue should be sent to the server.
- When you add queries to the queue, you should add an ID. The ID you specify is returned with the result set so that you identify each query and corresponding result set. The ID has no other purpose than as an identifier for correlating the query and result set. When operating in a passive mode, during profiling for example, you identify the original query and the corresponding result set so that the results expected by the client can be returned correctly.

- Unless your client is designed to cope with more result sets than queries, you should ensure that the number of queries from the client match the number of results sets returned to the client. Using the unique ID and removing result sets you inserted will help.

Normally, the `read_query()` and `read_query_result()` function are used in conjunction with each other to inject additional queries and remove the additional result sets. However, `read_query_result()` is only called if you populate the query queue within `read_query()`.

#### 15.7.4.8. Manipulating Results with `read_query_result()`

The `read_query_result()` is called for each result set returned by the server only if you have manually injected queries into the query queue. If you have not manipulated the query queue, this function is not called. The function supports a single argument, the result packet, which provides a number of properties:

- `id`: The ID of the result set, which corresponds to the ID that was set when the query packet was submitted to the server when using `append(id)` on the query queue. You must have set the `resultset_is_needed` flag to `append` to intercept the result set before it is returned to the client. See [proxy.queries \[1909\]](#).
- `query`: The text of the original query.
- `query_time`: The number of microseconds required to receive the first row of a result set since the query was sent to the server.
- `response_time`: The number of microseconds required to receive the last row of the result set since the query was sent to the server.
- `resultset`: The content of the result set data.

By accessing the result information from the MySQL server, you can extract the results that match the queries that you injected, return different result sets (for example, from a modified query), and even create your own result sets.

The following Lua script, for example, will output the query, followed by the query time and response time (that is, the time to execute the query and the time to return the data for the query) for each query sent to the server:

```
function read_query( packet )
    if packet:byte() == proxy.COM_QUERY then
        print("we got a normal query: " .. packet:sub(2))

        proxy.queries:append(1, packet )

        return proxy.PROXY_SEND_QUERY
    end
end

function read_query_result(inj)
    print("query-time: " .. (inj.query_time / 1000) .. "ms")
    print("response-time: " .. (inj.response_time / 1000) .. "ms")
end
```

You can access the rows of returned results from the result set by accessing the `rows` property of the `resultset` property of the result that is exposed through `read_query_result()`. For example, you can iterate over the results showing the first column from each row using this Lua fragment:

```
for row in inj.resultset.rows do
    print("injected query returned: " .. row[1])
end
```

Just like `read_query()`, `read_query_result()` can return different values for each result according to the result returned. If you have injected additional queries into the query queue, for example, remove the results returned from those additional queries and return only the results from the query originally submitted by the client.

The following example injects additional `SELECT NOW()` statements into the query queue, giving them a different ID to the ID of the original query. Within `read_query_result()`, if the ID for the injected queries is identified, we display the result row, and return the `proxy.PROXY_IGNORE_RESULT` from the function so that the result is not returned to the client. If the result is from any other query, we print out the query time information for the query and return the default, which passes on the result set unchanged. We could also have explicitly returned `proxy.PROXY_IGNORE_RESULT` to the MySQL client.

```
function read_query( packet )
    if packet:byte() == proxy.COM_QUERY then
        proxy.queries:append(2, string.char(proxy.COM_QUERY) .. "SELECT NOW()", {resultset_is_needed = true})
        proxy.queries:append(1, packet, {resultset_is_needed = true})
        proxy.queries:append(2, string.char(proxy.COM_QUERY) .. "SELECT NOW()", {resultset_is_needed = true})

        return proxy.PROXY_SEND_QUERY
    end
end

function read_query_result(inj)
    if inj.id == 2 then
        for row in inj.resultset.rows do
            print("injected query returned: " .. row[1])
        end
        return proxy.PROXY_IGNORE_RESULT
    else
        print("query-time: " .. (inj.query_time / 1000) .. "ms")
        print("response-time: " .. (inj.response_time / 1000) .. "ms")
    end
end
```

For further examples, see [Section 15.7.5, “Using MySQL Proxy”](#).

## 15.7.5. Using MySQL Proxy

There are a number of different ways to use MySQL Proxy. At the most basic level, you can allow MySQL Proxy to pass queries from clients to a single server. To use MySQL Proxy in this mode, you just have to specify on the command line the backend server to which the proxy should connect:

```
shell> mysql-proxy --proxy-backend-addresses=sakila:3306
```

If you specify multiple backend MySQL servers, the proxy connects each client to each server in a round-robin fashion. Suppose that you have two MySQL servers, A and B. The first client to connect is connected to server A, the second to server B, the third to server A. For example:

```
shell> mysql-proxy \
--proxy-backend-addresses=narcissus:3306 \
--proxy-backend-addresses=nostromo:3306
```

When you specify multiple servers in this way, the proxy automatically identifies when a MySQL server has become unavailable and marks it accordingly. New connections are automatically attached to a server that is available, and a warning is reported to the standard output from `mysql-proxy`:

```
network-mysqld.c.367: connect(nostromo:3306) failed: Connection refused
network-mysqld-proxy.c.2405: connecting to backend (nostromo:3306) failed, marking it as down for ...
```

Lua scripts enable a finer level of control, both over the connections and their distribution and how queries and result sets are processed. When using an Lua script, you must specify the name of the script on the command line using the `--proxy-lua-script [1903]` option:

```
shell> mysql-proxy --proxy-lua-script=mc.lua --proxy-backend-addresses=sakila:3306
```

When you specify a script, the script is not executed until a connection is made. This means that faults with the script are not raised until the script is executed. Script faults will not affect the distribution of queries to backend MySQL servers.

**Note**

Because a script is not read until the connection is made, you can modify the contents of the Lua script file while the proxy is still running and the modified script is automatically used for the next connection. This ensures that MySQL Proxy remains available because it need not be restarted for the changes to take effect.

### 15.7.5.1. Using the Administration Interface

The `mysql-proxy` administration interface can be accessed using any MySQL client using the standard protocols. You can use the administration interface to gain information about the proxy server as a whole - standard connections to the proxy are isolated to operate as if you were connected directly to the backend MySQL server.

In `mysql-proxy` 0.8.0 and earlier, a rudimentary interface was built into the proxy. In later versions this was replaced so that you must specify an administration script to be used when users connect to the administration interface.

To use the administration interface, specify the user name and password required to connect to the admin service, using the `--admin-username` [1898] and `--admin-password` [1898] options. You must also specify the Lua script to be used as the interface to the administration service by using the `admin-lua-script` [1898] script option to point to a Lua script.

For example, you can create a basic interface to the internal components of the `mysql-proxy` system using the following script, written by Diego Medina:

```
--[[
    Copyright 2008, 2010, Oracle and/or its affiliates. All rights reserved.

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; version 2 of the License.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
--]]

-- admin.lua

--[[
    See http://www.chriscalender.com/?p=41
    (Thanks to Chris Calender)
    See http://datacharmer.blogspot.com/2009/01/mysql-proxy-is-back.html
    (Thanks Giuseppe Maxia)
--]]

function set_error(errmsg)
    proxy.response = {
        type = proxy.MYSQLD_PACKET_ERR,
        errmsg = errmsg or "error"
    }
end

function read_query(packet)
    if packet:byte() ~= proxy.COM_QUERY then
        set_error("[admin] we only handle text-based queries (COM_QUERY)")
        return proxy.PROXY_SEND_RESULT
    end
end
```

```

end

local query = packet:sub(2)
local rows = { }
local fields = { }

-- try to match the string up to the first non-alphanum
local f_s, f_e, command = string.find(packet, "^%s*(%w+)", 2)
local option

if f_e then
    -- if that match, take the next sub-string as option
    f_s, f_e, option = string.find(packet, "^%s+(%w+)", f_e + 1)
end

-- we got our commands, execute it
if command == "show" and option == "querycounter" then
    ---
    -- proxy.PROXY_SEND_RESULT requires
    --
    -- proxy.response.type to be either
    -- * proxy.MYSQLD_PACKET_OK or
    -- * proxy.MYSQLD_PACKET_ERR
    --
    -- for proxy.MYSQLD_PACKET_OK you need a resultset
    -- * fields
    -- * rows
    --
    -- for proxy.MYSQLD_PACKET_ERR
    -- * errmsg
    proxy.response.type = proxy.MYSQLD_PACKET_OK
    proxy.response.resultset = {
        fields = {
            { type = proxy.MYSQL_TYPE_LONG, name = "query_counter", },
        },
        rows = {
            { proxy.global.query_counter }
        }
    }

    -- we have our result, send it back
    return proxy.PROXY_SEND_RESULT
elseif command == "show" and option == "myerror" then
    proxy.response.type = proxy.MYSQLD_PACKET_ERR
    proxy.response.errmsg = "my first error"

    return proxy.PROXY_SEND_RESULT

elseif string.sub(packet, 2):lower() == 'select help' then
    return show_process_help()

elseif string.sub(packet, 2):lower() == 'show proxy processlist' then
    return show_process_table()

elseif query == "SELECT * FROM backends" then
    fields = {
        { name = "backend_ndx",
          type = proxy.MYSQL_TYPE_LONG },

        { name = "address",
          type = proxy.MYSQL_TYPE_STRING },
        { name = "state",
          type = proxy.MYSQL_TYPE_STRING },
        { name = "type",
          type = proxy.MYSQL_TYPE_STRING },
    }

    for i = 1, #proxy.global.backends do
        local b = proxy.global.backends[i]

        rows[#rows + 1] = {
            i, b.dst.name, b.state, b.type
        }
    end

```

```

    }
    end
else
    set_error()
    return proxy.PROXY_SEND_RESULT
end

proxy.response = {
    type = proxy.MYSQLD_PACKET_OK,
    resultset = {
        fields = fields,
        rows = rows
    }
}
return proxy.PROXY_SEND_RESULT
end

function make_dataset (header, dataset)
    proxy.response.type = proxy.MYSQLD_PACKET_OK

    proxy.response.resultset = {
        fields = {},
        rows = {}
    }
    for i,v in pairs (header) do
        table.insert(proxy.response.resultset.fields, {type = proxy.MYSQL_TYPE_STRING, name = v})
    end
    for i,v in pairs (dataset) do
        table.insert(proxy.response.resultset.rows, v )
    end
    return proxy.PROXY_SEND_RESULT
end

function show_process_table()
    local dataset = {}
    local header = { 'Id', 'IP Address', 'Time' }
    local rows = {}
    for t_i, t_v in pairs (proxy.global.process) do
        for s_i, s_v in pairs ( t_v ) do
            table.insert(rows, { t_i, s_v.ip, os.date('%c',s_v.ts) })
        end
    end
    return make_dataset(header,rows)
end

function show_process_help()
    local dataset = {}
    local header = { 'command', 'description' }
    local rows = {
        {'SELECT HELP', 'This command.'},
        {'SHOW PROXY PROCESSLIST', 'Show all connections and their true IP Address.'},
    }
    return make_dataset(header,rows)
end

function dump_process_table()
    proxy.global.initialize_process_table()
    print('current contents of process table')
    for t_i, t_v in pairs (proxy.global.process) do
        print ('session id: ', t_i)
        for s_i, s_v in pairs ( t_v ) do
            print ( '\t', s_i, s_v.ip, s_v.ts )
        end
    end
    print ('---END PROCESS TABLE---')
end

--[ Help

we use a simple string-match to split commands are word-boundaries

```



```
mysql> show querycounter
```

```
is split into  
command = "show"  
option  = "querycounter"
```

spaces are ignored, the case has to be as is.

```
mysql> show myerror
```

returns a error-packet

```
--]]
```

The script works in combination with a main proxy script, [reporter.lua](#):

```
--[[
```

```
Copyright 2008, 2010, Oracle and/or its affiliates. All rights reserved.
```

```
This program is free software; you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation; version 2 of the License.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program; if not, write to the Free Software  
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
```

```
--]]
```

```
-- reporter.lua
```

```
--[[
```

```
See http://www.chriscalender.com/?p=41  
(Thanks to Chris Calender)  
See http://datacharmer.blogspot.com/2009/01/mysql-proxy-is-back.html  
(Thanks Giuseppe Maxia)
```

```
--]]
```

```
proxy.global.query_counter = proxy.global.query_counter or 0
```

```
function proxy.global.initialize_process_table()  
    if proxy.global.process == nil then  
        proxy.global.process = {}  
    end  
    if proxy.global.process[proxy.connection.server.thread_id] == nil then  
        proxy.global.process[proxy.connection.server.thread_id] = {}  
    end  
end
```

```
function read_auth_result( auth )  
    local state = auth.packet:byte()  
    if state == proxy.MYSQLD_PACKET_OK then  
        proxy.global.initialize_process_table()  
        table.insert( proxy.global.process[proxy.connection.server.thread_id],  
            { ip = proxy.connection.client.src.name, ts = os.time() } )  
    end  
end
```

```
function disconnect_client()  
    local connection_id = proxy.connection.server.thread_id  
    if connection_id then  
        -- client has disconnected, set this to nil  
        proxy.global.process[connection_id] = nil  
    end  
end
```

```

end

---
-- read_query() can return a resultset
--
-- You can use read_query() to return a result-set.
--
-- @param packet the mysql-packet sent by the client
--
-- @return
-- * nothing to pass on the packet as is,
-- * proxy.PROXY_SEND_QUERY to send the queries from the proxy.queries queue
-- * proxy.PROXY_SEND_RESULT to send your own result-set
--
function read_query( packet )
    -- a new query came in in this connection
    -- using proxy.global.* to make it available to the admin plugin
    proxy.global.query_counter = proxy.global.query_counter + 1
end

```

To use the script, save the first script to a file ([admin.lua](#) in the following example) and the other to [reporter.lua](#), then run `mysql-proxy` specifying the admin script and a backend MySQL server:

```

shell> mysql-proxy --admin-lua-script=admin.lua --admin-password=password \ »
--admin-username=root --proxy-backend-addresses=127.0.0.1:3306 -proxy-lua-script=reporter.lua

```

In a different window, connect to the MySQL server through the proxy:

```

shell> mysql --user=root --password=password --port=4040
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1798669
Server version: 5.0.70-log Gentoo Linux mysql-5.0.70-r1

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>

```

In another different window, connect to the `mysql-proxy` admin service using the specified user name and password:

```

shell> mysql --user=root --password=password --port=4041 --host=localhost
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1
Server version: 5.0.99-agent-admin

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>

```

To monitor the status of the proxy, ask for a list of the current active processes:

```

mysql> show proxy processlist;
+-----+-----+-----+
| Id      | IP Address          | Time                |
+-----+-----+-----+
| 1798669 | 192.168.0.112:52592 | Wed Jan 20 16:58:00 2010 |
+-----+-----+-----+
1 row in set (0.00 sec)

mysql>

```

For more information on the example, see [MySQL Proxy Admin Example](#).

## 15.7.6. MySQL Proxy FAQ

### Questions

- [16.7.6.1: \[1925\]](#) In load balancing, how can I separate reads from writes?

- [16.7.6.2: \[1925\]](#) How do I use a socket with MySQL Proxy? Proxy change logs mention that support for UNIX sockets has been added.
- [16.7.6.3: \[1925\]](#) Can I use MySQL Proxy with all versions of MySQL?
- [16.7.6.4: \[1925\]](#) Can I run MySQL Proxy as a daemon?
- [16.7.6.5: \[1925\]](#) Do proxy applications run on a separate server? If not, what is the overhead incurred by Proxy on the DB server side?
- [16.7.6.6: \[1925\]](#) With load balancing, what happens to transactions? Are all queries sent to the same server?
- [16.7.6.7: \[1925\]](#) Is it possible to use MySQL Proxy with updating a Lucene index (or Solr) by making TCP calls to that server to update?
- [16.7.6.8: \[1926\]](#) Is the system context switch expensive, how much overhead does the Lua script add?
- [16.7.6.9: \[1926\]](#) How much latency does a proxy add to a connection?
- [16.7.6.10: \[1926\]](#) Do you have to make one large script and call it at proxy startup, can I change scripts without stopping and restarting (interrupting) the proxy?
- [16.7.6.11: \[1926\]](#) If MySQL Proxy has to live on same machine as MySQL, are there any tuning considerations to ensure both perform optimally?
- [16.7.6.12: \[1926\]](#) I currently use SQL Relay for efficient connection pooling with a number of Apache processes connecting to a MySQL server. Can MySQL Proxy currently accomplish this? My goal is to minimize connection latency while keeping temporary tables available.
- [16.7.6.13: \[1926\]](#) Are these reserved function names (for example, `error_result()`) that get automatically called?
- [16.7.6.14: \[1926\]](#) As the script is re-read by MySQL Proxy, does it cache this or is it looking at the file system with each request?
- [16.7.6.15: \[1926\]](#) Given that there is a `connect_server()` function, can a Lua script link up with multiple servers?
- [16.7.6.16: \[1926\]](#) Is the MySQL Proxy an API?
- [16.7.6.17: \[1926\]](#) The global namespace variable example with quotas does not persist after a reboot, is that correct?
- [16.7.6.18: \[1926\]](#) Can MySQL Proxy handle SSL connections?
- [16.7.6.19: \[1926\]](#) Could MySQL Proxy be used to capture passwords?
- [16.7.6.20: \[1927\]](#) Are there tools for isolating problems? How can someone figure out whether a problem is in the client, the database, or the proxy?
- [16.7.6.21: \[1927\]](#) Is MySQL Proxy similar to what is provided by Java connection pools?
- [16.7.6.22: \[1927\]](#) So authentication with connection pooling has to be done at every connection? What is the authentication latency?
- [16.7.6.23: \[1927\]](#) If you have multiple databases on the same box, can you use proxy to connect to databases on default port 3306?
- [16.7.6.24: \[1927\]](#) What about caching the authorization information so clients connecting are given back-end connections that were established with identical authorization information, thus saving a few more round trips?

- [16.7.6.25: \[1927\]](#) Is there any big web site using MySQL Proxy? For what purpose and what transaction rate have they achieved?
- [16.7.6.26: \[1927\]](#) How does MySQL Proxy compare to DBSLayer?
- [16.7.6.27: \[1927\]](#) I tried using MySQL Proxy without any Lua script to try a round-robin type load balancing. In this case, if the first database in the list is down, MySQL Proxy would not connect the client to the second database in the list.
- [16.7.6.28: \[1927\]](#) Is it “safe” to use `LuaSocket` with proxy scripts?
- [16.7.6.29: \[1927\]](#) How different is MySQL Proxy from DBCP (Database connection pooling) for Apache in terms of connection pooling?
- [16.7.6.30: \[1927\]](#) MySQL Proxy can handle about 5000 connections, what is the limit on a MySQL server?
- [16.7.6.31: \[1928\]](#) Would the Java-only connection pooling solution work for multiple web servers? With this, I would assume that you can pool across many web servers at once?

### Questions and Answers

#### 16.7.6.1: In load balancing, how can I separate reads from writes?

There is no automatic separation of queries that perform reads or writes to the different backend servers. However, you can specify to `mysql-proxy` that one or more of the “backend” MySQL servers are read only.

```
shell> mysql-proxy \  
--proxy-backend-addresses=10.0.1.2:3306 \  
--proxy-read-only-backend-addresses=10.0.1.3:3306 &
```

#### 16.7.6.2: How do I use a socket with MySQL Proxy? Proxy change logs mention that support for UNIX sockets has been added.

Specify the path to the socket:

```
--proxy-backend-addresses=/path/to/socket
```

#### 16.7.6.3: Can I use MySQL Proxy with all versions of MySQL?

MySQL Proxy is designed to work with MySQL 5.0 or higher, and supports the MySQL network protocol for 5.0 and higher.

#### 16.7.6.4: Can I run MySQL Proxy as a daemon?

Use the `--daemon [1899]` option. To keep track of the process ID, the daemon can be started with the `--pid-file=file [1903]` option to save the PID to a known file name. On version 0.5.x, the Proxy cannot be started natively as a daemon.

#### 16.7.6.5: Do proxy applications run on a separate server? If not, what is the overhead incurred by Proxy on the DB server side?

You can run the proxy on the application server, on its own box, or on the DB-server depending on the use case.

#### 16.7.6.6: With load balancing, what happens to transactions? Are all queries sent to the same server?

Without any special customization the whole connection is sent to the same server. That keeps the whole connection state intact.

#### 16.7.6.7: Is it possible to use MySQL Proxy with updating a Lucene index (or Solr) by making TCP calls to that server to update?

Yes, but it is not advised for now.

**16.7.6.8: Is the system context switch expensive, how much overhead does the Lua script add?**

Lua is fast and the overhead should be small enough for most applications. The raw packet overhead is around 400 microseconds.

**16.7.6.9: How much latency does a proxy add to a connection?**

In the range of 400 microseconds per request.

**16.7.6.10: Do you have to make one large script and call it at proxy startup, can I change scripts without stopping and restarting (interrupting) the proxy?**

You can just change the script and the proxy will reload it when a client connects.

**16.7.6.11: If MySQL Proxy has to live on same machine as MySQL, are there any tuning considerations to ensure both perform optimally?**

MySQL Proxy can live on any box: application, database, or its own box. MySQL Proxy uses comparatively little CPU or RAM, with negligible additional requirements or overhead.

**16.7.6.12: I currently use SQL Relay for efficient connection pooling with a number of Apache processes connecting to a MySQL server. Can MySQL Proxy currently accomplish this? My goal is to minimize connection latency while keeping temporary tables available.**

Yes.

**16.7.6.13: Are these reserved function names (for example, `error_result()`) that get automatically called?**

Only functions and values starting with `proxy.*` are provided by the proxy. All others are user provided.

**16.7.6.14: As the script is re-read by MySQL Proxy, does it cache this or is it looking at the file system with each request?**

It looks for the script at client-connect and reads it if it has changed, otherwise it uses the cached version.

**16.7.6.15: Given that there is a `connect_server()` function, can a Lua script link up with multiple servers?**

MySQL Proxy provides some tutorials in the source package; one is `examples/tutorial-keepalive.lua`.

**16.7.6.16: Is the MySQL Proxy an API?**

No, MySQL Proxy is an application that forwards packets from a client to a server using the MySQL network protocol. The MySQL Proxy provides a API allowing you to change its behavior.

**16.7.6.17: The global namespace variable example with quotas does not persist after a reboot, is that correct?**

Yes. If you restart the proxy, you lose the results, unless you save them in a file.

**16.7.6.18: Can MySQL Proxy handle SSL connections?**

No, being the man-in-the-middle, Proxy cannot handle encrypted sessions because it cannot share the SSL information.

**16.7.6.19: Could MySQL Proxy be used to capture passwords?**

The MySQL network protocol does not allow passwords to be sent in cleartext, all you could capture is the encrypted version.

**16.7.6.20: Are there tools for isolating problems? How can someone figure out whether a problem is in the client, the database, or the proxy?**

You can set a debug script in the proxy, which is an exceptionally good tool for this purpose. You can see very clearly which component is causing the problem, if you set the right breakpoints.

**16.7.6.21: Is MySQL Proxy similar to what is provided by Java connection pools?**

Yes and no. Java connection pools are specific to Java applications, MySQL Proxy works with any client API that talks the MySQL network protocol. Also, connection pools do not provide any functionality for intelligently examining the network packets and modifying the contents.

**16.7.6.22: So authentication with connection pooling has to be done at every connection? What is the authentication latency?**

You can skip the round-trip and use the connection as it was added to the pool. As long as the application cleans up the temporary tables it used. The overhead is (as always) around 400 microseconds.

**16.7.6.23: If you have multiple databases on the same box, can you use proxy to connect to databases on default port 3306?**

Yes, MySQL Proxy can listen on any port, provided that none of the MySQL servers are listening on the same port.

**16.7.6.24: What about caching the authorization information so clients connecting are given back-end connections that were established with identical authorization information, thus saving a few more round trips?**

There is an `--proxy-pool-no-change-user` [1903] option that provides this functionality.

**16.7.6.25: Is there any big web site using MySQL Proxy? For what purpose and what transaction rate have they achieved?**

Yes, [gaiaonline](#). They have tested MySQL Proxy and seen it handle 2400 queries per second through the proxy.

**16.7.6.26: How does MySQL Proxy compare to DBSLayer?**

DBSLayer is a REST->MySQL tool, MySQL Proxy is transparent to your application. No change to the application is needed.

**16.7.6.27: I tried using MySQL Proxy without any Lua script to try a round-robin type load balancing. In this case, if the first database in the list is down, MySQL Proxy would not connect the client to the second database in the list.**

This issue is fixed in version 0.7.0.

**16.7.6.28: Is it “safe” to use LuaSocket with proxy scripts?**

You can, but it is not advised because it may block.

**16.7.6.29: How different is MySQL Proxy from DBCP (Database connection pooling) for Apache in terms of connection pooling?**

Connection Pooling is just one use case of the MySQL Proxy. You can use it for a lot more and it works in cases where you cannot use DBCP (for example, if you do not have Java).

**16.7.6.30: MySQL Proxy can handle about 5000 connections, what is the limit on a MySQL server?**

The server limit is given by the value of the `max_connections` [503] system variable. The default value is version dependent.

**16.7.6.31: Would the Java-only connection pooling solution work for multiple web servers? With this, I would assume that you can pool across many web servers at once?**

Yes. But you can also start one proxy on each application server to get a similar behavior as you have it already.



---

# Chapter 16. Replication

## Table of Contents

16.1. Replication Configuration .....	1930
16.1.1. How to Set Up Replication .....	1931
16.1.2. Replication Formats .....	1940
16.1.3. Replication with Global Transaction Identifiers .....	1947
16.1.4. Replication and Binary Logging Options and Variables .....	1955
16.1.5. Common Replication Administration Tasks .....	2024
16.2. Replication Implementation .....	2026
16.2.1. Replication Implementation Details .....	2027
16.2.2. Replication Relay and Status Logs .....	2028
16.2.3. How Servers Evaluate Replication Filtering Rules .....	2034
16.3. Replication Solutions .....	2040
16.3.1. Using Replication for Backups .....	2041
16.3.2. Using Replication with Different Master and Slave Storage Engines .....	2044
16.3.3. Using Replication for Scale-Out .....	2045
16.3.4. Replicating Different Databases to Different Slaves .....	2046
16.3.5. Improving Replication Performance .....	2047
16.3.6. Switching Masters During Failover .....	2048
16.3.7. Setting Up Replication Using SSL .....	2051
16.3.8. Semisynchronous Replication .....	2052
16.3.9. Delayed Replication .....	2057
16.4. Replication Notes and Tips .....	2057
16.4.1. Replication Features and Issues .....	2057
16.4.2. Replication Compatibility Between MySQL Versions .....	2080
16.4.3. Upgrading a Replication Setup .....	2081
16.4.4. Troubleshooting Replication .....	2082
16.4.5. How to Report Replication Bugs or Problems .....	2083

Replication enables data from one MySQL database server (the master) to be replicated to one or more MySQL database servers (the slaves). Replication is asynchronous by default - slaves need not to be connected permanently to receive updates from the master. This means that updates can occur over long-distance connections and even over temporary or intermittent connections such as a dial-up service. Depending on the configuration, you can replicate all databases, selected databases, or even selected tables within a database.

For answers to some questions often asked by those who are new to MySQL Replication, see [Section B.13, “MySQL 5.6 FAQ: Replication”](#).

The target uses for replication in MySQL include:

- Scale-out solutions - spreading the load among multiple slaves to improve performance. In this environment, all writes and updates must take place on the master server. Reads, however, may take place on one or more slaves. This model can improve the performance of writes (since the master is dedicated to updates), while dramatically increasing read speed across an increasing number of slaves.
- Data security - because data is replicated to the slave, and the slave can pause the replication process, it is possible to run backup services on the slave without corrupting the corresponding master data.
- Analytics - live data can be created on the master, while the analysis of the information can take place on the slave without affecting the performance of the master.

- Long-distance data distribution - if a branch office would like to work with a copy of your main data, you can use replication to create a local copy of the data for their use without requiring permanent access to the master.

Replication in MySQL features support for one-way, asynchronous replication, in which one server acts as the master, while one or more other servers act as slaves. This is in contrast to the *synchronous* replication which is a characteristic of MySQL Cluster (see [MySQL Cluster NDB 7.2](#)). In MySQL 5.6, an interface to semisynchronous replication is supported in addition to the built-in asynchronous replication. With semisynchronous replication, a commit performed on the master side blocks before returning to the session that performed the transaction until at least one slave acknowledges that it has received and logged the events for the transaction. See [Section 16.3.8, “Semisynchronous Replication”](#). MySQL 5.6 also supports delayed replication such that a slave server deliberately lags behind the master by at least a specified amount of time. See [Section 16.3.9, “Delayed Replication”](#).

There are a number of solutions available for setting up replication between two servers, but the best method to use depends on the presence of data and the engine types you are using. For more information on the available options, see [Section 16.1.1, “How to Set Up Replication”](#).

There are two core types of replication format, Statement Based Replication (SBR), which replicates entire SQL statements, and Row Based Replication (RBR), which replicates only the changed rows. You may also use a third variety, Mixed Based Replication (MBR). For more information on the different replication formats, see [Section 16.1.2, “Replication Formats”](#). In MySQL 5.6, statement-based format is the default.

MySQL 5.6.5 and later supports transactional replication based on *global transaction identifiers* (GTIDs). When using this type of replication, it is not necessary to work directly with log files or positions within these files, which greatly simplifies many common replication tasks. Because replication using GTIDs is entirely transactional, consistency between master and slave is guaranteed as long as all transactions committed on the master have also been applied on the slave. For more information about GTIDs and GTID-based replication, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

Replication is controlled through a number of different options and variables. These control the core operation of the replication, timeouts, and the databases and filters that can be applied on databases and tables. For more information on the available options, see [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).

You can use replication to solve a number of different problems, including problems with performance, supporting the backup of different databases, and as part of a larger solution to alleviate system failures. For information on how to address these issues, see [Section 16.3, “Replication Solutions”](#).

For notes and tips on how different data types and statements are treated during replication, including details of replication features, version compatibility, upgrades, and problems and their resolution, including an FAQ, see [Section 16.4, “Replication Notes and Tips”](#).

For detailed information on the implementation of replication, how replication works, the process and contents of the binary log, background threads and the rules used to decide how statements are recorded and replication, see [Section 16.2, “Replication Implementation”](#).

## 16.1. Replication Configuration

Replication between servers in MySQL is based on the binary logging mechanism. The MySQL instance operating as the master (the source of the database changes) writes updates and changes as “events” to the binary log. The information in the binary log is stored in different logging formats according to the database changes being recorded. Slaves are configured to read the binary log from the master and to execute the events in the binary log on the slave’s local database.

The master is “dumb” in this scenario. Once binary logging has been enabled, all statements are recorded in the binary log. Each slave receives a copy of the entire contents of the binary log. It is the responsibility of the slave to decide which statements in the binary log should be executed; you cannot configure the master to log only certain events. If you do not specify otherwise, all events in the master

binary log are executed on the slave. If required, you can configure the slave to process only events that apply to particular databases or tables.

Each slave keeps a record of the binary log coordinates: The file name and position within the file that it has read and processed from the master. This means that multiple slaves can be connected to the master and executing different parts of the same binary log. Because the slaves control this process, individual slaves can be connected and disconnected from the server without affecting the master's operation. Also, because each slave remembers the position within the binary log, it is possible for slaves to be disconnected, reconnect and then “catch up” by continuing from the recorded position.

Both the master and each slave must be configured with a unique ID (using the `server-id` [1955] option). In addition, each slave must be configured with information about the master host name, log file name, and position within that file. These details can be controlled from within a MySQL session using the `CHANGE MASTER TO` statement on the slave. The details are stored within the slave's `master.info` file.

This section describes the setup and configuration required for a replication environment, including step-by-step instructions for creating a new replication environment. The major components of this section are:

- For a guide to setting up two or more servers for replication, [Section 16.1.1, “How to Set Up Replication”](#), deals with the configuration of the systems and provides methods for copying data between the master and slaves.
- Events in the binary log are recorded using a number of formats. These are referred to as statement-based replication (SBR) or row-based replication (RBR). A third type, mixed-format replication (MIXED), uses SBR or RBR replication automatically to take advantage of the benefits of both SBR and RBR formats when appropriate. The different formats are discussed in [Section 16.1.2, “Replication Formats”](#).
- Detailed information on the different configuration options and variables that apply to replication is provided in [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).
- Once started, the replication process should require little administration or monitoring. However, for advice on common tasks that you may want to execute, see [Section 16.1.5, “Common Replication Administration Tasks”](#).

### 16.1.1. How to Set Up Replication

This section describes how to set up complete replication of a MySQL server. There are a number of different methods for setting up replication, and the exact method to use depends on how you are setting up replication, and whether you already have data within your master database.

There are some generic tasks that are common to all replication setups:

- On the master, you must enable binary logging and configure a unique server ID. This might require a server restart. See [Section 16.1.1.1, “Setting the Replication Master Configuration”](#).
- On each slave that you want to connect to the master, you must configure a unique server ID. This might require a server restart. See [Section 16.1.1.2, “Setting the Replication Slave Configuration”](#).
- You may want to create a separate user that will be used by your slaves to authenticate with the master to read the binary log for replication. The step is optional. See [Section 16.1.1.3, “Creating a User for Replication”](#).
- Before creating a data snapshot or starting the replication process, you should record the position of the binary log on the master. You will need this information when configuring the slave so that the slave knows where within the binary log to start executing events. See [Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”](#).
- If you already have data on your master and you want to use it to synchronize your slave, you will need to create a data snapshot. You can create a snapshot using `mysqldump` (see

Section 16.1.1.5, “Creating a Data Snapshot Using `mysqldump`”) or by copying the data files directly (see Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”).

- You will need to configure the slave with settings for connecting to the master, such as the host name, login credentials, and binary log file name and position. See Section 16.1.1.10, “Setting the Master Configuration on the Slave”.

Once you have configured the basic options, you will need to follow the instructions for your replication setup. A number of alternatives are provided:

- If you are establishing a new MySQL master and one or more slaves, you need only set up the configuration, as you have no data to exchange. For guidance on setting up replication in this situation, see Section 16.1.1.7, “Setting Up Replication with New Master and Slaves”.
- If you are already running a MySQL server, and therefore already have data that must be transferred to your slaves before replication starts, have not previously configured the binary log and are able to shut down your MySQL server for a short period during the process, see Section 16.1.1.8, “Setting Up Replication with Existing Data”.
- If you are adding slaves to an existing replication environment, you can set up the slaves without affecting the master. See Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”.

If you will be administering MySQL replication servers, we suggest that you read this entire chapter through and try all statements mentioned in Section 13.4.1, “SQL Statements for Controlling Master Servers”, and Section 13.4.2, “SQL Statements for Controlling Slave Servers”. You should also familiarize yourself with the replication startup options described in Section 16.1.4, “Replication and Binary Logging Options and Variables”.



#### Note

Note that certain steps within the setup process require the `SUPER` [756] privilege. If you do not have this privilege, it might not be possible to enable replication.

### 16.1.1.1. Setting the Replication Master Configuration

On a replication master, you must enable binary logging and establish a unique server ID. If this has not already been done, this part of master setup requires a server restart.

Binary logging *must* be enabled on the master because the binary log is the basis for sending data changes from the master to its slaves. If binary logging is not enabled, replication will not be possible.

Each server within a replication group must be configured with a unique server ID. This ID is used to identify individual servers within the group, and must be a positive integer between 1 and  $(2^{32})-1$ . How you organize and select the numbers is entirely up to you.

To configure the binary log and server ID options, you will need to shut down your MySQL server and edit the `my.cnf` or `my.ini` file. Add the following options to the configuration file within the `[mysqld]` section. If these options already exist, but are commented out, uncomment the options and alter them according to your needs. For example, to enable binary logging using a log file name prefix of `mysql-bin`, and configure a server ID of 1, use these lines:

```
[mysqld]
log-bin=mysql-bin
server-id=1
```

After making the changes, restart the server.



#### Note

If you omit `server-id` [1955] (or set it explicitly to its default value of 0), a master refuses connections from all slaves.

**Note**

For the greatest possible durability and consistency in a replication setup using InnoDB with transactions, you should use `innodb_flush_log_at_trx_commit=1` and `sync_binlog=1` in the master `my.cnf` file.

**Note**

Ensure that the `skip-networking` [440] option is not enabled on your replication master. If networking has been disabled, your slave will not be able to communicate with the master and replication will fail.

### 16.1.1.2. Setting the Replication Slave Configuration

On a replication slave, you must establish a unique server ID. If this has not already been done, this part of slave setup requires a server restart.

If the slave server ID is not already set, or the current value conflicts with the value that you have chosen for the master server, you should shut down your slave server and edit the configuration to specify a unique server ID. For example:

```
[mysqld]
server-id=2
```

After making the changes, restart the server.

If you are setting up multiple slaves, each one must have a unique `server-id` [1955] value that differs from that of the master and from each of the other slaves. Think of `server-id` [1955] values as something similar to IP addresses: These IDs uniquely identify each server instance in the community of replication partners.

**Note**

If you omit `server-id` [1955] (or set it explicitly to its default value of 0), a slave refuses to connect to a master.

You do not have to enable binary logging on the slave for replication to be enabled. However, if you enable binary logging on the slave, you can use the binary log for data backups and crash recovery on the slave, and also use the slave as part of a more complex replication topology (for example, where the slave acts as a master to other slaves).

### 16.1.1.3. Creating a User for Replication

Each slave must connect to the master using a MySQL user name and password, so there must be a user account on the master that the slave can use to connect. Any account can be used for this operation, providing it has been granted the `REPLICATION SLAVE` [755] privilege. You may wish to create a different account for each slave, or connect to the master using the same account for each slave.

You need not create an account specifically for replication. However, you should be aware that the user name and password will be stored in plain text within the `master.info` file (see [Section 16.2.2.2, “Slave Status Logs”](#)). Therefore, you may want to create a separate account that has privileges only for the replication process, to minimize the possibility of compromise to other accounts.

To create a new account, use `CREATE USER`. To grant this account the privileges required for replication, use the `GRANT` statement. If you create an account solely for the purposes of replication, that account needs only the `REPLICATION SLAVE` [755] privilege. For example, to set up a new user, `repl`, that can connect for replication from any host within the `mydomain.com` domain, issue these statements on the master:

```
mysql> CREATE USER 'repl'@'%mydomain.com' IDENTIFIED BY 'slavepass';
mysql> GRANT REPLICATION SLAVE ON *.* TO 'repl'@'%mydomain.com';
```

See [Section 13.7.1, “Account Management Statements”](#), for more information on statements for manipulation of user accounts.

#### 16.1.1.4. Obtaining the Replication Master Binary Log Coordinates

To configure replication on the slave you must determine the master's current coordinates within its binary log. You will need this information so that when the slave starts the replication process, it is able to start processing events from the binary log at the correct point.

If you have existing data on your master that you want to synchronize on your slaves before starting the replication process, you must stop processing statements on the master, and then obtain its current binary log coordinates and dump its data, before permitting the master to continue executing statements. If you do not stop the execution of statements, the data dump and the master status information that you use will not match and you will end up with inconsistent or corrupted databases on the slaves.

To obtain the master binary log coordinates, follow these steps:

1. Start a session on the master by connecting to it with the command-line client, and flush all tables and block write statements by executing the `FLUSH TABLES WITH READ LOCK` statement:

```
mysql> FLUSH TABLES WITH READ LOCK;
```

For `InnoDB` tables, note that `FLUSH TABLES WITH READ LOCK` also blocks `COMMIT` operations.



#### Warning

Leave the client from which you issued the `FLUSH TABLES` statement running so that the read lock remains in effect. If you exit the client, the lock is released.

2. In a different session on the master, use the `SHOW MASTER STATUS` statement to determine the current binary log file name and position:

```
mysql > SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| mysql-bin.000003 | 73      | test        | manual,mysql     |
+-----+-----+-----+-----+
```

The `File` column shows the name of the log file and `Position` shows the position within the file. In this example, the binary log file is `mysql-bin.000003` and the position is 73. Record these values. You need them later when you are setting up the slave. They represent the replication coordinates at which the slave should begin processing new updates from the master.

If the master has been running previously without binary logging enabled, the log file name and position values displayed by `SHOW MASTER STATUS` or `mysqldump --master-data` will be empty. In that case, the values that you need to use later when specifying the slave's log file and position are the empty string ( `' '` ) and `4`.

You now have the information you need to enable the slave to start reading from the binary log in the correct place to start replication.

If you have existing data that needs be to synchronized with the slave before you start replication, leave the client running so that the lock remains in place and then proceed to [Section 16.1.1.5, “Creating a Data Snapshot Using `mysqldump`”](#), or [Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”](#). The idea here is to prevent any further changes so that the data copied to the slaves is in synchrony with the master.

If you are setting up a brand new master and slave replication group, you can exit the first session to release the read lock.



### 16.1.1.5. Creating a Data Snapshot Using `mysqldump`

One way to create a snapshot of the data in an existing master database is to use the `mysqldump` tool. Once the data dump has been completed, you then import this data into the slave before starting the replication process.

To obtain a snapshot of the data using `mysqldump`:

1. If you have not already locked the tables on the server to prevent statements that update data from executing:

Start a session on the server by connecting to it with the command-line client, and flush all tables and block write statements by executing the `FLUSH TABLES WITH READ LOCK` statement:

```
mysql> FLUSH TABLES WITH READ LOCK;
```

Remember to use `SHOW MASTER STATUS` and record the binary log details for use when starting up the slave. The point in time of your snapshot and the binary log position must match. See [Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”](#).

2. In another session, use `mysqldump` to create a dump either of all the databases you want to replicate, or of selected individual databases. For example:

```
shell> mysqldump --all-databases --lock-all-tables >dbdump.db
```

An alternative to using a bare dump, is to use the `--master-data` option, which automatically appends the `CHANGE MASTER TO` statement required on the slave to start the replication process.

```
shell> mysqldump --all-databases --master-data >dbdump.db
```

3. In the client where you acquired the read lock, release the lock:

```
mysql> UNLOCK TABLES;
```

Alternatively, exit the first session to release the read lock.

When choosing databases to include in the dump, remember to filter out databases on each slave that you do not want to include in the replication process.

Either to copy the dump file to the slave, or use the file from the master when connecting remotely to the slave to import the data.

### 16.1.1.6. Creating a Data Snapshot Using Raw Data Files

If your database is large, copying the raw data files can be more efficient than using `mysqldump` and importing the file on each slave. This technique skips the overhead of updating indexes as the `INSERT` statements are replayed.

Using this method with tables in storage engines with complex caching or logging algorithms requires extra steps to produce a perfect “point in time” snapshot: the initial copy command might leave out cache information and logging updates, even if you have acquired a global read lock. How the storage engine responds to this depends on its crash recovery abilities.

This method also does not work reliably if the master and slave have different values for `ft_stopword_file` [485], `ft_min_word_len` [484], or `ft_max_word_len` [484] and you are copying tables having full-text indexes.

If you use `InnoDB` tables, you can use the `mysqlbackup` command from the MySQL Enterprise Backup component to produce a consistent snapshot. This command records the log name and offset corresponding to the snapshot to be later used on the slave. MySQL Enterprise Backup is a commercial product that is included as part of a MySQL Enterprise subscription. See [Section 23.2, “MySQL Enterprise Backup”](#) for detailed information.



Otherwise, use the [cold backup](#) technique to obtain a reliable binary snapshot of [InnoDB](#) tables: copy all data files after doing a [slow shutdown](#) of the MySQL Server.

To create a raw data snapshot of [MyISAM](#) tables, you can use standard copy tools such as [cp](#) or [copy](#), a remote copy tool such as [scp](#) or [rsync](#), an archiving tool such as [zip](#) or [tar](#), or a file system snapshot tool such as [dump](#), providing that your MySQL data files exist on a single file system. If you are replicating only certain databases, copy only those files that relate to those tables. (For [InnoDB](#), all tables in all databases are stored in the [system tablespace](#) files, unless you have the [innodb\\_file\\_per\\_table \[1718\]](#) option enabled.)

You might want to specifically exclude the following files from your archive:

- Files relating to the [mysql](#) database.
- The [master.info](#) file.
- The master's binary log files.
- Any relay log files.

To get the most consistent results with a raw data snapshot, shut down the master server during the process, as follows:

1. Acquire a read lock and get the master's status. See [Section 16.1.1.4, "Obtaining the Replication Master Binary Log Coordinates"](#).
2. In a separate session, shut down the master server:

```
shell> mysqladmin shutdown
```

3. Make a copy of the MySQL data files. The following examples show common ways to do this. You need to choose only one of them:

```
shell> tar cf /tmp/db.tar ./data
shell> zip -r /tmp/db.zip ./data
shell> rsync --recursive ./data /tmp/dbdata
```

4. Restart the master server.

If you are not using [InnoDB](#) tables, you can get a snapshot of the system from a master without shutting down the server as described in the following steps:

1. Acquire a read lock and get the master's status. See [Section 16.1.1.4, "Obtaining the Replication Master Binary Log Coordinates"](#).
2. Make a copy of the MySQL data files. The following examples show common ways to do this. You need to choose only one of them:

```
shell> tar cf /tmp/db.tar ./data
shell> zip -r /tmp/db.zip ./data
shell> rsync --recursive ./data /tmp/dbdata
```

3. In the client where you acquired the read lock, release the lock:

```
mysql> UNLOCK TABLES;
```

Once you have created the archive or copy of the database, copy the files to each slave before starting the slave replication process.

### 16.1.1.7. Setting Up Replication with New Master and Slaves

The easiest and most straightforward method for setting up replication is to use new master and slave servers.

You can also use this method if you are setting up new servers but have an existing dump of the databases from a different server that you want to load into your replication configuration. By loading the data into a new master, the data will be automatically replicated to the slaves.

To set up replication between a new master and slave:

1. Configure the MySQL master with the necessary configuration properties. See [Section 16.1.1.1, “Setting the Replication Master Configuration”](#).
2. Start up the MySQL master.
3. Set up a user. See [Section 16.1.1.3, “Creating a User for Replication”](#).
4. Obtain the master status information. See [Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”](#).
5. On the master, release the read lock:

```
mysql> UNLOCK TABLES;
```

6. On the slave, edit the MySQL configuration. See [Section 16.1.1.2, “Setting the Replication Slave Configuration”](#).
7. Start up the MySQL slave.
8. Execute a `CHANGE MASTER TO` statement to set the master replication server configuration. See [Section 16.1.1.10, “Setting the Master Configuration on the Slave”](#).

Perform the slave setup steps on each slave.

Because there is no data to load or exchange on a new server configuration you do not need to copy or import any information.

If you are setting up a new replication environment using the data from a different existing database server, you will now need to run the dump file generated from that server on the new master. The database updates will automatically be propagated to the slaves:

```
shell> mysql -h master < fulldb.dump
```

### 16.1.1.8. Setting Up Replication with Existing Data

When setting up replication with existing data, you will need to decide how best to get the data from the master to the slave before starting the replication service.

The basic process for setting up replication with existing data is as follows:

1. With the MySQL master running, create a user to be used by the slave when connecting to the master during replication. See [Section 16.1.1.3, “Creating a User for Replication”](#).
2. If you have not already configured the `server-id` [1955] and enabled binary logging on the master server, you will need to shut it down to configure these options. See [Section 16.1.1.1, “Setting the Replication Master Configuration”](#).

If you have to shut down your master server, this is a good opportunity to take a snapshot of its databases. You should obtain the master status (see [Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”](#)) before taking down the master, updating the configuration and taking a snapshot. For information on how to create a snapshot using raw data files, see [Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”](#).

3. If your master server is already correctly configured, obtain its status (see [Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”](#)) and then use `mysqldump` to take a snapshot (see [Section 16.1.1.5, “Creating a Data Snapshot Using mysqldump”](#)) or take a raw

snapshot of the live server using the guide in [Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”](#).

4. Update the configuration of the slave. See [Section 16.1.1.2, “Setting the Replication Slave Configuration”](#).
5. The next step depends on how you created the snapshot of data on the master.

If you used `mysqldump`:

- a. Start the slave, using the `--skip-slave-start [1984]` option so that replication does not start.
- b. Import the dump file:

```
shell> mysql < fulldb.dump
```

If you created a snapshot using the raw data files:

- a. Extract the data files into your slave data directory. For example:

```
shell> tar xvf dbdump.tar
```

You may need to set permissions and ownership on the files so that the slave server can access and modify them.

- b. Start the slave, using the `--skip-slave-start [1984]` option so that replication does not start.
6. Configure the slave with the replication coordinates from the master. This tells the slave the binary log file and position within the file where replication needs to start. Also, configure the slave with the login credentials and host name of the master. For more information on the `CHANGE MASTER TO` statement required, see [Section 16.1.1.10, “Setting the Master Configuration on the Slave”](#).
7. Start the slave threads:

```
mysql> START SLAVE;
```

After you have performed this procedure, the slave should connect to the master and catch up on any updates that have occurred since the snapshot was taken.

If you have forgotten to set the `server-id [1955]` option for the master, slaves cannot connect to it.

If you have forgotten to set the `server-id [1955]` option for the slave, you get the following error in the slave's error log:

```
Warning: You should set server-id to a non-0 value if master_host is set; we will force server id to 2, but this MySQL server will not act as a slave.
```

You also find error messages in the slave's error log if it is not able to replicate for any other reason.

Once a slave is replicating, you can find in its data directory one file named `master.info` and another named `relay-log.info`. The slave uses these two files to keep track of how much of the master's binary log it has processed. Do *not* remove or edit these files unless you know exactly what you are doing and fully understand the implications. Even in that case, it is preferred that you use the `CHANGE MASTER TO` statement to change replication parameters. The slave will use the values specified in the statement to update the status files automatically.



#### Note

The content of `master.info` overrides some of the server options specified on the command line or in `my.cnf`. See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#), for more details.

A single snapshot of the master suffices for multiple slaves. To set up additional slaves, use the same master snapshot and follow the slave portion of the procedure just described.

### 16.1.1.9. Introducing Additional Slaves to an Existing Replication Environment

To add another slave to an existing replication configuration, you can do so without stopping the master. Instead, set up the new slave by making a copy of an existing slave, except that you configure the new slave with a different `server-id` [1955] value.

To duplicate an existing slave:

1. Shut down the existing slave:

```
shell> mysqladmin shutdown
```

2. Copy the data directory from the existing slave to the new slave. You can do this by creating an archive using `tar` or `WinZip`, or by performing a direct copy using a tool such as `cp` or `rsync`. Ensure that you also copy the log files and relay log files.

A common problem that is encountered when adding new replication slaves is that the new slave fails with a series of warning and error messages like these:

```
071118 16:44:10 [Warning] Neither --relay-log nor --relay-log-index were used; so
replication may break when this MySQL server acts as a slave and has his hostname
changed!! Please use '--relay-log=new_slave_hostname-relay-bin' to avoid this problem.
071118 16:44:10 [ERROR] Failed to open the relay log './old_slave_hostname-relay-bin.003525'
(relay_log_pos 22940879)
071118 16:44:10 [ERROR] Could not find target log during relay log initialization
071118 16:44:10 [ERROR] Failed to initialize the master info structure
```

This is due to the fact that, if the `--relay-log` [1974] option is not specified, the relay log files contain the host name as part of their file names. (This is also true of the relay log index file if the `--relay-log-index` [1974] option is not used. See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#), for more information about these options.)

To avoid this problem, use the same value for `--relay-log` [1974] on the new slave that was used on the existing slave. (If this option was not set explicitly on the existing slave, use `existing_slave_hostname-relay-bin`.) If this is not feasible, copy the existing slave's relay log index file to the new slave and set the `--relay-log-index` [1974] option on the new slave to match what was used on the existing slave. (If this option was not set explicitly on the existing slave, use `existing_slave_hostname-relay-bin.index`.) Alternatively—if you have already tried to start the new slave (after following the remaining steps in this section) and have encountered errors like those described previously—then perform the following steps:

- a. If you have not already done so, issue a `STOP SLAVE` on the new slave.
 

If you have already started the existing slave again, issue a `STOP SLAVE` on the existing slave as well.
  - b. Copy the contents of the existing slave's relay log index file into the new slave's relay log index file, making sure to overwrite any content already in the file.
  - c. Proceed with the remaining steps in this section.
3. Copy the `master.info` and `relay-log.info` files from the existing slave to the new slave if they were not located in the data directory. These files hold the current log coordinates for the master's binary log and the slave's relay log.
  4. Start the existing slave.
  5. On the new slave, edit the configuration and give the new slave a unique `server-id` [1955] not used by the master or any of the existing slaves.

6. Start the new slave. The slave will use the information in its `master.info` file to start the replication process.

### 16.1.1.10. Setting the Master Configuration on the Slave

To set up the slave to communicate with the master for replication, you must tell the slave the necessary connection information. To do this, execute the following statement on the slave, replacing the option values with the actual values relevant to your system:

```
mysql> CHANGE MASTER TO
->     MASTER_HOST='master_host_name',
->     MASTER_USER='replication_user_name',
->     MASTER_PASSWORD='replication_password',
->     MASTER_LOG_FILE='recorded_log_file_name',
->     MASTER_LOG_POS=recorded_log_position;
```



#### Note

Replication cannot use Unix socket files. You must be able to connect to the master MySQL server using TCP/IP.

The `CHANGE MASTER TO` statement has other options as well. For example, it is possible to set up secure replication using SSL. For a full list of options, and information about the maximum permissible length for the string-valued options, see [Section 13.4.2.1, “CHANGE MASTER TO Syntax”](#).

## 16.1.2. Replication Formats

Replication works because events written to the binary log are read from the master and then processed on the slave. The events are recorded within the binary log in different formats according to the type of event. The different replication formats used correspond to the binary logging format used when the events were recorded in the master's binary log. The correlation between binary logging formats and the terms used during replication are:

- Replication capabilities in MySQL originally were based on propagation of SQL statements from master to slave. This is called *statement-based replication* (often abbreviated as *SBR*), which corresponds to the standard statement-based binary logging format. In older versions of MySQL (5.1.4 and earlier), binary logging and replication used this format exclusively.
- Row-based binary logging logs changes in individual table rows. When used with MySQL replication, this is known as *row-based replication* (often abbreviated as *RBR*). In row-based replication, the master writes *events* to the binary log that indicate how individual table rows are changed.
- The server can change the binary logging format in real time according to the type of event using *mixed-format logging*.

When the mixed format is in effect, statement-based logging is used by default, but automatically switches to row-based logging in particular cases as described later. Replication using the mixed format is often referred to as *mixed-based replication* or *mixed-format replication*. For more information, see [Section 5.2.4.3, “Mixed Binary Logging Format”](#).

In MySQL 5.6, statement-based format is the default.

When using `MIXED` format, the binary logging format is determined in part by the storage engine being used and the statement being executed. For more information on mixed-format logging and the rules governing the support of different logging formats, see [Section 5.2.4.3, “Mixed Binary Logging Format”](#).

The logging format in a running MySQL server is controlled by setting the `binlog_format` [2010] server system variable. This variable can be set with session or global scope. The rules governing when and how the new setting takes effect are the same as for other MySQL server system variables—setting the variable for the current session lasts only until the end of that session, and the change is not visible to other sessions; setting the variable globally requires a restart of the server to take effect. For more information, see [Section 13.7.4, “SET Syntax”](#).

There are conditions under which you cannot change the binary logging format at runtime or doing so causes replication to fail. See [Section 5.2.4.2, “Setting The Binary Log Format”](#).

You must have the `SUPER` [756] privilege to set either the global or session `binlog_format` [2010] value.

The statement-based and row-based replication formats have different issues and limitations. For a comparison of their relative advantages and disadvantages, see [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

With statement-based replication, you may encounter issues with replicating stored routines or triggers. You can avoid these issues by using row-based replication instead. For more information, see [Section 18.7, “Binary Logging of Stored Programs”](#).

### 16.1.2.1. Advantages and Disadvantages of Statement-Based and Row-Based Replication

Each binary logging format has advantages and disadvantages. For most users, the mixed replication format should provide the best combination of data integrity and performance. If, however, you want to take advantage of the features specific to the statement-based or row-based replication format when performing certain tasks, you can use the information in this section, which provides a summary of their relative advantages and disadvantages, to determine which is best for your needs.

- [Advantages of statement-based replication](#)
- [Disadvantages of statement-based replication](#)
- [Advantages of row-based replication](#)
- [Disadvantages of row-based replication](#)

#### Advantages of statement-based replication

- Proven technology that has existed in MySQL since 3.23.
- Less data written to log files. When updates or deletes affect many rows, this results in *much* less storage space required for log files. This also means that taking and restoring from backups can be accomplished more quickly.
- Log files contain all statements that made any changes, so they can be used to audit the database.

#### Disadvantages of statement-based replication

- **Statements that are unsafe for SBR.**  
Not all statements which modify data (such as `INSERT DELETE`, `UPDATE`, and `REPLACE` statements) can be replicated using statement-based replication. Any nondeterministic behavior is difficult to replicate when using statement-based replication. Examples of such DML (Data Modification Language) statements include the following:
  - A statement that depends on a UDF or stored program that is nondeterministic, since the value returned by such a UDF or stored program or depends on factors other than the parameters supplied to it. (Row-based replication, however, simply replicates the value returned by the UDF or stored program, so its effect on table rows and data is the same on both the master and slave.) See [Section 16.4.1.10, “Replication of Invoked Features”](#), for more information.
  - `DELETE` and `UPDATE` statements that use a `LIMIT` clause without an `ORDER BY` are nondeterministic. See [Section 16.4.1.15, “Replication and LIMIT”](#).
  - Statements using any of the following functions cannot be replicated properly using statement-based replication:
    - `LOAD_FILE()` [1181]

- `UUID()` [1280], `UUID_SHORT()` [1281]
- `USER()` [1273]
- `FOUND_ROWS()` [1268]
- `SYSDATE()` [1223] (unless both the master and the slave are started with the `--sysdate-is-now` [444] option)
- `GET_LOCK()` [1275]
- `IS_FREE_LOCK()` [1278]
- `IS_USED_LOCK()` [1279]
- `MASTER_POS_WAIT()` [1279]
- `RAND()` [1206]
- `RELEASE_LOCK()` [1280]
- `SLEEP()` [1280]
- `VERSION()` [1273]

However, all other functions are replicated correctly using statement-based replication, including `NOW()` [1220] and so forth.

For more information, see [Section 16.4.1.14, “Replication and System Functions”](#).

Statements that cannot be replicated correctly using statement-based replication are logged with a warning like the one shown here:

```
090213 16:58:54 [Warning] Statement is not safe to log in statement format.
```

A similar warning is also issued to the client in such cases. The client can display it using `SHOW WARNINGS`.

- `INSERT ... SELECT` requires a greater number of row-level locks than with row-based replication.
- `UPDATE` statements that require a table scan (because no index is used in the `WHERE` clause) must lock a greater number of rows than with row-based replication.
- For `InnoDB`: An `INSERT` statement that uses `AUTO_INCREMENT` blocks other nonconflicting `INSERT` statements.
- For complex statements, the statement must be evaluated and executed on the slave before the rows are updated or inserted. With row-based replication, the slave only has to modify the affected rows, not execute the full statement.
- If there is an error in evaluation on the slave, particularly when executing complex statements, statement-based replication may slowly increase the margin of error across the affected rows over time. See [Section 16.4.1.25, “Slave Errors During Replication”](#).
- Stored functions execute with the same `NOW()` [1220] value as the calling statement. However, this is not true of stored procedures.
- Deterministic UDFs must be applied on the slaves.
- Table definitions must be (nearly) identical on master and slave. See [Section 16.4.1.8, “Replication with Differing Table Definitions on Master and Slave”](#), for more information.

## Advantages of row-based replication



- All changes can be replicated. This is the safest form of replication.

The `mysql` database is not replicated. The `mysql` database is instead seen as a node-specific database. Row-based replication is not supported on tables in this database. Instead, statements that would normally update this information—such as `GRANT`, `REVOKE` and the manipulation of triggers, stored routines (including stored procedures), and views—are all replicated to slaves using statement-based replication.

For statements such as `CREATE TABLE ... SELECT`, a `CREATE` statement is generated from the table definition and replicated using statement-based format, while the row insertions are replicated using row-based format.

- The technology is the same as in most other database management systems; knowledge about other systems transfers to MySQL.
- Fewer row locks are required on the master, which thus achieves higher concurrency, for the following types of statements:
  - `INSERT ... SELECT`
  - `INSERT` statements with `AUTO_INCREMENT`
  - `UPDATE` or `DELETE` statements with `WHERE` clauses that do not use keys or do not change most of the examined rows.
- Fewer row locks are required on the slave for any `INSERT`, `UPDATE`, or `DELETE` statement.

### Disadvantages of row-based replication

- RBR tends to generate more data that must be logged. To replicate a DML statement (such as an `UPDATE` or `DELETE` statement), statement-based replication writes only the statement to the binary log. By contrast, row-based replication writes each changed row to the binary log. If the statement changes many rows, row-based replication may write significantly more data to the binary log; this is true even for statements that are rolled back. This also means that taking and restoring from backup can require more time. In addition, the binary log is locked for a longer time to write the data, which may cause concurrency problems.
- Deterministic UDFs that generate large `BLOB` values take longer to replicate with row-based replication than with statement-based replication. This is because the `BLOB` column value is logged, rather than the statement generating the data.
- You cannot examine the logs to see what statements were executed, nor can you see on the slave what statements were received from the master and executed.

However, you can see what data was changed using `mysqlbinlog` with the options `--base64-output=DECODE-ROWS` [348] and `--verbose` [354].

- For tables using the `MyISAM` storage engine, a stronger lock is required on the slave for `INSERT` statements when applying them as row-based events to the binary log than when applying them as statements. This means that concurrent inserts on `MyISAM` tables are not supported when using row-based replication.

### 16.1.2.2. Usage of Row-Based Logging and Replication

Major changes in the replication environment and in the behavior of applications can result from using row-based logging (RBL) or row-based replication (RBR) rather than statement-based logging or replication. This section describes a number of issues known to exist when using row-based logging or replication, and discusses some best practices for taking advantage of row-based logging and replication.

For additional information, see [Section 16.1.2, “Replication Formats”](#), and [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

- **RBL, RBR, and temporary tables.** As noted in [Section 16.4.1.21, “Replication and Temporary Tables”](#), temporary tables are not replicated when using row-based format. When mixed format is in effect, “safe” statements involving temporary tables are logged using statement-based format. For more information, see [Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”](#).

Temporary tables are not replicated when using row-based format because there is no need. In addition, because temporary tables can be read only from the thread which created them, there is seldom if ever any benefit obtained from replicating them, even when using statement-based format.

In MySQL 5.6, you can switch from statement-based to row-based binary logging mode even when temporary tables have been created. However, while using the row-based format, the MySQL server cannot determine the logging mode that was in effect when a given temporary table was created. For this reason, the server in such cases logs a `DROP TEMPORARY TABLE IF EXISTS` statement for each temporary table that still exists for a given client session when that session ends. While this means that it is possible that an unnecessary `DROP TEMPORARY TABLE` statement might be logged in some cases, the statement is harmless, and does not cause an error even if the table does not exist, due to the presence of the `IF NOT EXISTS` option.

In MySQL 5.6.6 and earlier, the `--disable-gtid-unsafe-statements` [2018] option caused any nontransactional DML statement involving temporary tables to fail with an error when using row-based logging, in spite of the fact that they are not written to the binary log. In MySQL 5.6.7 and later, such statements are allowed when using `binlog_format=ROW` [2010], as long as any nontransactional tables affected by the statements are temporary tables (Bug #14272672).

- **RBL and synchronization of nontransactional tables.** When many rows are affected, the set of changes is split into several events; when the statement commits, all of these events are written to the binary log. When executing on the slave, a table lock is taken on all tables involved, and then the rows are applied in batch mode. (This may or may not be effective, depending on the engine used for the slave's copy of the table.)
- **Latency and binary log size.** Because RBL writes changes for each row to the binary log, its size can increase quite rapidly. In a replication environment, this can significantly increase the time required to make changes on the slave that match those on the master. You should be aware of the potential for this delay in your applications.
- **Reading the binary log.** `mysqlbinlog` displays row-based events in the binary log using the `BINLOG` statement (see [Section 13.7.6.1, “BINLOG Syntax”](#)). This statement displays an event in printable form, but as a base 64-encoded string the meaning of which is not evident. When invoked with the `--base64-output=DECODE-ROWS` [348] and `--verbose` [354] options, `mysqlbinlog` formats the contents of the binary log in a manner that is easily human readable. This is helpful when binary log events were written in row-based format if you want to read or recover from a replication or database failure using the contents of the binary log. For more information, see [Section 4.6.8.2, “mysqlbinlog Row Event Display”](#).
- **Binary log execution errors and `slave_exec_mode`.** If `slave_exec_mode` [1992] is `IDEMPOTENT`, a failure to apply changes from RBL because the original row cannot be found does not trigger an error or cause replication to fail. This means that it is possible that updates are not applied on the slave, so that the master and slave are no longer synchronized. Latency issues and use of nontransactional tables with RBR when `slave_exec_mode` [1992] is `IDEMPOTENT` can cause the master and slave to diverge even further. For more information about `slave_exec_mode` [1992], see [Section 5.1.4, “Server System Variables”](#).



#### Note

`slave_exec_mode=IDEMPOTENT` [1992] is generally useful only for circular replication or multi-master replication with MySQL Cluster, for which `IDEMPOTENT` is the default value.

For other scenarios, setting `slave_exec_mode` [1992] to `STRICT` is normally sufficient; this is the default value for storage engines other than `NDB`.

The `NDBCLUSTER` storage engine is currently not supported in MySQL 5.6. MySQL Cluster users wishing to upgrade from MySQL 5.0 should instead migrate to MySQL Cluster NDB 7.1 or later. For more information about MySQL Cluster NDB 7.1, see [MySQL Cluster NDB 6.1 - 7.1](#); for more information about MySQL Cluster NDB 7.2, see [MySQL Cluster NDB 7.2](#).

- **Lack of binary log checksums.** RBL uses no checksums. This means that network, disk, and other errors may not be identified when processing the binary log. To ensure that data is transmitted without network corruption, you may want to consider using SSL, which adds another layer of checksumming, for replication connections. The `CHANGE MASTER TO` statement has options to enable replication over SSL. See also [Section 13.4.2.1, “CHANGE MASTER TO Syntax”](#), for general information about setting up MySQL with SSL.
- **Filtering based on server ID not supported.** A common practice is to filter out changes on some slaves by using a `WHERE` clause that includes the relation `@@server_id <> id_value` clause with `UPDATE` and `DELETE` statements, a simple example of such a clause being `WHERE @@server_id <> 1`. However, this does not work correctly with row-based logging. If you must use the `server_id` [539] system variable for statement filtering, you must also use `--binlog_format=STATEMENT` [2010].

In MySQL 5.6, you can do filtering based on server ID by using the `IGNORE_SERVER_IDS` option for the `CHANGE MASTER TO` statement. This option works with the statement-based and row-based logging formats.

- **Database-level replication options.** The effects of the `--replicate-do-db` [1976], `--replicate-ignore-db` [1977], and `--replicate-rewrite-db` [1979] options differ considerably depending on whether row-based or statement-based logging is used. Because of this, it is recommended to avoid database-level options and instead use table-level options such as `--replicate-do-table` [1978] and `--replicate-ignore-table` [1978]. For more information about these options and the impact that your choice of replication format has on how they operate, see [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).
- **RBL, nontransactional tables, and stopped slaves.** When using row-based logging, if the slave server is stopped while a slave thread is updating a nontransactional table, the slave database may reach an inconsistent state. For this reason, it is recommended that you use a transactional storage engine such as `InnoDB` for all tables replicated using the row-based format.

Use of `STOP SLAVE` or `STOP SLAVE SQL_THREAD` prior to shutting down the slave MySQL server helps prevent such issues from occurring, and is always recommended regardless of the logging format or storage engines employed.

### 16.1.2.3. Determination of Safe and Unsafe Statements in Binary Logging

When speaking of the “safeness” of a statement in MySQL Replication, we are referring to whether a statement and its effects can be replicated correctly using statement-based format. If this is true of the statement, we refer to the statement as *safe*; otherwise, we refer to it as *unsafe*.

In general, a statement is safe if it is deterministic, and unsafe if it is not. However, certain nondeterministic functions are *not* considered unsafe (see [Nondeterministic functions not considered unsafe](#), later in this section). In addition, statements using results from floating-point math functions—which are hardware-dependent—are always considered safe (see [Section 16.4.1.11, “Replication and Floating-Point Values”](#)).

**Handling of safe and unsafe statements.** A statement is treated differently depending on whether the statement is considered safe, and with respect to the binary logging format (that is, the current value of `binlog_format` [2010]).

- No distinction is made in the treatment of safe and unsafe statements when the binary logging mode is `ROW`.
- If the binary logging format is `MIXED`, statements flagged as unsafe are logged using the row-based format; statements regarded as safe are logged using the statement-based format.
- If the binary logging format is `STATEMENT`, statements flagged as being unsafe generate a warning to this effect. (Safe statements are logged normally.)

Each statement flagged as unsafe generates a warning. Formerly, in cases where a great many such statements were executed on the master, this could lead to very large error log files, sometimes even filling up an entire disk unexpectedly. To guard against this, MySQL 5.6.7 introduced a warning suppression mechanism, which behaves as follows: Whenever the 50 most recent `ER_BINLOG_UNSAFE_STATEMENT` [3698] warnings have been generated more than 50 times in any 50-second period, warning suppression is enabled. When activated, this causes such warnings not to be written to the error log; instead, for each 50 warnings of this type, a note `The last warning was repeated N times in last S seconds` is written to the error log. This continues as long as the 50 most recent such warnings were issued in 50 seconds or less; once the rate has decreased below this threshold, the warnings are once again logged normally. Warning suppression has no effect on how the safety of statements for statement-based logging is determined, nor on how warnings are sent to the client (MySQL clients still receive one warning for each such statement).

For more information, see [Section 16.1.2, “Replication Formats”](#).

#### Statements considered unsafe.

Statements having the following characteristics are considered unsafe:

- **Statements containing system functions that may return a different value on slave.** These functions include `FOUND_ROWS()` [1268], `GET_LOCK()` [1275], `IS_FREE_LOCK()` [1278], `IS_USED_LOCK()` [1279], `LOAD_FILE()` [1181], `MASTER_POS_WAIT()` [1279], `PASSWORD()` [1263], `RAND()` [1206], `RELEASE_LOCK()` [1280], `ROW_COUNT()` [1272], `SESSION_USER()` [1273], `SLEEP()` [1280], `SYSDATE()` [1223], `SYSTEM_USER()` [1273], `USER()` [1273], `UUID()` [1280], and `UUID_SHORT()` [1281].

**Nondeterministic functions not considered unsafe.** Although these functions are not deterministic, they are treated as safe for purposes of logging and replication:

`CONNECTION_ID()` [1267], `CURDATE()` [1212], `CURRENT_DATE()` [1212], `CURRENT_TIME()` [1212], `CURRENT_TIMESTAMP()` [1213], `CURTIME()` [1213], `LAST_INSERT_ID()` [1269], `LOCALTIME()` [1219], `LOCALTIMESTAMP()` [1219], `NOW()` [1220], `UNIX_TIMESTAMP()` [1226], `UTC_DATE()` [1227], `UTC_TIME()` [1227], and `UTC_TIMESTAMP()` [1227].

For more information, see [Section 16.4.1.14, “Replication and System Functions”](#).

- **References to system variables.** Most system variables are not replicated correctly using the statement-based format. For exceptions, see [Section 5.2.4.3, “Mixed Binary Logging Format”](#).

See [Section 16.4.1.32, “Replication and Variables”](#).

- **UDFs.** Since we have no control over what a UDF does, we must assume that it is executing unsafe statements.
- **Updates a table having an `AUTO_INCREMENT` column.** This is unsafe because the order in which the rows are updated may differ on the master and the slave.

In addition, an `INSERT` into a table that has a composite primary key containing an `AUTO_INCREMENT` column that is not the first column of this composite key is unsafe.

For more information, see [Section 16.4.1.1, “Replication and `AUTO\_INCREMENT`”](#).

- **`INSERT DELAYED` statement.** This statement is considered unsafe because the insertion of the rows may interleave with concurrently executing statements.

- **INSERT ... ON DUPLICATE KEY UPDATE statements on tables with multiple primary or unique keys.** When executed against a table that contains more than one primary or unique key, this statement is considered unsafe, being sensitive to the order in which the storage engine checks the keys, which is not deterministic, and on which the choice of rows updated by the MySQL Server depends.

An `INSERT ... ON DUPLICATE KEY UPDATE` statement against a table having more than one unique or primary key is marked as unsafe for statement-based replication beginning with MySQL 5.6.6. (Bug #11765650, Bug #58637)

- **Updates using LIMIT.** The order in which rows are retrieved is not specified.

See [Section 16.4.1.15, “Replication and LIMIT”](#).

- **Accesses or references log tables.** The contents of the system log table may differ between master and slave.
- **Nontransactional operations after transactional operations.** Within a transaction, allowing any nontransactional reads or writes to execute after any transactional reads or writes is considered unsafe.

For more information, see [Section 16.4.1.29, “Replication and Transactions”](#).

- **Accesses or references self-logging tables.** All reads and writes to self-logging tables are considered unsafe. Within a transaction, any statement following a read or write to self-logging tables is also considered unsafe.
- **LOAD DATA INFILE statements.** Beginning with MySQL 5.6, `LOAD DATA INFILE` is considered unsafe, it causes a warning in statement-based mode, and a switch to row-based format when using mixed-format logging. See [Section 16.4.1.16, “Replication and LOAD DATA INFILE”](#).

For additional information, see [Section 16.4.1, “Replication Features and Issues”](#).

### 16.1.3. Replication with Global Transaction Identifiers

In this section, we discuss transaction-based replication using *global transaction identifiers* (GTIDs), introduced in MySQL 5.6.5. When using GTIDs, each transaction can be identified and tracked as it is committed on the originating server and applied by any slaves; this means that it is not necessary when using GTIDs to refer to log files or positions within those files when starting a new slave or failing over to a new master, which greatly simplifies these tasks. Because GTID-based replication is completely transaction-based, it is simple to determine whether masters and slaves are consistent; as long as all transactions committed on a master are also committed on a slave, consistency between the two is guaranteed. You can use either statement-based or row-based replication with GTIDs (see [Section 16.1.2, “Replication Formats”](#)); however, for best results, we recommend that you use the row-based format.

The next few sections discuss the following topics:

- How GTIDs are defined and created, and how they are represented in the MySQL Server (see [Section 16.1.3.1, “GTID Concepts”](#)).
- A general procedure for setting up and starting GTID-based replication (see [Section 16.1.3.2, “Setting Up Replication Using GTIDs”](#)).
- Suggested methods for provisioning new replication servers when using GTIDs (see [Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”](#)).
- Restrictions and limitations that you should be aware of when using GTID-based replication (see [Section 16.1.3.4, “Restrictions on Replication with GTIDs”](#)).

For information about MySQL Server options and variables relating to GTID-based replication, see [Section 16.1.4.5, “Global Transaction ID Options and Variables”](#). See also [Section 12.15, “GTID Functions”](#), which describes SQL functions supported by MySQL 5.6 for use with GTIDs.

### 16.1.3.1. GTID Concepts

A global transaction identifier (GTID) is a unique identifier created and associated with each transaction when it is committed on the server of origin (master). This identifier is unique not only to the server on which it originated, but is unique across all servers in a given replication setup. There is a 1-to-1 mapping between all transactions and all GTIDs.

A GTID is represented as a pair of coordinates, separated by a colon character (:), as shown here:

```
GTID = source_id:transaction_id
```

The *source\_id* identifies the originating server. Normally, the server's *server\_uuid* [\[1955\]](#) is used for this purpose. (It is theoretically possible for it to be determined in a different manner if the source of the transaction is not a MySQL Server instance, but this is currently not supported.) The *transaction\_id* is a sequence number determined by the order in which the transaction was committed on this server; for example, the first transaction to be committed has `1` as its *transaction\_id*, and the tenth transaction to be committed on the same originating server is assigned a *transaction\_id* of `10`. (It is not possible for a transaction to have `0` as a sequence number in a GTID.) Thus, the twenty-third transaction to be committed originally on the server having the UUID `3E11FA47-71CA-11E1-9E33-C80AA9429562` has this GTID:

```
3E11FA47-71CA-11E1-9E33-C80AA9429562:23
```

This format is used to represent GTIDs in the output of statements such as `SHOW SLAVE STATUS` as well as in the binary log. They can also be seen when viewing the log file with `mysqlbinlog --base64-output=DECODE-ROWS` [\[348\]](#) or in the output from `SHOW BINLOG EVENTS`.

As written in the output of statements such as `SHOW MASTER STATUS`, `SHOW SLAVE STATUS`, a sequence of GTIDs originating from the same server may be collapsed into a single expression, as shown here.

```
3E11FA47-71CA-11E1-9E33-C80AA9429562:1-5
```

The example just shown represents the first through fifth transactions originating on the MySQL Server whose *server\_uuid* [\[1955\]](#) is `3E11FA47-71CA-11E1-9E33-C80AA9429562`.

In MySQL 5.6.6 and later, this format is also used to supply the argument required by the `START SLAVE` options `SQL_BEFORE_GTIDS` and `SQL_AFTER_GTIDS`.

**GTID sets.** A GTID set is a set of global transaction identifiers which is represented as shown here:

```
gtid_set:
  uuid_set [, uuid_set] ...
  | ''

uuid_set:
  uuid:interval[:interval]...

uuid:
  hhhhhhhh-hhhh-hhhh-hhhh-hhhhhhhhhhhh

h:
  [0-9,A-F]

interval:
  n[-n]

(n >= 1)
```

GTID sets are used in the MySQL Server in several ways. For example, the values stored by the `gtid_executed` [\[2021\]](#) and `gtid_purged` [\[2023\]](#) system variables are represented as GTID sets.



In addition, the functions `GTID_SUBSET()` [1274] and `GTID_SUBTRACT()` [1274] require GTID sets as input.

GTIDs are always preserved between master and slave. This means that you can always determine the source for any transaction applied on any slave by examining its binary log. In addition, once a transaction with a given GTID is committed on a given server, any subsequent transaction having the same GTID is ignored by that server. Thus, a transaction committed on the master can be applied no more than once on the slave, which helps to guarantee consistency.

When GTIDs are in use, the slave has no need for any nonlocal data, such as the name of a file on the master and a position within that file. All necessary information for synchronizing with the master is obtained directly from the replication data stream. From the perspective of the database administrator or developer, GTIDs entirely take the place of the file-offset pairs previously required to determine points for starting, stopping, or resuming the flow of data between master and slave. This means that, when you are using GTIDs for replication, you do not need (or want) to include `MASTER_LOG_FILE` or `MASTER_LOG_POS` options in the `CHANGE MASTER TO` statement used to direct a slave to replicate from a given master; in place of these options, it necessary only to enable the `MASTER_AUTO_POSITION` option introduced in MySQL 5.6.5. For the exact steps needed to configure and start masters and slaves using GTID-based replication, see [Section 16.1.3.2, “Setting Up Replication Using GTIDs”](#).

The generation and lifecycle of a GTID consists of the following steps:

1. A transaction is executed and committed on the master.

This transaction is assigned a GTID using the master's UUID and the smallest nonzero transaction sequence number not yet used on this server; the GTID is written to the master's binary log (immediately preceding the transaction itself in the log).

2. After the binary log data is transmitted to the slave and stored in the slave's relay log (using established mechanisms for this process—see [Section 16.2, “Replication Implementation”](#), for details), the slave reads the GTID and sets the value of its `gtid_next` [2022] system variable as this GTID. This tells the slave that the next transaction must be logged using this GTID.

It is important to note that the slave sets `gtid_next` in a session context.

3. The slave checks to make sure that this GTID has not already been used to log a transaction in its own binary log. If and only if this GTID has not been used, the slave then writes the GTID and applies the transaction (and writes the transaction to its binary log). By reading and checking the transaction's GTID first, before processing the transaction itself, the slave guarantees not only that no previous transaction having this GTID has been applied on the slave, but also that no other session has already read this GTID but has not yet committed the associated transaction. In other words, multiple clients are not permitted to apply the same transaction concurrently.
4. Because `gtid_next` [2022] is not empty, the slave does not attempt to generate a GTID for this transaction but instead writes the GTID stored in this variable—that is, the GTID obtained from the master—immediately preceding the transaction in its binary log.

### 16.1.3.2. Setting Up Replication Using GTIDs

This section describes a process for configuring and starting GTID-based replication in MySQL 5.6. This is a “cold start” procedure that assumes either that you are starting the replication master for the first time, or that it is possible to stop it; for information about provisioning replication slaves using GTIDs from a running master, see [Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”](#).

The key steps in this startup process for the simplest possible GTID replication topology—consisting of one master and one slave—are as follows:

1. If replication is already running, synchronize both servers by making them read-only.
2. Stop both servers.



- Restart both servers with GTIDs, binary logging, and slave update logging enabled, and with statements that are unsafe for GTID-based replication disabled. In addition, the servers should be started in read-only mode, and the slave SQL and I/O threads should be prevented from starting on the slave.

The `mysqld` options necessary to start the servers as described are discussed in the example that follows later in this section.

- Instruct the slave to use the master as the replication data source and to use auto-positioning, and then start the slave.

The SQL statements needed to accomplish this step are described in the example that follows later in this section.

- Disable read-only mode on both servers, so that they can once again accept updates.

We now present a more detailed example. We assume two servers already running as master and slave, using MySQL's "classic" file-based replication protocol.

Most of the steps that follow require the use of the MySQL `root` account or another MySQL user account that has the `SUPER` [756] privilege. `mysqladmin shutdown` requires either the `SUPER` privilege or the `SHUTDOWN` [756] privilege.

**Step 1: Synchronize the servers.** Make the servers read-only. To do this, enable the `read_only` [532] system variable by executing the following statement on both servers:

```
mysql> SET @@global.read_only = ON;
```

Then, allow the slave to catch up with the master. *It is extremely important that you make sure the slave has processed all updates before continuing.*

**Step 2: Stop both servers.** Stop each server using `mysqladmin` as shown here, where `username` and `password` are the user name and password for a MySQL user having sufficient privileges to shut down the server:

```
shell> mysqladmin -uusername -ppassword shutdown
```

**Step 3: Restart both servers with GTIDs enabled.** To enable binary logging with global transaction identifiers, each server must be started with GTID mode, binary logging, slave update logging enabled, and with statements that are unsafe for GTID-based replication disabled. In addition, you should prevent unwanted or accidental updates from being performed on either server by starting both in read-only mode as well. This means that both servers must be started with (at least) the options shown in the following invocation of `mysqld_safe`:

```
shell> mysqld_safe --gtid_mode=ON --log-bin --log-slave-updates --enforce-gtid-consistency &
```



#### Note

Prior to MySQL 5.6.9, `--enforce-gtid-consistency` [2018] was named `--disable-gtid-unsafe-statements` [2018].

In addition, you should start the slave with the `--skip-slave-start` [1984] option along with the other server options specified in the example just shown.

Although it may appear that `--gtid-mode` [2019] is a boolean, it is not (in fact, its values are enumerated). Use one of the values `ON` or `OFF` only, when setting this option. Using a numeric value such as 0 or 1 can lead to unexpected results.

For more information about the `--gtid-mode` [2019] and `--enforce-gtid-consistency` [2018] server options, see [Section 16.1.4.5, "Global Transaction ID Options and Variables"](#).

Depending on your circumstances, you may want or need to supply additional options to `mysqld_safe` (or other `mysqld` startup script).

**Step 4: Direct the slave to use the master.** Instruct the slave to use the master as the replication data source, and to use GTID-based auto-positioning rather than file-based positioning. You can do this by executing a `CHANGE MASTER TO` statement on the slave, using the `MASTER_AUTO_POSITION` option to tell the slave that transactions will be identified by GTIDs.

You may also need to supply appropriate values for the master's host name and port number as well as the user name and password for a replication user account which can be used by the slave to connect to the master; if these have already been set prior to Step 1 and no further changes need to be made, the corresponding options can safely be omitted from the statement shown here.

```
mysql> CHANGE MASTER TO
>     MASTER_HOST = host,
>     MASTER_PORT = port,
>     MASTER_USER = user,
>     MASTER_PASSWORD = password,
>     MASTER_AUTO_POSITION = 1;
```

Neither the `MASTER_LOG_FILE` option nor the `MASTER_LOG_POS` option may be used with `MASTER_AUTO_POSITION` set equal to 1. Attempting to do so causes the `CHANGE MASTER TO` statement to fail with an error. (If you need to revert from GTID-based replication to replication based on files and positions, you must use one or both of these options together with `MASTER_AUTO_POSITION = 0` in the `CHANGE MASTER TO` statement.)

Assuming that the `CHANGE MASTER TO` statement has succeeded, you can then start the slave, like this:

```
mysql> START SLAVE;
```

**Step 5: Disable read-only mode.** Allow both servers to begin accepting updates once again by running the following statement on each of them:

```
mysql> SET @@global.read_only = OFF;
```

GTID-based replication should now be running, and you can begin (or resume) activity on the master as before. [Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”](#), discusses creation of new slaves when using GTIDs.

### 16.1.3.3. Using GTIDs for Failover and Scaleout

There are a number of techniques when using MySQL Replication with Global Transaction Identifiers (GTIDs) in MySQL 5.6.9 and later for provisioning a new slave which can then be used for scaleout, being promoted to master as necessary for failover. In this section, we discuss the four techniques listed here:

- [Simple replication](#)
- [Copying data and transactions to the slave](#)
- [Injecting empty transactions](#)
- [Excluding transactions with `gtid\_purged`](#)

Global transaction identifiers were added to MySQL Replication for the purpose of simplifying in general management of the replication data flow and of failover activities in particular. Each identifier uniquely identifies a set of binary log events that together make up a transaction. GTIDs play a key role in applying changes to the database: the server automatically skips any transaction having an identifier which the server recognizes as one that it has processed before. This behavior is critical for automatic replication positioning and correct failover.

The mapping between identifiers and sets of events comprising a given transaction is captured in the binary log. This poses some challenges when a user wants to provision a new server with data from another existing server. To reproduce the identifier set on the new server, it is necessary, not only to copy the identifiers from the old server to the new one, but to preserve the relationship between the identifiers and the actual events as well, which is what is needed for restoring a slave that is immediately available as a candidate to become a new master on failover or switchover.

**Simple replication.** This is the easiest way to reproduce all identifiers and transactions on a new server; you simply make the new server into the slave of a master that has the entire execution history, and enable global transaction identifiers on both servers. (This requires that both master and slave are running with the options `--gtid-mode=ON` [2019] `--log-bin` [2002] `--log-slave-updates` [1971] `--enforce-gtid-consistency` [2018]; see Section 16.1.3.2, “Setting Up Replication Using GTIDs”, for more information.)

Once replication is started, the new server copies the entire binary log from the master and thus obtains all information about all GTIDs.

This method is simple and effective, but requires the slave to read the binary log from the master; it can sometimes take a comparatively long time for the new slave to catch up with the master, so this method is not suitable for fast failover or restoring from backup. We can obviate the need to fetch all of the execution history from the master by copying binary log files to the new server, as discussed in the next few paragraphs.

**Copying data and transactions to the slave.** Playing back the entire transaction history can be time-consuming, and represents a major bottleneck when setting up a new replication slave. To eliminate this requirement, we can take from the master a backup that includes, in addition to a dump containing a snapshot of the data set, the binary logs and the global transaction information they contain. Setting up the slave then consists of importing the snapshot, then playing back the binary log, after which replication can be started, allowing the slave to become current with any remaining transactions.

There are several variants of this method; these can be distinguished by the manner in which data (dumps) and transactions (binary logs) are shipped to the new slave, as outlined here:

Data Set	Transaction History
<ul style="list-style-type: none"> <li>• Use the <code>mysql</code> client to import a dump file created with <code>mysqldump</code>. Use the <code>--master-data</code> [288] option to include binary logging information and <code>--set-gtid-purged</code> [289] (available in MySQL 5.6.9 and later) to <code>AUTO</code> (the default) or <code>ON</code>, to include information about executed transactions. You should have <code>--gtid-mode=ON</code> [2019] while importing the dump on the slave. (Bug #14832472)</li> <li>• Stop the slave, copy the contents of the master's data directory to the slave's data directory, then restart the slave.</li> </ul>	<p>If <code>gtid_mode</code> [2022] is not <code>ON</code>, restart the server with GTID mode enabled.</p> <ul style="list-style-type: none"> <li>• Import the binary log using <code>mysqlbinlog</code>, with the <code>--read-from-remote-server</code> [352] and <code>--read-from-remote-master</code> [352] options.</li> <li>• Copy the master's binary log files to the slave. You can make copies from the slave using <code>mysqlbinlog --read-from-remote-server</code> [352] <code>--raw</code> [352]. These can be read in to the slave in either of the following ways: <ul style="list-style-type: none"> <li>• Update the slave's <code>binlog.index</code> file to point to the copied log files. Then execute a <code>CHANGE MASTER TO</code> statement in the <code>mysql</code> client to point to the first log file, and <code>START SLAVE</code> to read them.</li> <li>• Use <code>mysqlbinlog &gt; file</code> (without the <code>--raw</code> [352] option) to export the binary log files to SQL files that can be processed by the <code>mysql</code> client.</li> </ul> </li> </ul>

See also [Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”](#).

This method has the advantage that a new server is available almost immediately; only those transactions that were committed while the snapshot or dump file was being replayed still need to be obtained from the existing master. This means that the slave's availability is not instantaneous—but only a relatively short amount of time should be required for the slave to catch up with these few remaining transactions.

Copying over binary logs to the target server in advance is usually faster than reading the entire transaction execution history from the master in real time. However, due to size or other considerations, it may not always be feasible to move these files to the target when required. The two remaining methods for provisioning a new slave discussed in this section use other means to convey information about transactions to the new slave.

**Injecting empty transactions.** The master's global `gtid_executed` [2021] variable contains the set of all transactions executed on the master. Rather than copy the binary logs when taking a snapshot to provision a new server, you can instead note the content of `gtid_executed` on the server from which the snapshot was taken. Before adding the new server to the replication chain, simply commit an empty transaction on the new server for each transaction identifier contained in the master's `gtid_executed`, like this:

```
SET GTID_NEXT='aaa-bbb-ccc-ddd:N' ;

BEGIN ;
COMMIT ;

SET GTID_NEXT='AUTOMATIC' ;
```

Once all transaction identifiers have been reinstated in this way using empty transactions, you must flush and purge the slave's binary logs, as shown here, where `N` is the nonzero suffix of the current binary log file name:

```
FLUSH LOGS ;
PURGE BINARY LOGS TO 'master-bin.00000N' ;
```

You should do this to prevent this server from flooding the replication stream with false transactions in the event that it is later promoted to master. (The `FLUSH LOGS` statement forces the creation of a new binary log file; `PURGE BINARY LOGS` purges the empty transactions, but retains their identifiers.)

This method creates a server that is essentially a snapshot, but in time is able to become a master as its binary log history converges with that of the replication stream (that is, as it catches up with the master or masters). This outcome is similar in effect to that obtained using the remaining provisioning method, which we discuss in the next few paragraphs.

**Excluding transactions with `gtid_purged`.** The master's global `gtid_purged` [2023] variable contains the set of all transactions that have been purged from the master's binary log. As with the method discussed previously (see [Injecting empty transactions](#)), you can record the value of `gtid_executed` [2021] on the server from which the snapshot was taken (in place of copying the binary logs to the new server). Unlike the previous method, there is no need to commit empty transactions (or to issue `PURGE BINARY LOGS`); instead, you can set `gtid_purged` [2023] on the slave directly, based on the value of `gtid_executed` [2021] on the server from which the backup or snapshot was taken.



#### Note

Prior to MySQL 5.6.9, `gtid_purged` [2023] was not settable. (Bug #14797808)

As with the method using empty transactions, this method creates a server that is functionally a snapshot, but in time is able to become a master as its binary log history converges with that of the replication master or group.

### 16.1.3.4. Restrictions on Replication with GTIDs

Because GTID-based replication is dependent on transactions, some features otherwise available in MySQL are not supported when using it. This section provides information about restrictions on and limitations of replication with GTIDs.

**Updates involving nontransactional storage engines.** When using GTIDs, updates to tables using nontransactional storage engines such as [MyISAM](#) cannot be made in the same statement or transaction as updates to tables using transactional storage engines such as [InnoDB](#).

This restriction is due to the fact that updates to using a nontransactional storage engine mixed with updates to tables that use a transactional storage engine within the same transaction can result in multiple GTIDs being assigned to the same transaction. This problem can also occur in at least two other cases, listed here:

- When the master and the slave use different storage engines for their respective versions of the same table, where one storage engine is transactional and the other is not.
- When both the master and the slave use a nontransactional engine, but use different binary logging formats (for example, when the master has `binlog_format=ROW` [2010] and the slave has `binlog_format=STATEMENT`).

In any of the cases just mentioned, the one-to-one correspondence between transactions and GTIDs is broken, with the result that GTID-based replication cannot function correctly.

**CREATE TABLE ... SELECT statements.** `CREATE TABLE ... SELECT` is not safe for statement-based replication. When using row-based replication, this statement is actually logged as two separate events—one for the creation of the table, and another for the insertion of rows from the source table into the new table just created. When this statement is executed within a transaction, it is possible in some cases for these two events to receive the same transaction identifier, which means that the transaction containing the inserts is skipped by the slave. Therefore, `CREATE TABLE ... SELECT` is not supported when using GTID-based replication.

**Temporary tables.** `CREATE TEMPORARY TABLE` and `DROP TEMPORARY TABLE` statements are not supported inside transactions when using GTIDs (that is, when the server was started with the `--enforce-gtid-consistency` [2018] option). It is possible to use these statements with GTIDs enabled, but only outside of any transaction, and only with `autocommit=1` [464].

**Preventing execution of unsupported statements.** In order to prevent execution of statements that would cause GTID-based replication to fail, all servers must be started with the `--enforce-gtid-consistency` [2018] option when enabling GTIDs. This causes statements of any of the types discussed previously in this section to fail with an error.

For information about other required startup options when enabling GTIDs, see [Section 16.1.3.2, “Setting Up Replication Using GTIDs”](#).

`sql_slave_skip_counter` [1996] is not supported when using GTIDs. If you need to skip transactions, use the value of the master's `gtid_executed` [2021] variable instead; see [Injecting empty transactions](#), for more information.

**GTID mode and `mysqldump`.** In MySQL 5.6.9 and later, it is possible to import a dump made using `mysqldump` into a MySQL Server running with GTID mode enabled, provided that there are no GTIDs in the target server's binary log.

Prior to MySQL 5.6.9, `mysqldump` did not record global transaction IDs, and it was necessary to use the binary log and `mysqlbinlog` to restore GTIDs. (Bug #14797808, Bug #14832472)

**GTID mode and `mysql_upgrade`.** Prior to MySQL 5.6.7, `mysql_upgrade` could not connect to a MySQL Server that was running with `--gtid-mode=ON` [2019] unless `mysql_upgrade` was run with `--write-binlog=OFF` [244]. (Otherwise, `mysqld` had to be restarted with `--gtid-mode=OFF` before running `mysql_upgrade`, then restarted with `--gtid-mode=ON` afterwards.) This is not an issue in MySQL 5.6.7 and later, where `mysql_upgrade` runs with `--write-binlog=OFF` by

default. (Bug #13833710) However, it is not recommended to do so, since `mysql_upgrade` can make changes to system tables that use the `MyISAM` storage engine, which is nontransactional.

### 16.1.4. Replication and Binary Logging Options and Variables

The next few sections contain information about `mysqld` options and server variables that are used in replication and for controlling the binary log. Options and variables for use on replication masters and replication slaves are covered separately, as are options and variables relating to binary logging. A set of quick-reference tables providing basic information about these options and variables is also included (in the next section following this one).

Of particular importance is the `--server-id [1955]` option.

<b>Command-Line Format</b>	<code>--server-id=#</code>	
<b>Option-File Format</b>	<code>server-id</code>	
<b>Option Sets Variable</b>	Yes, <code>server_id</code> [539]	
<b>Variable Name</b>	<code>server_id</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 4294967295</code>

This option is common to both master and slave replication servers, and is used in replication to enable master and slave servers to identify themselves uniquely. For additional information, see [Section 16.1.4.2, “Replication Master Options and Variables”](#), and [Section 16.1.4.3, “Replication Slave Options and Variables”](#).

On the master and each slave, you *must* use the `--server-id [1955]` option to establish a unique replication ID in the range from 1 to  $2^{32} - 1$ . “Unique”, means that each ID must be different from every other ID in use by any other replication master or slave. Example: `server-id=3`.

If you omit `--server-id [1955]`, the default ID is 0, in which case a master refuses connections from all slaves, and a slave refuses to connect to a master. For more information, see [Section 16.1.1.2, “Setting the Replication Slave Configuration”](#).

`server_uuid [1955]`

Beginning with MySQL 5.6, the server generates a true UUID in addition to the `--server-id [1955]` supplied by the user. This is available as the global, read-only variable `server_uuid [1955]`.

<b>Version Introduced</b>	5.6.0	
<b>Variable Name</b>	<code>server_uuid</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

When starting, the MySQL server automatically obtains a UUID as follows:

1. Attempt to read and use the UUID written in the file `data_dir/auto.cnf` (where `data_dir` is the server's data directory); exit on success.
2. Otherwise, generate a new UUID and save it to this file, creating the file if necessary.



The `auto.cnf` file has a format similar to that used for `my.cnf` or `my.ini` files. In MySQL 5.6, `auto.cnf` has only a single `[auto]` section containing a single `server_uuid` [1955] setting and value; the file's contents appear similar to what is shown here:

```
[auto]
server_uuid=8a94f357-aab4-11df-86ab-c80aa9429562
```



### Important

The `auto.cnf` file is automatically generated; you should not attempt to write or modify this file.

Also beginning with MySQL 5.6, when using MySQL replication, masters and slaves know one another's UUIDs. The value of a slave's UUID can be seen in the output of `SHOW SLAVE HOSTS`. Once `START SLAVE` has been executed (but not before), the value of the master's UUID is available on the slave in the output of `SHOW SLAVE STATUS`.



### Note

Issuing a `STOP SLAVE` or `RESET SLAVE` statement does *not* reset the master's UUID as used on the slave.

In MySQL 5.6.5 and later, a server's `server_uuid` is also used in GTIDs for transactions originating on that server. For more information, see [Section 16.1.3, "Replication with Global Transaction Identifiers"](#).

When starting, the slave I/O thread generates an error and aborts if its master's UUID is equal to its own unless the `--replicate-same-server-id` [1979] option has been set. In addition, the slave I/O thread generates a warning if either of the following is true:

- No master having the expected `server_uuid` [1955] exists.
- The master's `server_uuid` [1955] has changed, although no `CHANGE MASTER TO` statement has ever been executed.



### Note

The addition of the `server_uuid` [1955] system variable in MySQL 5.6 does not change the requirement for setting a unique `--server-id` [1955] for each MySQL server as part of preparing and running MySQL replication, as described earlier in this section.

#### 16.1.4.1. Replication and Binary Logging Option and Variable Reference

The following tables list basic information about the MySQL command-line options and system variables applicable to replication and the binary log.

**Table 16.1. Replication Options and Variables: MySQL 5.6**

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
<code>Com_change_master</code> [586]	N	N	Both
	N	Y	N
<code>Com_show_master_status</code> [586]	N	N	Both
	N	Y	N
<code>Com_show_new_master</code> [586]	N	N	Both
	N	Y	N
<code>Com_show_slave_hosts</code> [586]	N	N	Both



Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	N	Y	N
Com_show_slave_status [586]	N	N	Both
	N	Y	N
Com_slave_start [586]	N	N	Both
	N	Y	N
Com_slave_stop [586]	N	N	Both
	N	Y	N
Rpl_semi_sync_master_clients [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_net_avg_wait_time [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_net_wait_time [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_net_waits [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_no_times [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_no_tx [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_status [596]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_timefunc_failures [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_tx_avg_wait_time [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_tx_wait_time [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_tx_waits [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_wait_pos_backtraverse [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_wait_sessions [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_master_yes_tx [597]	N	N	Global
	N	Y	N
Rpl_semi_sync_slave_status [597]	N	N	Global
	N	Y	N
slave_exec_mode [1992]	Y	Y	Global

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	Y	N	Y
Slave_open_temp_tables [598]	N	N	Global
	N	Y	N
Slave_retried_transactions [598]	N	N	Global
	N	Y	N
Slave_running [598]	N	N	Global
	N	Y	N
abort-slave-event-count [1971]	Y	N	Global
	Y	N	N
disable-gtid-unsafe-statements [2018]	Y	Y	Global
	Y	N	N
disable_gtid_unsafe_statements [2020]	Y	Y	Global
	Y	N	N
disconnect-slave-event-count [1971]	Y	N	Global
	Y	N	N
enforce-gtid-consistency [2018]	Y	Y	Global
	Y	N	N
enforce_gtid_consistency [2021]	Y	Y	Global
	Y	N	N
gtid_done [2020]	N	Y	Both
	N	N	N
gtid_executed [2021]	N	Y	Both
	N	N	N
gtid_lost [2022]	N	Y	Global
	N	N	N
gtid-mode [2019]	Y	Y	Global
	Y	N	N
gtid_mode [2022]	N	Y	Global
	N	N	N
gtid_next [2022]	N	Y	Session
	N	N	Y
gtid_owned [2023]	N	Y	Both
	N	N	N
gtid_purged [2023]	N	Y	Global
	N	N	Y
init_slave [1987]	Y	Y	Global
	Y	N	Y
log-slave-updates [1971]	Y	Y	Global

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	Y	N	N
log_slave_updates [2008]	Y	Y	Global
	Y	N	N
master-info-file [1973]	Y	N	Global
	Y	N	N
master-info-repository [2000]	Y	N	Global
	Y	N	N
master_info_repository [1988]	N	Y	Global
	N	N	Y
master-retry-count [1973]	Y	N	Global
	Y	N	N
relay-log [1974]	Y	Y	Global
	Y	N	N
relay_log_basename [1989]	N	Y	Global
	N	N	N
relay-log-index [1974]	Y	Y	Global
	Y	N	N
relay-log-info-file [1975]	Y	N	Global
	Y	N	N
relay_log_info_file [1989]	Y	Y	Global
	Y	N	N
relay-log-info-repository [2000]	Y	N	Global
	Y	N	N
relay_log_info_repository [1990]	N	Y	Global
	N	N	Y
relay_log_index [1989]	Y	Y	Global
	Y	N	N
relay_log_purge [534]	Y	Y	Global
	Y	N	Y
relay_log_recovery [1990]	Y	Y	Global
	Y	N	Y
relay_log_space_limit [534]	Y	Y	Global
	Y	N	N
replicate-do-db [1976]	Y	N	Global
	Y	N	N
replicate-do-table [1978]	Y	N	Global
	Y	N	N
replicate-ignore-db [1977]	Y	N	Global

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	Y	N	N
replicate-ignore-table [1978]	Y	N	Global
	Y	N	N
replicate-rewrite-db [1979]	Y	N	Global
	Y	N	N
replicate-same-server-id [1979]	Y	N	Global
	Y	N	N
replicate-wild-do-table [1979]	Y	N	Global
	Y	N	N
replicate-wild-ignore-table [1980]	Y	N	Global
	Y	N	N
report-host [1980]	Y	Y	Global
	Y	N	N
report-password [1981]	Y	Y	Global
	Y	N	N
report-port [1981]	Y	Y	Global
	Y	N	N
report-user [1982]	Y	Y	Global
	Y	N	N
rpl_semi_sync_master_enabled [536]	N	Y	Global
	N	N	Y
rpl_semi_sync_master_timeout [536]	N	Y	Global
	N	N	Y
rpl_semi_sync_master_trace_level [536]	N	Y	Global
	N	N	Y
rpl_semi_sync_master_wait_no_slave [537]	N	Y	Global
	N	N	Y
rpl_semi_sync_slave_enabled [537]	N	Y	Global
	N	N	Y
rpl_semi_sync_slave_trace_level [538]	N	Y	Global
	N	N	Y
server_uuid [1955]	N	Y	Global
	N	N	N
show-slave-auth-info [1982]	Y	N	Global
	Y	N	N
skip-slave-start [1984]	Y	N	Global
	Y	N	N
slave_allow_batching [1968]	Y	Y	Global

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	Y	N	Y
slave-load-tmpdir [1985]	Y	Y	Global
	Y	N	N
slave-skip-errors [1986]	Y	Y	Global
	Y	N	N
slave-checkpoint-group [1982]	Y	N	Global
	Y	N	N
slave_checkpoint_group [1991]	Y	Y	Global
	Y	N	Y
slave-checkpoint-period [1983]	Y	N	Global
	Y	N	N
slave_checkpoint_period [1991]	Y	Y	Global
	Y	N	Y
slave_compressed_protocol [1992]	Y	Y	Global
	Y	N	Y
slave-max-allowed-packet [1964]	Y	N	Global
	Y	N	N
slave_max_allowed_packet [1969]	N	Y	Global
	N	N	Y
slave_net_timeout [1985]	Y	Y	Global
	Y	N	Y
slave_parallel_workers [1993]	N	Y	Global
	N	N	Y
slave-parallel-workers [1983]	Y	N	Global
	Y	N	N
slave-pending-jobs-size-max [1984]	Y	N	Global
	N	N	N
slave_pending_jobs_size_max [1994]	N	Y	Global
	N	N	Y
slave-rows-search-algorithms [1964]	Y	N	Global
	Y	N	N
slave_rows_search_algorithms [1969]	N	Y	Global
	N	N	Y
slave_transaction_retries [1995]	Y	Y	Global
	Y	N	Y
slave_type_conversions [1996]	Y	Y	Global
	Y	N	N
sql_slave_skip_counter [1996]	N	Y	Global

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
	N	N	Y
sync_master_info [1997]	Y	Y	Global
	Y	N	Y
sync_relay_log [1998]	Y	Y	Global
	Y	N	Y
sync_relay_log_info [1999]	Y	Y	Global
	Y	N	Y

Section 16.1.4.2, “Replication Master Options and Variables”, provides more detailed information about options and variables relating to replication master servers. For more information about options and variables relating to replication slaves, see Section 16.1.4.3, “Replication Slave Options and Variables”.

**Table 16.2. Binary Logging Options and Variables: MySQL 5.6**

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
Binlog_cache_disk_use [586]	N	N	Global
	N	Y	N
binlog_row_image [2012]	Y	Y	Both
	N	N	Y
binlog_rows_query_log_events [2014]	N	Y	Both
	N	N	Y
Binlog_stmt_cache_disk_use [586]	N	N	Global
	N	Y	N
Binlog_cache_use [586]	N	N	Global
	N	Y	N
Binlog_stmt_cache_use [586]	N	N	Global
	N	Y	N
Com_show_binlog_events [586]	N	N	Both
	N	Y	N
Com_show_binlogs [586]	N	N	Both
	N	Y	N
binlog-checksum [2006]	Y	N	Global
	Y	N	N
binlog_checksum [2009]	N	Y	Global
	N	N	Y
binlog-do-db [2004]	Y	N	Global
	Y	N	N
binlog-ignore-db [2005]	Y	N	Global
	Y	N	N

Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
binlog-row-event-max-size [2001]	Y	N	Global
	Y	N	N
binlog_cache_size [2008]	Y	Y	Global
	Y	N	Y
binlog_max_flush_queue_time [2011]	N	Y	Global
	N	N	Y
binlog_order_commits [2012]	N	Y	Global
	N	N	Y
binlog_stmt_cache_size [2014]	Y	Y	Global
	Y	N	Y
binlog_format [415]	Y	Y	Both
	Y	N	Y
binlog_rows_query_log_events [2007]	Y	N	Global
	Y	N	N
binlog_direct_non_transactional_updates [2010]	Y	Y	Both
	Y	N	Y
log_bin_basename [2015]	N	Y	Global
	N	N	N
log-bin-use-v1-row-events [2003]	Y	Y	Global
	Y	N	N
log_bin_use_v1_row_events [2015]	Y	Y	Global
	Y	N	N
master-verify-checksum [2006]	Y	N	Global
	Y	N	N
master_verify_checksum [2015]	N	Y	Global
	N	N	Y
max-binlog-dump-events [2007]	Y	N	Global
	Y	N	N
max_binlog_cache_size [2016]	Y	Y	Global
	Y	N	Y
max_binlog_size [2016]	Y	Y	Global
	Y	N	Y
max_binlog_stmt_cache_size [2017]	Y	Y	Global
	Y	N	Y
slave-sql-verify-checksum [1987]	Y	N	Global
	Y	N	N
slave_sql_verify_checksum [1995]	N	Y	Global
	N	N	Y



Option or Variable Name	Command Line	System Variable	Scope
	Option File	Status Variable	Dynamic
<code>sporadic-binlog-dump-fail</code> [2007]	Y	N	Global
	Y	N	N

Section 16.1.4.4, “Binary Log Options and Variables”, provides more detailed information about options and variables relating to binary logging. For additional general information about the binary log, see Section 5.2.4, “The Binary Log”.

For a table showing *all* command-line options, system and status variables used with `mysqld`, see Section 5.1.1, “Server Option and Variable Reference”.

### 16.1.4.2. Replication Master Options and Variables

This section describes the server options and system variables that you can use on replication master servers. You can specify the options either on the [command line](#) or in an [option file](#). You can specify system variable values using [SET](#).

On the master and each slave, you must use the [server-id](#) [1955] option to establish a unique replication ID. For each server, you should pick a unique positive integer in the range from 1 to  $2^{32} - 1$ , and each ID must be different from every other ID in use by any other replication master or slave. Example: `server-id=3`.

For options used on the master for controlling binary logging, see Section 16.1.4.4, “Binary Log Options and Variables”.

**Startup options used on replication masters.** The following startup options are used for controlling replication masters:

- `slave-max-allowed-packet=bytes` [1964]

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--slave-max-allowed-packet=#</code>	
<b>Option-File Format</b>	<code>slave-max-allowed-packet</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_max_allowed_packet</code> [1969]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1073741824</code>
	<b>Range</b>	<code>1024 .. 1073741824</code>

In MySQL 5.6.6 and later, this option sets the maximum packet size in bytes for the slave SQL and I/O threads, so that large updates using row-based replication do not cause replication to fail because an update exceeded `max_allowed_packet` [502]. (Bug #12400221, BUG#60926)

The corresponding server variable `slave_max_allowed_packet` [1969] always has a value that is a positive integer multiple of 1024; if you set it to some value that is not such a multiple, the value is automatically rounded down to the next highest multiple of 1024. (For example, if you start the server with `--slave-max-allowed-packet=10000`, the value used is 9216; setting 0 as the value causes 1024 to be used.) A truncation warning is issued in such cases.

The maximum (and default) value is 1073741824 (1 GB); the minimum is 1024.

- `slave-rows-search-algorithms=list` [1964]

<b>Version Introduced</b>	5.6.6		
<b>Command-Line Format</b>	<code>--slave-rows-search-algorithms=list</code>		
<b>Option-File Format</b>	<code>slave-rows-search-algorithms</code>		
<b>Option Sets Variable</b>	Yes, <code>slave_rows_search_algorithms</code> [1969]		
	<b>Permitted Values</b>		
	<b>Type</b>	set	
	<b>Valid Values</b>	TABLE_SCAN, default	INDEX_SCAN
		INDEX_SCAN, HASH_SCAN	
		TABLE_SCAN, HASH_SCAN	
		TABLE_SCAN, equivalent to INDEX_SCAN, HASH_SCAN	

When preparing batches of rows for row-based logging and replication using `slave_allow_batching` [1968], this option controls how the rows are searched for matches—that is, whether or not hashing is used for searches using a primary or unique key, some other key, or no key at all. This option takes a comma-separated list of any 2 (or possibly 3) values from the list `INDEX_SCAN`, `TABLE_SCAN`, `HASH_SCAN`. The list need not be quoted, but must contain no spaces, whether or not quotes are used. Possible combinations (lists) and their effects are shown in the following table:

<b>Index used / option value</b>	<code>INDEX_SCAN, HASH_SCAN</code>	<code>INDEX_SCAN, TABLE_SCAN</code>	<code>TABLE_SCAN, HASH_SCAN</code>
<i>Primary key or unique key</i>	Index scan	index scan	Index hash
<i>(Other) Key</i>	Index hash	Index scan	Index hash
<i>No index</i>	Table hash	Table scan	Table hash

The order in which the algorithms are specified in the list does not make any difference in the order in which they are displayed by a `SELECT` or `SHOW VARIABLES` statement (which is the same as that used in the table just shown previously). The default value is `TABLE_SCAN, INDEX_SCAN`, which means that all searches that can use indexes do use them, and searches without any indexes use table scans.

Specifying `INDEX_SCAN, TABLE_SCAN, HASH_SCAN` has the same effect as specifying `INDEX_SCAN, HASH_SCAN`. To use hashing for any searches that does not use a primary or unique key, set this option to `INDEX_SCAN, HASH_SCAN`. To force hashing for *all* searches, set it to `TABLE_SCAN, HASH_SCAN`.

This option was added in MySQL 5.6.6.

**System variables used on replication masters.** The following system variables are used in controlling replication masters:

- `auto_increment_increment` [1965]

<b>Command-Line Format</b>	<code>--auto-increment-increment [=#]</code>
<b>Option-File Format</b>	<code>auto_increment_increment</code>
<b>Option Sets Variable</b>	Yes, <code>auto_increment_increment</code> [1965]
<b>Variable Name</b>	<code>auto_increment_increment</code>
<b>Variable Scope</b>	Global, Session

<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	numeric
<b>Default</b>	1
<b>Range</b>	1 .. 65535

`auto_increment_increment` [1965] and `auto_increment_offset` [1968] are intended for use with master-to-master replication, and can be used to control the operation of `AUTO_INCREMENT` columns. Both variables have global and session values, and each can assume an integer value between 1 and 65,535 inclusive. Setting the value of either of these two variables to 0 causes its value to be set to 1 instead. Attempting to set the value of either of these two variables to an integer greater than 65,535 or less than 0 causes its value to be set to 65,535 instead. Attempting to set the value of `auto_increment_increment` [1965] or `auto_increment_offset` [1968] to a noninteger value gives rise to an error, and the actual value of the variable remains unchanged.



### Note

`auto_increment_increment` [1965] is intended for use with MySQL Cluster, which is not currently supported in MySQL 5.6.

These two variables affect `AUTO_INCREMENT` column behavior as follows:

- `auto_increment_increment` [1965] controls the interval between successive column values. For example:

```
mysql> SHOW VARIABLES LIKE 'auto_inc%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| auto_increment_increment | 1 |
| auto_increment_offset | 1 |
+-----+-----+
2 rows in set (0.00 sec)

mysql> CREATE TABLE autoinc1
-> (col INT NOT NULL AUTO_INCREMENT PRIMARY KEY);
Query OK, 0 rows affected (0.04 sec)

mysql> SET @@auto_increment_increment=10;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE 'auto_inc%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| auto_increment_increment | 10 |
| auto_increment_offset | 1 |
+-----+-----+
2 rows in set (0.01 sec)

mysql> INSERT INTO autoinc1 VALUES (NULL), (NULL), (NULL), (NULL);
Query OK, 4 rows affected (0.00 sec)
Records: 4 Duplicates: 0 Warnings: 0

mysql> SELECT col FROM autoinc1;
+-----+
| col |
+-----+
| 1 |
| 11 |
| 21 |
| 31 |
+-----+
4 rows in set (0.00 sec)
```

- `auto_increment_offset` [1968] determines the starting point for the `AUTO_INCREMENT` column value. Consider the following, assuming that these statements are executed during the same session as the example given in the description for `auto_increment_increment` [1965]:

```
mysql> SET @@auto_increment_offset=5;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE 'auto_inc%';
+-----+-----+
| Variable_name          | Value |
+-----+-----+
| auto_increment_increment | 10    |
| auto_increment_offset   | 5     |
+-----+-----+
2 rows in set (0.00 sec)

mysql> CREATE TABLE autoinc2
-> (col INT NOT NULL AUTO_INCREMENT PRIMARY KEY);
Query OK, 0 rows affected (0.06 sec)

mysql> INSERT INTO autoinc2 VALUES (NULL), (NULL), (NULL), (NULL);
Query OK, 4 rows affected (0.00 sec)
Records: 4 Duplicates: 0 Warnings: 0

mysql> SELECT col FROM autoinc2;
+-----+
| col |
+-----+
| 5   |
| 15  |
| 25  |
| 35  |
+-----+
4 rows in set (0.02 sec)
```

If the value of `auto_increment_offset` [1968] is greater than that of `auto_increment_increment` [1965], the value of `auto_increment_offset` [1968] is ignored.

Should one or both of these variables be changed and then new rows inserted into a table containing an `AUTO_INCREMENT` column, the results may seem counterintuitive because the series of `AUTO_INCREMENT` values is calculated without regard to any values already present in the column, and the next value inserted is the least value in the series that is greater than the maximum existing value in the `AUTO_INCREMENT` column. In other words, the series is calculated like so:

$$\text{auto\_increment\_offset} + N \times \text{auto\_increment\_increment}$$

where  $N$  is a positive integer value in the series [1, 2, 3, ...]. For example:

```
mysql> SHOW VARIABLES LIKE 'auto_inc%';
+-----+-----+
| Variable_name          | Value |
+-----+-----+
| auto_increment_increment | 10    |
| auto_increment_offset   | 5     |
+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT col FROM autoinc1;
+-----+
| col |
+-----+
| 1   |
| 11  |
| 21  |
| 31  |
+-----+
4 rows in set (0.00 sec)
```

```
mysql> INSERT INTO autoinc1 VALUES (NULL), (NULL), (NULL), (NULL);
Query OK, 4 rows affected (0.00 sec)
Records: 4 Duplicates: 0 Warnings: 0

mysql> SELECT col FROM autoinc1;
+-----+
| col |
+-----+
| 1 |
| 11 |
| 21 |
| 31 |
| 35 |
| 45 |
| 55 |
| 65 |
+-----+
8 rows in set (0.00 sec)
```

The values shown for `auto_increment_increment` [1965] and `auto_increment_offset` [1968] generate the series  $5 + N \times 10$ , that is, [5, 15, 25, 35, 45, ...]. The greatest value present in the `col` column prior to the `INSERT` is 31, and the next available value in the `AUTO_INCREMENT` series is 35, so the inserted values for `col` begin at that point and the results are as shown for the `SELECT` query.

It is not possible to confine the effects of these two variables to a single table, and thus they do not take the place of the sequences offered by some other database management systems; these variables control the behavior of all `AUTO_INCREMENT` columns in *all* tables on the MySQL server. If the global value of either variable is set, its effects persist until the global value is changed or overridden by setting the session value, or until `mysqld` is restarted. If the local value is set, the new value affects `AUTO_INCREMENT` columns for all tables into which new rows are inserted by the current user for the duration of the session, unless the values are changed during that session.

The default value of `auto_increment_increment` [1965] is 1. See [Section 16.4.1.1, “Replication and AUTO\\_INCREMENT”](#).

- `auto_increment_offset` [1968]

<b>Command-Line Format</b>	<code>--auto-increment-offset [=#]</code>
<b>Option-File Format</b>	<code>auto_increment_offset</code>
<b>Option Sets Variable</b>	Yes, <code>auto_increment_offset</code> [1968]
<b>Variable Name</b>	<code>auto_increment_offset</code>
<b>Variable Scope</b>	Global, Session
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>1</code>
	<b>Range</b> <code>1 .. 65535</code>

This variable has a default value of 1. For particulars, see the description for `auto_increment_increment` [1965].



#### Note

`auto_increment_offset` [1968] is intended for use with MySQL Cluster, which is not currently supported in MySQL 5.6.

- `slave_allow_batching` [1968]

<b>Command-Line Format</b>	<code>--slave-allow-batching</code>
<b>Option-File Format</b>	<code>slave_allow_batching</code>
<b>Option Sets Variable</b>	Yes, <code>slave_allow_batching</code> [1968]
<b>Variable Name</b>	<code>slave_allow_batching</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>off</code>

Whether or not batched updates are enabled on replication slaves.

- `slave_max_allowed_packet` [1969]

<b>Version Introduced</b>	5.6.6
<b>Variable Name</b>	<code>slave_max_allowed_packet</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>1073741824</code>
	<b>Range</b> <code>1024 .. 1073741824</code>

In MySQL 5.6.6 and later, this variable sets the maximum packet size for the slave SQL and I/O threads, so that large updates using row-based replication do not cause replication to fail because an update exceeded `max_allowed_packet` [502].

This global variable always has a value that is a positive integer multiple of 1024; if you set it to some value that is not, the value is rounded down to the next highest multiple of 1024 for it is stored or used; setting `slave_max_allowed_packet` to 0 causes 1024 to be used. (A truncation warning is issued in all such cases.) The default and maximum value is 1073741824 (1 GB); the minimum is 1024.

`slave_max_allowed_packet` can also be set at startup, using the `--slave-max-allowed-packet` [1964] option.

- `slave_rows_search_algorithms` [1969]

<b>Version Introduced</b>	5.6.6
<b>Variable Name</b>	<code>slave_rows_search_algorithms=list</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>set</code>
	<b>Valid Values</b> <code>TABLE_SCAN, default, INDEX_SCAN</code>
	<code>INDEX_SCAN, HASH_SCAN</code>
	<code>TABLE_SCAN, HASH_SCAN</code>

		TABLE_SCAN, INDEX_SCAN, HASH_SCAN	equivalent to INDEX_SCAN, HASH_SCAN
--	--	-----------------------------------	-------------------------------------

When preparing batches of rows for row-based logging and replication using `slave_allow_batching` [1968], the `slave_rows_search_algorithms` variable controls how the rows are searched for matches—that is, whether or not hashing is used for searches using a primary or unique key, using some other key, or using no key at all. This option takes a comma-separated list of at least 2 values from the list `INDEX_SCAN`, `TABLE_SCAN`, `HASH_SCAN`. The value expected as a string, so the value must be quoted. In addition, the value must not contain any spaces. Possible combinations (lists) and their effects are shown in the following table:

Index used / option value	<code>INDEX_SCAN, HASH_SCAN</code> or <code>INDEX_SCAN, TABLE_SCAN, HASH_SCAN</code>	<code>INDEX_SCAN, TABLE_SCAN</code>	<code>TABLE_SCAN, HASH_SCAN</code>
Primary key or unique key	Index scan	index scan	Index hash
(Other) Key	Index hash	Index scan	Index hash
No index	Table hash	Table scan	Table hash

The order in which the algorithms are specified in the list does not make any difference in the order in which they are displayed by a `SELECT` or `SHOW VARIABLES` statement, as shown here:

```
mysql> SET GLOBAL slave_rows_search_algorithms = "INDEX_SCAN, TABLE_SCAN";
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE '%algorithms%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| slave_rows_search_algorithms | TABLE_SCAN, INDEX_SCAN |
+-----+-----+
1 row in set (0.00 sec)

mysql> SET GLOBAL slave_rows_search_algorithms = "TABLE_SCAN, INDEX_SCAN";
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE '%algorithms%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| slave_rows_search_algorithms | TABLE_SCAN, INDEX_SCAN |
+-----+-----+
1 row in set (0.00 sec)
```

The default value is `TABLE_SCAN, INDEX_SCAN`, which means that all searches that can use indexes do use them, and searches without any indexes use table scans.

Specifying `INDEX_SCAN, TABLE_SCAN, HASH_SCAN` has the same effect as specifying `INDEX_SCAN, HASH_SCAN`. To use hashing for any searches that does not use a primary or unique key, set this variable to `INDEX_SCAN, HASH_SCAN`. To force hashing for *all* searches, set it to `TABLE_SCAN, HASH_SCAN`.

This variable was added in MySQL 5.6.6.

### 16.1.4.3. Replication Slave Options and Variables

This section describes the server options and system variables that apply to slave replication servers. You can specify the options either on the [command line](#) or in an [option file](#). Many of the options can be set while the server is running by using the `CHANGE MASTER TO` statement. You can specify system variable values using `SET`.



**Server ID.** On the master and each slave, you must use the `server-id` [1955] option to establish a unique replication ID in the range from 1 to  $2^{32} - 1$ . “Unique” means that each ID must be different from every other ID in use by any other replication master or slave. Example `my.cnf` file:

```
[mysqld]
server-id=3
```

**Startup options for replication slaves.** The following list describes startup options for controlling replication slave servers. Many of these options can be set while the server is running by using the `CHANGE MASTER TO` statement. Others, such as the `--replicate-*` options, can be set only when the slave server starts. Replication-related system variables are discussed later in this section.

- `--abort-slave-event-count` [1971]

<b>Command-Line Format</b>	<code>--abort-slave-event-count=#</code>	
<b>Option-File Format</b>	<code>abort-slave-event-count</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Min Value</b>	<code>0</code>

When this option is set to some positive integer *value* other than 0 (the default) it affects replication behavior as follows: After the slave SQL thread has started, *value* log events are permitted to be executed; after that, the slave SQL thread does not receive any more events, just as if the network connection from the master were cut. The slave thread continues to run, and the output from `SHOW SLAVE STATUS` displays `Yes` in both the `Slave_IO_Running` and the `Slave_SQL_Running` columns, but no further events are read from the relay log.

This option is used internally by the MySQL test suite for replication testing and debugging. It is not intended for use in a production setting.

- `--disconnect-slave-event-count` [1971]

<b>Command-Line Format</b>	<code>--disconnect-slave-event-count=#</code>	
<b>Option-File Format</b>	<code>disconnect-slave-event-count</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

This option is used internally by the MySQL test suite for replication testing and debugging.

- `--log-slave-updates` [1971]

<b>Command-Line Format</b>	<code>--log-slave-updates</code>	
<b>Option-File Format</b>	<code>log-slave-updates</code>	
<b>Option Sets Variable</b>	Yes, <code>log_slave_updates</code> [2008]	
<b>Variable Name</b>	<code>log_slave_updates</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Normally, a slave does not log to its own binary log any updates that are received from a master server. This option tells the slave to log the updates performed by its SQL thread to its own binary log. For this option to have any effect, the slave must also be started with the `--log-bin` [2002] option to enable binary logging. Prior to MySQL 5.5, the server would not start when using the `--log-slave-updates` [1971] option without also starting the server with the `--log-bin` [2002] option, and would fail with an error; in MySQL 5.6, only a warning is generated. (Bug #44663) `--log-slave-updates` [1971] is used when you want to chain replication servers. For example, you might want to set up replication servers using this arrangement:

```
A -> B -> C
```

Here, `A` serves as the master for the slave `B`, and `B` serves as the master for the slave `C`. For this to work, `B` must be both a master *and* a slave. You must start both `A` and `B` with `--log-bin` [2002] to enable binary logging, and `B` with the `--log-slave-updates` [1971] option so that updates received from `A` are logged by `B` to its binary log.

- `--log-slow-slave-statements` [1972]

<b>Command-Line Format</b>	<code>--log-slow-slave-statements</code>	
<b>Option-File Format</b>	<code>log-slow-slave-statements</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

When the slow query log is enabled, this option enables logging for queries that have taken more than `long_query_time` [500] seconds to execute on the slave.

- `--log-warnings[=level]` [430]

<b>Command-Line Format</b>	<code>--log-warnings[=#]</code>	
	<code>-W [#]</code>	
<b>Option-File Format</b>	<code>log-warnings</code>	
<b>Option Sets Variable</b>	Yes, <code>log_warnings</code> [499]	
<b>Variable Name</b>	<code>log_warnings</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
<b>Disabled by</b>	<code>skip-log-warnings</code>	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	<code>64</code>
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1</code>
	<b>Range</b>	<code>0 .. 18446744073709547520</code>

This option causes a server to print more messages to the error log about what it is doing. With respect to replication, the server generates warnings that it succeeded in reconnecting after a network/connection failure, and informs you as to how each slave thread started. This option is enabled by default; to disable it, use `--skip-log-warnings` [430]. If the value is greater than 1, aborted connections are written to the error log, and access-denied errors for new connection attempts are written. See Section C.5.2.11, “Communication Errors and Aborted Connections”.

Note that the effects of this option are not limited to replication. It produces warnings across a spectrum of server activities.

- `--master-info-file=file_name` [1973]

<b>Command-Line Format</b>	<code>--master-info-file=file_name</code>	
<b>Option-File Format</b>	<code>master-info-file=file_name</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>master.info</code>

The name to use for the file in which the slave records information about the master. The default name is `master.info` in the data directory. For information about the format of this file, see [Section 16.2.2.2, “Slave Status Logs”](#).

- `--master-retry-count=count` [1973]

<b>Version Deprecated</b>	5.6.1	
<b>Command-Line Format</b>	<code>--master-retry-count=#</code>	
<b>Option-File Format</b>	<code>master-retry-count</code>	
<b>Deprecated</b>	5.6.1	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	86400
	<b>Range</b>	0 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	86400
	<b>Range</b>	0 .. 18446744073709551615

The number of times that the slave tries to connect to the master before giving up. Reconnects are attempted at intervals set by the `MASTER_CONNECT_RETRY` option of the `CHANGE MASTER TO` statement (default 60). Reconnects are triggered when data reads by the slave time out according to the `--slave-net-timeout` [1985] option. The default value is 86400. A value of 0 means “infinite”; the slave attempts to connect forever.

This option is deprecated as of MySQL 5.6.1 and will be removed in a future MySQL release. Applications should be updated to use the `MASTER_RETRY_COUNT` option of the `CHANGE MASTER TO` statement instead.

- `--max-relay-log-size=size` [1973]

The size at which the server rotates relay log files automatically. For more information, see [Section 16.2.2, “Replication Relay and Status Logs”](#). The default size is 1GB.

- `--read-only` [1973]

Cause the slave to permit no updates except from slave threads or from users having the [SUPER \[756\]](#) privilege. On a slave server, this can be useful to ensure that the slave accepts updates only from its master server and not from clients. This variable does not apply to [TEMPORARY](#) tables.

- `--relay-log=file_name` [\[1974\]](#)

<b>Command-Line Format</b>	<code>--relay-log=name</code>	
<b>Option-File Format</b>	<code>relay-log</code>	
<b>Option Sets Variable</b>	Yes, <code>relay_log</code> <a href="#">[1989]</a>	
<b>Variable Name</b>	<code>relay-log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The basename for the relay log. The default basename is `host_name-relay-bin`. The server writes the file in the data directory unless the basename is given with a leading absolute path name to specify a different directory. The server creates relay log files in sequence by adding a numeric suffix to the basename.

Due to the manner in which MySQL parses server options, if you specify this option, you must supply a value; *the default basename is used only if the option is not actually specified*. If you use the `--relay-log` [\[1974\]](#) option without specifying a value, unexpected behavior is likely to result; this behavior depends on the other options used, the order in which they are specified, and whether they are specified on the command line or in an option file. For more information about how MySQL handles server options, see [Section 4.2.3, “Specifying Program Options”](#).

If you specify this option, the value specified is also used as the basename for the relay log index file. You can override this behavior by specifying a different relay log index file basename using the `--relay-log-index` [\[1974\]](#) option.

Starting with MySQL 5.6.5, when the server reads an entry from the index file, it checks whether the entry contains a relative path. If it does, the relative part of the path is replaced with the absolute path set using the `--relay-log` option. An absolute path remains unchanged; in such a case, the index must be edited manually to enable the new path or paths to be used. Prior to MySQL 5.6.5, manual intervention was required whenever relocating the binary log or relay log files. (Bug #11745230, Bug #12133)

You may find the `--relay-log` [\[1974\]](#) option useful in performing the following tasks:

- Creating relay logs whose names are independent of host names.
- If you need to put the relay logs in some area other than the data directory because your relay logs tend to be very large and you do not want to decrease `max_relay_log_size` [\[507\]](#).
- To increase speed by using load-balancing between disks.

Beginning with MySQL 5.6.2, you can obtain the relay log filename (and path) from the `relay_log_basename` [\[1989\]](#) system variable.

- `--relay-log-index=file_name` [\[1974\]](#)

<b>Command-Line Format</b>	<code>--relay-log-index=name</code>	
<b>Option-File Format</b>	<code>relay-log-index</code>	
<b>Option Sets Variable</b>	Yes, <code>relay_log_index</code> <a href="#">[1989]</a>	

<b>Variable Name</b>	<code>relay-log-index</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The name to use for the relay log index file. The default name is `host_name-relay-bin.index` in the data directory, where `host_name` is the name of the slave server.

Due to the manner in which MySQL parses server options, if you specify this option, you must supply a value; *the default basename is used only if the option is not actually specified*. If you use the `--relay-log-index` [1974] option without specifying a value, unexpected behavior is likely to result; this behavior depends on the other options used, the order in which they are specified, and whether they are specified on the command line or in an option file. For more information about how MySQL handles server options, see [Section 4.2.3, “Specifying Program Options”](#).

If you specify this option, the value specified is also used as the basename for the relay logs. You can override this behavior by specifying a different relay log file basename using the `--relay-log` [1974] option.

- `--relay-log-info-file=file_name` [1975]

<b>Command-Line Format</b>	<code>--relay-log-info-file=file_name</code>	
<b>Option-File Format</b>	<code>relay-log-info-file</code>	
<b>Option Sets Variable</b>	Yes, <code>relay_log_info_file</code> [1989]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>relay-log.info</code>

The name to use for the file in which the slave records information about the relay logs. The default name is `relay-log.info` in the data directory. For information about the format of this file, see [Section 16.2.2.2, “Slave Status Logs”](#).

- `--relay-log-purge={0|1}` [1975]

Disable or enable automatic purging of relay logs as soon as they are no longer needed. The default value is 1 (enabled). This is a global variable that can be changed dynamically with `SET GLOBAL relay_log_purge = N`.

- `--relay-log-recovery={0|1}` [1975]

<b>Command-Line Format</b>	<code>--relay-log-recovery</code>	
<b>Option-File Format</b>	<code>relay-log-recovery</code>	
<b>Option Sets Variable</b>	Yes, <code>relay_log_recovery</code> [1990]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Enables automatic relay log recovery immediately following server startup, which means that the replication slave discards all unprocessed relay logs and retrieves them from the replication master. This should be used following a crash on the replication slave to ensure that no possibly corrupted relay logs are processed. The default value is 0 (disabled). This option must be enabled in order to provide a crash-proof slave.

Prior to MySQL 5.6.6, if this option is enabled for a multi-threaded slave, and the slave fails with errors, you cannot execute `CHANGE MASTER TO` on that slave. In MySQL 5.6.6 or later, you can use `START SLAVE UNTIL SQL_AFTER_MTS_GAPS` to ensure that any gaps in the relay log are processed; after running this statement, you can then use `CHANGE MASTER TO` to fail this slave over to a new master. (Bug #13893363)

- `--relay-log-space-limit=size` [1976]

This option places an upper limit on the total size in bytes of all relay logs on the slave. A value of 0 means “no limit.” This is useful for a slave server host that has limited disk space. When the limit is reached, the I/O thread stops reading binary log events from the master server until the SQL thread has caught up and deleted some unused relay logs. Note that this limit is not absolute: There are cases where the SQL thread needs more events before it can delete relay logs. In that case, the I/O thread exceeds the limit until it becomes possible for the SQL thread to delete some relay logs because not doing so would cause a deadlock. You should not set `--relay-log-space-limit` [1976] to less than twice the value of `--max-relay-log-size` [1973] (or `--max-binlog-size` [2016] if `--max-relay-log-size` [1973] is 0). In that case, there is a chance that the I/O thread waits for free space because `--relay-log-space-limit` [1976] is exceeded, but the SQL thread has no relay log to purge and is unable to satisfy the I/O thread. This forces the I/O thread to ignore `--relay-log-space-limit` [1976] temporarily.

- `--replicate-do-db=db_name` [1976]

<b>Command-Line Format</b>	<code>--replicate-do-db=name</code>	
<b>Option-File Format</b>	<code>replicate-do-db</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The effects of this option depend on whether statement-based or row-based replication is in use.

**Statement-based replication.** Tell the slave SQL thread to restrict replication to statements where the default database (that is, the one selected by `USE`) is `db_name`. To specify more than one database, use this option multiple times, once for each database; however, doing so does *not* replicate cross-database statements such as `UPDATE some_db.some_table SET foo='bar'` while a different database (or no database) is selected.



### Warning

To specify multiple databases you *must* use multiple instances of this option. Because database names can contain commas, if you supply a comma separated list then the list will be treated as the name of a single database.

An example of what does not work as you might expect when using statement-based replication: If the slave is started with `--replicate-do-db=sales` [1976] and you issue the following statements on the master, the `UPDATE` statement is *not* replicated:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

The main reason for this “check just the default database” behavior is that it is difficult from the statement alone to know whether it should be replicated (for example, if you are using multiple-table `DELETE` statements or multiple-table `UPDATE` statements that act across multiple databases). It is also faster to check only the default database rather than all databases if there is no need.

**Row-based replication.** Tells the slave SQL thread to restrict replication to database `db_name`. Only tables belonging to `db_name` are changed; the current database has no effect on this. Suppose that the slave is started with `--replicate-do-db=sales` [1976] and row-based replication is in effect, and then the following statements are run on the master:

```
USE prices;
UPDATE sales.february SET amount=amount+100;
```

The `february` table in the `sales` database on the slave is changed in accordance with the `UPDATE` statement; this occurs whether or not the `USE` statement was issued. However, issuing the following statements on the master has no effect on the slave when using row-based replication and `--replicate-do-db=sales` [1976]:

```
USE prices;
UPDATE prices.march SET amount=amount-25;
```

Even if the statement `USE prices` were changed to `USE sales`, the `UPDATE` statement's effects would still not be replicated.

Another important difference in how `--replicate-do-db` [1976] is handled in statement-based replication as opposed to row-based replication occurs with regard to statements that refer to multiple databases. Suppose that the slave is started with `--replicate-do-db=db1` [1976], and the following statements are executed on the master:

```
USE db1;
UPDATE db1.table1 SET col1 = 10, db2.table2 SET col2 = 20;
```

If you are using statement-based replication, then both tables are updated on the slave. However, when using row-based replication, only `table1` is affected on the slave; since `table2` is in a different database, `table2` on the slave is not changed by the `UPDATE`. Now suppose that, instead of the `USE db1` statement, a `USE db4` statement had been used:

```
USE db4;
UPDATE db1.table1 SET col1 = 10, db2.table2 SET col2 = 20;
```

In this case, the `UPDATE` statement would have no effect on the slave when using statement-based replication. However, if you are using row-based replication, the `UPDATE` would change `table1` on the slave, but not `table2`—in other words, only tables in the database named by `--replicate-do-db` [1976] are changed, and the choice of default database has no effect on this behavior.

If you need cross-database updates to work, use `--replicate-wild-do-table=db_name.%` [1979] instead. See Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”.



**Note**

This option affects replication in the same manner that `--binlog-do-db` [2004] affects binary logging, and the effects of the replication format on how `--replicate-do-db` [1976] affects replication behavior are the same as those of the logging format on the behavior of `--binlog-do-db` [2004].

This option has no effect on `BEGIN`, `COMMIT`, or `ROLLBACK` statements.

- `--replicate-ignore-db=db_name` [1977]

<b>Command-Line Format</b>	<code>--replicate-ignore-db=name</code>	
<b>Option-File Format</b>	<code>replicate-ignore-db</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

As with `--replicate-do-db` [1976], the effects of this option depend on whether statement-based or row-based replication is in use.

**Statement-based replication.** Tells the slave SQL thread not to replicate any statement where the default database (that is, the one selected by `USE`) is `db_name`.



**Row-based replication.** Tells the slave SQL thread not to update any tables in the database `db_name`. The default database has no effect.

When using statement-based replication, the following example does not work as you might expect. Suppose that the slave is started with `--replicate-ignore-db=sales [1977]` and you issue the following statements on the master:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

The `UPDATE` statement is replicated in such a case because `--replicate-ignore-db [1977]` applies only to the default database (determined by the `USE` statement). Because the `sales` database was specified explicitly in the statement, the statement has not been filtered. However, when using row-based replication, the `UPDATE` statement's effects are *not* propagated to the slave, and the slave's copy of the `sales.january` table is unchanged; in this instance, `--replicate-ignore-db=sales [1977]` causes *all* changes made to tables in the master's copy of the `sales` database to be ignored by the slave.

To specify more than one database to ignore, use this option multiple times, once for each database. Because database names can contain commas, if you supply a comma separated list then the list will be treated as the name of a single database.

You should not use this option if you are using cross-database updates and you do not want these updates to be replicated. See [Section 16.2.3, "How Servers Evaluate Replication Filtering Rules"](#).

If you need cross-database updates to work, use `--replicate-wild-ignore-table=db_name.% [1980]` instead. See [Section 16.2.3, "How Servers Evaluate Replication Filtering Rules"](#).



#### Note

This option affects replication in the same manner that `--binlog-ignore-db [2005]` affects binary logging, and the effects of the replication format on how `--replicate-ignore-db [1977]` affects replication behavior are the same as those of the logging format on the behavior of `--binlog-ignore-db [2005]`.

This option has no effect on `BEGIN`, `COMMIT`, or `ROLLBACK` statements.

- `--replicate-do-table=db_name.tbl_name [1978]`

<b>Command-Line Format</b>	<code>--replicate-do-table=name</code>
<b>Option-File Format</b>	<code>replicate-do-table</code>
	<b>Permitted Values</b>
	<b>Type</b> <code>string</code>

Tells the slave SQL thread to restrict replication to the specified table. To specify more than one table, use this option multiple times, once for each table. This works for both cross-database updates and default database updates, in contrast to `--replicate-do-db [1976]`. See [Section 16.2.3, "How Servers Evaluate Replication Filtering Rules"](#).

This option affects only statements that apply to tables. It does not affect statements that apply only to other database objects, such as stored routines. To filter statements operating on stored routines, use one or more of the `--replicate-*-db` options.

- `--replicate-ignore-table=db_name.tbl_name [1978]`

<b>Command-Line Format</b>	<code>--replicate-ignore-table=name</code>
<b>Option-File Format</b>	<code>replicate-ignore-table</code>

	<b>Permitted Values</b>	
	<b>Type</b>	string

Tells the slave SQL thread not to replicate any statement that updates the specified table, even if any other tables might be updated by the same statement. To specify more than one table to ignore, use this option multiple times, once for each table. This works for cross-database updates, in contrast to `--replicate-ignore-db` [1977]. See [Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#).

This option affects only statements that apply to tables. It does not affect statements that apply only to other database objects, such as stored routines. To filter statements operating on stored routines, use one or more of the `--replicate-*-db` options.

- `--replicate-rewrite-db=from_name->to_name` [1979]

<b>Command-Line Format</b>	<code>--replicate-rewrite-db=old_name-&gt;new_name</code>	
<b>Option-File Format</b>	<code>replicate-rewrite-db</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	string

Tells the slave to translate the default database (that is, the one selected by `USE`) to `to_name` if it was `from_name` on the master. Only statements involving tables are affected (not statements such as `CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE`), and only if `from_name` is the default database on the master. This does not work for cross-database updates. To specify multiple rewrites, use this option multiple times. The server uses the first one with a `from_name` value that matches. The database name translation is done *before* the `--replicate-*` rules are tested.

If you use this option on the command line and the “>” character is special to your command interpreter, quote the option value. For example:

```
shell> mysqld --replicate-rewrite-db="olddb->newdb"
```



#### Note

Prior to MySQL 5.6.7, multithreaded slaves did not honor this option correctly. (Bug #14232958)

- `--replicate-same-server-id` [1979]

<b>Command-Line Format</b>	<code>--replicate-same-server-id</code>	
<b>Option-File Format</b>	<code>replicate-same-server-id</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	FALSE

To be used on slave servers. Usually you should use the default setting of 0, to prevent infinite loops caused by circular replication. If set to 1, the slave does not skip events having its own server ID. Normally, this is useful only in rare configurations. Cannot be set to 1 if `--log-slave-updates` [1971] is used. By default, the slave I/O thread does not write binary log events to the relay log if they have the slave's server ID (this optimization helps save disk usage). If you want to use `--replicate-same-server-id` [1979], be sure to start the slave with this option before you make the slave read its own events that you want the slave SQL thread to execute.

- `--replicate-wild-do-table=db_name.tbl_name` [1979]

<b>Command-Line Format</b>	<code>--replicate-wild-do-table=name</code>
----------------------------	---------------------------------------------

<b>Option-File Format</b>	<code>replicate-wild-do-table</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

Tells the slave thread to restrict replication to statements where any of the updated tables match the specified database and table name patterns. Patterns can contain the “%” and “\_” wildcard characters, which have the same meaning as for the `LIKE` [1190] pattern-matching operator. To specify more than one table, use this option multiple times, once for each table. This works for cross-database updates. See [Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#).

This option applies to tables, views, and triggers. It does not apply to stored procedures and functions, or events. To filter statements operating on the latter objects, use one or more of the `--replicate-*-db` options.

Example: `--replicate-wild-do-table=foo%.bar%` [1979] replicates only updates that use a table where the database name starts with `foo` and the table name starts with `bar`.

If the table name pattern is `%`, it matches any table name and the option also applies to database-level statements (`CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE`). For example, if you use `--replicate-wild-do-table=foo%.%` [1979], database-level statements are replicated if the database name matches the pattern `foo%`.

To include literal wildcard characters in the database or table name patterns, escape them with a backslash. For example, to replicate all tables of a database that is named `my_own%db`, but not replicate tables from the `my1ownAABCdb` database, you should escape the “\_” and “%” characters like this: `--replicate-wild-do-table=my\_own\%db` [1979]. If you use the option on the command line, you might need to double the backslashes or quote the option value, depending on your command interpreter. For example, with the `bash` shell, you would need to type `--replicate-wild-do-table=my\\_own\\%db` [1979].

- `--replicate-wild-ignore-table=db_name.tbl_name` [1980]

<b>Command-Line Format</b>	<code>--replicate-wild-ignore-table=name</code>	
<b>Option-File Format</b>	<code>replicate-wild-ignore-table</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

Tells the slave thread not to replicate a statement where any table matches the given wildcard pattern. To specify more than one table to ignore, use this option multiple times, once for each table. This works for cross-database updates. See [Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#).

Example: `--replicate-wild-ignore-table=foo%.bar%` [1980] does not replicate updates that use a table where the database name starts with `foo` and the table name starts with `bar`.

For information about how matching works, see the description of the `--replicate-wild-do-table` [1979] option. The rules for including literal wildcard characters in the option value are the same as for `--replicate-wild-ignore-table` [1980] as well.

- `--report-host=host_name` [1980]

<b>Command-Line Format</b>	<code>--report-host=host_name</code>
<b>Option-File Format</b>	<code>report-host</code>
<b>Option Sets Variable</b>	Yes, <code>report_host</code> [535]
<b>Variable Name</b>	<code>report-host</code>

<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	string

The host name or IP address of the slave to be reported to the master during slave registration. This value appears in the output of `SHOW SLAVE HOSTS` on the master server. Leave the value unset if you do not want the slave to register itself with the master. Note that it is not sufficient for the master to simply read the IP address of the slave from the TCP/IP socket after the slave connects. Due to NAT and other routing issues, that IP may not be valid for connecting to the slave from the master or other hosts.

- `--report-password=password` [1981]

<b>Command-Line Format</b>	<code>--report-password=name</code>
<b>Option-File Format</b>	<code>report-password</code>
<b>Option Sets Variable</b>	Yes, <code>report_password</code> [535]
<b>Variable Name</b>	<code>report-password</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	string

The account password of the slave to be reported to the master during slave registration. This value appears in the output of `SHOW SLAVE HOSTS` on the master server if the `--show-slave-auth-info` [1982] option is given.

Although the name of this option might imply otherwise, `--report-password` is not connected to the MySQL user privilege system and so is not necessarily (or even likely to be) the same as the password for the MySQL replication user account.

- `--report-port=slave_port_num` [1981]

<b>Command-Line Format</b>	<code>--report-port=#</code>
<b>Option-File Format</b>	<code>report-port</code>
<b>Option Sets Variable</b>	Yes, <code>report_port</code> [535]
<b>Variable Name</b>	<code>report-port</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&gt;= 5.6.5)</b>
<b>Type</b>	numeric
<b>Default</b>	<code>[slave_port]</code>
<b>Range</b>	0 .. 65535

The TCP/IP port number for connecting to the slave, to be reported to the master during slave registration. Set this only if the slave is listening on a nondefault port or if you have a special tunnel from the master or other clients to the slave. If you are not sure, do not use this option.

Prior to MySQL 5.6.5, the default value for this option was 3306. In MySQL 5.6.5 and later, the value shown is the port number actually used by the slave (Bug #13333431). This change also affects the default value displayed by `SHOW SLAVE HOSTS`.

- `--report-user=user_name` [1982]

<b>Command-Line Format</b>	<code>--report-user=name</code>	
<b>Option-File Format</b>	<code>report-user</code>	
<b>Option Sets Variable</b>	Yes, <code>report_user</code> [536]	
<b>Variable Name</b>	<code>report-user</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The account user name of the slave to be reported to the master during slave registration. This value appears in the output of `SHOW SLAVE HOSTS` on the master server if the `--show-slave-auth-info` [1982] option is given.

Although the name of this option might imply otherwise, `--report-user` is not connected to the MySQL user privilege system and so is not necessarily (or even likely to be) the same as the name of the MySQL replication user account.

- `--show-slave-auth-info` [1982]

<b>Command-Line Format</b>	<code>--show-slave-auth-info</code>	
<b>Option-File Format</b>	<code>show-slave-auth-info</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Display slave user names and passwords in the output of `SHOW SLAVE HOSTS` on the master server for slaves started with the `--report-user` [1982] and `--report-password` [1981] options.

- `--slave-checkpoint-group=#` [1982]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-checkpoint-group=#</code>	
<b>Option-File Format</b>	<code>slave-checkpoint-group</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_checkpoint_group</code> [1991]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>512</code>
	<b>Range</b>	<code>32 .. 524280</code>
	<b>Block Size</b>	<code>8</code>

Sets the maximum number of transactions that can be processed by a multi-threaded slave before a checkpoint operation is called to update its status as shown by `SHOW SLAVE STATUS`. Setting this option has no effect on slaves for which multithreading is not enabled.

This option works in combination with the `--slave-checkpoint-period` [1983] option in such a way that, when either limit is exceeded, the checkpoint is executed and the counters tracking both the number of transactions and the time elapsed since the last checkpoint are reset.

The minimum allowed value for this option is 32, unless the server was built using `-DWITH_DEBUG` [136], in which case the minimum value is 1. The effective value is always a multiple of 8; you can set it to a value that is not such a multiple, but the server rounds it down to the next lower multiple of 8 before storing the value. (*Exception:* No such rounding is performed by the debug server.) Regardless of how the server was built, the default value is 512, and the maximum allowed value is 524280.

`--slave-checkpoint-group` [1982] was added in MySQL 5.6.3.

- `--slave-checkpoint-period=#` [1983]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-checkpoint-period=#</code>	
<b>Option-File Format</b>	<code>slave-checkpoint-period</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_checkpoint_period</code> [1991]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	300
	<b>Range</b>	1 .. 4G

Sets the maximum time (in milliseconds) that is allowed to pass before a checkpoint operation is called to update the status of a multi-threaded slave as shown by `SHOW SLAVE STATUS`. Setting this option has no effect on slaves for which multithreading is not enabled.

This option works in combination with the `--slave-checkpoint-group` [1982] option in such a way that, when either limit is exceeded, the checkpoint is executed and the counters tracking both the number of transactions and the time elapsed since the last checkpoint are reset.

The minimum allowed value for this option is 1, unless the server was built using `-DWITH_DEBUG` [136], in which case the minimum value is 0. Regardless of how the server was built, the default value is 300, and the maximum possible value is 4294967296 (4GB).

`--slave-checkpoint-period` [1983] was added in MySQL 5.6.3.

- `--slave-parallel-workers` [1983]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-parallel-workers=#</code>	
<b>Option-File Format</b>	<code>slave-parallel-workers</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_parallel_workers</code> [1993]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	0
	<b>Range</b>	0 .. 1024

Sets the number of slave worker threads for executing replication events (transactions) in parallel. Setting this variable to 0 (the default) disables parallel execution. The maximum is 1024.

When parallel execution is enabled, the slave SQL thread acts as the coordinator for the slave worker threads, among which transactions are distributed on a per-database basis. This means that a worker thread on the slave can process successive transactions on a given database without waiting for updates to other databases to complete. The current implementation of multithreading on the slave assumes that the data is partitioned per database, and that updates within a

given database occur in the same relative order as they do on the master, in order to work correctly. However, transactions do not need to be coordinated between any two databases.

Due to the fact that transactions on different databases can occur in a different order on the slave than on the master, checking for the most recently executed transaction does not guarantee that all previous transactions from the master have been executed on the slave. This has implications for logging and recovery when using a multi-threaded slave. For information about how to interpret binary logging information when using multi-threading on the slave, see [Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”](#). In addition, this means that `START SLAVE UNTIL` is not supported with a multi-threaded slave.

You should also note that, in MySQL 5.6.7 and later, enforcing foreign key relationships between tables in different databases causes multi-threaded slaves to use sequential rather than parallel mode, which can have a negative impact on performance. (Bug #14092635)

This option was added in MySQL 5.6.3.



**Note**

The value set for this option (or for the corresponding `slave_parallel_workers` [1993] system variable) was not always honored correctly in MySQL 5.6.3; this problem was fixed in MySQL 5.6.4 (Bug #13334470).

- `--slave-pending-jobs-size-max=#` [1984]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-pending-jobs-size-max=#</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_pending_jobs_size_max</code> [1994]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>16M</code>
	<b>Range</b>	<code>1024 .. 18EB</code>
	<b>Block Size</b>	<code>1024</code>

For multithreaded slaves, this option sets the maximum amount of memory (in bytes) available to slave worker queues holding events not yet applied. Setting this option has no effect on slaves for which multithreading is not enabled.

The minimum possible value for this option is 1024; the default is 16MB. The maximum possible value is 18446744073709551615 (16 exabytes). Values that are not exact multiples of 1024 are rounded down to the next-highest multiple of 1024 prior to being stored.



**Important**

The value for this option must not be less than the master's value for `max_allowed_packet` [502]; otherwise a slave worker queue may become full while there remain events coming from the master to be processed.

This option was added in MySQL 5.6.3.

- `--skip-slave-start` [1984]

<b>Command-Line Format</b>	<code>--skip-slave-start</code>
<b>Option-File Format</b>	<code>skip-slave-start</code>



	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	FALSE	

Tells the slave server not to start the slave threads when the server starts. To start the threads later, use a `START SLAVE` statement.

- `--slave_compressed_protocol={0|1}` [1992]

<b>Command-Line Format</b>	<code>--slave_compressed_protocol</code>	
<b>Option-File Format</b>	<code>slave_compressed_protocol</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_compressed_protocol</code> [1992]	
<b>Variable Name</b>	<code>slave_compressed_protocol</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	OFF	

If this option is set to 1, use compression for the slave/master protocol if both the slave and the master support it. The default is 0 (no compression).

- `--slave-load-tmpdir=file_name` [1985]

<b>Command-Line Format</b>	<code>--slave-load-tmpdir=path</code>	
<b>Option-File Format</b>	<code>slave-load-tmpdir</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_load_tmpdir</code> [1992]	
<b>Variable Name</b>	<code>slave_load_tmpdir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
<b>Type</b>	file name	
<b>Default</b>	/tmp	

The name of the directory where the slave creates temporary files. This option is by default equal to the value of the `tmpdir` [559] system variable. When the slave SQL thread replicates a `LOAD DATA INFILE` statement, it extracts the file to be loaded from the relay log into temporary files, and then loads these into the table. If the file loaded on the master is huge, the temporary files on the slave are huge, too. Therefore, it might be advisable to use this option to tell the slave to put temporary files in a directory located in some file system that has a lot of available space. In that case, the relay logs are huge as well, so you might also want to use the `--relay-log` [1974] option to place the relay logs in that file system.

The directory specified by this option should be located in a disk-based file system (not a memory-based file system) because the temporary files used to replicate `LOAD DATA INFILE` must survive machine restarts. The directory also should not be one that is cleared by the operating system during the system startup process.

- `--slave-net-timeout=seconds` [1985]

<b>Command-Line Format</b>	<code>--slave-net-timeout=#</code>
----------------------------	------------------------------------

<b>Option-File Format</b>	<code>slave-net-timeout</code>
<b>Option Sets Variable</b>	Yes, <code>slave_net_timeout</code> [1993]
<b>Variable Name</b>	<code>slave_net_timeout</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	<code>numeric</code>
<b>Default</b>	<code>3600</code>
<b>Min Value</b>	<code>1</code>

The number of seconds to wait for more data from the master before the slave considers the connection broken, aborts the read, and tries to reconnect. The first retry occurs immediately after the timeout. The interval between retries is controlled by the `MASTER_CONNECT_RETRY` option for the `CHANGE MASTER TO` statement, and the number of reconnection attempts is limited by the `--master-retry-count` [1973] option. The default is 3600 seconds (one hour).

- `--slave-skip-errors=[err_code1,err_code2,...|all]` [1986]

<b>Command-Line Format</b>	<code>--slave-skip-errors=name</code>
<b>Option-File Format</b>	<code>slave-skip-errors</code>
<b>Option Sets Variable</b>	Yes, <code>slave_skip_errors</code> [1994]
<b>Variable Name</b>	<code>slave_skip_errors</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>
<b>Default</b>	<code>OFF</code>
<b>Valid Values</b>	<code>[list of error codes]</code> <code>all</code> <code>ddl_exist_errors</code>

Normally, replication stops when an error occurs on the slave. This gives you the opportunity to resolve the inconsistency in the data manually. This option tells the slave SQL thread to continue replication when a statement returns any of the errors listed in the option value.

Do not use this option unless you fully understand why you are getting errors. If there are no bugs in your replication setup and client programs, and no bugs in MySQL itself, an error that stops replication should never occur. Indiscriminate use of this option results in slaves becoming hopelessly out of synchrony with the master, with you having no idea why this has occurred.

For error codes, you should use the numbers provided by the error message in your slave error log and in the output of `SHOW SLAVE STATUS`. [Appendix C, Errors, Error Codes, and Common Problems](#), lists server error codes.

You can also (but should not) use the very nonrecommended value of `all` to cause the slave to ignore all error messages and keeps going regardless of what happens. Needless to say, if you use `all`, there are no guarantees regarding the integrity of your data. Please do not complain (or file bug reports) in this case if the slave's data is not anywhere close to what it is on the master. *You have been warned.*

MySQL 5.6 supports an additional shorthand value `ddl_exists_errors`, which is equivalent to the error code list `1007,1008,4050,1051,1054,1060,1061,1068,1094,1146`.

Examples:

```
--slave-skip-errors=1062,1053
--slave-skip-errors=all
--slave-skip-errors=ddl_exist_errors
```

- `--slave-sql-verify-checksum={0|1}` [1987]

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--slave-sql-verify-checksum=value</code>	
<b>Option-File Format</b>	<code>slave-sql-verify-checksum</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	0
	<b>Valid Values</b>	0 1

When this option is enabled, the slave examines checksums read from the relay log, in the event of a mismatch, the slave stops with an error. Disabled by default.

This option was added in MySQL 5.6.2.

**Obsolete options.** *The following options are removed in MySQL 5.6. If you attempt to start `mysqld` with any of these options in MySQL 5.6, the server aborts with an `unknown variable` error. To set the replication parameters formerly associated with these options, you must use the `CHANGE MASTER TO ...` statement (see Section 13.4.2.1, “CHANGE MASTER TO Syntax”).*

The options affected are shown in this list:

- `--master-host`
- `--master-user`
- `--master-password`
- `--master-port`
- `--master-connect-retry`
- `--master-ssl`
- `--master-ssl-ca`
- `--master-ssl-capath`
- `--master-ssl-cert`
- `--master-ssl-cipher`
- `--master-ssl-key`

**System variables used on replication slaves.** The following list describes system variables for controlling replication slave servers. They can be set at server startup and some of them can be changed at runtime using `SET`. Server options used with replication slaves are listed earlier in this section.

- `init_slave` [1987]

<b>Command-Line Format</b>	<code>--init-slave=name</code>
<b>Option-File Format</b>	<code>init_slave</code>
<b>Option Sets Variable</b>	Yes, <code>init_slave</code> [1987]
<b>Variable Name</b>	<code>init_slave</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>string</code>

This variable is similar to `init_connect` [489], but is a string to be executed by a slave server each time the SQL thread starts. The format of the string is the same as for the `init_connect` [489] variable.



#### Note

The SQL thread sends an acknowledgment to the client before it executes `init_slave` [1987]. Therefore, it is not guaranteed that `init_slave` [1987] has been executed when `START SLAVE` returns. See [Section 13.4.2.5, “START SLAVE Syntax”](#), for more information.

- `log_slow_slave_statements` [1988]

<b>Version Introduced</b>	5.6.11
<b>Variable Name</b>	<code>log_slow_slave_statements</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>OFF</code>

When the slow query log is enabled, this variable enables logging for queries that have taken more than `long_query_time` [500] seconds to execute on the slave. This variable was added in MySQL 5.6.11.

- `master_info_repository` [1988]

<b>Version Introduced</b>	5.6.2
<b>Variable Name</b>	<code>master_info_repository</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>string</code>
	<b>Default</b> <code>FILE</code>
	<b>Valid Values</b> <code>FILE</code> <code>TABLE</code>

The setting of this variable determines whether the slave logs master status and connection information to a `FILE` (`master.info`), or to a `TABLE` (`mysql.slave_master_info`).

This variable was added in MySQL 5.6.2.

- [relay\\_log \[1989\]](#)

<b>Command-Line Format</b>	<code>--relay-log=name</code>	
<b>Option-File Format</b>	<code>relay-log</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">relay_log [1989]</a>	
<b>Variable Name</b>	<code>relay-log</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>

The name of the relay log file.

- [relay\\_log\\_basename \[1989\]](#)

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	<code>relay_log_basename</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>datadir + '/' + hostname + '-relay-bin'</code>

Holds the name and complete path to the relay log file.

The [relay\\_log\\_basename \[1989\]](#) system variable was added in MySQL 5.6.2.

- [relay\\_log\\_index \[1989\]](#)

<b>Command-Line Format</b>	<code>--relay-log-index</code>	
<b>Option-File Format</b>	<code>relay_log_index</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">relay_log_index [1989]</a>	
<b>Variable Name</b>	<code>relay_log_index</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>*host_name*-relay-bin.index</code>

The name of the relay log index file. The default name is `host_name-relay-bin.index` in the data directory, where `host_name` is the name of the slave server.

- [relay\\_log\\_info\\_file \[1989\]](#)

<b>Command-Line Format</b>	<code>--relay-log-info-file=file_name</code>	
<b>Option-File Format</b>	<code>relay_log_info_file</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">relay_log_info_file [1989]</a>	

<b>Variable Name</b>	<code>relay_log_info_file</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>
<b>Default</b>	<code>relay-log.info</code>

The name of the file in which the slave records information about the relay logs. The default name is `relay-log.info` in the data directory.

- [`relay\_log\_info\_repository` \[1990\]](#)

<b>Version Introduced</b>	5.6.2
<b>Variable Name</b>	<code>relay_log_info_repository</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	<code>string</code>
<b>Default</b>	<code>FILE</code>
<b>Valid Values</b>	<code>FILE</code> <code>TABLE</code>

This variable determines whether the slave's position in the relay logs is written to a `FILE` (`relay-log.info`) or to a `TABLE` (`mysql.slave_relay_log_info`).

This variable was added in MySQL 5.6.2.

- [`relay\_log\_recovery` \[1990\]](#)

<b>Command-Line Format</b>	<code>--relay-log-recovery</code>
<b>Option-File Format</b>	<code>relay_log_recovery</code>
<b>Option Sets Variable</b>	Yes, <a href="#"><code>relay_log_recovery</code> [1990]</a>
<b>Variable Name</b>	<code>relay_log_recovery</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
<b>Type</b>	<code>boolean</code>
<b>Default</b>	<code>FALSE</code>

Enables automatic relay log recovery immediately following server startup, which means that the replication slave discards all unprocessed relay logs and retrieves them from the replication master. In MySQL 5.6.5 and earlier, it was possible to change this global variable dynamically; beginning with MySQL 5.6.6, it is read-only. (Bug #13840948) Regardless of the MySQL Server version, its value can be changed by starting the slave with the `--relay-log-recovery` [\[1975\]](#) option, which should be used following a crash on the replication slave to ensure that no possibly corrupted relay logs are processed, and must be used in order to guarantee a crash-proof slave. The default value is 0 (disabled).

When `relay_log_recovery` is enabled and the slave has stopped due to errors encountered while running in multi-threaded mode, you cannot execute `CHANGE MASTER TO` if there are

any gaps in the log. Beginning with MySQL 5.6.6, you should use `START SLAVE UNTIL SQL_AFTER_MTS_GAPS` to ensure that all gaps are processed before switching back to single-threaded mode or executing a `CHANGE MASTER TO` statement.

- `slave_checkpoint_group` [1991]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-checkpoint-group=#</code>	
<b>Option-File Format</b>	<code>slave_checkpoint_group</code>	
<b>Variable Name</b>	<code>slave_checkpoint_group=#</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>512</code>
	<b>Range</b>	<code>32 .. 524280</code>
	<b>Block Size</b>	<code>8</code>

Sets the maximum number of transactions that can be processed by a multi-threaded slave before a checkpoint operation is called to update its status as shown by `SHOW SLAVE STATUS`. Setting this variable has no effect on slaves for which multithreading is not enabled.

This variable works in combination with the `slave_checkpoint_period` [1991] system variable in such a way that, when either limit is exceeded, the checkpoint is executed and the counters tracking both the number of transactions and the time elapsed since the last checkpoint are reset.

The minimum allowed value for this variable is 32, unless the server was built using `-DWITH_DEBUG` [136], in which case the minimum value is 1. The effective value is always a multiple of 8; you can set it to a value that is not such a multiple, but the server rounds it down to the next lower multiple of 8 before storing the value. (*Exception:* No such rounding is performed by the debug server.) Regardless of how the server was built, the default value is 512, and the maximum allowed value is 524280.

`slave_checkpoint_group` was added in MySQL 5.6.3.

- `slave_checkpoint_period` [1991]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--slave-checkpoint-period=#</code>	
<b>Option-File Format</b>	<code>slave_checkpoint_period</code>	
<b>Variable Name</b>	<code>slave_checkpoint_period=#</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>300</code>
	<b>Range</b>	<code>1 .. 4G</code>

Sets the maximum time (in milliseconds) that is allowed to pass before a checkpoint operation is called to update the status of a multi-threaded slave as shown by `SHOW SLAVE STATUS`. Setting this variable has no effect on slaves for which multithreading is not enabled.



This variable works in combination with the `slave_checkpoint_group` [1991] system variable in such a way that, when either limit is exceeded, the checkpoint is executed and the counters tracking both the number of transactions and the time elapsed since the last checkpoint are reset.

The minimum allowed value for this variable is 1, unless the server was built using `-DWITH_DEBUG` [136], in which case the minimum value is 0. Regardless of how the server was built, the default value is 300, and the maximum possible value is 4294967296 (4GB).

`slave_checkpoint_period` was added in MySQL 5.6.3.

- `slave_compressed_protocol` [1992]

<b>Command-Line Format</b>	<code>--slave_compressed_protocol</code>	
<b>Option-File Format</b>	<code>slave_compressed_protocol</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_compressed_protocol</code> [1992]	
<b>Variable Name</b>	<code>slave_compressed_protocol</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Whether to use compression of the slave/master protocol if both the slave and the master support it.

- `slave_exec_mode` [1992]

<b>Command-Line Format</b>	<code>--slave-exec-mode=mode</code>	
<b>Option-File Format</b>	<code>slave_exec_mode</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_exec_mode</code> [1992]	
<b>Variable Name</b>	<code>slave_exec_mode</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>STRICT (ALL)</code>
	<b>Default</b>	<code>IDEMPOTENT (NDB)</code>
	<b>Valid Values</b>	<code>IDEMPOTENT</code> <code>STRICT</code>

Controls whether `IDEMPOTENT` or `STRICT` mode is used in replication conflict resolution and error checking. `IDEMPOTENT` mode causes suppression of duplicate-key and no-key-found errors. This mode should be employed in multi-master replication, circular replication, and some other special replication scenarios. `STRICT` mode is the default, and is suitable for most other cases.

- `slave_load_tmpdir` [1992]

<b>Command-Line Format</b>	<code>--slave-load-tmpdir=path</code>
<b>Option-File Format</b>	<code>slave-load-tmpdir</code>
<b>Option Sets Variable</b>	Yes, <code>slave_load_tmpdir</code> [1992]

<b>Variable Name</b>	<code>slave_load_tmpdir</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>/tmp</code>

The name of the directory where the slave creates temporary files for replicating `LOAD DATA INFILE` statements.

- `slave_net_timeout` [1993]

<b>Command-Line Format</b>	<code>--slave-net-timeout=#</code>	
<b>Option-File Format</b>	<code>slave-net-timeout</code>	
<b>Option Sets Variable</b>	Yes, <code>slave_net_timeout</code> [1993]	
<b>Variable Name</b>	<code>slave_net_timeout</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>3600</code>
	<b>Min Value</b>	<code>1</code>

The number of seconds to wait for more data from a master/slave connection before aborting the read.

- `slave_parallel_workers` [1993]

<b>Version Introduced</b>	5.6.3	
<b>Variable Name</b>	<code>slave_parallel_workers</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>
	<b>Range</b>	<code>0 .. 1024</code>

Sets the number of slave worker threads for executing replication events (transactions) in parallel. Setting this variable to 0 (the default) disables parallel execution. The maximum is 1024.

When parallel execution is enabled, the slave SQL thread acts as the coordinator for the slave worker threads, among which transactions are distributed on a per-database basis. This means that a worker thread on the slave can process successive transactions on a given database without waiting for updates to other databases to complete. The current implementation of multi-threading on the slave assumes that the data is partitioned per database, and that updates within a given database occur in the same relative order as they do on the master, in order to work correctly. However, transactions do not need to be coordinated between any two databases.

Due to the fact that transactions on different databases can occur in a different order on the slave than on the master, checking for the most recently executed transaction does not guarantee that all previous transactions from the master have been executed on the slave. This has implications for logging and recovery when using a multi-threaded slave. For information about how to interpret binary logging information when using multi-threading on the slave, see [Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”](#). In addition, this means that `START SLAVE UNTIL` is not supported with a multi-threaded slave.

This variable was added in MySQL 5.6.3.



#### Note

The value set for this variable (or for the corresponding `--slave-parallel-workers` [1983] option) was not always honored correctly in MySQL 5.6.3; this problem was fixed in MySQL 5.6.4 (Bug #13334470).

- `slave_pending_jobs_size_max` [1994]

<b>Version Introduced</b>	5.6.3
<b>Variable Name</b>	<code>slave_pending_jobs_size_max</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>0</code>
	<b>Range</b> <code>1024 .. 18EB</code>
	<b>Block Size</b> <code>1024</code>

For multithreaded slaves, this variable sets the maximum amount of memory (in bytes) available to slave worker queues holding events not yet applied. Setting this variable has no effect on slaves for which multithreading is not enabled.

The minimum possible value for this variable is 1024; the default is 16MB. The maximum possible value is 18446744073709551615 (16 exabytes). Values that are not exact multiples of 1024 are rounded down to the next-highest multiple of 1024 prior to being stored.



#### Important

The value of this variable must not be less than the master's value for `max_allowed_packet` [502]; otherwise a slave worker queue may become full while there remain events coming from the master to be processed.

`slave_pending_jobs_size_max` was added in MySQL 5.6.3.

- `slave_skip_errors` [1994]

<b>Command-Line Format</b>	<code>--slave-skip-errors=name</code>
<b>Option-File Format</b>	<code>slave-skip-errors</code>
<b>Option Sets Variable</b>	Yes, <code>slave_skip_errors</code> [1994]
<b>Variable Name</b>	<code>slave_skip_errors</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

	<b>Permitted Values</b>	
<b>Type</b>	string	
<b>Default</b>	OFF	
<b>Valid Values</b>	[list of error codes] all ddl_exist_errors	

Normally, replication stops when an error occurs on the slave. This gives you the opportunity to resolve the inconsistency in the data manually. This variable tells the slave SQL thread to continue replication when a statement returns any of the errors listed in the variable value.

- [slave\\_sql\\_verify\\_checksum \[1995\]](#)

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	slave_sql_verify_checksum	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	0	
<b>Valid Values</b>	0 1	

Cause the slave SQL thread to verify data using the checksums read from the relay log. In the event of a mismatch, the slave stops with an error.



**Note**

The slave I/O thread always reads checksums if possible when accepting events from over the network.

slave\_sql\_verify\_checksum was added in MySQL 5.6.2.

- [slave\\_transaction\\_retries \[1995\]](#)

<b>Command-Line Format</b>	--slave_transaction_retries=#	
<b>Option-File Format</b>	slave_transaction_retries	
<b>Option Sets Variable</b>	Yes, slave_transaction_retries [1995]	
<b>Variable Name</b>	slave_transaction_retries	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Platform Bit Size</b>	32	
<b>Type</b>	numeric	
<b>Default</b>	10	
<b>Range</b>	0 .. 4294967295	
	<b>Permitted Values</b>	

<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	10
<b>Range</b>	0 .. 18446744073709547520

If a replication slave SQL thread fails to execute a transaction because of an InnoDB deadlock or because the transaction's execution time exceeded InnoDB's `innodb_lock_wait_timeout` [1729], it automatically retries `slave_transaction_retries` [1995] times before stopping with an error. The default value is 10.

- `slave_type_conversions` [1996]

<b>Command-Line Format</b>	<code>--slave_type_conversions=set</code>
<b>Option-File Format</b>	<code>slave_type_conversions</code>
<b>Option Sets Variable</b>	Yes, <code>slave_type_conversions</code> [1996]
<b>Variable Name</b>	<code>slave_type_conversions</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
	<b>Type</b> <code>set</code>
	<b>Default</b>
	<b>Valid Values</b> <code>ALL_LOSSY</code> <code>ALL_NON_LOSSY</code> <code>ALL_LOSSY,ALL_NON_LOSSY</code>

Controls the type conversion mode in effect on the slave when using row-based replication. Its value is a comma-delimited set of zero or more elements from the list: `ALL_LOSSY`, `ALL_NON_LOSSY`. Set this variable to an empty string to disallow type conversions between the master and the slave. Changes require a restart of the slave to take effect.

For additional information on type conversion modes applicable to attribute promotion and demotion in row-based replication, see [Row-based replication: attribute promotion and demotion](#).

- `sql_slave_skip_counter` [1996]

<b>Variable Name</b>	<code>sql_slave_skip_counter</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Type</b> <code>numeric</code>

The number of events from the master that a slave server should skip.

This option is incompatible with GTID-based replication, and must not be set to a nonzero value when `--gtid-mode=ON` [2019]. In MySQL 5.6.10 and later, trying to do so is specifically disallowed. (Bug #15833516) If you need to skip transactions when employing GTIDs, use `gtid_executed` [2021] from the master instead. See [Injecting empty transactions](#), for information about how to do this.



### Important

If skipping the number of events specified by setting this variable would cause the slave to begin in the middle of an event group, the slave continues to skip until it finds the beginning of the next event group and begins from that point. For more information, see [Section 13.4.2.4, “SET GLOBAL sql\\_slave\\_skip\\_counter Syntax”](#).

- [sync\\_master\\_info \[1997\]](#)

<b>Command-Line Format</b>	<code>--sync-master-info=#</code>
<b>Option-File Format</b>	<code>sync_master_info</code>
<b>Option Sets Variable</b>	Yes, <a href="#">sync_master_info [1997]</a>
<b>Variable Name</b>	<code>sync_master_info</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 0
	<b>Range</b> 0 .. 18446744073709547520
	<b>Permitted Values (&lt;= 5.6.5)</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 0
	<b>Range</b> 0 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>
	<b>Platform Bit Size</b> 64
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 10000
	<b>Range</b> 0 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.6)</b>
	<b>Platform Bit Size</b> 32
	<b>Type</b> <code>numeric</code>
	<b>Default</b> 10000
	<b>Range</b> 0 .. 4294967295

The effects of this variable on a replication slave depend on whether the slave's [master\\_info\\_repository \[1988\]](#) is set to `FILE` or `TABLE`, as explained in the following paragraphs.

**master\_info\_repository = FILE.** If the value of `sync_master_info` is greater than 0, the slave synchronizes its `master.info` file to disk (using `fdatasync()`) after every `sync_master_info` events. If it is 0, the MySQL server performs no synchronization of the

`master.info` file to disk; instead, the server relies on the operating system to flush its contents periodically as with any other file.

`master_info_repository = TABLE`. If the value of `sync_master_info` is greater than 0, the slave updates its master info repository table after every `sync_master_info` events. If it is 0, the table is never updated.

The default value for `sync_master_info` is 10000 as of MySQL 5.6.6, 0 before that.

- `sync_relay_log` [1998]

<b>Command-Line Format</b>	<code>--sync-relay-log=#</code>
<b>Option-File Format</b>	<code>sync_relay_log</code>
<b>Option Sets Variable</b>	Yes, <code>sync_relay_log</code> [1998]
<b>Variable Name</b>	<code>sync_relay_log</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	10000
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	10000
<b>Range</b>	0 .. 4294967295

If the value of this variable is greater than 0, the MySQL server synchronizes its relay log to disk (using `fdatasync()`) after every `sync_relay_log` [1998] writes to the relay log. There is one write to the relay log per statement if `autocommit` is enabled, and one write per transaction otherwise. A value of 0 does no synchronizing to disk—in this case, the server relies on the operating system to flush the relay log's contents from time to time as for any other file. A value of 1 is the safest choice because in the event of a crash you lose at most one statement or transaction from the relay log.



However, it is also the slowest choice (unless the disk has a battery-backed cache, which makes synchronization very fast). The default is 10000 as of MySQL 5.6.6, 0 before that.

- [sync\\_relay\\_log\\_info \[1999\]](#)

<b>Command-Line Format</b>	<code>--sync-relay-log-info=#</code>
<b>Option-File Format</b>	<code>sync_relay_log_info</code>
<b>Option Sets Variable</b>	Yes, <a href="#">sync_relay_log_info [1999]</a>
<b>Variable Name</b>	<code>sync_relay_log_info</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	0
<b>Range</b>	0 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	10000
<b>Range</b>	0 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	10000
<b>Range</b>	0 .. 4294967295

The effects of this variable on the slave depend on the server's [relay\\_log\\_info\\_repository \[1990\]](#) setting (`FILE` or `TABLE`), and if this is `TABLE`, additionally on whether the storage engine used by the relay log info table is transactional (such as `InnoDB`) or not (`MyISAM`). The effects of these factors on the behavior of the server for [sync\\_relay\\_log\\_info](#) values of zero and greater than zero are shown in the following table:

<code>sync_relay_log_info</code>	<code>relay_log_info_repository</code>		
	<b>FILE</b>	<b>TABLE</b>	
		<b>Transactional</b>	<b>Nontransactional</b>
$N > 0$	The slave synchronizes its <code>relay-log.info</code> file to disk (using <code>fdatasync()</code> ) after every $N$ transactions.	The table is updated after each transaction. ( $N$ is effectively ignored.)	The table is updated after every $N$ events.
0	The MySQL server performs no synchronization of the <code>relay-log.info</code> file to disk; instead, the server relies on the operating system to flush its contents periodically as with any other file.		The table is never updated.

The default value for `sync_relay_log_info` is 10000 as of MySQL 5.6.6, 0 before that.

**Options for logging slave status to tables.** MySQL 5.6 and later supports logging of replication slave status information to tables rather than files. Writing of the master info log and the relay log info can be configured separately using two server options added in MySQL 5.6.2 and listed here:

- `--master-info-repository={FILE|TABLE}` [2000]

<b>Version Introduced</b>	5.6.2
<b>Command-Line Format</b>	<code>--master-info-repository=FILE TABLE</code>
<b>Option-File Format</b>	<code>master-info-repository</code>
<b>Option Sets Variable</b>	Yes, <code>master_info_repository</code> [1988]
	<b>Permitted Values</b>
	<b>Type</b> <code>string</code>
	<b>Default</b> <code>FILE</code>
	<b>Valid Values</b> <code>FILE</code> <code>TABLE</code>

This option causes the server to write its master info log to a file or a table. The name of the file defaults to `master.info`; you can change the name of the file using the `--master-info-file` [1973] server option.

The default value for this option is `FILE`. If you use `TABLE`, the log is written to the `slave_master_info` table in the `mysql` database.

The `--master-info-repository` option was added in MySQL 5.6.2.

- `--relay-log-info-repository={FILE|TABLE}` [2000]

<b>Version Introduced</b>	5.6.2
<b>Command-Line Format</b>	<code>--relay-log-info-repository=FILE TABLE</code>
<b>Option-File Format</b>	<code>relay-log-info-repository</code>
<b>Option Sets Variable</b>	Yes, <code>relay_log_info_repository</code> [1990]
	<b>Permitted Values</b>
	<b>Type</b> <code>string</code>

	<b>Default</b>	FILE
	<b>Valid Values</b>	FILE TABLE

This option causes the server to log its relay log info to a file or a table. The name of the file defaults to `relay-log.info`; you can change the name of the file using the `--relay-log-info-file [1975]` server option.

The default value for this option is `FILE`. If you use `TABLE`, the log is written to the `slave_relay_log_info` table in the `mysql` database.

The `--relay-log-info-repository` option was added in MySQL 5.6.2.

The info log tables and their contents are considered local to a given MySQL Server. In MySQL 5.6.9 and later, they are not replicated, and changes to them are not written to the binary log. (Bug #14741537)

For more information, see [Section 16.2.2, “Replication Relay and Status Logs”](#).

#### 16.1.4.4. Binary Log Options and Variables

You can use the `mysqld` options and system variables that are described in this section to affect the operation of the binary log as well as to control which statements are written to the binary log. For additional information about the binary log, see [Section 5.2.4, “The Binary Log”](#). For additional information about using MySQL server options and system variables, see [Section 5.1.3, “Server Command Options”](#), and [Section 5.1.4, “Server System Variables”](#).

**Startup options used with binary logging.** The following list describes startup options for enabling and configuring the binary log. System variables used with binary logging are discussed later in this section.

- `--binlog-row-event-max-size=N [2001]`

<b>Command-Line Format</b>	<code>--binlog-row-event-max-size=#</code>
<b>Option-File Format</b>	<code>binlog-row-event-max-size</code>
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	1024
<b>Range</b>	256 .. 4294967295
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Platform Bit Size</b>	64
<b>Type</b>	numeric
<b>Default</b>	1024
<b>Range</b>	256 .. 18446744073709547520
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Platform Bit Size</b>	32
<b>Type</b>	numeric
<b>Default</b>	8192

	<b>Range</b>	256 .. 4294967295
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	8192
	<b>Range</b>	256 .. 18446744073709547520

Specify the maximum size of a row-based binary log event, in bytes. Rows are grouped into events smaller than this size if possible. The value should be a multiple of 256. The default is 8192 as of MySQL 5.6.6 and 1024 before that. See [Section 16.1.2, “Replication Formats”](#).

- `--log-bin[=base_name] [2002]`

<b>Command-Line Format</b>	<code>--log-bin</code>	
<b>Option-File Format</b>	<code>log-bin</code>	
<b>Variable Name</b>	<code>log_bin</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	file name
	<b>Default</b>	OFF

Enable binary logging. The server logs all statements that change data to the binary log, which is used for backup and replication. See [Section 5.2.4, “The Binary Log”](#).

The option value, if given, is the basename for the log sequence. The server creates binary log files in sequence by adding a numeric suffix to the basename. It is recommended that you specify a basename (see [Section C.5.8, “Known Issues in MySQL”](#), for the reason). Otherwise, MySQL uses `host_name-bin` as the basename.

In MySQL 5.6.5 and later, when the server reads an entry from the index file, it checks whether the entry contains a relative path, and if it does, the relative part of the path is replaced with the absolute path set using the `--log-bin` option. An absolute path remains unchanged; in such a case, the index must be edited manually to enable the new path or paths to be used. Previous to MySQL 5.6.5, manual intervention was required whenever relocating the binary log or relay log files. (Bug #11745230, Bug #12133)

Setting this option causes the `log_bin [2008]` system variable to be set to `ON` (or `1`), and not to the basename. Beginning with MySQL 5.6.2, the binary log filename (with path) is available as the `log_bin_basename [2015]` system variable.

- `--log-bin-index[=file_name] [2002]`

<b>Command-Line Format</b>	<code>--log-bin-index=name</code>	
<b>Option-File Format</b>	<code>log-bin-index</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	file name
	<b>Default</b>	OFF

The index file for binary log file names. See [Section 5.2.4, “The Binary Log”](#). If you omit the file name, and if you did not specify one with `--log-bin [2002]`, MySQL uses `host_name-bin.index` as the file name.

- `--log-bin-trust-function-creators[={0|1}]` [2003]

<b>Command-Line Format</b>	<code>--log-bin-trust-function-creators</code>	
<b>Option-File Format</b>	<code>log-bin-trust-function-creators</code>	
<b>Option Sets Variable</b>	Yes, <code>log_bin_trust_function_creators</code> [497]	
<b>Variable Name</b>	<code>log_bin_trust_function_creators</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

This option sets the corresponding `log_bin_trust_function_creators` [497] system variable. If no argument is given, the option sets the variable to 1. `log_bin_trust_function_creators` [497] affects how MySQL enforces restrictions on stored function and trigger creation. See [Section 18.7, “Binary Logging of Stored Programs”](#).

- `--log-bin-use-v1-row-events[={0|1}]` [2003]

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--log-bin-use-v1-row-events[={0 1}]</code>	
<b>Option-File Format</b>	<code>log-bin-use-v1-row-events</code>	
<b>Option Sets Variable</b>	Yes, <code>log_bin_use_v1_row_events</code> [2015]	
<b>Variable Name</b>	<code>log-bin-use-v1-row-events</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>0</code>

Version 2 binary log row events are available beginning with MySQL 5.6.6; however, Version 2 events cannot be read by previous MySQL Server releases. Setting this option to 1 causes `mysqld` to write the binary log using Version 1 logging events, which is the only version of binary log events used in previous releases, and thus produce binary logs that can be read by older slaves. Setting `--log-bin-use-v1-row-events` to 0 (the default) causes `mysqld` to use Version 2 binary log events.

The value used for this option can be obtained from the read-only `log_bin_use_v1_row_events` [2015] system variable.

- `--log-short-format` [428]

<b>Command-Line Format</b>	<code>--log-short-format</code>	
<b>Option-File Format</b>	<code>log-short-format</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Log less information to the binary log and slow query log, if they have been activated.

**Statement selection options.** The options in the following list affect which statements are written to the binary log, and thus sent by a replication master server to its slaves. There are also options for slave servers that control which statements received from the master should be executed or ignored. For details, see [Section 16.1.4.3, “Replication Slave Options and Variables”](#).

- `--binlog-do-db=db_name` [2004]

<b>Command-Line Format</b>	<code>--binlog-do-db=name</code>	
<b>Option-File Format</b>	<code>binlog-do-db</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This option affects binary logging in a manner similar to the way that `--replicate-do-db` [1976] affects replication.

The effects of this option depend on whether the statement-based or row-based logging format is in use, in the same way that the effects of `--replicate-do-db` [1976] depend on whether statement-based or row-based replication is in use. You should keep in mind that the format used to log a given statement may not necessarily be the same as that indicated by the value of `binlog_format` [2010]. For example, DDL statements such as `CREATE TABLE` and `ALTER TABLE` are always logged as statements, without regard to the logging format in effect, so the following statement-based rules for `--binlog-do-db` always apply in determining whether or not the statement is logged.

**Statement-based logging.** Only those statements are written to the binary log where the default database (that is, the one selected by `USE`) is `db_name`. To specify more than one database, use this option multiple times, once for each database; however, doing so does *not* cause cross-database statements such as `UPDATE some_db.some_table SET foo='bar'` to be logged while a different database (or no database) is selected.



### Warning

To specify multiple databases you *must* use multiple instances of this option. Because database names can contain commas, the list will be treated as the name of a single database if you supply a comma-separated list.

An example of what does not work as you might expect when using statement-based logging: If the server is started with `--binlog-do-db=sales` [2004] and you issue the following statements, the `UPDATE` statement is *not* logged:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

The main reason for this “just check the default database” behavior is that it is difficult from the statement alone to know whether it should be replicated (for example, if you are using multiple-table `DELETE` statements or multiple-table `UPDATE` statements that act across multiple databases). It is also faster to check only the default database rather than all databases if there is no need.

Another case which may not be self-evident occurs when a given database is replicated even though it was not specified when setting the option. If the server is started with `--binlog-do-db=sales`, the following `UPDATE` statement is logged even though `prices` was not included when setting `--binlog-do-db`:

```
USE sales;
UPDATE prices.discounts SET percentage = percentage + 10;
```

Because `sales` is the default database when the `UPDATE` statement is issued, the `UPDATE` is logged.

**Row-based logging.** Logging is restricted to database `db_name`. Only changes to tables belonging to `db_name` are logged; the default database has no effect on this. Suppose that the server is started with `--binlog-do-db=sales` [2004] and row-based logging is in effect, and then the following statements are executed:

```
USE prices;
UPDATE sales.february SET amount=amount+100;
```

The changes to the `february` table in the `sales` database are logged in accordance with the `UPDATE` statement; this occurs whether or not the `USE` statement was issued. However, when using the row-based logging format and `--binlog-do-db=sales` [2004], changes made by the following `UPDATE` are not logged:

```
USE prices;
UPDATE prices.march SET amount=amount-25;
```

Even if the `USE prices` statement were changed to `USE sales`, the `UPDATE` statement's effects would still not be written to the binary log.

Another important difference in `--binlog-do-db` [2004] handling for statement-based logging as opposed to the row-based logging occurs with regard to statements that refer to multiple databases. Suppose that the server is started with `--binlog-do-db=db1` [2004], and the following statements are executed:

```
USE db1;
UPDATE db1.table1 SET col1 = 10, db2.table2 SET col2 = 20;
```

If you are using statement-based logging, the updates to both tables are written to the binary log. However, when using the row-based format, only the changes to `table1` are logged; `table2` is in a different database, so it is not changed by the `UPDATE`. Now suppose that, instead of the `USE db1` statement, a `USE db4` statement had been used:

```
USE db4;
UPDATE db1.table1 SET col1 = 10, db2.table2 SET col2 = 20;
```

In this case, the `UPDATE` statement is not written to the binary log when using statement-based logging. However, when using row-based logging, the change to `table1` is logged, but not that to `table2`—in other words, only changes to tables in the database named by `--binlog-do-db` [2004] are logged, and the choice of default database has no effect on this behavior.

- `--binlog-ignore-db=db_name` [2005]

<b>Command-Line Format</b>	<code>--binlog-ignore-db=name</code>	
<b>Option-File Format</b>	<code>binlog-ignore-db</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This option affects binary logging in a manner similar to the way that `--replicate-ignore-db` [1977] affects replication.

The effects of this option depend on whether the statement-based or row-based logging format is in use, in the same way that the effects of `--replicate-ignore-db` [1977] depend on whether statement-based or row-based replication is in use. You should keep in mind that the format used to log a given statement may not necessarily be the same as that indicated by the value of `binlog_format` [2010]. For example, DDL statements such as `CREATE TABLE` and `ALTER TABLE` are always logged as statements, without regard to the logging format in effect, so the following statement-based rules for `--binlog-ignore-db` always apply in determining whether or not the statement is logged.



**Statement-based logging.** Tells the server to not log any statement where the default database (that is, the one selected by `USE`) is `db_name`.

**Row-based format.** Tells the server not to log updates to any tables in the database `db_name`. The current database has no effect.

When using statement-based logging, the following example does not work as you might expect. Suppose that the server is started with `--binlog-ignore-db=sales` [2005] and you issue the following statements:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

The `UPDATE` statement is logged in such a case because `--binlog-ignore-db` [2005] applies only to the default database (determined by the `USE` statement). Because the `sales` database was specified explicitly in the statement, the statement has not been filtered. However, when using row-based logging, the `UPDATE` statement's effects are *not* written to the binary log, which means that no changes to the `sales.january` table are logged; in this instance, `--binlog-ignore-db=sales` [2005] causes *all* changes made to tables in the master's copy of the `sales` database to be ignored for purposes of binary logging.

To specify more than one database to ignore, use this option multiple times, once for each database. Because database names can contain commas, the list will be treated as the name of a single database if you supply a comma-separated list.

You should not use this option if you are using cross-database updates and you do not want these updates to be logged.

**Checksum options.** Beginning with MySQL 5.6.2, MySQL supports reading and writing of binary log checksums. These are enabled using the two options listed here:

- `--binlog-checksum={NONE|CRC32}` [2006]

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--binlog-checksum=type</code>	
<b>Option-File Format</b>	<code>binlog-checksum</code>	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>NONE</code>
	<b>Valid Values</b>	<code>NONE</code> <code>CRC32</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>CRC32</code>
	<b>Valid Values</b>	<code>NONE</code> <code>CRC32</code>

Enabling this option causes the master to write checksums for events written to the binary log. Set to `NONE` to disable, or the name of the algorithm to be used for generating checksums; currently, only `CRC32` checksums are supported. As of MySQL 5.6.6, `CRC32` is the default.

This option was added in MySQL 5.6.2.

- `--master-verify-checksum={0|1}` [2006]

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--master-verify-checksum=name</code>	
<b>Option-File Format</b>	<code>master-verify-checksum</code>	
<b>Option Sets Variable</b>	Yes, <code>master_verify_checksum</code> [2015]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Enabling this option causes the master to verify events from the binary log using checksums, and to stop with an error in the event of a mismatch. Disabled by default.

This option was added in MySQL 5.6.2.

To control reading of checksums by the slave (from the relay) log, use the `--slave-sql-verify-checksum` [1987] option.

**Testing and debugging options.** The following binary log options are used in replication testing and debugging. They are not intended for use in normal operations.

- `--max-binlog-dump-events=N` [2007]

<b>Command-Line Format</b>	<code>--max-binlog-dump-events=#</code>	
<b>Option-File Format</b>	<code>max-binlog-dump-events</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>0</code>

This option is used internally by the MySQL test suite for replication testing and debugging.

- `--sporadic-binlog-dump-fail` [2007]

<b>Command-Line Format</b>	<code>--sporadic-binlog-dump-fail</code>	
<b>Option-File Format</b>	<code>sporadic-binlog-dump-fail</code>	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

This option is used internally by the MySQL test suite for replication testing and debugging.

- `--binlog-rows-query-log-events` [2007]

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--binlog-rows-query-log-events</code>	
<b>Option-File Format</b>	<code>binlog-rows-query-log-events</code>	
<b>Option Sets Variable</b>	Yes, <code>binlog_rows_query_log_events</code> [2014]	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

Added in MySQL 5.6.2, this option enables `binlog_rows_query_log_events` [2014]. Must be set to `OFF` (the default) when generating logs for a MySQL 5.6.1 or earlier slave server or version of `mysqlbinlog`.

**System variables used with the binary log.** The following list describes system variables for controlling binary logging. They can be set at server startup and some of them can be changed at runtime using `SET`. Server options used to control binary logging are listed earlier in this section.

- `log_bin` [2008]

<b>Variable Name</b>	<code>log_bin</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

Whether the binary log is enabled. If the `--log-bin` [2002] option is used, then the value of this variable is `ON`; otherwise it is `OFF`. This variable reports only on the status of binary logging (enabled or disabled); it does not actually report the value to which `--log-bin` [2002] is set.

See [Section 5.2.4, “The Binary Log”](#).

- `log_slave_updates` [2008]

<b>Command-Line Format</b>	<code>--log-slave-updates</code>
<b>Option-File Format</b>	<code>log_slave_updates</code>
<b>Option Sets Variable</b>	Yes, <code>log_slave_updates</code> [2008]
<b>Variable Name</b>	<code>log_slave_updates</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>
	<b>Type</b> <code>boolean</code>
	<b>Default</b> <code>FALSE</code>

Whether updates received by a slave server from a master server should be logged to the slave's own binary log. Binary logging must be enabled on the slave for this variable to have any effect. See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).

- `binlog_cache_size` [2008]

<b>Command-Line Format</b>	<code>--binlog_cache_size=#</code>
<b>Option-File Format</b>	<code>binlog_cache_size</code>
<b>Option Sets Variable</b>	Yes, <code>binlog_cache_size</code> [2008]
<b>Variable Name</b>	<code>binlog_cache_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes
	<b>Permitted Values</b>
	<b>Platform Bit Size</b> <code>32</code>
	<b>Type</b> <code>numeric</code>
	<b>Default</b> <code>32768</code>
	<b>Range</b> <code>4096 .. 4294967295</code>

	<b>Permitted Values</b>	
<b>Platform Bit Size</b>	64	
<b>Type</b>	numeric	
<b>Default</b>	32768	
<b>Range</b>	4096 .. 18446744073709547520	

The size of the cache to hold changes to the binary log during a transaction. A binary log cache is allocated for each client if the server supports any transactional storage engines and if the server has the binary log enabled (`--log-bin` [2002] option). If you often use large transactions, you can increase this cache size to get better performance. The `Binlog_cache_use` [586] and `Binlog_cache_disk_use` [586] status variables can be useful for tuning the size of this variable. See Section 5.2.4, “The Binary Log”.

`binlog_cache_size` sets the size for the transaction cache only; the size of the statement cache is governed by the `binlog_stmt_cache_size` [2014] system variable.

- `binlog_checksum` [2009]

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	<code>binlog_checksum</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>NONE</code>
	<b>Valid Values</b>	<code>NONE</code> <code>CRC32</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>string</code>
	<b>Default</b>	<code>CRC32</code>
	<b>Valid Values</b>	<code>NONE</code> <code>CRC32</code>

When enabled, this variable causes the master to write a checksum for each event in the binary log. `binlog_checksum` supports the values `NONE` (disabled) and `CRC32`. The default is `CRC32` as of MySQL 5.6.6, `NONE` before that.

When `binlog_checksum` is disabled (value `NONE`), the server verifies that it is writing only complete events to the binary log by writing and checking the event length (rather than a checksum) for each event.

Changing the value of this variable causes the binary log to be rotated; checksums are always written to an entire binary log file, and never to only part of one.

This variable was added in MySQL 5.6.2.

In MySQL 5.6.6 and later, setting this variable on the master to a value unrecognized by the slave causes the slave to set its own `binlog_checksum` value to `NONE`, and to stop replication with an error. (Bug #13553750, Bug #61096) If backward compatibility with older slaves is a concern, you may want to set the value explicitly to `NONE`.

- [binlog\\_direct\\_non\\_transactional\\_updates \[2010\]](#)

<b>Command-Line Format</b>	<code>--binlog_direct_non_transactional_updates[=value]</code>	
<b>Option-File Format</b>	<code>binlog_direct_non_transactional_updates</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">binlog_direct_non_transactional_updates [2010]</a>	
<b>Variable Name</b>	<code>binlog_direct_non_transactional_updates</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>OFF</code>

Due to concurrency issues, a slave can become inconsistent when a transaction contains updates to both transactional and nontransactional tables. MySQL tries to preserve causality among these statements by writing nontransactional statements to the transaction cache, which is flushed upon commit. However, problems arise when modifications done to nontransactional tables on behalf of a transaction become immediately visible to other connections because these changes may not be written immediately into the binary log.

The [binlog\\_direct\\_non\\_transactional\\_updates \[2010\]](#) variable offers one possible workaround to this issue. By default, this variable is disabled. Enabling [binlog\\_direct\\_non\\_transactional\\_updates \[2010\]](#) causes updates to nontransactional tables to be written directly to the binary log, rather than to the transaction cache.

*[binlog\\_direct\\_non\\_transactional\\_updates \[2010\]](#) works only for statements that are replicated using the statement-based binary logging format; that is, it works only when the value of [binlog\\_format \[2010\]](#) is `STATEMENT`, or when [binlog\\_format \[2010\]](#) is `MIXED` and a given statement is being replicated using the statement-based format. This variable has no effect when the binary log format is `ROW`, or when [binlog\\_format \[2010\]](#) is set to `MIXED` and a given statement is replicated using the row-based format.*



### Important

Before enabling this variable, you must make certain that there are no dependencies between transactional and nontransactional tables; an example of such a dependency would be the statement `INSERT INTO myisam_table SELECT * FROM innodb_table`. Otherwise, such statements are likely to cause the slave to diverge from the master.

In MySQL 5.6, this variable has no effect when the binary log format is `ROW` or `MIXED`. (Bug #51291)

- [binlog\\_format \[2010\]](#)

<b>Command-Line Format</b>	<code>--binlog-format=format</code>	
<b>Option-File Format</b>	<code>binlog-format=format</code>	
<b>Option Sets Variable</b>	Yes, <a href="#">binlog_format [2010]</a>	
<b>Variable Name</b>	<code>binlog_format</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>STATEMENT</code>

	<b>Valid Values</b>	ROW STATEMENT MIXED
--	---------------------	---------------------------

This variable sets the binary logging format, and can be any one of `STATEMENT`, `ROW`, or `MIXED`. See [Section 16.1.2, “Replication Formats”](#). `binlog_format [2010]` is set by the `--binlog-format [415]` option at startup, or by the `binlog_format [2010]` variable at runtime.



#### Note

While you can change the logging format at runtime, it is *not* recommended that you change it while replication is ongoing. This is due in part to the fact that slaves do not honor the master's `binlog_format [2010]` setting; a given MySQL Server can change only its own logging format.

In MySQL 5.6, the default format is `STATEMENT`.

You must have the `SUPER [756]` privilege to set either the global or session `binlog_format [2010]` value.

The rules governing when changes to this variable take effect and how long the effect lasts are the same as for other MySQL server system variables. See [Section 13.7.4, “SET Syntax”](#), for more information.

When `MIXED` is specified, statement-based replication is used, except for cases where only row-based replication is guaranteed to lead to proper results. For example, this happens when statements contain user-defined functions (UDF) or the `UUID() [1280]` function. An exception to this rule is that `MIXED` always uses statement-based replication for stored functions and triggers.

There are exceptions when you cannot switch the replication format at runtime:

- From within a stored function or a trigger.
- If the session is currently in row-based replication mode and has open temporary tables.
- From within a transaction.

Trying to switch the format in those cases results in an error.

The binary log format affects the behavior of the following server options:

- `--replicate-do-db [1976]`
- `--replicate-ignore-db [1977]`
- `--binlog-do-db [2004]`
- `--binlog-ignore-db [2005]`

These effects are discussed in detail in the descriptions of the individual options.

- `binlog_max_flush_queue_time [2011]`

<b>Version Introduced</b>	5.6.6
<b>Variable Name</b>	<code>binlog_max_flush_queue_time</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	Yes

	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	0	
<b>Range</b>	0 .. 100000	

How long in microseconds to keep reading transactions from the flush queue before proceeding with the group commit (and syncing the log to disk, if `sync_binlog` [2017] is greater than 0). If the value is 0 (the default), there is no timeout and the server keeps reading new transactions until the queue is empty.

Normally, `binlog_max_flush_queue_time` [2011] can remain set to 0. If the server processes a large number of connections (for example, 100 or more) and many short transactions with low-latency requirements, it may be useful to set the value larger than 0 to force more frequent flushes to disk.

This variable was added in MySQL 5.6.6.

- `binlog_order_commits` [2012]

<b>Version Introduced</b>	5.6.6	
<b>Variable Name</b>	<code>binlog_order_commits</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	boolean	
<b>Default</b>	ON	

If this variable is enabled (the default), transactions are committed in the same order they are written to the binary log. If disabled, transactions may be committed in parallel. In some cases, disabling this variable might produce a performance increment.

This variable was added in MySQL 5.6.6.

- `binlog_row_image` [2012]

<b>Version Introduced</b>	5.6.2	
<b>Command-Line Format</b>	<code>--binlog-row-image=image_type</code>	
<b>Variable Name</b>	<code>binlog_row_image=image_type</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
<b>Type</b>	enumeration	
<b>Default</b>	full	
<b>Valid Values</b>	full	Log all columns
	minimal	Log only changed columns, and columns needed to identify rows
	noblob	Log all columns, except for unneeded BLOB and TEXT columns



In MySQL row-based replication, each row change event contains two images, a “before” image whose columns are matched against when searching for the row to be updated, and an “after” image containing the changes. Normally, MySQL logs full rows (that is, all columns) for both the before and after images. However, it is not strictly necessary to include every column in both images, and we can often save disk, memory, and network usage by logging only those columns which are actually required.

**Note**

When deleting a row, only the before image is logged, since there are no changed values to propagate following the deletion. When inserting a row, only the after image is logged, since there is no existing row to be matched. Only when updating a row are both the before and after images required, and both written to the binary log.

For the before image, it is necessary only that the minimum set of columns required to uniquely identify rows is logged. If the table containing the row has a primary key, then only the primary key column or columns are written to the binary log. Otherwise, if the table has a unique key all of whose columns are `NOT NULL`, then only the columns in the unique key need be logged. (If the table has neither a primary key nor a unique key without any `NULL` columns, then all columns must be used in the before image, and logged.) In the after image, it is necessary to log only the columns which have actually changed.

In MySQL 5.6, you can cause the server to log full or minimal rows using the `binlog_row_image` system variable. This variable actually takes one of three possible values, as shown in the following list:

- `full`: Log all columns in both the before image and the after image.
- `minimal`: Log only those columns in the before image that are required to identify the row to be changed; log only those columns in the after image that are actually changed.
- `noblob`: Log all columns (same as `full`), except for `BLOB` and `TEXT` columns that are not required to identify rows, or that have not changed.

The default value is `full`. In MySQL 5.5 and earlier, full row images are always used for both before images and after images. If you need to replicate from a MySQL 5.6 (or later) master to a slave running a previous version of MySQL, the master should always use this value.

When using `minimal` or `noblob`, deletes and updates are guaranteed to work correctly for a given table if and only if the following conditions are true for both the source and destination tables:

- All columns must be present and in the same order; each column must use the same data type as its counterpart in the other table.
- The tables must have identical primary key definitions.

(In other words, the tables must be identical with the possible exception of indexes that are not part of the tables' primary keys.)

If these conditions are not met, it is possible that the primary key column values in the destination table may prove insufficient to provide a unique match for a delete or update. In this event, no warning or error is issued; the master and slave silently diverge, thus breaking consistency.

Setting this variable has no effect when the binary logging format is `STATEMENT`. When `binlog_format` [2010] is `MIXED`, the setting for `binlog_row_image` is applied to changes that are logged using row-based format, but this setting no effect on changes logged as statements.

Setting `binlog_row_image` on either the global or session level does not cause an implicit commit; this means that this variable can be changed while a transaction is in progress without affecting the transaction.

- `binlog_rows_query_log_events`

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	<code>binlog_rows_query_log_events</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>FALSE</code>

The `binlog_rows_query_log_events` [2014] system variable affects row-based logging only. When enabled, it causes a MySQL 5.6.2 or later server to write informational log events such as row query log events into its binary log. This information can be used for debugging and related purposes; such as obtaining the original query issued on the master when it cannot be reconstructed from the row updates.

These events are normally ignored by MySQL 5.6.2 and later programs reading the binary log and so cause no issues when replicating or restoring from backup. This is not true for a `mysqld` or `mysqlbinlog` from MySQL 5.6.1 or earlier: When the older version of the program reading the log encounters an informational log event, it fails, and stops reading at that point. To make the binary log readable by slave replication MySQL servers and other readers (for example, `mysqlbinlog`) from a MySQL 5.6.1 or earlier distribution, `binlog_rows_query_log_events` [2014] must be disabled during logging.

- `binlog_stmt_cache_size` [2014]

<b>Version Introduced</b>	5.6.1	
<b>Command-Line Format</b>	<code>--binlog_stmt_cache_size=#</code>	
<b>Option-File Format</b>	<code>binlog_stmt_cache_size</code>	
<b>Option Sets Variable</b>	Yes, <code>binlog_stmt_cache_size</code> [2014]	
<b>Variable Name</b>	<code>binlog_stmt_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	32
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	32768
	<b>Range</b>	4096 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	32768
	<b>Range</b>	4096 .. 18446744073709547520

This variable determines the size of the cache for the binary log to hold nontransactional statements issued during a transaction. Separate binary log transaction and statement caches are allocated for each client if the server supports any transactional storage engines and if the server has the binary log enabled (`--log-bin` [2002] option). If you often use large nontransactional statements during transactions, you can increase this cache size to get better performance. The `Binlog_stmt_cache_use` [586] and `Binlog_stmt_cache_disk_use` [586] status variables can be useful for tuning the size of this variable. See [Section 5.2.4, “The Binary Log”](#).

The `binlog_cache_size` [2008] system variable sets the size for the transaction cache.

- `log_bin_basename` [2015]

<b>Version Introduced</b>	5.6.2	
<b>Variable Name</b>	<code>log_bin_basename</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>file name</code>
	<b>Default</b>	<code>datadir + '/' + hostname + '-bin'</code>

Holds the name and complete path to the binary log file. Unlike the `log_bin` [2008] system variable, `log_bin_basename` [2015] reflects the name set with the `--log-bin` [2002] server option.

The `log_bin_basename` [2015] system variable was added in MySQL 5.6.2.

- `log_bin_use_v1_row_events` [2015]

<b>Version Introduced</b>	5.6.6	
<b>Command-Line Format</b>	<code>--log-bin-use-v1-row-events[={0 1}]</code>	
<b>Option-File Format</b>	<code>log_bin_use_v1_row_events</code>	
<b>Option Sets Variable</b>	Yes, <code>log_bin_use_v1_row_events</code> [2015]	
<b>Variable Name</b>	<code>log_bin_use_v1_row_events</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>0</code>

Shows whether Version 2 binary logging, available beginning with MySQL 5.6.6, is in use. A value of 1 shows that the server is writing the binary log using Version 1 logging events (the only version of binary log events used in MySQL 5.6.5 and previous MySQL Server releases), and thus producing a binary log that can be read by older slaves. 0 indicates that Version 2 binary log events are in use.

This variable is read-only. To switch between Version 1 and Version 2 binary event binary logging, it is necessary to restart `mysqld` with the `--log-bin-use-v1-row-events` [2003] option.

- `master_verify_checksum` [2015]

<b>Version Introduced</b>	5.6.2
<b>Variable Name</b>	<code>master_verify_checksum</code>
<b>Variable Scope</b>	Global

<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	OFF

Enabling this variable causes the master to examine checksums when reading from the binary log. `master_verify_checksum` is disabled by default; in this case, the master uses the event length from the binary log to verify events, so that only complete events are read from the binary log.

This variable was added in MySQL 5.6.2.

- `max_binlog_cache_size` [2016]

<b>Command-Line Format</b>	<code>--max_binlog_cache_size=#</code>	
<b>Option-File Format</b>	<code>max_binlog_cache_size</code>	
<b>Option Sets Variable</b>	Yes, <code>max_binlog_cache_size</code> [2016]	
<b>Variable Name</b>	<code>max_binlog_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	18446744073709547520
	<b>Range</b>	4096 .. 18446744073709547520

If a transaction requires more than this many bytes of memory, the server generates a `Multi-statement transaction required more than 'max_binlog_cache_size' bytes of storage` error. The minimum value is 4096. The maximum possible value is 16EB (exabytes). The maximum recommended value is 4GB; this is due to the fact that MySQL currently cannot work with binary log positions greater than 4GB.



#### Note

Prior to MySQL 5.6.7, 64-bit Windows platforms truncated the stored value for this variable to 4G, even when it was set to a greater value (Bug #13961678).

`max_binlog_cache_size` sets the size for the transaction cache only; the upper limit for the statement cache is governed by the `max_binlog_stmt_cache_size` [2017] system variable.

In MySQL 5.6, the visibility to sessions of `max_binlog_cache_size` matches that of the `binlog_cache_size` [2008] system variable; in other words, changing its value effects only new sessions that are started after the value is changed.

- `max_binlog_size` [2016]

<b>Command-Line Format</b>	<code>--max_binlog_size=#</code>	
<b>Option-File Format</b>	<code>max_binlog_size</code>	
<b>Option Sets Variable</b>	Yes, <code>max_binlog_size</code> [2016]	
<b>Variable Name</b>	<code>max_binlog_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric

	<b>Default</b>	1073741824
	<b>Range</b>	4096 .. 1073741824

If a write to the binary log causes the current log file size to exceed the value of this variable, the server rotates the binary logs (closes the current file and opens the next one). The minimum value is 4096 bytes. The maximum and default value is 1GB.

A transaction is written in one chunk to the binary log, so it is never split between several binary logs. Therefore, if you have big transactions, you might see binary log files larger than `max_binlog_size` [2016].

If `max_relay_log_size` [507] is 0, the value of `max_binlog_size` [2016] applies to relay logs as well.

- `max_binlog_stmt_cache_size` [2017]

<b>Version Introduced</b>	5.6.1	
<b>Command-Line Format</b>	<code>--max_binlog_stmt_cache_size=#</code>	
<b>Option-File Format</b>	<code>max_binlog_stmt_cache_size</code>	
<b>Option Sets Variable</b>	Yes, <code>max_binlog_stmt_cache_size</code> [2017]	
<b>Variable Name</b>	<code>max_binlog_stmt_cache_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	18446744073709547520
	<b>Range</b>	4096 .. 18446744073709547520

If nontransactional statements within a transaction require more than this many bytes of memory, the server generates an error. The minimum value is 4096. The maximum and default values are 4GB on 32-bit platforms and 16EB (exabytes) on 64-bit platforms.



#### Note

Prior to MySQL 5.6.7, 64-bit Windows platforms truncated the stored value for this variable to 4G, even when it was set to a greater value (Bug #13961678).

`max_binlog_stmt_cache_size` sets the size for the statement cache only; the upper limit for the transaction cache is governed exclusively by the `max_binlog_cache_size` [2016] system variable.

- `sync_binlog` [2017]

<b>Command-Line Format</b>	<code>--sync-binlog=#</code>	
<b>Option-File Format</b>	<code>sync_binlog</code>	
<b>Option Sets Variable</b>	Yes, <code>sync_binlog</code> [2017]	
<b>Variable Name</b>	<code>sync_binlog</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Platform</b>	32
	<b>Bit Size</b>	2017

	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 4294967295
	<b>Permitted Values</b>	
	<b>Platform Bit Size</b>	64
	<b>Type</b>	numeric
	<b>Default</b>	0
	<b>Range</b>	0 .. 18446744073709547520

If the value of this variable is greater than 0, the MySQL server synchronizes its binary log to disk (using `fdatasync()`) after `sync_binlog` [2017] commit groups are written to the binary log. The default value of `sync_binlog` [2017] is 0, which does no synchronizing to disk—in this case, the server relies on the operating system to flush the binary log's contents from time to time as for any other file. A value of 1 is the safest choice because in the event of a crash you lose at most one commit group from the binary log. However, it is also the slowest choice (unless the disk has a battery-backed cache, which makes synchronization very fast).

#### 16.1.4.5. Global Transaction ID Options and Variables

The MySQL Server options and system variables described in this section are used to monitor and control Global Transaction Identifiers (GTIDs), introduced in MySQL 5.6.5.



#### Note

Many of these options and variables were renamed in MySQL 5.6.9. See their descriptions in this section for more information.

For additional information, see [Section 16.1.3, “Replication with Global Transaction Identifiers”](#).

**Startup options used in GTID replication.** The following server startup options are used with GTID-based replication:

- `--disable-gtid-unsafe-statements`

<b>Version Introduced</b>	5.6.5	
<b>Version Removed</b>	5.6.9	
<b>Command-Line Format</b>	<code>--disable-gtid-unsafe-statements[=value]</code>	
<b>Option-File Format</b>	<code>disable-gtid-unsafe-statements</code>	
<b>Variable Name</b>	<code>disable-gtid-unsafe-statements</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	boolean
	<b>Default</b>	false

*Obsolete:* Replaced by `--enforce-gtid-consistency` [2018] in MySQL 5.6.9. (Bug #14775984)

- `--enforce-gtid-consistency`

<b>Version Introduced</b>	5.6.9	
<b>Command-Line Format</b>	<code>--enforce-gtid-consistency[=value]</code>	

<b>Option-File Format</b>	<code>enforce-gtid-consistency</code>	
<b>Variable Name</b>	<code>enforce-gtid-consistency</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>false</code>

When set, this option allows execution of only those statements that can be logged in a transactionally safe manner. This means that the following operations *cannot* be used when this option is enabled:

- `CREATE TABLE ... SELECT` statements
- `CREATE TEMPORARY TABLE` statements inside transactions
- Transactions or statements that update both transactional and nontransactional tables.

Prior to MySQL 5.6.9, this option was named `--disable-gtid-unsafe-statements` [2018]. (Bug #14775984)

Prior to MySQL 5.6.7, using this option caused nontransactional DML on temporary tables to fail, although changes to temporary tables are not logged when using row-based binary logging. In MySQL 5.6.7 and later, nontransactional DML statements are allowed on temporary tables with `--disable-gtid-unsafe-statements` [2018] (`--enforce-gtid-consistency` beginning with MySQL 5.6.9) as long as all affected tables are temporary tables (Bug #14272672).

Prior to MySQL 5.6.7, `mysql_upgrade` could not be used with a MySQL Server running with this option enabled, unless `mysql_upgrade` was running with `--write-binlog` [244] explicitly disabled. (Bug #13833710, Bug #14221043) In MySQL 5.6.7 and later, it is possible but not recommended to run `mysql_upgrade` on a server where `--gtid-mode=ON` [2019], since the MySQL system tables use the `MyISAM` storage engine, which is nontransactional.

In MySQL 5.6.8 and earlier, you could not use any statements affecting nontransactional tables when `--enforce-gtid-consistency` was used (the option was then called `--disable-gtid-unsafe-statements` [2018]). In MySQL 5.6.9 and later, this option allows single statements updating nontransactional tables. This is intended chiefly for use with programs such as `mysql_install_db` and `mysql_upgrade`. (Bug #14722659)

- `--gtid-mode` [2019]

<b>Version Introduced</b>	5.6.5	
<b>Command-Line Format</b>	<code>--gtid-mode=MODE</code>	
<b>Option-File Format</b>	<code>gtid-mode</code>	
<b>Option Sets Variable</b>	Yes, <code>gtid_mode</code> [2022]	
<b>Variable Name</b>	<code>gtid-mode</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>OFF</code>
	<b>Valid Values</b>	<code>OFF</code>



	UPGRADE_STEP_1
	UPGRADE_STEP_2
	ON

This option specifies whether GTIDs are enabled. Starting the server with `--gtid-mode=ON` requires that the server also be started with the `--log-bin` [2002] and `--log-slave-updates` [1971] options as well. (In addition, you should also use `--enforce-gtid-consistency` [2018].)

Setting this option to `OFF` when there are GTIDs in the binary log or in the relay log, or to `ON` when there remain anonymous transactions to be executed, causes an error.



### Important

This option does not employ boolean values; its values are in fact enumerated. You should not attempt to use numeric values when setting this option, as these may lead to unexpected results. The values `UPGRADE_STEP_1` and `UPGRADE_STEP_2` are reserved for future use, but currently are not supported in production; if you use one of these two values with `--gtid-mode`, the server refuses to start.

Prior to MySQL 5.6.7, `mysql_upgrade` could not be used with a MySQL Server running with this option enabled, unless `mysql_upgrade` was running with `--write-binlog` [244] explicitly disabled. (Bug #13833710, Bug #14221043) In MySQL 5.6.7 and later, it is possible but not recommended to run `mysql_upgrade` on a server where `--gtid-mode=ON` [2019], since it may make changes to MySQL system tables that use the `MyISAM` storage engine, which is nontransactional.

Prior to MySQL 5.6.10, setting the global value for the `sql_slave_skip_counter` [1996] variable to 1 had no effect when `--gtid-mode` was set to `ON`. (Bug #15833516) A workaround in MySQL 5.6.9 and earlier versions is to reset the slave's position using `CHANGE MASTER TO ... MASTER_LOG_FILE = ... MASTER_LOG_POS = ...`, including the `MASTER_AUTO_POSITION = 0` option with this statement if needed.

**System variables used on replication masters.** The following system variables are used with GTID-based replication:

- `disable_gtid_unsafe_statements`

<b>Version Introduced</b>	5.6.5	
<b>Version Removed</b>	5.6.9	
<b>Command-Line Format</b>	<code>--disable-gtid-unsafe-statements[=value]</code>	
<b>Option-File Format</b>	<code>disable_gtid_unsafe_statements</code>	
<b>Variable Name</b>	<code>disable_gtid_unsafe_statements</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>boolean</code>
	<b>Default</b>	<code>false</code>

*Obsolete:* Replaced by `enforce_gtid_consistency` [2021] in MySQL 5.6.9. (Bug #14775984)

- `gtid_done` [2020]

<b>Version Introduced</b>	5.6.5	
<b>Version Removed</b>	5.6.9	
<b>Variable Name</b>	<a href="#">gtid_done</a>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">string</a>

*Obsolete:* replaced in MySQL 5.6.9 by [gtid\\_executed](#) [2021]. (Bug #14775984)

- [enforce\\_gtid\\_consistency](#)

<b>Version Introduced</b>	5.6.9	
<b>Command-Line Format</b>	<code>--enforce-gtid-consistency[=value]</code>	
<b>Option-File Format</b>	<a href="#">enforce_gtid_consistency</a>	
<b>Variable Name</b>	<a href="#">enforce_gtid_consistency</a>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">boolean</a>
	<b>Default</b>	<a href="#">false</a>

When this variable is true, execution is allowed of only those statements that can be logged in a transactionally safe manner, which means that the following operations cannot be used:

- `CREATE TABLE ... SELECT` statements
- `CREATE TEMPORARY TABLE` statements inside transactions
- Transactions or statements that update both transactional and nontransactional tables.

This variable is read-only. To set it, use the `--enforce-gtid-consistency` [2018] option on the command line or in an option file when starting the MySQL Server.

Prior to MySQL 5.6.9, this variable was named `disable_gtid_unsafe_statements` [2020]. (Bug #14775984)

- [gtid\\_executed](#) [2021]

<b>Version Introduced</b>	5.6.9	
<b>Variable Name</b>	<a href="#">gtid_executed</a>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<a href="#">string</a>

When used with global scope, this variable contains a representation of the set of all transactions that are logged in the binary log. When used with session scope, it contains a representation of the set of transactions that are written to the cache in the current session.

Issuing `RESET MASTER` causes the global value (but not the session value) of this variable to be reset to an empty string.

Prior to MySQL 5.6.9, this variable was known as `gtid_done` [2020].

- `gtid_lost` [2022]

<b>Version Introduced</b>	5.6.5	
<b>Version Removed</b>	5.6.9	
<b>Variable Name</b>	<code>gtid_lost</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

*Obsolete:* Replaced by `gtid_purged` [2023] in MySQL 5.6.9. (Bug #14775984)

- `gtid_mode` [2022]

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<code>gtid_mode</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>OFF</code>
	<b>Valid Values</b>	<code>OFF</code> <code>UPGRADE_STEP_1</code> <code>UPGRADE_STEP_2</code> <code>ON</code>

Shows whether GTIDs are enabled. Read-only; set using `--gtid-mode` [2019].

- `gtid_next` [2022]

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<code>gtid_next</code>	
<b>Variable Scope</b>	Session	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>enumeration</code>
	<b>Default</b>	<code>AUTOMATIC</code>
	<b>Valid Values</b>	<code>AUTOMATIC</code> <code>ANONYMOUS</code> <code>UUID:NUMBER</code>

This variable is used to specify whether and how the next GTID is obtained. `gtid_next` can take any of the following values:

- **AUTOMATIC**: Use the next automatically-generated global transaction ID.
- **ANONYMOUS**: Transactions do not have global identifiers, and are identified by file and position only.
- A global transaction ID in `UUID:NUMBER` format.

You must have the **SUPER** [756] privilege to set this variable. Setting this variable has no effect if `gtid_mode` [2022] is `OFF`.

Beginning with MySQL 5.6.11, you cannot execute any of the statements `CHANGE MASTER TO`, `START SLAVE`, `STOP SLAVE`, `REPAIR TABLE`, `OPTIMIZE TABLE`, `ANALYZE TABLE`, `CHECK TABLE`, `CREATE SERVER`, `ALTER SERVER`, `DROP SERVER`, `CACHE INDEX`, `LOAD INDEX INTO CACHE`, `FLUSH`, or `RESET` when `gtid_next` [2022] is set to any value other than `AUTOMATIC`; in such cases, the statement fails with an error. (BUG#16062608)

- `gtid_owned` [2023]

<b>Version Introduced</b>	5.6.5	
<b>Variable Name</b>	<code>gtid_owned</code>	
<b>Variable Scope</b>	Global, Session	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

This read-only variable holds a list whose contents depend on its scope. When used with session scope, the list holds all GTIDs that are owned by this client; when used with global scope, it holds a list of all GTIDs along with their owners.

- `gtid_purged` [2023]

<b>Version Introduced</b>	5.6.9	
<b>Variable Name</b>	<code>gtid_purged</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	Yes	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>string</code>

The set of all transactions that have been purged from the binary log.

Issuing `RESET MASTER` causes the value of this variable to be reset to an empty string.

Prior to MySQL 5.6.9, this variable was known as `gtid_lost` [2022], and was read-only. In MySQL 5.6.9 and later, it is possible to update the value of this variable, but only by adding GTIDs to those already listed, and only when `gtid_executed` [2021] is unset—that is, on a new server. (Bug #14797808)

## 16.1.5. Common Replication Administration Tasks

Once replication has been started it should execute without requiring much regular administration. Depending on your replication environment, you will want to check the replication status of each slave periodically, daily, or even more frequently.

### 16.1.5.1. Checking Replication Status

The most common task when managing a replication process is to ensure that replication is taking place and that there have been no errors between the slave and the master. The primary statement for this is `SHOW SLAVE STATUS`, which you must execute on each slave:

```
mysql> SHOW SLAVE STATUS\G
***** 1. row *****
      Slave_IO_State: Waiting for master to send event
      Master_Host: master1
      Master_User: root
      Master_Port: 3306
      Connect_Retry: 60
      Master_Log_File: mysql-bin.000004
      Read_Master_Log_Pos: 931
      Relay_Log_File: slave1-relay-bin.000056
      Relay_Log_Pos: 950
      Relay_Master_Log_File: mysql-bin.000004
      Slave_IO_Running: Yes
      Slave_SQL_Running: Yes
      Replicate_Do_DB:
      Replicate_Ignore_DB:
      Replicate_Do_Table:
      Replicate_Ignore_Table:
      Replicate_Wild_Do_Table:
      Replicate_Wild_Ignore_Table:
      Last_Errno: 0
      Last_Error:
      Skip_Counter: 0
      Exec_Master_Log_Pos: 931
      Relay_Log_Space: 1365
      Until_Condition: None
      Until_Log_File:
      Until_Log_Pos: 0
      Master_SSL_Allowed: No
      Master_SSL_CA_File:
      Master_SSL_CA_Path:
      Master_SSL_Cert:
      Master_SSL_Cipher:
      Master_SSL_Key:
      Seconds_Behind_Master: 0
Master_SSL_Verify_Server_Cert: No
      Last_IO_Errno: 0
      Last_IO_Error:
      Last_SQL_Errno: 0
      Last_SQL_Error:
      Replicate_Ignore_Server_Ids: 0
```

The key fields from the status report to examine are:

- `Slave_IO_State`: The current status of the slave. See [Section 8.12.5.6, “Replication Slave I/O Thread States”](#), and [Section 8.12.5.7, “Replication Slave SQL Thread States”](#), for more information.
- `Slave_IO_Running`: Whether the I/O thread for reading the master's binary log is running. Normally, you want this to be `Yes` unless you have not yet started replication or have explicitly stopped it with `STOP SLAVE`.
- `Slave_SQL_Running`: Whether the SQL thread for executing events in the relay log is running. As with the I/O thread, this should normally be `Yes`.
- `Last_IO_Error`, `Last_SQL_Error`: The last errors registered by the I/O and SQL threads when processing the relay log. Ideally these should be blank, indicating no errors.

- `Seconds_Behind_Master`: The number of seconds that the slave SQL thread is behind processing the master binary log. A high number (or an increasing one) can indicate that the slave is unable to handle events from the master in a timely fashion.

A value of 0 for `Seconds_Behind_Master` can usually be interpreted as meaning that the slave has caught up with the master, but there are some cases where this is not strictly true. For example, this can occur if the network connection between master and slave is broken but the slave I/O thread has not yet noticed this—that is, `slave_net_timeout` [1993] has not yet elapsed.

It is also possible that transient values for `Seconds_Behind_Master` may not reflect the situation accurately. When the slave SQL thread has caught up on I/O, `Seconds_Behind_Master` displays 0; but when the slave I/O thread is still queuing up a new event, `Seconds_Behind_Master` may show a large value until the SQL thread finishes executing the new event. This is especially likely when the events have old timestamps; in such cases, if you execute `SHOW SLAVE STATUS` several times in a relatively short period, you may see this value change back and forth repeatedly between 0 and a relatively large value.

Several pairs of fields provide information about the progress of the slave in reading events from the master binary log and processing them in the relay log:

- (`Master_Log_File`, `Read_Master_Log_Pos`): Coordinates in the master binary log indicating how far the slave I/O thread has read events from that log.
- (`Relay_Master_Log_File`, `Exec_Master_Log_Pos`): Coordinates in the master binary log indicating how far the slave SQL thread has executed events received from that log.
- (`Relay_Log_File`, `Relay_Log_Pos`): Coordinates in the slave relay log indicating how far the slave SQL thread has executed the relay log. These correspond to the preceding coordinates, but are expressed in slave relay log coordinates rather than master binary log coordinates.

On the master, you can check the status of connected slaves using `SHOW PROCESSLIST` to examine the list of running processes. Slave connections have `Binlog Dump` in the `Command` field:

```
mysql> SHOW PROCESSLIST \G;
***** 4. row *****
  Id: 10
  User: root
  Host: slave1:58371
  db: NULL
Command: Binlog Dump
  Time: 777
  State: Has sent all binlog to slave; waiting for binlog to be updated
  Info: NULL
```

Because it is the slave that drives the replication process, very little information is available in this report.

For slaves that were started with the `--report-host` [1980] option and are connected to the master, the `SHOW SLAVE HOSTS` statement on the master shows basic information about the slaves. The output includes the ID of the slave server, the value of the `--report-host` [1980] option, the connecting port, and master ID:

```
mysql> SHOW SLAVE HOSTS;
+-----+-----+-----+-----+-----+
| Server_id | Host   | Port | Rpl_recovery_rank | Master_id |
+-----+-----+-----+-----+-----+
|          10 | slave1 | 3306 |          0         |          1 |
+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

### 16.1.5.2. Pausing Replication on the Slave

You can stop and start the replication of statements on the slave using the `STOP SLAVE` and `START SLAVE` statements.

To stop processing of the binary log from the master, use `STOP SLAVE`:

```
mysql> STOP SLAVE;
```

When replication is stopped, the slave I/O thread stops reading events from the master binary log and writing them to the relay log, and the SQL thread stops reading events from the relay log and executing them. You can pause the I/O or SQL thread individually by specifying the thread type:

```
mysql> STOP SLAVE IO_THREAD;
mysql> STOP SLAVE SQL_THREAD;
```

To start execution again, use the `START SLAVE` statement:

```
mysql> START SLAVE;
```

To start a particular thread, specify the thread type:

```
mysql> START SLAVE IO_THREAD;
mysql> START SLAVE SQL_THREAD;
```

For a slave that performs updates only by processing events from the master, stopping only the SQL thread can be useful if you want to perform a backup or other task. The I/O thread will continue to read events from the master but they are not executed. This makes it easier for the slave to catch up when you restart the SQL thread.

Stopping only the I/O thread enables the events in the relay log to be executed by the SQL thread up to the point where the relay log ends. This can be useful when you want to pause execution to catch up with events already received from the master, when you want to perform administration on the slave but also ensure that it has processed all updates to a specific point. This method can also be used to pause event receipt on the slave while you conduct administration on the master. Stopping the I/O thread but permitting the SQL thread to run helps ensure that there is not a massive backlog of events to be executed when replication is started again.

## 16.2. Replication Implementation

Replication is based on the master server keeping track of all changes to its databases (updates, deletes, and so on) in its binary log. The binary log serves as a written record of all events that modify database structure or content (data) from the moment the server was started. Typically, `SELECT` statements are not recorded because they modify neither database structure nor content.

Each slave that connects to the master requests a copy of the binary log. That is, it pulls the data from the master, rather than the master pushing the data to the slave. The slave also executes the events from the binary log that it receives. This has the effect of repeating the original changes just as they were made on the master. Tables are created or their structure modified, and data is inserted, deleted, and updated according to the changes that were originally made on the master.

Because each slave is independent, the replaying of the changes from the master's binary log occurs independently on each slave that is connected to the master. In addition, because each slave receives a copy of the binary log only by requesting it from the master, the slave is able to read and update the copy of the database at its own pace and can start and stop the replication process at will without affecting the ability to update to the latest database status on either the master or slave side.

For more information on the specifics of the replication implementation, see [Section 16.2.1, “Replication Implementation Details”](#).

Masters and slaves report their status in respect of the replication process regularly so that you can monitor them. See [Section 8.12.5, “Examining Thread Information”](#), for descriptions of all replicated-related states.

The master binary log is written to a local relay log on the slave before it is processed. The slave also records information about the current position with the master's binary log and the local relay log. See [Section 16.2.2, “Replication Relay and Status Logs”](#).



Database changes are filtered on the slave according to a set of rules that are applied according to the various configuration options and variables that control event evaluation. For details on how these rules are applied, see [Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#).

## 16.2.1. Replication Implementation Details

MySQL replication capabilities are implemented using three threads, one on the master server and two on the slave:

- **Binlog dump thread.** The master creates a thread to send the binary log contents to a slave when the slave connects. This thread can be identified in the output of `SHOW PROCESSLIST` on the master as the `Binlog Dump` thread.

The binlog dump thread acquires a lock on the master's binary log for reading each event that is to be sent to the slave. As soon as the event has been read, the lock is released, even before the event is sent to the slave.

- **Slave I/O thread.** When a `START SLAVE` statement is issued on a slave server, the slave creates an I/O thread, which connects to the master and asks it to send the updates recorded in its binary logs.

The slave I/O thread reads the updates that the master's `Binlog Dump` thread sends (see previous item) and copies them to local files that comprise the slave's relay log.

The state of this thread is shown as `Slave_IO_running` in the output of `SHOW SLAVE STATUS` or as `Slave_running [598]` in the output of `SHOW STATUS`.

- **Slave SQL thread.** The slave creates an SQL thread to read the relay log that is written by the slave I/O thread and execute the events contained therein.

In the preceding description, there are three threads per master/slave connection. A master that has multiple slaves creates one binlog dump thread for each currently connected slave, and each slave has its own I/O and SQL threads.

A slave uses two threads to separate reading updates from the master and executing them into independent tasks. Thus, the task of reading statements is not slowed down if statement execution is slow. For example, if the slave server has not been running for a while, its I/O thread can quickly fetch all the binary log contents from the master when the slave starts, even if the SQL thread lags far behind. If the slave stops before the SQL thread has executed all the fetched statements, the I/O thread has at least fetched everything so that a safe copy of the statements is stored locally in the slave's relay logs, ready for execution the next time that the slave starts. This enables the master server to purge its binary logs sooner because it no longer needs to wait for the slave to fetch their contents.

The `SHOW PROCESSLIST` statement provides information that tells you what is happening on the master and on the slave regarding replication. For information on master states, see [Section 8.12.5.5, “Replication Master Thread States”](#). For slave states, see [Section 8.12.5.6, “Replication Slave I/O Thread States”](#), and [Section 8.12.5.7, “Replication Slave SQL Thread States”](#).

The following example illustrates how the three threads show up in the output from `SHOW PROCESSLIST`.

On the master server, the output from `SHOW PROCESSLIST` looks like this:

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  Id: 2
  User: root
  Host: localhost:32931
  db: NULL
  Command: Binlog Dump
  Time: 94
```

```
State: Has sent all binlog to slave; waiting for binlog to
      be updated
Info: NULL
```

Here, thread 2 is a [Binlog Dump](#) replication thread that services a connected slave. The [State](#) information indicates that all outstanding updates have been sent to the slave and that the master is waiting for more updates to occur. If you see no [Binlog Dump](#) threads on a master server, this means that replication is not running; that is, no slaves are currently connected.

On a slave server, the output from `SHOW PROCESSLIST` looks like this:

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  Id: 10
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 11
  State: Waiting for master to send event
  Info: NULL
***** 2. row *****
  Id: 11
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 11
  State: Has read all relay log; waiting for the slave I/O
        thread to update it
  Info: NULL
```

The [State](#) information indicates that thread 10 is the I/O thread that is communicating with the master server, and thread 11 is the SQL thread that is processing the updates stored in the relay logs. At the time that `SHOW PROCESSLIST` was run, both threads were idle, waiting for further updates.

The value in the [Time](#) column can show how late the slave is compared to the master. See [Section B.13, “MySQL 5.6 FAQ: Replication”](#). If sufficient time elapses on the master side without activity on the [Binlog Dump](#) thread, the master determines that the slave is no longer connected. As for any other client connection, the timeouts for this depend on the values of `net_write_timeout` and `net_retry_count`; for more information about these, see [Section 5.1.4, “Server System Variables”](#).

The `SHOW SLAVE STATUS` statement provides additional information about replication processing on a slave server. See [Section 16.1.5.1, “Checking Replication Status”](#).

## 16.2.2. Replication Relay and Status Logs

During replication, a slave server creates several logs that hold the binary log events relayed from the master to the slave, and to record information about the current status and location within the relay log. There are three types of logs used in the process, listed here:

- The *relay log* consists of the events read from the binary log of the master and written by the slave I/O thread. Events in the relay log are executed on the slave as part of the SQL thread.
- The *master info log* contains status and current configuration information for the slave's connection to the master. This log holds information on the master host name, login credentials, and coordinates indicating how far the slave has read from the master's binary log.

Prior to MySQL 5.6, this log was always a file (`master.info`), but in MySQL 5.6 and later, this log can be written to the `mysql.slave_master_info` table instead of a file, by starting the slave with `--master-info-repository=TABLE [2000]`.

- The *relay log info log* holds status information about the execution point within the slave's relay log.

Prior to MySQL 5.6, this log was always a file (`relay-log.info`), but in MySQL 5.6 and later, this log can be written to the `mysql.slave_relay_log_info` table instead of a file by starting the slave with `--relay-log-info-repository=TABLE [2000]`.

Prior to MySQL 5.6.7, the `Master_id` column of the `slave_master_info` and `slave_relay_log_info` tables showed the slave's server ID instead of the master's server ID. (Bug #12334346)

In order for replication to be crash-safe when using tables for logging status and relay information, the tables must use a transactional storage engine, such as `InnoDB`. Beginning with MySQL 5.6.6, these tables are created using `InnoDB`. (Bug #13538891)



#### Note

In order to guarantee crash safety on the slave, you must also run the slave with `--relay-log-recovery [1975]` enabled.

Prior to MySQL 5.6.6, if `mysqld` was unable to initialize the replication logging tables, the slave refused to start. In MySQL 5.6.6 and later, a warning is given when this occurs, but the slave is allowed to continue starting. (Bug #13971348) This situation is most likely to occur when upgrading from a version of MySQL that does not support slave logging tables to one in which they are supported.

In MySQL 5.6.5 and earlier, the `slave_master_info` and `slave_relay_log_info` tables used `MyISAM` by default, which meant that it was necessary before starting replication to change the storage engine used by these tables by issuing `ALTER TABLE ... ENGINE=InnoDB`, as shown here:

```
ALTER TABLE mysql.slave_master_info ENGINE=InnoDB;
ALTER TABLE mysql.slave_relay_log_info ENGINE=InnoDB;
```

The `ALTER TABLE` statements must be executed by the MySQL `root` or other user account with the appropriate privileges on the `mysql` database. You should not attempt to do this while replication is running; beginning with MySQL 5.6.3, trying to execute an `ALTER TABLE` on either these tables while replication is ongoing is disallowed. Starting with MySQL 5.6.4, execution of any statement requiring a write lock on either or both of these tables is disallowed while replication is ongoing, while statements that perform only reads are permitted at any time.



#### Important

Do not attempt to update or insert rows in the `slave_master_info` or `slave_relay_log_info` table manually. Doing so can cause undefined behavior, and is not supported.

Prior to MySQL 5.6.4, `mysqldump` did not dump the replication log tables unless they were specified by name and the `--master-data [288]` option was used.

### 16.2.2.1. The Slave Relay Log

The relay log, like the binary log, consists of a set of numbered files containing events that describe database changes, and an index file that contains the names of all used relay log files.

The term “relay log file” generally denotes an individual numbered file containing database events. The term “relay log” collectively denotes the set of numbered relay log files plus the index file.

Relay log files have the same format as binary log files and can be read using `mysqlbinlog` (see [Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#)).

By default, relay log file names have the form `host_name-relay-bin.nnnnnn` in the data directory, where `host_name` is the name of the slave server host and `nnnnnn` is a sequence number. Successive relay log files are created using successive sequence numbers, beginning with `000001`. The slave uses an index file to track the relay log files currently in use. The default relay log index file name is `host_name-relay-bin.index` in the data directory.

The default relay log file and relay log index file names can be overridden with, respectively, the `--relay-log` [1974] and `--relay-log-index` [1974] server options (see Section 16.1.4, “Replication and Binary Logging Options and Variables”).

If a slave uses the default host-based relay log file names, changing a slave's host name after replication has been set up can cause replication to fail with the errors `Failed to open the relay log` and `Could not find target log during relay log initialization`. This is a known issue (see Bug #2122). If you anticipate that a slave's host name might change in the future (for example, if networking is set up on the slave such that its host name can be modified using DHCP), you can avoid this issue entirely by using the `--relay-log` [1974] and `--relay-log-index` [1974] options to specify relay log file names explicitly when you initially set up the slave. This will make the names independent of server host name changes.

If you encounter the issue after replication has already begun, one way to work around it is to stop the slave server, prepend the contents of the old relay log index file to the new one, and then restart the slave. On a Unix system, this can be done as shown here:

```
shell> cat new_relay_log_name.index >> old_relay_log_name.index
shell> mv old_relay_log_name.index new_relay_log_name.index
```

A slave server creates a new relay log file under the following conditions:

- Each time the I/O thread starts.
- When the logs are flushed; for example, with `FLUSH LOGS` or `mysqladmin flush-logs`.
- When the size of the current relay log file becomes “too large,” determined as follows:
  - If the value of `max_relay_log_size` [507] is greater than 0, that is the maximum relay log file size.
  - If the value of `max_relay_log_size` [507] is 0, `max_binlog_size` [2016] determines the maximum relay log file size.

The SQL thread automatically deletes each relay log file as soon as it has executed all events in the file and no longer needs it. There is no explicit mechanism for deleting relay logs because the SQL thread takes care of doing so. However, `FLUSH LOGS` rotates relay logs, which influences when the SQL thread deletes them.

### 16.2.2.2. Slave Status Logs

A replication slave server creates two logs. By default, these logs are files named `master.info` and `relay-log.info` and created in the data directory. The names and locations of these files can be changed by using the `--master-info-file` [1973] and `--relay-log-info-file` [1975] options, respectively. In MySQL 5.6 and later, either or both of these logs can also be written to tables in the `mysql` database by starting the server with the appropriate option: use `--master-info-repository` [2000] to have the master info log written to the `mysql.slave_master_info` table, and use `--relay-log-info-repository` [2000] to have the relay log info log written to the `mysql.slave_relay_log_info` table. See Section 16.1.4, “Replication and Binary Logging Options and Variables”.

The two status logs contain information like that shown in the output of the `SHOW SLAVE STATUS` statement, which is discussed in Section 13.4.2, “SQL Statements for Controlling Slave Servers”. Because the status logs are stored on disk, they survive a slave server's shutdown. The next time the slave starts up, it reads the two logs to determine how far it has proceeded in reading binary logs from the master and in processing its own relay logs.

The master info log file or table should be protected because it contains the password for connecting to the master. See Section 6.1.2.3, “Passwords and Logging”.

The slave I/O thread updates the master info log. The following table shows the correspondence between the lines in the `master.info` file, the columns in the `mysql.slave_master_info` table, and the columns displayed by `SHOW SLAVE STATUS`.

Line in <code>master.info</code> File	<code>slave_master_info</code> Table Column	<code>SHOW SLAVE STATUS</code> Column	Description
1	<code>Number_of_lines</code>	[None]	Number of lines in the file
2	<code>Master_log_name</code>	<code>Master_Log_File</code>	The name of the master binary log currently being read from the master
3	<code>Master_log_pos</code>	<code>Read_Master_Log_Pos</code>	The current position within the master binary log that have been read from the master
4	<code>Host</code>	<code>Master_Host</code>	The host name of the master
5	<code>User</code>	<code>Master_User</code>	The user name used to connect to the master
6	<code>User_password</code>	Password (not shown by <code>SHOW SLAVE STATUS</code> )	The password used to connect to the master
7	<code>Port</code>	<code>Master_Port</code>	The network port used to connect to the master
8	<code>Connect_retry</code>	<code>Connect_Retry</code>	The period (in seconds) that the slave will wait before trying to reconnect to the master
9	<code>Enabled_ssl</code>	<code>Master_SSL_Allowed</code>	Indicates whether the server supports SSL connections
10	<code>Ssl_ca</code>	<code>Master_SSL_CA_File</code>	The file used for the Certificate Authority (CA) certificate
11	<code>Ssl_capath</code>	<code>Master_SSL_CA_Path</code>	The path to the Certificate Authority (CA) certificates

Line in <code>master.info</code> File	<code>slave_master_info</code> Table Column	<code>SHOW SLAVE STATUS</code> Column	Description
12	<code>Ssl_cert</code>	<code>Master_SSL_Cert</code>	The name of the SSL certificate file
13	<code>Ssl_cipher</code>	<code>Master_SSL_Cipher</code>	The list of possible ciphers used in the handshake for the SSL connection
14	<code>Ssl_key</code>	<code>Master_SSL_Key</code>	The name of the SSL key file
15	<code>Ssl_verify_server_cert</code>	<code>Master_SSL_Verify_Server_Cert</code>	Whether to verify the server certificate
16	<code>Heartbeat</code>	[None]	Interval between replication heartbeats, in seconds
17	<code>Bind</code>	<code>Master_Bind</code>	Which of the slave's network interfaces should be used for connecting to the master
18	<code>Ignored_server_ids</code>	<code>Replicate_Ignore_Server_Ids</code>	The number of server IDs to be ignored, followed by the actual server IDs
19	<code>Uuid</code>	<code>Master_UUID</code>	The master's unique ID
20	<code>Retry_count</code>	<code>Master_Retry_Count</code>	Maximum number of reconnection attempts permitted (Added in MySQL 5.6.1)


**Note**

Prior to MySQL 5.6.3, the name of the `Ssl_verify_server_cert` column was `Ssl_verify_servert_cert`. (Bug #12407446, Bug #60988)

The slave SQL thread updates the relay log info log. In MySQL 5.6, the `relay-log.info` file includes a line count and a replication delay value. The following table shows the correspondence between the lines in the `relay-log.info` file, the columns in the `mysql.slave_relay_log_info` table, and the columns displayed by `SHOW SLAVE STATUS`.

Line in relay-log.info	slave_relay_log_info Table Column	SHOW SLAVE STATUS Column	Description
1	Number_of_lines	[None]	Number of lines in the file or rows in the table
2	Relay_log_name	Relay_Log_File	The name of the current relay log file
3	Relay_log_pos	Relay_Log_Pos	The current position within the relay log file; events up to this position have been executed on the slave database
4	Master_log_name	Relay_Master_Log_File	The name of the master binary log file from which the events in the relay log file were read
5	Master_log_pos	Exec_Master_Log_Pos	The equivalent position within the master's binary log file of events that have already been executed
5	Sql_delay	SQL_Delay	The number of seconds that the slave must lag the master

Prior to MySQL 5.6, the `relay-log.info` file does not include a line count or a delay value (and the `slave_relay_log_info` table is not available).

Line	Status Column	Description
1	Relay_Log_File	The name of the current relay log file
2	Relay_Log_Pos	The current position within the relay log file; events up to this position have been executed on the slave database
3	Relay_Master_Log_File	The name of the master binary log file from which the events in the relay log file were read
4	Exec_Master_Log_Pos	The equivalent position within the master's binary log file of events that have already been executed



**Note**

If you downgrade a slave server to a version older than MySQL 5.6, the older server does not read the `relay-log.info` file correctly. To address this, modify the file in a text editor by deleting the initial line containing the number of lines.

The contents of the `relay-log.info` file and the states shown by the `SHOW SLAVE STATUS` statement might not match if the `relay-log.info` file has not been flushed to disk. Ideally, you



should only view `relay-log.info` on a slave that is offline (that is, `mysqld` is not running). For a running system, you can use `SHOW SLAVE STATUS`, or query the `slave_master_info` and `slave_relay_log_info` tables if you are writing the status logs to tables.

When you back up the slave's data, you should back up these two status logs, along with the relay log files. The status logs are needed to resume replication after you restore the data from the slave. If you lose the relay logs but still have the relay log info log, you can check it to determine how far the SQL thread has executed in the master binary logs. Then you can use `CHANGE MASTER TO` with the `MASTER_LOG_FILE` and `MASTER_LOG_POS` options to tell the slave to re-read the binary logs from that point. Of course, this requires that the binary logs still exist on the master.

### 16.2.3. How Servers Evaluate Replication Filtering Rules

If a master server does not write a statement to its binary log, the statement is not replicated. If the server does log the statement, the statement is sent to all slaves and each slave determines whether to execute it or ignore it.

On the master, you can control which databases to log changes for by using the `--binlog-do-db` [2004] and `--binlog-ignore-db` [2005] options to control binary logging. For a description of the rules that servers use in evaluating these options, see [Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”](#). You should not use these options to control which databases and tables are replicated. Instead, use filtering on the slave to control the events that are executed on the slave.

On the slave side, decisions about whether to execute or ignore statements received from the master are made according to the `--replicate-*` options that the slave was started with. (See [Section 16.1.4, “Replication and Binary Logging Options and Variables”](#).)

In the simplest case, when there are no `--replicate-*` options, the slave executes all statements that it receives from the master. Otherwise, the result depends on the particular options given.

Database-level options (`--replicate-do-db` [1976], `--replicate-ignore-db` [1977]) are checked first; see [Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”](#), for a description of this process. If no matching database-level options are found, option checking proceeds to any table-level options that may be in use, as discussed in [Section 16.2.3.2, “Evaluation of Table-Level Replication Options”](#).

For statements affecting databases only (that is, `CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE`), database-level options always take precedence over any `--replicate-wild-do-table` [1979] options. In other words, for such statements, `--replicate-wild-do-table` [1979] options are checked if and only if there are no database-level options that apply. This is a change in behavior from previous versions of MySQL, where the statement `CREATE DATABASE dbx` was not replicated if the slave had been started with `--replicate-do-db=dbx` [1976] `--replicate-wild-do-table=db%.t1` [1979]. (Bug #46110)

To make it easier to determine what effect an option set will have, it is recommended that you avoid mixing “do” and “ignore” options, or wildcard and nonwildcard options.

If any `--replicate-rewrite-db` [1979] options were specified, they are applied before the `--replicate-*` filtering rules are tested.



#### Note

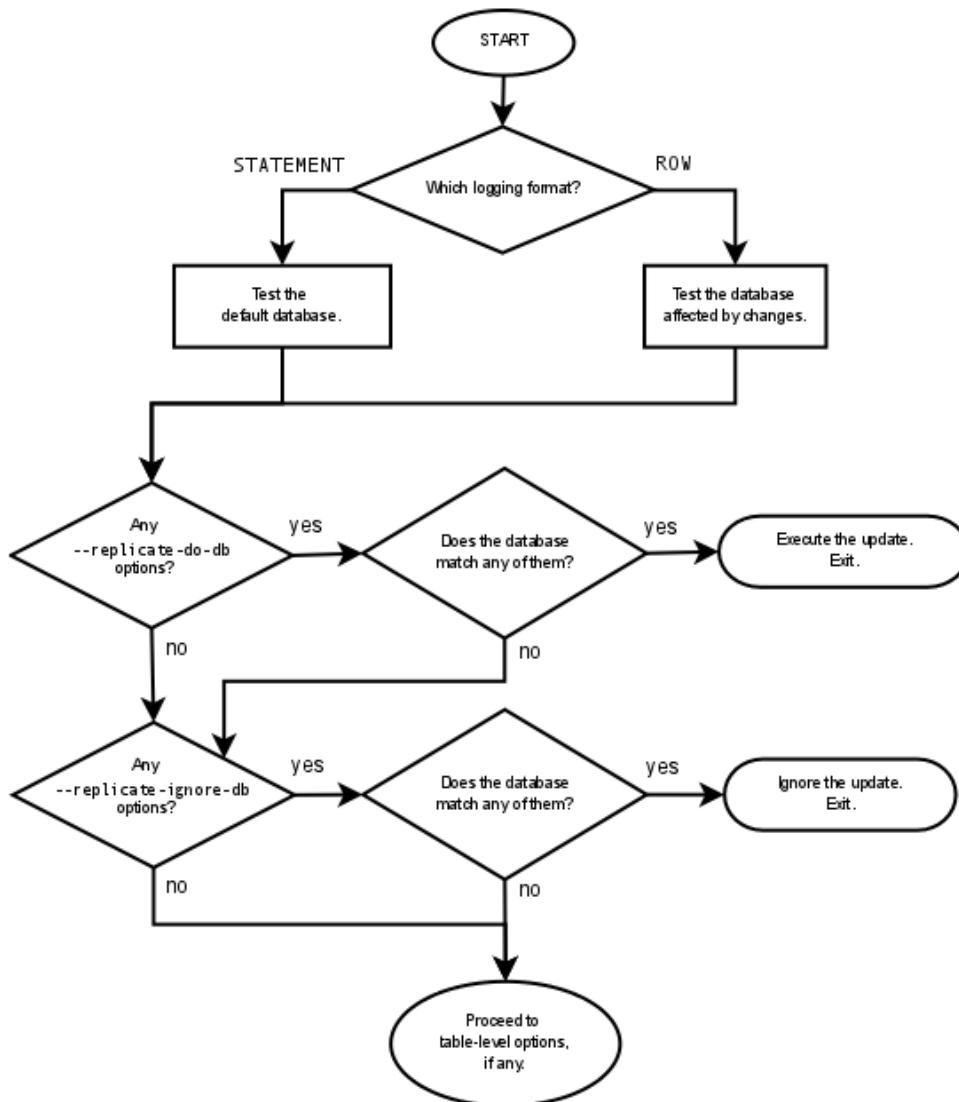
In MySQL 5.6, all replication filtering options follow the same rules for case sensitivity that apply to names of databases and tables elsewhere in the MySQL server, including the effects of the `lower_case_table_names` [501] system variable.

This is a change from previous versions of MySQL. (Bug #51639)

### 16.2.3.1. Evaluation of Database-Level Replication and Binary Logging Options

When evaluating replication options, the slave begins by checking to see whether there are any `--replicate-do-db` [1976] or `--replicate-ignore-db` [1977] options that apply. When using `--binlog-do-db` [2004] or `--binlog-ignore-db` [2005], the process is similar, but the options are checked on the master.

With statement-based replication, the default database is checked for a match. With row-based replication, the database where data is to be changed is the database that is checked. Regardless of the binary logging format, checking of database-level options proceeds as shown in the following diagram.



The steps involved are listed here:

1. Are there any `--replicate-do-db` [1976] options?
  - **Yes.** Do any of them match the database?
    - **Yes.** Execute the statement and exit.
    - **No.** Continue to step 2.
  - **No.** Continue to step 2.
2. Are there any `--replicate-ignore-db` [1977] options?

- **Yes.** Do any of them match the database?
    - **Yes.** Ignore the statement and exit.
    - **No.** Continue to step 3.
  - **No.** Continue to step 3.
3. Proceed to checking the table-level replication options, if there are any. For a description of how these options are checked, see [Section 16.2.3.2, “Evaluation of Table-Level Replication Options”](#).



**Important**

A statement that is still permitted at this stage is not yet actually executed. The statement is not executed until all table-level options (if any) have also been checked, and the outcome of that process permits execution of the statement.

For binary logging, the steps involved are listed here:

1. Are there any `--binlog-do-db [2004]` or `--binlog-ignore-db [2005]` options?
  - **Yes.** Continue to step 2.
  - **No.** Log the statement and exit.
2. Is there a default database (has any database been selected by `USE`)?
  - **Yes.** Continue to step 3.
  - **No.** Ignore the statement and exit.
3. There is a default database. Are there any `--binlog-do-db [2004]` options?
  - **Yes.** Do any of them match the database?
    - **Yes.** Log the statement and exit.
    - **No.** Ignore the statement and exit.
  - **No.** Continue to step 4.
4. Do any of the `--binlog-ignore-db [2005]` options match the database?
  - **Yes.** Ignore the statement and exit.
  - **No.** Log the statement and exit.



**Important**

For statement-based logging, an exception is made in the rules just given for the `CREATE DATABASE`, `ALTER DATABASE`, and `DROP DATABASE` statements. In those cases, the database being *created*, *altered*, or *dropped* replaces the default database when determining whether to log or ignore updates.

`--binlog-do-db [2004]` can sometimes mean “ignore other databases”. For example, when using statement-based logging, a server running with only `--binlog-do-db=sales [2004]` does not write to the binary log statements for which the default database differs from `sales`. When using row-based logging with the same option, the server logs only those updates that change data in `sales`.

### 16.2.3.2. Evaluation of Table-Level Replication Options

The slave checks for and evaluates table options only if no matching database options were found (see [Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”](#)).

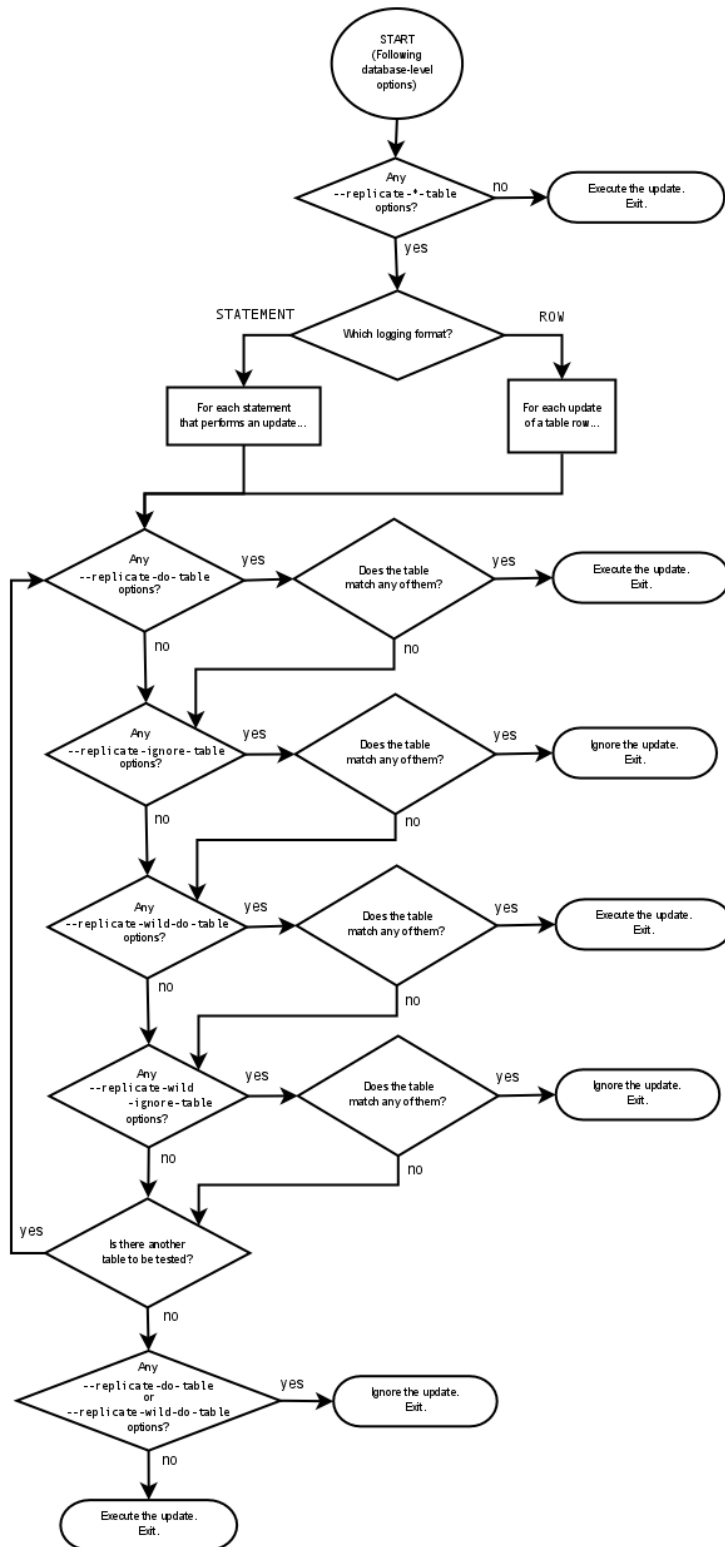
First, as a preliminary condition, the slave checks whether statement-based replication is enabled. If so, and the statement occurs within a stored function, the slave executes the statement and exits. If row-based replication is enabled, the slave does not know whether a statement occurred within a stored function on the master, so this condition does not apply.



#### Note

For statement-based replication, replication events represent statements (all changes making up a given event are associated with a single SQL statement); for row-based replication, each event represents a change in a single table row (thus a single statement such as `UPDATE mytable SET mycol = 1` may yield many row-based events). When viewed in terms of events, the process of checking table options is the same for both row-based and statement-based replication.

Having reached this point, if there are no table options, the slave simply executes all events. If there are any `--replicate-do-table` [\[1978\]](#) or `--replicate-wild-do-table` [\[1979\]](#) options, the event must match one of these if it is to be executed; otherwise, it is ignored. If there are any `--replicate-ignore-table` [\[1978\]](#) or `--replicate-wild-ignore-table` [\[1980\]](#) options, all events are executed except those that match any of these options. This process is illustrated in the following diagram.



The following steps describe this evaluation in more detail:

1. Are there any table options?
  - **Yes.** Continue to step 2.
  - **No.** Execute the event and exit.
2. Are there any `--replicate-do-table` [1978] options?

- **Yes.** Does the table match any of them?
    - **Yes.** Execute the event and exit.
    - **No.** Continue to step 3.
  - **No.** Continue to step 3.
3. Are there any `--replicate-ignore-table` [1978] options?
- **Yes.** Does the table match any of them?
    - **Yes.** Ignore the event and exit.
    - **No.** Continue to step 4.
  - **No.** Continue to step 4.
4. Are there any `--replicate-wild-do-table` [1979] options?
- **Yes.** Does the table match any of them?
    - **Yes.** Execute the event and exit.
    - **No.** Continue to step 5.
  - **No.** Continue to step 5.
5. Are there any `--replicate-wild-ignore-table` [1980] options?
- **Yes.** Does the table match any of them?
    - **Yes.** Ignore the event and exit.
    - **No.** Continue to step 6.
  - **No.** Continue to step 6.
6. Are there any `--replicate-do-table` [1978] or `--replicate-wild-do-table` [1979] options?
- **Yes.** Ignore the event and exit.
  - **No.** Execute the event and exit.

### 16.2.3.3. Replication Rule Application

This section provides additional explanation and examples of usage for different combinations of replication filtering options.

Some typical combinations of replication filter rule types are given in the following table:

Condition (Types of Options)	Outcome
No <code>--replicate-*</code> options at all:	The slave executes all events that it receives from the master.
<code>--replicate-*-db</code> options, but no table options:	The slave accepts or ignores events using the database options. It executes all events permitted by those options because there are no table restrictions.
<code>--replicate-*-table</code> options, but no database options:	All events are accepted at the database-checking stage because there are no database conditions. The slave

Condition (Types of Options)	Outcome
	executes or ignores events based solely on the table options.
A combination of database and table options:	The slave accepts or ignores events using the database options. Then it evaluates all events permitted by those options according to the table options. This can sometimes lead to results that seem counterintuitive, and that may be different depending on whether you are using statement-based or row-based replication; see the text for an example.

A more complex example follows, in which we examine the outcomes for both statement-based and row-based settings.

Suppose that we have two tables `mytbl1` in database `db1` and `mytbl2` in database `db2` on the master, and the slave is running with the following options (and no other replication filtering options):

```
replicate-ignore-db = db1
replicate-do-table = db2.tbl2
```

Now we execute the following statements on the master:

```
USE db1;
INSERT INTO db2.tbl2 VALUES (1);
```

The results on the slave vary considerably depending on the binary log format, and may not match initial expectations in either case.

**Statement-based replication.** The `USE` statement causes `db1` to be the default database. Thus the `--replicate-ignore-db [1977]` option matches, and the `INSERT` statement is ignored. The table options are not checked.

**Row-based replication.** The default database has no effect on how the slave reads database options when using row-based replication. Thus, the `USE` statement makes no difference in how the `--replicate-ignore-db [1977]` option is handled: the database specified by this option does not match the database where the `INSERT` statement changes data, so the slave proceeds to check the table options. The table specified by `--replicate-do-table [1978]` matches the table to be updated, and the row is inserted.

## 16.3. Replication Solutions

Replication can be used in many different environments for a range of purposes. This section provides general notes and advice on using replication for specific solution types.

For information on using replication in a backup environment, including notes on the setup, backup procedure, and files to back up, see [Section 16.3.1, “Using Replication for Backups”](#).

For advice and tips on using different storage engines on the master and slaves, see [Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”](#).

Using replication as a scale-out solution requires some changes in the logic and operation of applications that use the solution. See [Section 16.3.3, “Using Replication for Scale-Out”](#).

For performance or data distribution reasons, you may want to replicate different databases to different replication slaves. See [Section 16.3.4, “Replicating Different Databases to Different Slaves”](#)

As the number of replication slaves increases, the load on the master can increase and lead to reduced performance (because of the need to replicate the binary log to each slave). For tips on improving your replication performance, including using a single secondary server as an replication master, see [Section 16.3.5, “Improving Replication Performance”](#).

For guidance on switching masters, or converting slaves into masters as part of an emergency failover solution, see [Section 16.3.6, “Switching Masters During Failover”](#).



To secure your replication communication, you can use SSL to encrypt the communication channel. For step-by-step instructions, see [Section 16.3.7, “Setting Up Replication Using SSL”](#).

## 16.3.1. Using Replication for Backups

To use replication as a backup solution, replicate data from the master to a slave, and then back up the data slave. The slave can be paused and shut down without affecting the running operation of the master, so you can produce an effective snapshot of “live” data that would otherwise require the master to be shut down.

How you back up a database depends on its size and whether you are backing up only the data, or the data and the replication slave state so that you can rebuild the slave in the event of failure. There are therefore two choices:

- If you are using replication as a solution to enable you to back up the data on the master, and the size of your database is not too large, the `mysqldump` tool may be suitable. See [Section 16.3.1.1, “Backing Up a Slave Using `mysqldump`”](#).
- For larger databases, where `mysqldump` would be impractical or inefficient, you can back up the raw data files instead. Using the raw data files option also means that you can back up the binary and relay logs that will enable you to recreate the slave in the event of a slave failure. For more information, see [Section 16.3.1.2, “Backing Up Raw Data from a Slave”](#).

Another backup strategy, which can be used for either master or slave servers, is to put the server in a read-only state. The backup is performed against the read-only server, which then is changed back to its usual read/write operational status. See [Section 16.3.1.3, “Backing Up a Master or Slave by Making It Read Only”](#).

### 16.3.1.1. Backing Up a Slave Using `mysqldump`

Using `mysqldump` to create a copy of a database enables you to capture all of the data in the database in a format that enables the information to be imported into another instance of MySQL Server (see [Section 4.5.4, “`mysqldump` — A Database Backup Program”](#)). Because the format of the information is SQL statements, the file can easily be distributed and applied to running servers in the event that you need access to the data in an emergency. However, if the size of your data set is very large, `mysqldump` may be impractical.

When using `mysqldump`, you should stop replication on the slave before starting the dump process to ensure that the dump contains a consistent set of data:

1. Stop the slave from processing requests. You can stop replication completely on the slave using `mysqladmin`:

```
shell> mysqladmin stop-slave
```

Alternatively, you can stop only the slave SQL thread to pause event execution:

```
shell> mysql -e 'STOP SLAVE SQL_THREAD;'
```

This enables the slave to continue to receive data change events from the master's binary log and store them in the relay logs using the I/O thread, but prevents the slave from executing these events and changing its data. Within busy replication environments, permitting the I/O thread to run during backup may speed up the catch-up process when you restart the slave SQL thread.

2. Run `mysqldump` to dump your databases. You may either dump all databases or select databases to be dumped. For example, to dump all databases:

```
shell> mysqldump --all-databases > fulldb.dump
```

3. Once the dump has completed, start slave operations again:

```
shell> mysqladmin start-slave
```

In the preceding example, you may want to add login credentials (user name, password) to the commands, and bundle the process up into a script that you can run automatically each day.

If you use this approach, make sure you monitor the slave replication process to ensure that the time taken to run the backup does not affect the slave's ability to keep up with events from the master. See [Section 16.1.5.1, “Checking Replication Status”](#). If the slave is unable to keep up, you may want to add another slave and distribute the backup process. For an example of how to configure this scenario, see [Section 16.3.4, “Replicating Different Databases to Different Slaves”](#).

### 16.3.1.2. Backing Up Raw Data from a Slave

To guarantee the integrity of the files that are copied, backing up the raw data files on your MySQL replication slave should take place while your slave server is shut down. If the MySQL server is still running, background tasks may still be updating the database files, particularly those involving storage engines with background processes such as [InnoDB](#). With [InnoDB](#), these problems should be resolved during crash recovery, but since the slave server can be shut down during the backup process without affecting the execution of the master it makes sense to take advantage of this capability.

To shut down the server and back up the files:

1. Shut down the slave MySQL server:

```
shell> mysqladmin shutdown
```

2. Copy the data files. You can use any suitable copying or archive utility, including [cp](#), [tar](#) or [WinZip](#). For example, assuming that the data directory is located under the current directory, you can archive the entire directory as follows:

```
shell> tar cf /tmp/dbbackup.tar ./data
```

3. Start the MySQL server again. Under Unix:

```
shell> mysqld_safe &
```

Under Windows:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld"
```

Normally you should back up the entire data directory for the slave MySQL server. If you want to be able to restore the data and operate as a slave (for example, in the event of failure of the slave), then in addition to the slave's data, you should also back up the slave status files, [master.info](#) and [relay-log.info](#), along with the relay log files. These files are needed to resume replication after you restore the slave's data.

If you lose the relay logs but still have the [relay-log.info](#) file, you can check it to determine how far the SQL thread has executed in the master binary logs. Then you can use [CHANGE MASTER TO](#) with the [MASTER\\_LOG\\_FILE](#) and [MASTER\\_LOG\\_POS](#) options to tell the slave to re-read the binary logs from that point. This requires that the binary logs still exist on the master server.

If your slave is replicating [LOAD DATA INFILE](#) statements, you should also back up any [SQL\\_LOAD-\\*](#) files that exist in the directory that the slave uses for this purpose. The slave needs these files to resume replication of any interrupted [LOAD DATA INFILE](#) operations. The location of this directory is the value of the [--slave-load-tmpdir \[1985\]](#) option. If the server was not started with that option, the directory location is the value of the [tmpdir \[559\]](#) system variable.

### 16.3.1.3. Backing Up a Master or Slave by Making It Read Only

It is possible to back up either master or slave servers in a replication setup by acquiring a global read lock and manipulating the [read\\_only \[532\]](#) system variable to change the read-only state of the server to be backed up:

1. Make the server read-only, so that it processes only retrievals and blocks updates.
2. Perform the backup.

3. Change the server back to its normal read/write state.

**Note**

The instructions in this section place the server to be backed up in a state that is safe for backup methods that get the data from the server, such as `mysqldump` (see [Section 4.5.4, “mysqldump — A Database Backup Program”](#)). You should not attempt to use these instructions to make a binary backup by copying files directly because the server may still have modified data cached in memory and not flushed to disk.

The following instructions describe how to do this for a master server and for a slave server. For both scenarios discussed here, suppose that you have the following replication setup:

- A master server M1
- A slave server S1 that has M1 as its master
- A client C1 connected to M1
- A client C2 connected to S1

In either scenario, the statements to acquire the global read lock and manipulate the `read_only` [532] variable are performed on the server to be backed up and do not propagate to any slaves of that server.

**Scenario 1: Backup with a Read-Only Master**

Put the master M1 in a read-only state by executing these statements on it:

```
mysql> FLUSH TABLES WITH READ LOCK;  
mysql> SET GLOBAL read_only = ON;
```

While M1 is in a read-only state, the following properties are true:

- Requests for updates sent by C1 to M1 will block because the server is in read-only mode.
- Requests for query results sent by C1 to M1 will succeed.
- Making a backup on M1 is safe.
- Making a backup on S1 is not safe. This server is still running, and might be processing the binary log or update requests coming from client C2

While M1 is read only, perform the backup. For example, you can use `mysqldump`.

After the backup operation on M1 completes, restore M1 to its normal operational state by executing these statements:

```
mysql> SET GLOBAL read_only = OFF;  
mysql> UNLOCK TABLES;
```

Although performing the backup on M1 is safe (as far as the backup is concerned), it is not optimal for performance because clients of M1 are blocked from executing updates.

This strategy applies to backing up a master server in a replication setup, but can also be used for a single server in a nonreplication setting.

**Scenario 2: Backup with a Read-Only Slave**

Put the slave S1 in a read-only state by executing these statements on it:

```
mysql> FLUSH TABLES WITH READ LOCK;  
mysql> SET GLOBAL read_only = ON;
```

While S1 is in a read-only state, the following properties are true:

- The master M1 will continue to operate, so making a backup on the master is not safe.
- The slave S1 is stopped, so making a backup on the slave S1 is safe.

These properties provide the basis for a popular backup scenario: Having one slave busy performing a backup for a while is not a problem because it does not affect the entire network, and the system is still running during the backup. In particular, clients can still perform updates on the master server, which remains unaffected by backup activity on the slave.

While S1 is read only, perform the backup. For example, you can use `mysqldump`.

After the backup operation on S1 completes, restore S1 to its normal operational state by executing these statements:

```
mysql> SET GLOBAL read_only = OFF;
mysql> UNLOCK TABLES;
```

After the slave is restored to normal operation, it again synchronizes to the master by catching up with any outstanding updates from the binary log of the master.

## 16.3.2. Using Replication with Different Master and Slave Storage Engines

It does not matter for the replication process whether the source table on the master and the replicated table on the slave use different engine types. In fact, the `default_storage_engine` [474] and `storage_engine` [549] system variables are not replicated.

This provides a number of benefits in the replication process in that you can take advantage of different engine types for different replication scenarios. For example, in a typical scale-out scenario (see [Section 16.3.3, “Using Replication for Scale-Out”](#)), you want to use `InnoDB` tables on the master to take advantage of the transactional functionality, but use `MyISAM` on the slaves where transaction support is not required because the data is only read. When using replication in a data-logging environment you may want to use the `Archive` storage engine on the slave.

Configuring different engines on the master and slave depends on how you set up the initial replication process:

- If you used `mysqldump` to create the database snapshot on your master, you could edit the dump file text to change the engine type used on each table.

Another alternative for `mysqldump` is to disable engine types that you do not want to use on the slave before using the dump to build the data on the slave. For example, you can add the `--skip-innodb` [1698] option on your slave to disable the `InnoDB` engine. If a specific engine does not exist for a table to be created, MySQL will use the default engine type, usually `MyISAM`. (This requires that the `NO_ENGINE_SUBSTITUTION` [605] SQL mode is not enabled.) If you want to disable additional engines in this way, you may want to consider building a special binary to be used on the slave that only supports the engines you want.

- If you are using raw data files (a binary backup) to set up the slave, you will be unable to change the initial table format. Instead, use `ALTER TABLE` to change the table types after the slave has been started.
- For new master/slave replication setups where there are currently no tables on the master, avoid specifying the engine type when creating new tables.

If you are already running a replication solution and want to convert your existing tables to another engine type, follow these steps:

1. Stop the slave from running replication updates:

```
mysql> STOP SLAVE;
```

This will enable you to change engine types without interruptions.

2. Execute an `ALTER TABLE ... ENGINE='engine_type'` for each table to be changed.

- Start the slave replication process again:

```
mysql> START SLAVE;
```

Although the `default_storage_engine` [474] variable is not replicated, be aware that `CREATE TABLE` and `ALTER TABLE` statements that include the engine specification will be correctly replicated to the slave. For example, if you have a CSV table and you execute:

```
mysql> ALTER TABLE csvtable Engine='MyISAM';
```

The above statement will be replicated to the slave and the engine type on the slave will be converted to `MyISAM`, even if you have previously changed the table type on the slave to an engine other than CSV. If you want to retain engine differences on the master and slave, you should be careful to use the `default_storage_engine` [474] variable on the master when creating a new table. For example, instead of:

```
mysql> CREATE TABLE tablea (columna int) Engine=MyISAM;
```

Use this format:

```
mysql> SET default_storage_engine=MyISAM;
mysql> CREATE TABLE tablea (columna int);
```

When replicated, the `default_storage_engine` [474] variable will be ignored, and the `CREATE TABLE` statement will execute on the slave using the slave's default engine.

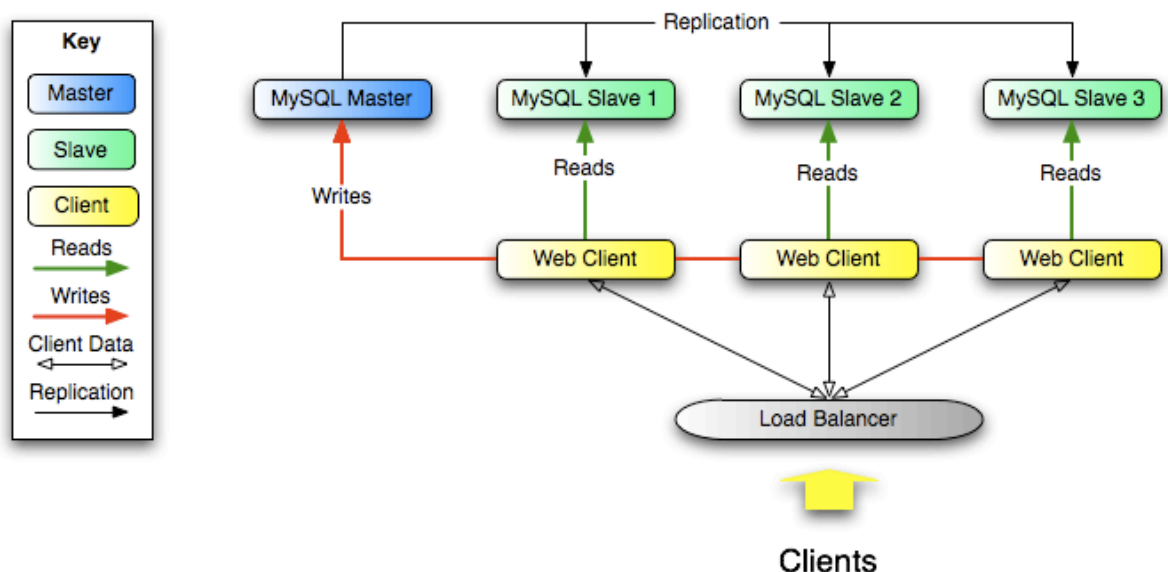
### 16.3.3. Using Replication for Scale-Out

You can use replication as a scale-out solution; that is, where you want to split up the load of database queries across multiple database servers, within some reasonable limitations.

Because replication works from the distribution of one master to one or more slaves, using replication for scale-out works best in an environment where you have a high number of reads and low number of writes/updates. Most Web sites fit into this category, where users are browsing the Web site, reading articles, posts, or viewing products. Updates only occur during session management, or when making a purchase or adding a comment/message to a forum.

Replication in this situation enables you to distribute the reads over the replication slaves, while still enabling your web servers to communicate with the replication master when a write is required. You can see a sample replication layout for this scenario in [Figure 16.1, "Using Replication to Improve Performance During Scale-Out"](#).

**Figure 16.1. Using Replication to Improve Performance During Scale-Out**



If the part of your code that is responsible for database access has been properly abstracted/modularized, converting it to run with a replicated setup should be very smooth and easy. Change the implementation of your database access to send all writes to the master, and to send reads to either the master or a slave. If your code does not have this level of abstraction, setting up a replicated system gives you the opportunity and motivation to clean it up. Start by creating a wrapper library or module that implements the following functions:

- `safe_writer_connect()`
- `safe_reader_connect()`
- `safe_reader_statement()`
- `safe_writer_statement()`

`safe_` in each function name means that the function takes care of handling all error conditions. You can use different names for the functions. The important thing is to have a unified interface for connecting for reads, connecting for writes, doing a read, and doing a write.


Then convert your client code to use the wrapper library. This may be a painful and scary process at first, but it pays off in the long run. All applications that use the approach just described are able to take advantage of a master/slave configuration, even one involving multiple slaves. The code is much easier to maintain, and adding troubleshooting options is trivial. You need modify only one or two functions; for example, to log how long each statement took, or which statement among those issued gave you an error.

If you have written a lot of code, you may want to automate the conversion task by using the `replace` utility that comes with standard MySQL distributions, or write your own conversion script. Ideally, your code uses consistent programming style conventions. If not, then you are probably better off rewriting it anyway, or at least going through and manually regularizing it to use a consistent style.

#### 16.3.4. Replicating Different Databases to Different Slaves

There may be situations where you have a single master and want to replicate different databases to different slaves. For example, you may want to distribute different sales data to different departments to help spread the load during data analysis. A sample of this layout is shown in [Figure 16.2, "Using Replication to Replicate Databases to Separate Replication Slaves"](#).

**Figure 16.2. Using Replication to Replicate Databases to Separate Replication Slaves**



You can achieve this separation by configuring the master and slaves as normal, and then limiting the binary log statements that each slave processes by using the `--replicate-wild-do-table` [1979] configuration option on each slave.



### Important

You should *not* use `--replicate-do-db [1976]` for this purpose when using statement-based replication, since statement-based replication causes this option's affects to vary according to the database that is currently selected. This applies to mixed-format replication as well, since this enables some updates to be replicated using the statement-based format.

However, it should be safe to use `--replicate-do-db [1976]` for this purpose if you are using row-based replication only, since in this case the currently selected database has no effect on the option's operation.

For example, to support the separation as shown in [Figure 16.2, “Using Replication to Replicate Databases to Separate Replication Slaves”](#), you should configure each replication slave as follows, before executing `START SLAVE`:

- Replication slave 1 should use `--replicate-wild-do-table=databaseA.%`.
- Replication slave 2 should use `--replicate-wild-do-table=databaseB.%`.
- Replication slave 3 should use `--replicate-wild-do-table=databaseC.%`.

Each slave in this configuration receives the entire binary log from the master, but executes only those events from the binary log that apply to the databases and tables included by the `--replicate-wild-do-table [1979]` option in effect on that slave.

If you have data that must be synchronized to the slaves before replication starts, you have a number of choices:

- Synchronize all the data to each slave, and delete the databases, tables, or both that you do not want to keep.
- Use `mysqldump` to create a separate dump file for each database and load the appropriate dump file on each slave.
- Use a raw data file dump and include only the specific files and databases that you need for each slave.



### Note

This does not work with `InnoDB` databases unless you use `innodb_file_per_table [1718]`.


## 16.3.5. Improving Replication Performance

As the number of slaves connecting to a master increases, the load, although minimal, also increases, as each slave uses a client connection to the master. Also, as each slave must receive a full copy of the master binary log, the network load on the master may also increase and create a bottleneck.

If you are using a large number of slaves connected to one master, and that master is also busy processing requests (for example, as part of a scale-out solution), then you may want to improve the performance of the replication process.

One way to improve the performance of the replication process is to create a deeper replication structure that enables the master to replicate to only one slave, and for the remaining slaves to connect to this primary slave for their individual replication requirements. A sample of this structure is shown in [Figure 16.3, “Using an Additional Replication Host to Improve Performance”](#).



**Figure 16.3. Using an Additional Replication Host to Improve Performance**

For this to work, you must configure the MySQL instances as follows:

- Master 1 is the primary master where all changes and updates are written to the database. Binary logging should be enabled on this machine.
- Master 2 is the slave to the Master 1 that provides the replication functionality to the remainder of the slaves in the replication structure. Master 2 is the only machine permitted to connect to Master 1. Master 2 also has binary logging enabled, and the `--log-slave-updates` [1971] option so that replication instructions from Master 1 are also written to Master 2's binary log so that they can then be replicated to the true slaves.
- Slave 1, Slave 2, and Slave 3 act as slaves to Master 2, and replicate the information from Master 2, which actually consists of the upgrades logged on Master 1.

The above solution reduces the client load and the network interface load on the primary master, which should improve the overall performance of the primary master when used as a direct database solution.

If your slaves are having trouble keeping up with the replication process on the master, there are a number of options available:

- If possible, put the relay logs and the data files on different physical drives. To do this, use the `--relay-log` [1974] option to specify the location of the relay log.
- If the slaves are significantly slower than the master, you may want to divide up the responsibility for replicating different databases to different slaves. See [Section 16.3.4, “Replicating Different Databases to Different Slaves”](#).
- If your master makes use of transactions and you are not concerned about transaction support on your slaves, use `MyISAM` or another nontransactional engine on the slaves. See [Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”](#).
- If your slaves are not acting as masters, and you have a potential solution in place to ensure that you can bring up a master in the event of failure, then you can switch off `--log-slave-updates` [1971]. This prevents “dumb” slaves from also logging events they have executed into their own binary log.

### 16.3.6. Switching Masters During Failover

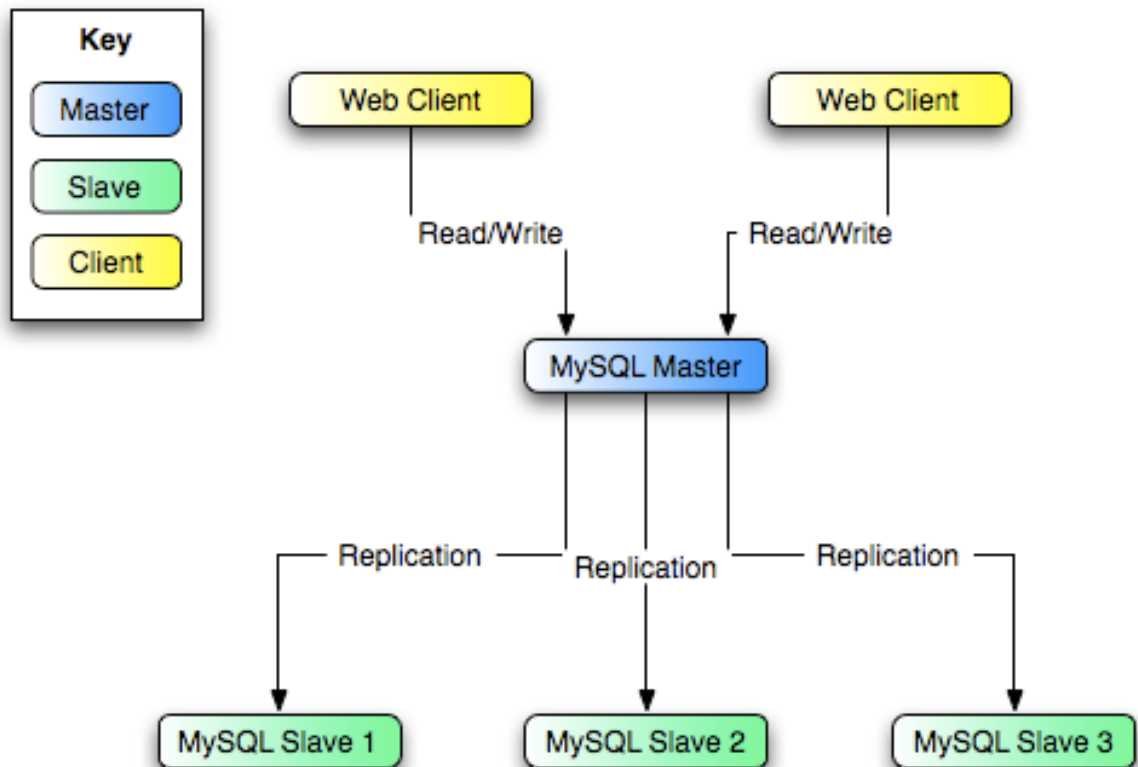
There is currently no official solution for providing failover between master and slaves in the event of a failure. With the currently available features, you would have to set up a master and a slave (or several slaves), and to write a script that monitors the master to check whether it is up. Then instruct your applications and the slaves to change master in case of failure.

Remember that you can tell a slave to change its master at any time, using the `CHANGE MASTER TO` statement. The slave will not check whether the databases on the master are compatible with the

slave, it will just start reading and executing events from the specified binary log coordinates on the new master. In a failover situation, all the servers in the group are typically executing the same events from the same binary log file, so changing the source of the events should not affect the database structure or integrity providing you are careful.

Run your slaves with the `--log-bin` [2002] option and without `--log-slave-updates` [1971]. In this way, the slave is ready to become a master as soon as you issue `STOP SLAVE; RESET MASTER,` and `CHANGE MASTER TO` statement on the other slaves. For example, assume that you have the structure shown in Figure 16.4, “Redundancy Using Replication, Initial Structure”.

**Figure 16.4. Redundancy Using Replication, Initial Structure**



In this diagram, the `MySQL Master` holds the master database, the `MySQL Slave` hosts are replication slaves, and the `Web Client` machines are issuing database reads and writes. Web clients that issue only reads (and would normally be connected to the slaves) are not shown, as they do not need to switch to a new server in the event of failure. For a more detailed example of a read/write scale-out replication structure, see Section 16.3.3, “Using Replication for Scale-Out”.

Each `MySQL Slave` (`Slave 1`, `Slave 2`, and `Slave 3`) is a slave running with `--log-bin` [2002] and without `--log-slave-updates` [1971]. Because updates received by a slave from the master are not logged in the binary log unless `--log-slave-updates` [1971] is specified, the binary log on each slave is empty initially. If for some reason `MySQL Master` becomes unavailable, you can pick one of the slaves to become the new master. For example, if you pick `Slave 1`, all `Web Clients` should be redirected to `Slave 1`, which will log updates to its binary log. `Slave 2` and `Slave 3` should then replicate from `Slave 1`.

The reason for running the slave without `--log-slave-updates` [1971] is to prevent slaves from receiving updates twice in case you cause one of the slaves to become the new master. Suppose that `Slave 1` has `--log-slave-updates` [1971] enabled. Then it will write updates that it receives from `Master` to its own binary log. When `Slave 2` changes from `Master` to `Slave 1` as its master, it may receive updates from `Slave 1` that it has already received from `Master`.

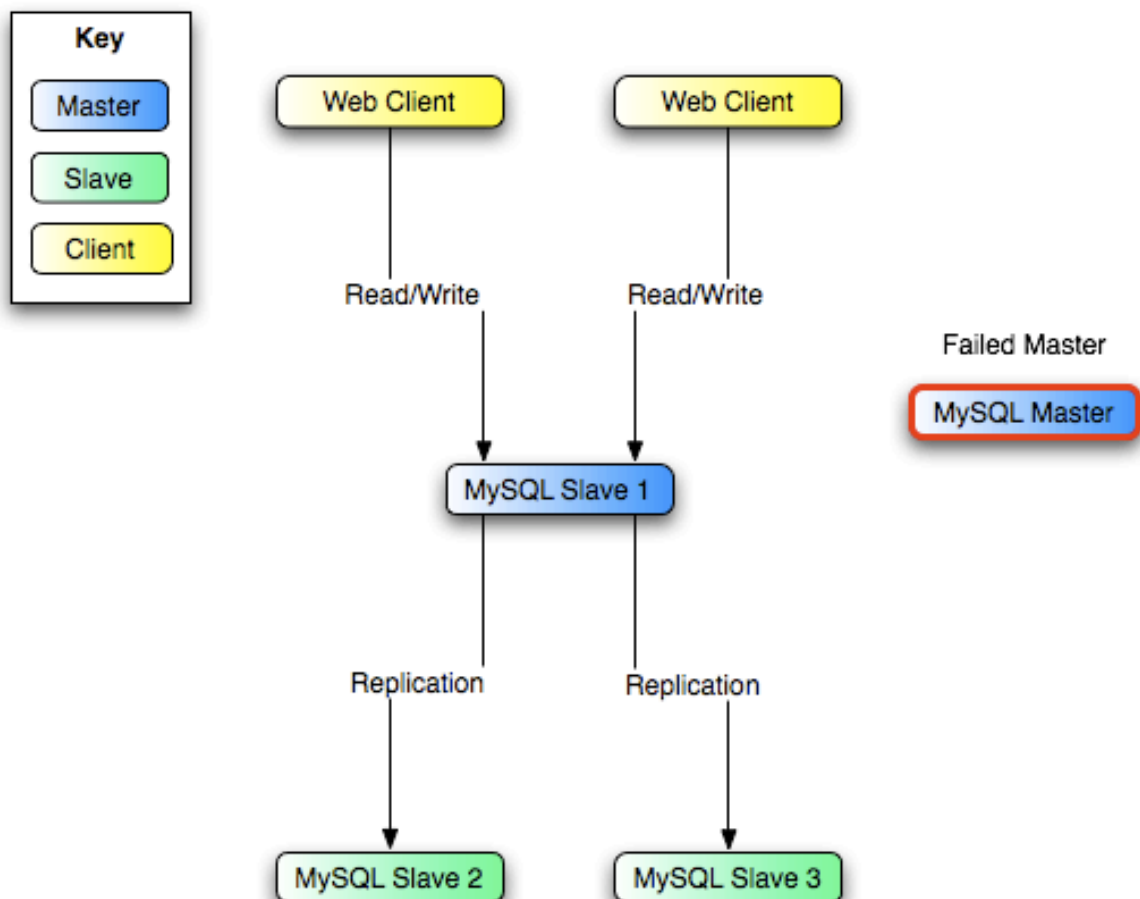
Make sure that all slaves have processed any statements in their relay log. On each slave, issue `STOP SLAVE IO_THREAD`, then check the output of `SHOW PROCESSLIST` until you see `Has read all relay log`. When this is true for all slaves, they can be reconfigured to the new setup. On the slave `Slave 1` being promoted to become the master, issue `STOP SLAVE` and `RESET MASTER`.

On the other slaves `Slave 2` and `Slave 3`, use `STOP SLAVE` and `CHANGE MASTER TO MASTER_HOST='Slave1'` (where `'Slave1'` represents the real host name of `Slave 1`). To use `CHANGE MASTER TO`, add all information about how to connect to `Slave 1` from `Slave 2` or `Slave 3` (`user, password, port`). In `CHANGE MASTER TO`, there is no need to specify the name of the `Slave 1` binary log file or log position to read from: We know it is the first binary log file and position 4, which are the defaults for `CHANGE MASTER TO`. Finally, use `START SLAVE` on `Slave 2` and `Slave 3`.

Once the new replication is in place, you will then need to instruct each `Web Client` to direct its statements to `Slave 1`. From that point on, all updates statements sent by `Web Client` to `Slave 1` are written to the binary log of `Slave 1`, which then contains every update statement sent to `Slave 1` since `Master` died.

The resulting server structure is shown in [Figure 16.5, "Redundancy Using Replication, After Master Failure"](#).

**Figure 16.5. Redundancy Using Replication, After Master Failure**



When `Master` is up again, you must issue on it the same `CHANGE MASTER TO` as that issued on `Slave 2` and `Slave 3`, so that `Master` becomes a slave of `S1` and picks up each `Web Client` writes that it missed while it was down.

To make `Master` a master again (for example, because it is the most powerful machine), use the preceding procedure as if `Slave 1` was unavailable and `Master` was to be the new master. During this procedure, do not forget to run `RESET MASTER` on `Master` before making `Slave 1`, `Slave 2`, and `Slave 3` slaves of `Master`. Otherwise, they may pick up old `Web Client` writes from before the point at which `Master` became unavailable.

Note that there is no synchronization between the different slaves to a master. Some slaves might be ahead of others. This means that the concept outlined in the previous example might not work. In practice, however, the relay logs of different slaves will most likely not be far behind the master, so it would work, anyway (but there is no guarantee).

A good way to keep your applications informed as to the location of the master is by having a dynamic DNS entry for the master. With `bind` you can use `nsupdate` to dynamically update your DNS.

### 16.3.7. Setting Up Replication Using SSL

To use SSL for encrypting the transfer of the binary log required during replication, both the master and the slave must support SSL network connections. If either host does not support SSL connections (because it has not been compiled or configured for SSL), replication through an SSL connection is not possible.

Setting up replication using an SSL connection is similar to setting up a server and client using SSL. You must obtain (or create) a suitable security certificate that you can use on the master, and a similar certificate (from the same certificate authority) on each slave.

For more information on setting up a server and client for SSL connectivity, see [Section 6.3.9.2, “Configuring MySQL for SSL”](#).

To enable SSL on the master you must create or obtain suitable certificates, and then add the following configuration options to the master's configuration within the `[mysqld]` section of the master's `my.cnf` file:

```
[mysqld]
ssl-ca=cacert.pem
ssl-cert=server-cert.pem
ssl-key=server-key.pem
```

The paths to the certificates may be relative or absolute; we recommend that you always use complete paths for this purpose.

The options are as follows:

- `ssl-ca` identifies the Certificate Authority (CA) certificate.
- `ssl-cert` identifies the server public key. This can be sent to the client and authenticated against the CA certificate that it has.
- `ssl-key` identifies the server private key.

On the slave, you have two options available for setting the SSL information. You can either add the slave certificates to the `[client]` section of the slave's `my.cnf` file, or you can explicitly specify the SSL information using the `CHANGE MASTER TO` statement:

- To add the slave certificates using an option file, add the following lines to the `[client]` section of the slave's `my.cnf` file:

```
[client]
ssl-ca=cacert.pem
ssl-cert=client-cert.pem
ssl-key=client-key.pem
```

Restart the slave server, using the `--skip-slave-start` [1984] option to prevent the slave from connecting to the master. Use `CHANGE MASTER TO` to specify the master configuration, using the `MASTER_SSL` option to enable SSL connectivity:

```
mysql> CHANGE MASTER TO
-> MASTER_HOST='master_hostname',
-> MASTER_USER='replicate',
-> MASTER_PASSWORD='password',
-> MASTER_SSL=1;
```

- To specify the SSL certificate options using the `CHANGE MASTER TO` statement, append the SSL options:

```
mysql> CHANGE MASTER TO
-> MASTER_HOST='master_hostname',
-> MASTER_USER='replicate',
-> MASTER_PASSWORD='password',
-> MASTER_SSL=1,
-> MASTER_SSL_CA = 'ca_file_name',
-> MASTER_SSL_CAPATH = 'ca_directory_name',
-> MASTER_SSL_CERT = 'cert_file_name',
-> MASTER_SSL_KEY = 'key_file_name';
```

After the master information has been updated, start the slave replication process:

```
mysql> START SLAVE;
```

You can use the `SHOW SLAVE STATUS` statement to confirm that the SSL connection was established successfully.

For more information on the `CHANGE MASTER TO` statement, see [Section 13.4.2.1, “CHANGE MASTER TO Syntax”](#).

If you want to enforce the use of SSL connections during replication, then create a user with the `REPLICATION SLAVE [755]` privilege and use the `REQUIRE SSL` option for that user. For example:

```
mysql> CREATE USER 'repl'@'%.mydomain.com' IDENTIFIED BY 'slavepass';
mysql> GRANT REPLICATION SLAVE ON *.*
-> TO 'repl'@'%.mydomain.com' REQUIRE SSL;
```

If the account already exists, you can add `REQUIRE SSL` to it with this statement:

```
mysql> GRANT USAGE ON *.*
-> TO 'repl'@'%.mydomain.com' REQUIRE SSL;
```

## 16.3.8. Semisynchronous Replication

MySQL 5.6 supports an interface to semisynchronous replication in addition to the built-in asynchronous replication. This section discusses what semisynchronous replication is and how it works. The following sections cover the administrative interface to semisynchronous replication and how to install, configure, and monitor it.

MySQL replication by default is asynchronous. The master writes events to its binary log but does not know whether or when a slave has retrieved and processed them. With asynchronous replication, if the master crashes, transactions that it has committed might not have been transmitted to any slave. Consequently, failover from master to slave in this case may result in failover to a server that is missing transactions relative to the master.

Semisynchronous replication can be used as an alternative to asynchronous replication:

- A slave indicates whether it is semisynchronous-capable when it connects to the master.
- If semisynchronous replication is enabled on the master side and there is at least one semisynchronous slave, a thread that performs a transaction commit on the master blocks after the commit is done and waits until at least one semisynchronous slave acknowledges that it has received all events for the transaction, or until a timeout occurs.

- The slave acknowledges receipt of a transaction's events only after the events have been written to its relay log and flushed to disk.
- If a timeout occurs without any slave having acknowledged the transaction, the master reverts to asynchronous replication. When at least one semisynchronous slave catches up, the master returns to semisynchronous replication.
- Semisynchronous replication must be enabled on both the master and slave sides. If semisynchronous replication is disabled on the master, or enabled on the master but on no slaves, the master uses asynchronous replication.

While the master is blocking (waiting for acknowledgment from a slave after having performed a commit), it does not return to the session that performed the transaction. When the block ends, the master returns to the session, which then can proceed to execute other statements. At this point, the transaction has committed on the master side, and receipt of its events has been acknowledged by at least one slave.

Blocking also occurs after rollbacks that are written to the binary log, which occurs when a transaction that modifies nontransactional tables is rolled back. The rolled-back transaction is logged even though it has no effect for transactional tables because the modifications to the nontransactional tables cannot be rolled back and must be sent to slaves.

For statements that do not occur in transactional context (that is, when no transaction has been started with `START TRANSACTION` or `SET autocommit = 0`), autocommit is enabled and each statement commits implicitly. With semisynchronous replication, the master blocks after committing each such statement, just as it does for explicit transaction commits.

To understand what the “semi” in “semisynchronous replication” means, compare it with asynchronous and fully synchronous replication:

- With asynchronous replication, the master writes events to its binary log and slaves request them when they are ready. There is no guarantee that any event will ever reach any slave.
- With fully synchronous replication, when a master commits a transaction, all slaves also will have committed the transaction before the master returns to the session that performed the transaction. The drawback of this is that there might be a lot of delay to complete a transaction.
- Semisynchronous replication falls between asynchronous and fully synchronous replication. The master waits after commit only until at least one slave has received and logged the events. It does not wait for all slaves to acknowledge receipt, and it requires only receipt, not that the events have been fully executed and committed on the slave side.

Compared to asynchronous replication, semisynchronous replication provides improved data integrity. When a commit returns successfully, it is known that the data exists in at least two places (on the master and at least one slave). If the master commits but a crash occurs while the master is waiting for acknowledgment from a slave, it is possible that the transaction may not have reached any slave.

Semisynchronous replication also places a rate limit on busy sessions by constraining the speed at which binary log events can be sent from master to slave. When one user is too busy, this will slow it down, which is useful in some deployment situations.

Semisynchronous replication does have some performance impact because commits are slower due to the need to wait for slaves. This is the tradeoff for increased data integrity. The amount of slowdown is at least the TCP/IP roundtrip time to send the commit to the slave and wait for the acknowledgment of receipt by the slave. This means that semisynchronous replication works best for close servers communicating over fast networks, and worst for distant servers communicating over slow networks.

### 16.3.8.1. Semisynchronous Replication Administrative Interface

The administrative interface to semisynchronous replication has several components:

- Two plugins implement semisynchronous capability. There is one plugin for the master side and one for the slave side.
- System variables control plugin behavior. Some examples:

- `rpl_semi_sync_master_enabled` [536]

Controls whether semisynchronous replication is enabled on the master. To enable or disable the plugin, set this variable to 1 or 0, respectively. The default is 1.

- `rpl_semi_sync_master_timeout` [536]

A value in milliseconds that controls how long the master waits on a commit for acknowledgment from a slave before timing out and reverting to asynchronous replication. The default value is 10000 (10 seconds).

- `rpl_semi_sync_slave_enabled` [537]

Similar to `rpl_semi_sync_master_enabled` [536], but controls the slave plugin.

All `rpl_semi_sync_xxx` system variables are described at [Section 5.1.4, “Server System Variables”](#).

- Status variables enable semisynchronous replication monitoring. Some examples:

- `Rpl_semi_sync_master_clients` [596]

The number of semisynchronous slaves.

- `Rpl_semi_sync_master_status` [596]

Whether semisynchronous replication currently is operational on the master. The value is 1 if the plugin has been enabled and a commit acknowledgment has not occurred. It is 0 if the plugin is not enabled or the master has fallen back to asynchronous replication due to commit acknowledgment timeout.

- `Rpl_semi_sync_master_no_tx` [596]

The number of commits that were not acknowledged successfully by a slave.

- `Rpl_semi_sync_master_yes_tx` [597]

The number of commits that were acknowledged successfully by a slave.

- `Rpl_semi_sync_slave_status` [597]

Whether semisynchronous replication currently is operational on the slave. This is 1 if the plugin has been enabled and the slave I/O thread is running, 0 otherwise.

All `Rpl_semi_sync_xxx` status variables are described at [Section 5.1.6, “Server Status Variables”](#).

The system and status variables are available only if the appropriate master or slave plugin has been installed with `INSTALL PLUGIN`.

### 16.3.8.2. Semisynchronous Replication Installation and Configuration

Semisynchronous replication is implemented using plugins, so the plugins must be installed into the server to make them available. After a plugin has been installed, you control it by means of the system variables associated with it. These system variables are unavailable until the associated plugin has been installed.

To use semisynchronous replication, the following requirements must be satisfied:



- MySQL 5.5 or higher must be installed.
- The capability of installing plugins requires a MySQL server that supports dynamic loading. To verify this, check that the value of the `have_dynamic_loading` [487] system variable is `YES`. Binary distributions should support dynamic loading.
- Replication must already be working. For information on creating a master/slave relationship, see [Section 16.1.1, “How to Set Up Replication”](#).

To set up semisynchronous replication, use the following instructions. The `INSTALL PLUGIN`, `SET GLOBAL`, `STOP SLAVE`, and `START SLAVE` statements mentioned here require the `SUPER` [756] privilege.

The semisynchronous replication plugins are included with MySQL distributions.

Unpack the component distribution, which contains files for the master side and the slave side.

Install the component files in the plugin directory of the appropriate server. Install the `semisync_master*` files in the plugin directory of the master server. Install the `semisync_slave*` files in the plugin directory of each slave server. The location of the plugin directory is available as the value of the server's `plugin_dir` [525] system variable.

To load the plugins, use the `INSTALL PLUGIN` statement on the master and on each slave that is to be semisynchronous.

On the master:

```
mysql> INSTALL PLUGIN rpl_semi_sync_master SONAME 'semisync_master.so';
```

On each slave:

```
mysql> INSTALL PLUGIN rpl_semi_sync_slave SONAME 'semisync_slave.so';
```

The preceding commands use a plugin file name suffix of `.so`. A different suffix might apply on your system. If you are not sure about the plugin file name, look for the plugins in the server's plugin directory.

If an attempt to install a plugin results in an error on Linux similar to that shown here, you will need to install `libimf`:

```
mysql> INSTALL PLUGIN rpl_semi_sync_master SONAME 'semisync_master.so';
ERROR 1126 (HY000): Can't open shared library
'/usr/local/mysql/lib/plugin/semisync_master.so' (errno: 22 libimf.so: cannot open
shared object file: No such file or directory)
```

You can obtain `libimf` from <http://dev.mysql.com/downloads/os-linux.html>.

To see which plugins are installed, use the `SHOW PLUGINS` statement, or query the `INFORMATION_SCHEMA.PLUGINS` table.

After a semisynchronous replication plugin has been installed, it is disabled by default. The plugins must be enabled both on the master side and the slave side to enable semisynchronous replication. If only one side is enabled, replication will be asynchronous.

To control whether an installed plugin is enabled, set the appropriate system variables. You can set these variables at runtime using `SET GLOBAL`, or at server startup on the command line or in an option file.

At runtime, these master-side system variables are available:

```
mysql> SET GLOBAL rpl_semi_sync_master_enabled = {0|1};
mysql> SET GLOBAL rpl_semi_sync_master_timeout = N;
```

On the slave side, this system variable is available:

```
mysql> SET GLOBAL rpl_semi_sync_slave_enabled = {0|1};
```

For `rpl_semi_sync_master_enabled` [536] or `rpl_semi_sync_slave_enabled` [537], the value should be 1 to enable semisynchronous replication or 0 to disable it. By default, these variables are set to 1.

For `rpl_semi_sync_master_timeout` [536], the value *N* is given in milliseconds. The default value is 10000 (10 seconds).

If you enable semisynchronous replication on a slave at runtime, you must also start the slave I/O thread (stopping it first if it is already running) to cause the slave to connect to the master and register as a semisynchronous slave:

```
mysql> STOP SLAVE IO_THREAD; START SLAVE IO_THREAD;
```

If the I/O thread is already running and you do not restart it, the slave continues to use asynchronous replication.

At server startup, the variables that control semisynchronous replication can be set as command-line options or in an option file. A setting listed in an option file takes effect each time the server starts. For example, you can set the variables in `my.cnf` files on the master and slave sides as follows.

On the master:

```
[mysqld]
rpl_semi_sync_master_enabled=1
rpl_semi_sync_master_timeout=1000 # 1 second
```

On each slave:

```
[mysqld]
rpl_semi_sync_slave_enabled=1
```

### 16.3.8.3. Semisynchronous Replication Monitoring

The plugins for the semisynchronous replication capability expose several system and status variables that you can examine to determine its configuration and operational state.

The system variable reflect how semisynchronous replication is configured. To check their values, use `SHOW VARIABLES`:

```
mysql> SHOW VARIABLES LIKE 'rpl_semi_sync%';
```

The status variables enable you to monitor the operation of semisynchronous replication. To check their values, use `SHOW STATUS`:

```
mysql> SHOW STATUS LIKE 'Rpl_semi_sync%';
```

When the master switches between asynchronous or semisynchronous replication due to commit-blocking timeout or a slave catching up, it sets the value of the `Rpl_semi_sync_master_status` [596] status variable appropriately. Automatic fallback from semisynchronous to asynchronous replication on the master means that it is possible for the `rpl_semi_sync_master_enabled` [536] system variable to have a value of 1 on the master side even when semisynchronous replication is in fact not operational at the moment. You can monitor the `Rpl_semi_sync_master_status` [596] status variable to determine whether the master currently is using asynchronous or semisynchronous replication.

To see how many semisynchronous slaves are connected, check `Rpl_semi_sync_master_clients` [596].

The number of commits that have been acknowledged successfully or unsuccessfully by slaves are indicated by the `Rpl_semi_sync_master_yes_tx` [597] and `Rpl_semi_sync_master_no_tx` [596] variables.

On the slave side, `Rpl_semi_sync_slave_status` [597] indicates whether semisynchronous replication currently is operational.

### 16.3.9. Delayed Replication

MySQL 5.6 supports delayed replication such that a slave server deliberately lags behind the master by at least a specified amount of time. The default delay is 0 seconds. Use the `MASTER_DELAY` option for `CHANGE MASTER TO` to set the delay to `N` seconds:

```
CHANGE MASTER TO MASTER_DELAY = N;
```

An event received from the master is not executed until at least `N` seconds later than its execution on the master. The exceptions are that there is no delay for format description events or log file rotation events, which affect only the internal state of the SQL thread.

Delayed replication can be used for several purposes:

- To protect against user mistakes on the master. A DBA can roll back a delayed slave to the time just before the disaster.
- To test how the system behaves when there is a lag. For example, in an application, a lag might be caused by a heavy load on the slave. However, it can be difficult to generate this load level. Delayed replication can simulate the lag without having to simulate the load. It can also be used to debug conditions related to a lagging slave.
- To inspect what the database looked like long ago, without having to reload a backup. For example, if the delay is one week and the DBA needs to see what the database looked like before the last few days' worth of development, the delayed slave can be inspected.

`START SLAVE` and `STOP SLAVE` take effect immediately and ignore any delay. `RESET SLAVE` resets the delay to 0.

`SHOW SLAVE STATUS` has three fields that provide information about the delay:

- `SQL_Delay`: A nonnegative integer indicating the number of seconds that the slave must lag the master.
- `SQL_Remaining_Delay`: When `Slave_SQL_Running_State` is `Waiting until MASTER_DELAY seconds after master executed event`, this field contains an integer indicating the number of seconds left of the delay. At other times, this field is `NULL`.
- `Slave_SQL_Running_State`: A string indicating the state of the SQL thread (analogous to `Slave_IO_State`). The value is identical to the `State` value of the SQL thread as displayed by `SHOW PROCESSLIST`.

When the slave SQL thread is waiting for the delay to elapse before executing an event, `SHOW PROCESSLIST` displays its `State` value as `Waiting until MASTER_DELAY seconds after master executed event`.

The `relay-log.info` file now contains the delay value, so the file format has changed. See [Section 16.2.2.2, “Slave Status Logs”](#). In particular, the first line of the file now indicates how many lines are in the file. If you downgrade a slave server to a version older than MySQL 5.6, the older server will not read the file correctly. To address this, modify the file in a text editor to delete the initial line containing the number of lines.

## 16.4. Replication Notes and Tips

### 16.4.1. Replication Features and Issues

The following sections provide information about what is supported and what is not in MySQL replication, and about specific issues and situations that may occur when replicating certain statements.

Statement-based replication depends on compatibility at the SQL level between the master and slave. In others, successful SBR requires that any SQL features used be supported by both the master and the slave servers. For example, if you use a feature on the master server that is available only in MySQL 5.6 (or later), you cannot replicate to a slave that uses MySQL 5.5 (or earlier).

Such incompatibilities also can occur within a release series when using pre-production releases of MySQL. For example, the `SLEEP()` [1280] function is available beginning with MySQL 5.0.12. If you use this function on the master, you cannot replicate to a slave that uses MySQL 5.0.11 or earlier.

For this reason, use Generally Available (GA) releases of MySQL for statement-based replication in a production setting, since we do not introduce new SQL statements or change their behavior within a given release series once that series reaches GA release status.

If you are planning to use statement-based replication between MySQL 5.6 and a previous MySQL release series, it is also a good idea to consult the edition of the *MySQL Reference Manual* corresponding to the earlier release series for information regarding the replication characteristics of that series.

With MySQL's statement-based replication, there may be issues with replicating stored routines or triggers. You can avoid these issues by using MySQL's row-based replication instead. For a detailed list of issues, see [Section 18.7, “Binary Logging of Stored Programs”](#). For more information about row-based logging and row-based replication, see [Section 5.2.4.1, “Binary Logging Formats”](#), and [Section 16.1.2, “Replication Formats”](#).

For additional information specific to replication and `InnoDB`, see [Section 14.2.2.6, “InnoDB and MySQL Replication”](#). For information relating to replication with MySQL Cluster, see [MySQL Cluster Replication](#).

#### 16.4.1.1. Replication and `AUTO_INCREMENT`

Statement-based replication of `AUTO_INCREMENT`, `LAST_INSERT_ID()` [1269], and `TIMESTAMP` values is done correctly, subject to the following exceptions:

- When using statement-based replication prior to MySQL 5.6.10, `AUTO_INCREMENT` columns in tables on the slave must match the same columns on the master; that is, `AUTO_INCREMENT` columns must be replicated to `AUTO_INCREMENT` columns. (Bug #12669186)
- A statement invoking a trigger or function that causes an update to an `AUTO_INCREMENT` column is not replicated correctly using statement-based replication. In MySQL 5.6, such statements are marked as unsafe. (Bug #45677)
- An `INSERT` into a table that has a composite primary key that includes an `AUTO_INCREMENT` column that is not the first column of this composite key is not safe for statement-based logging or replication. Beginning with MySQL 5.6.6, such statements are marked as unsafe. (Bug #11754117, Bug #45670)

This issue does not affect tables using the `InnoDB` storage engine, since an `InnoDB` table with an `AUTO_INCREMENT` column requires at least one key where the auto-increment column is the only or leftmost column.

- Adding an `AUTO_INCREMENT` column to a table with `ALTER TABLE` might not produce the same ordering of the rows on the slave and the master. This occurs because the order in which the rows are numbered depends on the specific storage engine used for the table and the order in which the rows were inserted. If it is important to have the same order on the master and slave, the rows must be ordered before assigning an `AUTO_INCREMENT` number. Assuming that you want to add an `AUTO_INCREMENT` column to a table `t1` that has columns `col1` and `col2`, the following statements produce a new table `t2` identical to `t1` but with an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 LIKE t1;
ALTER TABLE t2 ADD id INT AUTO_INCREMENT PRIMARY KEY;
INSERT INTO t2 SELECT * FROM t1 ORDER BY col1, col2;
```



### Important

To guarantee the same ordering on both master and slave, the `ORDER BY` clause must name *all* columns of `t1`.

The instructions just given are subject to the limitations of `CREATE TABLE ... LIKE`: Foreign key definitions are ignored, as are the `DATA DIRECTORY` and `INDEX DIRECTORY` table options. If a table definition includes any of those characteristics, create `t2` using a `CREATE TABLE` statement that is identical to the one used to create `t1`, but with the addition of the `AUTO_INCREMENT` column.

Regardless of the method used to create and populate the copy having the `AUTO_INCREMENT` column, the final step is to drop the original table and then rename the copy:

```
DROP t1;
ALTER TABLE t2 RENAME t1;
```

See also [Section C.5.7.1, “Problems with ALTER TABLE”](#).

## 16.4.1.2. Replication and Character Sets

The following applies to replication between MySQL servers that use different character sets:

- If the master uses MySQL 4.1, you must *always* use the same *global* character set and collation on the master and the slave, regardless of the slave MySQL version. (These are controlled by the `--character-set-server` [417] and `--collation-server` [417] options.) Otherwise, you may get duplicate-key errors on the slave, because a key that is unique in the master character set might not be unique in the slave character set. Note that this is not a cause for concern when master and slave are both MySQL 5.0 or later.
- If the master is older than MySQL 4.1.3, the character set of any client should never be made different from its global value because this character set change is not known to the slave. In other words, clients should not use `SET NAMES`, `SET CHARACTER SET`, and so forth. If both the master and the slave are 4.1.3 or newer, clients can freely set session values for character set variables because these settings are written to the binary log and so are known to the slave. That is, clients can use `SET NAMES` or `SET CHARACTER SET` or can set variables such as `collation_client` or `collation_server` [470]. However, clients are prevented from changing the *global* value of these variables; as stated previously, the master and slave must always have identical global character set values. This is true whether you are using statement-based or row-based replication.
- If the master has databases with a character set different from the global `character_set_server` [468] value, you should design your `CREATE TABLE` statements so that they do not implicitly rely on the database default character set. A good workaround is to state the character set and collation explicitly in `CREATE TABLE` statements.

## 16.4.1.3. Replication of CREATE ... IF NOT EXISTS Statements

MySQL applies these rules when various `CREATE ... IF NOT EXISTS` statements are replicated:

- Every `CREATE DATABASE IF NOT EXISTS` statement is replicated, whether or not the database already exists on the master.
- Similarly, every `CREATE TABLE IF NOT EXISTS` statement without a `SELECT` is replicated, whether or not the table already exists on the master. This includes `CREATE TABLE IF NOT EXISTS ... LIKE`. Replication of `CREATE TABLE IF NOT EXISTS ... SELECT` follows somewhat different rules; see [Section 16.4.1.4, “Replication of CREATE TABLE ... SELECT Statements”](#), for more information.

- `CREATE EVENT IF NOT EXISTS` is always replicated in MySQL 5.6, whether or not the event named in the statement already exists on the master.

See also Bug #45574.

#### 16.4.1.4. Replication of `CREATE TABLE ... SELECT` Statements

This section discusses how MySQL replicates `CREATE TABLE ... SELECT` statements.

MySQL 5.6 does not allow a `CREATE TABLE ... SELECT` statement to make any changes in tables other than the table that is created by the statement. This is a change in behavior from previous versions of MySQL, which permitted these statements to do so. This means that, when using statement-based replication between a MySQL 5.6 or later slave and a master running a previous version of MySQL, a `CREATE TABLE ... SELECT` statement causing changes in other tables on the master fails on the slave, causing replication to stop. To keep this from happening, you should use row-based replication, rewrite the offending statement before running it on the master, or upgrade the master to MySQL 5.6 (or later). (If you choose to upgrade the master, keep in mind that such a `CREATE TABLE ... SELECT` statement will fail following the upgrade unless it is rewritten to remove any side effects on other tables.) This is not an issue when using row-based replication, because the statement is logged as a `CREATE TABLE` statement with any changes to table data logged as row-insert events, rather than as the entire `CREATE TABLE ... SELECT`.

These behaviors are not dependent on MySQL version:

- `CREATE TABLE ... SELECT` always performs an implicit commit ([Section 13.3.3, “Statements That Cause an Implicit Commit”](#)).
- If destination table does not exist, logging occurs as follows. It does not matter whether `IF NOT EXISTS` is present.
  - `STATEMENT` or `MIXED` format: The statement is logged as written.
  - `ROW` format: The statement is logged as a `CREATE TABLE` statement followed by a series of insert-row events.
- If the statement fails, nothing is logged. This includes the case that the destination table exists and `IF NOT EXISTS` is not given.

When the destination table exists and `IF NOT EXISTS` is given, MySQL handles the statement in a version-dependent way.

In MySQL 5.1 before 5.1.51 and in MySQL 5.5 before 5.5.6 (this is the original behavior):

- `STATEMENT` or `MIXED` format: The statement is logged as written.
- `ROW` format: The statement is logged as a `CREATE TABLE` statement followed by a series of insert-row events.

In MySQL 5.1 as of 5.1.51:

- `STATEMENT` or `MIXED` format: The statement is logged as the equivalent pair of `CREATE TABLE` and `INSERT INTO ... SELECT` statements.
- `ROW` format: The statement is logged as a `CREATE TABLE` statement followed by a series of insert-row events.

In MySQL 5.5 as of 5.5.6:

- Nothing is inserted or logged.

These version dependencies arise due to a change in MySQL 5.5.6 in handling of `CREATE TABLE ... SELECT` not to insert rows if the destination table already exists, and a change made in

MySQL 5.1.51 to preserve forward compatibility in replication of such statements from a 5.1 master to a 5.5 slave. For details, see [Section 13.1.14.1](#), “`CREATE TABLE ... SELECT` Syntax”.

#### 16.4.1.5. Replication of `CREATE SERVER`, `ALTER SERVER`, and `DROP SERVER`

In MySQL 5.6, the statements `CREATE SERVER`, `ALTER SERVER`, and `DROP SERVER` are not written to the binary log, regardless of the binary logging format that is in use.

#### 16.4.1.6. Replication of `CURRENT_USER()`

The following statements support use of the `CURRENT_USER()` [\[1267\]](#) function to take the place of the name of (and, possibly, the host for) an affected user or a definer; in such cases, `CURRENT_USER()` [\[1267\]](#) is expanded where and as needed:

- `DROP USER`
- `RENAME USER`
- `GRANT`
- `REVOKE`
- `CREATE FUNCTION`
- `CREATE PROCEDURE`
- `CREATE TRIGGER`
- `CREATE EVENT`
- `CREATE VIEW`
- `ALTER EVENT`
- `ALTER VIEW`
- `SET PASSWORD`

When `CURRENT_USER()` [\[1267\]](#) or `CURRENT_USER` [\[1267\]](#) is used as the definer in any of the statements `CREATE FUNCTION`, `CREATE PROCEDURE`, `CREATE TRIGGER`, `CREATE EVENT`, `CREATE VIEW`, or `ALTER VIEW` when binary logging is enabled, the function reference is expanded before it is written to the binary log, so that the statement refers to the same user on both the master and the slave when the statement is replicated. `CURRENT_USER()` [\[1267\]](#) or `CURRENT_USER` [\[1267\]](#) is also expanded prior to being written to the binary log when used in `DROP USER`, `RENAME USER`, `GRANT`, `REVOKE`, or `ALTER EVENT`.

#### 16.4.1.7. Replication of `DROP ... IF EXISTS` Statements

The `DROP DATABASE IF EXISTS`, `DROP TABLE IF EXISTS`, and `DROP VIEW IF EXISTS` statements are always replicated, even if the database, table, or view to be dropped does not exist on the master. This is to ensure that the object to be dropped no longer exists on either the master or the slave, once the slave has caught up with the master.

`DROP ... IF EXISTS` statements for stored programs (stored procedures and functions, triggers, and events) are also replicated, even if the stored program to be dropped does not exist on the master.

#### 16.4.1.8. Replication with Differing Table Definitions on Master and Slave

Source and target tables for replication do not have to be identical. A table on the master can have more or fewer columns than the slave's copy of the table. In addition, corresponding table columns on the master and the slave can use different data types, subject to certain conditions.



In all cases where the source and target tables do not have identical definitions, the database and table names must be the same on both the master and the slave. Additional conditions are discussed, with examples, in the following two sections.

#### 16.4.1.8.1. Replication with More Columns on Master or Slave

You can replicate a table from the master to the slave such that the master and slave copies of the table have differing numbers of columns, subject to the following conditions:

- Columns common to both versions of the table must be defined in the same order on the master and the slave.

(This is true even if both tables have the same number of columns.)

- Columns common to both versions of the table must be defined before any additional columns.

This means that executing an `ALTER TABLE` statement on the slave where a new column is inserted into the table within the range of columns common to both tables causes replication to fail, as shown in the following example:

Suppose that a table `t`, existing on the master and the slave, is defined by the following `CREATE TABLE` statement:

```
CREATE TABLE t (
  c1 INT,
  c2 INT,
  c3 INT
);
```

Suppose that the `ALTER TABLE` statement shown here is executed on the slave:

```
ALTER TABLE t ADD COLUMN cnew1 INT AFTER c3;
```

The previous `ALTER TABLE` is permitted on the slave because the columns `c1`, `c2`, and `c3` that are common to both versions of table `t` remain grouped together in both versions of the table, before any columns that differ.

However, the following `ALTER TABLE` statement cannot be executed on the slave without causing replication to break:

```
ALTER TABLE t ADD COLUMN cnew2 INT AFTER c3;
```

Replication fails after execution on the slave of the `ALTER TABLE` statement just shown, because the new column `cnew2` comes between columns common to both versions of `t`.

- Each “extra” column in the version of the table having more columns must have a default value.

A column's default value is determined by a number of factors, including its type, whether it is defined with a `DEFAULT` option, whether it is declared as `NULL`, and the server SQL mode in effect at the time of its creation; for more information, see [Section 11.5, “Data Type Default Values”](#)).

In addition, when the slave's copy of the table has more columns than the master's copy, each column common to the tables must use the same data type in both tables.

**Examples.** The following examples illustrate some valid and invalid table definitions:

**More columns on the master.** The following table definitions are valid and replicate correctly:

```
master> CREATE TABLE t1 (c1 INT, c2 INT, c3 INT);
slave> CREATE TABLE t1 (c1 INT, c2 INT);
```

The following table definitions would raise Error 1532 (`ER_BINLOG_ROW_RBR_TO_SBR [3695]`) because the definitions of the columns common to both versions of the table are in a different order on the slave than they are on the master:

```
master> CREATE TABLE t1 (c1 INT, c2 INT, c3 INT);
slave> CREATE TABLE t1 (c2 INT, c1 INT);
```

The following table definitions would also raise Error 1532 because the definition of the extra column on the master appears before the definitions of the columns common to both versions of the table:

```
master> CREATE TABLE t1 (c3 INT, c1 INT, c2 INT);
slave> CREATE TABLE t1 (c1 INT, c2 INT);
```

**More columns on the slave.** The following table definitions are valid and replicate correctly:

```
master> CREATE TABLE t1 (c1 INT, c2 INT);
slave> CREATE TABLE t1 (c1 INT, c2 INT, c3 INT);
```

The following definitions raise Error 1532 because the columns common to both versions of the table are not defined in the same order on both the master and the slave:

```
master> CREATE TABLE t1 (c1 INT, c2 INT);
slave> CREATE TABLE t1 (c2 INT, c1 INT, c3 INT);
```

The following table definitions also raise Error 1532 because the definition for the extra column in the slave's version of the table appears before the definitions for the columns which are common to both versions of the table:

```
master> CREATE TABLE t1 (c1 INT, c2 INT);
slave> CREATE TABLE t1 (c3 INT, c1 INT, c2 INT);
```

The following table definitions fail because the slave's version of the table has additional columns compared to the master's version, and the two versions of the table use different data types for the common column `c2`:

```
master> CREATE TABLE t1 (c1 INT, c2 BIGINT);
slave> CREATE TABLE t1 (c1 INT, c2 INT, c3 INT);
```

#### 16.4.1.8.2. Replication of Columns Having Different Data Types

Corresponding columns on the master's and the slave's copies of the same table ideally should have the same data type. However, beginning with MySQL 5.1.21, this is not always strictly enforced, as long as certain conditions are met.

All other things being equal, it is always possible to replicate from a column of a given data type to another column of the same type and same size or width, where applicable, or larger. For example, you can replicate from a `CHAR(10)` column to another `CHAR(10)`, or from a `CHAR(10)` column to a `CHAR(25)` column without any problems. In certain cases, it is also possible to replicate from a column having one data type (on the master) to a column having a different data type (on the slave); when the data type of the master's version of the column is promoted to a type that is the same size or larger on the slave, this is known as *attribute promotion*.

Attribute promotion can be used with both statement-based and row-based replication, and is not dependent on the storage engine used by either the master or the slave. However, the choice of logging format does have an effect on the type conversions that are permitted; the particulars are discussed later in this section.



#### Important

Whether you use statement-based or row-based replication, the slave's copy of the table cannot contain more columns than the master's copy if you wish to employ attribute promotion.

**Statement-based replication.** When using statement-based replication, a simple rule of thumb to follow is, "If the statement run on the master would also execute successfully on the slave, it should also replicate successfully". In other words, if the statement uses a value that is compatible with the

type of a given column on the slave, the statement can be replicated. For example, you can insert any value that fits in a `TINYINT` column into a `BIGINT` column as well; it follows that, even if you change the type of a `TINYINT` column in the slave's copy of a table to `BIGINT`, any insert into that column on the master that succeeds should also succeed on the slave, since it is impossible to have a legal `TINYINT` value that is large enough to exceed a `BIGINT` column.

Prior to MySQL 5.6.10, when using statement-based replication, `AUTO_INCREMENT` columns were required to be the same on both the master and the slave; otherwise, updates could be applied to the wrong table on the slave. (Bug #12669186)

**Row-based replication: attribute promotion and demotion.** Formerly, due to the fact that in row-based replication changes rather than statements are replicated, and that these changes are transmitted using formats that do not always map directly to MySQL server column data types, you could not replicate between different subtypes of the same general type (for example, from `TINYINT` to `BIGINT`, both `INT` subtypes). However, beginning with MySQL 5.5.3, MySQL Replication supports attribute promotion and demotion between smaller data types and larger types. It is also possible to specify whether or not to permit lossy (truncated) or non-lossy conversions of demoted column values, as explained later in this section.

**Lossy and non-lossy conversions.** In the event that the target type cannot represent the value being inserted, a decision must be made on how to handle the conversion. If we permit the conversion but truncate (or otherwise modify) the source value to achieve a "fit" in the target column, we make what is known as a *lossy conversion*. A conversion which does not require truncation or similar modifications to fit the source column value in the target column is a *non-lossy conversion*.

**Type conversion modes (`slave_type_conversions` variable).** The setting of the `slave_type_conversions` global server variable controls the type conversion mode used on the slave. This variable takes a set of values from the following table, which shows the effects of each mode on the slave's type-conversion behavior:

Mode	Effect
<code>ALL_LOSSY</code>	<p>In this mode, type conversions that would mean loss of information are permitted.</p> <p>This does not imply that non-lossy conversions are permitted, merely that only cases requiring either lossy conversions or no conversion at all are permitted; for example, enabling <i>only</i> this mode permits an <code>INT</code> column to be converted to <code>TINYINT</code> (a lossy conversion), but not a <code>TINYINT</code> column to an <code>INT</code> column (non-lossy). Attempting the latter conversion in this case would cause replication to stop with an error on the slave.</p>
<code>ALL_NON_LOSSY</code>	<p>This mode permits conversions that do not require truncation or other special handling of the source value; that is, it permits conversions where the target type has a wider range than the source type.</p> <p>Setting this mode has no bearing on whether lossy conversions are permitted; this is controlled with the <code>ALL_LOSSY</code> mode. If only <code>ALL_NON_LOSSY</code> is set, but not <code>ALL_LOSSY</code>, then attempting a conversion that would result in the loss of data (such as <code>INT</code> to <code>TINYINT</code>, or <code>CHAR(25)</code> to <code>VARCHAR(20)</code>) causes the slave to stop with an error.</p>
<code>ALL_LOSSY, ALL_NON_LOSSY</code>	When this mode is set, all supported type conversions are permitted, whether or not they are lossy conversions.
[empty]	When <code>slave_type_conversions</code> is not set, no attribute promotion or demotion is permitted; this means that all columns in the source and target tables must be of the same types.

Mode	Effect
	This mode is the default.

Changing the type conversion mode requires restarting the slave with the new `slave_type_conversions` setting.

**Supported conversions.** Supported conversions between different but similar data types are shown in the following list:

- Between any of the integer types `TINYINT`, `SMALLINT`, `MEDIUMINT`, `INT`, and `BIGINT`.

This includes conversions between the signed and unsigned versions of these types.

Lossy conversions are made by truncating the source value to the maximum (or minimum) permitted by the target column. For insuring non-lossy conversions when going from unsigned to signed types, the target column must be large enough to accommodate the range of values in the source column. For example, you can demote `TINYINT UNSIGNED` non-lossily to `SMALLINT`, but not to `TINYINT`.

- Between any of the decimal types `DECIMAL`, `FLOAT`, `DOUBLE`, and `NUMERIC`.

`FLOAT` to `DOUBLE` is a non-lossy conversion; `DOUBLE` to `FLOAT` can only be handled lossily. A conversion from `DECIMAL(M,D)` to `DECIMAL(M',D')` where  $M' \Rightarrow M$  and  $D' \Rightarrow D$  is non-lossy; for any case where  $M' < M$ ,  $D' < D$ , or both, only a lossy conversion can be made.

For any of the decimal types, if a value to be stored cannot be fit in the target type, the value is rounded down according to the rounding rules defined for the server elsewhere in the documentation. See [Section 12.19.4, “Rounding Behavior”](#), for information about how this is done for decimal types.

- Between any of the string types `CHAR`, `VARCHAR`, and `TEXT`, including conversions between different widths.

Conversion of a `CHAR`, `VARCHAR`, or `TEXT` to a `CHAR`, `VARCHAR`, or `TEXT` column the same size or larger is never lossy. Lossy conversion is handled by inserting only the first  $N$  characters of the string on the slave, where  $N$  is the width of the target column.



#### Important

Replication between columns using different character sets is not supported.

- Between any of the binary data types `BINARY`, `VARBINARY`, and `BLOB`, including conversions between different widths.

Conversion of a `BINARY`, `VARBINARY`, or `BLOB` to a `BINARY`, `VARBINARY`, or `BLOB` column the same size or larger is never lossy. Lossy conversion is handled by inserting only the first  $N$  bytes of the string on the slave, where  $N$  is the width of the target column.

- Between any 2 `BIT` columns of any 2 sizes.

When inserting a value from a `BIT(M)` column into a `BIT(M')` column, where  $M' > M$ , the most significant bits of the `BIT(M')` columns are cleared (set to zero) and the  $M$  bits of the `BIT(M)` value are set as the least significant bits of the `BIT(M')` column.

When inserting a value from a source `BIT(M)` column into a target `BIT(M')` column, where  $M' < M$ , the maximum possible value for the `BIT(M')` column is assigned; in other words, an “all-set” value is assigned to the target column.

Conversions between types not in the previous list are not permitted.

**Replication type conversions in MySQL 5.5.3 and earlier.** Prior to MySQL 5.5.3, with row-based binary logging, you could not replicate between different `INT` subtypes, such as from `TINYINT` to

`BIGINT`, because changes to columns of these types were represented differently from one another in the binary log when using row-based logging. (However, you could replicate from `BLOB` to `TEXT` using row-based replication because changes to `BLOB` and `TEXT` columns were represented using the same format in the binary log.)

Supported conversions for attribute promotion when using row-based replication prior to MySQL 5.5.3 are shown in the following table:

From (Master)	To (Slave)
<code>BINARY</code>	<code>CHAR</code>
<code>BLOB</code>	<code>TEXT</code>
<code>CHAR</code>	<code>BINARY</code>
<code>DECIMAL</code>	<code>NUMERIC</code>
<code>NUMERIC</code>	<code>DECIMAL</code>
<code>TEXT</code>	<code>BLOB</code>
<code>VARBINARY</code>	<code>VARCHAR</code>
<code>VARCHAR</code>	<code>VARBINARY</code>



**Note**

In all cases, the size or width of the column on the slave must be equal to or greater than that of the column on the master. For example, you could replicate from a `CHAR(10)` column on the master to a column that used `BINARY(10)` or `BINARY(25)` on the slave, but you could not replicate from a `CHAR(10)` column on the master to `BINARY(5)` column on the slave.

Any unique index (including primary keys) having a prefix must use a prefix of the same length on both master and slave; in such cases, differing prefix lengths are disallowed. It is possible to use a nonunique index whose prefix length differs between master and slave, but this can cause serious performance issues, particularly when the prefix used on the master is longer. This is due to the fact that 2 unique prefixes of a given length may no longer be unique at a shorter length; for example, the words *catalogue* and *catamount* have the 5-character prefixes *catal* and *catam*, respectively, but share the same 4-character prefix (*cata*). This can lead to queries that use such indexes executing less efficiently on the slave, when a shorter prefix is employed in the slave's definition of the same index than on the master.

For `DECIMAL` and `NUMERIC` columns, both the mantissa (*M*) and the number of decimals (*D*) must be the same size or larger on the slave as compared with the master. For example, replication from a `NUMERIC(5,4)` to a `DECIMAL(6,4)` worked, but not from a `NUMERIC(5,4)` to a `DECIMAL(5,3)`.

Prior to MySQL 5.5.3, MySQL replication did not support attribute promotion of any of the following data types to or from any other data type when using row-based replication:

- `INT` (including `TINYINT`, `SMALLINT`, `MEDIUMINT`, `BIGINT`).

Promotion between `INT` subtypes—for example, from `SMALLINT` to `BIGINT`—was also not supported prior to MySQL 5.5.3.

- `SET` or `ENUM`.
- `FLOAT` or `DOUBLE`.
- All of the data types relating to dates, times, or both: `DATE`, `TIME`, `DATETIME`, `TIMESTAMP`, and `YEAR`.

### 16.4.1.9. Replication and `DIRECTORY` Table Options

If a `DATA DIRECTORY` or `INDEX DIRECTORY` table option is used in a `CREATE TABLE` statement on the master server, the table option is also used on the slave. This can cause problems if no corresponding directory exists in the slave host file system or if it exists but is not accessible to the slave server. This can be overridden by using the `NO_DIR_IN_CREATE` [604] server SQL mode on the slave, which causes the slave to ignore the `DATA DIRECTORY` and `INDEX DIRECTORY` table options when replicating `CREATE TABLE` statements. The result is that `MyISAM` data and index files are created in the table's database directory.

For more information, see [Section 5.1.7, “Server SQL Modes”](#).

### 16.4.1.10. Replication of Invoked Features

Replication of invoked features such as user-defined functions (UDFs) and stored programs (stored procedures and functions, triggers, and events) provides the following characteristics:

- The effects of the feature are always replicated.
- The following statements are replicated using statement-based replication:
  - `CREATE EVENT`
  - `ALTER EVENT`
  - `DROP EVENT`
  - `CREATE PROCEDURE`
  - `DROP PROCEDURE`
  - `CREATE FUNCTION`
  - `DROP FUNCTION`
  - `CREATE TRIGGER`
  - `DROP TRIGGER`

However, the *effects* of features created, modified, or dropped using these statements are replicated using row-based replication.



#### Note

Attempting to replicate invoked features using statement-based replication produces the warning `Statement is not safe to log in statement format`. For example, trying to replicate a UDF with statement-based replication generates this warning because it currently cannot be determined by the MySQL server whether the UDF is deterministic. If you are absolutely certain that the invoked feature's effects are deterministic, you can safely disregard such warnings.

- In the case of `CREATE EVENT` and `ALTER EVENT`:
  - The status of the event is set to `SLAVESIDE_DISABLED` on the slave regardless of the state specified (this does not apply to `DROP EVENT`).
  - The master on which the event was created is identified on the slave by its server ID. The `ORIGINATOR` column in `INFORMATION_SCHEMA.EVENTS` and the `originator` column in `mysql.event` store this information. See [Section 19.7, “The `INFORMATION\_SCHEMA.EVENTS` Table”](#), and [Section 13.7.5.19, “`SHOW EVENTS` Syntax”](#), for more information.

- The feature implementation resides on the slave in a renewable state so that if the master fails, the slave can be used as the master without loss of event processing.

To determine whether there are any scheduled events on a MySQL server that were created on a different server (that was acting as a replication master), query the `INFORMATION_SCHEMA.EVENTS` table in a manner similar to what is shown here:

```
SELECT EVENT_SCHEMA, EVENT_NAME
FROM INFORMATION_SCHEMA.EVENTS
WHERE STATUS = 'SLAVESIDE_DISABLED';
```

Alternatively, you can use the `SHOW EVENTS` statement, like this:

```
SHOW EVENTS
WHERE STATUS = 'SLAVESIDE_DISABLED';
```

When promoting a replication slave having such events to a replication master, you must enable each event using `ALTER EVENT event_name ENABLED`, where `event_name` is the name of the event.

If more than one master was involved in creating events on this slave, and you wish to identify events that were created only on a given master having the server ID `master_id`, modify the previous query on the `EVENTS` table to include the `ORIGINATOR` column, as shown here:

```
SELECT EVENT_SCHEMA, EVENT_NAME, ORIGINATOR
FROM INFORMATION_SCHEMA.EVENTS
WHERE STATUS = 'SLAVESIDE_DISABLED'
AND ORIGINATOR = 'master_id'
```

You can employ `ORIGINATOR` with the `SHOW EVENTS` statement in a similar fashion:

```
SHOW EVENTS
WHERE STATUS = 'SLAVESIDE_DISABLED'
AND ORIGINATOR = 'master_id'
```

Before enabling events that were replicated from the master, you should disable the MySQL Event Scheduler on the slave (using a statement such as `SET GLOBAL event_scheduler = OFF;`), run any necessary `ALTER EVENT` statements, restart the server, then re-enable the Event Scheduler on the slave afterward (using a statement such as `SET GLOBAL event_scheduler = ON;`)-

If you later demote the new master back to being a replication slave, you must disable manually all events enabled by the `ALTER EVENT` statements. You can do this by storing in a separate table the event names from the `SELECT` statement shown previously, or using `ALTER EVENT` statements to rename the events with a common prefix such as `replicated_` to identify them.

If you rename the events, then when demoting this server back to being a replication slave, you can identify the events by querying the `EVENTS` table, as shown here:

```
SELECT CONCAT(EVENT_SCHEMA, '.', EVENT_NAME) AS 'Db.Event'
FROM INFORMATION_SCHEMA.EVENTS
WHERE INSTR(EVENT_NAME, 'replicated_') = 1;
```

### 16.4.1.11. Replication and Floating-Point Values

With statement-based replication, values are converted from decimal to binary. Because conversions between decimal and binary representations of them may be approximate, comparisons involving floating-point values are inexact. This is true for operations that use floating-point values explicitly, or that use values that are converted to floating-point implicitly. Comparisons of floating-point values might yield different results on master and slave servers due to differences in computer architecture, the compiler used to build MySQL, and so forth. See [Section 12.2, “Type Conversion in Expression Evaluation”](#), and [Section C.5.5.8, “Problems with Floating-Point Values”](#).

### 16.4.1.12. Replication and Fractional Seconds Support

MySQL 5.6.4 and up permits fractional seconds for `TIME`, `DATETIME`, and `TIMESTAMP` values, with up to microseconds (6 digits) precision. See [Section 11.3.6, “Fractional Seconds in Time Values”](#).



There may be problems replicating from a master server that understands fractional seconds to an older slave that does not:

- For `CREATE TABLE` statements containing columns that have an *fsp* (fractional seconds precision) value greater than 0, replication will fail due to parser errors.
- Statements that use temporal data types with an *fsp* value of 0 will work for with statement-based logging but not row-based logging. In the latter case, the data types have binary formats and type codes on the master that differ from those on the slave.
- Some expression results will differ on master and slave. Examples: On the master, the `timestamp` system variable returns a value that includes a microseconds fractional part; on the slave, it returns an integer. On the master, functions that return a result that includes the current time (such as `CURTIME()` [1213], `SYSDATE()` [1223], or `UTC_TIMESTAMP()` [1227]) interpret an argument as an *fsp* value and the return value includes a fractional seconds part of that many digits. On the slave, these functions permit an argument but ignore it.

### 16.4.1.13. Replication and `FLUSH`

Some forms of the `FLUSH` statement are not logged because they could cause problems if replicated to a slave: `FLUSH LOGS`, `FLUSH MASTER`, `FLUSH SLAVE`, and `FLUSH TABLES WITH READ LOCK`. For a syntax example, see [Section 13.7.6.3, “FLUSH Syntax”](#). The `FLUSH TABLES`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements are written to the binary log and thus replicated to slaves. This is not normally a problem because these statements do not modify table data.

However, this behavior can cause difficulties under certain circumstances. If you replicate the privilege tables in the `mysql` database and update those tables directly without using `GRANT`, you must issue a `FLUSH PRIVILEGES` on the slaves to put the new privileges into effect. In addition, if you use `FLUSH TABLES` when renaming a `MyISAM` table that is part of a `MERGE` table, you must issue `FLUSH TABLES` manually on the slaves. These statements are written to the binary log unless you specify `NO_WRITE_TO_BINLOG` or its alias `LOCAL`.

### 16.4.1.14. Replication and System Functions

Certain functions do not replicate well under some conditions:

- The `USER()` [1273], `CURRENT_USER()` [1267] (or `CURRENT_USER` [1267]), `UUID()` [1280], `VERSION()` [1273], and `LOAD_FILE()` [1181] functions are replicated without change and thus do not work reliably on the slave unless row-based replication is enabled. (See [Section 16.1.2, “Replication Formats”](#).)

`USER()` [1273] and `CURRENT_USER()` [1267] are automatically replicated using row-based replication when using `MIXED` mode, and generate a warning in `STATEMENT` mode. (See also [Section 16.4.1.6, “Replication of `CURRENT\_USER\(\)`”](#).) This is also true for `VERSION()` [1273] and `RAND()` [1206].

- For `NOW()` [1220], the binary log includes the timestamp. This means that the value *as returned by the call to this function on the master* is replicated to the slave. This can lead to a possibly unexpected result when replicating between MySQL servers in different time zones. Suppose that the master is located in New York, the slave is located in Stockholm, and both servers are using local time. Suppose further that, on the master, you create a table `mytable`, perform an `INSERT` statement on this table, and then select from the table, as shown here:

```
mysql> CREATE TABLE mytable (mycol TEXT);
Query OK, 0 rows affected (0.06 sec)

mysql> INSERT INTO mytable VALUES ( NOW() );
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM mytable;
+-----+
| mycol |
```

```
+-----+
| 2009-09-01 12:00:00 |
+-----+
1 row in set (0.00 sec)
```

Local time in Stockholm is 6 hours later than in New York; so, if you issue `SELECT NOW()` on the slave at that exact same instant, the value `2009-09-01 18:00:00` is returned. For this reason, if you select from the slave's copy of `mytable` after the `CREATE TABLE` and `INSERT` statements just shown have been replicated, you might expect `mycol` to contain the value `2009-09-01 18:00:00`. However, this is not the case; when you select from the slave's copy of `mytable`, you obtain exactly the same result as on the master:

```
mysql> SELECT * FROM mytable;
+-----+
| mycol |
+-----+
| 2009-09-01 12:00:00 |
+-----+
1 row in set (0.00 sec)
```

Unlike `NOW()` [1220], the `SYSDATE()` [1223] function is not replication-safe because it is not affected by `SET TIMESTAMP` statements in the binary log and is nondeterministic if statement-based logging is used. This is not a problem if row-based logging is used.

An alternative is to use the `--sysdate-is-now` [444] option to cause `SYSDATE()` [1223] to be an alias for `NOW()` [1220]. This must be done on the master and the slave to work correctly. In such cases, a warning is still issued by this function, but can safely be ignored as long as `--sysdate-is-now` [444] is used on both the master and the slave.

Beginning with MySQL 5.5.1, `SYSDATE()` [1223] is automatically replicated using row-based replication when using `MIXED` mode, and generates a warning in `STATEMENT` mode. (Bug #47995)

See also [Section 16.4.1.28, "Replication and Time Zones"](#).

- *The following restriction applies to statement-based replication only, not to row-based replication.* The `GET_LOCK()` [1275], `RELEASE_LOCK()` [1280], `IS_FREE_LOCK()` [1278], and `IS_USED_LOCK()` [1279] functions that handle user-level locks are replicated without the slave knowing the concurrency context on the master. Therefore, these functions should not be used to insert into a master table because the content on the slave would differ. For example, do not issue a statement such as `INSERT INTO mytable VALUES(GET_LOCK(...))`.

Beginning with MySQL 5.5.1, these functions are automatically replicated using row-based replication when using `MIXED` mode, and generate a warning in `STATEMENT` mode. (Bug #47995)

As a workaround for the preceding limitations when statement-based replication is in effect, you can use the strategy of saving the problematic function result in a user variable and referring to the variable in a later statement. For example, the following single-row `INSERT` is problematic due to the reference to the `UUID()` [1280] function:

```
INSERT INTO t VALUES(UUID());
```

To work around the problem, do this instead:

```
SET @my_uuid = UUID();
INSERT INTO t VALUES(@my_uuid);
```

That sequence of statements replicates because the value of `@my_uuid` is stored in the binary log as a user-variable event prior to the `INSERT` statement and is available for use in the `INSERT`.

The same idea applies to multiple-row inserts, but is more cumbersome to use. For a two-row insert, you can do this:

```
SET @my_uuid1 = UUID(); @my_uuid2 = UUID();
INSERT INTO t VALUES(@my_uuid1),(@my_uuid2);
```

However, if the number of rows is large or unknown, the workaround is difficult or impracticable. For example, you cannot convert the following statement to one in which a given individual user variable is associated with each row:

```
INSERT INTO t2 SELECT UUID(), * FROM t1;
```

Within a stored function, `RAND()` [1206] replicates correctly as long as it is invoked only once during the execution of the function. (You can consider the function execution timestamp and random number seed as implicit inputs that are identical on the master and slave.)

The `FOUND_ROWS()` [1268] and `ROW_COUNT()` [1272] functions are not replicated reliably using statement-based replication. A workaround is to store the result of the function call in a user variable, and then use that in the `INSERT` statement. For example, if you wish to store the result in a table named `mytable`, you might normally do so like this:

```
SELECT SQL_CALC_FOUND_ROWS FROM mytable LIMIT 1;
INSERT INTO mytable VALUES( FOUND_ROWS() );
```

However, if you are replicating `mytable`, you should use `SELECT ... INTO`, and then store the variable in the table, like this:

```
SELECT SQL_CALC_FOUND_ROWS INTO @found_rows FROM mytable LIMIT 1;
INSERT INTO mytable VALUES(@found_rows);
```

In this way, the user variable is replicated as part of the context, and applied on the slave correctly.

These functions are automatically replicated using row-based replication when using `MIXED` mode, and generate a warning in `STATEMENT` mode. (Bug #12092, Bug #30244)

#### 16.4.1.15. Replication and `LIMIT`

Statement-based replication of `LIMIT` clauses in `DELETE`, `UPDATE`, and `INSERT ... SELECT` statements is unsafe since the order of the rows affected is not defined. (Such statements can be replicated correctly with statement-based replication only if they also contain an `ORDER BY` clause.) When such a statement is encountered:

- When using `STATEMENT` mode, a warning that the statement is not safe for statement-based replication is now issued.

Currently, when using `STATEMENT` mode, warnings are issued for DML statements containing `LIMIT` even when they also have an `ORDER BY` clause (and so are made deterministic). This is a known issue. (Bug #42851)

- When using `MIXED` mode, the statement is now automatically replicated using row-based mode.

#### 16.4.1.16. Replication and `LOAD DATA INFILE`

The `LOAD DATA INFILE` statement was not always replicated correctly to a slave running MySQL 5.5.0 or earlier from a master running MySQL 4.0 or earlier. When using statement-based replication, the `LOAD DATA INFILE` statement `CONCURRENT` option was not replicated. This issue was fixed in MySQL 5.5.0. This issue does not have any impact on `CONCURRENT` option handling when using row-based replication in MySQL 5.1 or later. (Bug #34628)

As of MySQL 5.6, `LOAD DATA INFILE` is considered unsafe (see [Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”](#)). It causes a warning when using statement-based logging format, and is logged using row-based format when using mixed-format logging.

#### 16.4.1.17. Replication and `REPAIR TABLE`

When used on a corrupted or otherwise damaged table, it is possible for the `REPAIR TABLE` statement to delete rows that cannot be recovered. However, any such modifications of table data performed

by this statement are not replicated, which can cause master and slave to lose synchronization. For this reason, in the event that a table on the master becomes damaged and you use `REPAIR TABLE` to repair it, you should first stop replication (if it is still running) before using `REPAIR TABLE`, then afterward compare the master's and slave's copies of the table and be prepared to correct any discrepancies manually, before restarting replication.

#### 16.4.1.18. Replication and Master or Slave Shutdowns

It is safe to shut down a master server and restart it later. When a slave loses its connection to the master, the slave tries to reconnect immediately and retries periodically if that fails. The default is to retry every 60 seconds. This may be changed with the `CHANGE MASTER TO` statement. A slave also is able to deal with network connectivity outages. However, the slave notices the network outage only after receiving no data from the master for `slave_net_timeout` [1993] seconds. If your outages are short, you may want to decrease `slave_net_timeout` [1993]. See Section 5.1.4, “Server System Variables”.

An unclean shutdown (for example, a crash) on the master side can result in the master binary log having a final position less than the most recent position read by the slave, due to the master binary log file not being flushed. This can cause the slave not to be able to replicate when the master comes back up. Setting `sync_binlog=1` [2017] in the master `my.cnf` file helps to minimize this problem because it causes the master to flush its binary log more frequently.

Shutting down a slave cleanly is safe because it keeps track of where it left off. However, be careful that the slave does not have temporary tables open; see Section 16.4.1.21, “Replication and Temporary Tables”. Unclean shutdowns might produce problems, especially if the disk cache was not flushed to disk before the problem occurred:

- For transactions, the slave commits and then updates `relay-log.info`. If a crash occurs between these two operations, relay log processing will have proceeded further than the information file indicates and the slave will re-execute the events from the last transaction in the relay log after it has been restarted.
- A similar problem can occur if the slave updates `relay-log.info` but the server host crashes before the write has been flushed to disk. To minimize the chance of this occurring, set `sync_relay_log_info=1` [1999] in the slave `my.cnf` file. The default value of `sync_relay_log_info` [1999] is 0, which does not cause writes to be forced to disk; the server relies on the operating system to flush the file from time to time.

The fault tolerance of your system for these types of problems is greatly increased if you have a good uninterruptible power supply.

#### 16.4.1.19. Replication and `max_allowed_packet`

`max_allowed_packet` [502] sets an upper limit on the size of any single message between the MySQL server and clients, including replication slaves. If you are replicating large column values (such as might be found in `TEXT` or `BLOB` columns) and `max_allowed_packet` [502] is too small on the master, the master fails with an error, and the slave shuts down the I/O thread. If `max_allowed_packet` [502] is too small on the slave, this also causes the slave to stop the I/O thread.

Row-based replication currently sends all columns and column values for updated rows from the master to the slave, including values of columns that were not actually changed by the update. This means that, when you are replicating large column values using row-based replication, you must take care to set `max_allowed_packet` [502] large enough to accommodate the largest row in any table to be replicated, even if you are replicating updates only, or you are inserting only relatively small values.

#### 16.4.1.20. Replication and `MEMORY` Tables

When a master server shuts down and restarts, its `MEMORY` tables become empty. To replicate this effect to slaves, the first time that the master uses a given `MEMORY` table after startup, it logs an event

that notifies slaves that the table must to be emptied by writing a `DELETE` statement for that table to the binary log.

When a slave server shuts down and restarts, its `MEMORY` tables become empty. This causes the slave to be out of synchrony with the master and may lead to other failures or cause the slave to stop:

- Row-format updates and deletes received from the master may fail with `Can't find record in 'memory_table'`.
- Statements such as `INSERT INTO ... SELECT FROM memory_table` may insert a different set of rows on the master and slave.

The safe way to restart a slave that is replicating `MEMORY` tables is to first drop or delete all rows from the `MEMORY` tables on the master and wait until those changes have replicated to the slave. Then it is safe to restart the slave.

An alternative restart method may apply in some cases. When `binlog_format=ROW` [2010], you can prevent the slave from stopping if you set `slave_exec_mode=IDEMPOTENT` [1992] before you start the slave again. This allows the slave to continue to replicate, but its `MEMORY` tables will still be different from those on the master. This can be okay if the application logic is such that the contents of `MEMORY` tables can be safely lost (for example, if the `MEMORY` tables are used for caching). `slave_exec_mode=IDEMPOTENT` [1992] applies globally to all tables, so it may hide other replication errors in non-`MEMORY` tables.

The size of `MEMORY` tables is limited by the value of the `max_heap_table_size` [504] system variable, which is not replicated (see Section 16.4.1.32, “Replication and Variables”). A change in `max_heap_table_size` takes effect for `MEMORY` tables that are created or updated using `ALTER TABLE ... ENGINE = MEMORY` or `TRUNCATE TABLE` following the change, or for all `MEMORY` tables following a server restart. If you increase the value of this variable on the master without doing so on the slave, it becomes possible for a table on the master to grow larger than its counterpart on the slave, leading to inserts that succeed on the master but fail on the slave with `Table is full` errors. This is a known issue (Bug #48666). In such cases, you must set the global value of `max_heap_table_size` [504] on the slave as well as on the master, then restart replication. It is also recommended that you restart both the master and slave MySQL servers, to insure that the new value takes complete (global) effect on each of them.

See Section 14.4, “The `MEMORY` Storage Engine”, for more information about `MEMORY` tables.

### 16.4.1.21. Replication and Temporary Tables

This section does not apply when row-based replication is in use because in that case temporary tables are not replicated. It does apply with mixed-format replication for statements involving temporary tables that can be logged safely using statement-based format. For information about row-based replication and temporary tables, see [RBL, RBR, and temporary tables](#).

**Safe slave shutdown when using temporary tables.** Temporary tables are replicated except in the case where you stop the slave server (not just the slave threads) and you have replicated temporary tables that are open for use in updates that have not yet been executed on the slave. If you stop the slave server, the temporary tables needed by those updates are no longer available when the slave is restarted. To avoid this problem, do not shut down the slave while it has temporary tables open. Instead, use the following procedure:

1. Issue a `STOP SLAVE SQL_THREAD` statement.
2. Use `SHOW STATUS` to check the value of the `Slave_open_temp_tables` [598] variable.
3. If the value is not 0, restart the slave SQL thread with `START SLAVE SQL_THREAD` and repeat the procedure later.
4. When the value is 0, issue a `mysqladmin shutdown` command to stop the slave.

**Temporary tables and replication options.** By default, all temporary tables are replicated; this happens whether or not there are any matching `--replicate-do-db` [1976], `--replicate-do-table` [1978], or `--replicate-wild-do-table` [1979] options in effect. However, the `--replicate-ignore-table` [1978] and `--replicate-wild-ignore-table` [1980] options are honored for temporary tables.

A recommended practice when using statement-based or mixed-format replication is to designate a prefix for exclusive use in naming temporary tables that you do not want replicated, then employ a `--replicate-wild-ignore-table` [1980] option to match that prefix. For example, you might give all such tables names beginning with `norep` (such as `norepmytable`, `norepyourtable`, and so on), then use `--replicate-wild-ignore-table=norep%` [1980] to prevent them from being replicated.

#### 16.4.1.22. Replication of the `mysql` System Database

Data modification statements made to tables in the `mysql` database are replicated according to the value of `binlog_format` [2010]; if this value is `MIXED`, these statements are replicated using row-based format. However, statements that would normally update this information indirectly—such as `GRANT`, `REVOKE`, and statements manipulating triggers, stored routines, and views—are replicated to slaves using statement-based replication.

#### 16.4.1.23. Replication and the Query Optimizer

It is possible for the data on the master and slave to become different if a statement is written in such a way that the data modification is nondeterministic; that is, left up to the query optimizer. (In general, this is not a good practice, even outside of replication.) Examples of nondeterministic statements include `DELETE` or `UPDATE` statements that use `LIMIT` with no `ORDER BY` clause; see [Section 16.4.1.15](#), “Replication and `LIMIT`”, for a detailed discussion of these.

#### 16.4.1.24. Replication and Reserved Words

You can encounter problems when you attempt to replicate from an older master to a newer slave and you make use of identifiers on the master that are reserved words in the newer MySQL version running on the slave. An example of this is using a table column named `current_user` on a 4.0 master that is replicating to a 4.1 or higher slave because `CURRENT_USER` is a reserved word beginning in MySQL 4.1. Replication can fail in such cases with Error 1064 *You have an error in your SQL syntax. . . , even if a database or table named using the reserved word or a table having a column named using the reserved word is excluded from replication*. This is due to the fact that each SQL event must be parsed by the slave prior to execution, so that the slave knows which database object or objects would be affected; only after the event is parsed can the slave apply any filtering rules defined by `--replicate-do-db` [1976], `--replicate-do-table` [1978], `--replicate-ignore-db` [1977], and `--replicate-ignore-table` [1978].

To work around the problem of database, table, or column names on the master which would be regarded as reserved words by the slave, do one of the following:

- Use one or more `ALTER TABLE` statements on the master to change the names of any database objects where these names would be considered reserved words on the slave, and change any SQL statements that use the old names to use the new names instead.
- In any SQL statements using these database object names, write the names as quoted identifiers using backtick characters (```).

For listings of reserved words by MySQL version, see [Reserved Words](#), in the *MySQL Server Version Reference*. For identifier quoting rules, see [Section 9.2](#), “Schema Object Names”.

#### 16.4.1.25. Slave Errors During Replication

If a statement produces the same error (identical error code) on both the master and the slave, the error is logged, but replication continues.



If a statement produces different errors on the master and the slave, the slave SQL thread terminates, and the slave writes a message to its error log and waits for the database administrator to decide what to do about the error. This includes the case that a statement produces an error on the master or the slave, but not both. To address the issue, connect to the slave manually and determine the cause of the problem. `SHOW SLAVE STATUS` is useful for this. Then fix the problem and run `START SLAVE`. For example, you might need to create a nonexistent table before you can start the slave again.

If this error code validation behavior is not desirable, some or all errors can be masked out (ignored) with the `--slave-skip-errors` [1986] option.

For nontransactional storage engines such as `MyISAM`, it is possible to have a statement that only partially updates a table and returns an error code. This can happen, for example, on a multiple-row insert that has one row violating a key constraint, or if a long update statement is killed after updating some of the rows. If that happens on the master, the slave expects execution of the statement to result in the same error code. If it does not, the slave SQL thread stops as described previously.

If you are replicating between tables that use different storage engines on the master and slave, keep in mind that the same statement might produce a different error when run against one version of the table, but not the other, or might cause an error for one version of the table, but not the other. For example, since `MyISAM` ignores foreign key constraints, an `INSERT` or `UPDATE` statement accessing an `InnoDB` table on the master might cause a foreign key violation but the same statement performed on a `MyISAM` version of the same table on the slave would produce no such error, causing replication to stop.

#### 16.4.1.26. Replication and Server SQL Mode

Using different server SQL mode settings on the master and the slave may cause the same `INSERT` statements to be handled differently on the master and the slave, leading the master and slave to diverge. For best results, you should always use the same server SQL mode on the master and on the slave. This advice applies whether you are using statement-based or row-based replication.

If you are replicating partitioned tables, using different SQL modes on the master and the slave is likely to cause issues. At a minimum, this is likely to cause the distribution of data among partitions to be different in the master's and slave's copies of a given table. It may also cause inserts into partitioned tables that succeed on the master to fail on the slave.

For more information, see [Section 5.1.7, “Server SQL Modes”](#).

#### 16.4.1.27. Replication Retries and Timeouts

The global system variable `slave_transaction_retries` [1995] affects replication as follows: If the slave SQL thread fails to execute a transaction because of an `InnoDB` deadlock or because it exceeded the `InnoDB innodb_lock_wait_timeout` [1729] value, or the `NDB TransactionDeadlockDetectionTimeout` or `TransactionInactiveTimeout` value, the slave automatically retries the transaction `slave_transaction_retries` [1995] times before stopping with an error. The default value is 10. The total retry count can be seen in the output of `SHOW STATUS`; see [Section 5.1.6, “Server Status Variables”](#).

#### 16.4.1.28. Replication and Time Zones

The same system time zone should be set for both master and slave. Otherwise, statements depending on the local time on the master are not replicated properly, such as statements that use the `NOW()` [1220] or `FROM_UNIXTIME()` [1218] functions. You can set the time zone in which MySQL server runs by using the `--timezone=timezone_name` [229] option of the `mysqld_safe` script or by setting the `TZ` environment variable. See also [Section 16.4.1.14, “Replication and System Functions”](#).

If the master is MySQL 4.1 or earlier, both master and slave should also use the same default connection time zone. That is, the `--default-time-zone` [420] parameter should have the same value for both master and slave.



`CONVERT_TZ(..., ..., @session.time_zone)` [1212] is properly replicated only if both master and slave are running MySQL 5.0.4 or newer.

### 16.4.1.29. Replication and Transactions

**Mixing transactional and nontransactional statements within the same transaction.** In general, you should avoid transactions that update both transactional and nontransactional tables in a replication environment. You should also avoid using any statement that accesses both transactional (or temporary) and nontransactional tables and writes to any of them.

As of MySQL 5.5.2, the server uses these rules for binary logging:

- If the initial statements in a transaction are nontransactional, they are written to the binary log immediately. The remaining statements in the transaction are cached and not written to the binary log until the transaction is committed. (If the transaction is rolled back, the cached statements are written to the binary log only if they make nontransactional changes that cannot be rolled back. Otherwise, they are discarded.)
- For statement-based logging, logging of nontransactional statements is affected by the `binlog_direct_non_transactional_updates` [2010] system variable. When this variable is `OFF` (the default), logging is as just described. When this variable is `ON`, logging occurs immediately for nontransactional statements occurring anywhere in the transaction (not just initial nontransactional statements). Other statements are kept in the transaction cache and logged when the transaction commits. `binlog_direct_non_transactional_updates` [2010] has no effect for row-format or mixed-format binary logging.

#### Transactional, nontransactional, and mixed statements.

To apply those rules, the server considers a statement nontransactional if it changes only nontransactional tables, and transactional if it changes only transactional tables. In MySQL 5.6, a statement that references both nontransactional and transactional tables and updates *any* of the tables involved, is considered a “mixed” statement. (In previous MySQL release series, a statement that changed both nontransactional and transactional tables was considered mixed.) Mixed statements, like transactional statements, are cached and logged when the transaction commits.

A mixed statement that updates a transactional table is considered unsafe if the statement also performs either of the following actions:

- Updates or reads a transactional table
- Reads a nontransactional table and the transaction isolation level is less than `REPEATABLE_READ`

A mixed statement following the update of a transactional table within a transaction is considered unsafe if it performs either of the following actions:

- Updates any table and reads from any temporary table
- Updates a nontransactional table and `binlog_direct_non_trans_update` is `OFF`

For more information, see [Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”](#).



#### Note

A mixed statement is unrelated to mixed binary logging format.

In situations where transactions mix updates to transactional and nontransactional tables, the order of statements in the binary log is correct, and all needed statements are written to the binary log even in case of a `ROLLBACK`. However, when a second connection updates the nontransactional table before the first connection transaction is complete, statements can be logged out of order because the second connection update is written immediately after it is performed, regardless of the state of the transaction being performed by the first connection.

**Using different storage engines on master and slave.** It is possible to replicate transactional tables on the master using nontransactional tables on the slave. For example, you can replicate an [InnoDB](#) master table as a [MyISAM](#) slave table. However, if you do this, there are problems if the slave is stopped in the middle of a `BEGIN ... COMMIT` block because the slave restarts at the beginning of the `BEGIN` block.

In MySQL 5.6, it is also safe to replicate transactions from [MyISAM](#) tables on the master to transactional tables—such as tables that use the [InnoDB](#) storage engine—on the slave. In such cases (beginning with MySQL 5.5.0), an `AUTOCOMMIT=1` [\[464\]](#) statement issued on the master is replicated, thus enforcing `AUTOCOMMIT` mode on the slave.

When the storage engine type of the slave is nontransactional, transactions on the master that mix updates of transactional and nontransactional tables should be avoided because they can cause inconsistency of the data between the master transactional table and the slave nontransactional table. That is, such transactions can lead to master storage engine-specific behavior with the possible effect of replication going out of synchrony. MySQL does not issue a warning about this currently, so extra care should be taken when replicating transactional tables from the master to nontransactional tables on the slaves.

**Changing the binary logging format within transactions.** Beginning with MySQL 5.5.3, the `binlog_format` [\[2010\]](#) system variable is read-only as long as a transaction is in progress. (Bug #47863)

Every transaction (including `autocommit` [\[464\]](#) transactions) is recorded in the binary log as though it starts with a `BEGIN` statement, and ends with either a `COMMIT` or a `ROLLBACK` statement. In MySQL 5.6, this true is even for statements affecting tables that use a nontransactional storage engine (such as [MyISAM](#)).

### 16.4.1.30. Replication and Triggers

With statement-based replication, triggers executed on the master also execute on the slave. With row-based replication, triggers executed on the master do not execute on the slave. Instead, the row changes on the master resulting from trigger execution are replicated and applied on the slave.

This behavior is by design. If under row-based replication the slave applied the triggers as well as the row changes caused by them, the changes would in effect be applied twice on the slave, leading to different data on the master and the slave.

If you want triggers to execute on both the master and the slave—perhaps because you have different triggers on the master and slave—you must use statement-based replication. However, to enable slave-side triggers, it is not necessary to use statement-based replication exclusively. It is sufficient to switch to statement-based replication only for those statements where you want this effect, and to use row-based replication the rest of the time.

A statement invoking a trigger (or function) that causes an update to an `AUTO_INCREMENT` column is not replicated correctly using statement-based replication. MySQL 5.6 marks such statements as unsafe. (Bug #45677)

### 16.4.1.31. Replication and `TRUNCATE TABLE`

`TRUNCATE TABLE` is normally regarded as a DML statement, and so would be expected to be logged and replicated using row-based format when the binary logging mode is `ROW` or `MIXED`. However this caused issues when logging or replicating, in `STATEMENT` or `MIXED` mode, tables that used transactional storage engines such as [InnoDB](#) when the transaction isolation level was `READ COMMITTED` or `READ UNCOMMITTED`, which precludes statement-based logging.

`TRUNCATE TABLE` is treated for purposes of logging and replication as DDL rather than DML so that it can be logged and replicated as a statement. However, the effects of the statement as applicable to [InnoDB](#) and other transactional tables on replication slaves still follow the rules described in [Section 13.1.27](#), “`TRUNCATE TABLE Syntax`” governing such tables. (Bug #36763)

### 16.4.1.32. Replication and Variables

System variables are not replicated correctly when using `STATEMENT` mode, except for the following variables when they are used with session scope:

- `auto_increment_increment` [1965]
- `auto_increment_offset` [1968]
- `character_set_client` [467]
- `character_set_connection` [467]
- `character_set_database` [468]
- `character_set_server` [468]
- `collation_connection` [469]
- `collation_database` [469]
- `collation_server` [470]
- `foreign_key_checks` [483]
- `identity` [488]
- `last_insert_id` [495]
- `lc_time_names` [495]
- `pseudo_thread_id` [527]
- `sql_auto_is_null` [543]
- `time_zone` [557]
- `timestamp` [558]
- `unique_checks` [561]

When `MIXED` mode is used, the variables in the preceding list, when used with session scope, cause a switch from statement-based to row-based logging. See [Section 5.2.4.3, “Mixed Binary Logging Format”](#).

`sql_mode` [545] is also replicated except for the `NO_DIR_IN_CREATE` [604] mode; the slave always preserves its own value for `NO_DIR_IN_CREATE` [604], regardless of changes to it on the master. This is true for all replication formats.

However, when `mysqlbinlog` parses a `SET @@sql_mode = mode` statement, the full `mode` value, including `NO_DIR_IN_CREATE` [604], is passed to the receiving server. For this reason, replication of such a statement may not be safe when `STATEMENT` mode is in use.

The `default_storage_engine` [474] and `storage_engine` [549] system variables are not replicated, regardless of the logging mode; this is intended to facilitate replication between different storage engines.

The `read_only` [532] system variable is not replicated. In addition, the enabling this variable has different effects with regard to temporary tables, table locking, and the `SET PASSWORD` statement in different MySQL versions.

The `max_heap_table_size` [504] system variable is not replicated. Increasing the value of this variable on the master without doing so on the slave can lead eventually to `Table is full` errors on the slave when trying to execute `INSERT` statements on a `MEMORY` table on the master that is thus

permitted to grow larger than its counterpart on the slave. For more information, see [Section 16.4.1.20, “Replication and MEMORY Tables”](#).

In statement-based replication, session variables are not replicated properly when used in statements that update tables. For example, the following sequence of statements will not insert the same data on the master and the slave:

```
SET max_join_size=1000;
INSERT INTO mytable VALUES(@@max_join_size);
```

This does not apply to the common sequence:

```
SET time_zone=...;
INSERT INTO mytable VALUES(CONVERT_TZ(..., ..., @@time_zone));
```

Replication of session variables is not a problem when row-based replication is being used, in which case, session variables are always replicated safely. See [Section 16.1.2, “Replication Formats”](#).

In MySQL 5.6, the following session variables are written to the binary log and honored by the replication slave when parsing the binary log, regardless of the logging format:

- `sql_mode` [\[545\]](#)
- `foreign_key_checks` [\[483\]](#)
- `unique_checks` [\[561\]](#)
- `character_set_client` [\[467\]](#)
- `collation_connection` [\[469\]](#)
- `collation_database` [\[469\]](#)
- `collation_server` [\[470\]](#)
- `sql_auto_is_null` [\[543\]](#)



#### Important

Even though session variables relating to character sets and collations are written to the binary log, replication between different character sets is not supported.

To help reduce possible confusion, we recommend that you always use the same setting for the `lower_case_table_names` [\[501\]](#) system variable on both master and slave, especially when you are running MySQL on platforms with case-sensitive file systems.



#### Note

In previous versions of MySQL, when a case-sensitive file system was in use, setting this variable to 1 on the slave and to a different value on the master could lead to replication failure. This issue is fixed in MySQL 5.6.1. (Bug [#37656](#))

### 16.4.1.33. Replication and Views

Views are always replicated to slaves. Views are filtered by their own name, not by the tables they refer to. This means that a view can be replicated to the slave even if the view contains a table that would normally be filtered out by `replication-ignore-table` rules. Care should therefore be taken to ensure that views do not replicate table data that would normally be filtered for security reasons.

MySQL does not support replication from tables to views; you can alter the target table and still replicate to it in many cases (see [Section 16.4.1.8.2, “Replication of Columns Having Different Data](#)

Types”, and [Section 16.4.1.8, “Replication with Differing Table Definitions on Master and Slave”](#)), but you cannot replace the target table with a view. In MySQL 5.6.11 and later, trying to do so causes an error. (Bug #11752707, Bug #43975)

## 16.4.2. Replication Compatibility Between MySQL Versions

MySQL supports replication from one major version to the next higher major version. For example, you can replicate from a master running MySQL 4.1 to a slave running MySQL 5.0, from a master running MySQL 5.0 to a slave running MySQL 5.1, and so on.

However, one may encounter difficulties when replicating from an older master to a newer slave if the master uses statements or relies on behavior no longer supported in the version of MySQL used on the slave. For example, in MySQL 5.5, `CREATE TABLE . . . SELECT` statements are permitted to change tables other than the one being created, but are no longer allowed to do so in MySQL 5.6 (see [Section 16.4.1.4, “Replication of CREATE TABLE . . . SELECT Statements”](#)).

The use of more than 2 MySQL Server versions is not supported in replication setups involving multiple masters, regardless of the number of master or slave MySQL servers. This restriction applies not only to major versions, but to minor versions within the same major version as well. For example, if you are using a chained or circular replication setup, you cannot use MySQL 5.6.1, MySQL 5.6.2, and MySQL 5.6.4 concurrently, although you could use any 2 of these releases together.

In some cases, it is also possible to replicate between a master and a slave that is more than one major version newer than the master. However, there are known issues with trying to replicate from a master running MySQL 4.1 or earlier to a slave running MySQL 5.1 or later. To work around such problems, you can insert a MySQL server running an intermediate version between the two; for example, rather than replicating directly from a MySQL 4.1 master to a MySQL 5.1 slave, it is possible to replicate from a MySQL 4.1 server to a MySQL 5.0 server, and then from the MySQL 5.0 server to a MySQL 5.1 server.



### Important

It is strongly recommended to use the most recent release available within a given MySQL major version because replication (and other) capabilities are continually being improved. It is also recommended to upgrade masters and slaves that use early releases of a major version of MySQL to GA (production) releases when the latter become available for that major version.

Replication from newer masters to older slaves may be possible, but is generally not supported. This is due to a number of factors:

- **Binary log format changes.** The binary log format can change between major releases. While we attempt to maintain backward compatibility, this is not always possible. For example, the binary log format implemented in MySQL 5.0 changed considerably from that used in previous versions, especially with regard to handling of character sets, `LOAD DATA INFILE`, and time zones. This means that replication from a MySQL 5.0 (or later) master to a MySQL 4.1 (or earlier) slave is generally not supported.

This also has significant implications for upgrading replication servers; see [Section 16.4.3, “Upgrading a Replication Setup”](#), for more information.

- **Use of row-based replication.** Row-based replication was implemented in MySQL 5.1.5, so you cannot replicate using row-based replication from any MySQL 5.6 or later master to a slave older than MySQL 5.1.5.

For more information about row-based replication, see [Section 16.1.2, “Replication Formats”](#).

- **SQL incompatibilities.** You cannot replicate from a newer master to an older slave using statement-based replication if the statements to be replicated use SQL features available on the master but not on the slave.

However, if both the master and the slave support row-based replication, and there are no data definition statements to be replicated that depend on SQL features found on the master but not on the slave, you can use row-based replication to replicate the effects of data modification statements even if the DDL run on the master is not supported on the slave.

For more information on potential replication issues, see [Section 16.4.1, “Replication Features and Issues”](#).

### 16.4.3. Upgrading a Replication Setup

When you upgrade servers that participate in a replication setup, the procedure for upgrading depends on the current server versions and the version to which you are upgrading.

This section applies to upgrading replication from older versions of MySQL to MySQL 5.6. A 4.0 server should be 4.0.3 or newer.

When you upgrade a master to 5.6 from an earlier MySQL release series, you should first ensure that all the slaves of this master are using the same 5.6.x release. If this is not the case, you should first upgrade the slaves. To upgrade each slave, shut it down, upgrade it to the appropriate 5.6.x version, restart it, and restart replication. The 5.6 slave is able to read the old relay logs written prior to the upgrade and to execute the statements they contain. Relay logs created by the slave after the upgrade are in 5.6 format.

After the slaves have been upgraded, shut down the master, upgrade it to the same 5.6.x release as the slaves, and restart it. The 5.6 master is able to read the old binary logs written prior to the upgrade and to send them to the 5.6 slaves. The slaves recognize the old format and handle it properly. Binary logs created by the master subsequent to the upgrade are in 5.6 format. These too are recognized by the 5.6 slaves.

In other words, when upgrading to MySQL 5.6, the slaves must be MySQL 5.6 before you can upgrade the master to 5.6. Note that downgrading from 5.6 to older versions does not work so simply: You must ensure that any 5.6 binary log or relay log has been fully processed, so that you can remove it before proceeding with the downgrade.

Downgrading a replication setup to a previous version cannot be done once you have switched from statement-based to row-based replication, and after the first row-based statement has been written to the binlog. See [Section 16.1.2, “Replication Formats”](#).

Some upgrades may require that you drop and re-create database objects when you move from one MySQL series to the next. For example, collation changes might require that table indexes be rebuilt. Such operations, if necessary, will be detailed at [Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”](#). It is safest to perform these operations separately on the slaves and the master, and to disable replication of these operations from the master to the slave. To achieve this, use the following procedure:

1. Stop all the slaves and upgrade them. Restart them with the `--skip-slave-start [1984]` option so that they do not connect to the master. Perform any table repair or rebuilding operations needed to re-create database objects, such as use of `REPAIR TABLE` or `ALTER TABLE`, or dumping and reloading tables or triggers.
2. Disable the binary log on the master. To do this without restarting the master, execute a `SET sql_log_bin = 0` statement. Alternatively, stop the master and restart it without the `--log-bin [2002]` option. If you restart the master, you might also want to disallow client connections. For example, if all clients connect using TCP/IP, use the `--skip-networking [440]` option when you restart the master.
3. With the binary log disabled, perform any table repair or rebuilding operations needed to re-create database objects. The binary log must be disabled during this step to prevent these operations from being logged and sent to the slaves later.



4. Re-enable the binary log on the master. If you set `sql_log_bin` [544] to 0 earlier, execute a `SET sql_log_bin = 1` statement. If you restarted the master to disable the binary log, restart it with `--log-bin` [2002], and without `--skip-networking` [440] so that clients and slaves can connect.
5. Restart the slaves, this time without the `--skip-slave-start` [1984] option.

Replication with global transaction identifiers was introduced in MySQL 5.6.7. If you are upgrading an existing replication setup from a version of MySQL that does not support GTIDs to a version that does, you should not enable GTIDs on either the master or the slave before making sure that the setup meets all the requirements for GTID-based replication. See [Section 16.1.3.2, “Setting Up Replication Using GTIDs”](#), which contains information about converting existing replication setups to use GTID-based replication.

## 16.4.4. Troubleshooting Replication

If you have followed the instructions but your replication setup is not working, the first thing to do is *check the error log for messages*. Many users have lost time by not doing this soon enough after encountering problems.

If you cannot tell from the error log what the problem was, try the following techniques:

- Verify that the master has binary logging enabled by issuing a `SHOW MASTER STATUS` statement. If logging is enabled, `Position` is nonzero. If binary logging is not enabled, verify that you are running the master with the `--log-bin` [2002] option.
- Verify that the master and slave both were started with the `--server-id` [1955] option and that the ID value is unique on each server.
- Verify that the slave is running. Use `SHOW SLAVE STATUS` to check whether the `Slave_IO_Running` and `Slave_SQL_Running` values are both `Yes`. If not, verify the options that were used when starting the slave server. For example, `--skip-slave-start` [1984] prevents the slave threads from starting until you issue a `START SLAVE` statement.
- If the slave is running, check whether it established a connection to the master. Use `SHOW PROCESSLIST`, find the I/O and SQL threads and check their `State` column to see what they display. See [Section 16.2.1, “Replication Implementation Details”](#). If the I/O thread state says `Connecting to master`, check the following:
  - Verify the privileges for the user being used for replication on the master.
  - Check that the host name of the master is correct and that you are using the correct port to connect to the master. The port used for replication is the same as used for client network communication (the default is 3306). For the host name, ensure that the name resolves to the correct IP address.
  - Check that networking has not been disabled on the master or slave. Look for the `skip-networking` [440] option in the configuration file. If present, comment it out or remove it.
  - If the master has a firewall or IP filtering configuration, ensure that the network port being used for MySQL is not being filtered.
  - Check that you can reach the master by using `ping` or `traceroute/tracert` to reach the host.
- If the slave was running previously but has stopped, the reason usually is that some statement that succeeded on the master failed on the slave. This should never happen if you have taken a proper snapshot of the master, and never modified the data on the slave outside of the slave thread. If the slave stops unexpectedly, it is a bug or you have encountered one of the known replication limitations described in [Section 16.4.1, “Replication Features and Issues”](#). If it is a bug, see [Section 16.4.5, “How to Report Replication Bugs or Problems”](#), for instructions on how to report it.



- If a statement that succeeded on the master refuses to run on the slave, try the following procedure if it is not feasible to do a full database resynchronization by deleting the slave's databases and copying a new snapshot from the master:

1. Determine whether the affected table on the slave is different from the master table. Try to understand how this happened. Then make the slave's table identical to the master's and run `START SLAVE`.
2. If the preceding step does not work or does not apply, try to understand whether it would be safe to make the update manually (if needed) and then ignore the next statement from the master.
3. If you decide that the slave can skip the next statement from the master, issue the following statements:

```
mysql> SET GLOBAL sql_slave_skip_counter = N;  
mysql> START SLAVE;
```

The value of *N* should be 1 if the next statement from the master does not use `AUTO_INCREMENT` or `LAST_INSERT_ID()` [1269]. Otherwise, the value should be 2. The reason for using a value of 2 for statements that use `AUTO_INCREMENT` or `LAST_INSERT_ID()` [1269] is that they take two events in the binary log of the master.

See also [Section 13.4.2.4](#), “`SET GLOBAL sql_slave_skip_counter Syntax`”.

4. If you are sure that the slave started out perfectly synchronized with the master, and that no one has updated the tables involved outside of the slave thread, then presumably the discrepancy is the result of a bug. If you are running the most recent version of MySQL, please report the problem. If you are running an older version, try upgrading to the latest production release to determine whether the problem persists.

## 16.4.5. How to Report Replication Bugs or Problems

When you have determined that there is no user error involved, and replication still either does not work at all or is unstable, it is time to send us a bug report. We need to obtain as much information as possible from you to be able to track down the bug. Please spend some time and effort in preparing a good bug report.

If you have a repeatable test case that demonstrates the bug, please enter it into our bugs database using the instructions given in [Section 1.7](#), “[How to Report Bugs or Problems](#)”. If you have a “phantom” problem (one that you cannot duplicate at will), use the following procedure:

1. Verify that no user error is involved. For example, if you update the slave outside of the slave thread, the data goes out of synchrony, and you can have unique key violations on updates. In this case, the slave thread stops and waits for you to clean up the tables manually to bring them into synchrony. *This is not a replication problem. It is a problem of outside interference causing replication to fail.*
2. Run the slave with the `--log-slave-updates` [1971] and `--log-bin` [2002] options. These options cause the slave to log the updates that it receives from the master into its own binary logs.
3. Save all evidence before resetting the replication state. If we have no information or only sketchy information, it becomes difficult or impossible for us to track down the problem. The evidence you should collect is:
  - All binary log files from the master
  - All binary log files from the slave
  - The output of `SHOW MASTER STATUS` from the master at the time you discovered the problem
  - The output of `SHOW SLAVE STATUS` from the slave at the time you discovered the problem

- Error logs from the master and the slave
4. Use `mysqlbinlog` to examine the binary logs. The following should be helpful to find the problem statement. `log_file` and `log_pos` are the `Master_Log_File` and `Read_Master_Log_Pos` values from `SHOW SLAVE STATUS`.

```
shell> mysqlbinlog --start-position=log_pos log_file | head
```

After you have collected the evidence for the problem, try to isolate it as a separate test case first. Then enter the problem with as much information as possible into our bugs database using the instructions at [Section 1.7, “How to Report Bugs or Problems”](#).

# Chapter 17. Partitioning

## Table of Contents

17.1. Overview of Partitioning in MySQL .....	2086
17.2. Partitioning Types .....	2089
17.2.1. <code>RANGE</code> Partitioning .....	2091
17.2.2. <code>LIST</code> Partitioning .....	2094
17.2.3. <code>COLUMNS</code> Partitioning .....	2097
17.2.4. <code>HASH</code> Partitioning .....	2104
17.2.5. <code>KEY</code> Partitioning .....	2107
17.2.6. Subpartitioning .....	2108
17.2.7. How MySQL Partitioning Handles <code>NULL</code> .....	2111
17.3. Partition Management .....	2115
17.3.1. Management of <code>RANGE</code> and <code>LIST</code> Partitions .....	2116
17.3.2. Management of <code>HASH</code> and <code>KEY</code> Partitions .....	2121
17.3.3. Exchanging Partitions and Subpartitions with Tables .....	2122
17.3.4. Maintenance of Partitions .....	2127
17.3.5. Obtaining Information About Partitions .....	2128
17.4. Partition Pruning .....	2130
17.5. Partition Selection .....	2133
17.6. Restrictions and Limitations on Partitioning .....	2139
17.6.1. Partitioning Keys, Primary Keys, and Unique Keys .....	2144
17.6.2. Partitioning Limitations Relating to Storage Engines .....	2147
17.6.3. Partitioning Limitations Relating to Functions .....	2148
17.6.4. Partitioning and Locking .....	2149

This chapter discusses MySQL's implementation of *user-defined partitioning*. You can determine whether your MySQL Server supports partitioning by checking the output of the `SHOW PLUGINS` statement, as shown here:

```
mysql> SHOW PLUGINS;
```

Name	Status	Type	Library	License
binlog	ACTIVE	STORAGE ENGINE	NULL	GPL
<b>partition</b>	<b>ACTIVE</b>	<b>STORAGE ENGINE</b>	<b>NULL</b>	<b>GPL</b>
ARCHIVE	ACTIVE	STORAGE ENGINE	NULL	GPL
BLACKHOLE	ACTIVE	STORAGE ENGINE	NULL	GPL
CSV	ACTIVE	STORAGE ENGINE	NULL	GPL
FEDERATED	DISABLED	STORAGE ENGINE	NULL	GPL
MEMORY	ACTIVE	STORAGE ENGINE	NULL	GPL
InnoDB	ACTIVE	STORAGE ENGINE	NULL	GPL
MRG_MYISAM	ACTIVE	STORAGE ENGINE	NULL	GPL
MyISAM	ACTIVE	STORAGE ENGINE	NULL	GPL
ndbcluster	DISABLED	STORAGE ENGINE	NULL	GPL

```
11 rows in set (0.00 sec)
```

If you do not see the `partition` plugin listed with the value `ACTIVE` for the `Status` column in the output of `SHOW PLUGINS` (show in bold text in the example just given), then your version of MySQL was not built with partitioning support.

MySQL 5.6 Community binaries provided by Oracle include partitioning support. For information about partitioning support offered in commercial MySQL Server binaries, see *MySQL Enterprise Server 5.1* on the MySQL Web site at <http://www.mysql.com/products/enterprise/server.html>.

To enable partitioning if you are compiling MySQL 5.6 from source, the build must be configured with the `-DWITH_PARTITION_STORAGE_ENGINE` option. For more information about building MySQL, see [Section 2.9, "Installing MySQL from Source"](#).

If your MySQL binary is built with partitioning support, nothing further needs to be done to enable it (for example, no special entries are required in your `my.cnf` file).

If you want to disable partitioning support, you can start the MySQL Server with the `--skip-partition` [440] option, in which case the value of `have_partitioning` [487] is `DISABLED`. When partitioning support is disabled, you can see any existing partitioned tables and drop them (although doing this is not advised), but you cannot otherwise manipulate them or access their data.

An introduction to partitioning and partitioning concepts may be found in [Section 17.1, “Overview of Partitioning in MySQL”](#).

MySQL supports several types of partitioning, which are discussed in [Section 17.2, “Partitioning Types”](#), as well as subpartitioning, which is described in [Section 17.2.6, “Subpartitioning”](#).

Methods of adding, removing, and altering partitions in existing partitioned tables are covered in [Section 17.3, “Partition Management”](#).

Table maintenance commands for use with partitioned tables are discussed in [Section 17.3.4, “Maintenance of Partitions”](#).

The `PARTITIONS` table in the `INFORMATION_SCHEMA` database provides information about partitions and partitioned tables. See [Section 19.14, “The INFORMATION\\_SCHEMA PARTITIONS Table”](#), for more information; for some examples of queries against this table, see [Section 17.2.7, “How MySQL Partitioning Handles NULL”](#).

For known issues with partitioning in MySQL 5.6, see [Section 17.6, “Restrictions and Limitations on Partitioning”](#).

You may also find the following resources to be useful when working with partitioned tables.

**Additional Resources.** Other sources of information about user-defined partitioning in MySQL include the following:

- [MySQL Partitioning Forum](#)

This is the official discussion forum for those interested in or experimenting with MySQL Partitioning technology. It features announcements and updates from MySQL developers and others. It is monitored by members of the Partitioning Development and Documentation Teams.

- [Mikael Ronström's Blog](#)

MySQL Partitioning Architect and Lead Developer Mikael Ronström frequently posts articles here concerning his work with MySQL Partitioning and MySQL Cluster.

- [PlanetMySQL](#)

A MySQL news site featuring MySQL-related blogs, which should be of interest to anyone using my MySQL. We encourage you to check here for links to blogs kept by those working with MySQL Partitioning, or to have your own blog added to those covered.

MySQL 5.6 binaries are available from <http://dev.mysql.com/downloads/mysql/5.6.html>. However, for the latest partitioning bugfixes and feature additions, you can obtain the source from our Bazaar repository. To enable partitioning, the build must be configured with the `-DWITH_PARTITION_STORAGE_ENGINE` option. For more information about building MySQL, see [Section 2.9, “Installing MySQL from Source”](#). If you have problems compiling a partitioning-enabled MySQL 5.6 build, check the [MySQL Partitioning Forum](#) and ask for assistance there if you do not find a solution to your problem already posted.

## 17.1. Overview of Partitioning in MySQL

This section provides a conceptual overview of partitioning in MySQL 5.6.

For information on partitioning restrictions and feature limitations, see [Section 17.6, “Restrictions and Limitations on Partitioning”](#).

The SQL standard does not provide much in the way of guidance regarding the physical aspects of data storage. The SQL language itself is intended to work independently of any data structures or media underlying the schemas, tables, rows, or columns with which it works. Nonetheless, most advanced database management systems have evolved some means of determining the physical location to be used for storing specific pieces of data in terms of the file system, hardware or even both. In MySQL, the [InnoDB](#) storage engine has long supported the notion of a tablespace, and the MySQL Server, even prior to the introduction of partitioning, could be configured to employ different physical directories for storing different databases (see [Section 8.11.3.1, “Using Symbolic Links”](#), for an explanation of how this is done).

*Partitioning* takes this notion a step further, by enabling you to distribute portions of individual tables across a file system according to rules which you can set largely as needed. In effect, different portions of a table are stored as separate tables in different locations. The user-selected rule by which the division of data is accomplished is known as a *partitioning function*, which in MySQL can be the modulus, simple matching against a set of ranges or value lists, an internal hashing function, or a linear hashing function. The function is selected according to the partitioning type specified by the user, and takes as its parameter the value of a user-supplied expression. This expression can be a column value, a function acting on one or more column values, or a set of one or more column values, depending on the type of partitioning that is used.

In the case of [RANGE](#), [LIST](#), and [\[LINEAR\] HASH](#) partitioning, the value of the partitioning column is passed to the partitioning function, which returns an integer value representing the number of the partition in which that particular record should be stored. This function must be nonconstant and nonrandom. It may not contain any queries, but may use an SQL expression that is valid in MySQL, as long as that expression returns either [NULL](#) or an integer *intval* such that

```
-MAXVALUE <= intval <= MAXVALUE
```

([MAXVALUE](#) is used to represent the least upper bound for the type of integer in question. [-MAXVALUE](#) represents the greatest lower bound.)

For [\[LINEAR\] KEY](#), [RANGE COLUMNS](#), and [LIST COLUMNS](#) partitioning, the partitioning expression consists of a list of one or more columns.

For [\[LINEAR\] KEY](#) partitioning, the partitioning function is supplied by MySQL.

For more information about permitted partitioning column types and partitioning functions, see [Section 17.2, “Partitioning Types”](#), as well as [Section 13.1.14, “CREATE TABLE Syntax”](#), which provides partitioning syntax descriptions and additional examples. For information about restrictions on partitioning functions, see [Section 17.6.3, “Partitioning Limitations Relating to Functions”](#).

This is known as *horizontal partitioning*—that is, different rows of a table may be assigned to different physical partitions. MySQL 5.6 does not support *vertical partitioning*, in which different columns of a table are assigned to different physical partitions. There are not at this time any plans to introduce vertical partitioning into MySQL 5.6.

For information about determining whether your MySQL Server binary supports user-defined partitioning, see [Chapter 17, Partitioning](#).

For creating partitioned tables, you can use most storage engines that are supported by your MySQL server; the MySQL partitioning engine runs in a separate layer and can interact with any of these. In MySQL 5.6, all partitions of the same partitioned table must use the same storage engine; for example, you cannot use [MyISAM](#) for one partition and [InnoDB](#) for another. However, there is nothing preventing you from using different storage engines for different partitioned tables on the same MySQL server or even in the same database.

MySQL partitioning cannot be used with the [MERGE](#), [CSV](#), or [FEDERATED](#) storage engines.

To employ a particular storage engine for a partitioned table, it is necessary only to use the `[STORAGE] ENGINE` option just as you would for a nonpartitioned table. However, you should keep in mind that `[STORAGE] ENGINE` (and other table options) need to be listed *before* any partitioning options are used in a `CREATE TABLE` statement. This example shows how to create a table that is partitioned by hash into 6 partitions and which uses the `InnoDB` storage engine:

```
CREATE TABLE ti (id INT, amount DECIMAL(7,2), tr_date DATE)
ENGINE=INNODB
PARTITION BY HASH( MONTH(tr_date) )
PARTITIONS 6;
```

Each `PARTITION` clause can include a `[STORAGE] ENGINE` option, but in MySQL 5.6 this has no effect.



### Important

Partitioning applies to all data and indexes of a table; you cannot partition only the data and not the indexes, or *vice versa*, nor can you partition only a portion of the table.

Data and indexes for each partition can be assigned to a specific directory using the `DATA DIRECTORY` and `INDEX DIRECTORY` options for the `PARTITION` clause of the `CREATE TABLE` statement used to create the partitioned table.

The `DATA DIRECTORY` and `INDEX DIRECTORY` options have no effect when defining partitions for tables using the `InnoDB` storage engine.

`DATA DIRECTORY` and `INDEX DIRECTORY` are not supported for individual partitions or subpartitions on Windows. These options are ignored on Windows, except that a warning is generated.

In addition, `MAX_ROWS` and `MIN_ROWS` can be used to determine the maximum and minimum numbers of rows, respectively, that can be stored in each partition. See [Section 17.3, “Partition Management”](#), for more information on these options.

Some advantages of partitioning are listed here:

- Partitioning makes it possible to store more data in one table than can be held on a single disk or file system partition.
- Data that loses its usefulness can often be easily removed from a partitioned table by dropping the partition (or partitions) containing only that data. Conversely, the process of adding new data can in some cases be greatly facilitated by adding one or more new partitions for storing specifically that data.
- Some queries can be greatly optimized in virtue of the fact that data satisfying a given `WHERE` clause can be stored only on one or more partitions, which automatically excluding any remaining partitions from the search. Because partitions can be altered after a partitioned table has been created, you can reorganize your data to enhance frequent queries that may not have been often used when the partitioning scheme was first set up. This ability to exclude non-matching partitions (and thus any rows they contain) is often referred to as *partition pruning*. For more information, see [Section 17.4, “Partition Pruning”](#).

In addition, MySQL 5.6 supports explicit partition selection for queries. For example, `SELECT * FROM t PARTITION (p0,p1) WHERE c < 5` selects only those rows in partitions `p0` and `p1` that match the `WHERE` condition. In this case, MySQL does not check any other partitions of table `t`; this can greatly speed up queries when you already know which partition or partitions you wish to examine. Partition selection is also supported for the data modification statements `DELETE`, `INSERT`, `REPLACE`, `UPDATE`, and `LOAD DATA`, `LOAD XML`. See the descriptions of these statements for more information and examples.

Other benefits usually associated with partitioning include those in the following list. These features are not currently implemented in MySQL Partitioning, but are high on our list of priorities.

- Queries involving aggregate functions such as `SUM()` [1285] and `COUNT()` [1283] can easily be parallelized. A simple example of such a query might be `SELECT salesperson_id, COUNT(orders) as order_total FROM sales GROUP BY salesperson_id;` By “parallelized,” we mean that the query can be run simultaneously on each partition, and the final result obtained merely by summing the results obtained for all partitions.
- Achieving greater query throughput in virtue of spreading data seeks over multiple disks.

Be sure to check this section and chapter frequently for updates as MySQL Partitioning development continues.

## 17.2. Partitioning Types

This section discusses the types of partitioning which are available in MySQL 5.6. These include the types listed here:

- **RANGE partitioning.** This type of partitioning assigns rows to partitions based on column values falling within a given range. See [Section 17.2.1, “RANGE Partitioning”](#). For information about an extension to this type, `RANGE COLUMNS`, see [Section 17.2.3.1, “RANGE COLUMNS partitioning”](#).
- **LIST partitioning.** Similar to partitioning by `RANGE`, except that the partition is selected based on columns matching one of a set of discrete values. See [Section 17.2.2, “LIST Partitioning”](#). For information about an extension to this type, `LIST COLUMNS`, see [Section 17.2.3.2, “LIST COLUMNS partitioning”](#).
- **HASH partitioning.** With this type of partitioning, a partition is selected based on the value returned by a user-defined expression that operates on column values in rows to be inserted into the table. The function may consist of any expression valid in MySQL that yields a nonnegative integer value. An extension to this type, `LINEAR HASH`, is also available. See [Section 17.2.4, “HASH Partitioning”](#).
- **KEY partitioning.** This type of partitioning is similar to partitioning by `HASH`, except that only one or more columns to be evaluated are supplied, and the MySQL server provides its own hashing function. These columns can contain other than integer values, since the hashing function supplied by MySQL guarantees an integer result regardless of the column data type. An extension to this type, `LINEAR KEY`, is also available. See [Section 17.2.5, “KEY Partitioning”](#).

A very common use of database partitioning is to segregate data by date. Some database systems support explicit date partitioning, which MySQL does not implement in 5.6. However, it is not difficult in MySQL to create partitioning schemes based on `DATE`, `TIME`, or `DATETIME` columns, or based on expressions making use of such columns.

When partitioning by `KEY` or `LINEAR KEY`, you can use a `DATE`, `TIME`, or `DATETIME` column as the partitioning column without performing any modification of the column value. For example, this table creation statement is perfectly valid in MySQL:

```
CREATE TABLE members (
  firstname VARCHAR(25) NOT NULL,
  lastname VARCHAR(25) NOT NULL,
  username VARCHAR(16) NOT NULL,
  email VARCHAR(35),
  joined DATE NOT NULL
)
PARTITION BY KEY(joined)
PARTITIONS 6;
```

In MySQL 5.6, it is also possible to use a `DATE` or `DATETIME` column as the partitioning column using `RANGE COLUMNS` and `LIST COLUMNS` partitioning.

MySQL's other partitioning types, however, require a partitioning expression that yields an integer value or `NULL`. If you wish to use date-based partitioning by `RANGE`, `LIST`, `HASH`, or `LINEAR HASH`, you can



simply employ a function that operates on a `DATE`, `TIME`, or `DATETIME` column and returns such a value, as shown here:

```
CREATE TABLE members (
  firstname VARCHAR(25) NOT NULL,
  lastname VARCHAR(25) NOT NULL,
  username VARCHAR(16) NOT NULL,
  email VARCHAR(35),
  joined DATE NOT NULL
)
PARTITION BY RANGE( YEAR(joined) ) (
  PARTITION p0 VALUES LESS THAN (1960),
  PARTITION p1 VALUES LESS THAN (1970),
  PARTITION p2 VALUES LESS THAN (1980),
  PARTITION p3 VALUES LESS THAN (1990),
  PARTITION p4 VALUES LESS THAN MAXVALUE
);
```

Additional examples of partitioning using dates may be found in the following sections of this chapter:

- [Section 17.2.1, “RANGE Partitioning”](#)
- [Section 17.2.4, “HASH Partitioning”](#)
- [Section 17.2.4.1, “LINEAR HASH Partitioning”](#)

For more complex examples of date-based partitioning, see the following sections:

- [Section 17.4, “Partition Pruning”](#)
- [Section 17.2.6, “Subpartitioning”](#)

MySQL partitioning is optimized for use with the `TO_DAYS()` [1225], `YEAR()` [1229], and `TO_SECONDS()` [1226] functions. However, you can use other date and time functions that return an integer or `NULL`, such as `WEEKDAY()` [1228], `DAYOFYEAR()` [1217], or `MONTH()` [1220]. See [Section 12.7, “Date and Time Functions”](#), for more information about such functions.

It is important to remember—regardless of the type of partitioning that you use—that partitions are always numbered automatically and in sequence when created, starting with 0. When a new row is inserted into a partitioned table, it is these partition numbers that are used in identifying the correct partition. For example, if your table uses 4 partitions, these partitions are numbered 0, 1, 2, and 3. For the `RANGE` and `LIST` partitioning types, it is necessary to ensure that there is a partition defined for each partition number. For `HASH` partitioning, the user function employed must return an integer value greater than 0. For `KEY` partitioning, this issue is taken care of automatically by the hashing function which the MySQL server employs internally.

Names of partitions generally follow the rules governing other MySQL identifiers, such as those for tables and databases. However, you should note that partition names are not case-sensitive. For example, the following `CREATE TABLE` statement fails as shown:

```
mysql> CREATE TABLE t2 (val INT)
-> PARTITION BY LIST(val)(
-> PARTITION mypart VALUES IN (1,3,5),
-> PARTITION MyPart VALUES IN (2,4,6)
-> );
ERROR 1488 (HY000): Duplicate partition name mypart
```

Failure occurs because MySQL sees no difference between the partition names `mypart` and `MyPart`.

When you specify the number of partitions for the table, this must be expressed as a positive, nonzero integer literal with no leading zeros, and may not be an expression such as `0.8E+01` or `6-2`, even if it evaluates to an integer value. Decimal fractions are not permitted.

In the sections that follow, we do not necessarily provide all possible forms for the syntax that can be used for creating each partition type; this information may be found in [Section 13.1.14, “CREATE TABLE Syntax”](#).

## 17.2.1. RANGE Partitioning

A table that is partitioned by range is partitioned in such a way that each partition contains rows for which the partitioning expression value lies within a given range. Ranges should be contiguous but not overlapping, and are defined using the `VALUES LESS THAN` operator. For the next few examples, suppose that you are creating a table such as the following to hold personnel records for a chain of 20 video stores, numbered 1 through 20:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT NOT NULL,
  store_id INT NOT NULL
);
```

This table can be partitioned by range in a number of ways, depending on your needs. One way would be to use the `store_id` column. For instance, you might decide to partition the table 4 ways by adding a `PARTITION BY RANGE` clause as shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT NOT NULL,
  store_id INT NOT NULL
)
PARTITION BY RANGE (store_id) (
  PARTITION p0 VALUES LESS THAN (6),
  PARTITION p1 VALUES LESS THAN (11),
  PARTITION p2 VALUES LESS THAN (16),
  PARTITION p3 VALUES LESS THAN (21)
);
```

In this partitioning scheme, all rows corresponding to employees working at stores 1 through 5 are stored in partition `p0`, to those employed at stores 6 through 10 are stored in partition `p1`, and so on. Note that each partition is defined in order, from lowest to highest. This is a requirement of the `PARTITION BY RANGE` syntax; you can think of it as being analogous to a series of `if ... elseif ...` statements in C or Java in this regard.

It is easy to determine that a new row containing the data (72, 'Michael', 'Widenius', '1998-06-25', NULL, 13) is inserted into partition `p2`, but what happens when your chain adds a 21<sup>st</sup> store? Under this scheme, there is no rule that covers a row whose `store_id` is greater than 20, so an error results because the server does not know where to place it. You can keep this from occurring by using a “catchall” `VALUES LESS THAN` clause in the `CREATE TABLE` statement that provides for all values greater than the highest value explicitly named:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT NOT NULL,
  store_id INT NOT NULL
)
PARTITION BY RANGE (store_id) (
  PARTITION p0 VALUES LESS THAN (6),
  PARTITION p1 VALUES LESS THAN (11),
  PARTITION p2 VALUES LESS THAN (16),
  PARTITION p3 VALUES LESS THAN MAXVALUE
);
```

**Note**

Another way to avoid an error when no matching value is found is to use the `IGNORE` keyword as part of the `INSERT` statement. For an example, see [Section 17.2.2, “LIST Partitioning”](#). Also see [Section 13.2.5, “INSERT Syntax”](#), for general information about `IGNORE`.

`MAXVALUE` represents an integer value that is always greater than the largest possible integer value (in mathematical language, it serves as a *least upper bound*). Now, any rows whose `store_id` column value is greater than or equal to 16 (the highest value defined) are stored in partition `p3`. At some point in the future—when the number of stores has increased to 25, 30, or more—you can use an `ALTER TABLE` statement to add new partitions for stores 21-25, 26-30, and so on (see [Section 17.3, “Partition Management”](#), for details of how to do this).

In much the same fashion, you could partition the table based on employee job codes—that is, based on ranges of `job_code` column values. For example—assuming that two-digit job codes are used for regular (in-store) workers, three-digit codes are used for office and support personnel, and four-digit codes are used for management positions—you could create the partitioned table using the following statement:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT NOT NULL,
  store_id INT NOT NULL
)
PARTITION BY RANGE (job_code) (
  PARTITION p0 VALUES LESS THAN (100),
  PARTITION p1 VALUES LESS THAN (1000),
  PARTITION p2 VALUES LESS THAN (10000)
);
```

In this instance, all rows relating to in-store workers would be stored in partition `p0`, those relating to office and support staff in `p1`, and those relating to managers in partition `p2`.

It is also possible to use an expression in `VALUES LESS THAN` clauses. However, MySQL must be able to evaluate the expression's return value as part of a `LESS THAN (<)` comparison.

Rather than splitting up the table data according to store number, you can use an expression based on one of the two `DATE` columns instead. For example, let us suppose that you wish to partition based on the year that each employee left the company; that is, the value of `YEAR(separated)` [1229]. An example of a `CREATE TABLE` statement that implements such a partitioning scheme is shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
  store_id INT
)
PARTITION BY RANGE ( YEAR(separated) ) (
  PARTITION p0 VALUES LESS THAN (1991),
  PARTITION p1 VALUES LESS THAN (1996),
  PARTITION p2 VALUES LESS THAN (2001),
  PARTITION p3 VALUES LESS THAN MAXVALUE
);
```

In this scheme, for all employees who left before 1991, the rows are stored in partition `p0`; for those who left in the years 1991 through 1995, in `p1`; for those who left in the years 1996 through 2000, in `p2`; and for any workers who left after the year 2000, in `p3`.

It is also possible to partition a table by [RANGE](#), based on the value of a [TIMESTAMP](#) column, using the [UNIX\\_TIMESTAMP\(\)](#) [\[1226\]](#) function, as shown in this example:

```
CREATE TABLE quarterly_report_status (
  report_id INT NOT NULL,
  report_status VARCHAR(20) NOT NULL,
  report_updated TIMESTAMP NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
)
PARTITION BY RANGE ( UNIX_TIMESTAMP(report_updated) ) (
  PARTITION p0 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-01-01 00:00:00') ),
  PARTITION p1 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-04-01 00:00:00') ),
  PARTITION p2 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-07-01 00:00:00') ),
  PARTITION p3 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-10-01 00:00:00') ),
  PARTITION p4 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-01-01 00:00:00') ),
  PARTITION p5 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-04-01 00:00:00') ),
  PARTITION p6 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-07-01 00:00:00') ),
  PARTITION p7 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-10-01 00:00:00') ),
  PARTITION p8 VALUES LESS THAN ( UNIX_TIMESTAMP('2010-01-01 00:00:00') ),
  PARTITION p9 VALUES LESS THAN (MAXVALUE)
);
```

Any other expressions involving [TIMESTAMP](#) values are not permitted. (See Bug #42849.)

Range partitioning is particularly useful when one or more of the following conditions is true:

- You want or need to delete “old” data. If you are using the partitioning scheme shown immediately above, you can simply use `ALTER TABLE employees DROP PARTITION p0;` to delete all rows relating to employees who stopped working for the firm prior to 1991. (See [Section 13.1.6](#), “[ALTER TABLE Syntax](#)”, and [Section 17.3](#), “[Partition Management](#)”, for more information.) For a table with a great many rows, this can be much more efficient than running a `DELETE` query such as `DELETE FROM employees WHERE YEAR(separated) <= 1990;`
- You want to use a column containing date or time values, or containing values arising from some other series.
- You frequently run queries that depend directly on the column used for partitioning the table. For example, when executing a query such as `EXPLAIN PARTITIONS SELECT COUNT(*) FROM employees WHERE separated BETWEEN '2000-01-01' AND '2000-12-31' GROUP BY store_id;`, MySQL can quickly determine that only partition `p2` needs to be scanned because the remaining partitions cannot contain any records satisfying the `WHERE` clause. See [Section 17.4](#), “[Partition Pruning](#)”, for more information about how this is accomplished.

A variant on this type of partitioning is [RANGE COLUMNS](#) partitioning. Partitioning by [RANGE COLUMNS](#) makes it possible to employ multiple columns for defining partitioning ranges that apply both to placement of rows in partitions and for determining the inclusion or exclusion of specific partitions when performing partition pruning. See [Section 17.2.3.1](#), “[RANGE COLUMNS partitioning](#)”, for more information.

**Partitioning schemes based on time intervals.** If you wish to implement a partitioning scheme based on ranges or intervals of time in MySQL 5.6, you have two options:

1. Partition the table by [RANGE](#), and for the partitioning expression, employ a function operating on a [DATE](#), [TIME](#), or [DATETIME](#) column and returning an integer value, as shown here:

```
CREATE TABLE members (
  firstname VARCHAR(25) NOT NULL,
  lastname VARCHAR(25) NOT NULL,
  username VARCHAR(16) NOT NULL,
  email VARCHAR(35),
  joined DATE NOT NULL
)
PARTITION BY RANGE( YEAR(joined) ) (
  PARTITION p0 VALUES LESS THAN (1960),
  PARTITION p1 VALUES LESS THAN (1970),
  PARTITION p2 VALUES LESS THAN (1980),
```

```

PARTITION p3 VALUES LESS THAN (1990),
PARTITION p4 VALUES LESS THAN MAXVALUE
);

```

In MySQL 5.6, it is also possible to partition a table by [RANGE](#) based on the value of a [TIMESTAMP](#) column, using the [UNIX\\_TIMESTAMP\(\)](#) [1226] function, as shown in this example:

```

CREATE TABLE quarterly_report_status (
  report_id INT NOT NULL,
  report_status VARCHAR(20) NOT NULL,
  report_updated TIMESTAMP NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
)
PARTITION BY RANGE ( UNIX_TIMESTAMP(report_updated) ) (
  PARTITION p0 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-01-01 00:00:00') ),
  PARTITION p1 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-04-01 00:00:00') ),
  PARTITION p2 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-07-01 00:00:00') ),
  PARTITION p3 VALUES LESS THAN ( UNIX_TIMESTAMP('2008-10-01 00:00:00') ),
  PARTITION p4 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-01-01 00:00:00') ),
  PARTITION p5 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-04-01 00:00:00') ),
  PARTITION p6 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-07-01 00:00:00') ),
  PARTITION p7 VALUES LESS THAN ( UNIX_TIMESTAMP('2009-10-01 00:00:00') ),
  PARTITION p8 VALUES LESS THAN ( UNIX_TIMESTAMP('2010-01-01 00:00:00') ),
  PARTITION p9 VALUES LESS THAN (MAXVALUE)
);

```

In MySQL 5.6, any other expressions involving [TIMESTAMP](#) values are not permitted. (See Bug #42849.)



#### Note

It is also possible in MySQL 5.6 to use [UNIX\\_TIMESTAMP\(timestamp\\_column\)](#) [1226] as a partitioning expression for tables that are partitioned by [LIST](#). However, it is usually not practical to do so.

- Partition the table by [RANGE COLUMNS](#), using a [DATE](#) or [DATETIME](#) column as the partitioning column. For example, the `members` table could be defined using the `joined` column directly, as shown here:

```

CREATE TABLE members (
  firstname VARCHAR(25) NOT NULL,
  lastname VARCHAR(25) NOT NULL,
  username VARCHAR(16) NOT NULL,
  email VARCHAR(35),
  joined DATE NOT NULL
)
PARTITION BY RANGE COLUMNS(joined) (
  PARTITION p0 VALUES LESS THAN ('1960-01-01'),
  PARTITION p1 VALUES LESS THAN ('1970-01-01'),
  PARTITION p2 VALUES LESS THAN ('1980-01-01'),
  PARTITION p3 VALUES LESS THAN ('1990-01-01'),
  PARTITION p4 VALUES LESS THAN MAXVALUE
);

```



#### Note

The use of partitioning columns employing date or time types other than [DATE](#) or [DATETIME](#) is not supported with [RANGE COLUMNS](#).

## 17.2.2. LIST Partitioning

List partitioning in MySQL is similar to range partitioning in many ways. As in partitioning by [RANGE](#), each partition must be explicitly defined. The chief difference between the two types of partitioning is that, in list partitioning, each partition is defined and selected based on the membership of a column value in one of a set of value lists, rather than in one of a set of contiguous ranges of values. This is done by using [PARTITION BY LIST\(expr\)](#) where `expr` is a column value or an expression based

on a column value and returning an integer value, and then defining each partition by means of a `VALUES IN (value_list)`, where `value_list` is a comma-separated list of integers.



### Note

In MySQL 5.6, it is possible to match against only a list of integers (and possibly `NULL`—see [Section 17.2.7, “How MySQL Partitioning Handles NULL”](#)) when partitioning by `LIST`.

However, other column types may be used in value lists when employing `LIST COLUMN` partitioning, which is described later in this section.

Unlike the case with partitions defined by range, list partitions do not need to be declared in any particular order. For more detailed syntactical information, see [Section 13.1.14, “CREATE TABLE Syntax”](#).

For the examples that follow, we assume that the basic definition of the table to be partitioned is provided by the `CREATE TABLE` statement shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
  store_id INT
);
```

(This is the same table used as a basis for the examples in [Section 17.2.1, “RANGE Partitioning”](#).)

Suppose that there are 20 video stores distributed among 4 franchises as shown in the following table.

Region	Store ID Numbers
North	3, 5, 6, 9, 17
East	1, 2, 10, 11, 19, 20
West	4, 12, 13, 14, 18
Central	7, 8, 15, 16

To partition this table in such a way that rows for stores belonging to the same region are stored in the same partition, you could use the `CREATE TABLE` statement shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
  store_id INT
)
PARTITION BY LIST(store_id) (
  PARTITION pNorth VALUES IN (3,5,6,9,17),
  PARTITION pEast VALUES IN (1,2,10,11,19,20),
  PARTITION pWest VALUES IN (4,12,13,14,18),
  PARTITION pCentral VALUES IN (7,8,15,16)
);
```

This makes it easy to add or drop employee records relating to specific regions to or from the table. For instance, suppose that all stores in the West region are sold to another company. In MySQL 5.6, all rows relating to employees working at stores in that region can be deleted with the query `ALTER TABLE employees TRUNCATE PARTITION pWest`, which can be executed much more efficiently than the equivalent `DELETE` statement `DELETE FROM employees WHERE store_id IN`

(4,12,13,14,18);. (Using `ALTER TABLE employees DROP PARTITION pWest` would also delete all of these rows, but would also remove the partition `pWest` from the definition of the table; you would need to use an `ALTER TABLE ... ADD PARTITION` statement to restore the table's original partitioning scheme.)

As with `RANGE` partitioning, it is possible to combine `LIST` partitioning with partitioning by hash or key to produce a composite partitioning (subpartitioning). See [Section 17.2.6, "Subpartitioning"](#).

Unlike the case with `RANGE` partitioning, there is no "catch-all" such as `MAXVALUE`; all expected values for the partitioning expression should be covered in `PARTITION ... VALUES IN (...)` clauses. An `INSERT` statement containing an unmatched partitioning column value fails with an error, as shown in this example:

```
mysql> CREATE TABLE h2 (
->   c1 INT,
->   c2 INT
-> )
-> PARTITION BY LIST(c1) (
->   PARTITION p0 VALUES IN (1, 4, 7),
->   PARTITION p1 VALUES IN (2, 5, 8)
-> );
Query OK, 0 rows affected (0.11 sec)

mysql> INSERT INTO h2 VALUES (3, 5);
ERROR 1525 (HY000): Table has no partition for value 3
```

When inserting multiple rows using a single `INSERT` statement, any rows coming before the row containing the unmatched value are inserted, but any coming after it are not:

```
mysql> SELECT * FROM h2;
Empty set (0.00 sec)

mysql> INSERT INTO h2 VALUES (4, 7), (3, 5), (6, 0);
ERROR 1525 (HY000): Table has no partition for value 3
mysql> SELECT * FROM h2;
+-----+-----+
| c1   | c2   |
+-----+-----+
|    4 |    7 |
+-----+-----+
1 row in set (0.00 sec)
```

You can cause this type of error to be ignored by using the `IGNORE` keyword. If you do so, rows containing unmatched partitioning column values are not inserted, but any rows with matching values are inserted, and no errors are reported:

```
mysql> TRUNCATE h2;
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM h2;
Empty set (0.00 sec)

mysql> INSERT IGNORE INTO h2 VALUES (2, 5), (6, 10), (7, 5), (3, 1), (1, 9);
Query OK, 3 rows affected (0.00 sec)
Records: 5  Duplicates: 2  Warnings: 0

mysql> SELECT * FROM h2;
+-----+-----+
| c1   | c2   |
+-----+-----+
|    7 |    5 |
|    1 |    9 |
|    2 |    5 |
+-----+-----+
3 rows in set (0.00 sec)
```

MySQL 5.6 provides support for `LIST COLUMNS` partitioning. This is a variant of `LIST` partitioning that enables you to use columns of types other than integer types for partitioning columns, as well as to use



multiple columns as partitioning keys. For more information, see [Section 17.2.3.2, “LIST COLUMNS partitioning”](#).

### 17.2.3. COLUMNS Partitioning

The next two sections discuss *COLUMNS partitioning*, which are variants on [RANGE](#) and [LIST](#) partitioning. [COLUMNS](#) partitioning enables the use of multiple columns in partitioning keys. All of these columns are taken into account both for the purpose of placing rows in partitions and for the determination of which partitions are to be checked for matching rows in partition pruning.

In addition, both [RANGE COLUMNS](#) partitioning and [LIST COLUMNS](#) partitioning support the use of non-integer columns for defining value ranges or list members. The permitted data types are shown in the following list:

- All integer types: [TINYINT](#), [SMALLINT](#), [MEDIUMINT](#), [INT \(INTEGER\)](#), and [BIGINT](#). (This is the same as with partitioning by [RANGE](#) and [LIST](#).)

Other numeric data types (such as [DECIMAL](#) or [FLOAT](#)) are not supported as partitioning columns.

- [DATE](#) and [DATETIME](#).

Columns using other data types relating to dates or times are not supported as partitioning columns.

- The following string types: [CHAR](#), [VARCHAR](#), [BINARY](#), and [VARBINARY](#).

[TEXT](#) and [BLOB](#) columns are not supported as partitioning columns.

The discussions of [RANGE COLUMNS](#) and [LIST COLUMNS](#) partitioning in the next two sections assume that you are already familiar with partitioning based on ranges and lists as supported in MySQL 5.1 and later; for more information about these, see [Section 17.2.1, “RANGE Partitioning”](#), and [Section 17.2.2, “LIST Partitioning”](#), respectively.

#### 17.2.3.1. RANGE COLUMNS partitioning

Range columns partitioning is similar to range partitioning, but enables you to define partitions using ranges based on multiple column values. In addition, you can define the ranges using columns of types other than integer types.

[RANGE COLUMNS](#) partitioning differs significantly from [RANGE](#) partitioning in the following ways:

- [RANGE COLUMNS](#) does not accept expressions, only names of columns.
- [RANGE COLUMNS](#) accepts a list of one or more columns.

[RANGE COLUMNS](#) partitions are based on comparisons between *tuples* (lists of column values) rather than comparisons between scalar values. Placement of rows in [RANGE COLUMNS](#) partitions is also based on comparisons between tuples; this is discussed further later in this section.

- [RANGE COLUMNS](#) partitioning columns are not restricted to integer columns; string, [DATE](#) and [DATETIME](#) columns can also be used as partitioning columns. (See [Section 17.2.3, “COLUMNS Partitioning”](#), for details.)

The basic syntax for creating a table partitioned by [RANGE COLUMNS](#) is shown here:

```
CREATE TABLE table_name
PARTITIONED BY RANGE COLUMNS(column_list) (
  PARTITION partition_name VALUES LESS THAN (value_list)[,
  PARTITION partition_name VALUES LESS THAN (value_list)][,
  ...]
)

column_list:
  column_name[, column_name][, ...]
```

```
value_list:
  value[, value][, ...]
```

**Note**

Not all `CREATE TABLE` options that can be used when creating partitioned tables are shown here. For complete information, see [Section 13.1.14, “CREATE TABLE Syntax”](#).

In the syntax just shown, *column\_list* is a list of one or more columns (sometimes called a *partitioning column list*), and *value\_list* is a list of values (that is, it is a *partition definition value list*). A *value\_list* must be supplied for each partition definition, and each *value\_list* must have the same number of values as the *column\_list* has columns. Generally speaking, if you use *N* columns in the `COLUMNS` clause, then each `VALUES LESS THAN` clause must also be supplied with a list of *N* values.

The elements in the partitioning column list and in the value list defining each partition must occur in the same order. In addition, each element in the value list must be of the same data type as the corresponding element in the column list. However, the order of the column names in the partitioning column list and the value lists does not have to be the same as the order of the table column definitions in the main part of the `CREATE TABLE` statement. As with table partitioned by `RANGE`, you can use `MAXVALUE` to represent a value such that any legal value inserted into a given column is always less than this value. Here is an example of a `CREATE TABLE` statement that helps to illustrate all of these points:

```
mysql> CREATE TABLE rcx (
->   a INT,
->   b INT,
->   c CHAR(3),
->   d INT
-> )
-> PARTITION BY RANGE COLUMNS(a,d,c) (
->   PARTITION p0 VALUES LESS THAN (5,10,'ggg'),
->   PARTITION p1 VALUES LESS THAN (10,20,'mmm'),
->   PARTITION p2 VALUES LESS THAN (15,30,'sss'),
->   PARTITION p3 VALUES LESS THAN (MAXVALUE,MAXVALUE,MAXVALUE)
-> );
Query OK, 0 rows affected (0.15 sec)
```

Table `rcx` contains the columns `a`, `b`, `c`, `d`. The partitioning column list supplied to the `COLUMNS` clause uses 3 of these columns, in the order `a`, `d`, `c`. Each value list used to define a partition contains 3 values in the same order; that is, each value list tuple has the form `(INT, INT, CHAR(3))`, which corresponds to the data types used by columns `a`, `d`, and `c` (in that order).

Placement of rows into partitions is determined by comparing the tuple from a row to be inserted that matches the column list in the `COLUMNS` clause with the tuples used in the `VALUES LESS THAN` clauses to define partitions of the table. Because we are comparing tuples (that is, lists or sets of values) rather than scalar values, the semantics of `VALUES LESS THAN` as used with `RANGE COLUMNS` partitions differs somewhat from the case with simple `RANGE` partitions. In `RANGE` partitioning, a row generating an expression value that is equal to a limiting value in a `VALUES LESS THAN` is never placed in the corresponding partition; however, when using `RANGE COLUMNS` partitioning, it is sometimes possible for a row whose partitioning column list's first element is equal in value to the that of the first element in a `VALUES LESS THAN` value list to be placed in the corresponding partition.

Consider the `RANGE` partitioned table created by this statement:

```
CREATE TABLE r1 (
  a INT,
  b INT
)
PARTITION BY RANGE (a) (
  PARTITION p0 VALUES LESS THAN (5),
  PARTITION p1 VALUES LESS THAN (MAXVALUE)
);
```

If we insert 3 rows into this table such that the column value for `a` is 5 for each row, all 3 rows are stored in partition `p1` because the `a` column value is in each case not less than 5, as we can see by executing the proper query against the `INFORMATION_SCHEMA.PARTITIONS` table:

```
mysql> INSERT INTO r1 VALUES (5,10), (5,11), (5,12);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT PARTITION_NAME, TABLE_ROWS
  -> FROM INFORMATION_SCHEMA.PARTITIONS
  -> WHERE TABLE_NAME = 'r1';
+-----+-----+
| PARTITION_NAME | TABLE_ROWS |
+-----+-----+
| p0              |            0 |
| p1              |            3 |
+-----+-----+
2 rows in set (0.00 sec)
```

Now consider a similar table `rc1` that uses `RANGE COLUMNS` partitioning with both columns `a` and `b` referenced in the `COLUMNS` clause, created as shown here:

```
CREATE TABLE rc1 (
  a INT,
  b INT
)
PARTITION BY RANGE COLUMNS(a, b) (
  PARTITION p0 VALUES LESS THAN (5, 12),
  PARTITION p3 VALUES LESS THAN (MAXVALUE, MAXVALUE)
);
```

If we insert exactly the same rows into `rc1` as we just inserted into `r1`, the distribution of the rows is quite different:

```
mysql> INSERT INTO rc1 VALUES (5,10), (5,11), (5,12);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT PARTITION_NAME, TABLE_ROWS
  -> FROM INFORMATION_SCHEMA.PARTITIONS
  -> WHERE TABLE_NAME = 'rc1';
+-----+-----+-----+
| TABLE_SCHEMA | PARTITION_NAME | TABLE_ROWS |
+-----+-----+-----+
| p              | p0              |            2 |
| p              | p1              |            1 |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

This is because we are comparing rows rather than scalar values. We can compare the row values inserted with the limiting row value from the `VALUES THAN LESS THAN` clause used to define partition `p0` in table `rc1`, like this:

```
mysql> SELECT (5,10) < (5,12), (5,11) < (5,12), (5,12) < (5,12);
+-----+-----+-----+
| (5,10) < (5,12) | (5,11) < (5,12) | (5,12) < (5,12) |
+-----+-----+-----+
|                1 |                1 |                0 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

The 2 tuples `(5,10)` and `(5,11)` evaluate as less than `(5,12)`, so they are stored in partition `p0`. Since 5 is not less than 5 and 12 is not less than 12, `(5,12)` is considered not less than `(5,12)`, and is stored in partition `p1`.

The `SELECT` statement in the preceding example could also have been written using explicit row constructors, like this:

```
SELECT ROW(5,10) < ROW(5,12), ROW(5,11) < ROW(5,12), ROW(5,12) < ROW(5,12);
```

For more information about the use of row constructors in MySQL, see [Section 13.2.10.5, “Row Subqueries”](#).

For a table partitioned by `RANGE COLUMNS` using only a single partitioning column, the storing of rows in partitions is the same as that of an equivalent table that is partitioned by `RANGE`. The following `CREATE TABLE` statement creates a table partitioned by `RANGE COLUMNS` using 1 partitioning column:

```
CREATE TABLE rx (
  a INT,
  b INT
)
PARTITION BY RANGE COLUMNS (a) (
  PARTITION p0 VALUES LESS THAN (5),
  PARTITION p1 VALUES LESS THAN (MAXVALUE)
);
```

If we insert the rows `(5,10)`, `(5,11)`, and `(5,12)` into this table, we can see that their placement is the same as it is for the table `r` we created and populated earlier:

```
mysql> INSERT INTO rx VALUES (5,10), (5,11), (5,12);
Query OK, 3 rows affected (0.00 sec)
Records: 3  Duplicates: 0  Warnings: 0

mysql> SELECT PARTITION_NAME, TABLE_ROWS
-> FROM INFORMATION_SCHEMA.PARTITIONS
-> WHERE TABLE_NAME = 'rx';
+-----+-----+-----+
| TABLE_SCHEMA | PARTITION_NAME | TABLE_ROWS |
+-----+-----+-----+
| p             | p0             | 0           |
| p             | p1             | 3           |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

It is also possible to create tables partitioned by `RANGE COLUMNS` where limiting values for one or more columns are repeated in successive partition definitions. You can do this as long as the tuples of column values used to define the partitions are strictly increasing. For example, each of the following `CREATE TABLE` statements is valid:

```
CREATE TABLE rc2 (
  a INT,
  b INT
)
PARTITION BY RANGE COLUMNS(a,b) (
  PARTITION p0 VALUES LESS THAN (0,10),
  PARTITION p1 VALUES LESS THAN (10,20),
  PARTITION p2 VALUES LESS THAN (10,30),
  PARTITION p3 VALUES LESS THAN (MAXVALUE,MAXVALUE)
);

CREATE TABLE rc3 (
  a INT,
  b INT
)
PARTITION BY RANGE COLUMNS(a,b) (
  PARTITION p0 VALUES LESS THAN (0,10),
  PARTITION p1 VALUES LESS THAN (10,20),
  PARTITION p2 VALUES LESS THAN (10,30),
  PARTITION p3 VALUES LESS THAN (10,35),
  PARTITION p4 VALUES LESS THAN (20,40),
  PARTITION p5 VALUES LESS THAN (MAXVALUE,MAXVALUE)
);
```

The following statement also succeeds, even though it might appear at first glance that it would not, since the limiting value of column `b` is 25 for partition `p0` and 20 for partition `p1`, and the limiting value of column `c` is 100 for partition `p1` and 50 for partition `p2`:

```
CREATE TABLE rc4 (
  a INT,
```

```

    b INT,
    c INT
)
PARTITION BY RANGE COLUMNS(a,b,c) (
    PARTITION p0 VALUES LESS THAN (0,25,50),
    PARTITION p1 VALUES LESS THAN (10,20,100),
    PARTITION p2 VALUES LESS THAN (10,30,50)
    PARTITION p3 VALUES LESS THAN (MAXVALUE,MAXVALUE,MAXVALUE)
);

```

When designing tables partitioned by `RANGE COLUMNS`, you can always test successive partition definitions by comparing the desired tuples using the `mysql` client, like this:

```

mysql> SELECT (0,25,50) < (10,20,100), (10,20,100) < (10,30,50);
+-----+-----+
| (0,25,50) < (10,20,100) | (10,20,100) < (10,30,50) |
+-----+-----+
| 1 | 1 |
+-----+-----+
1 row in set (0.00 sec)

```

If a `CREATE TABLE` statement contains partition definitions that are not in strictly increasing order, it fails with an error, as shown in this example:

```

mysql> CREATE TABLE rcf (
->   a INT,
->   b INT,
->   c INT
-> )
-> PARTITION BY RANGE COLUMNS(a,b,c) (
->   PARTITION p0 VALUES LESS THAN (0,25,50),
->   PARTITION p1 VALUES LESS THAN (20,20,100),
->   PARTITION p2 VALUES LESS THAN (10,30,50),
->   PARTITION p3 VALUES LESS THAN (MAXVALUE,MAXVALUE,MAXVALUE)
-> );
ERROR 1493 (HY000): VALUES LESS THAN value must be strictly increasing for each partition

```

When you get such an error, you can deduce which partition definitions are invalid by making “less than” comparisons between their column lists. In this case, the problem is with the definition of partition `p2` because the tuple used to define it is not less than the tuple used to define partition `p3`, as shown here:

```

mysql> SELECT (0,25,50) < (20,20,100), (20,20,100) < (10,30,50);
+-----+-----+
| (0,25,50) < (20,20,100) | (20,20,100) < (10,30,50) |
+-----+-----+
| 1 | 0 |
+-----+-----+
1 row in set (0.00 sec)

```

It is also possible for `MAXVALUE` to appear for the same column in more than one `VALUES LESS THAN` clause when using `RANGE COLUMNS`. However, the limiting values for individual columns in successive partition definitions should otherwise be increasing, there should be no more than one partition defined where `MAXVALUE` is used as the upper limit for all column values, and this partition definition should appear last in the list of `PARTITION ... VALUES LESS THAN` clauses. In addition, you cannot use `MAXVALUE` as the limiting value for the first column in more than one partition definition.

As stated previously, it is also possible with `RANGE COLUMNS` partitioning to use non-integer columns as partitioning columns. (See [Section 17.2.3, “COLUMNS Partitioning”](#), for a complete listing of these.) Consider a table named `employees` (which is not partitioned), created using the following statement:

```

CREATE TABLE employees (
    id INT NOT NULL,
    fname VARCHAR(30),
    lname VARCHAR(30),
    hired DATE NOT NULL DEFAULT '1970-01-01',
    separated DATE NOT NULL DEFAULT '9999-12-31',
    job_code INT NOT NULL,

```

```
store_id INT NOT NULL
);
```

Using [RANGE COLUMNS](#) partitioning, you can create a version of this table that stores each row in one of four partitions based on the employee's last name, like this:

```
CREATE TABLE employees_by_lname (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT NOT NULL,
  store_id INT NOT NULL
)
PARTITION BY RANGE COLUMNS (lname) (
  PARTITION p0 VALUES LESS THAN ('g'),
  PARTITION p1 VALUES LESS THAN ('m'),
  PARTITION p2 VALUES LESS THAN ('t'),
  PARTITION p3 VALUES LESS THAN (MAXVALUE)
);
```

Alternatively, you could cause the `employees` table as created previously to be partitioned using this scheme by executing the following [ALTER TABLE](#) statement:

```
ALTER TABLE employees PARTITION BY RANGE COLUMNS (lname) (
  PARTITION p0 VALUES LESS THAN ('g'),
  PARTITION p1 VALUES LESS THAN ('m'),
  PARTITION p2 VALUES LESS THAN ('t'),
  PARTITION p3 VALUES LESS THAN (MAXVALUE)
);
```



#### Note

Because different character sets and collations have different sort orders, the character sets and collations in use may effect which partition of a table partitioned by [RANGE COLUMNS](#) a given row is stored in when using string columns as partitioning columns. In addition, changing the character set or collation for a given database, table, or column after such a table is created may cause changes in how rows are distributed. For example, when using a case-sensitive collation, `'and'` sorts before `'Andersen'`, but when using a collation that is case insensitive, the reverse is true.

For information about how MySQL handles character sets and collations, see [Section 10.1, “Character Set Support”](#).

Similarly, you can cause the `employees` table to be partitioned in such a way that each row is stored in one of several partitions based on the decade in which the corresponding employee was hired using the [ALTER TABLE](#) statement shown here:

```
ALTER TABLE employees PARTITION BY RANGE COLUMNS (hired) (
  PARTITION p0 VALUES LESS THAN ('1970-01-01'),
  PARTITION p1 VALUES LESS THAN ('1980-01-01'),
  PARTITION p2 VALUES LESS THAN ('1990-01-01'),
  PARTITION p3 VALUES LESS THAN ('2000-01-01'),
  PARTITION p4 VALUES LESS THAN ('2010-01-01'),
  PARTITION p5 VALUES LESS THAN (MAXVALUE)
);
```

See [Section 13.1.14, “CREATE TABLE Syntax”](#), for additional information about [PARTITION BY RANGE COLUMNS](#) syntax.

### 17.2.3.2. LIST COLUMNS partitioning

MySQL 5.6 provides support for [LIST COLUMNS](#) partitioning. This is a variant of [LIST](#) partitioning that enables the use of multiple columns as partition keys, and for columns of data types other than

integer types to be used as partitioning columns; you can use string types, [DATE](#), and [DATETIME](#) columns. (For more information about permitted data types for [COLUMNS](#) partitioning columns, see [Section 17.2.3, "COLUMNS Partitioning"](#).)

Suppose that you have a business that has customers in 12 cities which, for sales and marketing purposes, you organize into 4 regions of 3 cities each as shown in the following table:

Region	Cities
1	Oskarshamn, Högsby, Mönsterås
2	Vimmerby, Hultsfred, Västervik
3	Nässjö, Eksjö, Vetlanda
4	Uppvidinge, Alvesta, Växjö

With [LIST COLUMNS](#) partitioning, you can create a table for customer data that assigns a row to any of 4 partitions corresponding to these regions based on the name of the city where a customer resides, as shown here:

```
CREATE TABLE customers_1 (
  first_name VARCHAR(25),
  last_name VARCHAR(25),
  street_1 VARCHAR(30),
  street_2 VARCHAR(30),
  city VARCHAR(15),
  renewal DATE
)
PARTITION BY LIST COLUMNS(city) (
  PARTITION pRegion_1 VALUES IN('Oskarshamn', 'Högsby', 'Mönsterås'),
  PARTITION pRegion_2 VALUES IN('Vimmerby', 'Hultsfred', 'Västervik'),
  PARTITION pRegion_3 VALUES IN('Nässjö', 'Eksjö', 'Vetlanda'),
  PARTITION pRegion_4 VALUES IN('Uppvidinge', 'Alvesta', 'Växjö')
);
```

As with partitioning by [RANGE COLUMNS](#), you do not need to use expressions in the [COLUMNS\(\)](#) clause to convert column values into integers. (In fact, the use of expressions other than column names is not permitted with [COLUMNS\(\)](#).)

It is also possible to use [DATE](#) and [DATETIME](#) columns, as shown in the following example that uses the same name and columns as the `customers_1` table shown previously, but employs [LIST COLUMNS](#) partitioning based on the `renewal` column to store rows in one of 4 partitions depending on the week in February 2010 the customer's account is scheduled to renew:

```
CREATE TABLE customers_2 (
  first_name VARCHAR(25),
  last_name VARCHAR(25),
  street_1 VARCHAR(30),
  street_2 VARCHAR(30),
  city VARCHAR(15),
  renewal DATE
)
PARTITION BY LIST COLUMNS(renewal) (
  PARTITION pWeek_1 VALUES IN('2010-02-01', '2010-02-02', '2010-02-03',
    '2010-02-04', '2010-02-05', '2010-02-06', '2010-02-07'),
  PARTITION pWeek_2 VALUES IN('2010-02-08', '2010-02-09', '2010-02-10',
    '2010-02-11', '2010-02-12', '2010-02-13', '2010-02-14'),
  PARTITION pWeek_3 VALUES IN('2010-02-15', '2010-02-16', '2010-02-17',
    '2010-02-18', '2010-02-19', '2010-02-20', '2010-02-21'),
  PARTITION pWeek_4 VALUES IN('2010-02-22', '2010-02-23', '2010-02-24',
    '2010-02-25', '2010-02-26', '2010-02-27', '2010-02-28')
);
```

This works, but becomes cumbersome to define and maintain if the number of dates involved grows very large; in such cases, it is usually more practical to employ [RANGE](#) or [RANGE COLUMNS](#) partitioning instead. In this case, since the column we wish to use as the partitioning key is a [DATE](#) column, we use [RANGE COLUMNS](#) partitioning, as shown here:



```
CREATE TABLE customers_3 (
  first_name VARCHAR(25),
  last_name VARCHAR(25),
  street_1 VARCHAR(30),
  street_2 VARCHAR(30),
  city VARCHAR(15),
  renewal DATE
)
PARTITION BY RANGE COLUMNS(renewal) (
  PARTITION pWeek_1 VALUES LESS THAN('2010-02-09'),
  PARTITION pWeek_2 VALUES LESS THAN('2010-02-15'),
  PARTITION pWeek_3 VALUES LESS THAN('2010-02-22'),
  PARTITION pWeek_4 VALUES LESS THAN('2010-03-01')
);
```

See [Section 17.2.3.1, “RANGE COLUMNS partitioning”](#), for more information.

In addition (as with [RANGE COLUMNS](#) partitioning), you can use multiple columns in the [COLUMNS\(\)](#) clause.

See [Section 13.1.14, “CREATE TABLE Syntax”](#), for additional information about [PARTITION BY LIST COLUMNS\(\)](#) syntax.

## 17.2.4. HASH Partitioning

Partitioning by [HASH](#) is used primarily to ensure an even distribution of data among a predetermined number of partitions. With range or list partitioning, you must specify explicitly into which partition a given column value or set of column values is to be stored; with hash partitioning, MySQL takes care of this for you, and you need only specify a column value or expression based on a column value to be hashed and the number of partitions into which the partitioned table is to be divided.

To partition a table using [HASH](#) partitioning, it is necessary to append to the [CREATE TABLE](#) statement a [PARTITION BY HASH \(expr\)](#) clause, where *expr* is an expression that returns an integer. This can simply be the name of a column whose type is one of MySQL's integer types. In addition, you will most likely want to follow this with a [PARTITIONS num](#) clause, where *num* is a positive integer representing the number of partitions into which the table is to be divided.

For example, the following statement creates a table that uses hashing on the `store_id` column and is divided into 4 partitions:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
  store_id INT
)
PARTITION BY HASH(store_id)
PARTITIONS 4;
```

If you do not include a [PARTITIONS](#) clause, the number of partitions defaults to 1.

Using the [PARTITIONS](#) keyword without a number following it results in a syntax error.

You can also use an SQL expression that returns an integer for *expr*. For instance, you might want to partition based on the year in which an employee was hired. This can be done as shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
```

```

    store_id INT
)
PARTITION BY HASH( YEAR(hired) )
PARTITIONS 4;

```

*expr* must return a nonconstant, nonrandom integer value (in other words, it should be varying but deterministic), and must not contain any prohibited constructs as described in [Section 17.6, “Restrictions and Limitations on Partitioning”](#). You should also keep in mind that this expression is evaluated each time a row is inserted or updated (or possibly deleted); this means that very complex expressions may give rise to performance issues, particularly when performing operations (such as batch inserts) that affect a great many rows at one time.

The most efficient hashing function is one which operates upon a single table column and whose value increases or decreases consistently with the column value, as this allows for “pruning” on ranges of partitions. That is, the more closely that the expression varies with the value of the column on which it is based, the more efficiently MySQL can use the expression for hash partitioning.

For example, where *date\_col* is a column of type `DATE`, then the expression `TO_DAYS(date_col)` [1225] is said to vary directly with the value of *date\_col*, because for every change in the value of *date\_col*, the value of the expression changes in a consistent manner. The variance of the expression `YEAR(date_col)` [1229] with respect to *date\_col* is not quite as direct as that of `TO_DAYS(date_col)` [1225], because not every possible change in *date\_col* produces an equivalent change in `YEAR(date_col)` [1229]. Even so, `YEAR(date_col)` [1229] is a good candidate for a hashing function, because it varies directly with a portion of *date\_col* and there is no possible change in *date\_col* that produces a disproportionate change in `YEAR(date_col)` [1229].

By way of contrast, suppose that you have a column named *int\_col* whose type is `INT`. Now consider the expression `POW(5-int_col,3) + 6` [1206]. This would be a poor choice for a hashing function because a change in the value of *int\_col* is not guaranteed to produce a proportional change in the value of the expression. Changing the value of *int\_col* by a given amount can produce by widely different changes in the value of the expression. For example, changing *int\_col* from 5 to 6 produces a change of `-1` in the value of the expression, but changing the value of *int\_col* from 6 to 7 produces a change of `-7` in the expression value.

In other words, the more closely the graph of the column value *versus* the value of the expression follows a straight line as traced by the equation  $y=cx$  where *c* is some nonzero constant, the better the expression is suited to hashing. This has to do with the fact that the more nonlinear an expression is, the more uneven the distribution of data among the partitions it tends to produce.

In theory, pruning is also possible for expressions involving more than one column value, but determining which of such expressions are suitable can be quite difficult and time-consuming. For this reason, the use of hashing expressions involving multiple columns is not particularly recommended.

When `PARTITION BY HASH` is used, MySQL determines which partition of *num* partitions to use based on the modulus of the result of the user function. In other words, for an expression *expr*, the partition in which the record is stored is partition number *N*, where  $N = \text{MOD}(\text{expr}, \text{num})$ . Suppose that table *t1* is defined as follows, so that it has 4 partitions:

```

CREATE TABLE t1 (col1 INT, col2 CHAR(5), col3 DATE)
PARTITION BY HASH( YEAR(col3) )
PARTITIONS 4;

```

If you insert a record into *t1* whose *col3* value is `'2005-09-15'`, then the partition in which it is stored is determined as follows:

```

MOD(YEAR('2005-09-01'),4)
= MOD(2005,4)
= 1

```

MySQL 5.6 also supports a variant of `HASH` partitioning known as *linear hashing* which employs a more complex algorithm for determining the placement of new rows inserted into the partitioned table. See [Section 17.2.4.1, “LINEAR HASH Partitioning”](#), for a description of this algorithm.

The user function is evaluated each time a record is inserted or updated. It may also—depending on the circumstances—be evaluated when records are deleted.



### Note

If a table to be partitioned has a [UNIQUE](#) key, then any columns supplied as arguments to the [HASH](#) user function or to the [KEY](#)'s *column\_list* must be part of that key.

#### 17.2.4.1. LINEAR HASH Partitioning

MySQL also supports linear hashing, which differs from regular hashing in that linear hashing utilizes a linear powers-of-two algorithm whereas regular hashing employs the modulus of the hashing function's value.

Syntactically, the only difference between linear-hash partitioning and regular hashing is the addition of the [LINEAR](#) keyword in the [PARTITION BY](#) clause, as shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30),
  hired DATE NOT NULL DEFAULT '1970-01-01',
  separated DATE NOT NULL DEFAULT '9999-12-31',
  job_code INT,
  store_id INT
)
PARTITION BY LINEAR HASH( YEAR(hired) )
PARTITIONS 4;
```

Given an expression *expr*, the partition in which the record is stored when linear hashing is used is partition number *N* from among *num* partitions, where *N* is derived according to the following algorithm:

1. Find the next power of 2 greater than *num*. We call this value *V*; it can be calculated as:

```
V = POWER(2, CEILING(LOG(2, num)))
```

(Suppose that *num* is 13. Then  $\text{LOG}(2, 13)$  [1204] is 3.7004397181411.  $\text{CEILING}(3.7004397181411)$  [1203] is 4, and  $V = \text{POWER}(2, 4)$  [1206], which is 16.)

2. Set  $N = F(\text{column\_list}) \& (V - 1)$ .
3. While  $N \geq \text{num}$ :
  - Set  $V = \text{CEIL}(V / 2)$
  - Set  $N = N \& (V - 1)$

Suppose that the table `t1`, using linear hash partitioning and having 6 partitions, is created using this statement:

```
CREATE TABLE t1 (col1 INT, col2 CHAR(5), col3 DATE)
PARTITION BY LINEAR HASH( YEAR(col3) )
PARTITIONS 6;
```

Now assume that you want to insert two records into `t1` having the `col3` column values '2003-04-14' and '1998-10-19'. The partition number for the first of these is determined as follows:

```
V = POWER(2, CEILING( LOG(2,6) )) = 8
N = YEAR('2003-04-14') & (8 - 1)
  = 2003 & 7
  = 3
```

```
(3 >= 6 is FALSE: record stored in partition #3)
```

The number of the partition where the second record is stored is calculated as shown here:

```
V = 8
N = YEAR('1998-10-19') & (8-1)
  = 1998 & 7
  = 6

(6 >= 6 is TRUE: additional step required)

N = 6 & CEILING(8 / 2)
  = 6 & 3
  = 2

(2 >= 6 is FALSE: record stored in partition #2)
```

The advantage in partitioning by linear hash is that the adding, dropping, merging, and splitting of partitions is made much faster, which can be beneficial when dealing with tables containing extremely large amounts (terabytes) of data. The disadvantage is that data is less likely to be evenly distributed between partitions as compared with the distribution obtained using regular hash partitioning.

## 17.2.5. KEY Partitioning

Partitioning by key is similar to partitioning by hash, except that where hash partitioning employs a user-defined expression, the hashing function for key partitioning is supplied by the MySQL server. This internal hashing function is based on the same algorithm as `PASSWORD()` [1263].

The syntax rules for `CREATE TABLE ... PARTITION BY KEY` are similar to those for creating a table that is partitioned by hash. The major differences are listed here:

- `KEY` is used rather than `HASH`.
- `KEY` takes only a list of zero or more column names. Any columns used as the partitioning key must comprise part or all of the table's primary key, if the table has one. Where no column name is specified as the partitioning key, the table's primary key is used, if there is one. For example, the following `CREATE TABLE` statement is valid in MySQL 5.6:

```
CREATE TABLE k1 (
  id INT NOT NULL PRIMARY KEY,
  name VARCHAR(20)
)
PARTITION BY KEY()
PARTITIONS 2;
```

If there is no primary key but there is a unique key, then the unique key is used for the partitioning key:

```
CREATE TABLE k1 (
  id INT NOT NULL,
  name VARCHAR(20),
  UNIQUE KEY (id)
)
PARTITION BY KEY()
PARTITIONS 2;
```

However, if the unique key column were not defined as `NOT NULL`, then the previous statement would fail.

In both of these cases, the partitioning key is the `id` column, even though it is not shown in the output of `SHOW CREATE TABLE` or in the `PARTITION_EXPRESSION` column of the `INFORMATION_SCHEMA.PARTITIONS` table.

Unlike the case with other partitioning types, columns used for partitioning by `KEY` are not restricted to integer or `NULL` values. For example, the following `CREATE TABLE` statement is valid:

```
CREATE TABLE tm1 (
  s1 CHAR(32) PRIMARY KEY
)
PARTITION BY KEY(s1)
PARTITIONS 10;
```

The preceding statement would *not* be valid, were a different partitioning type to be specified. (In this case, simply using `PARTITION BY KEY()` would also be valid and have the same effect as `PARTITION BY KEY(s1)`, since `s1` is the table's primary key.)

For additional information about this issue, see [Section 17.6, “Restrictions and Limitations on Partitioning”](#).



### Important

For a key-partitioned table, you cannot execute an `ALTER TABLE DROP PRIMARY KEY`, as doing so generates the error `ERROR 1466 (HY000): Field in list of fields for partition function not found in table.`

It is also possible to partition a table by linear key. Here is a simple example:

```
CREATE TABLE tk (
  col1 INT NOT NULL,
  col2 CHAR(5),
  col3 DATE
)
PARTITION BY LINEAR KEY (col1)
PARTITIONS 3;
```

Using `LINEAR` has the same effect on `KEY` partitioning as it does on `HASH` partitioning, with the partition number being derived using a powers-of-two algorithm rather than modulo arithmetic. See [Section 17.2.4.1, “LINEAR HASH Partitioning”](#), for a description of this algorithm and its implications.

## 17.2.6. Subpartitioning

Subpartitioning—also known as *composite partitioning*—is the further division of each partition in a partitioned table. Consider the following `CREATE TABLE` statement:

```
CREATE TABLE ts (id INT, purchased DATE)
PARTITION BY RANGE( YEAR(purchased) )
SUBPARTITION BY HASH( TO_DAYS(purchased) )
SUBPARTITIONS 2 (
  PARTITION p0 VALUES LESS THAN (1990),
  PARTITION p1 VALUES LESS THAN (2000),
  PARTITION p2 VALUES LESS THAN MAXVALUE
);
```

Table `ts` has 3 `RANGE` partitions. Each of these partitions—`p0`, `p1`, and `p2`—is further divided into 2 subpartitions. In effect, the entire table is divided into  $3 * 2 = 6$  partitions. However, due to the action of the `PARTITION BY RANGE` clause, the first 2 of these store only those records with a value less than 1990 in the `purchased` column.

In MySQL 5.6, it is possible to subpartition tables that are partitioned by `RANGE` or `LIST`. Subpartitions may use either `HASH` or `KEY` partitioning. This is also known as *composite partitioning*.



### Note

`SUBPARTITION BY HASH` and `SUBPARTITION BY KEY` generally follow the same syntax rules as `PARTITION BY HASH` and `PARTITION BY KEY`, respectively. An exception to this is that `SUBPARTITION BY KEY` (unlike `PARTITION BY KEY`) does not currently support a default column, so the

column used for this purpose must be specified, even if the table has an explicit primary key. This is a known issue which we are working to address; see [Issues with subpartitions](#), for more information and an example.

It is also possible to define subpartitions explicitly using `SUBPARTITION` clauses to specify options for individual subpartitions. For example, a more verbose fashion of creating the same table `ts` as shown in the previous example would be:

```
CREATE TABLE ts (id INT, purchased DATE)
  PARTITION BY RANGE( YEAR(purchased) )
  SUBPARTITION BY HASH( TO_DAYS(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990) (
      SUBPARTITION s0,
      SUBPARTITION s1
    ),
    PARTITION p1 VALUES LESS THAN (2000) (
      SUBPARTITION s2,
      SUBPARTITION s3
    ),
    PARTITION p2 VALUES LESS THAN MAXVALUE (
      SUBPARTITION s4,
      SUBPARTITION s5
    )
  );
```

Some syntactical items of note are listed here:

- Each partition must have the same number of subpartitions.
- If you explicitly define any subpartitions using `SUBPARTITION` on any partition of a partitioned table, you must define them all. In other words, the following statement will fail:

```
CREATE TABLE ts (id INT, purchased DATE)
  PARTITION BY RANGE( YEAR(purchased) )
  SUBPARTITION BY HASH( TO_DAYS(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990) (
      SUBPARTITION s0,
      SUBPARTITION s1
    ),
    PARTITION p1 VALUES LESS THAN (2000),
    PARTITION p2 VALUES LESS THAN MAXVALUE (
      SUBPARTITION s2,
      SUBPARTITION s3
    )
  );
```

This statement would still fail even if it included a `SUBPARTITIONS 2` clause.

- Each `SUBPARTITION` clause must include (at a minimum) a name for the subpartition. Otherwise, you may set any desired option for the subpartition or allow it to assume its default setting for that option.
- Subpartition names must be unique across the entire table. For example, the following `CREATE TABLE` statement is valid in MySQL 5.6:

```
CREATE TABLE ts (id INT, purchased DATE)
  PARTITION BY RANGE( YEAR(purchased) )
  SUBPARTITION BY HASH( TO_DAYS(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990) (
      SUBPARTITION s0,
      SUBPARTITION s1
    ),
    PARTITION p1 VALUES LESS THAN (2000) (
      SUBPARTITION s2,
      SUBPARTITION s3
    ),
    PARTITION p2 VALUES LESS THAN MAXVALUE (
      SUBPARTITION s4,
```

```

        SUBPARTITION s5
    )
);

```

Subpartitions can be used with especially large tables to distribute data and indexes across many disks. Suppose that you have 6 disks mounted as `/disk0`, `/disk1`, `/disk2`, and so on. Now consider the following example:

```

CREATE TABLE ts (id INT, purchased DATE)
PARTITION BY RANGE( YEAR(purchased) )
SUBPARTITION BY HASH( TO_DAYS(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990) (
        SUBPARTITION s0
            DATA DIRECTORY = '/disk0/data'
            INDEX DIRECTORY = '/disk0/idx',
        SUBPARTITION s1
            DATA DIRECTORY = '/disk1/data'
            INDEX DIRECTORY = '/disk1/idx'
    ),
    PARTITION p1 VALUES LESS THAN (2000) (
        SUBPARTITION s2
            DATA DIRECTORY = '/disk2/data'
            INDEX DIRECTORY = '/disk2/idx',
        SUBPARTITION s3
            DATA DIRECTORY = '/disk3/data'
            INDEX DIRECTORY = '/disk3/idx'
    ),
    PARTITION p2 VALUES LESS THAN MAXVALUE (
        SUBPARTITION s4
            DATA DIRECTORY = '/disk4/data'
            INDEX DIRECTORY = '/disk4/idx',
        SUBPARTITION s5
            DATA DIRECTORY = '/disk5/data'
            INDEX DIRECTORY = '/disk5/idx'
    )
);

```

In this case, a separate disk is used for the data and for the indexes of each `RANGE`. Many other variations are possible; another example might be:

```

CREATE TABLE ts (id INT, purchased DATE)
PARTITION BY RANGE(YEAR(purchased))
SUBPARTITION BY HASH( TO_DAYS(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990) (
        SUBPARTITION s0a
            DATA DIRECTORY = '/disk0'
            INDEX DIRECTORY = '/disk1',
        SUBPARTITION s0b
            DATA DIRECTORY = '/disk2'
            INDEX DIRECTORY = '/disk3'
    ),
    PARTITION p1 VALUES LESS THAN (2000) (
        SUBPARTITION s1a
            DATA DIRECTORY = '/disk4/data'
            INDEX DIRECTORY = '/disk4/idx',
        SUBPARTITION s1b
            DATA DIRECTORY = '/disk5/data'
            INDEX DIRECTORY = '/disk5/idx'
    ),
    PARTITION p2 VALUES LESS THAN MAXVALUE (
        SUBPARTITION s2a,
        SUBPARTITION s2b
    )
);

```

Here, the storage is as follows:

- Rows with `purchased` dates from before 1990 take up a vast amount of space, so are split up 4 ways, with a separate disk dedicated to the data and to the indexes for each of the two subpartitions (`s0a` and `s0b`) making up partition `p0`. In other words:



- The data for subpartition `s0a` is stored on `/disk0`.
- The indexes for subpartition `s0a` are stored on `/disk1`.
- The data for subpartition `s0b` is stored on `/disk2`.
- The indexes for subpartition `s0b` are stored on `/disk3`.
- Rows containing dates ranging from 1990 to 1999 (partition `p1`) do not require as much room as those from before 1990. These are split between 2 disks (`/disk4` and `/disk5`) rather than 4 disks as with the legacy records stored in `p0`:
  - Data and indexes belonging to `p1`'s first subpartition (`s1a`) are stored on `/disk4`—the data in `/disk4/data`, and the indexes in `/disk4/idx`.
  - Data and indexes belonging to `p1`'s second subpartition (`s1b`) are stored on `/disk5`—the data in `/disk5/data`, and the indexes in `/disk5/idx`.
- Rows reflecting dates from the year 2000 to the present (partition `p2`) do not take up as much space as required by either of the two previous ranges. Currently, it is sufficient to store all of these in the default location.

In future, when the number of purchases for the decade beginning with the year 2000 grows to a point where the default location no longer provides sufficient space, the corresponding rows can be moved using an `ALTER TABLE ... REORGANIZE PARTITION` statement. See [Section 17.3, “Partition Management”](#), for an explanation of how this can be done.

The `DATA DIRECTORY` and `INDEX DIRECTORY` options are not permitted in partition definitions when the `NO_DIR_IN_CREATE` [604] server SQL mode is in effect. In MySQL 5.6, these options are also not permitted when defining subpartitions (Bug #42954).

### 17.2.7. How MySQL Partitioning Handles `NULL`

Partitioning in MySQL does nothing to disallow `NULL` as the value of a partitioning expression, whether it is a column value or the value of a user-supplied expression. Even though it is permitted to use `NULL` as the value of an expression that must otherwise yield an integer, it is important to keep in mind that `NULL` is not a number. MySQL's partitioning implementation treats `NULL` as being less than any non-`NULL` value, just as `ORDER BY` does.

This means that treatment of `NULL` varies between partitioning of different types, and may produce behavior which you do not expect if you are not prepared for it. This being the case, we discuss in this section how each MySQL partitioning type handles `NULL` values when determining the partition in which a row should be stored, and provide examples for each.

**Handling of `NULL` with `RANGE` partitioning.** If you insert a row into a table partitioned by `RANGE` such that the column value used to determine the partition is `NULL`, the row is inserted into the lowest partition. Consider these two tables in a database named `p`, created as follows:

```
mysql> CREATE TABLE t1 (
->   c1 INT,
->   c2 VARCHAR(20)
-> )
-> PARTITION BY RANGE(c1) (
->   PARTITION p0 VALUES LESS THAN (0),
->   PARTITION p1 VALUES LESS THAN (10),
->   PARTITION p2 VALUES LESS THAN MAXVALUE
-> );
```

Query OK, 0 rows affected (0.09 sec)

```
mysql> CREATE TABLE t2 (
->   c1 INT,
->   c2 VARCHAR(20)
```

```
-> )
-> PARTITION BY RANGE(c1) (
->     PARTITION p0 VALUES LESS THAN (-5),
->     PARTITION p1 VALUES LESS THAN (0),
->     PARTITION p2 VALUES LESS THAN (10),
->     PARTITION p3 VALUES LESS THAN MAXVALUE
-> );
Query OK, 0 rows affected (0.09 sec)
```

You can see the partitions created by these two `CREATE TABLE` statements using the following query against the `PARTITIONS` table in the `INFORMATION_SCHEMA` database:

```
mysql> SELECT TABLE_NAME, PARTITION_NAME, TABLE_ROWS, AVG_ROW_LENGTH, DATA_LENGTH
> FROM INFORMATION_SCHEMA.PARTITIONS
> WHERE TABLE_SCHEMA = 'p' AND TABLE_NAME LIKE 't_';
+-----+-----+-----+-----+-----+
| TABLE_NAME | PARTITION_NAME | TABLE_ROWS | AVG_ROW_LENGTH | DATA_LENGTH |
+-----+-----+-----+-----+-----+
| t1          | p0             | 0           | 0              | 0            |
| t1          | p1             | 0           | 0              | 0            |
| t1          | p2             | 0           | 0              | 0            |
| t2          | p0             | 0           | 0              | 0            |
| t2          | p1             | 0           | 0              | 0            |
| t2          | p2             | 0           | 0              | 0            |
| t2          | p3             | 0           | 0              | 0            |
+-----+-----+-----+-----+-----+
7 rows in set (0.00 sec)
```

(For more information about this table, see [Section 19.14, “The `INFORMATION\_SCHEMA PARTITIONS` Table”](#).) Now let us populate each of these tables with a single row containing a `NULL` in the column used as the partitioning key, and verify that the rows were inserted using a pair of `SELECT` statements:

```
mysql> INSERT INTO t1 VALUES (NULL, 'mothra');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO t2 VALUES (NULL, 'mothra');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM t1;
+-----+-----+
| id   | name |
+-----+-----+
| NULL | mothra |
+-----+-----+
1 row in set (0.00 sec)

mysql> SELECT * FROM t2;
+-----+-----+
| id   | name |
+-----+-----+
| NULL | mothra |
+-----+-----+
1 row in set (0.00 sec)
```

You can see which partitions are used to store the inserted rows by rerunning the previous query against `INFORMATION_SCHEMA.PARTITIONS` and inspecting the output:

```
mysql> SELECT TABLE_NAME, PARTITION_NAME, TABLE_ROWS, AVG_ROW_LENGTH, DATA_LENGTH
> FROM INFORMATION_SCHEMA.PARTITIONS
> WHERE TABLE_SCHEMA = 'p' AND TABLE_NAME LIKE 't_';
+-----+-----+-----+-----+-----+
| TABLE_NAME | PARTITION_NAME | TABLE_ROWS | AVG_ROW_LENGTH | DATA_LENGTH |
+-----+-----+-----+-----+-----+
| t1          | p0             | 1           | 20             | 20           |
| t1          | p1             | 0           | 0              | 0            |
| t1          | p2             | 0           | 0              | 0            |
| t2          | p0             | 1           | 20             | 20           |
| t2          | p1             | 0           | 0              | 0            |
| t2          | p2             | 0           | 0              | 0            |
| t2          | p3             | 0           | 0              | 0            |
+-----+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+
7 rows in set (0.01 sec)
```

You can also demonstrate that these rows were stored in the lowest partition of each table by dropping these partitions, and then re-running the `SELECT` statements:

```
mysql> ALTER TABLE t1 DROP PARTITION p0;
Query OK, 0 rows affected (0.16 sec)

mysql> ALTER TABLE t2 DROP PARTITION p0;
Query OK, 0 rows affected (0.16 sec)

mysql> SELECT * FROM t1;
Empty set (0.00 sec)

mysql> SELECT * FROM t2;
Empty set (0.00 sec)
```

(For more information on `ALTER TABLE ... DROP PARTITION`, see [Section 13.1.6, “ALTER TABLE Syntax”](#).)

`NULL` is also treated in this way for partitioning expressions that use SQL functions. Suppose that we define a table using a `CREATE TABLE` statement such as this one:

```
CREATE TABLE tndate (
  id INT,
  dt DATE
)
PARTITION BY RANGE( YEAR(dt) ) (
  PARTITION p0 VALUES LESS THAN (1990),
  PARTITION p1 VALUES LESS THAN (2000),
  PARTITION p2 VALUES LESS THAN MAXVALUE
);
```

As with other MySQL functions, `YEAR(NULL)` [1229] returns `NULL`. A row with a `dt` column value of `NULL` is treated as though the partitioning expression evaluated to a value less than any other value, and so is inserted into partition `p0`.

**Handling of `NULL` with `LIST` partitioning.** A table that is partitioned by `LIST` admits `NULL` values if and only if one of its partitions is defined using that value-list that contains `NULL`. The converse of this is that a table partitioned by `LIST` which does not explicitly use `NULL` in a value list rejects rows resulting in a `NULL` value for the partitioning expression, as shown in this example:

```
mysql> CREATE TABLE ts1 (
->   c1 INT,
->   c2 VARCHAR(20)
-> )
-> PARTITION BY LIST(c1) (
->   PARTITION p0 VALUES IN (0, 3, 6),
->   PARTITION p1 VALUES IN (1, 4, 7),
->   PARTITION p2 VALUES IN (2, 5, 8)
-> );
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO ts1 VALUES (9, 'mothra');
ERROR 1504 (HY000): Table has no partition for value 9

mysql> INSERT INTO ts1 VALUES (NULL, 'mothra');
ERROR 1504 (HY000): Table has no partition for value NULL
```

Only rows having a `c1` value between 0 and 8 inclusive can be inserted into `ts1`. `NULL` falls outside this range, just like the number 9. We can create tables `ts2` and `ts3` having value lists containing `NULL`, as shown here:

```
mysql> CREATE TABLE ts2 (
->   c1 INT,
->   c2 VARCHAR(20)
```

```

-> )
-> PARTITION BY LIST(c1) (
->     PARTITION p0 VALUES IN (0, 3, 6),
->     PARTITION p1 VALUES IN (1, 4, 7),
->     PARTITION p2 VALUES IN (2, 5, 8),
->     PARTITION p3 VALUES IN (NULL)
-> );
Query OK, 0 rows affected (0.01 sec)

mysql> CREATE TABLE ts3 (
->     c1 INT,
->     c2 VARCHAR(20)
-> )
-> PARTITION BY LIST(c1) (
->     PARTITION p0 VALUES IN (0, 3, 6),
->     PARTITION p1 VALUES IN (1, 4, 7, NULL),
->     PARTITION p2 VALUES IN (2, 5, 8)
-> );
Query OK, 0 rows affected (0.01 sec)
    
```

When defining value lists for partitioning, you can (and should) treat `NULL` just as you would any other value. For example, both `VALUES IN (NULL)` and `VALUES IN (1, 4, 7, NULL)` are valid, as are `VALUES IN (1, NULL, 4, 7)`, `VALUES IN (NULL, 1, 4, 7)`, and so on. You can insert a row having `NULL` for column `c1` into each of the tables `ts2` and `ts3`:

```

mysql> INSERT INTO ts2 VALUES (NULL, 'mothra');
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO ts3 VALUES (NULL, 'mothra');
Query OK, 1 row affected (0.00 sec)
    
```

By issuing the appropriate query against `INFORMATION_SCHEMA.PARTITIONS`, you can determine which partitions were used to store the rows just inserted (we assume, as in the previous examples, that the partitioned tables were created in the `p` database):

```

mysql> SELECT TABLE_NAME, PARTITION_NAME, TABLE_ROWS, AVG_ROW_LENGTH, DATA_LENGTH
> FROM INFORMATION_SCHEMA.PARTITIONS
> WHERE TABLE_SCHEMA = 'p' AND TABLE_NAME LIKE 'ts_';
+-----+-----+-----+-----+-----+
| TABLE_NAME | PARTITION_NAME | TABLE_ROWS | AVG_ROW_LENGTH | DATA_LENGTH |
+-----+-----+-----+-----+-----+
| ts2         | p0              | 0           | 0              | 0            |
| ts2         | p1              | 0           | 0              | 0            |
| ts2         | p2              | 0           | 0              | 0            |
| / ts2       | / p3            | / 1         | / 20           | / 20         |
| ts3         | p0              | 0           | 0              | 0            |
| / ts3       | / p1            | / 1         | / 20           | / 20         |
| ts3         | p2              | 0           | 0              | 0            |
+-----+-----+-----+-----+-----+
7 rows in set (0.01 sec)
    
```

As shown earlier in this section, you can also verify which partitions were used for storing the rows by deleting these partitions and then performing a `SELECT`.

**Handling of `NULL` with `HASH` and `KEY` partitioning.** `NULL` is handled somewhat differently for tables partitioned by `HASH` or `KEY`. In these cases, any partition expression that yields a `NULL` value is treated as though its return value were zero. We can verify this behavior by examining the effects on the file system of creating a table partitioned by `HASH` and populating it with a record containing appropriate values. Suppose that you have a table `th` (also in the `p` database) created using the following statement:

```

mysql> CREATE TABLE th (
->     c1 INT,
->     c2 VARCHAR(20)
-> )
-> PARTITION BY HASH(c1)
-> PARTITIONS 2;
Query OK, 0 rows affected (0.00 sec)
    
```

The partitions belonging to this table can be viewed using the query shown here:

```
mysql> SELECT TABLE_NAME, PARTITION_NAME, TABLE_ROWS, AVG_ROW_LENGTH, DATA_LENGTH
> FROM INFORMATION_SCHEMA.PARTITIONS
> WHERE TABLE_SCHEMA = 'p' AND TABLE_NAME = 'th';
+-----+-----+-----+-----+-----+
| TABLE_NAME | PARTITION_NAME | TABLE_ROWS | AVG_ROW_LENGTH | DATA_LENGTH |
+-----+-----+-----+-----+-----+
| th          | p0             | 0           | 0              | 0            |
| th          | p1             | 0           | 0              | 0            |
+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

Note that `TABLE_ROWS` for each partition is 0. Now insert two rows into `th` whose `c1` column values are `NULL` and 0, and verify that these rows were inserted, as shown here:

```
mysql> INSERT INTO th VALUES (NULL, 'mothra'), (0, 'gigan');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM th;
+-----+-----+
| c1   | c2     |
+-----+-----+
| NULL | mothra |
+-----+-----+
| 0    | gigan  |
+-----+-----+
2 rows in set (0.01 sec)
```

Recall that for any integer  $N$ , the value of `NULL MOD N` is always `NULL`. For tables that are partitioned by `HASH` or `KEY`, this result is treated for determining the correct partition as 0. Checking the `INFORMATION_SCHEMA.PARTITIONS` table once again, we can see that both rows were inserted into partition `p0`:

```
mysql> SELECT TABLE_NAME, PARTITION_NAME, TABLE_ROWS, AVG_ROW_LENGTH, DATA_LENGTH
> FROM INFORMATION_SCHEMA.PARTITIONS
> WHERE TABLE_SCHEMA = 'p' AND TABLE_NAME = 'th';
+-----+-----+-----+-----+-----+
| TABLE_NAME | PARTITION_NAME | TABLE_ROWS | AVG_ROW_LENGTH | DATA_LENGTH |
+-----+-----+-----+-----+-----+
| th          | p0             | 2           | 20             | 20           |
| th          | p1             | 0           | 0              | 0            |
+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

If you repeat this example using `PARTITION BY KEY` in place of `PARTITION BY HASH` in the definition of the table, you can verify easily that `NULL` is also treated like 0 for this type of partitioning.

## 17.3. Partition Management

MySQL 5.6 provides a number of ways to modify partitioned tables. It is possible to add, drop, redefine, merge, or split existing partitions. All of these actions can be carried out using the partitioning extensions to the `ALTER TABLE` statement. There are also ways to obtain information about partitioned tables and partitions. We discuss these topics in the sections that follow.

- For information about partition management in tables partitioned by `RANGE` or `LIST`, see [Section 17.3.1, “Management of RANGE and LIST Partitions”](#).
- For a discussion of managing `HASH` and `KEY` partitions, see [Section 17.3.2, “Management of HASH and KEY Partitions”](#).
- See [Section 17.3.5, “Obtaining Information About Partitions”](#), for a discussion of mechanisms provided in MySQL 5.6 for obtaining information about partitioned tables and partitions.
- For a discussion of performing maintenance operations on partitions, see [Section 17.3.4, “Maintenance of Partitions”](#).

**Note**

In MySQL 5.6, all partitions of a partitioned table must have the same number of subpartitions, and it is not possible to change the subpartitioning once the table has been created.

To change a table's partitioning scheme, it is necessary only to use the `ALTER TABLE` statement with a `partition_options` clause. This clause has the same syntax as that as used with `CREATE TABLE` for creating a partitioned table, and always begins with the keywords `PARTITION BY`. Suppose that you have a table partitioned by range using the following `CREATE TABLE` statement:

```
CREATE TABLE trb3 (id INT, name VARCHAR(50), purchased DATE)
  PARTITION BY RANGE( YEAR(purchased) ) (
    PARTITION p0 VALUES LESS THAN (1990),
    PARTITION p1 VALUES LESS THAN (1995),
    PARTITION p2 VALUES LESS THAN (2000),
    PARTITION p3 VALUES LESS THAN (2005)
  );
```

To repartition this table so that it is partitioned by key into two partitions using the `id` column value as the basis for the key, you can use this statement:

```
ALTER TABLE trb3 PARTITION BY KEY(id) PARTITIONS 2;
```

This has the same effect on the structure of the table as dropping the table and re-creating it using `CREATE TABLE trb3 PARTITION BY KEY(id) PARTITIONS 2;`.

`ALTER TABLE ... ENGINE = ...` changes only the storage engine used by the table, and leaves the table's partitioning scheme intact. Use `ALTER TABLE ... REMOVE PARTITIONING` to remove a table's partitioning. See [Section 13.1.6, "ALTER TABLE Syntax"](#).

**Important**

Only a single `PARTITION BY`, `ADD PARTITION`, `DROP PARTITION`, `REORGANIZE PARTITION`, or `COALESCE PARTITION` clause can be used in a given `ALTER TABLE` statement. If you (for example) wish to drop a partition and reorganize a table's remaining partitions, you must do so in two separate `ALTER TABLE` statements (one using `DROP PARTITION` and then a second one using `REORGANIZE PARTITIONS`).

In MySQL 5.6, it is possible to delete all rows from one or more selected partitions using `ALTER TABLE ... TRUNCATE PARTITION`.

### 17.3.1. Management of RANGE and LIST Partitions

Range and list partitions are very similar with regard to how the adding and dropping of partitions are handled. For this reason we discuss the management of both sorts of partitioning in this section. For information about working with tables that are partitioned by hash or key, see [Section 17.3.2, "Management of HASH and KEY Partitions"](#). Dropping a `RANGE` or `LIST` partition is more straightforward than adding one, so we discuss this first.

Dropping a partition from a table that is partitioned by either `RANGE` or by `LIST` can be accomplished using the `ALTER TABLE` statement with a `DROP PARTITION` clause. Here is a very basic example, which supposes that you have already created a table which is partitioned by range and then populated with 10 records using the following `CREATE TABLE` and `INSERT` statements:

```
mysql> CREATE TABLE tr (id INT, name VARCHAR(50), purchased DATE)
->   PARTITION BY RANGE( YEAR(purchased) ) (
->     PARTITION p0 VALUES LESS THAN (1990),
->     PARTITION p1 VALUES LESS THAN (1995),
->     PARTITION p2 VALUES LESS THAN (2000),
->     PARTITION p3 VALUES LESS THAN (2005)
->   );
Query OK, 0 rows affected (0.01 sec)
```

```
mysql> INSERT INTO tr VALUES
-> (1, 'desk organiser', '2003-10-15'),
-> (2, 'CD player', '1993-11-05'),
-> (3, 'TV set', '1996-03-10'),
-> (4, 'bookcase', '1982-01-10'),
-> (5, 'exercise bike', '2004-05-09'),
-> (6, 'sofa', '1987-06-05'),
-> (7, 'popcorn maker', '2001-11-22'),
-> (8, 'aquarium', '1992-08-04'),
-> (9, 'study desk', '1984-09-16'),
-> (10, 'lava lamp', '1998-12-25');
Query OK, 10 rows affected (0.01 sec)
```

You can see which items should have been inserted into partition `p2` as shown here:

```
mysql> SELECT * FROM tr
-> WHERE purchased BETWEEN '1995-01-01' AND '1999-12-31';
+-----+-----+-----+
| id  | name      | purchased |
+-----+-----+-----+
| 3   | TV set    | 1996-03-10 |
| 10  | lava lamp | 1998-12-25 |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

To drop the partition named `p2`, execute the following command:

```
mysql> ALTER TABLE tr DROP PARTITION p2;
Query OK, 0 rows affected (0.03 sec)
```

It is very important to remember that, *when you drop a partition, you also delete all the data that was stored in that partition*. You can see that this is the case by re-running the previous `SELECT` query:

```
mysql> SELECT * FROM tr WHERE purchased
-> BETWEEN '1995-01-01' AND '1999-12-31';
Empty set (0.00 sec)
```

Because of this, you must have the `DROP [754]` privilege for a table before you can execute `ALTER TABLE ... DROP PARTITION` on that table.

If you wish to drop all data from all partitions while preserving the table definition and its partitioning scheme, use the `TRUNCATE TABLE` statement. (See [Section 13.1.27](#), “`TRUNCATE TABLE Syntax`”.)

If you intend to change the partitioning of a table *without* losing data, use `ALTER TABLE ... REORGANIZE PARTITION` instead. See below or in [Section 13.1.6](#), “`ALTER TABLE Syntax`”, for information about `REORGANIZE PARTITION`.

If you now execute a `SHOW CREATE TABLE` statement, you can see how the partitioning makeup of the table has been changed:

```
mysql> SHOW CREATE TABLE tr\G
***** 1. row *****
Table: tr
Create Table: CREATE TABLE `tr` (
  `id` int(11) default NULL,
  `name` varchar(50) default NULL,
  `purchased` date default NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
PARTITION BY RANGE ( YEAR(purchased) ) (
  PARTITION p0 VALUES LESS THAN (1990) ENGINE = MyISAM,
  PARTITION p1 VALUES LESS THAN (1995) ENGINE = MyISAM,
  PARTITION p3 VALUES LESS THAN (2005) ENGINE = MyISAM
)
1 row in set (0.01 sec)
```

When you insert new rows into the changed table with `purchased` column values between `'1995-01-01'` and `'2004-12-31'` inclusive, those rows will be stored in partition `p3`. You can verify this as follows:



```
mysql> INSERT INTO tr VALUES (11, 'pencil holder', '1995-07-12');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM tr WHERE purchased
-> BETWEEN '1995-01-01' AND '2004-12-31';
+-----+-----+-----+
| id   | name           | purchased |
+-----+-----+-----+
| 11  | pencil holder  | 1995-07-12 |
| 1   | desk organiser | 2003-10-15 |
| 5   | exercise bike  | 2004-05-09 |
| 7   | popcorn maker  | 2001-11-22 |
+-----+-----+-----+
4 rows in set (0.00 sec)

mysql> ALTER TABLE tr DROP PARTITION p3;
Query OK, 0 rows affected (0.03 sec)

mysql> SELECT * FROM tr WHERE purchased
-> BETWEEN '1995-01-01' AND '2004-12-31';
Empty set (0.00 sec)
```

Note that the number of rows dropped from the table as a result of `ALTER TABLE ... DROP PARTITION` is not reported by the server as it would be by the equivalent `DELETE` query.

Dropping `LIST` partitions uses exactly the same `ALTER TABLE ... DROP PARTITION` syntax as used for dropping `RANGE` partitions. However, there is one important difference in the effect this has on your use of the table afterward: You can no longer insert into the table any rows having any of the values that were included in the value list defining the deleted partition. (See [Section 17.2.2](#), “`LIST Partitioning`”, for an example.)

To add a new range or list partition to a previously partitioned table, use the `ALTER TABLE ... ADD PARTITION` statement. For tables which are partitioned by `RANGE`, this can be used to add a new range to the end of the list of existing partitions. Suppose that you have a partitioned table containing membership data for your organization, which is defined as follows:

```
CREATE TABLE members (
  id INT,
  fname VARCHAR(25),
  lname VARCHAR(25),
  dob DATE
)
PARTITION BY RANGE( YEAR(dob) ) (
  PARTITION p0 VALUES LESS THAN (1970),
  PARTITION p1 VALUES LESS THAN (1980),
  PARTITION p2 VALUES LESS THAN (1990)
);
```

Suppose further that the minimum age for members is 16. As the calendar approaches the end of 2005, you realize that you will soon be admitting members who were born in 1990 (and later in years to come). You can modify the `members` table to accommodate new members born in the years 1990 to 1999 as shown here:

```
ALTER TABLE members ADD PARTITION (PARTITION p3 VALUES LESS THAN (2000));
```

With tables that are partitioned by range, you can use `ADD PARTITION` to add new partitions to the high end of the partitions list only. Trying to add a new partition in this manner between or before existing partitions results in an error as shown here:

```
mysql> ALTER TABLE members
> ADD PARTITION (
> PARTITION n VALUES LESS THAN (1960));
ERROR 1463 (HY000): VALUES LESS THAN value must be strictly >
increasing for each partition
```

You can work around this problem by reorganizing the first partition into two new ones that split the range between them, like this:

```
ALTER TABLE members
  REORGANIZE PARTITION p0 INTO (
    PARTITION n0 VALUES LESS THAN (1960),
    PARTITION n1 VALUES LESS THAN (1970)
  );
```

Using `SHOW CREATE TABLE` you can see that the `ALTER TABLE` statement has had the desired effect:

```
mysql> SHOW CREATE TABLE members\G
***** 1. row *****
      Table: members
Create Table: CREATE TABLE `members` (
  `id` int(11) DEFAULT NULL,
  `fname` varchar(25) DEFAULT NULL,
  `lname` varchar(25) DEFAULT NULL,
  `dob` date DEFAULT NULL
) ENGINE=InnoDB DEFAULT CHARSET=latin1
/*!50100 PARTITION BY RANGE ( YEAR(dob))
(PARTITION n0 VALUES LESS THAN (1960) ENGINE = InnoDB,
PARTITION n1 VALUES LESS THAN (1970) ENGINE = InnoDB,
PARTITION p1 VALUES LESS THAN (1980) ENGINE = InnoDB,
PARTITION p2 VALUES LESS THAN (1990) ENGINE = InnoDB,
PARTITION p3 VALUES LESS THAN (2000) ENGINE = InnoDB) */
1 row in set (0.00 sec)
```

See also [Section 13.1.6.1, “ALTER TABLE Partition Operations”](#).

You can also use `ALTER TABLE ... ADD PARTITION` to add new partitions to a table that is partitioned by `LIST`. Suppose a table `tt` is defined using the following `CREATE TABLE` statement:

```
CREATE TABLE tt (
  id INT,
  data INT
)
PARTITION BY LIST(data) (
  PARTITION p0 VALUES IN (5, 10, 15),
  PARTITION p1 VALUES IN (6, 12, 18)
);
```

You can add a new partition in which to store rows having the `data` column values `7`, `14`, and `21` as shown:

```
ALTER TABLE tt ADD PARTITION (PARTITION p2 VALUES IN (7, 14, 21));
```

Note that you *cannot* add a new `LIST` partition encompassing any values that are already included in the value list of an existing partition. If you attempt to do so, an error will result:

```
mysql> ALTER TABLE tt ADD PARTITION
> (PARTITION np VALUES IN (4, 8, 12));
ERROR 1465 (HY000): Multiple definition of same constant »
in list partitioning
```

Because any rows with the `data` column value `12` have already been assigned to partition `p1`, you cannot create a new partition on table `tt` that includes `12` in its value list. To accomplish this, you could drop `p1`, and add `np` and then a new `p1` with a modified definition. However, as discussed earlier, this would result in the loss of all data stored in `p1`—and it is often the case that this is not what you really want to do. Another solution might appear to be to make a copy of the table with the new partitioning and to copy the data into it using `CREATE TABLE ... SELECT ...`, then drop the old table and rename the new one, but this could be very time-consuming when dealing with a large amounts of data. This also might not be feasible in situations where high availability is a requirement.

You can add multiple partitions in a single `ALTER TABLE ... ADD PARTITION` statement as shown here:

```
CREATE TABLE employees (
  id INT NOT NULL,
```

```

    fname VARCHAR(50) NOT NULL,
    lname VARCHAR(50) NOT NULL,
    hired DATE NOT NULL
)
PARTITION BY RANGE( YEAR(hired) ) (
    PARTITION p1 VALUES LESS THAN (1991),
    PARTITION p2 VALUES LESS THAN (1996),
    PARTITION p3 VALUES LESS THAN (2001),
    PARTITION p4 VALUES LESS THAN (2005)
);

ALTER TABLE employees ADD PARTITION (
    PARTITION p5 VALUES LESS THAN (2010),
    PARTITION p6 VALUES LESS THAN MAXVALUE
);

```

Fortunately, MySQL's partitioning implementation provides ways to redefine partitions without losing data. Let us look first at a couple of simple examples involving [RANGE](#) partitioning. Recall the [members](#) table which is now defined as shown here:

```

mysql> SHOW CREATE TABLE members\G
***** 1. row *****
      Table: members
Create Table: CREATE TABLE `members` (
  `id` int(11) default NULL,
  `fname` varchar(25) default NULL,
  `lname` varchar(25) default NULL,
  `dob` date default NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
PARTITION BY RANGE ( YEAR(dob) ) (
  PARTITION p0 VALUES LESS THAN (1970) ENGINE = MyISAM,
  PARTITION p1 VALUES LESS THAN (1980) ENGINE = MyISAM,
  PARTITION p2 VALUES LESS THAN (1990) ENGINE = MyISAM,
  PARTITION p3 VALUES LESS THAN (2000) ENGINE = MyISAM
)

```

Suppose that you would like to move all rows representing members born before 1960 into a separate partition. As we have already seen, this cannot be done using [ALTER TABLE ... ADD PARTITION](#). However, you can use another partition-related extension to [ALTER TABLE](#) to accomplish this:

```

ALTER TABLE members REORGANIZE PARTITION p0 INTO (
    PARTITION s0 VALUES LESS THAN (1960),
    PARTITION s1 VALUES LESS THAN (1970)
);

```

In effect, this command splits partition [p0](#) into two new partitions [s0](#) and [s1](#). It also moves the data that was stored in [p0](#) into the new partitions according to the rules embodied in the two [PARTITION ... VALUES ...](#) clauses, so that [s0](#) contains only those records for which [YEAR\(dob\)](#) [[1229](#)] is less than 1960 and [s1](#) contains those rows in which [YEAR\(dob\)](#) [[1229](#)] is greater than or equal to 1960 but less than 1970.

A [REORGANIZE PARTITION](#) clause may also be used for merging adjacent partitions. You can return the [members](#) table to its previous partitioning as shown here:

```

ALTER TABLE members REORGANIZE PARTITION s0,s1 INTO (
    PARTITION p0 VALUES LESS THAN (1970)
);

```

No data is lost in splitting or merging partitions using [REORGANIZE PARTITION](#). In executing the above statement, MySQL moves all of the records that were stored in partitions [s0](#) and [s1](#) into partition [p0](#).

The general syntax for [REORGANIZE PARTITION](#) is shown here:

```

ALTER TABLE tbl_name
    REORGANIZE PARTITION partition_list
    INTO (partition_definitions);

```

Here, *tbl\_name* is the name of the partitioned table, and *partition\_list* is a comma-separated list of names of one or more existing partitions to be changed. *partition\_definitions* is a comma-separated list of new partition definitions, which follow the same rules as for the *partition\_definitions* list used in **CREATE TABLE** (see [Section 13.1.14](#), “**CREATE TABLE Syntax**”). It should be noted that you are not limited to merging several partitions into one, or to splitting one partition into many, when using **REORGANIZE PARTITION**. For example, you can reorganize all four partitions of the *members* table into two, as follows:

```
ALTER TABLE members REORGANIZE PARTITION p0,p1,p2,p3 INTO (
  PARTITION m0 VALUES LESS THAN (1980),
  PARTITION m1 VALUES LESS THAN (2000)
);
```

You can also use **REORGANIZE PARTITION** with tables that are partitioned by **LIST**. Let us return to the problem of adding a new partition to the list-partitioned *tt* table and failing because the new partition had a value that was already present in the value-list of one of the existing partitions. We can handle this by adding a partition that contains only nonconflicting values, and then reorganizing the new partition and the existing one so that the value which was stored in the existing one is now moved to the new one:

```
ALTER TABLE tt ADD PARTITION (PARTITION np VALUES IN (4, 8));
ALTER TABLE tt REORGANIZE PARTITION p1,np INTO (
  PARTITION p1 VALUES IN (6, 18),
  PARTITION np VALUES in (4, 8, 12)
);
```

Here are some key points to keep in mind when using **ALTER TABLE ... REORGANIZE PARTITION** to repartition tables that are partitioned by **RANGE** or **LIST**:

- The **PARTITION** clauses used to determine the new partitioning scheme are subject to the same rules as those used with a **CREATE TABLE** statement.

Most importantly, you should remember that the new partitioning scheme cannot have any overlapping ranges (applies to tables partitioned by **RANGE**) or sets of values (when reorganizing tables partitioned by **LIST**).

- The combination of partitions in the *partition\_definitions* list should account for the same range or set of values overall as the combined partitions named in the *partition\_list*.

For instance, in the *members* table used as an example in this section, partitions *p1* and *p2* together cover the years 1980 through 1999. Therefore, any reorganization of these two partitions should cover the same range of years overall.

- For tables partitioned by **RANGE**, you can reorganize only adjacent partitions; you cannot skip over range partitions.

For instance, you could not reorganize the *members* table used as an example in this section using a statement beginning with **ALTER TABLE members REORGANIZE PARTITION p0,p2 INTO ...** because *p0* covers the years prior to 1970 and *p2* the years from 1990 through 1999 inclusive, and thus the two are not adjacent partitions.

- You cannot use **REORGANIZE PARTITION** to change the table's partitioning type; that is, you cannot (for example) change **RANGE** partitions to **HASH** partitions or *vice versa*. You also cannot use this command to change the partitioning expression or column. To accomplish either of these tasks without dropping and re-creating the table, you can use **ALTER TABLE ... PARTITION BY ...**. For example:

```
ALTER TABLE members
  PARTITION BY HASH( YEAR(dob) )
  PARTITIONS 8;
```

## 17.3.2. Management of **HASH** and **KEY** Partitions

Tables which are partitioned by hash or by key are very similar to one another with regard to making changes in a partitioning setup, and both differ in a number of ways from tables which have been partitioned by range or list. For that reason, this section addresses the modification of tables partitioned by hash or by key only. For a discussion of adding and dropping of partitions of tables that are partitioned by range or list, see [Section 17.3.1, “Management of RANGE and LIST Partitions”](#).

You cannot drop partitions from tables that are partitioned by `HASH` or `KEY` in the same way that you can from tables that are partitioned by `RANGE` or `LIST`. However, you can merge `HASH` or `KEY` partitions using the `ALTER TABLE ... COALESCE PARTITION` statement. Suppose that you have a table containing data about clients, which is divided into twelve partitions. The `clients` table is defined as shown here:

```
CREATE TABLE clients (
  id INT,
  fname VARCHAR(30),
  lname VARCHAR(30),
  signed DATE
)
PARTITION BY HASH( MONTH(signed) )
PARTITIONS 12;
```

To reduce the number of partitions from twelve to eight, execute the following `ALTER TABLE` command:

```
mysql> ALTER TABLE clients COALESCE PARTITION 4;
Query OK, 0 rows affected (0.02 sec)
```

`COALESCE` works equally well with tables that are partitioned by `HASH`, `KEY`, `LINEAR HASH`, or `LINEAR KEY`. Here is an example similar to the previous one, differing only in that the table is partitioned by `LINEAR KEY`:

```
mysql> CREATE TABLE clients_lk (
->   id INT,
->   fname VARCHAR(30),
->   lname VARCHAR(30),
->   signed DATE
-> )
-> PARTITION BY LINEAR KEY(signed)
-> PARTITIONS 12;
Query OK, 0 rows affected (0.03 sec)

mysql> ALTER TABLE clients_lk COALESCE PARTITION 4;
Query OK, 0 rows affected (0.06 sec)
Records: 0 Duplicates: 0 Warnings: 0
```

Note that the number following `COALESCE PARTITION` is the number of partitions to merge into the remainder—in other words, it is the number of partitions to remove from the table.

If you attempt to remove more partitions than the table has, the result is an error like the one shown:

```
mysql> ALTER TABLE clients COALESCE PARTITION 18;
ERROR 1478 (HY000): Cannot remove all partitions, use DROP TABLE instead
```

To increase the number of partitions for the `clients` table from 12 to 18, use `ALTER TABLE ... ADD PARTITION` as shown here:

```
ALTER TABLE clients ADD PARTITION PARTITIONS 6;
```

### 17.3.3. Exchanging Partitions and Subpartitions with Tables

In MySQL 5.6, it is possible to exchange a table partition or subpartition with a table using the `ALTER TABLE ... EXCHANGE PARTITION` statement—that is, to move any existing rows in the partition or subpartition to the nonpartitioned table, and any existing rows in the nonpartitioned table to the table partition or subpartition.

Such operations are subject to the following conditions:

- The table to be exchanged must not be partitioned, but must otherwise have the same table structure as the partitioned table.
- The table to be exchanged must not be a temporary table.
- Any rows existing in the nonpartitioned table prior to the exchange must lie within the range defined for the partition or subpartition.
- The table to be exchanged may not have any foreign keys, nor may any other tables have foreign keys which reference this table.
- In addition to the [ALTER \[753\]](#), [INSERT \[755\]](#), and [CREATE \[754\]](#) privileges usually required for [ALTER TABLE](#) statements, you must have the [DROP \[754\]](#) privilege to perform [ALTER TABLE ... EXCHANGE PARTITION](#).

In addition, you should also be aware of the following effects of [ALTER TABLE ... EXCHANGE PARTITION](#):

- Executing this statement does *not* invoke any triggers on either the partitioned table or the exchanged table.
- Any [AUTO\\_INCREMENT](#) columns in the exchanged table are reset.

The complete syntax of the [ALTER TABLE ... EXCHANGE PARTITION](#) statement is shown here, where *pt* is the partitioned table, *p* is the partition or subpartition to be exchanged, and *t* is the nonpartitioned table to be exchanged with *p*:

```
ALTER TABLE pt
  EXCHANGE PARTITION p
  WITH TABLE t;
```

One and only one partition or subpartition may be exchanged with one and only one nonpartitioned table in a single [ALTER TABLE EXCHANGE PARTITION](#) statement. To exchange multiple partitions or subpartitions, use multiple [ALTER TABLE EXCHANGE PARTITION](#) statements. [EXCHANGE PARTITION](#) may not be combined with other [ALTER TABLE](#) options. The partitioning and (if applicable) subpartitioning used by the partitioned table may be of any type or types supported in MySQL 5.6.

Suppose that a partitioned table *e* has been created and populated using the following SQL statements:

```
CREATE TABLE e (
  id INT NOT NULL,
  fname VARCHAR(30),
  lname VARCHAR(30)
)
PARTITION BY RANGE (id) (
  PARTITION p0 VALUES LESS THAN (50),
  PARTITION p1 VALUES LESS THAN (100),
  PARTITION p2 VALUES LESS THAN (150),
  PARTITION p3 VALUES LESS THAN (MAXVALUE)
);

INSERT INTO e VALUES
  (1669, "Jim", "Smith"),
  (337, "Mary", "Jones"),
  (16, "Frank", "White"),
  (2005, "Linda", "Black");
```

Now we create a nonpartitioned copy of *e* named *e2*. This can be done using the [mysql](#) client as shown here:

```
mysql> CREATE TABLE e2 LIKE e;
```

```
Query OK, 0 rows affected (1.34 sec)
```

```
mysql> ALTER TABLE e2 REMOVE PARTITIONING;
Query OK, 0 rows affected (0.90 sec)
Records: 0 Duplicates: 0 Warnings: 0
```

You can see which partitions in table `e` contain rows by querying the `INFORMATION_SCHEMA.PARTITIONS` table, like this:

```
mysql> SELECT PARTITION_NAME, TABLE_ROWS
-> FROM INFORMATION_SCHEMA.PARTITIONS
-> WHERE TABLE_NAME = 'e';
+-----+-----+
| PARTITION_NAME | TABLE_ROWS |
+-----+-----+
| p0              | 1           |
| p1              | 0           |
| p2              | 0           |
| p3              | 3           |
+-----+-----+
4 rows in set (0.00 sec)
```

To exchange partition `p0` in table `e` with table `e2`, you can use the `ALTER TABLE` statement shown here:

```
mysql> ALTER TABLE e EXCHANGE PARTITION p0 WITH TABLE e2;
Query OK, 0 rows affected (0.28 sec)
```

More precisely, the statement just issued causes any rows found in the partition to be swapped with those found in the table. You can observe how this has happened by querying the `INFORMATION_SCHEMA.PARTITIONS` table, as before. The table row that was previously found in partition `p0` is no longer present:

```
mysql> SELECT PARTITION_NAME, TABLE_ROWS
-> FROM INFORMATION_SCHEMA.PARTITIONS
-> WHERE TABLE_NAME = 'e';
+-----+-----+
| PARTITION_NAME | TABLE_ROWS |
+-----+-----+
| p0              | 0           |
| p1              | 0           |
| p2              | 0           |
| p3              | 3           |
+-----+-----+
4 rows in set (0.00 sec)
```

If you query table `e2`, you can see that the “missing” row can now be found there:

```
mysql> SELECT * FROM e2;
+-----+-----+
| id | fname | lname |
+-----+-----+
| 16 | Frank | White |
+-----+-----+
1 row in set (0.00 sec)
```

The table to be exchanged with the partition does not necessarily have to be empty. To demonstrate this, we first insert a new row into table `e`, making sure that this row is stored in partition `p0` by choosing an `id` column value that is less than 50, and verifying this afterwards by querying the `PARTITIONS` table:

```
mysql> INSERT INTO e VALUES (41, "Michael", "Green");
Query OK, 1 row affected (0.05 sec)

mysql> SELECT PARTITION_NAME, TABLE_ROWS
-> FROM INFORMATION_SCHEMA.PARTITIONS
-> WHERE TABLE_NAME = 'e';
```



```

+-----+-----+
| PARTITION_NAME | TABLE_ROWS |
+-----+-----+
| p0              |             1 |
| p1              |             0 |
| p2              |             0 |
| p3              |             3 |
+-----+-----+
4 rows in set (0.00 sec)
    
```

Now we once again exchange partition `p0` with table `e2` using the same `ALTER TABLE` statement as previously:

```

mysql> ALTER TABLE e EXCHANGE PARTITION p0 WITH TABLE e2;
Query OK, 0 rows affected (0.28 sec)
    
```

The output of the following queries shows that the table row that was stored in partition `p0` and the table row that was stored in table `e2`, prior to issuing the `ALTER TABLE` statement, have now switched places:

```

mysql> SELECT * FROM e;
+-----+-----+-----+
| id  | fname | lname |
+-----+-----+-----+
| 16  | Frank | White |
| 1669 | Jim   | Smith |
| 337 | Mary  | Jones |
| 2005 | Linda | Black |
+-----+-----+-----+
4 rows in set (0.00 sec)

mysql> SELECT PARTITION_NAME, TABLE_ROWS
-> FROM INFORMATION_SCHEMA.PARTITIONS
-> WHERE TABLE_NAME = 'e';
+-----+-----+
| PARTITION_NAME | TABLE_ROWS |
+-----+-----+
| p0              |             1 |
| p1              |             0 |
| p2              |             0 |
| p3              |             3 |
+-----+-----+
4 rows in set (0.00 sec)

mysql> SELECT * FROM e2;
+-----+-----+-----+
| id | fname  | lname |
+-----+-----+-----+
| 41 | Michael | Green |
+-----+-----+-----+
1 row in set (0.00 sec)
    
```

You should keep in mind that any rows found in the nonpartitioned table prior to issuing the `ALTER TABLE ... EXCHANGE PARTITION` statement must meet any conditions required for them to be stored in the target partition; otherwise, the statement fails. To see how this occurs, first insert a row into `e2` that cannot be stored in partition `p0` of table `e` because its `id` column value is too large; then, try to exchange the table with the partition again:

```

mysql> INSERT INTO e2 VALUES (51, "Ellen", "McDonald");
Query OK, 1 row affected (0.08 sec)

mysql> ALTER TABLE e EXCHANGE PARTITION p0 WITH TABLE e2;
ERROR 1707 (HY000): Found row that does not match the partition
    
```

The `IGNORE` keyword is accepted, but has no effect when used with `EXCHANGE PARTITION`, as shown here:

```

mysql> ALTER IGNORE TABLE e EXCHANGE PARTITION p0 WITH TABLE e2;
ERROR 1707 (HY000): Found row that does not match the partition
    
```

You can also exchange a subpartition of a subpartitioned table (see [Section 17.2.6, “Subpartitioning”](#)) with a nonpartitioned table using an `ALTER TABLE ... EXCHANGE PARTITION` statement. In the following example, we first create a table `es` that is partitioned by `RANGE` and subpartitioned by `KEY`, populate this table as we did table `e`, and then create an empty, nonpartitioned copy `es2` of the table, as shown here:

```
mysql> CREATE TABLE es (
->   id INT NOT NULL,
->   fname VARCHAR(30),
->   lname VARCHAR(30)
-> )
->   PARTITION BY RANGE (id)
->   SUBPARTITION BY KEY (lname)
->   SUBPARTITIONS 2 (
->     PARTITION p0 VALUES LESS THAN (50),
->     PARTITION p1 VALUES LESS THAN (100),
->     PARTITION p2 VALUES LESS THAN (150),
->     PARTITION p3 VALUES LESS THAN (MAXVALUE)
-> );
Query OK, 0 rows affected (2.76 sec)

mysql> INSERT INTO es VALUES
->   (1669, "Jim", "Smith"),
->   (337, "Mary", "Jones"),
->   (16, "Frank", "White"),
->   (2005, "Linda", "Black");
Query OK, 4 rows affected (0.04 sec)
Records: 4 Duplicates: 0 Warnings: 0

mysql> CREATE TABLE es2 LIKE es;
Query OK, 0 rows affected (1.27 sec)

mysql> ALTER TABLE es2 REMOVE PARTITIONING;
Query OK, 0 rows affected (0.70 sec)
Records: 0 Duplicates: 0 Warnings: 0
```

Although we did not explicitly name any of the subpartitions when creating table `es`, we can obtain generated names for these by including the `SUBPARTITION_NAME` of the `PARTITIONS` table from `INFORMATION_SCHEMA` when selecting from that table, as shown here:

```
mysql> SELECT PARTITION_NAME, SUBPARTITION_NAME, TABLE_ROWS
->   FROM INFORMATION_SCHEMA.PARTITIONS
->   WHERE TABLE_NAME = 'es';
+-----+-----+-----+
| PARTITION_NAME | SUBPARTITION_NAME | TABLE_ROWS |
+-----+-----+-----+
| p0              | p0sp0             | 1           |
| p0              | p0sp1             | 0           |
| p1              | p1sp0             | 0           |
| p1              | p1sp1             | 0           |
| p2              | p2sp0             | 0           |
| p2              | p2sp1             | 0           |
| p3              | p3sp0             | 3           |
| p3              | p3sp1             | 0           |
+-----+-----+-----+
8 rows in set (0.00 sec)
```

The following `ALTER TABLE` statement exchanges subpartition `p3sp0` table `es` with the nonpartitioned table `es2`:

```
mysql> ALTER TABLE es EXCHANGE PARTITION p3sp0 WITH TABLE es2;
Query OK, 0 rows affected (0.29 sec)
```

You can verify that the rows were exchanged by issuing the following queries:

```
mysql> SELECT PARTITION_NAME, SUBPARTITION_NAME, TABLE_ROWS
->   FROM INFORMATION_SCHEMA.PARTITIONS
->   WHERE TABLE_NAME = 'es';
+-----+-----+-----+
| PARTITION_NAME | SUBPARTITION_NAME | TABLE_ROWS |
```

```

+-----+-----+-----+
| p0      | p0sp0      |      1 |
| p0      | p0sp1      |      0 |
| p1      | p1sp0      |      0 |
| p1      | p1sp1      |      0 |
| p2      | p2sp0      |      0 |
| p2      | p2sp1      |      0 |
| p3      | p3sp0      |      0 |
| p3      | p3sp1      |      0 |
+-----+-----+-----+
8 rows in set (0.00 sec)

mysql> SELECT * FROM es2;
+-----+-----+-----+
| id  | fname | lname |
+-----+-----+-----+
| 1669 | Jim   | Smith |
| 337  | Mary  | Jones |
| 2005 | Linda | Black |
+-----+-----+-----+
3 rows in set (0.00 sec)

```

If a table is subpartitioned, you can exchange only a subpartition of the table—not an entire partition—with an unpartitioned table, as shown here:

```

mysql> ALTER TABLE es EXCHANGE PARTITION p3 WITH TABLE es2;
ERROR 1704 (HY000): Subpartitioned table, use subpartition instead of partition

```

The comparison of table structures used by MySQL is very strict. The number, order, names, and types of columns and indexes of the partitioned table and the nonpartitioned table must match exactly. In addition, both tables must use the same storage engine:

```

mysql> CREATE TABLE es3 LIKE e;
Query OK, 0 rows affected (1.31 sec)

mysql> ALTER TABLE es3 REMOVE PARTITIONING;
Query OK, 0 rows affected (0.53 sec)
Records: 0 Duplicates: 0 Warnings: 0

mysql> SHOW CREATE TABLE es3\G
***** 1. row *****
      Table: es3
Create Table: CREATE TABLE `es3` (
  `id` int(11) NOT NULL,
  `fname` varchar(30) DEFAULT NULL,
  `lname` varchar(30) DEFAULT NULL
) ENGINE=InnoDB DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

mysql> ALTER TABLE es3 ENGINE = MyISAM;
Query OK, 0 rows affected (0.15 sec)
Records: 0 Duplicates: 0 Warnings: 0

mysql> ALTER TABLE es EXCHANGE PARTITION p3sp0 WITH TABLE es3;
ERROR 1497 (HY000): The mix of handlers in the partitions is not allowed in this version of MySQL

```

### 17.3.4. Maintenance of Partitions

A number of table and partition maintenance tasks can be carried out using SQL statements intended for such purposes on partitioned tables in MySQL 5.6.

Table maintenance of partitioned tables can be accomplished using the statements `CHECK TABLE`, `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `REPAIR TABLE`, which are supported for partitioned tables.

You can use a number of extensions to `ALTER TABLE` for performing operations of this type on one or more partitions directly, as described in the following list:

- **Rebuilding partitions.** Rebuilds the partition; this has the same effect as dropping all records stored in the partition, then reinserting them. This can be useful for purposes of defragmentation.

Example:

```
ALTER TABLE t1 REBUILD PARTITION p0, p1;
```

- **Optimizing partitions.** If you have deleted a large number of rows from a partition or if you have made many changes to a partitioned table with variable-length rows (that is, having `VARCHAR`, `BLOB`, or `TEXT` columns), you can use `ALTER TABLE ... OPTIMIZE PARTITION` to reclaim any unused space and to defragment the partition data file.

Example:

```
ALTER TABLE t1 OPTIMIZE PARTITION p0, p1;
```

Using `OPTIMIZE PARTITION` on a given partition is equivalent to running `CHECK PARTITION`, `ANALYZE PARTITION`, and `REPAIR PARTITION` on that partition.

- **Analyzing partitions.** This reads and stores the key distributions for partitions.

Example:

```
ALTER TABLE t1 ANALYZE PARTITION p3;
```

- **Repairing partitions.** This repairs corrupted partitions.

Example:

```
ALTER TABLE t1 REPAIR PARTITION p0,p1;
```

- **Checking partitions.** You can check partitions for errors in much the same way that you can use `CHECK TABLE` with nonpartitioned tables.

Example:

```
ALTER TABLE trb3 CHECK PARTITION p1;
```

This command will tell you if the data or indexes in partition `p1` of table `t1` are corrupted. If this is the case, use `ALTER TABLE ... REPAIR PARTITION` to repair the partition.

Each of the statements in the list just shown also supports the keyword `ALL` in place of the list of partition names. Using `ALL` causes the statement to act on all partitions in the table.

The use of `mysqlcheck` and `myisamchk` is not supported with partitioned tables.

In MySQL 5.6, you can also truncate partitions using `ALTER TABLE ... TRUNCATE PARTITION`. This statement can be used to delete all rows from one or more partitions in much the same way that `TRUNCATE TABLE` deletes all rows from a table.

```
ALTER TABLE ... TRUNCATE PARTITION ALL truncates all partitions in the table.
```

## 17.3.5. Obtaining Information About Partitions

This section discusses obtaining information about existing partitions, which can be done in a number of ways. Methods of obtaining such information include the following:

- Using the `SHOW CREATE TABLE` statement to view the partitioning clauses used in creating a partitioned table.
- Using the `SHOW TABLE STATUS` statement to determine whether a table is partitioned.
- Querying the `INFORMATION_SCHEMA.PARTITIONS` table.
- Using the statement `EXPLAIN PARTITIONS SELECT` to see which partitions are used by a given `SELECT`.

As discussed elsewhere in this chapter, `SHOW CREATE TABLE` includes in its output the `PARTITION BY` clause used to create a partitioned table. For example:

```
mysql> SHOW CREATE TABLE trb3\G
***** 1. row *****
      Table: trb3
Create Table: CREATE TABLE `trb3` (
  `id` int(11) default NULL,
  `name` varchar(50) default NULL,
  `purchased` date default NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
PARTITION BY RANGE (YEAR(purchased)) (
  PARTITION p0 VALUES LESS THAN (1990) ENGINE = MyISAM,
  PARTITION p1 VALUES LESS THAN (1995) ENGINE = MyISAM,
  PARTITION p2 VALUES LESS THAN (2000) ENGINE = MyISAM,
  PARTITION p3 VALUES LESS THAN (2005) ENGINE = MyISAM
)
1 row in set (0.00 sec)
```

The output from `SHOW TABLE STATUS` for partitioned tables is the same as that for nonpartitioned tables, except that the `Create_options` column contains the string `partitioned`. The `Engine` column contains the name of the storage engine used by all partitions of the table. (See [Section 13.7.5.37, “SHOW TABLE STATUS Syntax”](#), for more information about this statement.)

You can also obtain information about partitions from `INFORMATION_SCHEMA`, which contains a `PARTITIONS` table. See [Section 19.14, “The INFORMATION\\_SCHEMA PARTITIONS Table”](#).

It is possible to determine which partitions of a partitioned table are involved in a given `SELECT` query using `EXPLAIN PARTITIONS`. The `PARTITIONS` keyword adds a `partitions` column to the output of `EXPLAIN` listing the partitions from which records would be matched by the query.

Suppose that you have a table `trb1` created and populated as follows:

```
CREATE TABLE trb1 (id INT, name VARCHAR(50), purchased DATE)
  PARTITION BY RANGE(id)
  (
    PARTITION p0 VALUES LESS THAN (3),
    PARTITION p1 VALUES LESS THAN (7),
    PARTITION p2 VALUES LESS THAN (9),
    PARTITION p3 VALUES LESS THAN (11)
  );

INSERT INTO trb1 VALUES
  (1, 'desk organiser', '2003-10-15'),
  (2, 'CD player', '1993-11-05'),
  (3, 'TV set', '1996-03-10'),
  (4, 'bookcase', '1982-01-10'),
  (5, 'exercise bike', '2004-05-09'),
  (6, 'sofa', '1987-06-05'),
  (7, 'popcorn maker', '2001-11-22'),
  (8, 'aquarium', '1992-08-04'),
  (9, 'study desk', '1984-09-16'),
  (10, 'lava lamp', '1998-12-25');
```

You can see which partitions are used in a query such as `SELECT * FROM trb1`; as shown here:

```
mysql> EXPLAIN PARTITIONS SELECT * FROM trb1\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: trb1
  partitions: p0,p1,p2,p3
         type: ALL
possible_keys: NULL
          key: NULL
        key_len: NULL
         ref: NULL
         rows: 10
      Extra: Using filesort
```

In this case, all four partitions are searched. However, when a limiting condition making use of the partitioning key is added to the query, you can see that only those partitions containing matching values are scanned, as shown here:

```
mysql> EXPLAIN PARTITIONS SELECT * FROM trb1 WHERE id < 5\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: trb1
  partitions: p0,p1
         type: ALL
possible_keys: NULL
          key: NULL
        key_len: NULL
         ref: NULL
         rows: 10
      Extra: Using where
```

`EXPLAIN PARTITIONS` provides information about keys used and possible keys, just as with the standard `EXPLAIN SELECT` statement:

```
mysql> ALTER TABLE trb1 ADD PRIMARY KEY (id);
Query OK, 10 rows affected (0.03 sec)
Records: 10  Duplicates: 0  Warnings: 0

mysql> EXPLAIN PARTITIONS SELECT * FROM trb1 WHERE id < 5\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: trb1
  partitions: p0,p1
         type: range
possible_keys: PRIMARY
          key: PRIMARY
        key_len: 4
         ref: NULL
         rows: 7
      Extra: Using where
```

You should take note of the following restrictions and limitations on `EXPLAIN PARTITIONS`:

- You cannot use the `PARTITIONS` and `EXTENDED` keywords together in the same `EXPLAIN . . . SELECT` statement. Attempting to do so produces a syntax error.
- If `EXPLAIN PARTITIONS` is used to examine a query against a nonpartitioned table, no error is produced, but the value of the `partitions` column is always `NULL`.

The `rows` column of `EXPLAIN PARTITIONS` output displays the total number of rows in the table.

See also [Section 13.8.2, “EXPLAIN Syntax”](#).

## 17.4. Partition Pruning

This section discusses an optimization known as *partition pruning*. The core concept behind partition pruning is relatively simple, and can be described as “Do not scan partitions where there can be no matching values”. Suppose that you have a partitioned table `t1` defined by this statement:

```
CREATE TABLE t1 (
  fname VARCHAR(50) NOT NULL,
  lname VARCHAR(50) NOT NULL,
  region_code TINYINT UNSIGNED NOT NULL,
  dob DATE NOT NULL
)
PARTITION BY RANGE( region_code ) (
  PARTITION p0 VALUES LESS THAN (64),
  PARTITION p1 VALUES LESS THAN (128),
  PARTITION p2 VALUES LESS THAN (192),
  PARTITION p3 VALUES LESS THAN MAXVALUE
```

```
);
```

Consider the case where you wish to obtain results from a query such as this one:

```
SELECT fname, lname, region_code, dob
   FROM t1
  WHERE region_code > 125 AND region_code < 130;
```

It is easy to see that none of the rows which ought to be returned will be in either of the partitions `p0` or `p3`; that is, we need to search only in partitions `p1` and `p2` to find matching rows. By doing so, it is possible to expend much less time and effort in finding matching rows than would be required to scan all partitions in the table. This “cutting away” of unneeded partitions is known as *pruning*. When the optimizer can make use of partition pruning in performing a query, execution of the query can be an order of magnitude faster than the same query against a nonpartitioned table containing the same column definitions and data.

The query optimizer can perform pruning whenever a `WHERE` condition can be reduced to either one of the following two cases:

- `partition_column = constant`
- `partition_column IN (constant1, constant2, ..., constantN)`

In the first case, the optimizer simply evaluates the partitioning expression for the value given, determines which partition contains that value, and scans only this partition. In many cases, the equal sign can be replaced with another arithmetic comparison, including `<`, `>`, `<=`, `>=`, and `<>`. Some queries using `BETWEEN` in the `WHERE` clause can also take advantage of partition pruning. See the examples later in this section.

In the second case, the optimizer evaluates the partitioning expression for each value in the list, creates a list of matching partitions, and then scans only the partitions in this partition list.

Pruning can also be applied to short ranges, which the optimizer can convert into equivalent lists of values. For instance, in the previous example, the `WHERE` clause can be converted to `WHERE region_code IN (126, 127, 128, 129)`. Then the optimizer can determine that the first three values in the list are found in partition `p1`, the remaining three values in partition `p2`, and that the other partitions contain no relevant values and so do not need to be searched for matching rows.

In MySQL 5.6, the optimizer can also perform pruning for queries that involve comparisons of the preceding types on multiple columns for tables that use `RANGE COLUMNS` or `LIST COLUMNS` partitioning.

This type of optimization can be applied whenever the partitioning expression consists of an equality or a range which can be reduced to a set of equalities, or when the partitioning expression represents an increasing or decreasing relationship. Pruning can also be applied for tables partitioned on a `DATE` or `DATETIME` column when the partitioning expression uses the `YEAR()` [1229] or `TO_DAYS()` [1225] function. In addition, in MySQL 5.6, pruning can be applied for such tables when the partitioning expression uses the `TO_SECONDS()` [1226] function.

Suppose that table `t2`, defined as shown here, is partitioned on a `DATE` column:

```
CREATE TABLE t2 (
  fname VARCHAR(50) NOT NULL,
  lname VARCHAR(50) NOT NULL,
  region_code TINYINT UNSIGNED NOT NULL,
  dob DATE NOT NULL
)
PARTITION BY RANGE( YEAR(dob) ) (
  PARTITION d0 VALUES LESS THAN (1970),
  PARTITION d1 VALUES LESS THAN (1975),
  PARTITION d2 VALUES LESS THAN (1980),
  PARTITION d3 VALUES LESS THAN (1985),
  PARTITION d4 VALUES LESS THAN (1990),
  PARTITION d5 VALUES LESS THAN (2000),
```



```
PARTITION d6 VALUES LESS THAN (2005),
PARTITION d7 VALUES LESS THAN MAXVALUE
);
```

The following queries on `t2` can make use of partition pruning:

```
SELECT * FROM t2 WHERE dob = '1982-06-23';

SELECT * FROM t2 WHERE dob BETWEEN '1991-02-15' AND '1997-04-25';

SELECT * FROM t2 WHERE dob >= '1984-06-21' AND dob <= '1999-06-21'
```

In the case of the last query, the optimizer can also act as follows:

1. Find the partition containing the low end of the range.

`YEAR('1984-06-21')` [1229] yields the value `1984`, which is found in partition `d3`.

2. Find the partition containing the high end of the range.

`YEAR('1999-06-21')` [1229] evaluates to `1999`, which is found in partition `d5`.

3. Scan only these two partitions and any partitions that may lie between them.

In this case, this means that only partitions `d3`, `d4`, and `d5` are scanned. The remaining partitions may be safely ignored (and are ignored).



#### Important

Invalid `DATE` and `DATETIME` values referenced in the `WHERE` clause of a query on a partitioned table are treated as `NULL`. This means that a query such as `SELECT * FROM partitioned_table WHERE date_column < '2008-12-00'` does not return any values (see Bug #40972).

So far, we have looked only at examples using `RANGE` partitioning, but pruning can be applied with other partitioning types as well.

Consider a table that is partitioned by `LIST`, where the partitioning expression is increasing or decreasing, such as the table `t3` shown here. (In this example, we assume for the sake of brevity that the `region_code` column is limited to values between 1 and 10 inclusive.)

```
CREATE TABLE t3 (
  fname VARCHAR(50) NOT NULL,
  lname VARCHAR(50) NOT NULL,
  region_code TINYINT UNSIGNED NOT NULL,
  dob DATE NOT NULL
)
PARTITION BY LIST(region_code) (
  PARTITION r0 VALUES IN (1, 3),
  PARTITION r1 VALUES IN (2, 5, 8),
  PARTITION r2 VALUES IN (4, 9),
  PARTITION r3 VALUES IN (6, 7, 10)
);
```

For a query such as `SELECT * FROM t3 WHERE region_code BETWEEN 1 AND 3`, the optimizer determines in which partitions the values 1, 2, and 3 are found (`r0` and `r1`) and skips the remaining ones (`r2` and `r3`).

For tables that are partitioned by `HASH` or `KEY`, partition pruning is also possible in cases in which the `WHERE` clause uses a simple `=` relation against a column used in the partitioning expression. Consider a table created like this:

```
CREATE TABLE t4 (
  fname VARCHAR(50) NOT NULL,
  lname VARCHAR(50) NOT NULL,
  region_code TINYINT UNSIGNED NOT NULL,
```

```

    dob DATE NOT NULL
)
PARTITION BY KEY(region_code)
PARTITIONS 8;

```

Any query, such as this one, that compares a column value with a constant can be pruned:

```
SELECT * FROM t4 WHERE region_code = 7;
```

Pruning can also be employed for short ranges, because the optimizer can turn such conditions into [IN](#) relations. For example, using the same table `t4` as defined previously, queries such as these can be pruned:

```
SELECT * FROM t4 WHERE region_code > 2 AND region_code < 6;
```

```
SELECT * FROM t4 WHERE region_code BETWEEN 3 AND 5;
```

In both these cases, the `WHERE` clause is transformed by the optimizer into `WHERE region_code IN (3, 4, 5)`.



### Important

This optimization is used only if the range size is smaller than the number of partitions. Consider this query:

```
SELECT * FROM t4 WHERE region_code BETWEEN 4 AND 12;
```

The range in the `WHERE` clause covers 9 values (4, 5, 6, 7, 8, 9, 10, 11, 12), but `t4` has only 8 partitions. This means that the previous query cannot be pruned.

Pruning can be used only on integer columns of tables partitioned by [HASH](#) or [KEY](#). For example, this query on table `t4` cannot use pruning because `dob` is a [DATE](#) column:

```
SELECT * FROM t4 WHERE dob >= '2001-04-14' AND dob <= '2005-10-15';
```

However, if the table stores year values in an [INT](#) column, then a query having `WHERE year_col >= 2001 AND year_col <= 2005` can be pruned.

In MySQL 5.6.8 and later, partition pruning is disabled for all tables using a storage that provides automatic partitioning, such as the [NDB](#) storage engine used by MySQL Cluster (not currently supported in MySQL 5.6). (Bug #14672885) Beginning with MySQL 5.6.10, such tables can be pruned if they are explicitly partitioned. (Bug #14827952)

## 17.5. Partition Selection

MySQL 5.6 supports explicit selection of partitions and subpartitions that, when executing a statement, should be checked for rows matching a given `WHERE` condition. Partition selection is similar to partition pruning, in that only specific partitions are checked for matches, but differs in two key respects:

1. The partitions to be checked are specified by the issuer of the statement, unlike partition pruning, which is automatic.
2. Whereas partition pruning applies only to queries, explicit selection of partitions is supported for both queries and a number of DML statements.

SQL statements supporting explicit partition selection are listed here:

- [SELECT](#)
- [DELETE](#)
- [INSERT](#)
- [REPLACE](#)

- UPDATE
- LOAD DATA.
- LOAD XML.

The remainder of this section discusses explicit partition selection as it applies generally to the statements just listed, and provides some examples.

Explicit partition selection is implemented using a `PARTITION` option. For all supported statements, this option uses the syntax shown here:

```
PARTITION (partition_names)

partition_names:
    partition_name, ...
```

This option always follows the name of the table to which the partition or partitions belong. `partition_names` is a comma-separated list of partitions or subpartitions to be used. Each name in this list must be the name of an existing partition or subpartition of the specified table; if any of the partitions or subpartitions are not found, the statement fails with an error (`partition 'partition_name' doesn't exist`). Partitions and subpartitions named in `partition_names` may be listed in any order, and may overlap.

When the `PARTITION` option is used, only the partitions and subpartitions listed are checked for matching rows. This option can be used in a `SELECT` statement to determine which rows belong to a given partition. Consider a partitioned table named `employees`, created and populated using the statements shown here:

```
SET @@SQL_MODE = '';

CREATE TABLE employees (
    id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
    fname VARCHAR(25) NOT NULL,
    lname VARCHAR(25) NOT NULL,
    store_id INT NOT NULL,
    department_id INT NOT NULL
)
PARTITION BY RANGE(id) (
    PARTITION p0 VALUES LESS THAN (5),
    PARTITION p1 VALUES LESS THAN (10),
    PARTITION p2 VALUES LESS THAN (15),
    PARTITION p3 VALUES LESS THAN MAXVALUE
);

INSERT INTO employees VALUES
    ('', 'Bob', 'Taylor', 3, 2), ('', 'Frank', 'Williams', 1, 2),
    ('', 'Ellen', 'Johnson', 3, 4), ('', 'Jim', 'Smith', 2, 4),
    ('', 'Mary', 'Jones', 1, 1), ('', 'Linda', 'Black', 2, 3),
    ('', 'Ed', 'Jones', 2, 1), ('', 'June', 'Wilson', 3, 1),
    ('', 'Andy', 'Smith', 1, 3), ('', 'Lou', 'Waters', 2, 4),
    ('', 'Jill', 'Stone', 1, 4), ('', 'Roger', 'White', 3, 2),
    ('', 'Howard', 'Andrews', 1, 2), ('', 'Fred', 'Goldberg', 3, 3),
    ('', 'Barbara', 'Brown', 2, 3), ('', 'Alice', 'Rogers', 2, 2),
    ('', 'Mark', 'Morgan', 3, 3), ('', 'Karen', 'Cole', 3, 2);
```

You can see which rows are stored in partition `p1` like this:

```
mysql> SELECT * FROM employees PARTITION (p1);
+----+-----+-----+-----+-----+
| id | fname | lname | store_id | department_id |
+----+-----+-----+-----+-----+
|  5 | Mary  | Jones |         1 |              1 |
|  6 | Linda | Black |         2 |              3 |
|  7 | Ed    | Jones |         2 |              1 |
|  8 | June  | Wilson|         3 |              1 |
|  9 | Andy  | Smith |         1 |              3 |
+----+-----+-----+-----+-----+
5 rows in set (0.00 sec)
```

The result is the same as obtained by the query `SELECT * FROM employees WHERE id BETWEEN 5 AND 9`.

To obtain rows from multiple partitions, supply their names as a comma-delimited list. For example, `SELECT * FROM employees PARTITION (p1, p2)` returns all rows from partitions `p1` and `p2` while excluding rows from the remaining partitions.

Any valid query against a partitioned table can be rewritten with a `PARTITION` option to restrict the result to one or more desired partitions. You can use `WHERE` conditions, `ORDER BY` and `LIMIT` options, and so on. You can also use aggregate functions with `HAVING` and `GROUP BY` options. Each of the following queries produces a valid result when run on the `employees` table as previously defined:

```
mysql> SELECT * FROM employees PARTITION (p0, p2)
->     WHERE lname LIKE 'S%';
+-----+-----+-----+-----+-----+
| id | fname | lname | store_id | department_id |
+-----+-----+-----+-----+-----+
|  4 | Jim   | Smith |         2 |              4 |
| 11 | Jill  | Stone |         1 |              4 |
+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT id, CONCAT(fname, ' ', lname) AS name
->     FROM employees PARTITION (p0) ORDER BY lname;
+-----+-----+
| id | name          |
+-----+-----+
|  3 | Ellen Johnson |
|  4 | Jim Smith     |
|  1 | Bob Taylor    |
|  2 | Frank Williams |
+-----+-----+
4 rows in set (0.06 sec)

mysql> SELECT store_id, COUNT(department_id) AS c
->     FROM employees PARTITION (p1,p2,p3)
->     GROUP BY store_id HAVING c > 4;
+-----+-----+
| c | store_id |
+-----+-----+
|  5 |         2 |
|  5 |         3 |
+-----+-----+
2 rows in set (0.00 sec)
```

Statements using partition selection can be employed with tables using any of the partitioning types supported in MySQL 5.6. When a table is created using `[LINEAR] HASH` or `[LINEAR] KEY` partitioning and the names of the partitions are not specified, MySQL automatically names the partitions `p0`, `p1`, `p2`, ..., `pN-1`, where `N` is the number of partitions. For subpartitions not explicitly named, MySQL assigns automatically to the subpartitions in each partition `pX` the names `pXsp0`, `pXsp1`, `pXsp2`, ..., `pXspM-1`, where `M` is the number of subpartitions. When executing against this table a `SELECT` (or other SQL statement for which explicit partition selection is allowed), you can use these generated names in a `PARTITION` option, as shown here:

```
mysql> CREATE TABLE employees_sub (
->     id INT NOT NULL AUTO_INCREMENT,
->     fname VARCHAR(25) NOT NULL,
->     lname VARCHAR(25) NOT NULL,
->     store_id INT NOT NULL,
->     department_id INT NOT NULL,
->     PRIMARY KEY pk (id, lname)
-> )
->     PARTITION BY RANGE(id)
->     SUBPARTITION BY KEY (lname)
->     SUBPARTITIONS 2 (
->         PARTITION p0 VALUES LESS THAN (5),
->         PARTITION p1 VALUES LESS THAN (10),
->         PARTITION p2 VALUES LESS THAN (15),
```

```

->         PARTITION p3 VALUES LESS THAN MAXVALUE
-> );
Query OK, 0 rows affected (1.14 sec)

mysql> INSERT INTO employees_sub # re-use data in employees table
->     SELECT * FROM employees;
Query OK, 18 rows affected (0.09 sec)
Records: 18 Duplicates: 0 Warnings: 0

mysql> SELECT id, CONCAT(fname, ' ', lname) AS name
->     FROM employees_sub PARTITION (p2sp1);
+-----+-----+
| id | name |
+-----+-----+
| 10 | Lou Waters |
| 14 | Fred Goldberg |
+-----+-----+
2 rows in set (0.00 sec)

```

You may also use a `PARTITION` option in the `SELECT` portion of an `INSERT ... SELECT` statement, as shown here:

```

mysql> CREATE TABLE employees_copy LIKE employees;
Query OK, 0 rows affected (0.28 sec)

mysql> INSERT INTO employees_copy
->     SELECT * FROM employees PARTITION (p2);
Query OK, 5 rows affected (0.04 sec)
Records: 5 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM employees_copy;
+-----+-----+-----+-----+-----+
| id | fname | lname | store_id | department_id |
+-----+-----+-----+-----+-----+
| 10 | Lou | Waters | 2 | 4 |
| 11 | Jill | Stone | 1 | 4 |
| 12 | Roger | White | 3 | 2 |
| 13 | Howard | Andrews | 1 | 2 |
| 14 | Fred | Goldberg | 3 | 3 |
+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)

```

Partition selection can also be used with joins. Suppose we create and populate two tables using the statements shown here:

```

CREATE TABLE stores (
  id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  city VARCHAR(30) NOT NULL
)
PARTITION BY HASH(id)
PARTITIONS 2;

INSERT INTO stores VALUES
  ('', 'Nambucca'), ('', 'Uranga'),
  ('', 'Bellinghen'), ('', 'Grafton');

CREATE TABLE departments (
  id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  name VARCHAR(30) NOT NULL
)
PARTITION BY KEY(id)
PARTITIONS 2;

INSERT INTO departments VALUES
  ('', 'Sales'), ('', 'Customer Service'),
  ('', 'Delivery'), ('', 'Accounting');

```

You can explicitly select partitions (or subpartitions, or both) from any or all of the tables in a join. (Note that the `PARTITION` option used to select partitions from a given table immediately follows the name of the table, before all other options, including any table alias.) For example, the following query gets the name, employee ID, department, and city of all employees who work in the Sales or Delivery

department (partition `p1` of the `departments` table) at the stores in either of the cities of Nambucca and Bellingin (partition `p0` of the `stores` table):

```
mysql> SELECT
->     e.id AS 'Employee ID', CONCAT(e.fname, ' ', e.lname) AS Name,
->     s.city AS City, d.name AS department
-> FROM employees AS e
-> JOIN stores PARTITION (p1) AS s ON e.store_id=s.id
-> JOIN departments PARTITION (p0) AS d ON e.department_id=d.id
-> ORDER BY e.lname;
```

Employee ID	Name	City	department
14	Fred Goldberg	Bellingin	Delivery
5	Mary Jones	Nambucca	Sales
17	Mark Morgan	Bellingin	Delivery
9	Andy Smith	Nambucca	Delivery
8	June Wilson	Bellingin	Sales

5 rows in set (0.00 sec)

For general information about joins in MySQL, see [Section 13.2.9.2, “JOIN Syntax”](#).

When the `PARTITION` option is used with `DELETE` statements, only those partitions (and subpartitions, if any) listed with the option are checked for rows to be deleted. Any other partitions are ignored, as shown here:

```
mysql> SELECT * FROM employees WHERE fname LIKE 'j%';
```

id	fname	lname	store_id	department_id
4	Jim	Smith	2	4
8	June	Wilson	3	1
11	Jill	Stone	1	4

3 rows in set (0.00 sec)

```
mysql> DELETE FROM employees PARTITION (p0, p1)
->     WHERE fname LIKE 'j%';
Query OK, 2 rows affected (0.09 sec)
```

```
mysql> SELECT * FROM employees WHERE fname LIKE 'j%';
```

id	fname	lname	store_id	department_id
11	Jill	Stone	1	4

1 row in set (0.00 sec)

Only the two rows in partitions `p0` and `p1` matching the `WHERE` condition were deleted. As you can see from the result when the `SELECT` is run a second time, there remains a row in the table matching the `WHERE` condition, but residing in a different partition (`p2`).

`UPDATE` statements using explicit partition selection behave in the same way; only rows in the partitions referenced by the `PARTITION` option are considered when determining the rows to be updated, as can be seen by executing the following statements:

```
mysql> UPDATE employees PARTITION (p0)
->     SET store_id = 2 WHERE fname = 'Jill';
Query OK, 0 rows affected (0.00 sec)
Rows matched: 0  Changed: 0  Warnings: 0
```

```
mysql> SELECT * FROM employees WHERE fname = 'Jill';
```

id	fname	lname	store_id	department_id
11	Jill	Stone	1	4

1 row in set (0.00 sec)

```
mysql> UPDATE employees PARTITION (p2)
->     SET store_id = 2 WHERE fname = 'Jill';
Query OK, 1 row affected (0.09 sec)
Rows matched: 1  Changed: 1  Warnings: 0

mysql> SELECT * FROM employees WHERE fname = 'Jill';
+-----+-----+-----+-----+-----+
| id | fname | lname | store_id | department_id |
+-----+-----+-----+-----+-----+
| 11 | Jill  | Stone |         2 |              4 |
+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

In the same way, when `PARTITION` is used with `DELETE`, only rows in the partition or partitions named in the partition list are checked for deletion.

For statements that insert rows, the behavior differs in that failure to find a suitable partition causes the statement to fail. This is true for both `INSERT` and `REPLACE` statements, as shown here:

```
mysql> INSERT INTO employees PARTITION (p2) VALUES (20, 'Jan', 'Jones', 1, 3);
ERROR 1729 (HY000): Found a row not matching the given partition set
mysql> INSERT INTO employees PARTITION (p3) VALUES (20, 'Jan', 'Jones', 1, 3);
Query OK, 1 row affected (0.07 sec)

mysql> REPLACE INTO employees PARTITION (p0) VALUES (20, 'Jan', 'Jones', 3, 2);
ERROR 1729 (HY000): Found a row not matching the given partition set

mysql> REPLACE INTO employees PARTITION (p3) VALUES (20, 'Jan', 'Jones', 3, 2);
Query OK, 2 rows affected (0.09 sec)
```

For statements that write multiple rows to a partitioned table that uses the `InnoDB` storage engine: If any row in the list following `VALUES` cannot be written to one of the partitions specified in the `partition_names` list, the entire statement fails and no rows are written. This is shown for `INSERT` statements in the following example, reusing the `employees` table created previously:

```
mysql> ALTER TABLE employees
->     REORGANIZE PARTITION p3 INTO (
->         PARTITION p3 VALUES LESS THAN (20),
->         PARTITION p4 VALUES LESS THAN (25),
->         PARTITION p5 VALUES LESS THAN MAXVALUE
->     );
Query OK, 6 rows affected (2.09 sec)
Records: 6  Duplicates: 0  Warnings: 0

mysql> SHOW CREATE TABLE employees\G
***** 1. row *****
      Table: employees
Create Table: CREATE TABLE `employees` (
  `id` int(11) NOT NULL AUTO_INCREMENT,
  `fname` varchar(25) NOT NULL,
  `lname` varchar(25) NOT NULL,
  `store_id` int(11) NOT NULL,
  `department_id` int(11) NOT NULL,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB AUTO_INCREMENT=27 DEFAULT CHARSET=latin1
/*!50100 PARTITION BY RANGE (id)
(PARTITION p0 VALUES LESS THAN (5) ENGINE = InnoDB,
PARTITION p1 VALUES LESS THAN (10) ENGINE = InnoDB,
PARTITION p2 VALUES LESS THAN (15) ENGINE = InnoDB,
PARTITION p3 VALUES LESS THAN (20) ENGINE = InnoDB,
PARTITION p4 VALUES LESS THAN (25) ENGINE = InnoDB,
PARTITION p5 VALUES LESS THAN MAXVALUE ENGINE = InnoDB) */
1 row in set (0.00 sec)

mysql> INSERT INTO employees PARTITION (p3, p4) VALUES
->     (24, 'Tim', 'Greene', 3, 1), (26, 'Linda', 'Mills', 2, 1);
ERROR 1729 (HY000): Found a row not matching the given partition set

mysql> INSERT INTO employees PARTITION (p3, p4, p5) VALUES
->     (24, 'Tim', 'Greene', 3, 1), (26, 'Linda', 'Mills', 2, 1);
```



```
Query OK, 2 rows affected (0.06 sec)
Records: 2 Duplicates: 0 Warnings: 0
```

The preceding is true for both `INSERT` statements and `REPLACE` statements that write multiple rows.

In MySQL 5.6.10 and later, partition selection is disabled for tables employing a storage engine that supplies automatic partitioning, such as `NDB`. (Bug #14827952)

## 17.6. Restrictions and Limitations on Partitioning

This section discusses current restrictions and limitations on MySQL partitioning support.

**Prohibited constructs.** The following constructs are not permitted in partitioning expressions:

- Stored procedures, stored functions, UDFs, or plugins.
- Declared variables or user variables.

For a list of SQL functions which are permitted in partitioning expressions, see [Section 17.6.3, “Partitioning Limitations Relating to Functions”](#).

**Arithmetic and logical operators.** Use of the arithmetic operators `+ [1200]`, `- [1200]`, and `* [1200]` is permitted in partitioning expressions. However, the result must be an integer value or `NULL` (except in the case of `[LINEAR] KEY` partitioning, as discussed elsewhere in this chapter; see [Section 17.2, “Partitioning Types”](#), for more information).

The `DIV [1201]` operator is also supported, and the `/ [1200]` operator is not permitted. (Bug #30188, Bug #33182)

The bit operators `| [1258]`, `& [1258]`, `^ [1258]`, `<< [1258]`, `>> [1258]`, and `~ [1259]` are not permitted in partitioning expressions.

**Server SQL mode.** Tables employing user-defined partitioning do not preserve the SQL mode in effect at the time that they were created. As discussed in [Section 5.1.7, “Server SQL Modes”](#), the results of many MySQL functions and operators may change according to the server SQL mode. Therefore, a change in the SQL mode at any time after the creation of partitioned tables may lead to major changes in the behavior of such tables, and could easily lead to corruption or loss of data. For these reasons, *it is strongly recommended that you never change the server SQL mode after creating partitioned tables.*

**Examples.** The following examples illustrate some changes in behavior of partitioned tables due to a change in the server SQL mode:

1. **Error handling.** Suppose that you create a partitioned table whose partitioning expression is one such as `column DIV 0` or `column MOD 0`, as shown here:

```
mysql> CREATE TABLE tn (c1 INT)
->     PARTITION BY LIST(1 DIV c1) (
->     PARTITION p0 VALUES IN (NULL),
->     PARTITION p1 VALUES IN (1)
-> );
Query OK, 0 rows affected (0.05 sec)
```

The default behavior for MySQL is to return `NULL` for the result of a division by zero, without producing any errors:

```
mysql> SELECT @@sql_mode;
+-----+
| @@sql_mode |
+-----+
|             |
+-----+
1 row in set (0.00 sec)
```

```
mysql> INSERT INTO tn VALUES (NULL), (0), (1);
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0
```

However, changing the server SQL mode to treat division by zero as an error and to enforce strict error handling causes the same `INSERT` statement to fail, as shown here:

```
mysql> SET sql_mode='STRICT_ALL_TABLES,ERROR_FOR_DIVISION_BY_ZERO';
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO tn VALUES (NULL), (0), (1);
ERROR 1365 (22012): Division by 0
```

2. **Table accessibility.** Sometimes a change in the server SQL mode can make partitioned tables unusable. The following `CREATE TABLE` statement can be executed successfully only if the `NO_UNSIGNED_SUBTRACTION` [605] mode is in effect:

```
mysql> SELECT @@sql_mode;
+-----+
| @@sql_mode |
+-----+
|           |
+-----+
1 row in set (0.00 sec)

mysql> CREATE TABLE tu (c1 BIGINT UNSIGNED)
-> PARTITION BY RANGE(c1 - 10) (
-> PARTITION p0 VALUES LESS THAN (-5),
-> PARTITION p1 VALUES LESS THAN (0),
-> PARTITION p2 VALUES LESS THAN (5),
-> PARTITION p3 VALUES LESS THAN (10),
-> PARTITION p4 VALUES LESS THAN (MAXVALUE)
-> );
ERROR 1563 (HY000): Partition constant is out of partition function domain

mysql> SET sql_mode='NO_UNSIGNED_SUBTRACTION';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @@sql_mode;
+-----+
| @@sql_mode |
+-----+
| NO_UNSIGNED_SUBTRACTION |
+-----+
1 row in set (0.00 sec)

mysql> CREATE TABLE tu (c1 BIGINT UNSIGNED)
-> PARTITION BY RANGE(c1 - 10) (
-> PARTITION p0 VALUES LESS THAN (-5),
-> PARTITION p1 VALUES LESS THAN (0),
-> PARTITION p2 VALUES LESS THAN (5),
-> PARTITION p3 VALUES LESS THAN (10),
-> PARTITION p4 VALUES LESS THAN (MAXVALUE)
-> );
Query OK, 0 rows affected (0.05 sec)
```

If you remove the `NO_UNSIGNED_SUBTRACTION` [605] server SQL mode after creating `tu`, you may no longer be able to access this table:

```
mysql> SET sql_mode='';
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT * FROM tu;
ERROR 1563 (HY000): Partition constant is out of partition function domain
mysql> INSERT INTO tu VALUES (20);
ERROR 1563 (HY000): Partition constant is out of partition function domain
```

Server SQL modes also impact replication of partitioned tables. Differing SQL modes on master and slave can lead to partitioning expressions being evaluated differently; this can cause the distribution of

data among partitions to be different in the master's and slave's copies of a given table, and may even cause inserts into partitioned tables that succeed on the master to fail on the slave. For best results, you should always use the same server SQL mode on the master and on the slave.

**Performance considerations.** Some affects of partitioning operations on performance are given in the following list:

- **File system operations.** Partitioning and repartitioning operations (such as `ALTER TABLE` with `PARTITION BY ...`, `REORGANIZE PARTITIONS`, or `REMOVE PARTITIONING`) depend on file system operations for their implementation. This means that the speed of these operations is affected by such factors as file system type and characteristics, disk speed, swap space, file handling efficiency of the operating system, and MySQL server options and variables that relate to file handling. In particular, you should make sure that `large_files_support` [494] is enabled and that `open_files_limit` [518] is set properly. For partitioned tables using the `MyISAM` storage engine, increasing `myisam_max_sort_file_size` [511] may improve performance; partitioning and repartitioning operations involving `InnoDB` tables may be made more efficient by enabling `innodb_file_per_table` [1718].

See also [Maximum number of partitions](#).

- **Table locks.** The process executing a partitioning operation on a table takes a write lock on the table. Reads from such tables are relatively unaffected; pending `INSERT` and `UPDATE` operations are performed as soon as the partitioning operation has completed.
- **Storage engine.** Partitioning operations, queries, and update operations generally tend to be faster with `MyISAM` tables than with `InnoDB` tables.
- **Use of indexes and partition pruning.** As with nonpartitioned tables, proper use of indexes can speed up queries on partitioned tables significantly. In addition, designing partitioned tables and queries on these tables to take advantage of *partition pruning* can improve performance dramatically. See [Section 17.4, "Partition Pruning"](#), for more information.
- **Performance with `LOAD DATA`.** In MySQL 5.6, `LOAD DATA` uses buffering to improve performance. You should be aware that the buffer uses 130 KB memory per partition to achieve this.

#### **Maximum number of partitions.**

Prior to MySQL 5.6.7, the maximum possible number of partitions for a given table was 1024. Beginning with MySQL 5.6.7, this limit is increased to 8192 partitions. Regardless of the MySQL Server version, this maximum includes subpartitions.

If, when creating tables with a large number of partitions (but less than the maximum), you encounter an error message such as `Got error ... from storage engine: Out of resources when opening file`, you may be able to address the issue by increasing the value of the `open_files_limit` [518] system variable. However, this is dependent on the operating system, and may not be possible or advisable on all platforms; see [Section C.5.2.18, "'File' Not Found and Similar Errors"](#), for more information. In some cases, using large numbers (hundreds) of partitions may also not be advisable due to other concerns, so using more partitions does not automatically lead to better results.

See also [File system operations](#).

#### **Query cache not supported.**

The query cache is not supported for partitioned tables. Beginning with MySQL 5.6.5, the query cache is automatically disabled for queries involving partitioned tables, and cannot be enabled for such queries. (Bug #53775)

#### **Per-partition key caches.**

In MySQL 5.6, key caches are supported for partitioned `MyISAM` tables, using the `CACHE INDEX` and `LOAD INDEX INTO CACHE` statements. Key caches may be defined for one, several, or all partitions, and indexes for one, several, or all partitions may be preloaded into key caches.

**Foreign keys not supported for partitioned InnoDB tables.**

Partitioned tables using the [InnoDB](#) storage engine do not support foreign keys. More specifically, this means that the following two statements are true:

1. No definition of an [InnoDB](#) table employing user-defined partitioning may contain foreign key references; no [InnoDB](#) table whose definition contains foreign key references may be partitioned.
2. No [InnoDB](#) table definition may contain a foreign key reference to a user-partitioned table; no [InnoDB](#) table with user-defined partitioning may contain columns referenced by foreign keys.

The scope of the restrictions just listed includes all tables that use the [InnoDB](#) storage engine. [CREATE TABLE](#) and [ALTER TABLE](#) statements that would result in tables violating these restrictions are not allowed.

**ALTER TABLE ... ORDER BY.** An [ALTER TABLE ... ORDER BY column](#) statement run against a partitioned table causes ordering of rows only within each partition.

**FULLTEXT indexes.**

Partitioned tables do not support [FULLTEXT](#) indexes or searches, even for partitioned tables employing the [InnoDB](#) or [MyISAM](#) storage engine.

**Spatial columns.** Columns with spatial data types such as [POINT](#) or [GEOMETRY](#) cannot be used in partitioned tables.

**Temporary tables.**

Temporary tables cannot be partitioned. (Bug #17497)

**Log tables.** It is not possible to partition the log tables; an [ALTER TABLE ... PARTITION BY ...](#) statement on such a table fails with an error.

**Data type of partitioning key.**

A partitioning key must be either an integer column or an expression that resolves to an integer. The column or expression value may also be [NULL](#). (See [Section 17.2.7](#), “How MySQL Partitioning Handles NULL”.)

There are two exceptions to this restriction:

1. When partitioning by [\[LINEAR\] KEY](#), it is possible to use columns of other types as partitioning keys, because MySQL's internal key-hashing functions produce the correct data type from these types. For example, the following [CREATE TABLE](#) statement is valid:

```
CREATE TABLE tkc (c1 CHAR)
PARTITION BY KEY(c1)
PARTITIONS 4;
```

2. When partitioning by [RANGE COLUMNS](#) or [LIST COLUMNS](#), it is possible to use string, [DATE](#), and [DATETIME](#) columns. For example, each of the following [CREATE TABLE](#) statements is valid:

```
CREATE TABLE rc (c1 INT, c2 DATE)
PARTITION BY RANGE COLUMNS(c2) (
  PARTITION p0 VALUES LESS THAN('1990-01-01'),
  PARTITION p1 VALUES LESS THAN('1995-01-01'),
  PARTITION p2 VALUES LESS THAN('2000-01-01'),
  PARTITION p3 VALUES LESS THAN('2005-01-01'),
  PARTITION p4 VALUES LESS THAN(MAXVALUE)
);

CREATE TABLE lc (c1 INT, c2 CHAR(1))
PARTITION BY LIST COLUMNS(c2) (
  PARTITION p0 VALUES IN('a', 'd', 'g', 'j', 'm', 'p', 's', 'v', 'y'),
  PARTITION p1 VALUES IN('b', 'e', 'h', 'k', 'n', 'q', 't', 'w', 'z'),
  PARTITION p2 VALUES IN('c', 'f', 'i', 'l', 'o', 'r', 'u', 'x', NULL)
);
```

Neither of the preceding exceptions applies to [BLOB](#) or [TEXT](#) column types.

**Subqueries.**

A partitioning key may not be a subquery, even if that subquery resolves to an integer value or `NULL`.

**Issues with subpartitions.**

Subpartitions must use `HASH` or `KEY` partitioning. Only `RANGE` and `LIST` partitions may be subpartitioned; `HASH` and `KEY` partitions cannot be subpartitioned.

Currently, `SUBPARTITION BY KEY` requires that the subpartitioning column or columns be specified explicitly, unlike the case with `PARTITION BY KEY`, where it can be omitted (in which case the table's primary key column is used by default). Consider the table created by this statement:

```
CREATE TABLE ts (
  id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  name VARCHAR(30)
);
```

You can create a table having the same columns, partitioned by `KEY`, using a statement such as this one:

```
CREATE TABLE ts (
  id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  name VARCHAR(30)
)
PARTITION BY KEY()
PARTITIONS 4;
```

The previous statement is treated as though it had been written like this, with the table's primary key column used as the partitioning column:

```
CREATE TABLE ts (
  id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  name VARCHAR(30)
)
PARTITION BY KEY(id)
PARTITIONS 4;
```

However, the following statement that attempts to create a subpartitioned table using the default column as the subpartitioning column fails, and the column must be specified for the statement to succeed, as shown here:

```
mysql> CREATE TABLE ts (
->   id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   name VARCHAR(30)
-> )
-> PARTITION BY RANGE(id)
-> SUBPARTITION BY KEY()
-> SUBPARTITIONS 4
-> (
->   PARTITION p0 VALUES LESS THAN (100),
->   PARTITION p1 VALUES LESS THAN (MAXVALUE)
-> );
ERROR 1064 (42000): You have an error in your SQL syntax; check the manual that
corresponds to your MySQL server version for the right syntax to use near ')'

mysql> CREATE TABLE ts (
->   id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   name VARCHAR(30)
-> )
-> PARTITION BY RANGE(id)
-> SUBPARTITION BY KEY(id)
-> SUBPARTITIONS 4
-> (
->   PARTITION p0 VALUES LESS THAN (100),
->   PARTITION p1 VALUES LESS THAN (MAXVALUE)
-> );
Query OK, 0 rows affected (0.07 sec)
```

This is a known issue (see Bug #51470).

**DELAYED option not supported.** Use of `INSERT DELAYED` to insert rows into a partitioned table is not supported. Attempting to do so fails with an error.

**DATA DIRECTORY and INDEX DIRECTORY options.** `DATA DIRECTORY` and `INDEX DIRECTORY` are subject to the following restrictions when used with partitioned tables:

- Table-level `DATA DIRECTORY` and `INDEX DIRECTORY` options are ignored (see Bug #32091).
- On Windows, the `DATA DIRECTORY` and `INDEX DIRECTORY` options are not supported for individual partitions or subpartitions of `MyISAM` tables (Bug #30459). However, you can use `DATA DIRECTORY` for individual partitions or subpartitions of `InnoDB` tables.

**Repairing and rebuilding partitioned tables.** The statements `CHECK TABLE`, `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `REPAIR TABLE` are supported for partitioned tables.

In addition, you can use `ALTER TABLE ... REBUILD PARTITION` to rebuild one or more partitions of a partitioned table; `ALTER TABLE ... REORGANIZE PARTITION` also causes partitions to be rebuilt. See [Section 13.1.6, “ALTER TABLE Syntax”](#), for more information about these two statements.

`mysqlcheck` and `myisamchk` are not supported with partitioned tables.

### 17.6.1. Partitioning Keys, Primary Keys, and Unique Keys

This section discusses the relationship of partitioning keys with primary keys and unique keys. The rule governing this relationship can be expressed as follows: All columns used in the partitioning expression for a partitioned table must be part of every unique key that the table may have.

In other words, *every unique key on the table must use every column in the table's partitioning expression*. (This also includes the table's primary key, since it is by definition a unique key. This particular case is discussed later in this section.) For example, each of the following table creation statements is invalid:

```
CREATE TABLE t1 (
  col1 INT NOT NULL,
  col2 DATE NOT NULL,
  col3 INT NOT NULL,
  col4 INT NOT NULL,
  UNIQUE KEY (col1, col2)
)
PARTITION BY HASH(col3)
PARTITIONS 4;

CREATE TABLE t2 (
  col1 INT NOT NULL,
  col2 DATE NOT NULL,
  col3 INT NOT NULL,
  col4 INT NOT NULL,
  UNIQUE KEY (col1),
  UNIQUE KEY (col3)
)
PARTITION BY HASH(col1 + col3)
PARTITIONS 4;
```

In each case, the proposed table would have at least one unique key that does not include all columns used in the partitioning expression.

Each of the following statements is valid, and represents one way in which the corresponding invalid table creation statement could be made to work:

```
CREATE TABLE t1 (
  col1 INT NOT NULL,
  col2 DATE NOT NULL,
  col3 INT NOT NULL,
```

```

    col4 INT NOT NULL,
    UNIQUE KEY (col1, col2, col3)
)
PARTITION BY HASH(col3)
PARTITIONS 4;

CREATE TABLE t2 (
    col1 INT NOT NULL,
    col2 DATE NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    UNIQUE KEY (col1, col3)
)
PARTITION BY HASH(col1 + col3)
PARTITIONS 4;

```

This example shows the error produced in such cases:

```

mysql> CREATE TABLE t3 (
->     col1 INT NOT NULL,
->     col2 DATE NOT NULL,
->     col3 INT NOT NULL,
->     col4 INT NOT NULL,
->     UNIQUE KEY (col1, col2),
->     UNIQUE KEY (col3)
-> )
-> PARTITION BY HASH(col1 + col3)
-> PARTITIONS 4;
ERROR 1491 (HY000): A PRIMARY KEY must include all columns in the table's partitioning function

```

The `CREATE TABLE` statement fails because both `col1` and `col3` are included in the proposed partitioning key, but neither of these columns is part of both of unique keys on the table. This shows one possible fix for the invalid table definition:

```

mysql> CREATE TABLE t3 (
->     col1 INT NOT NULL,
->     col2 DATE NOT NULL,
->     col3 INT NOT NULL,
->     col4 INT NOT NULL,
->     UNIQUE KEY (col1, col2, col3),
->     UNIQUE KEY (col3)
-> )
-> PARTITION BY HASH(col3)
-> PARTITIONS 4;
Query OK, 0 rows affected (0.05 sec)

```

In this case, the proposed partitioning key `col3` is part of both unique keys, and the table creation statement succeeds.

The following table cannot be partitioned at all, because there is no way to include in a partitioning key any columns that belong to both unique keys:

```

CREATE TABLE t4 (
    col1 INT NOT NULL,
    col2 INT NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    UNIQUE KEY (col1, col3),
    UNIQUE KEY (col2, col4)
);

```

Since every primary key is by definition a unique key, this restriction also includes the table's primary key, if it has one. For example, the next two statements are invalid:

```

CREATE TABLE t5 (
    col1 INT NOT NULL,
    col2 DATE NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    PRIMARY KEY(col1, col2)

```



```

)
PARTITION BY HASH(col3)
PARTITIONS 4;

CREATE TABLE t6 (
    col1 INT NOT NULL,
    col2 DATE NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    PRIMARY KEY(col1, col3),
    UNIQUE KEY(col2)
)
PARTITION BY HASH( YEAR(col2) )
PARTITIONS 4;

```

In both cases, the primary key does not include all columns referenced in the partitioning expression. However, both of the next two statements are valid:

```

CREATE TABLE t7 (
    col1 INT NOT NULL,
    col2 DATE NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    PRIMARY KEY(col1, col2)
)
PARTITION BY HASH(col1 + YEAR(col2))
PARTITIONS 4;

CREATE TABLE t8 (
    col1 INT NOT NULL,
    col2 DATE NOT NULL,
    col3 INT NOT NULL,
    col4 INT NOT NULL,
    PRIMARY KEY(col1, col2, col4),
    UNIQUE KEY(col2, col1)
)
PARTITION BY HASH(col1 + YEAR(col2))
PARTITIONS 4;

```

If a table has no unique keys—this includes having no primary key—then this restriction does not apply, and you may use any column or columns in the partitioning expression as long as the column type is compatible with the partitioning type.

For the same reason, you cannot later add a unique key to a partitioned table unless the key includes all columns used by the table's partitioning expression. Consider the partitioned table created as shown here:

```

mysql> CREATE TABLE t_no_pk (c1 INT, c2 INT)
->     PARTITION BY RANGE(c1) (
->         PARTITION p0 VALUES LESS THAN (10),
->         PARTITION p1 VALUES LESS THAN (20),
->         PARTITION p2 VALUES LESS THAN (30),
->         PARTITION p3 VALUES LESS THAN (40)
->     );
Query OK, 0 rows affected (0.12 sec)

```

It is possible to add a primary key to `t_no_pk` using either of these `ALTER TABLE` statements:

```

# possible PK
mysql> ALTER TABLE t_no_pk ADD PRIMARY KEY(c1);
Query OK, 0 rows affected (0.13 sec)
Records: 0 Duplicates: 0 Warnings: 0

# drop this PK
mysql> ALTER TABLE t_no_pk DROP PRIMARY KEY;
Query OK, 0 rows affected (0.10 sec)
Records: 0 Duplicates: 0 Warnings: 0

# use another possible PK
mysql> ALTER TABLE t_no_pk ADD PRIMARY KEY(c1, c2);

```

```
Query OK, 0 rows affected (0.12 sec)
Records: 0 Duplicates: 0 Warnings: 0

# drop this PK
mysql> ALTER TABLE t_no_pk DROP PRIMARY KEY;
Query OK, 0 rows affected (0.09 sec)
Records: 0 Duplicates: 0 Warnings: 0
```

However, the next statement fails, because `c1` is part of the partitioning key, but is not part of the proposed primary key:

```
# fails with error 1503
mysql> ALTER TABLE t_no_pk ADD PRIMARY KEY(c2);
ERROR 1503 (HY000): A PRIMARY KEY must include all columns in the table's partitioning function
```

Since `t_no_pk` has only `c1` in its partitioning expression, attempting to adding a unique key on `c2` alone fails. However, you can add a unique key that uses both `c1` and `c2`.

These rules also apply to existing nonpartitioned tables that you wish to partition using `ALTER TABLE ... PARTITION BY`. Consider a table `np_pk` created as shown here:

```
mysql> CREATE TABLE np_pk (
->   id INT NOT NULL AUTO_INCREMENT,
->   name VARCHAR(50),
->   added DATE,
->   PRIMARY KEY (id)
-> );
Query OK, 0 rows affected (0.08 sec)
```

The following `ALTER TABLE` statement fails with an error, because the `added` column is not part of any unique key in the table:

```
mysql> ALTER TABLE np_pk
->   PARTITION BY HASH( TO_DAYS(added) )
->   PARTITIONS 4;
ERROR 1503 (HY000): A PRIMARY KEY must include all columns in the table's partitioning function
```

However, this statement using the `id` column for the partitioning column is valid, as shown here:

```
mysql> ALTER TABLE np_pk
->   PARTITION BY HASH(id)
->   PARTITIONS 4;
Query OK, 0 rows affected (0.11 sec)
Records: 0 Duplicates: 0 Warnings: 0
```

In the case of `np_pk`, the only column that may be used as part of a partitioning expression is `id`; if you wish to partition this table using any other column or columns in the partitioning expression, you must first modify the table, either by adding the desired column or columns to the primary key, or by dropping the primary key altogether.

## 17.6.2. Partitioning Limitations Relating to Storage Engines

The following limitations apply to the use of storage engines with user-defined partitioning of tables.

**MERGE storage engine.** User-defined partitioning and the `MERGE` storage engine are not compatible. Tables using the `MERGE` storage engine cannot be partitioned. Partitioned tables cannot be merged.

**FEDERATED storage engine.** Partitioning of `FEDERATED` tables is not supported; it is not possible to create partitioned `FEDERATED` tables.

**CSV storage engine.** Partitioned tables using the `CSV` storage engine are not supported; it is not possible to create partitioned `CSV` tables.

**InnoDB storage engine.** `InnoDB` foreign keys and MySQL partitioning are not compatible. Partitioned `InnoDB` tables cannot have foreign key references, nor can they have columns referenced

by foreign keys. [InnoDB](#) tables which have or which are referenced by foreign keys cannot be partitioned.

In addition, `ALTER TABLE ... OPTIMIZE PARTITION` does not work correctly with partitioned tables that use the [InnoDB](#) storage engine. Use `ALTER TABLE ... REBUILD PARTITION` and `ALTER TABLE ... ANALYZE PARTITION`, instead, for such tables. For more information, see [Section 13.1.6.1, "ALTER TABLE Partition Operations"](#).

**Upgrading partitioned tables.** When performing an upgrade, tables which are partitioned by `KEY` must be dumped and reloaded.

**Same storage engine for all partitions.** All partitions of a partitioned table must use the same storage engine and it must be the same storage engine used by the table as a whole. In addition, if one does not specify an engine on the table level, then one must do either of the following when creating or altering a partitioned table:

- Do *not* specify any engine for *any* partition or subpartition
- Specify the engine for *all* partitions or subpartitions

### 17.6.3. Partitioning Limitations Relating to Functions

This section discusses limitations in MySQL Partitioning relating specifically to functions used in partitioning expressions.

Only the MySQL functions shown in the following table are supported in partitioning expressions.

<a href="#">ABS() [1202]</a>	<a href="#">CEILING() [1203]</a> (see <a href="#">CEILING()</a> and <a href="#">FLOOR()</a> )	<a href="#">DAY()</a> [1217]
<a href="#">DAYOFMONTH()</a> [1217]	<a href="#">DAYOFWEEK()</a> [1217]	<a href="#">DAYOFYEAR()</a> [1217]
<a href="#">DATEDIFF()</a> [1213]	<a href="#">EXTRACT()</a> [1217] (see <a href="#">EXTRACT()</a> function with <code>WEEK</code> specifier)	<a href="#">FLOOR()</a> [1204] (see <a href="#">CEILING()</a> and <a href="#">FLOOR()</a> )
<a href="#">HOUR()</a> [1219]	<a href="#">MICROSECOND()</a> [1220]	<a href="#">MINUTE()</a> [1220]
<a href="#">MOD()</a> [1205]	<a href="#">MONTH()</a> [1220]	<a href="#">QUARTER()</a> [1221]
<a href="#">SECOND()</a> [1221]	<a href="#">TIME_TO_SEC()</a> [1225]	<a href="#">TO_DAYS()</a> [1225]
<a href="#">TO_SECONDS()</a> [1226]	<a href="#">UNIX_TIMESTAMP()</a> [1226] (permitted beginning with MySQL 5.6.1 and fully supported beginning with MySQL 5.6.3, with <a href="#">TIMESTAMP</a> columns)	<a href="#">WEEKDAY()</a> [1228]
<a href="#">YEAR()</a> [1229]		<a href="#">YEARWEEK()</a> [1229]

In MySQL 5.6, partition pruning is supported for the [TO\\_DAYS\(\)](#) [1225], [TO\\_SECONDS\(\)](#) [1226], and [YEAR\(\)](#) [1229] functions. In addition, beginning with MySQL 5.6.3, [UNIX\\_TIMESTAMP\(\)](#) [1226] is treated as monotonic in partitioning expressions. See [Section 17.4, "Partition Pruning"](#), for more information.

**CEILING() and FLOOR().** Each of these functions returns an integer only if it is passed an argument of an exact numeric type, such as one of the [INT](#) types or [DECIMAL](#). This means, for example, that the following `CREATE TABLE` statement fails with an error, as shown here:

```
mysql> CREATE TABLE t (c FLOAT) PARTITION BY LIST( FLOOR(c) )(
->   PARTITION p0 VALUES IN (1,3,5),
->   PARTITION p1 VALUES IN (2,4,6)
-> );
ERROR 1490 (HY000): The PARTITION function returns the wrong type
```

**EXTRACT() function with WEEK specifier.** The value returned by the `EXTRACT()` [1217] function, when used as `EXTRACT(WEEK FROM col)` [1217], depends on the value of the `default_week_format` [475] system variable. For this reason, beginning with MySQL 5.6.2, `EXTRACT()` [1217] is longer permitted as a partitioning function when it specifies the unit as `WEEK`. (Bug #54483)

See [Section 12.6.2, “Mathematical Functions”](#), for more information about the return types of these functions, as well as [Section 11.2, “Numeric Types”](#).

## 17.6.4. Partitioning and Locking

In MySQL 5.6.5 and earlier, for storage engines such as `MyISAM` that actually execute table-level locks when executing DML or DDL statements, such a statement affecting a partitioned table imposed a lock on the table as a whole; that is, all partitions were locked until the statement was finished. MySQL 5.6.6 implements *partition lock pruning*, which eliminates unneeded locks in many cases. In MySQL 5.6.6 and later, most statements reading from or updating a partitioned `MyISAM` table cause only the effected partitions to be locked. For example, prior to MySQL 5.6.6, a `SELECT` from a partitioned `MyISAM` table caused a lock on the entire table; in MySQL 5.6.6 and later, only those partitions actually containing rows that satisfy the `SELECT` statement's `WHERE` condition are locked. This has the effect of increasing the speed and efficiency of concurrent operations on partitioned `MyISAM` tables. This improvement becomes particularly noticeable when working with `MyISAM` tables that have many (32 or more) partitions.

This change in behavior does not have any impact on statements effecting partitioned tables using storage engines such as `InnoDB`, that employ row-level locking and do not actually perform (or need to perform) the locks prior to partition pruning.

The next few paragraphs discuss the effects of partition lock pruning for various MySQL statements on tables using storage engines that employ table-level locks.

### Affects on DML statements

`SELECT` statements (including those containing unions or joins) now lock only those partitions that actually need to be read. This also applies to `SELECT ... PARTITION`.

An `UPDATE` prunes locks only for tables on which no partitioning columns are updated.

`REPLACE` and `INSERT` now lock only those partitions having rows to be inserted or replaced. However, if an `AUTO_INCREMENT` value is generated for any partitioning column then all partitions are locked.

`INSERT ... ON DUPLICATE KEY UPDATE` is pruned as long as no partitioning column is updated.

`INSERT ... SELECT` now locks only those partitions in the source table that need to be read, although all partitions in the target table are locked.



#### Note

`INSERT DELAYED` is not supported for partitioned tables.

Locks imposed by `LOAD DATA` statements on partitioned tables cannot be pruned.

The presence of `BEFORE INSERT` or `BEFORE UPDATE` triggers using any partitioning column of a partitioned table means that locks on `INSERT` and `UPDATE` statements updating this table cannot be pruned, since the trigger can alter its values: A `BEFORE INSERT` trigger on any of the table's partitioning columns means that locks set by `INSERT` or `REPLACE` cannot be pruned, since the `BEFORE INSERT` trigger may change a row's partitioning columns before the row is inserted, forcing the row into a different partition than it would be otherwise. A `BEFORE UPDATE` trigger on a partitioning column means that locks imposed by `UPDATE` or `INSERT ... ON DUPLICATE KEY UPDATE` cannot be pruned.

## Affected DDL statements

`CREATE VIEW` no longer causes any locks.

`ALTER TABLE ... EXCHANGE PARTITION` now prunes locks; only the exchanged table and the exchanged partition are locked.

`ALTER TABLE ... TRUNCATE PARTITION` now prunes locks; only the partitions to be emptied are locked.

`ALTER TABLE` statements still take metadata locks on the table level.

## Other statements

`LOCK TABLES` cannot prune partition locks.

`CALL stored_procedure(expr)` supports lock pruning, but evaluating `expr` does not.

`DO` and `SET` statements do not support partitioning lock pruning.

---

# Chapter 18. Stored Programs and Views

## Table of Contents

18.1. Defining Stored Programs .....	2152
18.2. Using Stored Routines (Procedures and Functions) .....	2153
18.2.1. Stored Routine Syntax .....	2153
18.2.2. Stored Routines and MySQL Privileges .....	2154
18.2.3. Stored Routine Metadata .....	2154
18.2.4. Stored Procedures, Functions, Triggers, and <code>LAST_INSERT_ID()</code> .....	2155
18.3. Using Triggers .....	2155
18.3.1. Trigger Syntax .....	2156
18.3.2. Trigger Metadata .....	2158
18.4. Using the Event Scheduler .....	2158
18.4.1. Event Scheduler Overview .....	2158
18.4.2. Event Scheduler Configuration .....	2159
18.4.3. Event Syntax .....	2161
18.4.4. Event Metadata .....	2161
18.4.5. Event Scheduler Status .....	2162
18.4.6. The Event Scheduler and MySQL Privileges .....	2163
18.5. Using Views .....	2165
18.5.1. View Syntax .....	2166
18.5.2. View Processing Algorithms .....	2166
18.5.3. Updatable and Insertable Views .....	2167
18.5.4. View Metadata .....	2169
18.6. Access Control for Stored Programs and Views .....	2169
18.7. Binary Logging of Stored Programs .....	2171

This chapter discusses stored programs and views, which are database objects defined in terms of SQL code that is stored on the server for later execution.

Stored programs include these objects:

- Stored routines, that is, stored procedures and functions. A stored procedure is invoked using the `CALL` statement. A procedure does not have a return value but can modify its parameters for later inspection by the caller. It can also generate result sets to be returned to the client program. A stored function is used much like a built-in function. you invoke it in an expression and it returns a value during expression evaluation.
- Triggers. A trigger is a named database object that is associated with a table and that is activated when a particular event occurs for the table, such as an insert or update.
- Events. An event is a task that the server runs according to schedule.

Views are stored queries that when referenced produce a result set. A view acts as a virtual table.

This chapter describes how to use stored programs and views. The following sections provide additional information about SQL syntax for statements related to these objects:

- For each object type, there are `CREATE`, `ALTER`, and `DROP` statements that control which objects exist and how they are defined. See [Section 13.1, “Data Definition Statements”](#).
- The `CALL` statement is used to invoke stored procedures. See [Section 13.2.1, “CALL Syntax”](#).
- Stored program definitions include a body that may use compound statements, loops, conditionals, and declared variables. See [Section 13.6, “MySQL Compound-Statement Syntax”](#).

As of MySQL 5.6.6, metadata changes to objects referred to by stored programs are detected and cause automatic reparsing of the affected statements when the program is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

## 18.1. Defining Stored Programs

Each stored program contains a body that consists of an SQL statement. This statement may be a compound statement made up of several statements separated by semicolon (;) characters. For example, the following stored procedure has a body made up of a `BEGIN ... END` block that contains a `SET` statement and a `REPEAT` loop that itself contains another `SET` statement:

```
CREATE PROCEDURE dorepeat(p1 INT)
BEGIN
  SET @x = 0;
  REPEAT SET @x = @x + 1; UNTIL @x > p1 END REPEAT;
END;
```

If you use the `mysql` client program to define a stored program containing semicolon characters, a problem arises. By default, `mysql` itself recognizes the semicolon as a statement delimiter, so you must redefine the delimiter temporarily to cause `mysql` to pass the entire stored program definition to the server.

To redefine the `mysql` delimiter, use the `delimiter` command. The following example shows how to do this for the `dorepeat()` procedure just shown. The delimiter is changed to `//` to enable the entire definition to be passed to the server as a single statement, and then restored to `;` before invoking the procedure. This enables the `;` delimiter used in the procedure body to be passed through to the server rather than being interpreted by `mysql` itself.

```
mysql> delimiter //

mysql> CREATE PROCEDURE dorepeat(p1 INT)
-> BEGIN
->   SET @x = 0;
->   REPEAT SET @x = @x + 1; UNTIL @x > p1 END REPEAT;
-> END
-> //
Query OK, 0 rows affected (0.00 sec)

mysql> delimiter ;

mysql> CALL dorepeat(1000);
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT @x;
+-----+
| @x   |
+-----+
| 1001 |
+-----+
1 row in set (0.00 sec)
```

You can redefine the delimiter to a string other than `//`, and the delimiter can consist of a single character or multiple characters. You should avoid the use of the backslash (“\”) character because that is the escape character for MySQL.

The following is an example of a function that takes a parameter, performs an operation using an SQL function, and returns the result. In this case, it is unnecessary to use `delimiter` because the function definition contains no internal `;` statement delimiters:

```
mysql> CREATE FUNCTION hello (s CHAR(20))
mysql> RETURNS CHAR(50) DETERMINISTIC
-> RETURN CONCAT('Hello, ',s,'!');
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT hello('world');
+-----+
| hello('world') |
```



```
+-----+
| Hello, world! |
+-----+
1 row in set (0.00 sec)
```

## 18.2. Using Stored Routines (Procedures and Functions)

Stored routines (procedures and functions) are supported in MySQL 5.6. A stored routine is a set of SQL statements that can be stored in the server. Once this has been done, clients don't need to keep reissuing the individual statements but can refer to the stored routine instead.

Stored routines require the `proc` table in the `mysql` database. This table is created during the MySQL 5.6 installation procedure. If you are upgrading to MySQL 5.6 from an earlier version, be sure to update your grant tables to make sure that the `proc` table exists. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

Stored routines can be particularly useful in certain situations:

- When multiple client applications are written in different languages or work on different platforms, but need to perform the same database operations.
- When security is paramount. Banks, for example, use stored procedures and functions for all common operations. This provides a consistent and secure environment, and routines can ensure that each operation is properly logged. In such a setup, applications and users would have no access to the database tables directly, but can only execute specific stored routines.

Stored routines can provide improved performance because less information needs to be sent between the server and the client. The tradeoff is that this does increase the load on the database server because more of the work is done on the server side and less is done on the client (application) side. Consider this if many client machines (such as Web servers) are serviced by only one or a few database servers.

Stored routines also enable you to have libraries of functions in the database server. This is a feature shared by modern application languages that enable such design internally (for example, by using classes). Using these client application language features is beneficial for the programmer even outside the scope of database use.

MySQL follows the SQL:2003 syntax for stored routines, which is also used by IBM's DB2. All syntax described here is supported and any limitations and extensions are documented where appropriate.

### Additional Resources

- You may find the [Stored Procedures User Forum](#) of use when working with stored procedures and functions.
- For answers to some commonly asked questions regarding stored routines in MySQL, see [Section B.4, “MySQL 5.6 FAQ: Stored Procedures and Functions”](#).
- There are some restrictions on the use of stored routines. See [Section E.1, “Restrictions on Stored Programs”](#).
- Binary logging for stored routines takes place as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

#### 18.2.1. Stored Routine Syntax

A stored routine is either a procedure or a function. Stored routines are created with the `CREATE PROCEDURE` and `CREATE FUNCTION` statements (see [Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#)). A procedure is invoked using a `CALL` statement (see [Section 13.2.1, “CALL Syntax”](#)), and can only pass back values using output variables. A function can be called from inside a statement just like any other function (that is, by invoking the function's name), and can return

a scalar value. The body of a stored routine can use compound statements (see [Section 13.6, “MySQL Compound-Statement Syntax”](#)).

Stored routines can be dropped with the `DROP PROCEDURE` and `DROP FUNCTION` statements (see [Section 13.1.21, “DROP PROCEDURE and DROP FUNCTION Syntax”](#)), and altered with the `ALTER PROCEDURE` and `ALTER FUNCTION` statements (see [Section 13.1.4, “ALTER PROCEDURE Syntax”](#)).

A stored procedure or function is associated with a particular database. This has several implications:

- When the routine is invoked, an implicit `USE db_name` is performed (and undone when the routine terminates). `USE` statements within stored routines are not permitted.
- You can qualify routine names with the database name. This can be used to refer to a routine that is not in the current database. For example, to invoke a stored procedure `p` or function `f` that is associated with the `test` database, you can say `CALL test.p()` or `test.f()`.
- When a database is dropped, all stored routines associated with it are dropped as well.

Stored functions cannot be recursive.

Recursion in stored procedures is permitted but disabled by default. To enable recursion, set the `max_sp_recursion_depth` [508] server system variable to a value greater than zero. Stored procedure recursion increases the demand on thread stack space. If you increase the value of `max_sp_recursion_depth` [508], it may be necessary to increase thread stack size by increasing the value of `thread_stack` [556] at server startup. See [Section 5.1.4, “Server System Variables”](#), for more information.

MySQL supports a very useful extension that enables the use of regular `SELECT` statements (that is, without using cursors or local variables) inside a stored procedure. The result set of such a query is simply sent directly to the client. Multiple `SELECT` statements generate multiple result sets, so the client must use a MySQL client library that supports multiple result sets. This means the client must use a client library from a version of MySQL at least as recent as 4.1. The client should also specify the `CLIENT_MULTI_RESULTS` option when it connects. For C programs, this can be done with the `mysql_real_connect()` C API function. See [Section 21.8.6.53, “mysql\\_real\\_connect\(\)”](#), and [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

## 18.2.2. Stored Routines and MySQL Privileges

The MySQL grant system takes stored routines into account as follows:

- The `CREATE ROUTINE` [754] privilege is needed to create stored routines.
- The `ALTER ROUTINE` [754] privilege is needed to alter or drop stored routines. This privilege is granted automatically to the creator of a routine if necessary, and dropped from the creator when the routine is dropped.
- The `EXECUTE` [754] privilege is required to execute stored routines. However, this privilege is granted automatically to the creator of a routine if necessary (and dropped from the creator when the routine is dropped). Also, the default `SQL SECURITY` characteristic for a routine is `DEFINER`, which enables users who have access to the database with which the routine is associated to execute the routine.
- If the `automatic_sp_privileges` [464] system variable is 0, the `EXECUTE` [754] and `ALTER ROUTINE` [754] privileges are not automatically granted to and dropped from the routine creator.
- The creator of a routine is the account used to execute the `CREATE` statement for it. This might not be the same as the account named as the `DEFINER` in the routine definition.

The server manipulates the `mysql.proc` table in response to statements that create, alter, or drop stored routines. It is not supported that the server will notice manual manipulation of this table.

## 18.2.3. Stored Routine Metadata

Metadata about stored routines can be obtained as follows:

- Query the `ROUTINES` table of the `INFORMATION_SCHEMA` database. See [Section 19.19](#), “[The INFORMATION\\_SCHEMA ROUTINES Table](#)”.
- Use the `SHOW CREATE PROCEDURE` and `SHOW CREATE FUNCTION` statements to see routine definitions. See [Section 13.7.5.11](#), “[SHOW CREATE PROCEDURE Syntax](#)”.
- Use the `SHOW PROCEDURE STATUS` and `SHOW FUNCTION STATUS` statements to see routine characteristics. See [Section 13.7.5.29](#), “[SHOW PROCEDURE STATUS Syntax](#)”.

### 18.2.4. Stored Procedures, Functions, Triggers, and `LAST_INSERT_ID()`

Within the body of a stored routine (procedure or function) or a trigger, the value of `LAST_INSERT_ID()` [1269] changes the same way as for statements executed outside the body of these kinds of objects (see [Section 12.14](#), “[Information Functions](#)”). The effect of a stored routine or trigger upon the value of `LAST_INSERT_ID()` [1269] that is seen by following statements depends on the kind of routine:

- If a stored procedure executes statements that change the value of `LAST_INSERT_ID()` [1269], the changed value is seen by statements that follow the procedure call.
- For stored functions and triggers that change the value, the value is restored when the function or trigger ends, so following statements do not see a changed value.

## 18.3. Using Triggers

A trigger is a named database object that is associated with a table, and that activates when a particular event occurs for the table. Some uses for triggers are to perform checks of values to be inserted into a table or to perform calculations on values involved in an update.

A trigger is defined to activate when an `INSERT`, `DELETE`, or `UPDATE` statement executes for the associated table. A trigger can be set to activate either before or after the triggering statement. For example, you can have a trigger activate before each row that is inserted into a table or after each row that is updated.



#### Important

MySQL triggers are activated by SQL statements *only*. They are not activated by changes in views, nor by changes to tables made by APIs that do not transmit SQL statements to the MySQL Server. This means that triggers are not activated by changes in `INFORMATION_SCHEMA` or `performance_schema` tables, because these tables are actually views.

To use triggers if you have upgraded to MySQL 5.6 from an older release that did not support triggers, you should upgrade your grant tables so that they contain the trigger-related privileges. See [Section 4.4.7](#), “[mysql\\_upgrade — Check and Upgrade MySQL Tables](#)”.

The following discussion describes the syntax for creating and dropping triggers, and shows some examples of how to use them.

### Additional Resources

- You may find the [Triggers User Forum](#) of use when working with views.
- For answers to some commonly asked questions regarding triggers in MySQL, see [Section B.5](#), “[MySQL 5.6 FAQ: Triggers](#)”.
- There are some restrictions on the use of triggers; see [Section E.1](#), “[Restrictions on Stored Programs](#)”.
- Binary logging for triggers takes place as described in [Section 18.7](#), “[Binary Logging of Stored Programs](#)”.

### 18.3.1. Trigger Syntax

To create a trigger or drop a trigger, use the `CREATE TRIGGER` or `DROP TRIGGER` statement. The syntax for these statements is described in [Section 13.1.15, “CREATE TRIGGER Syntax”](#), and [Section 13.1.24, “DROP TRIGGER Syntax”](#).

Here is a simple example that associates a trigger with a table for `INSERT` statements. The trigger acts as an accumulator, summing the values inserted into one of the columns of the table.

```
mysql> CREATE TABLE account (acct_num INT, amount DECIMAL(10,2));
Query OK, 0 rows affected (0.03 sec)

mysql> CREATE TRIGGER ins_sum BEFORE INSERT ON account
-> FOR EACH ROW SET @sum = @sum + NEW.amount;
Query OK, 0 rows affected (0.06 sec)
```

The `CREATE TRIGGER` statement creates a trigger named `ins_sum` that is associated with the `account` table. It also includes clauses that specify the trigger activation time, the triggering event, and what to do with the trigger activates:

- The keyword `BEFORE` indicates the trigger action time. In this case, the trigger should activate before each row inserted into the table. The other permissible keyword here is `AFTER`.
- The keyword `INSERT` indicates the event that activates the trigger. In the example, `INSERT` statements cause trigger activation. You can also create triggers for `DELETE` and `UPDATE` statements.
- The statement following `FOR EACH ROW` defines the statement to execute each time the trigger activates, which occurs once for each row affected by the triggering statement. In the example, the triggered statement is a simple `SET` that accumulates the values inserted into the `amount` column. The statement refers to the column as `NEW.amount` which means “the value of the `amount` column to be inserted into the new row.”

To use the trigger, set the accumulator variable to zero, execute an `INSERT` statement, and then see what value the variable has afterward:

```
mysql> SET @sum = 0;
mysql> INSERT INTO account VALUES(137,14.98),(141,1937.50),(97,-100.00);
mysql> SELECT @sum AS 'Total amount inserted';
+-----+
| Total amount inserted |
+-----+
| 1852.48                |
+-----+
```

In this case, the value of `@sum` after the `INSERT` statement has executed is `14.98 + 1937.50 - 100`, or `1852.48`.

To destroy the trigger, use a `DROP TRIGGER` statement. You must specify the schema name if the trigger is not in the default schema:

```
mysql> DROP TRIGGER test.ins_sum;
```

Triggers for a table are also dropped if you drop the table.

Trigger names exist in the schema namespace, meaning that all triggers must have unique names within a schema. Triggers in different schemas can have the same name.

In addition to the requirement that trigger names be unique for a schema, there are other limitations on the types of triggers you can create. In particular, you cannot have two triggers for a table that have the same activation time and activation event. For example, you cannot define two `BEFORE INSERT` triggers or two `AFTER UPDATE` triggers for a table. This should rarely be a significant limitation, because it is possible to define a trigger that executes multiple statements by using the `BEGIN ... END` compound statement construct after `FOR EACH ROW`. (An example appears later in this section.)

The `OLD` and `NEW` keywords enable you to access columns in the rows affected by a trigger. (`OLD` and `NEW` are not case sensitive.) In an `INSERT` trigger, only `NEW.col_name` can be used; there is no old row. In a `DELETE` trigger, only `OLD.col_name` can be used; there is no new row. In an `UPDATE` trigger, you can use `OLD.col_name` to refer to the columns of a row before it is updated and `NEW.col_name` to refer to the columns of the row after it is updated.

A column named with `OLD` is read only. You can refer to it (if you have the `SELECT` privilege), but not modify it. A column named with `NEW` can be referred to if you have the `SELECT` [755] privilege for it. In a `BEFORE` trigger, you can also change its value with `SET NEW.col_name = value` if you have the `UPDATE` [756] privilege for it. This means you can use a trigger to modify the values to be inserted into a new row or that are used to update a row.

In a `BEFORE` trigger, the `NEW` value for an `AUTO_INCREMENT` column is 0, not the automatically generated sequence number that will be generated when the new record actually is inserted.

`OLD` and `NEW` are MySQL extensions to triggers.

By using the `BEGIN ... END` construct, you can define a trigger that executes multiple statements. Within the `BEGIN` block, you also can use other syntax that is permitted within stored routines such as conditionals and loops. However, just as for stored routines, if you use the `mysql` program to define a trigger that executes multiple statements, it is necessary to redefine the `mysql` statement delimiter so that you can use the `;` statement delimiter within the trigger definition. The following example illustrates these points. It defines an `UPDATE` trigger that checks the new value to be used for updating each row, and modifies the value to be within the range from 0 to 100. This must be a `BEFORE` trigger because the value needs to be checked before it is used to update the row:

```
mysql> delimiter //
mysql> CREATE TRIGGER upd_check BEFORE UPDATE ON account
-> FOR EACH ROW
-> BEGIN
->     IF NEW.amount < 0 THEN
->         SET NEW.amount = 0;
->     ELSEIF NEW.amount > 100 THEN
->         SET NEW.amount = 100;
->     END IF;
-> END; //
mysql> delimiter ;
```

It can be easier to define a stored procedure separately and then invoke it from the trigger using a simple `CALL` statement. This is also advantageous if you want to invoke the same routine from within several triggers.

There are some limitations on what can appear in statements that a trigger executes when activated:

- The trigger cannot use the `CALL` statement to invoke stored procedures that return data to the client or that use dynamic SQL. (Stored procedures are permitted to return data to the trigger through `OUT` or `INOUT` parameters.)
- The trigger cannot use statements that explicitly or implicitly begin or end a transaction such as `START TRANSACTION`, `COMMIT`, or `ROLLBACK`.

MySQL handles errors during trigger execution as follows:

- If a `BEFORE` trigger fails, the operation on the corresponding row is not performed.
- A `BEFORE` trigger is activated by the *attempt* to insert or modify the row, regardless of whether the attempt subsequently succeeds.
- An `AFTER` trigger is executed only if the `BEFORE` trigger (if any) and the row operation both execute successfully.
- An error during either a `BEFORE` or `AFTER` trigger results in failure of the entire statement that caused trigger invocation.

- For transactional tables, failure of a statement should cause rollback of all changes performed by the statement. Failure of a trigger causes the statement to fail, so trigger failure also causes rollback. For nontransactional tables, such rollback cannot be done, so although the statement fails, any changes performed prior to the point of the error remain in effect.

### 18.3.2. Trigger Metadata

Metadata about triggers can be obtained as follows:

- Query the `TRIGGERS` table of the `INFORMATION_SCHEMA` database. See [Section 19.27, “The INFORMATION\\_SCHEMA TRIGGERS Table”](#).
- Use the `SHOW TRIGGERS` statement. See [Section 13.7.5.39, “SHOW TRIGGERS Syntax”](#).

## 18.4. Using the Event Scheduler

The *MySQL Event Scheduler* manages the scheduling and execution of events: Tasks that run according to schedule. The following discussion covers the Event Scheduler and is divided into the following sections:

- [Section 18.4.1, “Event Scheduler Overview”](#), provides an introduction to and conceptual overview of MySQL Events.
- [Section 18.4.3, “Event Syntax”](#), discusses the SQL statements for creating, altering, and dropping MySQL Events.
- [Section 18.4.4, “Event Metadata”](#), shows how to obtain information about events and how this information is stored by the MySQL Server.
- [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#), discusses the privileges required to work with events and the ramifications that events have with regard to privileges when executing.

Stored routines require the `event` table in the `mysql` database. This table is created during the MySQL 5.6 installation procedure. If you are upgrading to MySQL 5.6 from an earlier version, be sure to update your grant tables to make sure that the `event` table exists. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

### Additional Resources

- You may find the [MySQL Event Scheduler User Forum](#) of use when working with scheduled events.
- There are some restrictions on the use of events; see [Section E.1, “Restrictions on Stored Programs”](#).
- Binary logging for events takes place as described in [Section 18.7, “Binary Logging of Stored Programs”](#).

#### 18.4.1. Event Scheduler Overview

MySQL Events are tasks that run according to a schedule. Therefore, we sometimes refer to them as *scheduled* events. When you create an event, you are creating a named database object containing one or more SQL statements to be executed at one or more regular intervals, beginning and ending at a specific date and time. Conceptually, this is similar to the idea of the Unix `crontab` (also known as a “cron job”) or the Windows Task Scheduler.

Scheduled tasks of this type are also sometimes known as “temporal triggers”, implying that these are objects that are triggered by the passage of time. While this is essentially correct, we prefer to use the term *events* to avoid confusion with triggers of the type discussed in [Section 18.3, “Using Triggers”](#). Events should more specifically not be confused with “temporary triggers”. Whereas a trigger is a database object whose statements are executed in response to a specific type of event that occurs on a given table, a (scheduled) event is an object whose statements are executed in response to the passage of a specified time interval.



While there is no provision in the SQL Standard for event scheduling, there are precedents in other database systems, and you may notice some similarities between these implementations and that found in the MySQL Server.

MySQL Events have the following major features and properties:

- In MySQL 5.6, an event is uniquely identified by its name and the schema to which it is assigned.
- An event performs a specific action according to a schedule. This action consists of an SQL statement, which can be a compound statement in a `BEGIN . . . END` block if desired (see [Section 13.6, “MySQL Compound-Statement Syntax”](#)). An event's timing can be either *one-time* or *recurrent*. A one-time event executes one time only. A recurrent event repeats its action at a regular interval, and the schedule for a recurring event can be assigned a specific start day and time, end day and time, both, or neither. (By default, a recurring event's schedule begins as soon as it is created, and continues indefinitely, until it is disabled or dropped.)

If a repeating event does not terminate within its scheduling interval, the result may be multiple instances of the event executing simultaneously. If this is undesirable, you should institute a mechanism to prevent simultaneous instances. For example, you could use the `GET_LOCK( )` [\[1275\]](#) function, or row or table locking.

- Users can create, modify, and drop scheduled events using SQL statements intended for these purposes. Syntactically invalid event creation and modification statements fail with an appropriate error message. *A user may include statements in an event's action which require privileges that the user does not actually have.* The event creation or modification statement succeeds but the event's action fails. See [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#) for details.
- Many of the properties of an event can be set or modified using SQL statements. These properties include the event's name, timing, persistence (that is, whether it is preserved following the expiration of its schedule), status (enabled or disabled), action to be performed, and the schema to which it is assigned. See [Section 13.1.2, “ALTER EVENT Syntax”](#).

The default definer of an event is the user who created the event, unless the event has been altered, in which case the definer is the user who issued the last `ALTER EVENT` statement affecting that event. An event can be modified by any user having the `EVENT [754]` privilege on the database for which the event is defined. See [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#).

- An event's action statement may include most SQL statements permitted within stored routines. For restrictions, see [Section E.1, “Restrictions on Stored Programs”](#).

## 18.4.2. Event Scheduler Configuration

Events are executed by a special *event scheduler thread*; when we refer to the Event Scheduler, we actually refer to this thread. When running, the event scheduler thread and its current state can be seen by users having the `PROCESS [755]` privilege in the output of `SHOW PROCESSLIST`, as shown in the discussion that follows.

The global `event_scheduler [480]` system variable determines whether the Event Scheduler is enabled and running on the server. It has one of these 3 values, which affect event scheduling as described here:

- **OFF**: The Event Scheduler is stopped. The event scheduler thread does not run, is not shown in the output of `SHOW PROCESSLIST`, and no scheduled events are executed. **OFF** is the default value for `event_scheduler [480]`.

When the Event Scheduler is stopped (`event_scheduler [480]` is **OFF**), it can be started by setting the value of `event_scheduler [480]` to **ON**. (See next item.)

- **ON**: The Event Scheduler is started; the event scheduler thread runs and executes all scheduled events.



When the Event Scheduler is **ON**, the event scheduler thread is listed in the output of `SHOW PROCESSLIST` as a daemon process, and its state is represented as shown here:

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  Id: 1
  User: root
  Host: localhost
  db: NULL
Command: Query
  Time: 0
  State: NULL
  Info: show processlist
***** 2. row *****
  Id: 2
  User: event_scheduler
  Host: localhost
  db: NULL
Command: Daemon
  Time: 3
  State: Waiting for next activation
  Info: NULL
2 rows in set (0.00 sec)
```

Event scheduling can be stopped by setting the value of `event_scheduler` [480] to **OFF**.

- **DISABLED**: This value renders the Event Scheduler nonoperational. When the Event Scheduler is **DISABLED**, the event scheduler thread does not run (and so does not appear in the output of `SHOW PROCESSLIST`). In addition, the Event Scheduler state cannot be changed at runtime.

If the Event Scheduler status has not been set to **DISABLED**, `event_scheduler` [480] can be toggled between **ON** and **OFF** (using `SET`). It is also possible to use `0` for **OFF**, and `1` for **ON** when setting this variable. Thus, any of the following 4 statements can be used in the `mysql` client to turn on the Event Scheduler:

```
SET GLOBAL event_scheduler = ON;
SET @@global.event_scheduler = ON;
SET GLOBAL event_scheduler = 1;
SET @@global.event_scheduler = 1;
```

Similarly, any of these 4 statements can be used to turn off the Event Scheduler:

```
SET GLOBAL event_scheduler = OFF;
SET @@global.event_scheduler = OFF;
SET GLOBAL event_scheduler = 0;
SET @@global.event_scheduler = 0;
```

Although **ON** and **OFF** have numeric equivalents, the value displayed for `event_scheduler` [480] by `SELECT` or `SHOW VARIABLES` is always one of **OFF**, **ON**, or **DISABLED**. *DISABLED has no numeric equivalent*. For this reason, **ON** and **OFF** are usually preferred over `1` and `0` when setting this variable.

Note that attempting to set `event_scheduler` [480] without specifying it as a global variable causes an error:

```
mysql< SET @@event_scheduler = OFF;
ERROR 1229 (HY000): Variable 'event_scheduler' is a GLOBAL
variable and should be set with SET GLOBAL
```



### Important

It is possible to set the Event Scheduler to **DISABLED** only at server startup. If `event_scheduler` [480] is **ON** or **OFF**, you cannot set it to **DISABLED** at runtime. Also, if the Event Scheduler is set to **DISABLED** at startup, you cannot change the value of `event_scheduler` [480] at runtime.

To disable the event scheduler, use one of the following two methods:

- As a command-line option when starting the server:

```
--event-scheduler=DISABLED
```

- In the server configuration file (`my.cnf`, or `my.ini` on Windows systems), include the line where it will be read by the server (for example, in a `[mysqld]` section):

```
event_scheduler=DISABLED
```

To enable the Event Scheduler, restart the server without the `--event-scheduler=DISABLED` [422] command-line option, or after removing or commenting out the line containing `event_scheduler=DISABLED` [422] in the server configuration file, as appropriate. Alternatively, you can use `ON` (or `1`) or `OFF` (or `0`) in place of the `DISABLED` value when starting the server.



#### Note

You can issue event-manipulation statements when `event_scheduler` [480] is set to `DISABLED`. No warnings or errors are generated in such cases (provided that the statements are themselves valid). However, scheduled events cannot execute until this variable is set to `ON` (or `1`). Once this has been done, the event scheduler thread executes all events whose scheduling conditions are satisfied.

Starting the MySQL server with the `--skip-grant-tables` [439] option causes `event_scheduler` [480] to be set to `DISABLED`, overriding any other value set either on the command line or in the `my.cnf` or `my.ini` file (Bug #26807).

For SQL statements used to create, alter, and drop events, see [Section 18.4.3, “Event Syntax”](#).

MySQL 5.6 provides an `EVENTS` table in the `INFORMATION_SCHEMA` database. This table can be queried to obtain information about scheduled events which have been defined on the server. See [Section 18.4.4, “Event Metadata”](#), and [Section 19.7, “The INFORMATION\\_SCHEMA EVENTS Table”](#), for more information.

For information regarding event scheduling and the MySQL privilege system, see [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#).

### 18.4.3. Event Syntax

MySQL 5.6 provides several SQL statements for working with scheduled events:

- New events are defined using the `CREATE EVENT` statement. See [Section 13.1.9, “CREATE EVENT Syntax”](#).
- The definition of an existing event can be changed by means of the `ALTER EVENT` statement. See [Section 13.1.2, “ALTER EVENT Syntax”](#).
- When a scheduled event is no longer wanted or needed, it can be deleted from the server by its definer using the `DROP EVENT` statement. See [Section 13.1.18, “DROP EVENT Syntax”](#). Whether an event persists past the end of its schedule also depends on its `ON COMPLETION` clause, if it has one. See [Section 13.1.9, “CREATE EVENT Syntax”](#).

An event can be dropped by any user having the `EVENT` [754] privilege for the database on which the event is defined. See [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#).

### 18.4.4. Event Metadata

Metadata about events can be obtained as follows:

- Query the `event` table of the `mysql` database.
- Query the `EVENTS` table of the `INFORMATION_SCHEMA` database. See [Section 19.7, “The INFORMATION\\_SCHEMA EVENTS Table”](#).

- Use the `SHOW CREATE EVENT` statement. See [Section 13.7.5.9, “SHOW CREATE EVENT Syntax”](#).
- Use the `SHOW EVENTS` statement. See [Section 13.7.5.19, “SHOW EVENTS Syntax”](#).

### Event Scheduler Time Representation

Each session in MySQL has a session time zone (STZ). This is the session `time_zone` [557] value that is initialized from the server's global `time_zone` [557] value when the session begins but may be changed during the session.

The session time zone that is current when a `CREATE EVENT` or `ALTER EVENT` statement executes is used to interpret times specified in the event definition. This becomes the event time zone (ETZ); that is, the time zone that is used for event scheduling and is in effect within the event as it executes.

For representation of event information in the `mysql.event` table, the `execute_at`, `starts`, and `ends` times are converted to UTC and stored along with the event time zone. This enables event execution to proceed as defined regardless of any subsequent changes to the server time zone or daylight saving time effects. The `last_executed` time is also stored in UTC.

If you select information from `mysql.event`, the times just mentioned are retrieved as UTC values. These times can also be obtained by selecting from the `INFORMATION_SCHEMA.EVENTS` table or from `SHOW EVENTS`, but they are reported as ETZ values. Other times available from these sources indicate when an event was created or last altered; these are displayed as STZ values. The following table summarizes representation of event times.

Value	<code>mysql.event</code>	<code>INFORMATION_SCHEMA.EVENTS</code>	<code>SHOW EVENTS</code>
Execute at	UTC	ETZ	ETZ
Starts	UTC	ETZ	ETZ
Ends	UTC	ETZ	ETZ
Last executed	UTC	ETZ	n/a
Created	STZ	STZ	n/a
Last altered	STZ	STZ	n/a

## 18.4.5. Event Scheduler Status

The Event Scheduler writes information about event execution that terminates with an error or warning to the MySQL Server's error log. See [Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#) for an example.

To obtain information about the state of the Event Scheduler for debugging and troubleshooting purposes, run `mysqladmin debug` (see [Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”](#)); after running this command, the server's error log contains output relating to the Event Scheduler, similar to what is shown here:

```
Events status:
LLA = Last Locked At  LUA = Last Unlocked At
WOC = Waiting On Condition  DL = Data Locked

Event scheduler status:
State      : INITIALIZED
Thread id  : 0
LLA       : init_scheduler:313
LUA       : init_scheduler:318
WOC       : NO
Workers   : 0
Executed  : 0
Data locked: NO

Event queue status:
Element count : 1
Data locked   : NO
Attempting lock : NO
```

```
LLA      : init_queue:148
LUA      : init_queue:168
WOC      : NO
Next activation : 0000-00-00 00:00:00
```

In statements that occur as part of events executed by the Event Scheduler, diagnostics messages (not only errors, but also warnings) are written to the error log, and, on Windows, to the application event log. For frequently executed events, it is possible for this to result in many logged messages. For example, for `SELECT ... INTO var_list` statements, if the query returns no rows, a warning with error code 1329 occurs (`No data`), and the variable values remain unchanged. If the query returns multiple rows, error 1172 occurs (`Result consisted of more than one row`). For either condition, you can avoid having the warnings be logged by declaring a condition handler; see [Section 13.6.7.2, “DECLARE ... HANDLER Syntax”](#). For statements that may retrieve multiple rows, another strategy is to use `LIMIT 1` to limit the result set to a single row.

## 18.4.6. The Event Scheduler and MySQL Privileges

To enable or disable the execution of scheduled events, it is necessary to set the value of the global `event_scheduler` [\[480\]](#) system variable. This requires the `SUPER` [\[756\]](#) privilege.

The `EVENT` [\[754\]](#) privilege governs the creation, modification, and deletion of events. This privilege can be bestowed using `GRANT`. For example, this `GRANT` statement confers the `EVENT` [\[754\]](#) privilege for the schema named `myschema` on the user `jon@ghidora`:

```
GRANT EVENT ON myschema.* TO jon@ghidora;
```

(We assume that this user account already exists, and that we wish for it to remain unchanged otherwise.)

To grant this same user the `EVENT` [\[754\]](#) privilege on all schemas, use the following statement:

```
GRANT EVENT ON *.* TO jon@ghidora;
```

The `EVENT` [\[754\]](#) privilege has global or schema-level scope. Therefore, trying to grant it on a single table results in an error as shown:

```
mysql> GRANT EVENT ON myschema.mytable TO jon@ghidora;
ERROR 1144 (42000): Illegal GRANT/REVOKE command; please
consult the manual to see which privileges can be used
```

It is important to understand that an event is executed with the privileges of its definer, and that it cannot perform any actions for which its definer does not have the requisite privileges. For example, suppose that `jon@ghidora` has the `EVENT` [\[754\]](#) privilege for `myschema`. Suppose also that this user has the `SELECT` [\[755\]](#) privilege for `myschema`, but no other privileges for this schema. It is possible for `jon@ghidora` to create a new event such as this one:

```
CREATE EVENT e_store_ts
ON SCHEDULE
EVERY 10 SECOND
DO
INSERT INTO myschema.mytable VALUES (UNIX_TIMESTAMP());
```

The user waits for a minute or so, and then performs a `SELECT * FROM mytable;` query, expecting to see several new rows in the table. Instead, the table is empty. Since the user does not have the `INSERT` [\[755\]](#) privilege for the table in question, the event has no effect.

If you inspect the MySQL error log (`hostname.err`), you can see that the event is executing, but the action it is attempting to perform fails, as indicated by `RetCode=0`:

```
060209 22:39:44 [Note] EVEX EXECUTING event newdb.e [EXPR:10]
060209 22:39:44 [Note] EVEX EXECUTED event newdb.e [EXPR:10]. RetCode=0
060209 22:39:54 [Note] EVEX EXECUTING event newdb.e [EXPR:10]
060209 22:39:54 [Note] EVEX EXECUTED event newdb.e [EXPR:10]. RetCode=0
060209 22:40:04 [Note] EVEX EXECUTING event newdb.e [EXPR:10]
060209 22:40:04 [Note] EVEX EXECUTED event newdb.e [EXPR:10]. RetCode=0
```

Since this user very likely does not have access to the error log, it is possible to verify whether the event's action statement is valid by executing it directly:

```
mysql> INSERT INTO myschema.mytable VALUES (UNIX_TIMESTAMP());
ERROR 1142 (42000): INSERT command denied to user
'jon'@'ghidora' for table 'mytable'
```

Inspection of the `INFORMATION_SCHEMA.EVENTS` table shows that `e_store_ts` exists and is enabled, but its `LAST_EXECUTED` column is `NULL`:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.EVENTS
> WHERE EVENT_NAME='e_store_ts'
> AND EVENT_SCHEMA='myschema'\G
***** 1. row *****
EVENT_CATALOG: NULL
EVENT_SCHEMA: myschema
EVENT_NAME: e_store_ts
DEFINER: jon@ghidora
EVENT_BODY: SQL
EVENT_DEFINITION: INSERT INTO myschema.mytable VALUES (UNIX_TIMESTAMP())
EVENT_TYPE: RECURRING
EXECUTE_AT: NULL
INTERVAL_VALUE: 5
INTERVAL_FIELD: SECOND
SQL_MODE: NULL
STARTS: 0000-00-00 00:00:00
ENDS: 0000-00-00 00:00:00
STATUS: ENABLED
ON_COMPLETION: NOT PRESERVE
CREATED: 2006-02-09 22:36:06
LAST_ALTERED: 2006-02-09 22:36:06
LAST_EXECUTED: NULL
EVENT_COMMENT:
1 row in set (0.00 sec)
```

To rescind the `EVENT [754]` privilege, use the `REVOKE` statement. In this example, the `EVENT [754]` privilege on the schema `myschema` is removed from the `jon@ghidora` user account:

```
REVOKE EVENT ON myschema.* FROM jon@ghidora;
```



### Important

Revoking the `EVENT [754]` privilege from a user does not delete or disable any events that may have been created by that user.

An event is not migrated or dropped as a result of renaming or dropping the user who created it.

Suppose that the user `jon@ghidora` has been granted the `EVENT [754]` and `INSERT [755]` privileges on the `myschema` schema. This user then creates the following event:

```
CREATE EVENT e_insert
ON SCHEDULE
EVERY 7 SECOND
DO
INSERT INTO myschema.mytable;
```

After this event has been created, `root` revokes the `EVENT [754]` privilege for `jon@ghidora`. However, `e_insert` continues to execute, inserting a new row into `mytable` each seven seconds. The same would be true if `root` had issued either of these statements:

- `DROP USER jon@ghidora;`
- `RENAME USER jon@ghidora TO someotherguy@ghidora;`

You can verify that this is true by examining the `mysql.event` table (discussed later in this section) or the `INFORMATION_SCHEMA.EVENTS` table (see [Section 19.7, “The INFORMATION\\_SCHEMA EVENTS Table”](#)) before and after issuing a `DROP USER` or `RENAME USER` statement.

Event definitions are stored in the `mysql.event` table. To drop an event created by another user account, the MySQL `root` user (or another user with the necessary privileges) can delete rows from this table. For example, to remove the event `e_insert` shown previously, `root` can use the following statement:

```
DELETE FROM mysql.event
  WHERE db = 'myschema'
        AND definer = 'jon@ghidora'
        AND name = 'e_insert';
```

It is very important to match the event name, database schema name, and user account when deleting rows from the `mysql.event` table. This is because the same user can create different events of the same name in different schemas.

Users' `EVENT [754]` privileges are stored in the `Event_priv` columns of the `mysql.user` and `mysql.db` tables. In both cases, this column holds one of the values 'Y' or 'N'. 'N' is the default. `mysql.user.Event_priv` is set to 'Y' for a given user only if that user has the global `EVENT [754]` privilege (that is, if the privilege was bestowed using `GRANT EVENT ON *.*`). For a schema-level `EVENT [754]` privilege, `GRANT` creates a row in `mysql.db` and sets that row's `Db` column to the name of the schema, the `User` column to the name of the user, and the `Event_priv` column to 'Y'. There should never be any need to manipulate these tables directly, since the `GRANT EVENT` and `REVOKE EVENT` statements perform the required operations on them.

Five status variables provide counts of event-related operations (but *not* of statements executed by events; see [Section E.1, “Restrictions on Stored Programs”](#)). These are:

- `Com_create_event`: The number of `CREATE EVENT` statements executed since the last server restart.
- `Com_alter_event`: The number of `ALTER EVENT` statements executed since the last server restart.
- `Com_drop_event`: The number of `DROP EVENT` statements executed since the last server restart.
- `Com_show_create_event`: The number of `SHOW CREATE EVENT` statements executed since the last server restart.
- `Com_show_events`: The number of `SHOW EVENTS` statements executed since the last server restart.

You can view current values for all of these at one time by running the statement `SHOW STATUS LIKE '%event%';`.

## 18.5. Using Views

Views (including updatable views) are available in MySQL Server 5.6. Views are stored queries that when invoked produce a result set. A view acts as a virtual table.

To use views if you have upgraded to MySQL 5.6 from an older release that did not support views, you should upgrade your grant tables so that they contain the view-related privileges. See [Section 4.4.7, “mysql\\_upgrade — Check and Upgrade MySQL Tables”](#).

The following discussion describes the syntax for creating and dropping views, and shows some examples of how to use them.

### Additional Resources

- You may find the [Views User Forum](#) of use when working with views.
- For answers to some commonly asked questions regarding views in MySQL, see [Section B.6, “MySQL 5.6 FAQ: Views”](#).
- There are some restrictions on the use of views; see [Section E.5, “Restrictions on Views”](#).

## 18.5.1. View Syntax

The `CREATE VIEW` statement creates a new view (see [Section 13.1.16](#), “`CREATE VIEW Syntax`”). To alter the definition of a view or drop a view, use `ALTER VIEW` (see [Section 13.1.7](#), “`ALTER VIEW Syntax`”), or `DROP VIEW` (see [Section 13.1.25](#), “`DROP VIEW Syntax`”).

A view can be created from many kinds of `SELECT` statements. It can refer to base tables or other views. It can use joins, `UNION`, and subqueries. The `SELECT` need not even refer to any tables. The following example defines a view that selects two columns from another table, as well as an expression calculated from those columns:

```
mysql> CREATE TABLE t (qty INT, price INT);
mysql> INSERT INTO t VALUES(3, 50), (5, 60);
mysql> CREATE VIEW v AS SELECT qty, price, qty*price AS value FROM t;
mysql> SELECT * FROM v;
+-----+-----+-----+
| qty | price | value |
+-----+-----+-----+
| 3 | 50 | 150 |
| 5 | 60 | 300 |
+-----+-----+-----+
mysql> SELECT * FROM v WHERE qty = 5;
+-----+-----+-----+
| qty | price | value |
+-----+-----+-----+
| 5 | 60 | 300 |
+-----+-----+-----+
```

## 18.5.2. View Processing Algorithms

The optional `ALGORITHM` clause for `CREATE VIEW` or `ALTER VIEW` is a MySQL extension to standard SQL. It affects how MySQL processes the view. `ALGORITHM` takes three values: `MERGE`, `TEMPTABLE`, or `UNDEFINED`. The default algorithm is `UNDEFINED` if no `ALGORITHM` clause is present.

For `MERGE`, the text of a statement that refers to the view and the view definition are merged such that parts of the view definition replace corresponding parts of the statement.

For `TEMPTABLE`, the results from the view are retrieved into a temporary table, which then is used to execute the statement.

For `UNDEFINED`, MySQL chooses which algorithm to use. It prefers `MERGE` over `TEMPTABLE` if possible, because `MERGE` is usually more efficient and because a view cannot be updatable if a temporary table is used.

A reason to choose `TEMPTABLE` explicitly is that locks can be released on underlying tables after the temporary table has been created and before it is used to finish processing the statement. This might result in quicker lock release than the `MERGE` algorithm so that other clients that use the view are not blocked as long.

A view algorithm can be `UNDEFINED` for three reasons:

- No `ALGORITHM` clause is present in the `CREATE VIEW` statement.
- The `CREATE VIEW` statement has an explicit `ALGORITHM = UNDEFINED` clause.
- `ALGORITHM = MERGE` is specified for a view that can be processed only with a temporary table. In this case, MySQL generates a warning and sets the algorithm to `UNDEFINED`.

As mentioned earlier, `MERGE` is handled by merging corresponding parts of a view definition into the statement that refers to the view. The following examples briefly illustrate how the `MERGE` algorithm works. The examples assume that there is a view `v_merge` that has this definition:

```
CREATE ALGORITHM = MERGE VIEW v_merge (vc1, vc2) AS
SELECT c1, c2 FROM t WHERE c3 > 100;
```



Example 1: Suppose that we issue this statement:

```
SELECT * FROM v_merge;
```

MySQL handles the statement as follows:

- `v_merge` becomes `t`
- `*` becomes `vc1, vc2`, which corresponds to `c1, c2`
- The view `WHERE` clause is added

The resulting statement to be executed becomes:

```
SELECT c1, c2 FROM t WHERE c3 > 100;
```

Example 2: Suppose that we issue this statement:

```
SELECT * FROM v_merge WHERE vc1 < 100;
```

This statement is handled similarly to the previous one, except that `vc1 < 100` becomes `c1 < 100` and the view `WHERE` clause is added to the statement `WHERE` clause using an `AND` [1171] connective (and parentheses are added to make sure the parts of the clause are executed with correct precedence). The resulting statement to be executed becomes:

```
SELECT c1, c2 FROM t WHERE (c3 > 100) AND (c1 < 100);
```

Effectively, the statement to be executed has a `WHERE` clause of this form:

```
WHERE (select WHERE) AND (view WHERE)
```

If the `MERGE` algorithm cannot be used, a temporary table must be used instead. `MERGE` cannot be used if the view contains any of the following constructs:

- Aggregate functions (`SUM()` [1285], `MIN()` [1284], `MAX()` [1284], `COUNT()` [1283], and so forth)
- `DISTINCT`
- `GROUP BY`
- `HAVING`
- `LIMIT`
- `UNION` or `UNION ALL`
- Subquery in the select list
- Refers only to literal values (in this case, there is no underlying table)

### 18.5.3. Updatable and Insertable Views

Some views are updatable. That is, you can use them in statements such as `UPDATE`, `DELETE`, or `INSERT` to update the contents of the underlying table. For a view to be updatable, there must be a one-to-one relationship between the rows in the view and the rows in the underlying table. There are also certain other constructs that make a view nonupdatable. To be more specific, a view is not updatable if it contains any of the following:

- Aggregate functions (`SUM()` [1285], `MIN()` [1284], `MAX()` [1284], `COUNT()` [1283], and so forth)
- `DISTINCT`
- `GROUP BY`
- `HAVING`

- `UNION` or `UNION ALL`
- Subquery in the select list
- Certain joins (see additional join discussion later in this section)
- Nonupdatable view in the `FROM` clause
- A subquery in the `WHERE` clause that refers to a table in the `FROM` clause
- Refers only to literal values (in this case, there is no underlying table to update)
- Uses `ALGORITHM = TEMPTABLE` (use of a temporary table always makes a view nonupdatable)
- Multiple references to any column of a base table.

With respect to insertability (being updatable with `INSERT` statements), an updatable view is insertable if it also satisfies these additional requirements for the view columns:

- There must be no duplicate view column names.
- The view must contain all columns in the base table that do not have a default value.
- The view columns must be simple column references and not derived columns. A derived column is one that is not a simple column reference but is derived from an expression. These are examples of derived columns:

```
3.14159
col1 + 3
UPPER(col2)
col3 / col4
(subquery)
```

A view that has a mix of simple column references and derived columns is not insertable, but it can be updatable if you update only those columns that are not derived. Consider this view:

```
CREATE VIEW v AS SELECT col1, 1 AS col2 FROM t;
```

This view is not insertable because `col2` is derived from an expression. But it is updatable if the update does not try to update `col2`. This update is permissible:

```
UPDATE v SET col1 = 0;
```

This update is not permissible because it attempts to update a derived column:

```
UPDATE v SET col2 = 0;
```

It is sometimes possible for a multiple-table view to be updatable, assuming that it can be processed with the `MERGE` algorithm. For this to work, the view must use an inner join (not an outer join or a `UNION`). Also, only a single table in the view definition can be updated, so the `SET` clause must name only columns from one of the tables in the view. Views that use `UNION ALL` are not permitted even though they might be theoretically updatable, because the implementation uses temporary tables to process them.

For a multiple-table updatable view, `INSERT` can work if it inserts into a single table. `DELETE` is not supported.

`INSERT DELAYED` is not supported for views.

If a table contains an `AUTO_INCREMENT` column, inserting into an insertable view on the table that does not include the `AUTO_INCREMENT` column does not change the value of `LAST_INSERT_ID()` [1269], because the side effects of inserting default values into columns not part of the view should not be visible.

The `WITH CHECK OPTION` clause can be given for an updatable view to prevent inserts or updates to rows except those for which the `WHERE` clause in the `select_statement` is true.

In a `WITH CHECK OPTION` clause for an updatable view, the `LOCAL` and `CASCADED` keywords determine the scope of check testing when the view is defined in terms of another view. The `LOCAL` keyword restricts the `CHECK OPTION` only to the view being defined. `CASCADED` causes the checks for underlying views to be evaluated as well. When neither keyword is given, the default is `CASCADED`. Consider the definitions for the following table and set of views:

```
mysql> CREATE TABLE t1 (a INT);
mysql> CREATE VIEW v1 AS SELECT * FROM t1 WHERE a < 2
  -> WITH CHECK OPTION;
mysql> CREATE VIEW v2 AS SELECT * FROM v1 WHERE a > 0
  -> WITH LOCAL CHECK OPTION;
mysql> CREATE VIEW v3 AS SELECT * FROM v1 WHERE a > 0
  -> WITH CASCADED CHECK OPTION;
```

Here the `v2` and `v3` views are defined in terms of another view, `v1`. `v2` has a `LOCAL` check option, so inserts are tested only against the `v2` check. `v3` has a `CASCADED` check option, so inserts are tested not only against its own check, but against those of underlying views. The following statements illustrate these differences:

```
mysql> INSERT INTO v2 VALUES (2);
Query OK, 1 row affected (0.00 sec)
mysql> INSERT INTO v3 VALUES (2);
ERROR 1369 (HY000): CHECK OPTION failed 'test.v3'
```

MySQL sets a flag, called the view updatability flag, at `CREATE VIEW` time. The flag is set to `YES` (true) if `UPDATE` and `DELETE` (and similar operations) are legal for the view. Otherwise, the flag is set to `NO` (false). The `IS_UPDATABLE` column in the `INFORMATION_SCHEMA.VIEWS` table displays the status of this flag. It means that the server always knows whether a view is updatable. If the view is not updatable, statements such `UPDATE`, `DELETE`, and `INSERT` are illegal and will be rejected. (Note that even if a view is updatable, it might not be possible to insert into it, as described elsewhere in this section.)

The updatability of views may be affected by the value of the `updatable_views_with_limit` [562] system variable. See [Section 5.1.4, "Server System Variables"](#).

## 18.5.4. View Metadata

Metadata about views can be obtained as follows:

- Query the `VIEWS` table of the `INFORMATION_SCHEMA` database. See [Section 19.29, "The INFORMATION\\_SCHEMA VIEWS Table"](#).
- Use the `SHOW CREATE VIEW` statement. See [Section 13.7.5.14, "SHOW CREATE VIEW Syntax"](#).

## 18.6. Access Control for Stored Programs and Views

Stored programs and views are defined prior to use and, when referenced, execute within a security context that determines their privileges. These privileges are controlled by their `DEFINER` attribute, and, if there is one, their `SQL SECURITY` characteristic.

All stored programs (procedures, functions, triggers, and events) and views can have a `DEFINER` attribute that names a MySQL account. If the `DEFINER` attribute is omitted from a stored program or view definition, the default account is the user who creates the object.

In addition, stored routines (procedures and functions) and views can have a `SQL SECURITY` characteristic with a value of `DEFINER` or `INVOKER` to specify whether the object executes in definer or invoker context. If the `SQL SECURITY` characteristic is omitted, the default is definer context.

Triggers and events have no `SQL SECURITY` characteristic and always execute in definer context. The server invokes these objects automatically as necessary, so there is no invoking user.

Definer and invoker security contexts differ as follows:

- A stored program or view that executes in definer security context executes with the privileges of the account named by its `DEFINER` attribute. These privileges may be entirely different from those of the invoking user. The invoker must have appropriate privileges to reference the object (for example, `EXECUTE` [754] to call a stored procedure or `SELECT` [755] to select from a view), but when the object executes, the invoker's privileges are ignored and only the `DEFINER` account privileges matter. If this account has few privileges, the object is correspondingly limited in the operations it can perform. If the `DEFINER` account is highly privileged (such as a `root` account), the object can perform powerful operations *no matter who invokes it*.
- A stored routine or view that executes in invoker security context can perform only operations for which the invoker has privileges. The `DEFINER` attribute can be specified but has no effect for objects that execute in invoker context.

Consider the following stored procedure:

```
CREATE DEFINER = 'admin'@'localhost' PROCEDURE p1()
SQL SECURITY DEFINER
BEGIN
    UPDATE t1 SET counter = counter + 1;
END;
```

Any user who has the `EXECUTE` [754] privilege for `p1` can invoke it with a `CALL` statement. However, when `p1` executes, it does so in `DEFINER` security context and thus executes with the privileges of `'admin'@'localhost'`, the account named in the `DEFINER` attribute. This account must have the `EXECUTE` [754] privilege for `p1` as well as the `UPDATE` [756] privilege for the table `t1`. Otherwise, the procedure fails.

Now consider this stored procedure, which is identical to `p1` except that its `SQL SECURITY` characteristic is `INVOKER`:

```
CREATE DEFINER = 'admin'@'localhost' PROCEDURE p2()
SQL SECURITY INVOKER
BEGIN
    UPDATE t1 SET counter = counter + 1;
END;
```

`p2`, unlike `p1`, executes in `INVOKER` security context. The `DEFINER` attribute is irrelevant and `p2` executes with the privileges of the invoking user. `p2` fails if the invoker lacks the `EXECUTE` [754] privilege for `p2` or the `UPDATE` [756] privilege for the table `t1`.

MySQL uses the following rules to control which accounts a user can specify in an object `DEFINER` attribute:

- You can specify a `DEFINER` value other than your own account only if you have the `SUPER` [756] privilege.
- If you do not have the `SUPER` [756] privilege, the only legal user value is your own account, either specified literally or by using `CURRENT_USER` [1267]. You cannot set the definer to some other account.

To minimize the risk potential for stored program and view creation and use, follow these guidelines:

- For a stored routine or view, use `SQL SECURITY INVOKER` in the object definition when possible so that it can be used only by users with permissions appropriate for the operations performed by the object.
- If you create definer-context stored programs or views while using an account that has the `SUPER` [756] privilege, specify an explicit `DEFINER` attribute that names an account possessing only the privileges required for the operations performed by the object. Specify a highly privileged `DEFINER` account only when absolutely necessary.
- Administrators can prevent users from specifying highly privileged `DEFINER` accounts by not granting them the `SUPER` [756] privilege.

- Definer-context objects should be written keeping in mind that they may be able to access data for which the invoking user has no privileges. In some cases, you can prevent reference to these objects by not granting unauthorized users particular privileges:
  - A stored procedure or function cannot be referenced by a user who does not have the `EXECUTE` [754] privilege for it.
  - A view cannot be referenced by a user who does not have the appropriate privilege for it (`SELECT` [755] to select from it, `INSERT` [755] to insert into it, and so forth).

However, no such control exists for triggers because users do not reference them directly. A trigger always executes in `DEFINER` context and is activated by access to the table with which it is associated, even ordinary table accesses by users with no special privileges. If the `DEFINER` account is highly privileged, the trigger can perform sensitive or dangerous operations. This remains true if the `SUPER` [756] and `TRIGGER` [756] privileges needed to create the trigger are revoked from the account of the user who created it. Administrators should be especially careful about granting users that combination of privileges.

## 18.7. Binary Logging of Stored Programs

The binary log contains information about SQL statements that modify database contents. This information is stored in the form of “events” that describe the modifications. The binary log has two important purposes:

- For replication, the binary log is used on master replication servers as a record of the statements to be sent to slave servers. The master server sends the events contained in its binary log to its slaves, which execute those events to make the same data changes that were made on the master. See [Section 16.2, “Replication Implementation”](#).
- Certain data recovery operations require use of the binary log. After a backup file has been restored, the events in the binary log that were recorded after the backup was made are re-executed. These events bring databases up to date from the point of the backup. See [Section 7.3.2, “Using Backups for Recovery”](#).

However, there are certain binary logging issues that apply with respect to stored programs (stored procedures and functions, triggers, and events), if logging occurs at the statement level:

- In some cases, it is possible that a statement will affect different sets of rows on a master and a slave.
- Replicated statements executed on a slave are processed by the slave SQL thread, which has full privileges. It is possible for a procedure to follow different execution paths on master and slave servers, so a user can write a routine containing a dangerous statement that will execute only on the slave where it is processed by a thread that has full privileges.
- If a stored program that modifies data is nondeterministic, it is not repeatable. This can result in different data on a master and slave, or cause restored data to differ from the original data.

This section describes how MySQL 5.6 handles binary logging for stored programs. It states the current conditions that the implementation places on the use of stored programs, and what you can do to avoid problems. It also provides additional information about the reasons for these conditions.

In general, the issues described here result when binary logging occurs at the SQL statement level. If you use row-based binary logging, the log contains changes made to individual rows as a result of executing SQL statements. When routines or triggers execute, row changes are logged, not the statements that make the changes. For stored procedures, this means that the `CALL` statement is not logged. For stored functions, row changes made within the function are logged, not the function invocation. For triggers, row changes made by the trigger are logged. On the slave side, only the row changes are seen, not the stored program invocation. For general information about row-based logging, see [Section 16.1.2, “Replication Formats”](#).

Unless noted otherwise, the remarks here assume that you have enabled binary logging by starting the server with the `--log-bin` [2002] option. (See Section 5.2.4, “The Binary Log”.) If the binary log is not enabled, replication is not possible, nor is the binary log available for data recovery.

The current conditions on the use of stored functions in MySQL 5.6 can be summarized as follows. These conditions do not apply to stored procedures or Event Scheduler events and they do not apply unless binary logging is enabled.

- To create or alter a stored function, you must have the `SUPER` [756] privilege, in addition to the `CREATE ROUTINE` [754] or `ALTER ROUTINE` [754] privilege that is normally required. (Depending on the `DEFINER` value in the function definition, `SUPER` [756] might be required regardless of whether binary logging is enabled. See Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION Syntax`”.)
- When you create a stored function, you must declare either that it is deterministic or that it does not modify data. Otherwise, it may be unsafe for data recovery or replication.

By default, for a `CREATE FUNCTION` statement to be accepted, at least one of `DETERMINISTIC`, `NO SQL`, or `READS SQL DATA` must be specified explicitly. Otherwise an error occurs:

```
ERROR 1418 (HY000): This function has none of DETERMINISTIC, NO SQL,
or READS SQL DATA in its declaration and binary logging is enabled
(you *might* want to use the less safe log_bin_trust_function_creators
variable)
```

This function is deterministic (and does not modify data), so it is safe:

```
CREATE FUNCTION f1(i INT)
RETURNS INT
DETERMINISTIC
READS SQL DATA
BEGIN
    RETURN i;
END;
```

This function uses `UUID()` [1280], which is not deterministic, so the function also is not deterministic and is not safe:

```
CREATE FUNCTION f2()
RETURNS CHAR(36) CHARACTER SET utf8
BEGIN
    RETURN UUID();
END;
```

This function modifies data, so it may not be safe:

```
CREATE FUNCTION f3(p_id INT)
RETURNS INT
BEGIN
    UPDATE t SET modtime = NOW() WHERE id = p_id;
    RETURN ROW_COUNT();
END;
```

Assessment of the nature of a function is based on the “honesty” of the creator: MySQL does not check that a function declared `DETERMINISTIC` is free of statements that produce nondeterministic results.

- Although it is possible to create a deterministic stored function without specifying `DETERMINISTIC`, you cannot execute this function using statement-based binary logging. To execute such a function, you must use row-based or mixed binary logging. Alternatively, if you explicitly specify `DETERMINISTIC` in the function definition, you can use any kind of logging, including statement-based logging.
- To relax the preceding conditions on function creation (that you must have the `SUPER` [756] privilege and that a function must be declared deterministic or to not modify data), set the global

`log_bin_trust_function_creators` [497] system variable to 1. By default, this variable has a value of 0, but you can change it like this:

```
mysql> SET GLOBAL log_bin_trust_function_creators = 1;
```

You can also set this variable by using the `--log-bin-trust-function-creators=1` [2003] option when starting the server.

If binary logging is not enabled, `log_bin_trust_function_creators` [497] does not apply. `SUPER` [756] is not required for function creation unless, as described previously, the `DEFINER` value in the function definition requires it.

- For information about built-in functions that may be unsafe for replication (and thus cause stored functions that use them to be unsafe as well), see [Section 16.4.1, “Replication Features and Issues”](#).

Triggers are similar to stored functions, so the preceding remarks regarding functions also apply to triggers with the following exception: `CREATE TRIGGER` does not have an optional `DETERMINISTIC` characteristic, so triggers are assumed to be always deterministic. However, this assumption might in some cases be invalid. For example, the `UUID()` [1280] function is nondeterministic (and does not replicate). You should be careful about using such functions in triggers.

Triggers can update tables, so error messages similar to those for stored functions occur with `CREATE TRIGGER` if you do not have the required privileges. On the slave side, the slave uses the trigger `DEFINER` attribute to determine which user is considered to be the creator of the trigger.

The rest of this section provides additional detail about the logging implementation and its implications. You need not read it unless you are interested in the background on the rationale for the current logging-related conditions on stored routine use. This discussion applies only for statement-based logging, and not for row-based logging, with the exception of the first item: `CREATE` and `DROP` statements are logged as statements regardless of the logging mode.

- The server writes `CREATE EVENT`, `CREATE PROCEDURE`, `CREATE FUNCTION`, `ALTER EVENT`, `ALTER PROCEDURE`, `ALTER FUNCTION`, `DROP EVENT`, `DROP PROCEDURE`, and `DROP FUNCTION` statements to the binary log.
- A stored function invocation is logged as a `SELECT` statement if the function changes data and occurs within a statement that would not otherwise be logged. This prevents nonreplication of data changes that result from use of stored functions in nonlogged statements. For example, `SELECT` statements are not written to the binary log, but a `SELECT` might invoke a stored function that makes changes. To handle this, a `SELECT func_name()` statement is written to the binary log when the given function makes a change. Suppose that the following statements are executed on the master:

```
CREATE FUNCTION f1(a INT) RETURNS INT
BEGIN
  IF (a < 3) THEN
    INSERT INTO t2 VALUES (a);
  END IF;
  RETURN 0;
END;

CREATE TABLE t1 (a INT);
INSERT INTO t1 VALUES (1),(2),(3);

SELECT f1(a) FROM t1;
```

When the `SELECT` statement executes, the function `f1()` is invoked three times. Two of those invocations insert a row, and MySQL logs a `SELECT` statement for each of them. That is, MySQL writes the following statements to the binary log:

```
SELECT f1(1);
SELECT f1(2);
```

The server also logs a `SELECT` statement for a stored function invocation when the function invokes a stored procedure that causes an error. In this case, the server writes the `SELECT` statement to the



log along with the expected error code. On the slave, if the same error occurs, that is the expected result and replication continues. Otherwise, replication stops.

- Logging stored function invocations rather than the statements executed by a function has a security implication for replication, which arises from two factors:
  - It is possible for a function to follow different execution paths on master and slave servers.
  - Statements executed on a slave are processed by the slave SQL thread which has full privileges.

The implication is that although a user must have the `CREATE ROUTINE [754]` privilege to create a function, the user can write a function containing a dangerous statement that will execute only on the slave where it is processed by a thread that has full privileges. For example, if the master and slave servers have server ID values of 1 and 2, respectively, a user on the master server could create and invoke an unsafe function `unsafe_func()` as follows:

```
mysql> delimiter //
mysql> CREATE FUNCTION unsafe_func () RETURNS INT
-> BEGIN
->   IF @@server_id=2 THEN dangerous_statement; END IF;
->   RETURN 1;
-> END;
-> //
mysql> delimiter ;
mysql> INSERT INTO t VALUES(unsafe_func());
```

The `CREATE FUNCTION` and `INSERT` statements are written to the binary log, so the slave will execute them. Because the slave SQL thread has full privileges, it will execute the dangerous statement. Thus, the function invocation has different effects on the master and slave and is not replication-safe.

To guard against this danger for servers that have binary logging enabled, stored function creators must have the `SUPER [756]` privilege, in addition to the usual `CREATE ROUTINE [754]` privilege that is required. Similarly, to use `ALTER FUNCTION`, you must have the `SUPER [756]` privilege in addition to the `ALTER ROUTINE [754]` privilege. Without the `SUPER [756]` privilege, an error will occur:

```
ERROR 1419 (HY000): You do not have the SUPER privilege and
binary logging is enabled (you *might* want to use the less safe
log_bin_trust_function_creators variable)
```

If you do not want to require function creators to have the `SUPER [756]` privilege (for example, if all users with the `CREATE ROUTINE [754]` privilege on your system are experienced application developers), set the global `log_bin_trust_function_creators [497]` system variable to 1. You can also set this variable by using the `--log-bin-trust-function-creators=1 [2003]` option when starting the server. If binary logging is not enabled, `log_bin_trust_function_creators [497]` does not apply. `SUPER [756]` is not required for function creation unless, as described previously, the `DEFINER` value in the function definition requires it.

- If a function that performs updates is nondeterministic, it is not repeatable. This can have two undesirable effects:
  - It will make a slave different from the master.
  - Restored data will be different from the original data.

To deal with these problems, MySQL enforces the following requirement: On a master server, creation and alteration of a function is refused unless you declare the function to be deterministic or to not modify data. Two sets of function characteristics apply here:

- The `DETERMINISTIC` and `NOT DETERMINISTIC` characteristics indicate whether a function always produces the same result for given inputs. The default is `NOT DETERMINISTIC` if

neither characteristic is given. To declare that a function is deterministic, you must specify `DETERMINISTIC` explicitly.

- The `CONTAINS SQL`, `NO SQL`, `READS SQL DATA`, and `MODIFIES SQL DATA` characteristics provide information about whether the function reads or writes data. Either `NO SQL` or `READS SQL DATA` indicates that a function does not change data, but you must specify one of these explicitly because the default is `CONTAINS SQL` if no characteristic is given.

By default, for a `CREATE FUNCTION` statement to be accepted, at least one of `DETERMINISTIC`, `NO SQL`, or `READS SQL DATA` must be specified explicitly. Otherwise an error occurs:

```
ERROR 1418 (HY000): This function has none of DETERMINISTIC, NO SQL,
or READS SQL DATA in its declaration and binary logging is enabled
(you *might* want to use the less safe log_bin_trust_function_creators
variable)
```

If you set `log_bin_trust_function_creators` [497] to 1, the requirement that functions be deterministic or not modify data is dropped.

- Stored procedure calls are logged at the statement level rather than at the `CALL` level. That is, the server does not log the `CALL` statement, it logs those statements within the procedure that actually execute. As a result, the same changes that occur on the master will be observed on slave servers. This prevents problems that could result from a procedure having different execution paths on different machines.

In general, statements executed within a stored procedure are written to the binary log using the same rules that would apply were the statements to be executed in standalone fashion. Some special care is taken when logging procedure statements because statement execution within procedures is not quite the same as in nonprocedure context:

- A statement to be logged might contain references to local procedure variables. These variables do not exist outside of stored procedure context, so a statement that refers to such a variable cannot be logged literally. Instead, each reference to a local variable is replaced by this construct for logging purposes:

```
NAME_CONST(var_name, var_value)
```

`var_name` is the local variable name, and `var_value` is a constant indicating the value that the variable has at the time the statement is logged. `NAME_CONST()` [1279] has a value of `var_value`, and a "name" of `var_name`. Thus, if you invoke this function directly, you get a result like this:

```
mysql> SELECT NAME_CONST('myname', 14);
+-----+
| myname |
+-----+
|      14 |
+-----+
```

`NAME_CONST()` [1279] enables a logged standalone statement to be executed on a slave with the same effect as the original statement that was executed on the master within a stored procedure.

The use of `NAME_CONST()` [1279] can result in a problem for `CREATE TABLE ... SELECT` statements when the source column expressions refer to local variables. Converting these references to `NAME_CONST()` [1279] expressions can result in column names that are different on the master and slave servers, or names that are too long to be legal column identifiers. A workaround is to supply aliases for columns that refer to local variables. Consider this statement when `myvar` has a value of 1:

```
CREATE TABLE t1 SELECT myvar;
```

That will be rewritten as follows:

```
CREATE TABLE t1 SELECT NAME_CONST(myvar, 1);
```

To ensure that the master and slave tables have the same column names, write the statement like this:

```
CREATE TABLE t1 SELECT myvar AS myvar;
```

The rewritten statement becomes:

```
CREATE TABLE t1 SELECT NAME_CONST(myvar, 1) AS myvar;
```

- A statement to be logged might contain references to user-defined variables. To handle this, MySQL writes a `SET` statement to the binary log to make sure that the variable exists on the slave with the same value as on the master. For example, if a statement refers to a variable `@my_var`, that statement will be preceded in the binary log by the following statement, where *value* is the value of `@my_var` on the master:

```
SET @my_var = value;
```

- Procedure calls can occur within a committed or rolled-back transaction. Transactional context is accounted for so that the transactional aspects of procedure execution are replicated correctly. That is, the server logs those statements within the procedure that actually execute and modify data, and also logs `BEGIN`, `COMMIT`, and `ROLLBACK` statements as necessary. For example, if a procedure updates only transactional tables and is executed within a transaction that is rolled back, those updates are not logged. If the procedure occurs within a committed transaction, `BEGIN` and `COMMIT` statements are logged with the updates. For a procedure that executes within a rolled-back transaction, its statements are logged using the same rules that would apply if the statements were executed in standalone fashion:
  - Updates to transactional tables are not logged.
  - Updates to nontransactional tables are logged because rollback does not cancel them.
  - Updates to a mix of transactional and nontransactional tables are logged surrounded by `BEGIN` and `ROLLBACK` so that slaves will make the same changes and rollbacks as on the master.
- A stored procedure call is *not* written to the binary log at the statement level if the procedure is invoked from within a stored function. In that case, the only thing logged is the statement that invokes the function (if it occurs within a statement that is logged) or a `DO` statement (if it occurs within a statement that is not logged). For this reason, care should be exercised in the use of stored functions that invoke a procedure, even if the procedure is otherwise safe in itself.

---

## Chapter 19. INFORMATION\_SCHEMA Tables

### Table of Contents

19.1. The INFORMATION_SCHEMA CHARACTER_SETS Table .....	2180
19.2. The INFORMATION_SCHEMA COLLATIONS Table .....	2180
19.3. The INFORMATION_SCHEMA COLLATION_CHARACTER_SET_APPLICABILITY Table .....	2181
19.4. The INFORMATION_SCHEMA COLUMNS Table .....	2181
19.5. The INFORMATION_SCHEMA COLUMN_PRIVILEGES Table .....	2182
19.6. The INFORMATION_SCHEMA ENGINES Table .....	2183
19.7. The INFORMATION_SCHEMA EVENTS Table .....	2183
19.8. The INFORMATION_SCHEMA FILES Table .....	2186
19.9. The INFORMATION_SCHEMA GLOBAL_STATUS and SESSION_STATUS Tables .....	2188
19.10. The INFORMATION_SCHEMA GLOBAL_VARIABLES and SESSION_VARIABLES Tables ...	2188
19.11. The INFORMATION_SCHEMA KEY_COLUMN_USAGE Table .....	2188
19.12. The INFORMATION_SCHEMA OPTIMIZER_TRACE Table .....	2189
19.13. The INFORMATION_SCHEMA PARAMETERS Table .....	2189
19.14. The INFORMATION_SCHEMA PARTITIONS Table .....	2190
19.15. The INFORMATION_SCHEMA PLUGINS Table .....	2193
19.16. The INFORMATION_SCHEMA PROCESSLIST Table .....	2194
19.17. The INFORMATION_SCHEMA PROFILING Table .....	2195
19.18. The INFORMATION_SCHEMA REFERENTIAL_CONSTRAINTS Table .....	2196
19.19. The INFORMATION_SCHEMA ROUTINES Table .....	2196
19.20. The INFORMATION_SCHEMA SCHEMATA Table .....	2198
19.21. The INFORMATION_SCHEMA SCHEMA_PRIVILEGES Table .....	2198
19.22. The INFORMATION_SCHEMA STATISTICS Table .....	2198
19.23. The INFORMATION_SCHEMA TABLES Table .....	2199
19.24. The INFORMATION_SCHEMA TABLESPACES Table .....	2200
19.25. The INFORMATION_SCHEMA TABLE_CONSTRAINTS Table .....	2201
19.26. The INFORMATION_SCHEMA TABLE_PRIVILEGES Table .....	2201
19.27. The INFORMATION_SCHEMA TRIGGERS Table .....	2202
19.28. The INFORMATION_SCHEMA USER_PRIVILEGES Table .....	2203
19.29. The INFORMATION_SCHEMA VIEWS Table .....	2204
19.30. INFORMATION_SCHEMA Tables for InnoDB .....	2205
19.30.1. The INFORMATION_SCHEMA INNODB_CMP and INNODB_CMP_RESET Tables .....	2205
19.30.2. The INFORMATION_SCHEMA INNODB_CMP_PER_INDEX and INNODB_CMP_PER_INDEX_RESET Tables .....	2206
19.30.3. The INFORMATION_SCHEMA INNODB_CMPMEM and INNODB_CMPMEM_RESET Tables .....	2207
19.30.4. The INFORMATION_SCHEMA INNODB_TRX Table .....	2207
19.30.5. The INFORMATION_SCHEMA INNODB_LOCKS Table .....	2209
19.30.6. The INFORMATION_SCHEMA INNODB_LOCK_WAITS Table .....	2210
19.30.7. The INFORMATION_SCHEMA INNODB_SYS_TABLES Table .....	2210
19.30.8. The INFORMATION_SCHEMA INNODB_SYS_INDEXES Table .....	2211
19.30.9. The INFORMATION_SCHEMA INNODB_SYS_COLUMNS Table .....	2212
19.30.10. The INFORMATION_SCHEMA INNODB_SYS_FIELDS Table .....	2212
19.30.11. The INFORMATION_SCHEMA INNODB_SYS_FOREIGN Table .....	2213
19.30.12. The INFORMATION_SCHEMA INNODB_SYS_FOREIGN_COLS Table .....	2213
19.30.13. The INFORMATION_SCHEMA INNODB_SYS_TABLESTATS View .....	2214
19.30.14. The INFORMATION_SCHEMA INNODB_BUFFER_PAGE Table .....	2214
19.30.15. The INFORMATION_SCHEMA INNODB_BUFFER_PAGE_LRU Table .....	2215
19.30.16. The INFORMATION_SCHEMA INNODB_BUFFER_POOL_STATS Table .....	2216
19.30.17. The INFORMATION_SCHEMA INNODB_METRICS Table .....	2217
19.30.18. The INFORMATION_SCHEMA INNODB_FT_CONFIG Table .....	2222
19.30.19. The INFORMATION_SCHEMA INNODB_FT_DEFAULT_STOPWORD Table .....	2223
19.30.20. The INFORMATION_SCHEMA INNODB_FT_INDEX_TABLE Table .....	2223

19.30.21.	The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_INDEX_CACHE</code> Table .....	2224
19.30.22.	The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_DELETED</code> Table .....	2224
19.30.23.	The <code>INFORMATION_SCHEMA</code> <code>INNODB_FT_BEING_DELETED</code> Table .....	2225
19.31.	Thread Pool <code>INFORMATION_SCHEMA</code> Tables .....	2225
19.31.1.	The <code>INFORMATION_SCHEMA</code> <code>TP_THREAD_STATE</code> Table .....	2226
19.31.2.	The <code>INFORMATION_SCHEMA</code> <code>TP_THREAD_GROUP_STATE</code> Table .....	2226
19.31.3.	The <code>INFORMATION_SCHEMA</code> <code>TP_THREAD_GROUP_STATS</code> Table .....	2228
19.32.	Extensions to <code>SHOW</code> Statements .....	2229

`INFORMATION_SCHEMA` provides access to database *metadata*, information about the MySQL server such as the name of a database or table, the data type of a column, or access privileges. Other terms that are sometimes used for this information are *data dictionary* and *system catalog*.

## Usage Notes for the `INFORMATION_SCHEMA` Database

`INFORMATION_SCHEMA` is a database within each MySQL instance, the place that stores information about all the other databases that the MySQL server maintains. The `INFORMATION_SCHEMA` database contains several read-only tables. They are actually views, not base tables, so there are no files associated with them, and you cannot set triggers on them. Also, there is no database directory with that name.

Although you can select `INFORMATION_SCHEMA` as the default database with a `USE` statement, you can only read the contents of tables, not perform `INSERT`, `UPDATE`, or `DELETE` operations on them.

## Example

Here is an example of a statement that retrieves information from `INFORMATION_SCHEMA`:

```
mysql> SELECT table_name, table_type, engine
-> FROM information_schema.tables
-> WHERE table_schema = 'db5'
-> ORDER BY table_name;
```

table_name	table_type	engine
fk	BASE TABLE	InnoDB
fk2	BASE TABLE	InnoDB
goto	BASE TABLE	MyISAM
into	BASE TABLE	MyISAM
k	BASE TABLE	MyISAM
kurs	BASE TABLE	MyISAM
loop	BASE TABLE	MyISAM
pk	BASE TABLE	InnoDB
t	BASE TABLE	MyISAM
t2	BASE TABLE	MyISAM
t3	BASE TABLE	MyISAM
t7	BASE TABLE	MyISAM
tables	BASE TABLE	MyISAM
v	VIEW	NULL
v2	VIEW	NULL
v3	VIEW	NULL
v56	VIEW	NULL

17 rows in set (0.01 sec)

Explanation: The statement requests a list of all the tables in database `db5`, showing just three pieces of information: the name of the table, its type, and its storage engine.

## Character Set Considerations

The definition for character columns (for example, `TABLES.TABLE_NAME`) is generally `VARCHAR(N) CHARACTER SET utf8` where `N` is at least 64. MySQL uses the default collation for this character set (`utf8_general_ci`) for all searches, sorts, comparisons, and other string operations on such columns.

Because some MySQL objects are represented as files, searches in `INFORMATION_SCHEMA` string columns can be affected by file system case sensitivity. For more information, see [Section 10.1.7.9, “Collation and INFORMATION\\_SCHEMA Searches”](#).

## INFORMATION\_SCHEMA as Alternative to SHOW Statements

The `SELECT ... FROM INFORMATION_SCHEMA` statement is intended as a more consistent way to provide access to the information provided by the various `SHOW` statements that MySQL supports (`SHOW DATABASES`, `SHOW TABLES`, and so forth). Using `SELECT` has these advantages, compared to `SHOW`:

- It conforms to Codd's rules, because all access is done on tables.
- You can use the familiar syntax of the `SELECT` statement, and only need to learn some table and column names.
- The implementor need not worry about adding keywords.
- You can filter, sort, concatenate, and transform the results from `INFORMATION_SCHEMA` queries into whatever format your application needs, such as a data structure or a text representation to parse.
- This technique is more interoperable with other database systems. For example, Oracle Database users are familiar with querying tables in the Oracle data dictionary.

Because `SHOW` is familiar and widely used, the `SHOW` statements remain as an alternative. In fact, along with the implementation of `INFORMATION_SCHEMA`, there are enhancements to `SHOW` as described in [Section 19.32, “Extensions to SHOW Statements”](#).

## Privileges

Each MySQL user has the right to access these tables, but can see only the rows in the tables that correspond to objects for which the user has the proper access privileges. In some cases (for example, the `ROUTINE_DEFINITION` column in the `INFORMATION_SCHEMA.ROUTINES` table), users who have insufficient privileges see `NULL`. These restrictions do not apply for `InnoDB` tables; you can see them with only the `PROCESS [755]` privilege.

The same privileges apply to selecting information from `INFORMATION_SCHEMA` and viewing the same information through `SHOW` statements. In either case, you must have some privilege on an object to see information about it.

## Performance Considerations

`INFORMATION_SCHEMA` queries that search for information from more than one database might take a long time and impact performance. To check the efficiency of a query, you can use `EXPLAIN`. For information about using `EXPLAIN` output to tune `INFORMATION_SCHEMA` queries, see [Section 8.2.4, “Optimizing INFORMATION\\_SCHEMA Queries”](#).

## Standards Considerations

The implementation for the `INFORMATION_SCHEMA` table structures in MySQL follows the ANSI/ISO SQL:2003 standard Part 11 *Schemata*. Our intent is approximate compliance with SQL:2003 core feature F021 *Basic information schema*.

Users of SQL Server 2000 (which also follows the standard) may notice a strong similarity. However, MySQL has omitted many columns that are not relevant for our implementation, and added columns that are MySQL-specific. One such column is the `ENGINE` column in the `INFORMATION_SCHEMA.TABLES` table.

Although other DBMSs use a variety of names, like `syscat` or `system`, the standard name is `INFORMATION_SCHEMA`.



To avoid using any name that is reserved in the standard or in DB2, SQL Server, or Oracle, we changed the names of some columns marked “MySQL extension”. (For example, we changed `COLLATION` to `TABLE_COLLATION` in the `TABLES` table.) See the list of reserved words near the end of this article: [http://web.archive.org/web/20070409075643rn\\_1/www.dbazine.com/db2/db2-disarticles/gulutzan5](http://web.archive.org/web/20070409075643rn_1/www.dbazine.com/db2/db2-disarticles/gulutzan5).

## Conventions in the `INFORMATION_SCHEMA` Reference Sections

The following sections describe each of the tables and columns in `INFORMATION_SCHEMA`. For each column, there are three pieces of information:

- “`INFORMATION_SCHEMA` Name” indicates the name for the column in the `INFORMATION_SCHEMA` table. This corresponds to the standard SQL name unless the “Remarks” field says “MySQL extension.”
- “`SHOW` Name” indicates the equivalent field name in the closest `SHOW` statement, if there is one.
- “Remarks” provides additional information where applicable. If this field is `NULL`, it means that the value of the column is always `NULL`. If this field says “MySQL extension,” the column is a MySQL extension to standard SQL.

Many sections indicate what `SHOW` statement is equivalent to a `SELECT` that retrieves information from `INFORMATION_SCHEMA`. For `SHOW` statements that display information for the default database if you omit a `FROM db_name` clause, you can often select information for the default database by adding an `AND TABLE_SCHEMA = SCHEMA()` condition to the `WHERE` clause of a query that retrieves information from an `INFORMATION_SCHEMA` table.

For information about `INFORMATION_SCHEMA` tables specific to the `InnoDB` storage engine, see [Section 19.30, “`INFORMATION\_SCHEMA` Tables for `InnoDB`”](#). For information about `INFORMATION_SCHEMA` tables specific to the thread pool plugin, see [Section 19.31, “Thread Pool `INFORMATION\_SCHEMA` Tables”](#).

For answers to questions that are often asked concerning the `INFORMATION_SCHEMA` database, see [Section B.7, “MySQL 5.6 FAQ: `INFORMATION\_SCHEMA`”](#).

### 19.1. The `INFORMATION_SCHEMA CHARACTER_SETS` Table

The `CHARACTER_SETS` table provides information about available character sets.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CHARACTER_SET_NAME</code>	Charset	
<code>DEFAULT_COLLATE_NAME</code>	Default collation	
<code>DESCRIPTION</code>	Description	MySQL extension
<code>MAXLEN</code>	Maxlen	MySQL extension

The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.CHARACTER_SETS
  [WHERE CHARACTER_SET_NAME LIKE 'wild']

SHOW CHARACTER SET
  [LIKE 'wild']
```

### 19.2. The `INFORMATION_SCHEMA COLLATIONS` Table

The `COLLATIONS` table provides information about collations for each character set.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>COLLATION_NAME</code>	Collation	



<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CHARACTER_SET_NAME</code>	Charset	MySQL extension
<code>ID</code>	Id	MySQL extension
<code>IS_DEFAULT</code>	Default	MySQL extension
<code>IS_COMPILED</code>	Compiled	MySQL extension
<code>SORTLEN</code>	Sortlen	MySQL extension

- `COLLATION_NAME` is the collation name.
- `CHARACTER_SET_NAME` is the name of the character set with which the collation is associated.
- `ID` is the collation ID.
- `IS_DEFAULT` indicates whether the collation is the default for its character set.
- `IS_COMPILED` indicates whether the character set is compiled into the server.
- `SORTLEN` is related to the amount of memory required to sort strings expressed in the character set.

Collation information is also available from the `SHOW COLLATION` statement. The following statements are equivalent:

```
SELECT COLLATION_NAME FROM INFORMATION_SCHEMA.COLLATIONS
[WHERE COLLATION_NAME LIKE 'wild']

SHOW COLLATION
[LIKE 'wild']
```

### 19.3. The `INFORMATION_SCHEMA COLLATION_CHARACTER_SET_APPLICABILITY` Table

The `COLLATION_CHARACTER_SET_APPLICABILITY` table indicates what character set is applicable for what collation. The columns are equivalent to the first two display fields that we get from `SHOW COLLATION`.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>COLLATION_NAME</code>	Collation	
<code>CHARACTER_SET_NAME</code>	Charset	

### 19.4. The `INFORMATION_SCHEMA COLUMNS` Table

The `COLUMNS` table provides information about columns in tables.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>TABLE_CATALOG</code>		def
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>COLUMN_NAME</code>	Field	
<code>ORDINAL_POSITION</code>		see notes
<code>COLUMN_DEFAULT</code>	Default	
<code>IS_NULLABLE</code>	Null	
<code>DATA_TYPE</code>	Type	
<code>CHARACTER_MAXIMUM_LENGTH</code>	Type	
<code>CHARACTER_OCTET_LENGTH</code>		

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>NUMERIC_PRECISION</code>	Type	
<code>NUMERIC_SCALE</code>	Type	
<code>DATETIME_PRECISION</code>	Type	
<code>CHARACTER_SET_NAME</code>		
<code>COLLATION_NAME</code>	Collation	
<code>COLUMN_TYPE</code>	Type	MySQL extension
<code>COLUMN_KEY</code>	Key	MySQL extension
<code>EXTRA</code>	Extra	MySQL extension
<code>PRIVILEGES</code>	Privileges	MySQL extension
<code>COLUMN_COMMENT</code>	Comment	MySQL extension

**Notes:**

- In `SHOW`, the `Type` display includes values from several different `COLUMNS` columns.
- `ORDINAL_POSITION` is necessary because you might want to say `ORDER BY ORDINAL_POSITION`. Unlike `SHOW`, `SELECT` does not have automatic ordering.
- `CHARACTER_OCTET_LENGTH` should be the same as `CHARACTER_MAXIMUM_LENGTH`, except for multi-byte character sets.
- `CHARACTER_SET_NAME` can be derived from `Collation`. For example, if you say `SHOW FULL COLUMNS FROM t`, and you see in the `Collation` column a value of `latin1_swedish_ci`, the character set is what is before the first underscore: `latin1`.
- `DATETIME_PRECISION` was added in MySQL 5.6.4.

The following statements are nearly equivalent:

```
SELECT COLUMN_NAME, DATA_TYPE, IS_NULLABLE, COLUMN_DEFAULT
FROM INFORMATION_SCHEMA.COLUMNS
WHERE table_name = 'tbl_name'
[AND table_schema = 'db_name']
[AND column_name LIKE 'wild']

SHOW COLUMNS
FROM tbl_name
[FROM db_name]
[LIKE 'wild']
```

## 19.5. The `INFORMATION_SCHEMA COLUMN_PRIVILEGES` Table

The `COLUMN_PRIVILEGES` table provides information about column privileges. This information comes from the `mysql.columns_priv` grant table.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>GRANTEE</code>		' <i>user_name</i> '@' <i>host_name</i> ' value
<code>TABLE_CATALOG</code>		<code>def</code>
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>COLUMN_NAME</code>		
<code>PRIVILEGE_TYPE</code>		
<code>IS_GRANTABLE</code>		

**Notes:**

- In the output from `SHOW FULL COLUMNS`, the privileges are all in one field and in lowercase, for example, `select,insert,update,references`. In `COLUMN_PRIVILEGES`, there is one privilege per row, in uppercase.
- `PRIVILEGE_TYPE` can contain one (and only one) of these values: `SELECT [755]`, `INSERT [755]`, `UPDATE [756]`, `REFERENCES [755]`.
- If the user has `GRANT OPTION [755]` privilege, `IS_GRANTABLE` should be `YES`. Otherwise, `IS_GRANTABLE` should be `NO`. The output does not list `GRANT OPTION [755]` as a separate privilege.

The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.COLUMN_PRIVILEGES

SHOW GRANTS ...
```

## 19.6. The `INFORMATION_SCHEMA.ENGINES` Table

The `ENGINES` table provides information about storage engines.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>ENGINE</code>	<code>Engine</code>	MySQL extension
<code>SUPPORT</code>	<code>Support</code>	MySQL extension
<code>COMMENT</code>	<code>Comment</code>	MySQL extension
<code>TRANSACTIONS</code>	<code>Transactions</code>	MySQL extension
<code>XA</code>	<code>XA</code>	MySQL extension
<code>SAVEPOINTS</code>	<code>Savepoints</code>	MySQL extension

**Notes:**

- The `ENGINES` table is a nonstandard table.

See also [Section 13.7.5.17](#), “`SHOW ENGINES` Syntax”.

## 19.7. The `INFORMATION_SCHEMA.EVENTS` Table

The `EVENTS` table provides information about scheduled events, which are discussed in [Section 18.4](#), “Using the Event Scheduler”. The `SHOW` Name values correspond to column names of the `SHOW EVENTS` statement.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>EVENT_CATALOG</code>		<code>def</code> , MySQL extension
<code>EVENT_SCHEMA</code>	<code>Db</code>	MySQL extension
<code>EVENT_NAME</code>	<code>Name</code>	MySQL extension
<code>DEFINER</code>	<code>Definer</code>	MySQL extension
<code>TIME_ZONE</code>	<code>Time zone</code>	MySQL extension
<code>EVENT_BODY</code>		MySQL extension
<code>EVENT_DEFINITION</code>		MySQL extension
<code>EVENT_TYPE</code>	<code>Type</code>	MySQL extension
<code>EXECUTE_AT</code>	<code>Execute at</code>	MySQL extension

INFORMATION_SCHEMA Name	SHOW Name	Remarks
INTERVAL_VALUE	Interval value	MySQL extension
INTERVAL_FIELD	Interval field	MySQL extension
SQL_MODE		MySQL extension
STARTS	Starts	MySQL extension
ENDS	Ends	MySQL extension
STATUS	Status	MySQL extension
ON_COMPLETION		MySQL extension
CREATED		MySQL extension
LAST_ALTERED		MySQL extension
LAST_EXECUTED		MySQL extension
EVENT_COMMENT		MySQL extension
ORIGINATOR	Originator	MySQL extension
CHARACTER_SET_CLIENT	character_set_client	MySQL extension
COLLATION_CONNECTION	collation_connection [468]	MySQL extension
DATABASE_COLLATION	Database Collation	MySQL extension

**Notes:**

- The `EVENTS` table is a nonstandard table.
- `EVENT_CATALOG`: The value of this column is always `def`.
- `EVENT_SCHEMA`: The name of the schema (database) to which this event belongs.
- `EVENT_NAME`: The name of the event.
- `DEFINER`: The account of the user who created the event, in `'user_name'@'host_name'` format.
- `TIME_ZONE`: The event time zone, which is the time zone used for scheduling the event and that is in effect within the event as it executes. The default value is `SYSTEM`.
- `EVENT_BODY`: The language used for the statements in the event's `DO` clause; in MySQL 5.6, this is always `SQL`.

This column is not to be confused with the column of the same name (now named `EVENT_DEFINITION`) that existed in earlier MySQL versions.

- `EVENT_DEFINITION`: The text of the SQL statement making up the event's `DO` clause; in other words, the statement executed by this event.
- `EVENT_TYPE`: The event repetition type, either `ONE TIME` (transient) or `RECURRING` (repeating).
- `EXECUTE_AT`: For a one-time event, this is the `DATETIME` value specified in the `AT` clause of the `CREATE EVENT` statement used to create the event, or of the last `ALTER EVENT` statement that modified the event. The value shown in this column reflects the addition or subtraction of any `INTERVAL` value included in the event's `AT` clause. For example, if an event is created using `ON SCHEDULE AT CURRENT_TIMESTAMP + '1:6' DAY_HOUR`, and the event was created at 2006-02-09 14:05:30, the value shown in this column would be `'2006-02-10 20:05:30'`.

If the event's timing is determined by an `EVERY` clause instead of an `AT` clause (that is, if the event is recurring), the value of this column is `NULL`.

- `INTERVAL_VALUE`: For recurring events, this column contains the numeric portion of the event's `EVERY` clause.

For a one-time event (that is, an event whose timing is determined by an `AT` clause), this column is `NULL`.

- `INTERVAL_FIELD`: For recurring events, this column contains the units portion of the `EVERY` clause governing the timing of the event. Thus, this column contains a value such as `'YEAR'`, `'QUARTER'`, `'DAY'`, and so on.

For a one-time event (that is, an event whose timing is determined by an `AT` clause), this column is `NULL`.

- `SQL_MODE`: The SQL mode in effect at the time the event was created or altered.
- `STARTS`: For a recurring event whose definition includes a `STARTS` clause, this column contains the corresponding `DATETIME` value. As with the `EXECUTE_AT` column, this value resolves any expressions used.

If there is no `STARTS` clause affecting the timing of the event, this column is `NULL`.

- `ENDS`: For a recurring event whose definition includes a `ENDS` clause, this column contains the corresponding `DATETIME` value. As with the `EXECUTE_AT` column, this value resolves any expressions used.

If there is no `ENDS` clause affecting the timing of the event, this column is `NULL`.

- `STATUS`: One of the three values `ENABLED`, `DISABLED`, or `SLAVESIDE_DISABLED`.

`SLAVESIDE_DISABLED` indicates that the creation of the event occurred on another MySQL server acting as a replication master and was replicated to the current MySQL server which is acting as a slave, but the event is not presently being executed on the slave. See [Section 16.4.1.10, “Replication of Invoked Features”](#), for more information.

- `ON_COMPLETION`: One of the two values `PRESERVE` or `NOT PRESERVE`.
- `CREATED`: The date and time when the event was created. This is a `DATETIME` value.
- `LAST_ALTERED`: The date and time when the event was last modified. This is a `DATETIME` value. If the event has not been modified since its creation, this column holds the same value as the `CREATED` column.
- `LAST_EXECUTED`: The date and time when the event last executed. A `DATETIME` value. If the event has never executed, this column is `NULL`.

`LAST_EXECUTED` indicates when the event started. As a result, the `ENDS` column is never less than `LAST_EXECUTED`.

- `EVENT_COMMENT`: The text of a comment, if the event has one. If not, the value of this column is an empty string.
- `ORIGINATOR`: The server ID of the MySQL server on which the event was created; used in replication. The default value is 0.
- `CHARACTER_SET_CLIENT` is the session value of the `character_set_client` [467] system variable when the event was created. `COLLATION_CONNECTION` is the session value of the `collation_connection` [469] system variable when the event was created. `DATABASE_COLLATION` is the collation of the database with which the event is associated.

**Example:** Suppose that the user `jon@ghidora` creates an event named `e_daily`, and then modifies it a few minutes later using an `ALTER EVENT` statement, as shown here:

```
DELIMITER |
```

```

CREATE EVENT e_daily
ON SCHEDULE
EVERY 1 DAY
COMMENT 'Saves total number of sessions then clears the table each day'
DO
BEGIN
INSERT INTO site_activity.totals (time, total)
SELECT CURRENT_TIMESTAMP, COUNT(*)
FROM site_activity.sessions;
DELETE FROM site_activity.sessions;
END |
DELIMITER ;

ALTER EVENT e_daily
ENABLED;

```

(Note that comments can span multiple lines.)

This user can then run the following `SELECT` statement, and obtain the output shown:

```

mysql> SELECT * FROM INFORMATION_SCHEMA.EVENTS
> WHERE EVENT_NAME = 'e_daily'
> AND EVENT_SCHEMA = 'myschema'\G
***** 1. row *****
EVENT_CATALOG: def
EVENT_SCHEMA: test
EVENT_NAME: e_daily
DEFINER: paul@localhost
TIME_ZONE: SYSTEM
EVENT_BODY: SQL
EVENT_DEFINITION: BEGIN
INSERT INTO site_activity.totals (time, total)
SELECT CURRENT_TIMESTAMP, COUNT(*)
FROM site_activity.sessions;
DELETE FROM site_activity.sessions;
END
EVENT_TYPE: RECURRING
EXECUTE_AT: NULL
INTERVAL_VALUE: 1
INTERVAL_FIELD: DAY
SQL_MODE:
STARTS: 2008-09-03 12:13:39
ENDS: NULL
STATUS: ENABLED
ON_COMPLETION: NOT PRESERVE
CREATED: 2008-09-03 12:13:39
LAST_ALTERED: 2008-09-03 12:13:39
LAST_EXECUTED: NULL
EVENT_COMMENT: Saves total number of sessions then clears the
table each day
ORIGINATOR: 1
CHARACTER_SET_CLIENT: latin1
COLLATION_CONNECTION: latin1_swedish_ci
DATABASE_COLLATION: latin1_swedish_ci

```

Times in the `EVENTS` table are displayed using the event time zone or the current session time zone, as described in [Section 18.4.4, “Event Metadata”](#).

See also [Section 13.7.5.19, “SHOW EVENTS Syntax”](#).

## 19.8. The `INFORMATION_SCHEMA.FILES` Table

The `FILES` table provides information about the files in which MySQL tablespace data is stored.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>FILE_ID</code>		MySQL extension
<code>FILE_NAME</code>		MySQL extension

INFORMATION_SCHEMA Name	SHOW Name	Remarks
FILE_TYPE		MySQL extension
TABLESPACE_NAME		MySQL extension
TABLE_CATALOG		MySQL extension
TABLE_SCHEMA		MySQL extension
TABLE_NAME		MySQL extension
LOGFILE_GROUP_NAME		MySQL extension
LOGFILE_GROUP_NUMBER		MySQL extension
ENGINE		MySQL extension
FULLTEXT_KEYS		MySQL extension
DELETED_ROWS		MySQL extension
UPDATE_COUNT		MySQL extension
FREE_EXTENTS		MySQL extension
TOTAL_EXTENTS		MySQL extension
EXTENT_SIZE		MySQL extension
INITIAL_SIZE		MySQL extension
MAXIMUM_SIZE		MySQL extension
AUTOEXTEND_SIZE		MySQL extension
CREATION_TIME		MySQL extension
LAST_UPDATE_TIME		MySQL extension
LAST_ACCESS_TIME		MySQL extension
RECOVER_TIME		MySQL extension
TRANSACTION_COUNTER		MySQL extension
VERSION		MySQL extension
ROW_FORMAT		MySQL extension
TABLE_ROWS		MySQL extension
AVG_ROW_LENGTH		MySQL extension
DATA_LENGTH		MySQL extension
MAX_DATA_LENGTH		MySQL extension
INDEX_LENGTH		MySQL extension
DATA_FREE		MySQL extension
CREATE_TIME		MySQL extension
UPDATE_TIME		MySQL extension
CHECK_TIME		MySQL extension
CHECKSUM		MySQL extension
STATUS		MySQL extension
EXTRA		MySQL extension

**Notes:**

- `FILE_ID` column values are auto-generated.
- `FILE_NAME` is the name of a data file created by `CREATE TABLESPACE` or `ALTER TABLESPACE`.
- `FILE_TYPE` is the tablespace file type.



- `TABLESPACE_NAME` is the name of the tablespace with which the file is associated.
- Currently, the value of the `TABLESPACE_CATALOG` column is always `NULL`.
- `TABLE_NAME` is the name of the table with which the file is associated, if any.
- The `EXTENT_SIZE` is always 0.
- There are no `SHOW` statements associated with the `FILES` table.

## 19.9. The `INFORMATION_SCHEMA` `GLOBAL_STATUS` and `SESSION_STATUS` Tables

The `GLOBAL_STATUS` and `SESSION_STATUS` tables provide information about server status variables. Their contents correspond to the information produced by the `SHOW GLOBAL STATUS` and `SHOW SESSION STATUS` statements (see [Section 13.7.5.36](#), “`SHOW STATUS Syntax`”).

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>VARIABLE_NAME</code>	Variable_name	
<code>VARIABLE_VALUE</code>	Value	

### Notes:

- The `VARIABLE_VALUE` column for each of these tables is defined as `VARCHAR(20480)`.

## 19.10. The `INFORMATION_SCHEMA` `GLOBAL_VARIABLES` and `SESSION_VARIABLES` Tables

The `GLOBAL_VARIABLES` and `SESSION_VARIABLES` tables provide information about server status variables. Their contents correspond to the information produced by the `SHOW GLOBAL VARIABLES` and `SHOW SESSION VARIABLES` statements (see [Section 13.7.5.40](#), “`SHOW VARIABLES Syntax`”).

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>VARIABLE_NAME</code>	Variable_name	
<code>VARIABLE_VALUE</code>	Value	

### Notes:

- The `VARIABLE_VALUE` column for each of these tables is defined as `VARCHAR(20480)`.

## 19.11. The `INFORMATION_SCHEMA` `KEY_COLUMN_USAGE` Table

The `KEY_COLUMN_USAGE` table describes which key columns have constraints.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CONSTRAINT_CATALOG</code>		def
<code>CONSTRAINT_SCHEMA</code>		
<code>CONSTRAINT_NAME</code>		
<code>TABLE_CATALOG</code>		def
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>COLUMN_NAME</code>		
<code>ORDINAL_POSITION</code>		

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>POSITION_IN_UNIQUE_CONSTRAINT</code>		
<code>REFERENCED_TABLE_SCHEMA</code>		
<code>REFERENCED_TABLE_NAME</code>		
<code>REFERENCED_COLUMN_NAME</code>		

**Notes:**

- If the constraint is a foreign key, then this is the column of the foreign key, not the column that the foreign key references.
- The value of `ORDINAL_POSITION` is the column's position within the constraint, not the column's position within the table. Column positions are numbered beginning with 1.
- The value of `POSITION_IN_UNIQUE_CONSTRAINT` is `NULL` for unique and primary-key constraints. For foreign-key constraints, it is the ordinal position in key of the table that is being referenced.

Suppose that there are two tables name `t1` and `t3` that have the following definitions:

```
CREATE TABLE t1
(
  s1 INT,
  s2 INT,
  s3 INT,
  PRIMARY KEY(s3)
) ENGINE=InnoDB;

CREATE TABLE t3
(
  s1 INT,
  s2 INT,
  s3 INT,
  KEY(s1),
  CONSTRAINT CO FOREIGN KEY (s2) REFERENCES t1(s3)
) ENGINE=InnoDB;
```

For those two tables, the `KEY_COLUMN_USAGE` table has two rows:

- One row with `CONSTRAINT_NAME` = 'PRIMARY', `TABLE_NAME` = 't1', `COLUMN_NAME` = 's3', `ORDINAL_POSITION` = 1, `POSITION_IN_UNIQUE_CONSTRAINT` = `NULL`.
- One row with `CONSTRAINT_NAME` = 'CO', `TABLE_NAME` = 't3', `COLUMN_NAME` = 's2', `ORDINAL_POSITION` = 1, `POSITION_IN_UNIQUE_CONSTRAINT` = 1.

## 19.12. The `INFORMATION_SCHEMA OPTIMIZER_TRACE` Table

The `OPTIMIZER_TRACE` table provides information produced by the optimizer tracing capability. To enable tracking, use the `optimizer_trace` [523] system variable. For details, see [MySQL Internals: Tracing the Optimizer](#).

The `OPTIMIZER_TRACE` table was added in MySQL 5.6.3.

## 19.13. The `INFORMATION_SCHEMA PARAMETERS` Table

The `PARAMETERS` table provides information about stored procedure and function parameters, and about return values for stored functions. Parameter information is similar to the contents of the `param_list` column in the `mysql.proc` table.

<code>INFORMATION_SCHEMA</code> Name	<code>mysql.proc</code> Name	Remarks
<code>SPECIFIC_CATALOG</code>		<code>def</code>

<code>INFORMATION_SCHEMA</code> Name	<code>mysql.proc</code> Name	Remarks
<code>SPECIFIC_SCHEMA</code>	<code>db</code>	routine database
<code>SPECIFIC_NAME</code>	<code>name</code>	routine name
<code>ORDINAL_POSITION</code>		1, 2, 3, ... for parameters, 0 for function <code>RETURNS</code> clause
<code>PARAMETER_MODE</code>		<code>IN</code> , <code>OUT</code> , <code>INOUT</code> ( <code>NULL</code> for <code>RETURNS</code> )
<code>PARAMETER_NAME</code>		parameter name ( <code>NULL</code> for <code>RETURNS</code> )
<code>DATA_TYPE</code>		same as for <code>COLUMNS</code> table
<code>CHARACTER_MAXIMUM_LENGTH</code>		same as for <code>COLUMNS</code> table
<code>CHARACTER_OCTET_LENGTH</code>		same as for <code>COLUMNS</code> table
<code>NUMERIC_PRECISION</code>		same as for <code>COLUMNS</code> table
<code>NUMERIC_SCALE</code>		same as for <code>COLUMNS</code> table
<code>DATETIME_PRECISION</code>		same as for <code>COLUMNS</code> table
<code>CHARACTER_SET_NAME</code>		same as for <code>COLUMNS</code> table
<code>COLLATION_NAME</code>		same as for <code>COLUMNS</code> table
<code>DTD_IDENTIFIER</code>		same as for <code>COLUMNS</code> table
<code>ROUTINE_TYPE</code>	<code>type</code>	same as for <code>ROUTINES</code> table

**Notes:**

- For successive parameters of a stored procedure or function, the `ORDINAL_POSITION` values are 1, 2, 3, and so forth. For a stored function, there is also a row that describes the data type for the `RETURNS` clause. The return value is not a true parameter, so the row that describes it has these unique characteristics:
  - The `ORDINAL_POSITION` value is 0.
  - The `PARAMETER_NAME` and `PARAMETER_MODE` values are `NULL` because the return value has no name and the mode does not apply.
  - `DATETIME_PRECISION` was added in MySQL 5.6.4.

## 19.14. The `INFORMATION_SCHEMA PARTITIONS` Table

The `PARTITIONS` table provides information about table partitions. See [Chapter 17, Partitioning](#), for more information about partitioning tables.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>TABLE_CATALOG</code>		MySQL extension
<code>TABLE_SCHEMA</code>		MySQL extension
<code>TABLE_NAME</code>		MySQL extension
<code>PARTITION_NAME</code>		MySQL extension
<code>SUBPARTITION_NAME</code>		MySQL extension
<code>PARTITION_ORDINAL_POSITION</code>		MySQL extension
<code>SUBPARTITION_ORDINAL_POSITION</code>		MySQL extension
<code>PARTITION_METHOD</code>		MySQL extension
<code>SUBPARTITION_METHOD</code>		MySQL extension

INFORMATION_SCHEMA Name	SHOW Name	Remarks
PARTITION_EXPRESSION		MySQL extension
SUBPARTITION_EXPRESSION		MySQL extension
PARTITION_DESCRIPTION		MySQL extension
TABLE_ROWS		MySQL extension
AVG_ROW_LENGTH		MySQL extension
DATA_LENGTH		MySQL extension
MAX_DATA_LENGTH		MySQL extension
INDEX_LENGTH		MySQL extension
DATA_FREE		MySQL extension
CREATE_TIME		MySQL extension
UPDATE_TIME		MySQL extension
CHECK_TIME		MySQL extension
CHECKSUM		MySQL extension
PARTITION_COMMENT		MySQL extension
NODEGROUP		MySQL extension
TABLESPACE_NAME		MySQL extension

**Notes:**

- The `PARTITIONS` table is a nonstandard table.  
Each record in this table corresponds to an individual partition or subpartition of a partitioned table.
- `TABLE_CATALOG`: This column is always `def`.
- `TABLE_SCHEMA`: This column contains the name of the database to which the table belongs.
- `TABLE_NAME`: This column contains the name of the table containing the partition.
- `PARTITION_NAME`: The name of the partition.
- `SUBPARTITION_NAME`: If the `PARTITIONS` table record represents a subpartition, then this column contains the name of subpartition; otherwise it is `NULL`.
- `PARTITION_ORDINAL_POSITION`: All partitions are indexed in the same order as they are defined, with `1` being the number assigned to the first partition. The indexing can change as partitions are added, dropped, and reorganized; the number shown in this column reflects the current order, taking into account any indexing changes.
- `SUBPARTITION_ORDINAL_POSITION`: Subpartitions within a given partition are also indexed and reindexed in the same manner as partitions are indexed within a table.
- `PARTITION_METHOD`: One of the values `RANGE`, `LIST`, `HASH`, `LINEAR HASH`, `KEY`, or `LINEAR KEY`; that is, one of the available partitioning types as discussed in [Section 17.2, “Partitioning Types”](#).
- `SUBPARTITION_METHOD`: One of the values `HASH`, `LINEAR HASH`, `KEY`, or `LINEAR KEY`; that is, one of the available subpartitioning types as discussed in [Section 17.2.6, “Subpartitioning”](#).
- `PARTITION_EXPRESSION`: This is the expression for the partitioning function used in the `CREATE TABLE` or `ALTER TABLE` statement that created the table's current partitioning scheme.

For example, consider a partitioned table created in the `test` database using this statement:

```
CREATE TABLE tp (
  c1 INT,
```

```

    c2 INT,
    c3 VARCHAR(25)
)
PARTITION BY HASH(c1 + c2)
PARTITIONS 4;

```

The `PARTITION_EXPRESSION` column in a `PARTITIONS` table record for a partition from this table displays `c1 + c2`, as shown here:

```

mysql> SELECT DISTINCT PARTITION_EXPRESSION
>       FROM INFORMATION_SCHEMA.PARTITIONS
>       WHERE TABLE_NAME='tp' AND TABLE_SCHEMA='test';
+-----+
| PARTITION_EXPRESSION |
+-----+
| c1 + c2              |
+-----+
1 row in set (0.09 sec)

```

- **SUBPARTITION\_EXPRESSION**: This works in the same fashion for the subpartitioning expression that defines the subpartitioning for a table as `PARTITION_EXPRESSION` does for the partitioning expression used to define a table's partitioning.

If the table has no subpartitions, then this column is `NULL`.

- **PARTITION\_DESCRIPTION**: This column is used for `RANGE` and `LIST` partitions. For a `RANGE` partition, it contains the value set in the partition's `VALUES LESS THAN` clause, which can be either an integer or `MAXVALUE`. For a `LIST` partition, this column contains the values defined in the partition's `VALUES IN` clause, which is a comma-separated list of integer values.

For partitions whose `PARTITION_METHOD` is other than `RANGE` or `LIST`, this column is always `NULL`.

- **TABLE\_ROWS**: The number of table rows in the partition.

For partitioned `InnoDB` tables, the row count given in the `TABLE_ROWS` column is only an estimated value used in SQL optimization, and may not always be exact.

- **AVG\_ROW\_LENGTH**: The average length of the rows stored in this partition or subpartition, in bytes.

This is the same as `DATA_LENGTH` divided by `TABLE_ROWS`.

- **DATA\_LENGTH**: The total length of all rows stored in this partition or subpartition, in bytes—that is, the total number of bytes stored in the partition or subpartition.
- **MAX\_DATA\_LENGTH**: The maximum number of bytes that can be stored in this partition or subpartition.
- **INDEX\_LENGTH**: The length of the index file for this partition or subpartition, in bytes.
- **DATA\_FREE**: The number of bytes allocated to the partition or subpartition but not used.
- **CREATE\_TIME**: The time of the partition's or subpartition's creation.
- **UPDATE\_TIME**: The time that the partition or subpartition was last modified.
- **CHECK\_TIME**: The last time that the table to which this partition or subpartition belongs was checked.



#### Note

Some storage engines do not update this time; for tables using these storage engines, this value is always `NULL`.

- **CHECKSUM**: The checksum value, if any; otherwise, this column is `NULL`.
- **PARTITION\_COMMENT**: This column contains the text of any comment made for the partition.

Prior to MySQL 5.6.6, the display width of this column was 80 characters, and partition comments which exceeded this length were truncated to fit. As of MySQL 5.6.6, the maximum length for a partition comment is defined as 1024 characters, and the display width of the `PARTITION_COMMENT` column is increased to 1024 characters to match this limit (Bug #11748924, Bug #37728).

The default value for this column is an empty string.

- `NODEGROUP`: This is the nodegroup to which the partition belongs. This is relevant only to MySQL Cluster tables; otherwise the value of this column is always 0.
- `TABLESPACE_NAME`: This column contains the name of the tablespace to which the partition belongs. Currently, the value of this column is always `DEFAULT`.
- A nonpartitioned table has one record in `INFORMATION_SCHEMA.PARTITIONS`; however, the values of the `PARTITION_NAME`, `SUBPARTITION_NAME`, `PARTITION_ORDINAL_POSITION`, `SUBPARTITION_ORDINAL_POSITION`, `PARTITION_METHOD`, `SUBPARTITION_METHOD`, `PARTITION_EXPRESSION`, `SUBPARTITION_EXPRESSION`, and `PARTITION_DESCRIPTION` columns are all `NULL`. (The `PARTITION_COMMENT` column in this case is blank.)

## 19.15. The `INFORMATION_SCHEMA PLUGINS` Table

The `PLUGINS` table provides information about server plugins.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>PLUGIN_NAME</code>	Name	MySQL extension
<code>PLUGIN_VERSION</code>		MySQL extension
<code>PLUGIN_STATUS</code>	Status	MySQL extension
<code>PLUGIN_TYPE</code>	Type	MySQL extension
<code>PLUGIN_TYPE_VERSION</code>		MySQL extension
<code>PLUGIN_LIBRARY</code>	Library	MySQL extension
<code>PLUGIN_LIBRARY_VERSION</code>		MySQL extension
<code>PLUGIN_AUTHOR</code>		MySQL extension
<code>PLUGIN_DESCRIPTION</code>		MySQL extension
<code>PLUGIN_LICENSE</code>		MySQL extension
<code>LOAD_OPTION</code>		MySQL extension

### Notes:

- The `PLUGINS` table is a nonstandard table.
- `PLUGIN_NAME` is the name used to refer to the plugin in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`.
- `PLUGIN_VERSION` is the version from the plugin's general type descriptor.
- `PLUGIN_STATUS` indicates the plugin status, one of `ACTIVE`, `INACTIVE`, `DISABLED`, or `DELETED`.
- `PLUGIN_TYPE` indicates the type of plugin, such as `STORAGE ENGINE`, `INFORMATION_SCHEMA`, or `AUTHENTICATION`.
- `PLUGIN_TYPE_VERSION` is the version from the plugin's type-specific descriptor.
- `PLUGIN_LIBRARY` is the name of the plugin shared object file. This is the name used to refer to the plugin file in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`. This file is located in the directory named by the `plugin_dir` [525] system variable. If the library name is `NULL`, the plugin is compiled in and cannot be uninstalled with `UNINSTALL PLUGIN`.

- `PLUGIN_LIBRARY_VERSION` indicates the plugin API interface version.
- `PLUGIN_AUTHOR` names the plugin author.
- `PLUGIN_DESCRIPTION` provides a short description of the plugin.
- `PLUGIN_LICENSE` indicates how the plugin is licensed; for example, `GPL`.
- `LOAD_OPTION` indicates how the plugin was loaded. The value is `OFF`, `ON`, `FORCE`, or `FORCE_PLUS_PERMANENT`. See [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

For plugins installed with `INSTALL PLUGIN`, the `PLUGIN_NAME` and `PLUGIN_LIBRARY` values are also registered in the `mysql.plugin` table.

These statements are equivalent:

```
SELECT
  PLUGIN_NAME, PLUGIN_STATUS, PLUGIN_TYPE,
  PLUGIN_LIBRARY, PLUGIN_LICENSE
FROM INFORMATION_SCHEMA.PLUGINS;

SHOW PLUGINS;
```

For information about plugin data structures that form the basis of the information in the `PLUGINS` table, see [Section 22.2, “The MySQL Plugin API”](#).

Plugin information is also available using the `SHOW PLUGINS` statement. See [Section 13.7.5.26, “SHOW PLUGINS Syntax”](#).

## 19.16. The `INFORMATION_SCHEMA` `PROCESSLIST` Table

The `PROCESSLIST` table provides information about which threads are running.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>ID</code>	<code>Id</code>	MySQL extension
<code>USER</code>	<code>User</code>	MySQL extension
<code>HOST</code>	<code>Host</code>	MySQL extension
<code>DB</code>	<code>db</code>	MySQL extension
<code>COMMAND</code>	<code>Command</code>	MySQL extension
<code>TIME</code>	<code>Time</code>	MySQL extension
<code>STATE</code>	<code>State</code>	MySQL extension
<code>INFO</code>	<code>Info</code>	MySQL extension

For an extensive description of the table columns, see [Section 13.7.5.30, “SHOW PROCESSLIST Syntax”](#).

### Notes:

- The `PROCESSLIST` table is a nonstandard table.
- Like the output from the corresponding `SHOW` statement, the `PROCESSLIST` table will only show information about your own threads, unless you have the `PROCESS [755]` privilege, in which case you will see information about other threads, too. As an anonymous user, you cannot see any rows at all.
- If an SQL statement refers to `INFORMATION_SCHEMA.PROCESSLIST`, then MySQL will populate the entire table once, when statement execution begins, so there is read consistency during the statement. There is no read consistency for a multi-statement transaction, though.
- Process information is also available from the `performance_schema.threads` table. However, access to `threads` does not require a mutex and has minimal impact on server performance.



`INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` have negative performance consequences because they require a mutex. `threads` also shows information about background threads, which `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` do not. This means that `threads` can be used to monitor activity the other thread information sources cannot.

The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.PROCESSLIST
SHOW FULL PROCESSLIST
```

## 19.17. The `INFORMATION_SCHEMA.PROFILING` Table

The `PROFILING` table provides statement profiling information. Its contents correspond to the information produced by the `SHOW PROFILES` and `SHOW PROFILE` statements (see [Section 13.7.5.32](#), “`SHOW PROFILES` Syntax”). The table is empty unless the `profiling` [526] session variable is set to 1.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>QUERY_ID</code>	<code>Query_ID</code>	
<code>SEQ</code>		
<code>STATE</code>	<code>Status</code>	
<code>DURATION</code>	<code>Duration</code>	
<code>CPU_USER</code>	<code>CPU_user</code>	
<code>CPU_SYSTEM</code>	<code>CPU_system</code>	
<code>CONTEXT_VOLUNTARY</code>	<code>Context_voluntary</code>	
<code>CONTEXT_INVOLUNTARY</code>	<code>Context_involuntary</code>	
<code>BLOCK_OPS_IN</code>	<code>Block_ops_in</code>	
<code>BLOCK_OPS_OUT</code>	<code>Block_ops_out</code>	
<code>MESSAGES_SENT</code>	<code>Messages_sent</code>	
<code>MESSAGES_RECEIVED</code>	<code>Messages_received</code>	
<code>PAGE_FAULTS_MAJOR</code>	<code>Page_faults_major</code>	
<code>PAGE_FAULTS_MINOR</code>	<code>Page_faults_minor</code>	
<code>SWAPS</code>	<code>Swaps</code>	
<code>SOURCE_FUNCTION</code>	<code>Source_function</code>	
<code>SOURCE_FILE</code>	<code>Source_file</code>	
<code>SOURCE_LINE</code>	<code>Source_line</code>	

### Notes:

- `QUERY_ID` is a numeric statement identifier.
- `SEQ` is a sequence number indicating the display order for rows with the same `QUERY_ID` value.
- `STATE` is the profiling state to which the row measurements apply.
- `DURATION` indicates how long statement execution remained in the given state, in seconds.
- `CPU_USER` and `CPU_SYSTEM` indicate user and system CPU use, in seconds.
- `CONTEXT_VOLUNTARY` and `CONTEXT_INVOLUNTARY` indicate how many voluntary and involuntary context switches occurred.
- `BLOCK_OPS_IN` and `BLOCK_OPS_OUT` indicate the number of block input and output operations.

- `MESSAGES_SENT` and `MESSAGES_RECEIVED` indicate the number of communication messages sent and received.
- `PAGE_FAULTS_MAJOR` and `PAGE_FAULTS_MINOR` indicate the number of major and minor page faults.
- `SWAPS` indicates how many swaps occurred.
- `SOURCE_FUNCTION`, `SOURCE_FILE`, and `SOURCE_LINE` provide information indicating where in the source code the profiled state executes.

## 19.18. The `INFORMATION_SCHEMA REFERENTIAL_CONSTRAINTS` Table

The `REFERENTIAL_CONSTRAINTS` table provides information about foreign keys.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CONSTRAINT_CATALOG</code>		def
<code>CONSTRAINT_SCHEMA</code>		
<code>CONSTRAINT_NAME</code>		
<code>UNIQUE_CONSTRAINT_CATALOG</code>		def
<code>UNIQUE_CONSTRAINT_SCHEMA</code>		
<code>UNIQUE_CONSTRAINT_NAME</code>		
<code>MATCH_OPTION</code>		
<code>UPDATE_RULE</code>		
<code>DELETE_RULE</code>		
<code>TABLE_NAME</code>		
<code>REFERENCED_TABLE_NAME</code>		

### Notes:

- `TABLE_NAME` has the same value as `TABLE_NAME` in `INFORMATION_SCHEMA.TABLE_CONSTRAINTS`.
- `CONSTRAINT_SCHEMA` and `CONSTRAINT_NAME` identify the foreign key.
- `UNIQUE_CONSTRAINT_SCHEMA`, `UNIQUE_CONSTRAINT_NAME`, and `REFERENCED_TABLE_NAME` identify the referenced key.
- The only valid value at this time for `MATCH_OPTION` is `NONE`.
- The possible values for `UPDATE_RULE` or `DELETE_RULE` are `CASCADE`, `SET NULL`, `SET DEFAULT`, `RESTRICT`, `NO ACTION`.

## 19.19. The `INFORMATION_SCHEMA ROUTINES` Table

The `ROUTINES` table provides information about stored routines (both procedures and functions). The `ROUTINES` table does not include user-defined functions (UDFs).

The column named “`mysql.proc name`” indicates the `mysql.proc` table column that corresponds to the `INFORMATION_SCHEMA.ROUTINES` table column, if any.

<code>INFORMATION_SCHEMA</code> Name	<code>mysql.proc</code> Name	Remarks
<code>SPECIFIC_NAME</code>	<code>specific_name</code>	

INFORMATION_SCHEMA Name	mysql.proc Name	Remarks
ROUTINE_CATALOG		def
ROUTINE_SCHEMA	db	
ROUTINE_NAME	name	
ROUTINE_TYPE	type	{PROCEDURE   FUNCTION}
DATA_TYPE		same as for COLUMNS table
CHARACTER_MAXIMUM_LENGTH		same as for COLUMNS table
CHARACTER_OCTET_LENGTH		same as for COLUMNS table
NUMERIC_PRECISION		same as for COLUMNS table
NUMERIC_SCALE		same as for COLUMNS table
DATETIME_PRECISION		same as for COLUMNS table
CHARACTER_SET_NAME		same as for COLUMNS table
COLLATION_NAME		same as for COLUMNS table
DTD_IDENTIFIER		data type descriptor
ROUTINE_BODY		SQL
ROUTINE_DEFINITION	body_utf8	
EXTERNAL_NAME		NULL
EXTERNAL_LANGUAGE	language	NULL
PARAMETER_STYLE		SQL
IS_DETERMINISTIC	is_deterministic	
SQL_DATA_ACCESS	sql_data_access	
SQL_PATH		NULL
SECURITY_TYPE	security_type	
CREATED	created	
LAST_ALTERED	modified	
SQL_MODE	sql_mode [545]	MySQL extension
ROUTINE_COMMENT	comment	MySQL extension
DEFINER	definer	MySQL extension
CHARACTER_SET_CLIENT		MySQL extension
COLLATION_CONNECTION		MySQL extension
DATABASE_COLLATION		MySQL extension

**Notes:**

- MySQL calculates `EXTERNAL_LANGUAGE` thus:
  - If `mysql.proc.language = 'SQL'`, `EXTERNAL_LANGUAGE` is `NULL`
  - Otherwise, `EXTERNAL_LANGUAGE` is what is in `mysql.proc.language`. However, we do not have external languages yet, so it is always `NULL`.
- `CHARACTER_SET_CLIENT` is the session value of the `character_set_client` [467] system variable when the routine was created. `COLLATION_CONNECTION` is the session value of the `collation_connection` [469] system variable when the routine was created. `DATABASE_COLLATION` is the collation of the database with which the routine is associated.
- The `DATA_TYPE`, `CHARACTER_MAXIMUM_LENGTH`, `CHARACTER_OCTET_LENGTH`, `NUMERIC_PRECISION`, `NUMERIC_SCALE`, `DATETIME_PRECISION`, `CHARACTER_SET_NAME`,

and `COLLATION_NAME` columns provide information about the data type for the `RETURNS` clause of stored functions. If a stored routine is a stored procedure, these columns all are `NULL`. `DATETIME_PRECISION` was added in MySQL 5.6.4.

- Information about stored function `RETURNS` data types is also available in the `PARAMETERS` table. The return value data type row for a function can be identified as the row that has an `ORDINAL_POSITION` value of 0.

## 19.20. The `INFORMATION_SCHEMA.SCHEMATA` Table

A schema is a database, so the `SCHEMATA` table provides information about databases.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CATALOG_NAME</code>		<code>def</code>
<code>SCHEMA_NAME</code>		Database
<code>DEFAULT_CHARACTER_SET_NAME</code>		
<code>DEFAULT_COLLATION_NAME</code>		
<code>SQL_PATH</code>		<code>NULL</code>

The following statements are equivalent:

```
SELECT SCHEMA_NAME AS `Database`
FROM INFORMATION_SCHEMA.SCHEMATA
[WHERE SCHEMA_NAME LIKE 'wild']
```

```
SHOW DATABASES
[LIKE 'wild']
```

## 19.21. The `INFORMATION_SCHEMA.SCHEMA_PRIVILEGES` Table

The `SCHEMA_PRIVILEGES` table provides information about schema (database) privileges. This information comes from the `mysql.db` grant table.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>GRANTEE</code>		<code>'user_name'@'host_name'</code> value, MySQL extension
<code>TABLE_CATALOG</code>		<code>def</code> , MySQL extension
<code>TABLE_SCHEMA</code>		MySQL extension
<code>PRIVILEGE_TYPE</code>		MySQL extension
<code>IS_GRANTABLE</code>		MySQL extension

**Notes:**

- This is a nonstandard table. It takes its values from the `mysql.db` table.

## 19.22. The `INFORMATION_SCHEMA.STATISTICS` Table

The `STATISTICS` table provides information about table indexes.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>TABLE_CATALOG</code>		<code>def</code>
<code>TABLE_SCHEMA</code>		= Database

INFORMATION_SCHEMA Name	SHOW Name	Remarks
TABLE_NAME	Table	
NON_UNIQUE	Non_unique	
INDEX_SCHEMA		= Database
INDEX_NAME	Key_name	
SEQ_IN_INDEX	Seq_in_index	
COLUMN_NAME	Column_name	
COLLATION	Collation	
CARDINALITY	Cardinality	
SUB_PART	Sub_part	MySQL extension
PACKED	Packed	MySQL extension
NULLABLE	Null	MySQL extension
INDEX_TYPE	Index_type	MySQL extension
COMMENT	Comment	MySQL extension

**Notes:**

- There is no standard table for indexes. The preceding list is similar to what SQL Server 2000 returns for `sp_statistics`, except that we replaced the name `QUALIFIER` with `CATALOG` and we replaced the name `OWNER` with `SCHEMA`.

Clearly, the preceding table and the output from `SHOW INDEX` are derived from the same parent. So the correlation is already close.

The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.STATISTICS
  WHERE table_name = 'tbl_name'
     AND table_schema = 'db_name'

SHOW INDEX
  FROM tbl_name
  FROM db_name
```

## 19.23. The INFORMATION\_SCHEMA TABLES Table

The `TABLES` table provides information about tables in databases.

INFORMATION_SCHEMA Name	SHOW Name	Remarks
TABLE_CATALOG		def
TABLE_SCHEMA	Table_...	
TABLE_NAME	Table_...	
TABLE_TYPE		
ENGINE	Engine	MySQL extension
VERSION	Version	The version number of the table's <code>.frm</code> file, MySQL extension
ROW_FORMAT	Row_format	MySQL extension
TABLE_ROWS	Rows	MySQL extension
AVG_ROW_LENGTH	Avg_row_length	MySQL extension

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>DATA_LENGTH</code>	<code>Data_length</code>	MySQL extension
<code>MAX_DATA_LENGTH</code>	<code>Max_data_length</code>	MySQL extension
<code>INDEX_LENGTH</code>	<code>Index_length</code>	MySQL extension
<code>DATA_FREE</code>	<code>Data_free</code>	MySQL extension
<code>AUTO_INCREMENT</code>	<code>Auto_increment</code>	MySQL extension
<code>CREATE_TIME</code>	<code>Create_time</code>	MySQL extension
<code>UPDATE_TIME</code>	<code>Update_time</code>	MySQL extension
<code>CHECK_TIME</code>	<code>Check_time</code>	MySQL extension
<code>TABLE_COLLATION</code>	<code>Collation</code>	MySQL extension
<code>CHECKSUM</code>	<code>Checksum</code>	MySQL extension
<code>CREATE_OPTIONS</code>	<code>Create_options</code>	MySQL extension
<code>TABLE_COMMENT</code>	<code>Comment</code>	MySQL extension

**Notes:**

- `TABLE_SCHEMA` and `TABLE_NAME` are a single field in a `SHOW` display, for example `Table_in_db1`.
- `TABLE_TYPE` should be `BASE TABLE` or `VIEW`. Currently, the `TABLES` table does not list `TEMPORARY` tables.
- For partitioned tables, the `ENGINE` column shows the name of the storage engine used by all partitions. (Previously, this column showed `PARTITION` for such tables.)
- The `TABLE_ROWS` column is `NULL` if the table is in the `INFORMATION_SCHEMA` database.

For `InnoDB` tables, the row count is only a rough estimate used in SQL optimization. (This is also true if the `InnoDB` table is partitioned.)

- The `DATA_FREE` column shows the free space in bytes for `InnoDB` tables.
- We have nothing for the table's default character set. `TABLE_COLLATION` is close, because collation names begin with a character set name.
- The `CREATE_OPTIONS` column shows `partitioned` if the table is partitioned.

The following statements are equivalent:

```
SELECT table_name FROM INFORMATION_SCHEMA.TABLES
  WHERE table_schema = 'db_name'
  [AND table_name LIKE 'wild' ]

SHOW TABLES
  FROM db_name
  [LIKE 'wild' ]
```

## 19.24. The `INFORMATION_SCHEMA TABLESPACES` Table

The `TABLESPACES` table provides information about active tablespaces.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>TABLESPACE_NAME</code>		MySQL extension
<code>ENGINE</code>		MySQL extension
<code>TABLESPACE_TYPE</code>		MySQL extension

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>LOGFILE_GROUP_NAME</code>		MySQL extension
<code>EXTENT_SIZE</code>		MySQL extension
<code>AUTOEXTEND_SIZE</code>		MySQL extension
<code>MAXIMUM_SIZE</code>		MySQL extension
<code>NODEGROUP_ID</code>		MySQL extension
<code>TABLESPACE_COMMENT</code>		MySQL extension

## 19.25. The `INFORMATION_SCHEMA TABLE_CONSTRAINTS` Table

The `TABLE_CONSTRAINTS` table describes which tables have constraints.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>CONSTRAINT_CATALOG</code>		def
<code>CONSTRAINT_SCHEMA</code>		
<code>CONSTRAINT_NAME</code>		
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>CONSTRAINT_TYPE</code>		

### Notes:

- The `CONSTRAINT_TYPE` value can be `UNIQUE`, `PRIMARY KEY`, or `FOREIGN KEY`.
- The `UNIQUE` and `PRIMARY KEY` information is about the same as what you get from the `Key_name` field in the output from `SHOW INDEX` when the `Non_unique` field is 0.
- The `CONSTRAINT_TYPE` column can contain one of these values: `UNIQUE`, `PRIMARY KEY`, `FOREIGN KEY`, `CHECK`. This is a `CHAR` (not `ENUM`) column. The `CHECK` value is not available until we support `CHECK`.

## 19.26. The `INFORMATION_SCHEMA TABLE_PRIVILEGES` Table

The `TABLE_PRIVILEGES` table provides information about table privileges. This information comes from the `mysql.tables_priv` grant table.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>GRANTEE</code>		' <i>user_name</i> '@' <i>host_name</i> ' value
<code>TABLE_CATALOG</code>		def
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>PRIVILEGE_TYPE</code>		
<code>IS_GRANTABLE</code>		

### Notes:

- `PRIVILEGE_TYPE` can contain one (and only one) of these values: `SELECT` [755], `INSERT` [755], `UPDATE` [756], `REFERENCES` [755], `ALTER` [753], `INDEX` [755], `DROP` [754], `CREATE VIEW` [754].



The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.TABLE_PRIVILEGES
SHOW GRANTS ...
```

## 19.27. The INFORMATION\_SCHEMA TRIGGERS Table

The TRIGGERS table provides information about triggers. You can see results only for databases and tables for which you have the TRIGGER [756] privilege.

INFORMATION_SCHEMA Name	SHOW Name	Remarks
TRIGGER_CATALOG		def
TRIGGER_SCHEMA		
TRIGGER_NAME	Trigger	
EVENT_MANIPULATION	Event	
EVENT_OBJECT_CATALOG		def
EVENT_OBJECT_SCHEMA		
EVENT_OBJECT_TABLE	Table	
ACTION_ORDER		0
ACTION_CONDITION		NULL
ACTION_STATEMENT	Statement	
ACTION_ORIENTATION		ROW
ACTION_TIMING	Timing	
ACTION_REFERENCE_OLD_TABLE		NULL
ACTION_REFERENCE_NEW_TABLE		NULL
ACTION_REFERENCE_OLD_ROW		OLD
ACTION_REFERENCE_NEW_ROW		NEW
CREATED		NULL (0)
SQL_MODE		MySQL extension
DEFINER		MySQL extension
CHARACTER_SET_CLIENT		MySQL extension
COLLATION_CONNECTION		MySQL extension
DATABASE_COLLATION		MySQL extension

### Notes:

- The TRIGGER\_SCHEMA and TRIGGER\_NAME columns contain the name of the database in which the trigger occurs and the trigger name, respectively.
- The EVENT\_MANIPULATION column contains one of the values 'INSERT', 'DELETE', or 'UPDATE'.
- As noted in Section 18.3, “Using Triggers”, every trigger is associated with exactly one table. The EVENT\_OBJECT\_SCHEMA and EVENT\_OBJECT\_TABLE columns contain the database in which this table occurs, and the table's name.
- The ACTION\_ORDER column contains the ordinal position of the trigger's action within the list of all similar triggers on the same table. Currently, this value is always 0, because it is not possible to have more than one trigger with the same EVENT\_MANIPULATION and ACTION\_TIMING on the same table.

- The `ACTION_STATEMENT` column contains the statement to be executed when the trigger is invoked. This is the same as the text displayed in the `Statement` column of the output from `SHOW TRIGGERS`. Note that this text uses UTF-8 encoding.
- The `ACTION_ORIENTATION` column always contains the value `'ROW'`.
- The `ACTION_TIMING` column contains one of the two values `'BEFORE'` or `'AFTER'`.
- The columns `ACTION_REFERENCE_OLD_ROW` and `ACTION_REFERENCE_NEW_ROW` contain the old and new column identifiers, respectively. This means that `ACTION_REFERENCE_OLD_ROW` always contains the value `'OLD'` and `ACTION_REFERENCE_NEW_ROW` always contains the value `'NEW'`.
- The `SQL_MODE` column shows the server SQL mode that was in effect at the time when the trigger was created (and thus which remains in effect for this trigger whenever it is invoked, *regardless of the current server SQL mode*). The possible range of values for this column is the same as that of the `sql_mode` [545] system variable. See Section 5.1.7, “Server SQL Modes”.
- The `DEFINER` column indicates who defined the trigger.
- `CHARACTER_SET_CLIENT` is the session value of the `character_set_client` [467] system variable when the trigger was created. `COLLATION_CONNECTION` is the session value of the `collation_connection` [469] system variable when the trigger was created. `DATABASE_COLLATION` is the collation of the database with which the trigger is associated.
- The following columns currently always contain `NULL`: `ACTION_CONDITION`, `ACTION_REFERENCE_OLD_TABLE`, `ACTION_REFERENCE_NEW_TABLE`, and `CREATED`.

Example, using the `ins_sum` trigger defined in Section 18.3, “Using Triggers”:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.TRIGGERS\G
***** 1. row *****
      TRIGGER_CATALOG: def
      TRIGGER_SCHEMA: test
      TRIGGER_NAME: ins_sum
      EVENT_MANIPULATION: INSERT
      EVENT_OBJECT_CATALOG: def
      EVENT_OBJECT_SCHEMA: test
      EVENT_OBJECT_TABLE: account
      ACTION_ORDER: 0
      ACTION_CONDITION: NULL
      ACTION_STATEMENT: SET @sum = @sum + NEW.amount
      ACTION_ORIENTATION: ROW
      ACTION_TIMING: BEFORE
      ACTION_REFERENCE_OLD_TABLE: NULL
      ACTION_REFERENCE_NEW_TABLE: NULL
      ACTION_REFERENCE_OLD_ROW: OLD
      ACTION_REFERENCE_NEW_ROW: NEW
      CREATED: NULL
      SQL_MODE:
      DEFINER: me@localhost
```

See also Section 13.7.5.39, “`SHOW TRIGGERS` Syntax”.

## 19.28. The `INFORMATION_SCHEMA.USER_PRIVILEGES` Table

The `USER_PRIVILEGES` table provides information about global privileges. This information comes from the `mysql.user` grant table.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>GRANTEE</code>		<code>'user_name'@'host_name'</code> value, MySQL extension
<code>TABLE_CATALOG</code>		<code>def</code> , MySQL extension
<code>PRIVILEGE_TYPE</code>		MySQL extension

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>IS_GRANTABLE</code>		MySQL extension

**Notes:**

- This is a nonstandard table. It takes its values from the `mysql.user` table.

## 19.29. The `INFORMATION_SCHEMA VIEWS` Table

The `VIEWS` table provides information about views in databases. You must have the `SHOW VIEW` [756] privilege to access this table.

<code>INFORMATION_SCHEMA</code> Name	<code>SHOW</code> Name	Remarks
<code>TABLE_CATALOG</code>		def
<code>TABLE_SCHEMA</code>		
<code>TABLE_NAME</code>		
<code>VIEW_DEFINITION</code>		
<code>CHECK_OPTION</code>		
<code>IS_UPDATABLE</code>		
<code>DEFINER</code>		
<code>SECURITY_TYPE</code>		
<code>CHARACTER_SET_CLIENT</code>		MySQL extension
<code>COLLATION_CONNECTION</code>		MySQL extension

**Notes:**

- The `VIEW_DEFINITION` column has most of what you see in the `Create Table` field that `SHOW CREATE VIEW` produces. Skip the words before `SELECT` and skip the words `WITH CHECK OPTION`. Suppose that the original statement was:

```
CREATE VIEW v AS
  SELECT s2,s1 FROM t
  WHERE s1 > 5
  ORDER BY s1
  WITH CHECK OPTION;
```

Then the view definition looks like this:

```
SELECT s2,s1 FROM t WHERE s1 > 5 ORDER BY s1
```

- The `CHECK_OPTION` column has a value of `NONE`, `CASCADE`, or `LOCAL`.
- MySQL sets a flag, called the view updatability flag, at `CREATE VIEW` time. The flag is set to `YES` (true) if `UPDATE` and `DELETE` (and similar operations) are legal for the view. Otherwise, the flag is set to `NO` (false). The `IS_UPDATABLE` column in the `VIEWS` table displays the status of this flag. It means that the server always knows whether a view is updatable. If the view is not updatable, statements such `UPDATE`, `DELETE`, and `INSERT` are illegal and will be rejected. (Note that even if a view is updatable, it might not be possible to insert into it; for details, refer to [Section 13.1.16](#), “`CREATE VIEW Syntax`”.)
- The `DEFINER` column indicates who defined the view. `SECURITY_TYPE` has a value of `DEFINER` or `INVOKER`.
- `CHARACTER_SET_CLIENT` is the session value of the `character_set_client` [467] system variable when the view was created. `COLLATION_CONNECTION` is the session value of the `collation_connection` [469] system variable when the view was created.

MySQL lets you use different `sql_mode` [545] settings to tell the server the type of SQL syntax to support. For example, you might use the `ANSI` [608] SQL mode to ensure MySQL correctly interprets the standard SQL concatenation operator, the double bar (`||`), in your queries. If you then create a view that concatenates items, you might worry that changing the `sql_mode` [545] setting to a value different from `ANSI` [608] could cause the view to become invalid. But this is not the case. No matter how you write out a view definition, MySQL always stores it the same way, in a canonical form. Here is an example that shows how the server changes a double bar concatenation operator to a `CONCAT()` [1178] function:

```
mysql> SET sql_mode = 'ANSI';
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE VIEW test.v AS SELECT 'a' || 'b' as coll;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT VIEW_DEFINITION FROM INFORMATION_SCHEMA.VIEWS
       -> WHERE TABLE_SCHEMA = 'test' AND TABLE_NAME = 'v';
+-----+
| VIEW_DEFINITION |
+-----+
| select concat('a','b') AS `coll` |
+-----+
1 row in set (0.00 sec)
```

The advantage of storing a view definition in canonical form is that changes made later to the value of `sql_mode` [545] will not affect the results from the view. However an additional consequence is that comments prior to `SELECT` are stripped from the definition by the server.

## 19.30. INFORMATION\_SCHEMA Tables for InnoDB

The InnoDB tables related to the InnoDB storage engine serve two purposes:

- You can monitor ongoing InnoDB activity, to detect inefficiencies before they turn into issues, or to troubleshoot performance and capacity issues that do occur. As your database becomes bigger and busier, running up against the limits of your hardware capacity, you monitor and tune these aspects to keep the database running smoothly. The monitoring information deals with:
  - InnoDB table compression, a feature whose use depends on a balance between I/O reduction, CPU usage, buffer pool management, and how much compression is possible for your data.
  - Transactions and locks, features that balance high performance for a single operation, against the ability to run multiple operations concurrently. (Transactions are the high-level, user-visible aspect of concurrency. Locks are the low-level mechanism that transactions use to avoid reading or writing unreliable data.)
- You can extract information about schema objects managed by InnoDB, using the `INNODB_SYS_*` tables. This information comes from the InnoDB data dictionary, which cannot be queried directly like regular InnoDB tables. Traditionally, you would get this type of information using the techniques from Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the InnoDB Monitors”, setting up InnoDB monitors and parsing the output from the `SHOW ENGINE INNODB STATUS` command. The InnoDB interface offers a simpler, familiar technique to access this data.

### 19.30.1. The INFORMATION\_SCHEMA INNODB\_CMP and INNODB\_CMP\_RESET Tables

The `INNODB_CMP` and `INNODB_CMP_RESET` tables contain status information on operations related to compressed InnoDB tables.

**Table 19.1. Columns of `INNODB_CMP` and `INNODB_CMP_RESET`**

Column name	Description
<code>PAGE_SIZE</code>	Compressed page size in bytes.

Column name	Description
<code>COMPRESS_OPS</code>	Number of times a B-tree page of the size <code>PAGE_SIZE</code> has been compressed. Pages are compressed whenever an empty page is created or the space for the uncompressed modification log runs out.
<code>COMPRESS_OPS_OK</code>	Number of times a B-tree page of the size <code>PAGE_SIZE</code> has been successfully compressed. This count should never exceed <code>COMPRESS_OPS</code> .
<code>COMPRESS_TIME</code>	Total time in seconds spent in attempts to compress B-tree pages of the size <code>PAGE_SIZE</code> .
<code>UNCOMPRESS_OPS</code>	Number of times a B-tree page of the size <code>PAGE_SIZE</code> has been uncompressed. B-tree pages are uncompressed whenever compression fails or at first access when the uncompressed page does not exist in the buffer pool.
<code>UNCOMPRESS_TIME</code>	Total time in seconds spent in uncompressing B-tree pages of the size <code>PAGE_SIZE</code> .

**Notes:**

- Use these tables to measure the effectiveness of `InnoDB` table `compression` in your database.
- You must have the `PROCESS` privilege to query this table.
- For usage information, see [Section 5.4.6.4, “Monitoring Compression at Runtime”](#) and [Section 14.2.4.3.1.3, “Using the Compression Information Schema Tables”](#). For general information about `InnoDB` table compression, see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#).

## 19.30.2. The `INFORMATION_SCHEMA INNODB_CMP_PER_INDEX` and `INNODB_CMP_PER_INDEX_RESET` Tables

The `INNODB_CMP_PER_INDEX` and `INNODB_CMP_PER_INDEX_RESET` tables contain status information on operations related to `compressed` `InnoDB` tables and indexes, with separate statistics for each combination of database, table, and index, to help you evaluate the performance and usefulness of compression for specific tables.

For a compressed `InnoDB` table, both the table data and all the `secondary indexes` are compressed. In this context, the table data is treated as just another index, one that happens to contain all the columns: the `clustered index`.

**Table 19.2. Columns of `INNODB_CMP_PER_INDEX` and `INNODB_CMP_PER_INDEX_RESET`**

Column name	Description
<code>DATABASE_NAME</code>	Database containing the applicable table.
<code>TABLE_NAME</code>	Table to monitor for compression statistics.
<code>INDEX_NAME</code>	Index to monitor for compression statistics.
<code>COMPRESS_OPS</code>	Number of compression operations attempted. <code>Pages</code> are compressed whenever an empty page is created or the space for the uncompressed modification log runs out.
<code>COMPRESS_OPS_OK</code>	Number of successful compression operations. Subtract from the <code>COMPRESS_OPS</code> value to get the number of <code>compression failures</code> . Divide by the <code>COMPRESS_OPS</code> value to get the percentage of compression failures.
<code>COMPRESS_TIME</code>	Total amount of CPU time, in seconds, used for compressing data in this index.

Column name	Description
<code>UNCOMPRESS_OPS</code>	Number of uncompression operations performed. Compressed <code>InnoDB</code> pages are uncompressed whenever compression <code>fails</code> , or the first time a compressed page is accessed in the <code>buffer pool</code> and the uncompressed page does not exist.
<code>UNCOMPRESS_TIME</code>	Total amount of CPU time, in seconds, used for uncompressing data in this index.

**Notes:**

- Use these tables to measure the effectiveness of `InnoDB` table `compression` for specific tables, indexes, or both.
- You must have the `PROCESS` privilege to query these tables.
- Because collecting separate measurements for every index imposes substantial performance overhead, enable the `innodb_cmp_per_index_enabled` [1712] configuration option before performing the operations on compressed tables that you want to monitor.
- For usage information, see [Section 5.4.6.4, “Monitoring Compression at Runtime”](#) and [Section 14.2.4.3.1.3, “Using the Compression Information Schema Tables”](#). For general information about `InnoDB` table compression, see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#).

### 19.30.3. The `INFORMATION_SCHEMA INNODB_CMPMEM` and `INNODB_CMPMEM_RESET` Tables

The `INNODB_CMPMEM` and `INNODB_CMPMEM_RESET` tables contain status information on compressed pages within the `InnoDB buffer pool`.

**Table 19.3. Columns of `INNODB_CMPMEM` and `INNODB_CMPMEM_RESET`**

Column name	Description
<code>PAGE_SIZE</code>	Block size in bytes. Each record of this table describes blocks of this size.
<code>PAGES_USED</code>	Number of blocks of the size <code>PAGE_SIZE</code> that are currently in use.
<code>PAGES_FREE</code>	Number of blocks of the size <code>PAGE_SIZE</code> that are currently available for allocation. This column shows the external fragmentation in the memory pool. Ideally, these numbers should be at most 1.
<code>RELOCATION_OPS</code>	Number of times a block of the size <code>PAGE_SIZE</code> has been relocated. The buddy system can relocate the allocated “buddy neighbor” of a freed block when it tries to form a bigger freed block. Reading from the table <code>INNODB_CMPMEM_RESET</code> resets this count.
<code>RELOCATION_TIME</code>	Total time in microseconds spent in relocating blocks of the size <code>PAGE_SIZE</code> . Reading from the table <code>INNODB_CMPMEM_RESET</code> resets this count.

**Notes:**

- Use these tables to measure the effectiveness of `InnoDB` table `compression` in your database.
- You must have the `PROCESS` privilege to query this table.
- For usage information, see [Section 5.4.6.4, “Monitoring Compression at Runtime”](#) and [Section 14.2.4.3.1.3, “Using the Compression Information Schema Tables”](#). For general information about `InnoDB` table compression, see [Section 5.4.6, “Working with InnoDB Compressed Tables”](#).

### 19.30.4. The `INFORMATION_SCHEMA INNODB_TRX` Table

The `INNODB_TRX` table contains information about every transaction currently executing inside `InnoDB`, including whether the transaction is waiting for a lock, when the transaction started, and the SQL statement the transaction is executing.

**Table 19.4. INNODB\_TRX Columns**

Column name	Description
<code>TRX_ID</code>	Unique transaction ID number, internal to <code>InnoDB</code> . (Starting in MySQL 5.6, these IDs are not created for transactions that are read-only and non-locking. See <a href="#">Section 14.2.4.2.3, “Optimizations for Read-Only Transactions”</a> for details.)
<code>TRX_WEIGHT</code>	The weight of a transaction, reflecting (but not necessarily the exact count of) the number of rows altered and the number of rows locked by the transaction. To resolve a deadlock, <code>InnoDB</code> selects the transaction with the smallest weight as the “victim” to rollback. Transactions that have changed non-transactional tables are considered heavier than others, regardless of the number of altered and locked rows.
<code>TRX_STATE</code>	Transaction execution state. One of <code>RUNNING</code> , <code>LOCK WAIT</code> , <code>ROLLING BACK</code> or <code>COMMITTING</code> .
<code>TRX_STARTED</code>	Transaction start time.
<code>TRX_REQUESTED_LOCK_ID</code>	ID of the lock the transaction is currently waiting for (if <code>TRX_STATE</code> is <code>LOCK WAIT</code> , otherwise <code>NULL</code> ). Details about the lock can be found by joining with <code>INNODB_LOCKS</code> on <code>LOCK_ID</code> .
<code>TRX_WAIT_STARTED</code>	Time when the transaction started waiting on the lock (if <code>TRX_STATE</code> is <code>LOCK WAIT</code> , otherwise <code>NULL</code> ).
<code>TRX_MYSQL_THREAD_ID</code>	MySQL thread ID. Can be used for joining with <code>PROCESSLIST</code> on <code>ID</code> . See <a href="#">Section 14.2.4.3.4.3, “Possible Inconsistency with PROCESSLIST”</a> .
<code>TRX_QUERY</code>	The SQL query that is being executed by the transaction.
<code>TRX_OPERATION_STATE</code>	The transaction's current operation, or <code>NULL</code> .
<code>TRX_TABLES_IN_USE</code>	The number of <code>InnoDB</code> tables used while processing the current SQL statement of this transaction.
<code>TRX_TABLES_LOCKED</code>	Number of <code>InnoDB</code> tables that currently have any locks. (Because these are row locks, not table locks, the tables can usually still be read from and written to by multiple transactions, despite some rows being locked.)
<code>TRX_LOCK_STRUCTS</code>	The number of locks reserved by the transaction.
<code>TRX_LOCK_MEMORY_BYTES</code>	Total size taken up by the lock structures of this transaction in memory.
<code>TRX_ROWS_LOCKED</code>	Approximate number of rows locked by this transaction. The value might include delete-marked rows that are physically present but not visible to the transaction.
<code>TRX_ROWS_MODIFIED</code>	The number of modified and inserted rows in this transaction.
<code>TRX_CONCURRENCY_TICKETS</code>	A value indicating how much work the current transaction can do before being swapped out, as specified by the <code>innodb_concurrency_tickets</code> option.
<code>TRX_ISOLATION_LEVEL</code>	The isolation level of the current transaction.
<code>TRX_UNIQUE_CHECKS</code>	Whether unique checks are turned on or off for the current transaction. (They might be turned off during a bulk data load, for example.)



Column name	Description
<code>TRX_FOREIGN_KEY_CHECKS</code>	Whether foreign key checks are turned on or off for the current transaction. (They might be turned off during a bulk data load, for example.)
<code>TRX_LAST_FOREIGN_KEY_ERROR</code>	Detailed error message for last FK error, or <code>NULL</code> .
<code>TRX_ADAPTIVE_HASH_LATCHED</code>	Whether or not the adaptive hash index is locked by the current transaction. (Only a single transaction at a time can modify the adaptive hash index.)
<code>TRX_ADAPTIVE_HASH_TIMEOUT</code>	Whether to relinquish the search latch immediately for the adaptive hash index, or reserve it across calls from MySQL. When there is no AHI contention, this value remains zero and statements reserve the latch until they finish. During times of contention, it counts down to zero, and statements release the latch immediately after each row lookup.
<code>TRX_IS_READ_ONLY</code>	A value of 1 indicates the transaction is read-only. (5.6.4 and up.)
<code>TRX_AUTO_COMMIT_NON_LOCKING</code>	A value of 1 indicates the transaction is a <code>SELECT</code> statement that does not use the <code>FOR UPDATE</code> or <code>LOCK IN SHARED MODE</code> clauses, and is executing with the <code>autocommit</code> setting turned on so that the transaction will only contain this one statement. (5.6.4 and up.) When this column and <code>TRX_IS_READ_ONLY</code> are both 1, <code>InnoDB</code> optimizes the transaction to reduce the overhead associated with transactions that change table data.

**Notes:**

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load.
- You must have the `PROCESS` privilege to query this table.
- For usage information, see [Section 14.2.4.3.2.4, “Using the Transaction Information Schema Tables”](#).

### 19.30.5. The `INFORMATION_SCHEMA INNODB_LOCKS` Table

The `INNODB_LOCKS` table contains information about each lock that an `InnoDB` transaction has requested but not yet acquired, and each lock that a transaction holds that is blocking another transaction.

**Table 19.5. `INNODB_LOCKS` Columns**

Column name	Description
<code>LOCK_ID</code>	Unique lock ID number, internal to <code>InnoDB</code> . Treat it as an opaque string. Although <code>LOCK_ID</code> currently contains <code>TRX_ID</code> , the format of the data in <code>LOCK_ID</code> is not guaranteed to remain the same in future releases. Do not write programs that parse the <code>LOCK_ID</code> value.
<code>LOCK_TRX_ID</code>	ID of the transaction holding this lock. Details about the transaction can be found by joining with <code>INNODB_TRX</code> on <code>TRX_ID</code> .
<code>LOCK_MODE</code>	Mode of the lock. One of <code>S</code> , <code>X</code> , <code>IS</code> , <code>IX</code> , <code>S_GAP</code> , <code>X_GAP</code> , <code>IS_GAP</code> , <code>IX_GAP</code> , or <code>AUTO_INC</code> for shared, exclusive, intention shared, intention exclusive row locks, shared and exclusive gap locks, intention shared and intention exclusive gap locks, and auto-increment table level lock, respectively. Refer to the sections <a href="#">Section 14.2.3.2, “InnoDB Lock Modes”</a> and <a href="#">Section 14.2.3.1, “The InnoDB Transaction Model and Locking”</a> for information on <code>InnoDB</code> locking.
<code>LOCK_TYPE</code>	Type of the lock. One of <code>RECORD</code> or <code>TABLE</code> for record (row) level or table level locks, respectively.

Column name	Description
<code>LOCK_TABLE</code>	Name of the table that has been locked or contains locked records.
<code>LOCK_INDEX</code>	Name of the index if <code>LOCK_TYPE= 'RECORD'</code> , otherwise <code>NULL</code> .
<code>LOCK_SPACE</code>	Tablespace ID of the locked record if <code>LOCK_TYPE= 'RECORD'</code> , otherwise <code>NULL</code> .
<code>LOCK_PAGE</code>	Page number of the locked record if <code>LOCK_TYPE= 'RECORD'</code> , otherwise <code>NULL</code> .
<code>LOCK_REC</code>	Heap number of the locked record within the page if <code>LOCK_TYPE= 'RECORD'</code> , otherwise <code>NULL</code> .
<code>LOCK_DATA</code>	Primary key of the locked record if <code>LOCK_TYPE= 'RECORD'</code> , otherwise <code>NULL</code> . This column contains the value(s) of the primary key column(s) in the locked row, formatted as a valid SQL string (ready to be copied to SQL commands). If there is no primary key then the <code>InnoDB</code> internal unique row ID number is used. When the page containing the locked record is not in the buffer pool (in the case that it was paged out to disk while the lock was held), <code>InnoDB</code> does not fetch the page from disk, to avoid unnecessary disk operations. Instead, <code>LOCK_DATA</code> is set to <code>NULL</code> .

**Notes:**

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load.
- You must have the `PROCESS` privilege to query this table.
- For usage information, see [Section 14.2.4.3.2.4, “Using the Transaction Information Schema Tables”](#).

**19.30.6. The `INFORMATION_SCHEMA INNODB_LOCK_WAITS` Table**

The `INNODB_LOCK_WAITS` table contains one or more rows for each blocked `InnoDB` transaction, indicating the lock it has requested and any locks that are blocking that request.

**Table 19.6. `INNODB_LOCK_WAITS` Columns**

Column name	Description
<code>REQUESTING_TRX_ID</code>	ID of the requesting transaction.
<code>REQUESTED_LOCK_ID</code>	ID of the lock for which a transaction is waiting. Details about the lock can be found by joining with <code>INNODB_LOCKS</code> on <code>LOCK_ID</code> .
<code>BLOCKING_TRX_ID</code>	ID of the blocking transaction.
<code>BLOCKING_LOCK_ID</code>	ID of a lock held by a transaction blocking another transaction from proceeding. Details about the lock can be found by joining with <code>INNODB_LOCKS</code> on <code>LOCK_ID</code> .

**Notes:**

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load.
- You must have the `PROCESS` privilege to query this table.
- For usage information, see [Section 14.2.4.3.2.4, “Using the Transaction Information Schema Tables”](#).

**19.30.7. The `INFORMATION_SCHEMA INNODB_SYS_TABLES` Table**

The `INNODB_SYS_TABLES` table provides status information about `InnoDB` tables, equivalent to the information from the `SYS_TABLES` table in the `InnoDB` data dictionary.

**Table 19.7. `INNODB_SYS_TABLES` Columns**

Column name	Description
<code>TABLE_ID</code>	An identifier for each <code>InnoDB</code> table that is unique across all databases in the instance.
<code>NAME</code>	The name of the table. Preceded by the database name where appropriate, for example <code>test/t1</code> . <code>InnoDB</code> system table names are in all uppercase. Names of databases and user tables are in the same case as they were originally defined, possibly influenced by the <code>lower_case_table_names</code> [501] setting.
<code>FLAG</code>	0 = <code>InnoDB</code> system table, 1 = user table.
<code>N_COLS</code>	The number of columns in the table.
<code>SPACE</code>	An identifier for the tablespace where the table resides. 0 means the <code>InnoDB system tablespace</code> . Any other number represents a table created in <code>file-per-table</code> mode with a separate <code>.ibd</code> file. This identifier stays the same after a <code>TRUNCATE TABLE</code> statement. Other than the zero value, this identifier is unique for tables across all the databases in the instance.

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

**19.30.8. The `INFORMATION_SCHEMA INNODB_SYS_INDEXES` Table**

The `INNODB_SYS_INDEXES` table provides status information about `InnoDB` indexes, equivalent to the information from the `SYS_INDEXES` table in the `InnoDB` data dictionary.

**Table 19.8. `INNODB_SYS_INDEXES` Columns**

Column name	Description
<code>INDEX_ID</code>	An identifier for each index that is unique across all the databases in an instance.
<code>NAME</code>	The name of the index. User-created indexes have names in all lowercase. Indexes created implicitly by <code>InnoDB</code> have names in all lowercase. The index names are not necessarily unique. Indexes created implicitly by <code>InnoDB</code> have consistent names: <code>PRIMARY</code> for a primary key index, <code>GEN_CLUST_INDEX</code> for the index representing a primary key when one is not specified, <code>ID_IND</code> , <code>FOR_IND</code> for validating a foreign key constraint, and <code>REF_IND</code> .
<code>TABLE_ID</code>	An identifier representing the table associated with the index; the same value from <code>INNODB_SYS_TABLES.TABLE_ID</code> .
<code>TYPE</code>	A numeric identifier signifying the kind of index, in the range 0..3.
<code>N_FIELDS</code>	The number of columns in the index key. For the <code>GEN_CLUST_INDEX</code> indexes, this value is 0 because the index is created using an artificial value rather than a real table column.
<code>PAGE_NO</code>	
<code>SPACE</code>	An identifier for the tablespace where the index resides. 0 means the <code>InnoDB system tablespace</code> . Any other number represents a table created in <code>file-per-table</code> mode with a separate <code>.ibd</code> file. This identifier stays the same after a <code>TRUNCATE TABLE</code> statement. Because all indexes

Column name	Description
	for a table reside in the same tablespace as the table, this value is not necessarily unique.

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

### 19.30.9. The `INFORMATION_SCHEMA INNODB_SYS_COLUMNS` Table

The `INNODB_SYS_COLUMNS` table provides status information about `InnoDB` table columns, equivalent to the information from the `SYS_COLUMNS` table in the `InnoDB` data dictionary.

**Table 19.9. `INNODB_SYS_COLUMNS` Columns**

Column name	Description
<code>TABLE_ID</code>	An identifier representing the table associated with the column; the same value from <code>INNODB_SYS_TABLES.TABLE_ID</code> .
<code>NAME</code>	The name of each column in each table. These names can be uppercase or lowercase depending on the <code>lower_case_table_names</code> [501] setting. There are no special system-reserved names for columns.
<code>POS</code>	The ordinal position of the column within the table, starting from 0 and incrementing sequentially. When a column is dropped, the remaining columns are reordered so that the sequence has no gaps.
<code>MTYPE</code>	A numeric identifier for the column type. 1 = <code>VARCHAR</code> , 2 = <code>CHAR</code> , 3 = <code>FIXBINARY</code> , 4 = <code>BINARY</code> , 5 = <code>BLOB</code> , 6 = <code>INT</code> , 7 = <code>SYS_CHILD</code> , 8 = <code>SYS</code> , 9 = <code>FLOAT</code> , 10 = <code>DOUBLE</code> , 11 = <code>DECIMAL</code> , 12 = <code>VARMYSQL</code> , 13 = <code>MYSQL</code> .
<code>PRTYPE</code>	The <code>InnoDB</code> “precise type”, a binary value with bits representing MySQL data type, character set code, and nullability.
<code>LEN</code>	The column length, for example 4 for <code>INT</code> and 8 for <code>BIGINT</code> . For character columns in multi-byte character sets, this length value is the maximum length in bytes needed to represent a definition such as <code>VARCHAR(N)</code> ; that is, it might be $2*N$ , $3*N$ , and so on depending on the character encoding.

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

### 19.30.10. The `INFORMATION_SCHEMA INNODB_SYS_FIELDS` Table

The `INNODB_SYS_FIELDS` table provides status information about the key columns (fields) of `InnoDB` indexes, equivalent to the information from the `SYS_FIELDS` table in the `InnoDB` data dictionary.

**Table 19.10. `INNODB_SYS_FIELDS` Columns**

Column name	Description
<code>INDEX_ID</code>	An identifier for the index associated with this key field, using the same value as in <code>INNODB_SYS_INDEXES.INDEX_ID</code> .

Column name	Description
<code>NAME</code>	The name of the original column from the table, using the same value as in <code>INNODB_SYS_COLUMNS.NAME</code> .
<code>POS</code>	The ordinal position of the key field within the index, starting from 0 and incrementing sequentially. When a column is dropped, the remaining columns are reordered so that the sequence has no gaps.

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

### 19.30.11. The `INFORMATION_SCHEMA INNODB_SYS_FOREIGN` Table

The `INNODB_SYS_FOREIGN` table provides status information about `InnoDB` foreign keys, equivalent to the information from the `SYS_FOREIGN` table in the `InnoDB` data dictionary.

**Table 19.11. `INNODB_SYS_FOREIGN` Columns**

Column name	Description
<code>ID</code>	The name (not a numeric value) of the foreign key index. Preceded by the database name, for example, <code>test/products_fk</code> .
<code>FOR_NAME</code>	The name of the <code>child table</code> in this foreign key relationship.
<code>REF_NAME</code>	The name of the <code>parent table</code> in this foreign key relationship.
<code>N_COLS</code>	The number of columns in the foreign key index.
<code>TYPE</code>	A collection of bit flags with information about the foreign key column, ORed together. 1 = <code>ON DELETE CASCADE</code> , 2 = <code>ON UPDATE SET NULL</code> , 4 = <code>ON UPDATE CASCADE</code> , 8 = <code>ON UPDATE SET NULL</code> , 16 = <code>ON DELETE NO ACTION</code> , 32 = <code>ON UPDATE NO ACTION</code> .

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

### 19.30.12. The `INFORMATION_SCHEMA INNODB_SYS_FOREIGN_COLS` Table

The `INNODB_SYS_FOREIGN_COLS` table provides status information about the columns of `InnoDB` foreign keys, equivalent to the information from the `SYS_FOREIGN_COLS` table in the `InnoDB` data dictionary.

**Table 19.12. `INNODB_SYS_FOREIGN_COLS` Columns**

Column name	Description
<code>ID</code>	The foreign key index associated with this index key field, using the same value as <code>INNODB_SYS_FOREIGN.ID</code> .
<code>FOR_COL_NAME</code>	The name of the associated column in the child table.
<code>REF_COL_NAME</code>	The name of the associated column in the parent table.
<code>POS</code>	The ordinal position of this key field within the foreign key index, starting from 0.

**Notes:**

- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

**19.30.13. The `INFORMATION_SCHEMA INNODB_SYS_TABLESTATS` View**

The `INNODB_SYS_TABLESTATS` view provides status information about performance statistics for `InnoDB` tables. These statistics represent low-level information used by the MySQL optimizer to calculate which index to use when querying an `InnoDB` table. This information is derived from in-memory data structures rather than corresponding to data stored on disk.

`InnoDB` tables are represented in this view if they have been opened since the last server restart, and not aged out of the table cache. Tables for which persistent stats are available are always represented in this view.

**Table 19.13. `INNODB_SYS_TABLESTATS` Columns**

Column name	Description
<code>TABLE_ID</code>	An identifier representing the table for which statistics are available, using the same value as <code>INNODB_SYS_TABLES.TABLE_ID</code> .
<code>NAME</code>	The name of the table, using the same value as <code>INNODB_SYS_TABLES.NAME</code> .
<code>STATS_INITIALIZED</code>	The value is <code>Initialized</code> if the statistics are already collected, <code>Uninitialized</code> if not.
<code>NUM_ROWS</code>	The current estimated number of rows in the table. Updated after each DML operation. Could be imprecise if uncommitted transactions are inserting into or deleting from the table.
<code>CLUST_INDEX_SIZE</code>	Number of pages on disk that store the clustered index, which holds the <code>InnoDB</code> table data in primary key order. This value might be null if no statistics are collected yet for the table.
<code>OTHER_INDEX_SIZE</code>	Number of pages on disk that store all secondary indexes for the table. This value might be null if no statistics are collected yet for the table.
<code>MODIFIED_COUNTER</code>	The number of rows modified by DML operations, such as <code>INSERT</code> , <code>UPDATE</code> , <code>DELETE</code> , and also cascade operations from foreign keys.
<code>AUTOINC</code>	
<code>REF_COUNT</code>	When this counter reaches zero, the table metadata can be evicted from the table cache.

**Notes:**

- This table is primarily useful for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

**19.30.14. The `INFORMATION_SCHEMA INNODB_BUFFER_PAGE` Table**

The `INNODB_BUFFER_PAGE` table holds information about each `page` in the `InnoDB` buffer pool.

**Table 19.14. `INNODB_BUFFER_PAGE` Columns**

Column name	Description
<code>POOL_ID</code>	Buffer Pool ID. An identifier to distinguish between multiple buffer pool instances.
<code>BLOCK_ID</code>	Buffer Pool Block ID.
<code>SPACE</code>	Tablespace ID. Uses the same value as in <code>INNODB_SYS_TABLES.SPACE</code> .
<code>PAGE_NUMBER</code>	Page number.
<code>PAGE_TYPE</code>	Page type string. One of <code>allocated</code> (Freshly allocated page), <code>index</code> (B-tree node), <code>undo_log</code> (Undo log page), <code>inode</code> (Index node), <code>ibuf_free_list</code> (Insert buffer free list), <code>ibuf_bitmap</code> (Insert buffer bitmap), <code>system</code> (System page), <code>trx_system</code> (Transaction system data), <code>file_space_header</code> (File space header), <code>extent_descriptor</code> (Extent descriptor page), <code>blob</code> (Uncompressed BLOB page), <code>compressed_blob</code> (First compressed BLOB page), <code>compressed_blob2</code> (Subsequent comp BLOB page), <code>unknown</code> (unknown).
<code>FLUSH_TYPE</code>	Flush type.
<code>FIX_COUNT</code>	Number of threads using this block within the buffer pool. When zero, the block is eligible to be evicted.
<code>IS_HASHED</code>	Whether hash index has been built on this page.
<code>NEWEST_MODIFICATION</code>	Log Sequence Number of the youngest modification.
<code>OLDEST_MODIFICATION</code>	Log Sequence Number of the oldest modification.
<code>ACCESS_TIME</code>	Time of first access.
<code>TABLE_NAME</code>	Name of the table the page belongs to.
<code>INDEX_NAME</code>	Name of the index the page belongs to. It can be the name of a clustered index or a secondary index.
<code>NUMBER_RECORDS</code>	Number of records within the page.
<code>DATA_SIZE</code>	Sum of the sizes of the records.
<code>COMPRESSED_SIZE</code>	Compressed page size. Null for pages that are not compressed.
<code>PAGE_STATE</code>	Page state. A page with valid data has one of the following states: <code>FILE_PAGE</code> (buffers a page of data from a file), <code>MEMORY</code> (buffers a page from an in-memory object), Other possible states (managed by <code>InnoDB</code> ) are: <code>null</code> , <code>READY_FOR_USE</code> , <code>NOT_USED</code> , <code>REMOVE_HASH</code> .
<code>IO_FIX</code>	Specifies whether any I/O is pending for this page: <code>IO_NONE</code> = no pending I/O, <code>IO_READ</code> = read pending, <code>IO_WRITE</code> = write pending.
<code>IS_OLD</code>	<code>bpage-&gt;old</code> .
<code>FREE_PAGE_CLOCK</code>	<code>bpage-&gt;freed_page_clock</code> .

**Notes:**

- This table is primarily useful for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

**19.30.15. The `INFORMATION_SCHEMA INNODB_BUFFER_PAGE_LRU` Table**



The `INNODB_BUFFER_PAGE_LRU` table holds information about the pages in the `InnoDB` buffer pool, in particular how they are ordered in the LRU list that determines which pages to `evict` from the buffer pool when it becomes full.

The definition for this page is the same as for `INNODB_BUFFER_PAGE`, except this table has an `LRU_POSITION` column instead of `BLOCK_ID`.

### Notes:

- This table is primarily useful for expert-level performance monitoring, or when developing performance-related extensions for MySQL.

Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.

- Querying this table can require MySQL to allocate a large block of contiguous memory, more than 64 bytes time the number of active pages in the buffer pool. This allocation could potentially cause an out-of-memory error, especially for systems with multi-gigabyte buffer pools.
- Querying this table requires MySQL to lock the data structure representing the buffer pool while traversing the LRU list, which can reduce concurrency, especially for systems with multi-gigabyte buffer pools.

## 19.30.16. The `INFORMATION_SCHEMA INNODB_BUFFER_POOL_STATS` Table

The `INNODB_BUFFER_POOL_STATS` table represents the same kinds of `InnoDB` buffer pool information as the output from `SHOW ENGINE INNODB STATUS`.

The idea of making pages in the buffer pool “young” or “not young” refers to transferring them between the `sublists` at the head and tail of the buffer pool data structure. Pages made “young” take longer to age out of the buffer pool, while pages made “not young” are moved much closer to the point of `eviction`.

Column meanings in this table are intended to be self-explanatory for developers who formerly parsed the `SHOW ENGINE INNODB STATUS` to get these values.

**Table 19.15. `INNODB_BUFFER_POOL_STATS` Columns**

Column name	Description
<code>POOL_ID</code>	
<code>POOL_SIZE</code>	
<code>FREE_BUFFERS</code>	
<code>DATABASE_PAGES</code>	
<code>OLD_DATABASE_PAGES</code>	
<code>MODIFIED_DATABASE_PAGES</code>	
<code>PENDING_DECOMPRESS</code>	
<code>PENDING_READS</code>	
<code>PENDING_FLUSH_LRU</code>	
<code>PENDING_FLUSH_LIST</code>	
<code>PAGES_MADE_YOUNG</code>	
<code>PAGES_NOT_MADE_YOUNG</code>	
<code>PAGES_MADE_YOUNG_RATE</code>	
<code>PAGES_MADE_NOT_YOUNG_RATE</code>	

Column name	Description
<code>NUMBER_PAGES_READ</code>	
<code>NUMBER_PAGES_CREATED</code>	
<code>NUMBER_PAGES_WRITTEN</code>	
<code>PAGES_READ_RATE</code>	
<code>PAGES_CREATE_RATE</code>	
<code>PAGES_WRITTEN_RATE</code>	
<code>NUMBER_PAGES_GET</code>	
<code>HIT_RATE</code>	
<code>YOUNG_MAKE_PER_THOUSAND_GETS</code>	
<code>NOT_YOUNG_MAKE_PER_THOUSAND_GETS</code>	
<code>NUMBER_PAGES_READ_AHEAD</code>	
<code>NUMBER_READ_AHEAD_EVICTED</code>	
<code>READ_AHEAD_RATE</code>	
<code>READ_AHEAD_EVICTED_RATE</code>	
<code>LRU_IO_TOTAL</code>	
<code>LRU_IO_CURRENT</code>	
<code>UNCOMPRESS_TOTAL</code>	
<code>UNCOMPRESS_CURRENT</code>	

**Notes:**

- This table is primarily useful for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- Since the `INFORMATION_SCHEMA` is a general-purpose way to monitor the MySQL server, use this table rather than the corresponding `InnoDB` system table for any new monitoring application you develop.
- You must have the `PROCESS` privilege to query this table.

### 19.30.17. The `INFORMATION_SCHEMA INNODB_METRICS` Table

This `INFORMATION_SCHEMA` table presents a wide variety of `InnoDB` performance information, complementing the specific focus areas of the `PERFORMANCE_SCHEMA` tables for `InnoDB`. With simple queries, you can check the overall health of the system. With more detailed queries, you can diagnose issues such as performance bottlenecks, resource shortages, and application issues.

Each monitor represents a point within the `InnoDB` source code that is instrumented to gather counter information. Each counter can be started, stopped, and reset. You can also perform these actions for a group of counters using their common module name.

**Table 19.16. `INNODB_METRICS` Columns**

Column name	Description
<code>NAME</code>	Unique name for the counter.
<code>SUBSYSTEM</code>	The aspect of <code>InnoDB</code> that the metric applies to. See the list following the table for the corresponding module names to use with the <code>SET GLOBAL</code> syntax.
<code>COUNT</code>	Value since the counter is enabled.
<code>MAX_COUNT</code>	Maximum value since the counter is enabled.

Column name	Description
<a href="#">MIN_COUNT</a>	Minimum value since the counter is enabled.
<a href="#">AVG_COUNT</a>	Average value since the counter is enabled.
<a href="#">COUNT_RESET</a>	Counter value since it was last reset. (The <a href="#">_RESET</a> fields act like the lap counter on a stopwatch: you can measure the activity during some time interval, while the cumulative figures are still available in the <a href="#">COUNT</a> , <a href="#">MAX_COUNT</a> , and so on fields.)
<a href="#">MAX_COUNT_RESET</a>	Maximum counter value since it was last reset.
<a href="#">MIN_COUNT_RESET</a>	Minimum counter value since it was last reset.
<a href="#">AVG_COUNT_RESET</a>	Average counter value since it was last reset.
<a href="#">TIME_ENABLED</a>	Timestamp of last start.
<a href="#">TIME_DISABLED</a>	Timestamp of last stop.
<a href="#">TIME_ELAPSED</a>	Elapsed time in seconds since the counter started.
<a href="#">TIME_RESET</a>	Timestamp of last stop.
<a href="#">STATUS</a>	Whether the counter is still running (Y) or stopped (N).
<a href="#">TYPE</a>	Whether the item is a cumulative counter, or measures the current value of some resource.
<a href="#">COMMENT</a>	Additional description.

**Notes:**

- You must have the [PROCESS](#) privilege to query this table.
- By default, relatively little data is collected. To start, stop, and reset counters, you set one of the the configuration options [innodb\\_monitor\\_enable](#) [1737], [innodb\\_monitor\\_disable](#) [1736], [innodb\\_monitor\\_reset](#) [1737], or [innodb\\_monitor\\_reset\\_all](#) [1737], using the name of the counter, the name of the module, a wildcard match for such a name using the “%” character, or the special keyword `all`. For example:

```
-- Turn on a counter:
set global innodb_monitor_enable = counter_name;
-- Turn off a counter:
set global innodb_monitor_disable = counter_name;
-- Reset a counter:
set global innodb_monitor_reset = counter_name;
-- Reset all the counter values:
set global innodb_monitor_reset_all = counter_name;
-- Turn on a group of counters:
set global innodb_monitor_enable = module_name;
-- Turn off a group of counters:
set global innodb_monitor_disable = module_name;
-- Turn on monitor "dml_num_inserts":
set global innodb_monitor_enable = dml_num_inserts;
-- Reset all values for "dml_num_inserts":
set global innodb_monitor_reset_all = dml_num_inserts;
-- Reset the counters for all dml monitors
set global innodb_monitor_reset = "dml_%";
-- Turn on all monitors for the transaction subsystem
set global innodb_monitor_enable = module_trx;
-- Turn off all monitors
set global innodb_monitor_disable = all;
```

- The module names correspond to, but are not identical to, the values from the [SUBSYSTEM](#) column. Here are the values you can use for `module_name` with the [innodb\\_monitor\\_enable](#) [1737] and related configuration options, along with the corresponding [SUBSYSTEM](#) names:
  - `module_metadata` (subsystem = `metadata`)
  - `module_lock` (subsystem = `lock`)

- `module_buffer` (subsystem = `buffer`)
  - `module_buf_page` (subsystem = `buffer_page_io`)
  - `module_os` (subsystem = `os`)
  - `module_trx` (subsystem = `transaction`)
  - `module_purge` (subsystem = `purge`)
  - `module_compress` (subsystem = `compression`)
  - `module_file` (subsystem = `file_system`)
  - `module_index` (subsystem = `index`)
  - `module_adaptive_hash` (subsystem = `adaptive_hash_index`)
  - `module_ibuf_system` (subsystem = `change_buffer`)
  - `module_srv` (subsystem = `server`)
  - `module_ddl` (subsystem = `ddl`)
  - `module_dml` (subsystem = `dml`)
  - `module_log` (subsystem = `recovery`)
  - `module_icp` (subsystem = `icp`)
- Because each counter imposes some degree of runtime overhead on the server, typically you enable more counters on test and development servers during experimentation and benchmarking, and only enable counters on production servers to diagnose known issues or monitor aspects that are likely to be bottlenecks for a particular server and workload.
  - The items represented in the `innodb_metrics` table are subject to change, so for the most up-to-date list, query a running MySQL server. As of MySQL 5.6.10, the items include:

```
mysql: select name, subsystem from information_schema.innodb_metrics order by name;
```

name	subsystem
adaptive_hash_pages_added	adaptive_hash_index
adaptive_hash_pages_removed	adaptive_hash_index
adaptive_hash_rows_added	adaptive_hash_index
adaptive_hash_rows_deleted_no_hash_entry	adaptive_hash_index
adaptive_hash_rows_removed	adaptive_hash_index
adaptive_hash_rows_updated	adaptive_hash_index
adaptive_hash_searches	adaptive_hash_index
adaptive_hash_searches_btree	adaptive_hash_index
buffer_data_reads	buffer
buffer_data_written	buffer
buffer_flush_adaptive	buffer
buffer_flush_adaptive_pages	buffer
buffer_flush_adaptive_total_pages	buffer
buffer_flush_avg_page_rate	buffer
buffer_flush_background	buffer
buffer_flush_background_pages	buffer
buffer_flush_background_total_pages	buffer
buffer_flush_batches	buffer
buffer_flush_batch_num_scan	buffer
buffer_flush_batch_pages	buffer
buffer_flush_batch_scanned	buffer
buffer_flush_batch_scanned_per_call	buffer
buffer_flush_batch_total_pages	buffer
buffer_flush_lsn_avg_rate	buffer

The INFORMATION\_SCHEMA INNODB\_METRICS Table

buffer_flush_neighbor	buffer
buffer_flush_neighbor_pages	buffer
buffer_flush_neighbor_total_pages	buffer
buffer_flush_n_to_flush_requested	buffer
buffer_flush_pct_for_dirty	buffer
buffer_flush_pct_for_lsn	buffer
buffer_flush_sync	buffer
buffer_flush_sync_pages	buffer
buffer_flush_sync_total_pages	buffer
buffer_flush_sync_waits	buffer
buffer_LRU_batches	buffer
buffer_LRU_batch_num_scan	buffer
buffer_LRU_batch_pages	buffer
buffer_LRU_batch_scanned	buffer
buffer_LRU_batch_scanned_per_call	buffer
buffer_LRU_batch_total_pages	buffer
buffer_LRU_get_free_search	Buffer
buffer_LRU_search_num_scan	buffer
buffer_LRU_search_scanned	buffer
buffer_LRU_search_scanned_per_call	buffer
buffer_LRU_single_flush_failure_count	Buffer
buffer_LRU_single_flush_num_scan	buffer
buffer_LRU_single_flush_scanned	buffer
buffer_LRU_single_flush_scanned_per_call	buffer
buffer_LRU_unzip_search_num_scan	buffer
buffer_LRU_unzip_search_scanned	buffer
buffer_LRU_unzip_search_scanned_per_call	buffer
buffer_pages_created	buffer
buffer_pages_read	buffer
buffer_pages_written	buffer
buffer_page_read_blob	buffer_page_io
buffer_page_read_fsp_hdr	buffer_page_io
buffer_page_read_ibuf_bitmap	buffer_page_io
buffer_page_read_ibuf_free_list	buffer_page_io
buffer_page_read_index_ibuf_leaf	buffer_page_io
buffer_page_read_index_ibuf_non_leaf	buffer_page_io
buffer_page_read_index_inode	buffer_page_io
buffer_page_read_index_leaf	buffer_page_io
buffer_page_read_index_non_leaf	buffer_page_io
buffer_page_read_other	buffer_page_io
buffer_page_read_system_page	buffer_page_io
buffer_page_read_trx_system	buffer_page_io
buffer_page_read_undo_log	buffer_page_io
buffer_page_read_xdes	buffer_page_io
buffer_page_read_zblob	buffer_page_io
buffer_page_read_zblob2	buffer_page_io
buffer_page_written_blob	buffer_page_io
buffer_page_written_fsp_hdr	buffer_page_io
buffer_page_written_ibuf_bitmap	buffer_page_io
buffer_page_written_ibuf_free_list	buffer_page_io
buffer_page_written_index_ibuf_leaf	buffer_page_io
buffer_page_written_index_ibuf_non_leaf	buffer_page_io
buffer_page_written_index_inode	buffer_page_io
buffer_page_written_index_leaf	buffer_page_io
buffer_page_written_index_non_leaf	buffer_page_io
buffer_page_written_other	buffer_page_io
buffer_page_written_system_page	buffer_page_io
buffer_page_written_trx_system	buffer_page_io
buffer_page_written_undo_log	buffer_page_io
buffer_page_written_xdes	buffer_page_io
buffer_page_written_zblob	buffer_page_io
buffer_page_written_zblob2	buffer_page_io
buffer_pool_bytes_data	buffer
buffer_pool_bytes_dirty	buffer
buffer_pool_pages_data	buffer
buffer_pool_pages_dirty	buffer
buffer_pool_pages_free	buffer
buffer_pool_pages_misc	buffer
buffer_pool_pages_total	buffer
buffer_pool_reads	buffer
buffer_pool_read_ahead	buffer
buffer_pool_read_ahead_evicted	buffer

The INFORMATION\_SCHEMA INNODB\_METRICS Table

buffer_pool_read_requests	buffer
buffer_pool_size	server
buffer_pool_wait_free	buffer
buffer_pool_write_requests	buffer
compression_pad_decrements	compression
compression_pad_increments	compression
compress_pages_compressed	compression
compress_pages_decompressed	compression
ddl_background_drop_indexes	ddl
ddl_background_drop_tables	ddl
ddl_online_create_index	ddl
ddl_pending_alter_table	ddl
dml_deletes	dml
dml_inserts	dml
dml_reads	dml
dml_updates	dml
file_num_open_files	file_system
ibuf_merges	change_buffer
ibuf_merges_delete	change_buffer
ibuf_merges_delete_mark	change_buffer
ibuf_merges_discard_delete	change_buffer
ibuf_merges_discard_delete_mark	change_buffer
ibuf_merges_discard_insert	change_buffer
ibuf_merges_insert	change_buffer
ibuf_size	change_buffer
icp_attempts	icp
icp_match	icp
icp_no_match	icp
icp_out_of_range	icp
index_merges	index
index_splits	index
innodb_activity_count	server
innodb_background_drop_table_usec	server
innodb_checkpoint_usec	server
innodb_dblwr_pages_written	server
innodb_dblwr_writes	server
innodb_dict_lru_usec	server
innodb_ibuf_merge_usec	server
innodb_log_flush_usec	server
innodb_master_active_loops	server
innodb_master_idle_loops	server
innodb_master_purge_usec	server
innodb_master_thread_sleeps	server
innodb_mem_validate_usec	server
innodb_page_size	server
innodb_rwlock_s_os_waits	server
innodb_rwlock_s_spin_rounds	server
innodb_rwlock_s_spin_waits	server
innodb_rwlock_x_os_waits	server
innodb_rwlock_x_spin_rounds	server
innodb_rwlock_x_spin_waits	server
lock_deadlocks	lock
lock_rec_locks	lock
lock_rec_lock_created	lock
lock_rec_lock_removed	lock
lock_rec_lock_requests	lock
lock_rec_lock_waits	lock
lock_row_lock_current_waits	lock
lock_row_lock_time	lock
lock_row_lock_time_avg	lock
lock_row_lock_time_max	lock
lock_row_lock_waits	lock
lock_table_locks	lock
lock_table_lock_created	lock
lock_table_lock_removed	lock
lock_table_lock_waits	lock
lock_timeouts	lock
log_checkpoints	recovery
log_lsn_buf_pool_oldest	recovery
log_lsn_checkpoint_age	recovery
log_lsn_current	recovery
log_lsn_last_checkpoint	recovery

log_lsn_last_flush	recovery
log_max_modified_age_async	recovery
log_max_modified_age_sync	recovery
log_num_log_io	recovery
log_pending_checkpoint_writes	recovery
log_pending_log_writes	recovery
log_waits	recovery
log_writes	recovery
log_write_requests	recovery
metadata_mem_pool_size	metadata
metadata_table_handles_closed	metadata
metadata_table_handles_opened	metadata
metadata_table_reference_count	metadata
os_data_fsyncs	os
os_data_reads	os
os_data_writes	os
os_log_bytes_written	os
os_log_fsyncs	os
os_log_pending_fsyncs	os
os_log_pending_writes	os
os_pending_reads	os
os_pending_writes	os
purge_del_mark_records	purge
purge_dml_delay_usec	purge
purge_invoked	purge
purge_resume_count	purge
purge_stop_count	purge
purge_undo_log_pages	purge
purge_upd_exist_or_extern_records	purge
trx_active_transactions	transaction
trx_commits_insert_update	transaction
trx_nl_ro_commits	transaction
trx_rollback	transaction
trx_rollback_savepoint	transaction
trx_rollback_active	transaction
trx_ro_commits	transaction
trx_rseg_current_size	transaction
trx_rseg_history_len	transaction
trx_rw_commits	transaction
trx_undo_slots_cached	transaction
trx_undo_slots_used	transaction

209 rows in set (0.00 sec)

### 19.30.18. The `INFORMATION_SCHEMA INNODB_FT_CONFIG` Table

The `INNODB_FT_CONFIG` table displays metadata about the `FULLTEXT` index and associated processing for an `InnoDB` table.

This table is only accessible to users with the `SUPER` privilege.

**Table 19.17. `INNODB_FT_CONFIG` Columns**

Column name	Description
<code>KEY</code>	The name designating an item of metadata for an <code>InnoDB</code> table containing a <code>FULLTEXT</code> index.
<code>VALUE</code>	The value associated with the corresponding <code>KEY</code> column, reflecting some limit or current value for an aspect of a <code>FULLTEXT</code> index for an <code>InnoDB</code> table.

**Notes:**

- This table is only intended for internal configuration. It is not intended for statistical information purposes.
- You must have the `PROCESS` privilege to query this table.



- The values for the `KEY` column might evolve depending on the needs for performance tuning and debugging for InnoDB full-text processing. Currently, the key values include `optimize_checkpoint_limit`, `synced_doc_id`, `deleted_doc_count`, `stopword_table_name`, and `use_stopword`.
- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

### 19.30.19. The `INFORMATION_SCHEMA INNODB_FT_DEFAULT_STOPWORD` Table

The `INNODB_FT_DEFAULT_STOPWORD` table holds a list of [stopwords](#) that are used by default when creating a `FULLTEXT` index on an InnoDB table.

**Table 19.18. `INNODB_FT_DEFAULT_STOPWORD` Columns**

Column name	Description
<code>value</code>	A word that is used by default as a <a href="#">stopword</a> for <code>FULLTEXT</code> indexes on InnoDB tables. Not used if you override the default <a href="#">stopword processing</a> with either the <code>innodb_ft_server_stopword_table</code> [1725] or the <code>innodb_ft_user_stopword_table</code> [1726] option.

**Notes:**

- You must have the `PROCESS` privilege to query this table.
- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

### 19.30.20. The `INFORMATION_SCHEMA INNODB_FT_INDEX_TABLE` Table

The `INNODB_FT_INDEX_TABLE` table displays information about the inverted index used to process text searches against the `FULLTEXT` index of an InnoDB table.

This table is only accessible to users with the `SUPER` privilege. Before you query this table, set the configuration variable `innodb_ft_aux_table` [1723] to the name (including the database name) of the table that contains the `FULLTEXT` index, for example `test/articles`.

**Table 19.19. `INNODB_FT_INDEX_TABLE` Columns**

Column name	Description
<code>WORD</code>	A word extracted from the text of the columns that are part of a <code>FULLTEXT</code> .
<code>FIRST_DOC_ID</code>	The first document ID that this word appears in in the <code>FULLTEXT</code> index.
<code>LAST_DOC_ID</code>	The last document ID that this word appears in in the <code>FULLTEXT</code> index.
<code>DOC_COUNT</code>	The number of rows this word appears in in the <code>FULLTEXT</code> index. The same word can occur several times within the cache table, once for each combination of <code>DOC_ID</code> and <code>POSITION</code> values.
<code>DOC_ID</code>	The document ID of the row containing the word. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by InnoDB when the table does not contain a suitable column.
<code>POSITION</code>	The position of this particular instance of the word within the relevant document identified by the <code>DOC_ID</code> value.

**Notes:**

- This table initially appears empty, until you set the value of the configuration variable `innodb_ft_aux_table` [1723].
- You must have the `PROCESS` privilege to query this table.

- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

### 19.30.21. The `INFORMATION_SCHEMA.INNODB_FT_INDEX_CACHE` Table

The `INNODB_FT_INDEX_CACHE` table displays token information about newly inserted rows in a `FULLTEXT` index for an `InnoDB` table. To avoid expensive index reorganization during DML operations for an `InnoDB FULLTEXT` index, the information about newly indexed words is stored separately, and combined with the main search index only when you issue the `OPTIMIZE TABLE` statement for the `InnoDB` table.

This table is only accessible to users with the `SUPER` privilege. Before you query this table, set the configuration variable `innodb_ft_aux_table` [1723] to the name (including the database name) of the table that contains the `FULLTEXT` index, for example `test/articles`.

**Table 19.20. `INNODB_FT_INDEX_CACHE` Columns**

Column name	Description
<code>WORD</code>	A word extracted from the text of a newly inserted row.
<code>FIRST_DOC_ID</code>	The first document ID that this word appears in in the <code>FULLTEXT</code> index.
<code>LAST_DOC_ID</code>	The last document ID that this word appears in in the <code>FULLTEXT</code> index.
<code>DOC_COUNT</code>	The number of rows this word appears in in the <code>FULLTEXT</code> index. The same word can occur several times within the cache table, once for each combination of <code>DOC_ID</code> and <code>POSITION</code> values.
<code>DOC_ID</code>	The document ID of the newly inserted row. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by <code>InnoDB</code> when the table does not contain a suitable column.
<code>POSITION</code>	The position of this particular instance of the word within the relevant document identified by the <code>DOC_ID</code> value. The value does not represent an absolute position; it is an offset added to the <code>POSITION</code> of the previous instance of that word.

**Notes:**

- This table initially appears empty, until you set the value of the configuration variable `innodb_ft_aux_table` [1723].
- You must have the `PROCESS` privilege to query this table.
- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

### 19.30.22. The `INFORMATION_SCHEMA.INNODB_FT_DELETED` Table

The `INNODB_FT_DELETED` table records rows that are deleted from the `FULLTEXT` index for an `InnoDB` table. To avoid expensive index reorganization during DML operations for an `InnoDB FULLTEXT` index, the information about newly deleted words is stored separately, filtered out of search results when you do a text search, and removed from the main search index only when you issue the `OPTIMIZE TABLE` statement for the `InnoDB` table.

This table is only accessible to users with the `SUPER` privilege. Before you query this table, set the configuration variable `innodb_ft_aux_table` [1723] to the name (including the database name) of the table that contains the `FULLTEXT` index, for example `test/articles`.

**Table 19.21. `INNODB_FT_DELETED` Columns**

Column name	Description
<code>DOC_ID</code>	The document ID of the newly deleted row. This value might reflect the value of an ID column that you defined for the underlying table, or it

Column name	Description
	can be a sequence value generated by <code>InnoDB</code> when the table does not contain a suitable column. This value is used to skip rows in the <code>innodb_ft_index_table</code> table, when you do text searches before data for deleted rows is physically removed from the <code>FULLTEXT</code> index by an <code>OPTIMIZE TABLE</code> statement.

**Notes:**

- This table initially appears empty, until you set the value of the configuration variable `innodb_ft_aux_table` [1723].
- You must have the `PROCESS` privilege to query this table.
- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

**19.30.23. The `INFORMATION_SCHEMA INNODB_FT_BEING_DELETED` Table**

The `INNODB_FT_BEING_DELETED` table is a temporary work table while document IDs in the `INNODB_FT_DELETED` table are being removed from an `InnoDB FULLTEXT` index during an `OPTIMIZE TABLE` operation. Because its contents typically have a short lifetime, this table has limited utility for monitoring or debugging.

This table is only accessible to users with the `SUPER` privilege. This table initially appears empty, until you set the value of the configuration variable `innodb_ft_aux_table` [1723].

**Table 19.22. `INNODB_FT_BEING_DELETED` Columns**

Column name	Description
<code>DOC_ID</code>	The document ID of the row that is in the process of being deleted. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by <code>InnoDB</code> when the table does not contain a suitable column. This value is used to skip rows in the <code>innodb_ft_index_table</code> table, when you do text searches before data for deleted rows is physically removed from the <code>FULLTEXT</code> index by an <code>OPTIMIZE TABLE</code> statement.

**Notes:**

- This table initially appears empty, until you set the value of the configuration variable `innodb_ft_aux_table` [1723].
- You must have the `PROCESS` privilege to query this table.
- For more information, see [Section 12.9, “Full-Text Search Functions”](#).

**19.31. Thread Pool `INFORMATION_SCHEMA` Tables**

The following sections describe the `INFORMATION_SCHEMA` tables associated with the thread pool plugin. They provide information about thread pool operation:

- `TP_THREAD_STATE`: Information about thread pool thread states
- `TP_THREAD_GROUP_STATE`: Information about thread pool thread group states
- `TP_THREAD_GROUP_STATS`: Thread group statistics

Rows in these tables represent snapshots in time. In the case of `TP_THREAD_STATE`, all rows for a thread group comprise a snapshot in time. Thus, the MySQL server holds the mutex of the thread

group while producing the snapshot. But it does not hold mutexes on all thread groups at the same time, to prevent a statement against `TP_THREAD_STATE` from blocking the entire MySQL server.

The thread pool `INFORMATION_SCHEMA` tables are implemented by individual plugins and the decision whether to load one can be made independently of the others (see [Section 8.11.6.1, “Thread Pool Components and Installation”](#)). However, the content of all the tables depends on the thread pool plugin being enabled. If a table plugin is enabled but the thread pool plugin is not, the table becomes visible and can be accessed but will be empty.

### 19.31.1. The `INFORMATION_SCHEMA TP_THREAD_STATE` Table

This table has one row per thread created by the thread pool to handle connections. The table has these columns:

- `TP_GROUP_ID`

The thread group ID.

- `TP_THREAD_NUMBER`

The ID of the thread within its thread group. `TP_GROUP_ID` and `TP_THREAD_NUMBER` together provide a unique key within the table.

- `PROCESS_COUNT`

The 10ms interval in which the statement that uses this thread is currently executing. 0 means no statement is executing, 1 means it is in the first 10ms, and so forth.

- `WAIT_TYPE`

The type of wait for the thread. `NULL` means the thread is not blocked. Otherwise, the thread is blocked by a call to `thd_wait_begin()` and the value specifies the type of wait. The `xxx_WAIT` columns of the `TP_THREAD_GROUP_STATS` table accumulate counts for each wait type.

The `WAIT_TYPE` value is a string that describes the type of wait, as shown in the following table.

**Table 19.23. `WAIT_TYPE` Values**

Wait Type	Meaning
<code>THD_WAIT_SLEEP</code>	Waiting for sleep
<code>THD_WAIT_DISKIO</code>	Waiting for Disk IO
<code>THD_WAIT_ROW_LOCK</code>	Waiting for row lock
<code>THD_WAIT_GLOBAL_LOCK</code>	Waiting for global lock
<code>THD_WAIT_META_DATA_LOCK</code>	Waiting for metadata lock
<code>THD_WAIT_TABLE_LOCK</code>	Waiting for table lock
<code>THD_WAIT_USER_LOCK</code>	Waiting for user lock
<code>THD_WAIT_BINLOG</code>	Waiting for binlog
<code>THD_WAIT_GROUP_COMMIT</code>	Waiting for group commit
<code>THD_WAIT_SYNC</code>	Waiting for fsync

### 19.31.2. The `INFORMATION_SCHEMA TP_THREAD_GROUP_STATE` Table

This table has one row per thread group in the thread pool. Each row provides information about the current state of a group. The table has these columns:

- `TP_GROUP_ID`

The thread group ID. This is a unique key within the table.

- `CONSUMER_THREADS`

The number of consumer threads. There is at most one thread ready to start executing if the active threads become stalled or blocked.

- `RESERVE_THREADS`

The number of threads in the reserved state. This means that they will not be started until there is a need to wake a new thread and there is no consumer thread. This is where most threads end up when the thread group has created more threads than needed for normal operation. Often a thread group needs additional threads for a short while and then does not need them again for a while. In this case, they go into the reserved state and remain until needed again. They take up some extra memory resources, but no extra computing resources.

- `CONNECTION_COUNT`

The number of connections using this thread group.

- `QUEUED_QUERIES`

The number of statements waiting in the high-priority queue.

- `QUEUED_TRANSACTIONS`

The number of statements waiting in the low-priority queue. These are the initial statements for transactions that have not started, so they also represent queued transactions.

- `STALL_LIMIT`

The value of the `thread_pool_stall_limit` [556] variable on the thread group. This is the same value for all thread groups.

- `PRIO_KICKUP_TIMER`

The value of the `thread_pool_prio_kickup_timer` [555] on the thread group. This is the same value for all thread groups.

- `ALGORITHM`

The value of the `thread_pool_algorithm` [554] on the thread group. This is the same value for all thread groups.

- `THREAD_COUNT`

The number of threads started in the thread pool as part of this thread group.

- `ACTIVE_THREAD_COUNT`

The number of threads active executing statements.

- `MAX_THREAD_IDS_IN_GROUP`

The maximum thread ID of the threads in the group. This is the same as `MAX(TP_THREAD_NUMBER)` [1284] for the threads when selected from the `TP_THREAD_GROUP_STATE` table. That is, these two queries are equivalent:

```
SELECT TP_GROUP_ID, MAX_THREAD_IDS_IN_GROUP
FROM TP_THREAD_GROUP_STATE;

SELECT TP_GROUP_ID, MAX(TP_THREAD_NUMBER)
FROM TP_THREAD_STATE GROUP BY TP_GROUP_ID;
```

- `STALLED_THREAD_COUNT`

The number of stalled statements in the thread group. A stalled statement could be executing, but from a thread pool perspective it is stalled and making no progress. A long-running statement quickly ends up in this category.

- `WAITING_THREAD_NUMBER`

If there is a thread handling the polling of statements in the thread group, this specifies the thread number within this thread group. It is possible that this thread could be executing a statement.

- `OLDEST_QUEUED`

How long in milliseconds the oldest queued statement has been waiting for execution.

### 19.31.3. The `INFORMATION_SCHEMA TP_THREAD_GROUP_STATS` Table

This table reports statistics per thread group. There is one row per group. The table has these columns:

- `TP_GROUP_ID`

The thread group ID. This is a unique key within the table.

- `CONNECTIONS_STARTED`

The number of connections started.

- `CONNECTIONS_CLOSED`

The number of connections closed.

- `QUERIES_EXECUTED`

The number of statements executed. This number is incremented when a statement starts executing, not when it finishes.

- `QUERIES_QUEUED`

The number of statements received that were queued for execution. This does not count statements that the thread group was able to begin executing immediately without queuing, which can happen under the conditions described in [Section 8.11.6.2, “Thread Pool Operation”](#).

- `THREADS_STARTED`

The number of threads started.

- `PRIO_KICKUPS`

The number of statements that have been moved from low-priority queue to high-priority queue based on the value of the `thread_pool_prio_kickup_timer` [555] system variable. If this number increases quickly, consider increasing the value of that variable. A quickly increasing counter means that the priority system is not keeping transactions from starting too early. For `InnoDB`, this most likely means deteriorating performance due to too many concurrent transactions..

- `STALLED_QUERIES_EXECUTED`

The number of statements that have become defined as stalled due to executing for a time longer than the value of the `thread_pool_stall_limit` [556] system variable.

- `BECOME_CONSUMER_THREAD`

The number of times thread have been assigned the consumer thread role.

- `BECOME_RESERVE_THREAD`

The number of times threads have been assigned the reserve thread role.

- `BECOME_WAITING_THREAD`

The number of times threads have been assigned the waiter thread role. When statements are queued, this happens very often, even in normal operation, so rapid increases in this value are normal in the case of a highly loaded system where statements are queued up.

- `WAKE_THREAD_STALL_CHECKER`

The number of times the stall check thread decided to wake or create a thread to possibly handle some statements or take care of the waiter thread role.

- `SLEEP_WAITS`

The number of `THD_WAIT_SLEEP` waits. These occur when threads go to sleep; for example, by calling the `SLEEP( )` [1280] function.

- `DISK_IO_WAITS`

The number of `THD_WAIT_DISKIO` waits. These occur when threads perform disk I/O that is likely to not hit the file system cache. Such waits occur when the buffer pool reads and writes data to disk, not for normal reads from and writes to files.

- `ROW_LOCK_WAITS`

The number of `THD_WAIT_ROW_LOCK` waits for release of a row lock by another transaction.

- `GLOBAL_LOCK_WAITS`

The number of `THD_WAIT_GLOBAL_LOCK` waits for a global lock to be released.

- `META_DATA_LOCK_WAITS`

The number of `THD_WAIT_META_DATA_LOCK` waits for a metadata lock to be released.

- `TABLE_LOCK_WAITS`

The number of `THD_WAIT_TABLE_LOCK` waits for a table to be unlocked that the statement needs to access.

- `USER_LOCK_WAITS`

The number of `THD_WAIT_USER_LOCK` waits for a special lock constructed by the user thread.

- `BINLOG_WAITS`

The number of `THD_WAIT_BINLOG_WAITS` waits for the binary log to become free.

- `GROUP_COMMIT_WAITS`

The number of `THD_WAIT_GROUP_COMMIT` waits. These occur when a group commit must wait for the other parties to complete their part of a transaction.

- `FSYNC_WAITS`

The number of `THD_WAIT_SYNC` waits for a file sync operation.

## 19.32. Extensions to `SHOW` Statements

Some extensions to `SHOW` statements accompany the implementation of `INFORMATION_SCHEMA`:



- `SHOW` can be used to get information about the structure of `INFORMATION_SCHEMA` itself.
- Several `SHOW` statements accept a `WHERE` clause that provides more flexibility in specifying which rows to display.

`INFORMATION_SCHEMA` is an information database, so its name is included in the output from `SHOW DATABASES`. Similarly, `SHOW TABLES` can be used with `INFORMATION_SCHEMA` to obtain a list of its tables:

```
mysql> SHOW TABLES FROM INFORMATION_SCHEMA;
+-----+
| Tables_in_INFORMATION_SCHEMA |
+-----+
| CHARACTER_SETS                |
| COLLATIONS                    |
| COLLATION_CHARACTER_SET_APPLICABILITY |
| COLUMNS                      |
| COLUMN_PRIVILEGES             |
| ENGINES                      |
| EVENTS                        |
| FILES                         |
| GLOBAL_STATUS                 |
| GLOBAL_VARIABLES             |
| KEY_COLUMN_USAGE              |
| PARTITIONS                    |
| PLUGINS                      |
| PROCESSLIST                   |
| REFERENTIAL_CONSTRAINTS      |
| ROUTINES                     |
| SCHEMATA                      |
| SCHEMA_PRIVILEGES            |
| SESSION_STATUS               |
| SESSION_VARIABLES            |
| STATISTICS                    |
| TABLES                       |
| TABLE_CONSTRAINTS           |
| TABLE_PRIVILEGES            |
| TRIGGERS                     |
| USER_PRIVILEGES              |
| VIEWS                         |
+-----+
27 rows in set (0.00 sec)
```

`SHOW COLUMNS` and `DESCRIBE` can display information about the columns in individual `INFORMATION_SCHEMA` tables.

`SHOW` statements that accept a `LIKE [1190]` clause to limit the rows displayed also permit a `WHERE` clause that specifies more general conditions that selected rows must satisfy:

```
SHOW CHARACTER SET
SHOW COLLATION
SHOW COLUMNS
SHOW DATABASES
SHOW FUNCTION STATUS
SHOW INDEX
SHOW OPEN TABLES
SHOW PROCEDURE STATUS
SHOW STATUS
SHOW TABLE STATUS
SHOW TABLES
SHOW TRIGGERS
SHOW VARIABLES
```

The `WHERE` clause, if present, is evaluated against the column names displayed by the `SHOW` statement. For example, the `SHOW CHARACTER SET` statement produces these output columns:

```
mysql> SHOW CHARACTER SET;
+-----+-----+-----+-----+
| Charset | Description | Default collation | Maxlen |
+-----+-----+-----+-----+
```

big5	Big5 Traditional Chinese	big5_chinese_ci	2
dec8	DEC West European	dec8_swedish_ci	1
cp850	DOS West European	cp850_general_ci	1
hp8	HP West European	hp8_english_ci	1
koi8r	KOI8-R Relcom Russian	koi8r_general_ci	1
latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
...			

To use a `WHERE` clause with `SHOW CHARACTER SET`, you would refer to those column names. As an example, the following statement displays information about character sets for which the default collation contains the string 'japanese':

```
mysql> SHOW CHARACTER SET WHERE `Default collation` LIKE '%japanese%';
```

Charset	Description	Default collation	Maxlen
ujis	EUC-JP Japanese	ujis_japanese_ci	3
sjis	Shift-JIS Japanese	sjis_japanese_ci	2
cp932	SJIS for Windows Japanese	cp932_japanese_ci	2
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci	3

This statement displays the multi-byte character sets:

```
mysql> SHOW CHARACTER SET WHERE Maxlen > 1;
```

Charset	Description	Default collation	Maxlen
big5	Big5 Traditional Chinese	big5_chinese_ci	2
ujis	EUC-JP Japanese	ujis_japanese_ci	3
sjis	Shift-JIS Japanese	sjis_japanese_ci	2
euckr	EUC-KR Korean	euckr_korean_ci	2
gb2312	GB2312 Simplified Chinese	gb2312_chinese_ci	2
gbk	GBK Simplified Chinese	gbk_chinese_ci	2
utf8	UTF-8 Unicode	utf8_general_ci	3
ucs2	UCS-2 Unicode	ucs2_general_ci	2
cp932	SJIS for Windows Japanese	cp932_japanese_ci	2
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci	3



---

# Chapter 20. MySQL Performance Schema

## Table of Contents

20.1. Performance Schema Quick Start .....	2234
20.2. Performance Schema Configuration .....	2240
20.2.1. Performance Schema Build Configuration .....	2240
20.2.2. Performance Schema Startup Configuration .....	2241
20.2.3. Performance Schema Runtime Configuration .....	2243
20.3. Performance Schema Queries .....	2258
20.4. Performance Schema Instrument Naming Conventions .....	2259
20.5. Performance Schema Status Monitoring .....	2261
20.6. Performance Schema Atom and Molecule Events .....	2264
20.7. Performance Schema Statement Digests .....	2264
20.8. Performance Schema General Table Characteristics .....	2266
20.9. Performance Schema Table Descriptions .....	2266
20.9.1. Performance Schema Setup Tables .....	2266
20.9.2. Performance Schema Instance Tables .....	2270
20.9.3. Performance Schema Wait Event Tables .....	2275
20.9.4. Performance Schema Stage Event Tables .....	2278
20.9.5. Performance Schema Statement Event Tables .....	2279
20.9.6. Performance Schema Connection Tables .....	2284
20.9.7. Performance Schema Connection Attribute Tables .....	2286
20.9.8. Performance Schema Summary Tables .....	2286
20.9.9. Performance Schema Miscellaneous Tables .....	2299
20.10. Performance Schema Option and Variable Reference .....	2305
20.11. Performance Schema Command Options .....	2306
20.12. Performance Schema System Variables .....	2307
20.13. Performance Schema Status Variables .....	2318
20.14. Performance Schema and Plugins .....	2319
20.15. Using the Performance Schema to Diagnose Problems .....	2319

The MySQL Performance Schema is a feature for monitoring MySQL Server execution at a low level. The Performance Schema has these characteristics:

- The Performance Schema provides a way to inspect internal execution of the server at runtime. It is implemented using the [PERFORMANCE\\_SCHEMA](#) storage engine and the [performance\\_schema](#) database. The Performance Schema focuses primarily on performance data. This differs from [INFORMATION\\_SCHEMA](#), which serves for inspection of metadata.
- The Performance Schema monitors server events. An “event” is anything the server does that takes time and has been instrumented so that timing information can be collected. In general, an event could be a function call, a wait for the operating system, a stage of an SQL statement execution such as parsing or sorting, or an entire statement or group of statements. Currently, event collection provides access to information about synchronization calls (such as for mutexes) file and table I/O, table locks, and so forth for the server and for several storage engines.
- Performance Schema events are distinct from events written to the server's binary log (which describe data modifications) and Event Scheduler events (which are a type of stored program).
- Performance Schema events are specific to a given instance of the MySQL Server. In MySQL 5.6.9 and later, Performance Schema tables are considered local to the server, and changes to them are not replicated or written to the binary log. (Bug #14741537)
- Current events are available, as well as event histories and summaries. This enables you to determine how many times instrumented activities were performed and how much time they took.

Event information is available to show the activities of specific threads, or activity associated with particular objects such as a mutex or file.

- The `PERFORMANCE_SCHEMA` storage engine collects event data using “instrumentation points” in server source code.
- Collected events are stored in tables in the `performance_schema` database. These tables can be queried using `SELECT` statements like other tables.
- Performance Schema configuration can be modified dynamically by updating tables in the `performance_schema` database through SQL statements. Configuration changes affect data collection immediately.
- Tables in the `performance_schema` database are views or temporary tables that use no persistent on-disk storage.
- Monitoring is available on all platforms supported by MySQL.

Some limitations might apply: The types of timers might vary per platform. Instruments that apply to storage engines might not be implemented for all storage engines. Instrumentation of each third-party engine is the responsibility of the engine maintainer. See also [Section E.8, “Restrictions on Performance Schema”](#).

- Data collection is implemented by modifying the server source code to add instrumentation. There are no separate threads associated with the Performance Schema, unlike other features such as replication or the Event Scheduler.

The Performance Schema is intended to provide access to useful information about server execution while having minimal impact on server performance. The implementation follows these design goals:

- Activating the Performance Schema causes no changes in server behavior. For example, it does not cause thread scheduling to change, and it does not cause query execution plans (as shown by `EXPLAIN`) to change.
- No memory allocation is done beyond that which occurs during server startup. By using early allocation of structures with a fixed size, it is never necessary to resize or reallocate them, which is critical for achieving good runtime performance.
- Server monitoring occurs continuously and unobtrusively with very little overhead. Activating the Performance Schema does not make the server unusable.
- The parser is unchanged. There are no new keywords or statements.
- Execution of server code proceeds normally even if the Performance Schema fails internally.
- When there is a choice between performing processing during event collection initially or during event retrieval later, priority is given to making collection faster. This is because collection is ongoing whereas retrieval is on demand and might never happen at all.
- It is easy to add new instrumentation points.
- Instrumentation is versioned. If the instrumentation implementation changes, previously instrumented code will continue to work. This benefits developers of third-party plugins because it is not necessary to upgrade each plugin to stay synchronized with the latest Performance Schema changes.

## 20.1. Performance Schema Quick Start

This section briefly introduces the Performance Schema with examples that show how to use it. For additional examples, see [Section 20.15, “Using the Performance Schema to Diagnose Problems”](#).

For the Performance Schema to be available, support for it must have been configured when MySQL was built. You can verify whether this is the case by checking the server's help output. If the

Performance Schema is available, the output will mention several variables with names that begin with `performance_schema`:

```
shell> mysqld --verbose --help
...
--performance_schema
           Enable the performance schema.
--performance_schema_events_waits_history_long_size=#
           Number of rows in events_waits_history_long.
...
```

If such variables do not appear in the output, your server has not been built to support the Performance Schema. In this case, see [Section 20.2, “Performance Schema Configuration”](#).

Assuming that the Performance Schema is available, it is enabled by default as of MySQL 5.6.6. Before 5.6.6, it is disabled by default. To enable or disable it explicitly, start the server with the `performance_schema` [2307] variable set to an appropriate value. For example, use these lines in your `my.cnf` file:

```
[mysqld]
performance_schema=on
```

When the server starts, it sees `performance_schema` [2307] and attempts to initialize the Performance Schema. To verify successful initialization, use this statement:

```
mysql> SHOW VARIABLES LIKE 'performance_schema';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| performance_schema | ON   |
+-----+-----+
```

A value of `ON` means that the Performance Schema initialized successfully and is ready for use. A value of `OFF` means that some error occurred. Check the server error log for information about what went wrong.

The Performance Schema is implemented as a storage engine. If this engine is available (which you should already have checked earlier), you should see it listed with a `SUPPORT` value of `YES` in the output from the `INFORMATION_SCHEMA.ENGINES` table or the `SHOW ENGINES` statement:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.ENGINES
-> WHERE ENGINE='PERFORMANCE_SCHEMA'\G
***** 1. row *****
ENGINE: PERFORMANCE_SCHEMA
SUPPORT: YES
COMMENT: Performance Schema
TRANSACTIONS: NO
XA: NO
SAVEPOINTS: NO

mysql> SHOW ENGINES\G
...
Engine: PERFORMANCE_SCHEMA
Support: YES
Comment: Performance Schema
Transactions: NO
XA: NO
Savepoints: NO
...
```

The `PERFORMANCE_SCHEMA` storage engine operates on tables in the `performance_schema` database. You can make `performance_schema` the default database so that references to its tables need not be qualified with the database name:

```
mysql> USE performance_schema;
```

Many examples in this chapter assume that `performance_schema` is the default database.

Performance Schema tables are stored in the `performance_schema` database. Information about the structure of this database and its tables can be obtained, as for any other database, by selecting from the `INFORMATION_SCHEMA` database or by using `SHOW` statements. For example, use either of these statements to see what Performance Schema tables exist:

```
mysql> SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
-> WHERE TABLE_SCHEMA = 'performance_schema';
+-----+
| TABLE_NAME |
+-----+
| accounts |
| cond_instances |
| events_stages_current |
| events_stages_history |
| events_stages_history_long |
| events_stages_summary_by_account_by_event_name |
| events_stages_summary_by_host_by_event_name |
| events_stages_summary_by_thread_by_event_name |
| events_stages_summary_by_user_by_event_name |
| events_stages_summary_global_by_event_name |
| events_statements_current |
| events_statements_history |
| events_statements_history_long |
| ... |
| file_instances |
| file_summary_by_event_name |
| file_summary_by_instance |
| host_cache |
| hosts |
| mutex_instances |
| objects_summary_global_by_type |
| performance_timers |
| rwlock_instances |
| session_account_connect_attrs |
| session_connect_attrs |
| setup_actors |
| setup_consumers |
| setup_instruments |
| setup_objects |
| setup_timers |
| socket_instances |
| socket_summary_by_event_name |
| socket_summary_by_instance |
| table_io_waits_summary_by_index_usage |
| table_io_waits_summary_by_table |
| table_lock_waits_summary_by_table |
| threads |
| users |
+-----+

mysql> SHOW TABLES FROM performance_schema;
+-----+
| Tables_in_performance_schema |
+-----+
| accounts |
| cond_instances |
| events_stages_current |
| events_stages_history |
| events_stages_history_long |
| ... |
```

The number of Performance Schema tables is expected to increase over time as implementation of additional instrumentation proceeds.

The name of the `performance_schema` database is lowercase, as are the names of tables within it. Queries should specify the names in lowercase.

To see the structure of individual tables, use `SHOW CREATE TABLE`:

```
mysql> SHOW CREATE TABLE setup_timers\G
```



```

***** 1. row *****
      Table: setup_timers
Create Table: CREATE TABLE `setup_timers` (
  `NAME` varchar(64) NOT NULL,
  `TIMER_NAME` enum('CYCLE','NANOSECOND','MICROSECOND','MILLISECOND','TICK')
  NOT NULL
) ENGINE=PERFORMANCE_SCHEMA DEFAULT CHARSET=utf8

```

Table structure is also available by selecting from tables such as [INFORMATION\\_SCHEMA.COLUMNS](#) or by using statements such as [SHOW COLUMNS](#).

Tables in the [performance\\_schema](#) database can be grouped according to the type of information in them: Current events, event histories and summaries, object instances, and setup (configuration) information. The following examples illustrate a few uses for these tables. For detailed information about the tables in each group, see [Section 20.9, “Performance Schema Table Descriptions”](#).

To see what the server is doing at the moment, examine the [events\\_waits\\_current](#) table. It contains one row per thread showing each thread's most recent monitored event:

```

mysql> SELECT * FROM events_waits_current\G
***** 1. row *****
      THREAD_ID: 0
      EVENT_ID: 5523
      EVENT_NAME: wait/synch/mutex/mysys/THR_LOCK::mutex
      SOURCE: thr_lock.c:525
      TIMER_START: 201660494489586
      TIMER_END: 201660494576112
      TIMER_WAIT: 86526
      SPINS: NULL
      OBJECT_SCHEMA: NULL
      OBJECT_NAME: NULL
      OBJECT_TYPE: NULL
OBJECT_INSTANCE_BEGIN: 142270668
      NESTING_EVENT_ID: NULL
      OPERATION: lock
      NUMBER_OF_BYTES: NULL
      FLAGS: 0
...

```

This event indicates that thread 0 was waiting for 86,526 picoseconds to acquire a lock on [THR\\_LOCK::mutex](#), a mutex in the [mysys](#) subsystem. The first few columns provide the following information:

- The ID columns indicate which thread the event comes from and the event number.
- [EVENT\\_NAME](#) indicates what was instrumented and [SOURCE](#) indicates which source file contains the instrumented code.
- The timer columns show when the event started and stopped and how long it took. If an event is still in progress, the [TIMER\\_END](#) and [TIMER\\_WAIT](#) values are [NULL](#). Timer values are approximate and expressed in picoseconds. For information about timers and event time collection, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

The history tables contain the same kind of rows as the current-events table but have more rows and show what the server has been doing “recently” rather than “currently.” The [events\\_waits\\_history](#) and [events\\_waits\\_history\\_long](#) tables contain the most recent 10 events per thread and most recent 10,000 events, respectively. For example, to see information for recent events produced by thread 13, do this:

```

mysql> SELECT EVENT_ID, EVENT_NAME, TIMER_WAIT
-> FROM events_waits_history WHERE THREAD_ID = 13
-> ORDER BY EVENT_ID;
+-----+-----+-----+
| EVENT_ID | EVENT_NAME | TIMER_WAIT |
+-----+-----+-----+
| 86 | wait/synch/mutex/mysys/THR_LOCK::mutex | 686322 |
| 87 | wait/synch/mutex/mysys/THR_LOCK_malloc | 320535 |
+-----+-----+-----+

```

88	wait/synch/mutex/mysys/THR_LOCK_malloc	339390
89	wait/synch/mutex/mysys/THR_LOCK_malloc	377100
90	wait/synch/mutex/sql/LOCK_plugin	614673
91	wait/synch/mutex/sql/LOCK_open	659925
92	wait/synch/mutex/sql/THD::LOCK_thd_data	494001
93	wait/synch/mutex/mysys/THR_LOCK_malloc	222489
94	wait/synch/mutex/mysys/THR_LOCK_malloc	214947
95	wait/synch/mutex/mysys/LOCK_alarm	312993

As new events are added to a history table, older events are discarded if the table is full.

Summary tables provide aggregated information for all events over time. The tables in this group summarize event data in different ways. To see which instruments have been executed the most times or have taken the most wait time, sort the `events_waits_summary_global_by_event_name` table on the `COUNT_STAR` or `SUM_TIMER_WAIT` column, which correspond to a `COUNT(*)` or `SUM(TIMER_WAIT)` value, respectively, calculated over all events:

```
mysql> SELECT EVENT_NAME, COUNT_STAR
-> FROM events_waits_summary_global_by_event_name
-> ORDER BY COUNT_STAR DESC LIMIT 10;
```

EVENT_NAME	COUNT_STAR
wait/synch/mutex/mysys/THR_LOCK_malloc	6419
wait/io/file/sql/FRM	452
wait/synch/mutex/sql/LOCK_plugin	337
wait/synch/mutex/mysys/THR_LOCK_open	187
wait/synch/mutex/mysys/LOCK_alarm	147
wait/synch/mutex/sql/THD::LOCK_thd_data	115
wait/io/file/myisam/kfile	102
wait/synch/mutex/sql/LOCK_global_system_variables	89
wait/synch/mutex/mysys/THR_LOCK::mutex	89
wait/synch/mutex/sql/LOCK_open	88

```
mysql> SELECT EVENT_NAME, SUM_TIMER_WAIT
-> FROM events_waits_summary_global_by_event_name
-> ORDER BY SUM_TIMER_WAIT DESC LIMIT 10;
```

EVENT_NAME	SUM_TIMER_WAIT
wait/io/file/sql/MYSQL_LOG	1599816582
wait/synch/mutex/mysys/THR_LOCK_malloc	1530083250
wait/io/file/sql/binlog_index	1385291934
wait/io/file/sql/FRM	1292823243
wait/io/file/myisam/kfile	411193611
wait/io/file/myisam/dfile	322401645
wait/synch/mutex/mysys/LOCK_alarm	145126935
wait/io/file/sql/casetest	104324715
wait/synch/mutex/sql/LOCK_plugin	86027823
wait/io/file/sql/pid	72591750

These results show that the `THR_LOCK_malloc` mutex is “hot,” both in terms of how often it is used and amount of time that threads wait attempting to acquire it.



#### Note

The `THR_LOCK_malloc` mutex is used only in debug builds. In production builds it is not hot because it is nonexistent.

Instance tables document what types of objects are instrumented. An instrumented object, when used by the server, produces an event. These tables provide event names and explanatory notes or status information. For example, the `file_instances` table lists instances of instruments for file I/O operations and their associated files:

```
mysql> SELECT * FROM file_instances\G
***** 1. row *****
```

```

FILE_NAME: /opt/mysql-log/60500/binlog.000007
EVENT_NAME: wait/io/file/sql/binlog
OPEN_COUNT: 0
***** 2. row *****
FILE_NAME: /opt/mysql/60500/data/mysql/tables_priv.MYI
EVENT_NAME: wait/io/file/myisam/kfile
OPEN_COUNT: 1
***** 3. row *****
FILE_NAME: /opt/mysql/60500/data/mysql/columns_priv.MYI
EVENT_NAME: wait/io/file/myisam/kfile
OPEN_COUNT: 1
...

```

Setup tables are used to configure and display monitoring characteristics. For example, to see which event timers are selected, query the `setup_timers` tables:

```

mysql> SELECT * FROM setup_timers;
+-----+-----+
| NAME      | TIMER_NAME |
+-----+-----+
| idle      | MICROSECOND |
| wait      | CYCLE       |
| stage     | NANOSECOND  |
| statement | NANOSECOND  |
+-----+-----+

```

`setup_instruments` lists the set of instruments for which events can be collected and shows which of them are enabled:

```

mysql> SELECT * FROM setup_instruments;
+-----+-----+-----+
| NAME                                     | ENABLED | TIMED |
+-----+-----+-----+
...
| wait/synch/mutex/sql/LOCK_global_read_lock | YES     | YES   |
| wait/synch/mutex/sql/LOCK_global_system_variables | YES     | YES   |
| wait/synch/mutex/sql/LOCK_lock_db | YES     | YES   |
| wait/synch/mutex/sql/LOCK_manager | YES     | YES   |
...
| wait/synch/rwlock/sql/LOCK_grant | YES     | YES   |
| wait/synch/rwlock/sql/LOGGER::LOCK_logger | YES     | YES   |
| wait/synch/rwlock/sql/LOCK_sys_init_connect | YES     | YES   |
| wait/synch/rwlock/sql/LOCK_sys_init_slave | YES     | YES   |
...
| wait/io/file/sql/binlog | YES     | YES   |
| wait/io/file/sql/binlog_index | YES     | YES   |
| wait/io/file/sql/casetest | YES     | YES   |
| wait/io/file/sql/dbopt | YES     | YES   |
...

```

To understand how to interpret instrument names, see [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

To control whether events are collected for an instrument, set its `ENABLED` value to `YES` or `NO`. For example:

```

mysql> UPDATE setup_instruments SET ENABLED = 'NO'
-> WHERE NAME = 'wait/synch/mutex/sql/LOCK_mysql_create_db';

```

The Performance Schema uses collected events to update tables in the `performance_schema` database, which act as “consumers” of event information. The `setup_consumers` table lists the available consumers and which are enabled:

```

mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                                     | ENABLED |
+-----+-----+
| events_stages_current | NO      |
| events_stages_history | NO      |

```

events_stages_history_long	NO
events_statements_current	YES
events_statements_history	NO
events_statements_history_long	NO
events_waits_current	NO
events_waits_history	NO
events_waits_history_long	NO
global_instrumentation	YES
thread_instrumentation	YES
statements_digest	YES

To control whether the Performance Schema maintains a consumer as a destination for event information, set its `ENABLED` value.

For more information about the setup tables and how to use them to control event collection, see [Section 20.2.3.2, “Performance Schema Event Filtering”](#).

There are some miscellaneous tables that do not fall into any of the previous groups. For example, `performance_timers` lists the available event timers and their characteristics. For information about timers, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

## 20.2. Performance Schema Configuration

To use the MySQL Performance Schema, these configuration considerations apply:

- The Performance Schema must be configured into MySQL Server at build time to make it available. Performance Schema support is included in binary MySQL distributions. If you are building from source, you must ensure that it is configured into the build as described in [Section 20.2.1, “Performance Schema Build Configuration”](#).
- The Performance Schema must be enabled at server startup to enable event collection to occur. Specific Performance Schema features can be enabled at server startup or at runtime to control which types of event collection occur. See [Section 20.2.2, “Performance Schema Startup Configuration”](#), [Section 20.2.3, “Performance Schema Runtime Configuration”](#), and [Section 20.2.3.2, “Performance Schema Event Filtering”](#).

### 20.2.1. Performance Schema Build Configuration

For the Performance Schema to be available, it must be configured into the MySQL server at build time. Binary MySQL distributions provided by Oracle Corporation are configured to support the Performance Schema. If you use a binary MySQL distribution from another provider, check with the provider whether the distribution has been appropriately configured.

If you build MySQL from a source distribution, enable the Performance Schema by running `CMake` with the `WITH_PERFSCHEMA_STORAGE_ENGINE` option enabled:

```
shell> cmake . -DWITH_PERFSCHEMA_STORAGE_ENGINE=1
```

Configuring MySQL with the `-DWITHOUT_PERFSCHEMA_STORAGE_ENGINE=1` option prevents inclusion of the Performance Schema, so if you want it included, do not use this option. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

If you install MySQL over a previous installation that was configured without the Performance Schema (or with an older version of the Performance Schema that may not have all the current tables), run `mysql_upgrade` after starting the server to ensure that the `performance_schema` database exists with all current tables. Then restart the server. One indication that you need to do this is the presence of messages such as the following in the error log:

```
[ERROR] Native table 'performance_schema`.`events_waits_history'
has the wrong structure
[ERROR] Native table 'performance_schema`.`events_waits_history_long'
```

```
has the wrong structure
...
```

To verify whether a server was built with Performance Schema support, check its help output. If the Performance Schema is available, the output will mention several variables with names that begin with `performance_schema`:

```
shell> mysql --verbose --help
...
--performance_schema
           Enable the performance schema.
--performance_schema_events_waits_history_long_size=#
           Number of rows in events_waits_history_long.
...
```

You can also connect to the server and look for a line that names the `PERFORMANCE_SCHEMA` storage engine in the output from `SHOW ENGINES`:

```
mysql> SHOW ENGINES\G
...
  Engine: PERFORMANCE_SCHEMA
  Support: YES
  Comment: Performance Schema
  Transactions: NO
             XA: NO
  Savepoints: NO
...
```

If the Performance Schema was not configured into the server at build time, no row for `PERFORMANCE_SCHEMA` will appear in the output from `SHOW ENGINES`. You might see `performance_schema` listed in the output from `SHOW DATABASES`, but it will have no tables and you will not be able to use it.

A line for `PERFORMANCE_SCHEMA` in the `SHOW ENGINES` output means that the Performance Schema is available, not that it is enabled. To enable it, you must do so at server startup, as described in the next section.

## 20.2.2. Performance Schema Startup Configuration

Assuming that the Performance Schema is available, it is enabled by default as of MySQL 5.6.6. Before 5.6.6, it is disabled by default. To enable or disable it explicitly, start the server with the `performance_schema` [2307] variable set to an appropriate value. For example, use these lines in your `my.cnf` file:

```
[mysqld]
performance_schema=on
```

If the server is unable to allocate any internal buffer during Performance Schema initialization, the Performance Schema disables itself and sets `performance_schema` [2307] to `OFF`, and the server runs without instrumentation.

As of MySQL 5.6.4, the Performance Schema permits instrument and consumer configuration at server startup, which previously was possible only at runtime using `UPDATE` statements for the `setup_instruments` and `setup_consumers` tables. This change was made because configuration at runtime is too late to disable instruments that have already been initialized during server startup. For example, the `wait/synch/mutex/sql/LOCK_open` mutex is initialized once during server startup, so attempts to disable the corresponding instrument at runtime have no effect.

To control an instrument at server startup, use an option of this form:

```
--performance-schema-instrument='instrument_name=value'
```

Here, `instrument_name` is an instrument name such as `wait/synch/mutex/sql/LOCK_open`, and `value` is one of these values:

- `off`, `false`, or `0`: Disable the instrument
- `on`, `true`, or `1`: Enable and time the instrument
- `counted`: Enable and count (rather than time) the instrument

Each `--performance-schema-instrument` [2306] option can specify only one instrument name, but multiple instances of the option can be given to configure multiple instruments. In addition, patterns are permitted in instrument names to configure instruments that match the pattern. To configure all condition synchronization instruments as enabled and counted, use this option:

```
--performance-schema-instrument='wait/synch/cond/%=counted'
```

To disable all instruments, use this option:

```
--performance-schema-instrument='%=off'
```

Longer instrument name strings take precedence over shorter pattern names, regardless of order. For information about specifying patterns to select instruments, see [Section 20.2.3.2.2, “Naming Instruments or Consumers for Filtering Operations”](#).

An unrecognized instrument name is ignored. It is possible that a plugin installed later may create the instrument, at which time the name is recognized and configured.

To control a consumer at server startup, use an option of this form:

```
--performance-schema-consumer-consumer_name=value
```

Here, *consumer\_name* is a consumer name such as `events_waits_history`, and *value* is one of these values:

- `off`, `false`, or `0`: Do not collect events for the consumer
- `on`, `true`, or `1`: Collect events for the consumer

For example, to enable the `events_waits_history` consumer, use this option:

```
--performance-schema-consumer-events-waits-history=on
```

The permitted consumer names can be found by examining the `setup_consumers` table. Patterns are not permitted. Consumer names in the `setup_consumers` table use underscores, but for consumers set at startup, dashes and underscores within the name are equivalent.

The Performance Schema includes several system variables that provide configuration information:

```
mysql> SHOW VARIABLES LIKE 'perf%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| performance_schema | ON |
| performance_schema_accounts_size | 100 |
| performance_schema_digests_size | 200 |
| performance_schema_events_stages_history_long_size | 10000 |
| performance_schema_events_stages_history_size | 10 |
| performance_schema_events_statements_history_long_size | 10000 |
| performance_schema_events_statements_history_size | 10 |
| performance_schema_events_waits_history_long_size | 10000 |
| performance_schema_events_waits_history_size | 10 |
| performance_schema_hosts_size | 100 |
| performance_schema_max_cond_classes | 80 |
| performance_schema_max_cond_instances | 1000 |
| ... |
```

The `performance_schema` [2307] variable is `ON` or `OFF` to indicate whether the Performance Schema is enabled or disabled. The other variables indicate table sizes (number of rows) or memory allocation values.

**Note**

With the Performance Schema enabled, the number of Performance Schema instances affects the server memory footprint, perhaps to a large extent. It may be necessary to tune the values of Performance Schema system variables to find the number of instances that balances insufficient instrumentation against excessive memory consumption.

To change the value of Performance Schema system variables, set them at server startup. For example, put the following lines in a `my.cnf` file to change the sizes of the history tables:

```
[mysqld]
performance_schema
performance_schema_events_waits_history_size=20
performance_schema_events_waits_history_long_size=15000
```

As of MySQL 5.6.6, the Performance Schema automatically sizes the values of several of its parameters at server startup if they are not set explicitly. For example, it sizes the parameters that control the sizes of the events waits tables this way. To see which parameters are sized under this policy, use `mysqld --verbose --help` and look for those with a default value of `-1`, or see [Section 20.12, “Performance Schema System Variables”](#).

For each autosized parameter that is not set at server startup (or is set to `-1`), the Performance Schema determines how to set its value based on the value of the following system values, which are considered as “hints” about how you have configured your MySQL server:

```
max_connections
open_files_limit
table_definition_cache
table_open_cache
```

To override autosizing for a given parameter, set it a value other than `-1` at startup. In this case, the Performance Schema assigns it the specified value.

At runtime, `SHOW VARIABLES` displays the actual values that autosized parameters were set to.

If the Performance Schema is disabled, its autosized parameters remain set to `-1` and `SHOW VARIABLES` displays `-1`.

### 20.2.3. Performance Schema Runtime Configuration

Performance Schema setup tables contain information about monitoring configuration:

```
mysql> SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
  -> WHERE TABLE_SCHEMA = 'performance_schema'
  -> AND TABLE_NAME LIKE 'setup%';
+-----+
| TABLE_NAME |
+-----+
| setup_actors |
| setup_consumers |
| setup_instruments |
| setup_objects |
| setup_timers |
+-----+
```

You can examine the contents of these tables to obtain information about Performance Schema monitoring characteristics. If you have the `UPDATE` [756] privilege, you can change Performance Schema operation by modifying setup tables to affect how monitoring occurs. For additional details about these tables, see [Section 20.9.1, “Performance Schema Setup Tables”](#).

To see which event timers are selected, query the `setup_timers` tables:

```
mysql> SELECT * FROM setup_timers;
```



NAME	TIMER_NAME
idle	MICROSECOND
wait	CYCLE
stage	NANOSECOND
statement	NANOSECOND

The `NAME` value indicates the type of instrument to which the timer applies, and `TIMER_NAME` indicates which timer applies to those instruments. The timer applies to instruments where their name begins with a component matching the `NAME` value.

To change the timer, update the `NAME` value. For example, to use the `NANOSECOND` timer for the `wait` timer:

```
mysql> UPDATE setup_timers SET TIMER_NAME = 'NANOSECOND'
-> WHERE NAME = 'wait';
mysql> SELECT * FROM setup_timers;
```

NAME	TIMER_NAME
idle	MICROSECOND
wait	NANOSECOND
stage	NANOSECOND
statement	NANOSECOND

For discussion of timers, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

The `setup_instruments` and `setup_consumers` tables list the instruments for which events can be collected and the types of consumers for which event information actually is collected, respectively. Other setup tables enable further modification of the monitoring configuration. [Section 20.2.3.2, “Performance Schema Event Filtering”](#), discusses how you can modify these tables to affect event collection.

If there are Performance Schema configuration changes that must be made at runtime using SQL statements and you would like to these changes take effect each time the server starts, put the statements in a file and start the server with the `--init-file=file_name` [424] option. This strategy can also be useful if you have multiple monitoring configurations, each tailored to produce a different kind of monitoring, such as casual server health monitoring, incident investigation, application behavior troubleshooting, and so forth. Put the statements for each monitoring configuration into their own file and specify the appropriate file as the `--init-file` [424] argument when you start the server.

### 20.2.3.1. Performance Schema Event Timing

Events are collected by means of instrumentation added to the server source code. Instruments time events, which is how the Performance Schema provides an idea of how long events take. It is also possible to configure instruments not to collect timing information. This section discusses the available timers and their characteristics, and how timing values are represented in events.

Two tables provide timer information:

- `performance_timers` lists the available timers and their characteristics.
- `setup_timers` indicates which timers are used for which instruments.

Each timer row in `setup_timers` must refer to one of the timers listed in `performance_timers`.

Timers vary in precision and the amount of overhead they involve. To see what timers are available and their characteristics, check the `performance_timers` table:

```
mysql> SELECT * FROM performance_timers;
```

TIMER_NAME	TIMER_FREQUENCY	TIMER_RESOLUTION	TIMER_OVERHEAD
------------	-----------------	------------------	----------------

CYCLE	2389029850	1	72
NANOSECOND	NULL	NULL	NULL
MICROSECOND	1000000	1	585
MILLISECOND	1035	1	738
TICK	101	1	630

The `TIMER_NAME` column shows the names of the available timers. `CYCLE` refers to the timer that is based on the CPU (processor) cycle counter. If the values associated with a given timer name are `NULL`, that timer is not supported on your platform. The rows that do not have `NULL` indicate which timers you can use in `setup_timers`.

`TIMER_FREQUENCY` indicates the number of timer units per second. For a cycle timer, the frequency is generally related to the CPU speed. The value shown was obtained on a system with a 2.4GHz processor. The other timers are based on fixed fractions of seconds. For `TICK`, the frequency may vary by platform (for example, some use 100 ticks/second, others 1000 ticks/second).

`TIMER_RESOLUTION` indicates the number of timer units by which timer values increase at a time. If a timer has a resolution of 10, its value increases by 10 each time.

`TIMER_OVERHEAD` is the minimal number of cycles of overhead to obtain one timing with the given timer. The overhead per event is twice the value displayed because the timer is invoked at the beginning and end of the event.

To see which timer is in effect or to change the timer, access the `setup_timers` table:

```
mysql> SELECT * FROM setup_timers;
+-----+-----+
| NAME      | TIMER_NAME |
+-----+-----+
| idle      | MICROSECOND |
| wait      | NANOSECOND  |
| stage     | NANOSECOND  |
| statement | NANOSECOND  |
+-----+-----+

mysql> UPDATE setup_timers SET TIMER_NAME = 'MICROSECOND'
-> WHERE NAME = 'wait';
mysql> SELECT * FROM setup_timers;
+-----+-----+
| NAME      | TIMER_NAME |
+-----+-----+
| idle      | MICROSECOND |
| wait      | MICROSECOND |
| stage     | NANOSECOND  |
| statement | NANOSECOND  |
+-----+-----+
```

By default, the Performance Schema uses the best timer available for each instrument type, but you can select a different one.

To time waits, the most important criterion is to reduce overhead, at the possible expense of the timer accuracy, so using the `CYCLE` timer is the best.

The time a statement (or stage) takes to execute is in general orders of magnitude larger than the time it takes to execute a single wait. To time statements, the most important criterion is to have an accurate measure, which is not affected by changes in processor frequency, so using a timer which is not based on cycles is the best. The default timer for statements is `NANOSECOND`. The extra “overhead” compared to the cycle timer is not significant, because the overhead caused by calling a timer twice (once when the statement starts, once when it ends) is orders of magnitude less compared to the CPU time used to execute the statement itself. Using the `CYCLE` timer has no benefit here, only drawbacks.

The precision offered by the cycle counter depends on processor speed. If the processor runs at 1 GHz (one billion cycles/second) or higher, the cycle counter delivers sub-nanosecond precision. Using

the cycle counter is much cheaper than getting the actual time of day. For example, the standard `gettimeofday()` function can take hundreds of cycles, which is an unacceptable overhead for data gathering that may occur thousands or millions of times per second.

Cycle counters also have disadvantages:

- End users expect to see timings in wall-clock units, such as fractions of a second. Converting from cycles to fractions of seconds can be expensive. For this reason, the conversion is a quick and fairly rough multiplication operation.
- Processor cycle rate might change, such as when a laptop goes into power-saving mode or when a CPU slows down to reduce heat generation. If a processor's cycle rate fluctuates, conversion from cycles to real-time units is subject to error.
- Cycle counters might be unreliable or unavailable depending on the processor or the operating system. For example, on Pentiums, the instruction is `RDTSC` (an assembly-language rather than a C instruction) and it is theoretically possible for the operating system to prevent user-mode programs from using it.
- Some processor details related to out-of-order execution or multiprocessor synchronization might cause the counter to seem fast or slow by up to 1000 cycles.

Currently, MySQL works with cycle counters on x386 (Windows, Mac OS X, Linux, Solaris, and other Unix flavors), PowerPC, and IA-64.

The `setup_instruments` table has an `ENABLED` column to indicate the instruments for which to collect events. The table also has a `TIMED` column to indicate which instruments are timed. If an instrument is not enabled, it produces no events. If an enabled instrument is not timed, events produced by the instrument have `NULL` for the `TIMER_START`, `TIMER_END`, and `TIMER_WAIT` timer values. This in turn causes those values to be ignored when calculating the sum, minimum, maximum, and average time values in summary tables.

Within events, times are stored in picoseconds (trillionths of a second) to normalize them to a standard unit, regardless of which timer is selected. The timer used for an event is the one in effect when event timing begins. This timer is used to convert start and end values to picoseconds for storage in the event.

Modifications to the `setup_timers` table affect monitoring immediately. Events already in progress use the original timer for the begin time and the new timer for the end time, which leads to unpredictable results. If you make timer changes, you may want to use `TRUNCATE TABLE` to reset Performance Schema statistics.

The timer baseline (“time zero”) occurs at Performance Schema initialization during server startup. `TIMER_START` and `TIMER_END` values in events represent picoseconds since the baseline. `TIMER_WAIT` values are durations in picoseconds.

Picosecond values in events are approximate. Their accuracy is subject to the usual forms of error associated with conversion from one unit to another. If the `CYCLE` timer is used and the processor rate varies, there might be drift. For these reasons, it is not reasonable to look at the `TIMER_START` value for an event as an accurate measure of time elapsed since server startup. On the other hand, it is reasonable to use `TIMER_START` or `TIMER_WAIT` values in `ORDER BY` clauses to order events by start time or duration.

The choice of picoseconds in events rather than a value such as microseconds has a performance basis. One implementation goal was to show results in a uniform time unit, regardless of the timer. In an ideal world this time unit would look like a wall-clock unit and be reasonably precise; in other words, microseconds. But to convert cycles or nanoseconds to microseconds, it would be necessary to perform a division for every instrumentation. Division is expensive on many platforms. Multiplication is not expensive, so that is what is used. Therefore, the time unit is an integer multiple of the highest possible `TIMER_FREQUENCY` value, using a multiplier large enough to ensure that there is no major

precision loss. The result is that the time unit is “picoseconds.” This precision is spurious, but the decision enables overhead to be minimized.

### 20.2.3.2. Performance Schema Event Filtering

Events are processed in a producer/consumer fashion:

- Instrumented code is the source for events and produces events to be collected. The `setup_instruments` table lists the instruments for which events can be collected, whether they are enabled, and whether to collect timing information:

```
mysql> SELECT * FROM setup_instruments;
```

NAME	ENABLED	TIMED
wait/synch/mutex/sql/LOCK_global_read_lock	YES	YES
wait/synch/mutex/sql/LOCK_global_system_variables	YES	YES
wait/synch/mutex/sql/LOCK_lock_db	YES	YES
wait/synch/mutex/sql/LOCK_manager	YES	YES

The `setup_instruments` table provides the most basic form of control over event production. To further refine event production based on the type of object or thread being monitored, other tables may be used as described in [Section 20.2.3.2.1, “Event Pre-Filtering”](#).

- Performance Schema tables are the destinations for events and consume events. The `setup_consumers` table lists the types of consumers to which event information can be sent and which are enabled:

```
mysql> SELECT * FROM setup_consumers;
```

NAME	ENABLED
events_stages_current	NO
events_stages_history	NO
events_stages_history_long	NO
events_statements_current	YES
events_statements_history	NO
events_statements_history_long	NO
events_waits_current	NO
events_waits_history	NO
events_waits_history_long	NO
global_instrumentation	YES
thread_instrumentation	YES
statements_digest	YES

Filtering can be done at different stages of performance monitoring:

- Pre-filtering.** This is done by modifying Performance Schema configuration so that only certain types of events are collected from producers, and collected events update only certain consumers. This type of filtering is done by the Performance Schema and has a global effect that applies to all users.

Reasons to use pre-filtering:

- Pre-filtering reduces overhead. The overhead should be minimal even with all instruments enabled, but perhaps you want to reduce it further. Or you do not care about timing events and want to disable the timing code to eliminate timing overhead.
- You can avoid filling the current-events or history tables with events in which you have no interest. Pre-filtering leaves more “room” in these tables for instances of rows for enabled instrument types. If you enable only file instruments with pre-filtering, no rows are collected for nonfile instruments. With post-filtering, nonfile events are collected, leaving fewer rows for file events.

- You can avoid maintaining some kinds of event tables. If you disable a consumer, the server does not spend time maintaining it. For example, if you do not care about event histories, you can disable the history table consumers to improve performance.
- **Post-filtering.** This involves the use of [WHERE](#) clauses in queries that select information from Performance Schema tables, to specify which of the available events you want to see. This type of filtering is performed on a per-user basis because individual users select which of the available events are of interest.

Reasons to use post-filtering:

- To avoid making decisions for individual users about which event information is of interest.
- To use the Performance Schema to investigate a performance issue when the restrictions to impose using pre-filtering are not known in advance.

The following sections provide more detail about pre-filtering and provide guidelines for naming instruments or consumers in filtering operations. For information about writing queries to retrieve information (post-filtering), see [Section 20.3, “Performance Schema Queries”](#).

### 20.2.3.2.1. Event Pre-Filtering

Pre-filtering is done by the Performance Schema and has a global effect that applies to all users. Pre-filtering can be applied to either the producer or consumer stage of event processing:

- To configure pre-filtering at the producer stage, several tables can be used:
  - The [setup\\_instruments](#) table indicates which instruments are available. An instrument disabled in this table produces no events regardless of the contents of the other production-related setup tables. An instrument enabled in this table is permitted to produce events, subject to the contents of the other tables.
  - The [setup\\_objects](#) table determines whether particular objects are monitored. Currently, the Performance Schema uses it to control monitoring of table objects.
  - The [threads](#) table indicates whether monitoring is enabled for each server thread.
  - The [setup\\_actors](#) table determines the initial monitoring state for new foreground threads.
- To configure pre-filtering at the consumer stage, modify the [setup\\_consumers](#) table. This determines the destinations to which events are sent.

The [setup\\_consumers](#) table also implicitly affects event production. If a given event will not be sent to any destination (that is, will not be consumed), the Performance Schema does not produce it.

Modifications to any of these tables affect monitoring immediately, with the exception of [setup\\_actors](#). Modifications to [setup\\_actors](#) affect only foreground threads created thereafter.

When you change the monitoring configuration, the Performance Schema does not flush the history tables. Events already collected remain in the current-events and history tables until displaced by newer events. If you disable instruments, you might need to wait a while before events for them are displaced by newer events of interest. Alternatively, use [TRUNCATE TABLE](#) to empty the history tables.

After making instrumentation changes, you might want to truncate the summary tables as well to clear aggregate information for previously collected events. Except for [events\\_statements\\_summary\\_by\\_digest](#), the effect of [TRUNCATE TABLE](#) for summary tables is to reset the summary columns to 0 or [NULL](#), not to remove rows.

The following sections describe how to use specific tables to control how the Performance Schema performs pre-filtering.

### 20.2.3.2.1.1. Pre-Filtering by Instrument

The `setup_instruments` table lists the available instruments:

```
mysql> SELECT * FROM setup_instruments;
```

NAME	ENABLED	TIMED
...		
wait/synch/mutex/sql/LOCK_global_read_lock	YES	YES
wait/synch/mutex/sql/LOCK_global_system_variables	YES	YES
wait/synch/mutex/sql/LOCK_lock_db	YES	YES
wait/synch/mutex/sql/LOCK_manager	YES	YES
...		
wait/synch/rwlock/sql/LOCK_grant	YES	YES
wait/synch/rwlock/sql/LOGGER::LOCK_logger	YES	YES
wait/synch/rwlock/sql/LOCK_sys_init_connect	YES	YES
wait/synch/rwlock/sql/LOCK_sys_init_slave	YES	YES
...		
wait/io/file/sql/binlog	YES	YES
wait/io/file/sql/binlog_index	YES	YES
wait/io/file/sql/casetest	YES	YES
wait/io/file/sql/dbopt	YES	YES
...		

To control whether an instrument is enabled, set its `ENABLED` column to `YES` or `NO`. To configure whether to collect timing information for an instrument, set its `TIMED` value to `YES` or `NO`. Setting the `TIMED` column affects Performance Schema table contents as described in [Section 20.2.3.1, “Performance Schema Event Timing”](#).

The following examples demonstrate possible operations on the `setup_instruments` table. These changes, like other pre-filtering operations, affect all users.

- Disable all instruments:

```
mysql> UPDATE setup_instruments SET ENABLED = 'NO';
```

Now no events will be collected.

- Disable all file instruments, adding them to the current set of disabled instruments:

```
mysql> UPDATE setup_instruments SET ENABLED = 'NO'
-> WHERE NAME LIKE 'wait/io/file/%';
```

- Disable only file instruments, enable all other instruments:

```
mysql> UPDATE setup_instruments
-> SET ENABLED = IF(NAME LIKE 'wait/io/file/%', 'NO', 'YES');
```

The preceding queries use the `LIKE` [1190] operator and the pattern `'wait/io/file/%'` to match all instrument names that begin with `'wait/io/file/`. For additional information about specifying patterns to select instruments, see [Section 20.2.3.2.2, “Naming Instruments or Consumers for Filtering Operations”](#).

- Enable all but those instruments in the `mysys` library:

```
mysql> UPDATE setup_instruments
-> SET ENABLED = CASE WHEN NAME LIKE '%/mysys/%' THEN 'YES' ELSE 'NO' END;
```

- Disable a specific instrument:

```
mysql> UPDATE setup_instruments SET ENABLED = 'NO'
-> WHERE NAME = 'wait/synch/mutex/mysys/TMPDIR_mutex';
```

- To toggle the state of an instrument, “flip” its `ENABLED` value:

```
mysql> UPDATE setup_instruments
-> SET ENABLED = IF(ENABLED = 'YES', 'NO', 'YES');
```

```
-> WHERE NAME = 'wait/synch/mutex/mysys/TMPDIR_mutex';
```

- Disable timing for all events:

```
mysql> UPDATE setup_instruments SET TIMED = 'NO';
```

Modifications to the `setup_instruments` table affect monitoring immediately.

### 20.2.3.2.1.2. Pre-Filtering by Object

The `setup_objects` table controls whether particular objects are monitored. Currently, the Performance Schema uses it to control monitoring of table objects. The initial `setup_objects` contents look like this:

```
mysql> SELECT * FROM setup_objects;
+-----+-----+-----+-----+-----+
| OBJECT_TYPE | OBJECT_SCHEMA | OBJECT_NAME | ENABLED | TIMED |
+-----+-----+-----+-----+-----+
| TABLE     | mysql        | %           | NO      | NO    |
| TABLE     | performance_schema | %         | NO      | NO    |
| TABLE     | information_schema | %         | NO      | NO    |
| TABLE     | %            | %           | YES     | YES   |
+-----+-----+-----+-----+-----+
```

The `OBJECT_SCHEMA` and `OBJECT_NAME` columns should contain a literal schema or table name, or '%' to match any name.

When the Performance Schema checks for a match in `setup_objects`, it tries to find more specific matches first. For example, with a table `db1.t1`, it looks for a match for 'db1' and 't1', then for 'db1' and '%', then for '%' and '%'. The order in which matching occurs matters because different `setup_objects` rows can have different `ENABLED` and `TIMED` values.

The effect of the default object configuration is to instrument all tables except those in the `mysql`, `INFORMATION_SCHEMA`, and `performance_schema` databases. Tables in the `INFORMATION_SCHEMA` database are not instrumented regardless of the contents of `setup_objects`; the row for `information_schema.%` simply makes this default explicit.

For table-related events, the Performance Schema combines the contents of `setup_objects` with `setup_instruments` to determine whether to enable instruments and whether to time enabled instruments:

- For tables that match a row in `setup_objects`, table instruments produce events only if they are enabled in both `setup_instruments` and `setup_objects`.
- The `TIMED` values in the two tables are combined, so that timing information is collected only with both values are `YES`.

Suppose that `setup_objects` contains the following rows:

```
+-----+-----+-----+-----+-----+
| OBJECT_TYPE | OBJECT_SCHEMA | OBJECT_NAME | ENABLED | TIMED |
+-----+-----+-----+-----+-----+
| TABLE     | db1          | t1          | YES     | YES   |
| TABLE     | db1          | t2          | YES     | NO    |
| TABLE     | db2          | %           | YES     | YES   |
| TABLE     | db3          | %           | YES     | NO    |
| TABLE     | %            | %           | YES     | YES   |
+-----+-----+-----+-----+-----+
```

If a table-related instrument in `setup_instruments` has a `TIMED` value of `NO`, no events for the instrument are timed. If the `TIMED` value is `YES`, event timing occurs as follows:

- `db1.t1` events are timed
- `db1.t2` events are not timed



- `db2.t3` events are timed
- `db3.t4` events are not timed
- `db4.t5` events are timed

If a persistent table and a temporary table have the same name, matching against `setup_objects` rows occurs the same way for both. It is not possible to enable monitoring for one table but not the other. However, each table is instrumented separately.

Modifications to the `setup_objects` table affect object monitoring immediately.

The `ENABLED` column was added in MySQL 5.6.3. For earlier versions that have no `ENABLED` column, `setup_objects` is used only to enable monitoring for objects that match some row in the table. There is no way to explicitly disable instrumentation with the table.

### 20.2.3.2.1.3. Pre-Filtering by Thread

The `threads` table contains a row for each server thread. Each row contains information about a thread and indicates whether monitoring is enabled for it. For the Performance Schema to monitor a thread, these things must be true:

- The `thread_instrumentation` consumer in the `setup_consumers` table must be `YES`.
- The `thread.INSTRUMENTED` column must be `YES`.
- Monitoring occurs only for those thread events produced from instruments that are enabled, as specified in the `setup_instruments` table.

The `INSTRUMENTED` column in the `threads` table indicates the monitoring state for each thread. For foreground threads (resulting from client connections), the initial `INSTRUMENTED` value is determined by whether the user account associated with the thread matches any row in the `setup_actors` table. For background threads, `INSTRUMENTED` is `YES` by default. `setup_actors` is not consulted because there is no associated user for background threads. For any thread, its `INSTRUMENTED` value can be changed during the life of the thread.

The initial `setup_actors` contents look like this:

```
mysql> SELECT * FROM setup_actors;
+-----+-----+-----+
| HOST | USER | ROLE |
+-----+-----+-----+
| %    | %    | %    |
+-----+-----+-----+
```

The Performance Schema uses the `HOST` and `USER` columns to match each new foreground thread. (`ROLE` is unused.) The `INSTRUMENTED` value for the thread becomes `YES` if any row matches. This enables instrumenting to be applied selectively per host, user, or combination of host and user.

The `HOST` and `USER` columns should contain a literal host or user name, or `'%'` to match any name. By default, monitoring is enabled for all new foreground threads because the `setup_actors` table initially contains a row with `'%'` for both `HOST` and `USER`. To perform more limited matching such as to enable monitoring only for some foreground threads, you must delete this row because it matches any connection.

Suppose that you modify `setup_actors` as follows:

```
DELETE FROM setup_actors;
```

Now `setup_actors` is empty and there are no rows that could match incoming connections. Consequently, the Performance Schema will set the `INSTRUMENTED` column for all new foreground threads to `NO`.

Suppose that you further modify `setup_actors`:

```
INSERT INTO setup_actors (HOST,USER,ROLE) VALUES('localhost','joe','%');
INSERT INTO setup_actors (HOST,USER,ROLE) VALUES('%','sam','%');
```

Now the Performance Schema determines how to set the `INSTRUMENTED` value for new connection threads as follows:

- If `joe` connects from the local host, the connection matches the first inserted row.
- If `joe` connects from any other host, there is no match.
- If `sam` connects from any host, the connection matches the second inserted row.
- For any other connection, there is no match.

Modifications to the `setup_actors` table do not affect existing threads.

#### 20.2.3.2.1.4. Pre-Filtering by Consumer

The `setup_consumers` table lists the available consumer types and which are enabled:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                                | ENABLED |
+-----+-----+
| events_stages_current               | NO      |
| events_stages_history               | NO      |
| events_stages_history_long          | NO      |
| events_statements_current           | YES     |
| events_statements_history           | NO      |
| events_statements_history_long      | NO      |
| events_waits_current                | NO      |
| events_waits_history                | NO      |
| events_waits_history_long           | NO      |
| global_instrumentation              | YES     |
| thread_instrumentation              | YES     |
| statements_digest                   | YES     |
+-----+-----+
```

Modify the `setup_consumers` table to affect pre-filtering at the consumer stage and determine the destinations to which events are sent. To enable or disable a consumer, set its `ENABLED` value to `YES` or `NO`.

Modifications to the `setup_consumers` table affect monitoring immediately.

If you disable a consumer, the server does not spend time maintaining it. For example, you can disable history consumers if you do not care about historical event information:

```
mysql> UPDATE setup_consumers
-> SET ENABLED = 'NO' WHERE NAME LIKE '%history%';
```

The consumer settings in the `setup_consumers` table form a hierarchy from higher levels to lower. The following principles apply:

- Destinations associated with a consumer receive no events unless the Performance Schema checks the consumer and the consumer is enabled.
- A consumer is checked only if all consumers that it depends on (if any) are enabled.
- If a consumer is not checked, or is checked but is disabled, other consumers that depend on it are not checked.
- Dependent consumers may have their own dependent consumers.
- If an event would not be sent to any destination, the Performance Schema does not produce it.

The following lists describe the available consumer values. For discussion of several representative consumer configurations and their effect on instrumentation, see [Section 20.2.3.2.1.5, “Example Consumer Configurations”](#).

### Global and Thread Consumers

- `global_instrumentation` is the highest level consumer. If `global_instrumentation` is `NO`, it disables global instrumentation. All other settings are lower level and are not checked; it does not matter what they are set to. No global or per thread information is maintained and no individual events are collected in the current-events or event-history tables. If `global_instrumentation` is `YES`, the Performance Schema maintains information for global states and also checks the `thread_instrumentation` consumer.
- `thread_instrumentation` is checked only if `global_instrumentation` is `YES`. Otherwise, if `thread_instrumentation` is `NO`, it disables thread-specific instrumentation and all lower-level settings are ignored. No information is maintained per thread and no individual events are collected in the current-events or event-history tables. If `thread_instrumentation` is `YES`, the Performance Schema maintains thread-specific information and also checks `events_xxx_current` consumers.

### Statement Digest Consumer

This consumer requires `global_instrumentation` to be `YES` or it is not checked. There is no dependency on the Statement Event consumers. This means you can obtain statistics per digest without having to collect statistics in `events_statements_current`, which is advantageous in terms of overhead.

### Wait Event Consumers

These consumers require both `global_instrumentation` and `thread_instrumentation` to be `YES` or they are not checked.

- `events_waits_current`, if `NO`, disables collection of individual wait events in the `events_waits_current` table. If `YES`, it enables wait event collection and the Performance Schema checks the `events_waits_history` and `events_waits_history_long` consumers.
- `events_waits_history` is not checked if `event_waits_current` is `NO`. Otherwise, a `events_waits_history` value of `NO` or `YES` disables or enables collection of wait events in the `events_waits_history` table.
- `events_waits_history_long` is not checked if `event_waits_current` is `NO`. Otherwise, a `events_waits_history_long` value of `NO` or `YES` disables or enables collection of wait events in the `events_waits_history_long` table.

### Stage Event Consumers

These consumers require both `global_instrumentation` and `thread_instrumentation` to be `YES` or they are not checked.

- `events_stages_current`, if `NO`, disables collection of individual stage events in the `events_stages_current` table. If `YES`, it enables stage event collection and the Performance Schema checks the `events_stages_history` and `events_stages_history_long` consumers.
- `events_stages_history` is not checked if `event_stages_current` is `NO`. Otherwise, a `events_stages_history` value of `NO` or `YES` disables or enables collection of stage events in the `events_stages_history` table.
- `events_stages_history_long` is not checked if `event_stages_current` is `NO`. Otherwise, a `events_stages_history_long` value of `NO` or `YES` disables or enables collection of stage events in the `events_stages_history_long` table.

## Statement Event Consumers

These consumers require both `global_instrumentation` and `thread_instrumentation` to be `YES` or they are not checked.

- `events_statements_current`, if `NO`, disables collection of individual statement events in the `events_statements_current` table. If `YES`, it enables statement event collection and the Performance Schema checks the `events_statements_history` and `events_statements_history_long` consumers.
- `events_statements_history` is not checked if `events_statements_current` is `NO`. Otherwise, a `events_statements_history` value of `NO` or `YES` disables or enables collection of statement events in the `events_statements_history` table.
- `events_statements_history_long` is not checked if `events_statements_current` is `NO`. Otherwise, a `events_statements_history_long` value of `NO` or `YES` disables or enables collection of statement events in the `events_statements_history_long` table.

### 20.2.3.2.1.5. Example Consumer Configurations

The consumer settings in the `setup_consumers` table form a hierarchy from higher levels to lower. The following discussion describes how consumers work, showing specific configurations and their effects as consumer settings are enabled progressively from high to low. The consumer values shown are representative. The general principles described here apply to other consumer values that may be available.

The configuration descriptions occur in order of increasing functionality and overhead. If you do not need the information provided by enabling lower-level settings, disable them and the Performance Schema will execute less code on your behalf and you will have less information to sift through.

Suppose that the `setup_consumers` table contains the following hierarchy of values:

```
global_instrumentation
  thread_instrumentation
    events_waits_current
      events_waits_history
        events_waits_history_long
```

If a given consumer setting is `NO`, the Performance Schema disables the instrumentation associated with the consumer and ignores all lower-level settings. If a given setting is `YES`, the Performance Schema enables the instrumentation associated with it and checks the settings at the next lowest level.

Each configuration description indicates which setup elements the Performance Schema checks and which output tables it maintains (that is, for which tables it collects information).

#### No Instrumentation

Server configuration state:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                                | ENABLED |
+-----+-----+
| global_instrumentation              | NO      |
| ...                                  |         |
+-----+-----+
```

In this configuration, nothing is instrumented.

Setup elements checked:

- Table `setup_consumers`, consumer `global_instrumentation`

Output tables maintained:

- None

### Global Instrumentation Only

Server configuration state:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                | ENABLED |
+-----+-----+
| global_instrumentation | YES     |
| thread_instrumentation | NO      |
| ...                  |         |
+-----+-----+
```

In this configuration, instrumentation is maintained only for global states. Per-thread instrumentation is disabled.

Additional setup elements checked, relative to the preceding configuration:

- Table `setup_consumers`, consumer `thread_instrumentation`
- Table `setup_instruments`
- Table `setup_objects`
- Table `setup_timers`

Additional output tables maintained, relative to the preceding configuration:

- `mutex_instances`
- `rwlock_instances`
- `cond_instances`
- `file_instances`
- `file_summary_by_instance`
- `file_summary_by_event_name`
- `objects_summary_global_by_type`
- `table_lock_waits_summary_by_table`
- `table_io_waits_summary_by_index_usage`
- `table_io_waits_summary_by_table`
- `events_waits_summary_by_instance`
- `events_waits_summary_global_by_event_name`

### Global and Thread Instrumentation Only

Server configuration state:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                | ENABLED |
+-----+-----+
| global_instrumentation | YES     |
| thread_instrumentation | YES     |
| events_waits_current   | NO      |
+-----+-----+
```

```
...
+-----+-----+
```

In this configuration, instrumentation is maintained globally and per thread. No individual wait events are collected in the current-events or event-history tables.

Additional setup elements checked, relative to the preceding configuration:

- Table `setup_consumers`, consumer `events_waits_current`
- Table `setup_actors`
- Column `threads.instrumented`

Additional output tables maintained, relative to the preceding configuration:

- `events_waits_summary_by_thread_by_event_name`

### Global, Thread, and Current-Event Instrumentation

Server configuration state:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                                | ENABLED |
+-----+-----+
| global_instrumentation              | YES     |
| thread_instrumentation              | YES     |
| events_waits_current                 | YES     |
| events_waits_history                 | NO      |
| events_waits_history_long            | NO      |
...
+-----+-----+
```

In this configuration, instrumentation is maintained globally and per thread. Individual wait events are collected in the current-events table, but not in the event-history tables.

Additional setup elements checked, relative to the preceding configuration:

- Consumer `events_waits_history`
- Consumer `events_waits_history_long`

Additional output tables maintained, relative to the preceding configuration:

- `events_waits_current`

### Global, Thread, Current-Event, and Event-History instrumentation

The preceding configuration collects no wait event history because the `events_waits_history` and `events_waits_history_long` consumers are disabled. Those consumers can be enabled separately or together to collect event history per thread, globally, or both.

This configuration collects event history per thread, but not globally:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                                | ENABLED |
+-----+-----+
| global_instrumentation              | YES     |
| thread_instrumentation              | YES     |
| events_waits_current                 | YES     |
| events_waits_history                 | YES     |
| events_waits_history_long            | NO      |
...
+-----+-----+
```

Event-history tables maintained for this configuration:

- `events_waits_history`

This configuration collects event history globally, but not per thread:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                | ENABLED |
+-----+-----+
| global_instrumentation | YES     |
| thread_instrumentation | YES     |
| events_waits_current   | YES     |
| events_waits_history   | NO      |
| events_waits_history_long | YES    |
| ...                   |         |
+-----+-----+
```

Event-history tables maintained for this configuration:

- `events_waits_history_long`

This configuration collects event history per thread and globally:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME                | ENABLED |
+-----+-----+
| global_instrumentation | YES     |
| thread_instrumentation | YES     |
| events_waits_current   | YES     |
| events_waits_history   | YES     |
| events_waits_history_long | YES    |
| ...                   |         |
+-----+-----+
```

Event-history tables maintained for this configuration:

- `events_waits_history`
- `events_waits_history_long`

### 20.2.3.2.2. Naming Instruments or Consumers for Filtering Operations

Names given for filtering operations can be as specific or general as required. To indicate a single instrument or consumer, specify its name in full:

```
mysql> UPDATE setup_instruments
  -> SET ENABLED = 'NO'
  -> WHERE NAME = 'wait/synch/mutex/myisamrg/MYRG_INFO::mutex';

mysql> UPDATE setup_consumers
  -> SET ENABLED = 'NO' WHERE NAME = 'events_waits_current';
```

To specify a group of instruments or consumers, use a pattern that matches the group members:

```
mysql> UPDATE setup_instruments
  -> SET ENABLED = 'NO'
  -> WHERE NAME LIKE 'wait/synch/mutex/%';

mysql> UPDATE setup_consumers
  -> SET ENABLED = 'NO' WHERE NAME LIKE '%history%';
```

If you use a pattern, it should be chosen so that it matches all the items of interest and no others. For example, to select all file I/O instruments, it is better to use a pattern that includes the entire instrument name prefix:

```
... WHERE NAME LIKE 'wait/io/file/%';
```



A pattern of `'%/file/%'` will match other instruments that have a component of `'/file/'` anywhere in the name. Even less suitable is the pattern `'%file%'` because it will match instruments with `'file'` anywhere in the name, such as `wait/synch/mutex/sql/LOCK_des_key_file`.

To check which instrument or consumer names a pattern matches, perform a simple test:

```
mysql> SELECT NAME FROM setup_instruments WHERE NAME LIKE 'pattern';
mysql> SELECT NAME FROM setup_consumers WHERE NAME LIKE 'pattern';
```

For information about the types of names that are supported, see [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

### 20.2.3.3. Determining What Is Instrumented

It is always possible to determine what instruments the Performance Schema includes by checking the `setup_instruments` table. For example, to see what file-related events are instrumented for the InnoDB storage engine, use this query:

```
mysql> SELECT * FROM setup_instruments WHERE NAME LIKE 'wait/io/file/innodb/%';
```

NAME	ENABLED	TIMED
wait/io/file/innodb/innodb_data_file	YES	YES
wait/io/file/innodb/innodb_log_file	YES	YES
wait/io/file/innodb/innodb_temp_file	YES	YES

An exhaustive description of precisely what is instrumented is not given in this documentation, for several reasons:

- What is instrumented is the server code. Changes to this code occur often, which also affects the set of instruments.
- It is not practical to list all the instruments because there are hundreds of them.
- As described earlier, it is possible to find out by querying the `setup_instruments` table. This information is always up to date for your version of MySQL, also includes instrumentation for instrumented plugins you might have installed that are not part of the core server, and can be used by automated tools.

## 20.3. Performance Schema Queries

Pre-filtering limits which event information is collected and is independent of any particular user. By contrast, post-filtering is performed by individual users through the use of queries with appropriate `WHERE` clauses that restrict what event information to select from the events available after pre-filtering has been applied.

In [Section 20.2.3.2.1, “Event Pre-Filtering”](#), an example showed how to pre-filter for file instruments. If the event tables contain both file and nonfile information, post-filtering is another way to see information only for file events. Add a `WHERE` clause to queries to restrict event selection appropriately:

```
mysql> SELECT THREAD_ID, NUMBER_OF_BYTES
-> FROM events_waits_history
-> WHERE EVENT_NAME LIKE 'wait/io/file/%'
-> AND NUMBER_OF_BYTES IS NOT NULL;
```

THREAD_ID	NUMBER_OF_BYTES
11	66
11	47
11	139
5	24
5	834

+-----+-----+-----+

## 20.4. Performance Schema Instrument Naming Conventions

An instrument name consists of a sequence of components separated by '/' characters. Example names:

```
wait/io/file/myisam/log
wait/io/file/mysys/charset
wait/lock/table/sql/handler
wait/synch/cond/mysys/COND_alarm
wait/synch/cond/sql/BINLOG::update_cond
wait/synch/mutex/mysys/BITMAP_mutex
wait/synch/mutex/sql/LOCK_delete
wait/synch/rwlock/sql/Query_cache_query::lock
stage/sql/closing_tables
stage/sql/Sorting_result
statement/com/Execute
statement/com/Query
statement/sql/create_table
statement/sql/lock_tables
```

The instrument name space has a tree-like structure. The components of an instrument name from left to right provide a progression from more general to more specific. The number of components a name has depends on the type of instrument.

The interpretation of a given component in a name depends on the components to the left of it. For example, `myisam` appears in both of the following names, but `myisam` in the first name is related to file I/O, whereas in the second it is related to a synchronization instrument:

```
wait/io/file/myisam/log
wait/synch/cond/myisam/MI_SORT_INFO::cond
```

Instrument names consist of a prefix with a structure defined by the Performance Schema implementation and a suffix defined by the developer implementing the instrument code. The top-level component of an instrument prefix indicates the type of instrument. This component also determines which event timer in the `setup_timers` table applies to the instrument. For the prefix part of instrument names, the top level indicates the type of instrument.

The suffix part of instrument names comes from the code for the instruments themselves. Suffixes may include levels such as these:

- A name for the major component (a server module such as `myisam`, `innodb`, `mysys`, or `sql`) or a plugin name.
- The name of a variable in the code, in the form `XXX` (a global variable) or `CCC::MMM` (a member `MMM` in class `CCC`). Examples: `COND_thread_cache`, `THR_LOCK_myisam`, `BINLOG::LOCK_index`.

### Top-Level Instrument Components

- `idle`: An instrumented idle event. This instrument has no further components.
- `stage`: An instrumented stage event.
- `statement`: An instrumented statement event.
- `wait`: An instrumented wait event.

### Idle Instrument Components

- `idle`

The idle instrument. The Performance Schema generates idle events as discussed in the description of the `socket_instances.STATE` column in [Section 20.9.2.5, "The socket\\_instances Table"](#).

## Stage Instrument Components

Stage instruments have names of the form `stage/code_area/stage_name`, where `code_area` is a value such as `sql` or `myisam`, and `stage_name` indicates the stage of statement processing, such as `Sorting result` or `Sending data`. Stages correspond to the thread states displayed by `SHOW PROCESSLIST` or that are visible in the `INFORMATION_SCHEMA.PROCESSLIST` table.

## Statement Instrument Components

- `statement/com`: An instrumented command operation. These have names corresponding to `COM_xxx` operations (see the `mysql_com.h` header file and `sql/sql_parse.cc`). For example, the `statement/com/Connect` and `statement/com/Fetch` instruments correspond to the `COM_CONNECT` and `COM_FETCH` commands.
- `statement/sql`: An instrumented SQL statement operation. For example, the `statement/sql/create_db` and `statement/sql/select` instruments are used for `CREATE DATABASE` and `SELECT` statements.

## Wait Instrument Components

- `wait/io`

An instrumented I/O operation.

- `wait/io/file`

An instrumented file I/O operation. For files, the wait is the time waiting for the file operation to complete (for example, a call to `fwrite()`). Due to caching, the physical file I/O on the disk might not happen within this call.

- `wait/io/socket`

An instrumented socket operation. Socket instruments have names of the form `wait/io/socket/sql/socket_type`. The server has a listening socket for each network protocol that it supports. The instruments associated with listening sockets for TCP/IP or Unix socket file connections have a `socket_type` value of `server_tcpip_socket` or `server_unix_socket`, respectively. When a listening socket detects a connection, the server transfers the connection to a new socket managed by a separate thread. The instrument for the new connection thread has a `socket_type` value of `client_connection`.

- `wait/io/table`

An instrumented table I/O operation. These include row-level accesses to persistent base tables or temporary tables. Operations that affect rows are fetch, insert, update, and delete. For a view, waits are associated with base tables referenced by the view.

Unlike most waits, a table I/O wait can include other waits. For example, table I/O might include file I/O or memory operations. Thus, `events_waits_current` for a table I/O wait usually has two rows. For more information, see [Section 20.6, “Performance Schema Atom and Molecule Events”](#).

Some row operations might cause multiple table I/O waits. For example, an insert might activate a trigger that causes an update.

- `wait/lock`

An instrumented lock operation.

- `wait/lock/table`

An instrumented table lock operation.

- `wait/synch`

An instrumented synchronization object. For synchronization objects, the `TIMER_WAIT` time includes the amount of time blocked while attempting to acquire a lock on the object, if any.

- `wait/synch/cond`

A condition is used by one thread to signal to other threads that something they were waiting for has happened. If a single thread was waiting for a condition, it can wake up and proceed with its execution. If several threads were waiting, they can all wake up and compete for the resource for which they were waiting.

- `wait/synch/mutex`

A mutual exclusion object used to permit access to a resource (such as a section of executable code) while preventing other threads from accessing the resource.

- `wait/synch/rwlock`

A read/write lock object used to lock a specific variable for access while preventing its use by other threads. A shared read lock can be acquired simultaneously by multiple threads. An exclusive write lock can be acquired by only one thread at a time.

## 20.5. Performance Schema Status Monitoring

There are several status variables associated with the Performance Schema:

```
mysql> SHOW STATUS LIKE 'perf%';
```

Variable_name	Value
Performance_schema_cond_classes_lost	0
Performance_schema_cond_instances_lost	0
Performance_schema_file_classes_lost	0
Performance_schema_file_handles_lost	0
Performance_schema_file_instances_lost	0
Performance_schema_locker_lost	0
Performance_schema_mutex_classes_lost	0
Performance_schema_mutex_instances_lost	0
Performance_schema_rwlock_classes_lost	0
Performance_schema_rwlock_instances_lost	0
Performance_schema_table_handles_lost	0
Performance_schema_table_instances_lost	0
Performance_schema_thread_classes_lost	0
Performance_schema_thread_instances_lost	0

The Performance Schema status variables provide information about instrumentation that could not be loaded or created due to memory constraints. Names for these variables have several forms:

- `Performance_schema_xxx_classes_lost` indicates how many instruments of type `xxx` could not be loaded.
- `Performance_schema_xxx_instances_lost` indicates how many instances of object type `xxx` could not be created.
- `Performance_schema_xxx_handles_lost` indicates how many instances of object type `xxx` could not be opened.
- `Performance_schema_locker_lost` indicates how many events are “lost” or not recorded.

For example, if a mutex is instrumented in the server source but the server cannot allocate memory for the instrumentation at runtime, it increments `Performance_schema_mutex_classes_lost` [2318]. The mutex still functions as a synchronization object (that is, the server continues to function normally), but performance data for it will not be collected. If the instrument can be allocated, it can be used for

initializing instrumented mutex instances. For a singleton mutex such as a global mutex, there will be only one instance. Other mutexes have an instance per connection, or per page in various caches and data buffers, so the number of instances varies over time. Increasing the maximum number of connections or the maximum size of some buffers will increase the maximum number of instances that might be allocated at once. If the server cannot create a given instrumented mutex instance, it increments `Performance_schema_mutex_instances_lost` [2318].

Suppose that the following conditions hold:

- The server was started with the `--performance_schema_max_mutex_classes=200` [2312] option and thus has room for 200 mutex instruments.
- 150 mutex instruments have been loaded already.
- The plugin named `plugin_a` contains 40 mutex instruments.
- The plugin named `plugin_b` contains 20 mutex instruments.

The server allocates mutex instruments for the plugins depending on how many they need and how many are available, as illustrated by the following sequence of statements:

```
INSTALL PLUGIN plugin_a
```

The server now has  $150+40 = 190$  mutex instruments.

```
UNINSTALL PLUGIN plugin_a;
```

The server still has 190 instruments. All the historical data generated by the plugin code is still available, but new events for the instruments are not collected.

```
INSTALL PLUGIN plugin_a;
```

The server detects that the 40 instruments are already defined, so no new instruments are created, and previously assigned internal memory buffers are reused. The server still has 190 instruments.

```
INSTALL PLUGIN plugin_b;
```

The server has room for  $200-190 = 10$  instruments (in this case, mutex classes), and sees that the plugin contains 20 new instruments. 10 instruments are loaded, and 10 are discarded or “lost.” The `Performance_schema_mutex_classes_lost` [2318] indicates the number of instruments (mutex classes) lost:

```
mysql> SHOW STATUS LIKE "perf%mutex_classes_lost";
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Performance_schema_mutex_classes_lost | 10 |
+-----+-----+
1 row in set (0.10 sec)
```

The instrumentation still works and collects (partial) data for `plugin_b`.

When the server cannot create a mutex instrument, these results occur:

- No row for the instrument is inserted into the `setup_instruments` table.
- `Performance_schema_mutex_classes_lost` [2318] increases by 1.
- `Performance_schema_mutex_instances_lost` [2318] does not change. (When the mutex instrument is not created, it cannot be used to create instrumented mutex instances later.)

The pattern just described applies to all types of instruments, not just mutexes.

A value of `Performance_schema_mutex_classes_lost` [2318] greater than 0 can happen in two cases:

- To save a few bytes of memory, you start the server with `--performance_schema_max_mutex_classes=N` [2312], where *N* is less than the default value. The default value is chosen to be sufficient to load all the plugins provided in the MySQL distribution, but this can be reduced if some plugins are never loaded. For example, you might choose not to load some of the storage engines in the distribution.
- You load a third-party plugin that is instrumented for the Performance Schema but do not allow for the plugin's instrumentation memory requirements when you start the server. Because it comes from a third party, the instrument memory consumption of this engine is not accounted for in the default value chosen for `performance_schema_max_mutex_classes` [2312].

If the server has insufficient resources for the plugin's instruments and you do not explicitly allocate more using `--performance_schema_max_mutex_classes=N` [2312], loading the plugin leads to starvation of instruments.

If the value chosen for `performance_schema_max_mutex_classes` [2312] is too small, no error is reported in the error log and there is no failure at runtime. However, the content of the tables in the `performance_schema` database will miss events. The `Performance_schema_mutex_classes_lost` [2318] status variable is the only visible sign to indicate that some events were dropped internally due to failure to create instruments.

If an instrument is not lost, it is known to the Performance Schema, and is used when instrumenting instances. For example, `wait/synch/mutex/sql/LOCK_delete` is the name of a mutex instrument in the `setup_instruments` table. This single instrument is used when creating a mutex in the code (in `THD::LOCK_delete`) however many instances of the mutex are needed as the server runs. In this case, `LOCK_delete` is a mutex that is per connection (`THD`), so if a server has 1000 connections, there are 1000 threads, and 1000 instrumented `LOCK_delete` mutex instances (`THD::LOCK_delete`).

If the server does not have room for all these 1000 instrumented mutexes (instances), some mutexes are created with instrumentation, and some are created without instrumentation. If the server can create only 800 instances, 200 instances are lost. The server continues to run, but increments `Performance_schema_mutex_instances_lost` [2318] by 200 to indicate that instances could not be created.

A value of `Performance_schema_mutex_instances_lost` [2318] greater than 0 can happen when the code initializes more mutexes at runtime than were allocated for `--performance_schema_max_mutex_instances=N` [2313].

The bottom line is that if `SHOW STATUS LIKE 'perf%'` says that nothing was lost (all values are zero), the Performance Schema data is accurate and can be relied upon. If something was lost, the data is incomplete, and the Performance Schema could not record everything given the insufficient amount of memory it was given to use. In this case, the specific `Performance_schema_xxx_lost` variable indicates the problem area.

It might be appropriate in some cases to cause deliberate instrument starvation. For example, if you do not care about performance data for file I/O, you can start the server with all Performance Schema parameters related to file I/O set to 0. No memory will be allocated for file-related classes, instances, or handles, and all file events will be lost.

Use `SHOW ENGINE PERFORMANCE_SCHEMA STATUS` to inspect the internal operation of the Performance Schema code:

```
mysql> SHOW ENGINE PERFORMANCE_SCHEMA STATUS\G
...
***** 3. row *****
  Type: performance_schema
  Name: events_waits_history.row_size
  Status: 76
***** 4. row *****
  Type: performance_schema
  Name: events_waits_history.row_count
```

```

Status: 10000
***** 5. row *****
Type: performance_schema
Name: events_waits_history.memory
Status: 760000
...
***** 57. row *****
Type: performance_schema
Name: performance_schema.memory
Status: 26459600
...

```

This statement is intended to help the DBA understand the effects that different Performance Schema options have on memory requirements. For a description of the field meanings, see [Section 13.7.5.16](#), “`SHOW ENGINE Syntax`”.

## 20.6. Performance Schema Atom and Molecule Events

For a table I/O event, there are usually two rows in `events_waits_current`, not one. For example, a row fetch might result in rows like this:

Row#	EVENT_NAME	TIMER_START	TIMER_END
1	wait/io/file/myisam/dfile	10001	10002
2	wait/io/table/sql/handler	10000	NULL

The row fetch causes a file read. In the example, the table I/O fetch event started before the file I/O event but has not finished (its `TIMER_END` value is `NULL`). The file I/O event is “nested” within the table I/O event.

This occurs because, unlike other “atomic” wait events such as for mutexes or file I/O, table I/O events are “molecular” and include (overlap with) other events. In `events_waits_current`, the table I/O event usually has two rows:

- One row for the most recent table I/O wait event
- One row for the most recent wait event of any kind

Usually, but not always, the “of any kind” wait event differs from the table I/O event. As each subsidiary event completes, it disappears from `events_waits_current`. At this point, and until the next subsidiary event begins, the table I/O wait is also the most recent wait of any kind.

## 20.7. Performance Schema Statement Digests

As of MySQL 5.6.5, the Performance Schema maintains statement digest information. Digesting converts a SQL statement to normalized form and computes a hash value for the result. Normalization permits statements that are similar to be grouped and summarized to expose information about the types of statements the server is executing and how often they occur. This section describes how statement normalizing occurs and how it can be useful.

Statement digesting involves these Performance Schema components:

- A `statement_digest` consumer in the `setup_consumers` table controls whether the Performance Schema maintains digest information.
- The statement event tables (`events_statements_current`, `events_statements_history`, and `events_statements_history_long`) have `DIGEST` and `DIGEST_TEXT` columns that contain digest MD5 values and the corresponding normalized statement text strings.
- A `events_statements_summary_by_digest` table provides aggregated statement digest information.

Normalizing a statement transforms the statement text to a more standardized string representation that preserves the general statement structure while removing information not essential to the



structure. Object identifiers such as database and table names are preserved. Values and comments are removed, and whitespace is adjusted. The Performance Schema does not retain information such as names, passwords, dates, and so forth.

Consider these statements:

```
SELECT * FROM orders WHERE customer_id=10 AND quantity>20
SELECT * FROM orders WHERE customer_id = 20 AND quantity > 100
```

To normalize these statements, the Performance Schema replaces data values by ? and adjusts whitespace. Both statements yield the same normalized form and thus are considered “the same”:

```
SELECT * FROM orders WHERE customer_id = ? AND quantity > ?
```

The normalized statement contains less information but is still representative of the original statement. Other similar statements that have different comparison values have the same normalized form.

Now consider these statements:

```
SELECT * FROM customers WHERE customer_id = 1000
SELECT * FROM orders WHERE customer_id = 1000
```

In this case, the statements are not “the same.” The object identifiers differ, so the statements yield different normalized forms:

```
SELECT * FROM customers WHERE customer_id = ?
SELECT * FROM orders WHERE customer_id = ?
```

Normalized statements have a fixed length. The maximum length of a `DIGEST_TEXT` value is 1024 bytes. There is no option to change this maximum. If normalization produces a statement that exceeds this length, the text ends with “...”. Long statements that differ only in the part that occurs following the “...” are considered to be the same. Consider these statements:

```
SELECT * FROM mytable WHERE cola = 10 AND colb = 20
SELECT * FROM mytable WHERE cola = 10 AND colc = 20
```

If the cutoff happened to be right after the `AND`, both statements would have this normalized form:

```
SELECT * FROM mytable WHERE cola = ? AND ...
```

In this case, the difference in the second column name is lost and both statements are considered the same.

For each normalized statement, the Performance Schema computes a hash digest value and stores that value and the statement in the `DIGEST` and `DIGEST_TEXT` columns of the statement event tables (`events_statements_current`, `events_statements_history`, and `events_statements_history_long`). In addition, information for statements with the same `SCHEMA_NAME` and `DIGEST` values are aggregated in the `events_statements_summary_by_digest` summary table. The Performance Schema uses MD5 hash values because they are fast to compute and have a favorable statistical distribution that minimizes collisions.

The `events_statements_summary_by_digest` summary table has a fixed size, so when it becomes full, statements that have `SCHEMA_NAME` and `DIGEST` values not matching existing values in the table are grouped in a special row with `SCHEMA_NAME` and `DIGEST` set to `NULL`. This permits all statements to be counted. However, if the special row accounts for a significant percentage of the statements executed, it might be desirable to increase the size of the summary table. To do this, set the `performance_schema_digests_size` [2308] system variable to a larger value at server startup. If no `performance_schema_digests_size` [2308] value is given, the server estimates the value to use at startup. (Before MySQL 5.6.9, there is no `SCHEMA_NAME` column and the special row has `DIGEST` set to `NULL`.)

The statement digest summary table provides a profile of the statements executed by the server. It shows what kinds of statements an application is executing and how often. An application developer

can use this information together with other information in the table to assess the application's performance characteristics. For example, table columns that show wait times, lock times, or index use may highlight types of queries that are inefficient. This gives the developer insight into which parts of the application need attention.

## 20.8. Performance Schema General Table Characteristics

The name of the `performance_schema` database is lowercase, as are the names of tables within it. Queries should specify the names in lowercase.

Most tables in the `performance_schema` database are read only and cannot be modified. Some of the setup tables have columns that can be modified to affect Performance Schema operation; some also permit rows to be inserted or deleted. Truncation is permitted to clear collected events, so `TRUNCATE TABLE` can be used on tables containing those kinds of information, such as tables named with a prefix of `events_waits_`.

`TRUNCATE TABLE` can also be used with summary tables, but except for `events_statements_summary_by_digest`, the effect is to reset the summary columns to 0 or `NULL`, not to remove rows.

Privileges are as for other databases and tables:

- To retrieve from `performance_schema` tables, you must have the `SELECT` [755] privilege.
- To change those columns that can be modified, you must have the `UPDATE` [756] privilege.
- To truncate tables that can be truncated, you must have the `DROP` [754] privilege.

## 20.9. Performance Schema Table Descriptions

Tables in the `performance_schema` database can be grouped as follows:

- Setup tables. These tables are used to configure and display monitoring characteristics.
- Current events table. The `events_waits_current` table contains the most recent event for each thread.
- History tables. These tables have the same structure as `events_waits_current` but contain more rows. The `events_waits_history` table contains the most recent 10 events per thread. `events_waits_history_long` contains the most recent 10,000 events.

To change the sizes of these tables, set the

`performance_schema_events_waits_history_size` [2310] and `performance_schema_events_waits_history_long_size` [2310] system variables at server startup.

- Summary tables. These tables contain information aggregated over groups of events, including those that have been discarded from the history tables.
- Instance tables. These tables document what types of objects are instrumented. An instrumented object, when used by the server, produces an event. These tables provide event names and explanatory notes or status information.
- Miscellaneous tables. These do not fall into any of the other table groups.

### 20.9.1. Performance Schema Setup Tables

The setup tables provide information about the current instrumentation and enable the monitoring configuration to be changed. For this reason, some columns in these tables can be changed if you have the `UPDATE` [756] privilege.

The use of tables rather than individual variables for setup information provides a high degree of flexibility in modifying Performance Schema configuration. For example, you can use a single statement with standard SQL syntax to make multiple simultaneous configuration changes.

These setup tables are available:

- `setup_actors`: How to initialize monitoring for new foreground threads
- `setup_consumers`: The destinations to which event information can be sent and stored
- `setup_instruments`: The classes of instrumented objects for which events can be collected
- `setup_objects`: Which objects should be monitored
- `setup_timers`: The current event timer

### 20.9.1.1. The `setup_actors` Table

The `setup_actors` table contains information that determines whether to enable monitoring for new foreground server threads; that is, threads associated with client connections. This table has a maximum size of 100 rows by default. The size may be changed by modifying the `performance_schema_setup_actors_size` [2317] system variable at server startup.

The `setup_actors` table has these columns:

- `HOST`  
The host name. This should be a literal name, or `'%'` to mean “any host.”
- `USER`  
The user name. This should be a literal name, or `'%'` to mean “any user.”
- `ROLE`  
Unused.

For each new foreground thread, the Performance Schema matches the user and host for the thread against the rows of the `setup_actors` table. Based on whether any row matches, the `INSTRUMENTED` column of the `threads` table row for the thread is set to `YES` or `NO`. This enables instrumenting to be applied selectively per host, user, or combination of host and user.

The initial contents of the `setup_actors` table match any user and host combination, so monitoring for all foreground threads is enabled by default:

```
mysql> SELECT * FROM setup_actors;
+-----+-----+-----+
| HOST | USER | ROLE |
+-----+-----+-----+
| %    | %    | %    |
+-----+-----+-----+
```

Modifications to the `setup_actors` table do not affect existing threads.

For more information about the role of the `setup_actors` table in event monitoring, see [Section 20.2.3.2.1, “Event Pre-Filtering”](#).

### 20.9.1.2. The `setup_consumers` Table

The `setup_consumers` table lists the types of consumers for which event information can be stored and which are enabled:

```
mysql> SELECT * FROM setup_consumers;
+-----+-----+
| NAME | ENABLED |
+-----+-----+
```

events_stages_current	NO
events_stages_history	NO
events_stages_history_long	NO
events_statements_current	YES
events_statements_history	NO
events_statements_history_long	NO
events_waits_current	NO
events_waits_history	NO
events_waits_history_long	NO
global_instrumentation	YES
thread_instrumentation	YES

The consumer settings in the `setup_consumers` table form a hierarchy from higher levels to lower. For detailed information about the effect of enabling different consumers, see [Section 20.2.3.2.1.4, “Pre-Filtering by Consumer”](#).

The `setup_consumers` table has these columns:

- **NAME**

The consumer name.

- **ENABLED**

Whether the consumer is enabled. This column can be modified. If you disable a consumer, the server does not spend time adding event information to it.

Modifications to the `setup_consumers` table affect monitoring immediately.

### 20.9.1.3. The `setup_instruments` Table

The `setup_instruments` table lists classes of instrumented objects for which events can be collected:

```
mysql> SELECT * FROM setup_instruments;
```

NAME	ENABLED	TIMED
...		
wait/synch/mutex/sql/LOCK_global_read_lock	YES	YES
wait/synch/mutex/sql/LOCK_global_system_variables	YES	YES
wait/synch/mutex/sql/LOCK_lock_db	YES	YES
wait/synch/mutex/sql/LOCK_manager	YES	YES
...		
wait/synch/rwlock/sql/LOCK_grant	YES	YES
wait/synch/rwlock/sql/LOGGER::LOCK_logger	YES	YES
wait/synch/rwlock/sql/LOCK_sys_init_connect	YES	YES
wait/synch/rwlock/sql/LOCK_sys_init_slave	YES	YES
...		
wait/io/file/sql/binlog	YES	YES
wait/io/file/sql/binlog_index	YES	YES
wait/io/file/sql/casetest	YES	YES
wait/io/file/sql/dbopt	YES	YES
...		

Each instrument added to the source code provides a row for this table, even when the instrumented code is not executed. When an instrument is enabled and executed, instrumented instances are created, which are visible in the `*_instances` tables.

The `setup_instruments` table has these columns:

- **NAME**

The instrument name. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, “Performance Schema Instrument Naming Conventions”](#). Events produced from

execution of an instrument have an `EVENT_NAME` value that is taken from the instrument `NAME` value. (Events do not really have a “name,” but this provides a way to associate events with instruments.)

- `ENABLED`

Whether the instrument is enabled. This column can be modified. A disabled instrument produces no events.

- `TIMED`

Whether the instrument is timed. This column can be modified.

If an enabled instrument is not timed, the instrument code is enabled, but the timer is not. Events produced by the instrument have `NULL` for the `TIMER_START`, `TIMER_END`, and `TIMER_WAIT` timer values. This in turn causes those values to be ignored when calculating the sum, minimum, maximum, and average time values in summary tables.

Modifications to the `setup_instruments` table affect monitoring immediately.

For more information about the role of the `setup_instruments` table in event filtering, see [Section 20.2.3.2.1, “Event Pre-Filtering”](#).

#### 20.9.1.4. The `setup_objects` Table

The `setup_objects` table controls whether particular objects are monitored. This table has a maximum size of 100 rows by default. The size may be changed by modifying the `performance_schema_setup_objects_size` [2317] system variable at server startup.

The initial `setup_objects` contents look like this:

```
mysql> SELECT * FROM setup_objects;
+-----+-----+-----+-----+-----+
| OBJECT_TYPE | OBJECT_SCHEMA | OBJECT_NAME | ENABLED | TIMED |
+-----+-----+-----+-----+-----+
| TABLE     | mysql        | %           | NO      | NO    |
| TABLE     | performance_schema | %         | NO      | NO    |
| TABLE     | information_schema | %        | NO      | NO    |
| TABLE     | %           | %         | YES     | YES   |
+-----+-----+-----+-----+-----+
```

For object types listed in `setup_objects`, the Performance Schema uses the table to how to monitor them. Object matching is based on the `OBJECT_SCHEMA` and `OBJECT_NAME` columns. Objects for which there is no match are not monitored.

When the Performance Schema checks for a match in `setup_objects`, it tries to find more specific matches first. For example, with a table `db1.t1`, it looks for a match for `'db1'` and `'t1'`, then for `'db1'` and `'%'`, then for `'%'` and `'%'`. The order in which matching occurs matters because different `setup_objects` rows can have different `ENABLED` and `TIMED` values.

The effect of the default object configuration is to instrument all tables except those in the `mysql`, `INFORMATION_SCHEMA`, and `performance_schema` databases. Tables in the `INFORMATION_SCHEMA` database are not instrumented regardless of the contents of `setup_objects`; the row for `information_schema.%` simply makes this default explicit.

The `setup_objects` table has these columns:

- `OBJECT_TYPE`

The type of object to instrument. Currently, this is always `'TABLE'` (base table).

- `OBJECT_SCHEMA`

The schema that contains the object. This should be a literal name, or `'%'` to mean “any schema.”

- `OBJECT_NAME`

The name of the instrumented object. This should be a literal name, or `'%'` to mean “any object.”

- `ENABLED`

Whether events for the object are instrumented. This column can be modified.

This column was added in MySQL 5.6.3. For earlier versions in which it is not present, the Performance Schema enables monitoring only for objects matched by some row in the table; monitoring is implicitly disabled for nonmatching objects.

- `TIMED`

Whether events for the object are timed. This column can be modified.

Rows can be inserted into or deleted from `setup_objects` by users with the [INSERT \[755\]](#) or [DELETE \[754\]](#) privilege on the table. For existing rows, only the `ENABLED` and `TIMED` columns can be modified, by users with the [UPDATE \[756\]](#) privilege on the table.

Modifications to the `setup_objects` table affect object monitoring immediately.

For more information about the role of the `setup_objects` table in event filtering, see [Section 20.2.3.2.1, “Event Pre-Filtering”](#).

### 20.9.1.5. The `setup_timers` Table

The `setup_timers` table shows the currently selected event timers:

```
mysql> SELECT * FROM setup_timers;
+-----+-----+
| NAME      | TIMER_NAME |
+-----+-----+
| idle      | MICROSECOND |
| wait      | CYCLE       |
| stage     | NANOSECOND  |
| statement | NANOSECOND  |
+-----+-----+
```

The `setup_timers` table has these columns:

- `NAME`

The type of instrument the timer is used for.

- `TIMER_NAME`

The timer that applies to the instrument type. This column can be modified.

The `setup_timers.TIMER_NAME` value can be changed to select a different timer. The value can be any of the values in the `performance_timers.TIMER_NAME` column. For an explanation of how event timing occurs, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

Modifications to the `setup_timers` table affect monitoring immediately. Events already in progress may use the original timer for the begin time and the new timer for the end time, which may lead to unpredictable results. If you make timer changes, you may want to use `TRUNCATE TABLE` to reset Performance Schema statistics.

## 20.9.2. Performance Schema Instance Tables

Instance tables document what types of objects are instrumented. They provide event names and explanatory notes or status information:

- `cond_instances`: Condition synchronization object instances

- `file_instances`: File instances
- `mutex_instances`: Mutex synchronization object instances
- `rwlock_instances`: Lock synchronization object instances
- `socket_instances`: Active connection instances

These tables list instrumented synchronization objects, files, and connections. There are three types of synchronization objects: `cond`, `mutex`, and `rwlock`. Each instance table has an `EVENT_NAME` or `NAME` column to indicate the instrument associated with each row. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

The `mutex_instances.LOCKED_BY_THREAD_ID` and `rwlock_instances.WRITE_LOCKED_BY_THREAD_ID` columns are extremely important for investigating performance bottlenecks or deadlocks. For examples of how to use them for this purpose, see [Section 20.15, “Using the Performance Schema to Diagnose Problems”](#)

### 20.9.2.1. The `cond_instances` Table

The `cond_instances` table lists all the conditions seen by the Performance Schema while the server executes. A condition is a synchronization mechanism used in the code to signal that a specific event has happened, so that a thread waiting for this condition can resume work.

When a thread is waiting for something to happen, the condition name is an indication of what the thread is waiting for, but there is no immediate way to tell which other thread, or threads, will cause the condition to happen.

The `cond_instances` table has these columns:

- `NAME`  
The instrument name associated with the condition.
- `OBJECT_INSTANCE_BEGIN`  
The address in memory of the condition that was instrumented.

### 20.9.2.2. The `file_instances` Table

The `file_instances` table lists all the files seen by the Performance Schema when executing file I/O instrumentation. If a file on disk has never been opened, it will not be in `file_instances`. When a file is deleted from the disk, it is also removed from the `file_instances` table.

The `file_instances` table has these columns:

- `FILE_NAME`  
The file name.
- `EVENT_NAME`  
The instrument name associated with the file.
- `OPEN_COUNT`  
The count of open handles on the file. If a file was opened and then closed, it was opened 1 time, but `OPEN_COUNT` will be 0. To list all the files currently opened by the server, use `WHERE OPEN_COUNT > 0`.

### 20.9.2.3. The `mutex_instances` Table



The `mutex_instances` table lists all the mutexes seen by the Performance Schema while the server executes. A mutex is a synchronization mechanism used in the code to enforce that only one thread at a given time can have access to some common resource. The resource is said to be “protected” by the mutex.

When two threads executing in the server (for example, two user sessions executing a query simultaneously) do need to access the same resource (a file, a buffer, or some piece of data), these two threads will compete against each other, so that the first query to obtain a lock on the mutex will cause the other query to wait until the first is done and unlocks the mutex.

The work performed while holding a mutex is said to be in a “critical section,” and multiple queries do execute this critical section in a serialized way (one at a time), which is a potential bottleneck.

The `mutex_instances` table has these columns:

- `NAME`

The instrument name associated with the mutex.

- `OBJECT_INSTANCE_BEGIN`

The address in memory of the mutex that was instrumented.

- `LOCKED_BY_THREAD_ID`

When a thread currently has a mutex locked, `LOCKED_BY_THREAD_ID` is the `THREAD_ID` of the locking thread, otherwise it is `NULL`.

For every mutex instrumented in the code, the Performance Schema provides the following information.

- The `setup_instruments` table lists the name of the instrumentation point, with the prefix `wait/synch/mutex/`.
- When some code creates a mutex, a row is added to the `mutex_instances` table. The `OBJECT_INSTANCE_BEGIN` column is a property that uniquely identifies the mutex.
- When a thread attempts to lock a mutex, the `events_waits_current` table shows a row for that thread, indicating that it is waiting on a mutex (in the `EVENT_NAME` column), and indicating which mutex is waited on (in the `OBJECT_INSTANCE_BEGIN` column).
- When a thread succeeds in locking a mutex:
  - `events_waits_current` shows that the wait on the mutex is completed (in the `TIMER_END` and `TIMER_WAIT` columns)
  - The completed wait event is added to the `events_waits_history` and `events_waits_history_long` tables
  - `mutex_instances` shows that the mutex is now owned by the thread (in the `THREAD_ID` column).
- When a thread unlocks a mutex, `mutex_instances` shows that the mutex now has no owner (the `THREAD_ID` column is `NULL`).
- When a mutex object is destroyed, the corresponding row is removed from `mutex_instances`.

By performing queries on both of the following tables, a monitoring application or a DBA can detect bottlenecks or deadlocks between threads that involve mutexes:

- `events_waits_current`, to see what mutex a thread is waiting for
- `mutex_instances`, to see which other thread currently owns a mutex

### 20.9.2.4. The `rwlock_instances` Table

The `rwlock_instances` table lists all the `rwlock` instances (read write locks) seen by the Performance Schema while the server executes. An `rwlock` is a synchronization mechanism used in the code to enforce that threads at a given time can have access to some common resource following certain rules. The resource is said to be “protected” by the `rwlock`. The access is either shared (many threads can have a read lock at the same time) or exclusive (only one thread can have a write lock at a given time).

Depending on how many threads are requesting a lock, and the nature of the locks requested, access can be either granted in shared mode, granted in exclusive mode, or not granted at all, waiting for other threads to finish first.

The `rwlock_instances` table has these columns:

- `NAME`

The instrument name associated with the lock.

- `OBJECT_INSTANCE_BEGIN`

The address in memory of the lock that was instrumented.

- `WRITE_LOCKED_BY_THREAD_ID`

When a thread currently has an `rwlock` locked in exclusive (write) mode, `WRITE_LOCKED_BY_THREAD_ID` is the `THREAD_ID` of the locking thread, otherwise it is `NULL`.

- `READ_LOCKED_BY_COUNT`

When a thread currently has an `rwlock` locked in shared (read) mode, `READ_LOCKED_BY_COUNT` is incremented by 1. This is a counter only, so it cannot be used directly to find which thread holds a read lock, but it can be used to see whether there is a read contention on an `rwlock`, and see how many readers are currently active.

By performing queries on both of the following tables, a monitoring application or a DBA may detect some bottlenecks or deadlocks between threads that involve locks:

- `events_waits_current`, to see what `rwlock` a thread is waiting for
- `rwlock_instances`, to see which other thread currently owns an `rwlock`

There is a limitation: The `rwlock_instances` can be used only to identify the thread holding a write lock, but not the threads holding a read lock.

### 20.9.2.5. The `socket_instances` Table

The `socket_instances` table provides a real-time snapshot of the active connections to the MySQL server. The table contains one row per TCP/IP or Unix socket file connection. Information available in this table includes network activity such as socket instances, socket operations, and number of bytes transmitted and received.

```
mysql> SELECT * FROM socket_instances\G
***** 1. row *****
      EVENT_NAME: wait/io/socket/sql/server_unix_socket
OBJECT_INSTANCE_BEGIN: 4316619408
      THREAD_ID: 1
      SOCKET_ID: 16
        IP:
        PORT: 0
      STATE: ACTIVE
***** 2. row *****
      EVENT_NAME: wait/io/socket/sql/client_connection
OBJECT_INSTANCE_BEGIN: 4316644608
```

```

THREAD_ID: 21
SOCKET_ID: 39
  IP: 127.0.0.1
  PORT: 55233
  STATE: ACTIVE
***** 3. row *****
EVENT_NAME: wait/io/socket/sql/server_tcpip_socket
OBJECT_INSTANCE_BEGIN: 4316699040
  THREAD_ID: 1
  SOCKET_ID: 14
  IP: 0.0.0.0
  PORT: 50603
  STATE: ACTIVE

```

Socket instruments have names of the form `wait/io/socket/sql/socket_type` and are used like this:

1. The server has a listening socket for each network protocol that it supports. The instruments associated with listening sockets for TCP/IP or Unix socket file connections have a `socket_type` value of `server_tcpip_socket` or `server_unix_socket`, respectively.
2. When a listening socket detects a connection, the server transfers the connection to a new socket managed by a separate thread. The instrument for the new connection thread has a `socket_type` value of `client_connection`.
3. When a connection terminates, the row in `socket_instances` corresponding to it is deleted.

The `socket_instances` table has these columns:

- `EVENT_NAME`

The name of the `wait/io/socket/*` instrument that produced the event. This is a `setup_instruments.NAME` value. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

- `OBJECT_INSTANCE_BEGIN`

This column uniquely identifies the socket. The value is the address of an object in memory.

- `THREAD_ID`

The internal thread identifier assigned by the server. Each socket is managed by a single thread, so each socket can be mapped to a thread which can be mapped to a server process.

- `SOCKET_ID`

The internal file handle assigned to the socket.

- `IP`

The client IP address. The value may be either an IPv4 or IPv6 address, or blank to indicate a Unix socket file connection.

- `PORT`

The TCP/IP port number, in the range from 0 to 65535.

- `STATE`

The socket status, either `IDLE` or `ACTIVE`. Wait times for active sockets are tracked using the corresponding socket instrument. Wait times for idle sockets are tracked using the `idle` instrument.

A socket is idle if it is waiting for a request from the client. When a socket becomes idle, the event row in `socket_instances` that is tracking the socket switches from a status of `ACTIVE` to `IDLE`.

The `EVENT_NAME` value remains `wait/io/socket/*`, but timing for the instrument is suspended. Instead, an event is generated in the `events_waits_current` table with an `EVENT_NAME` value of `idle`.

When the next request is received, the `idle` event terminates, the socket instance switches from `IDLE` to `ACTIVE`, and timing of the socket instrument resumes.

The `IP:PORT` column combination value identifies the connection. This combination value is used in the `OBJECT_NAME` column of the `events_waits_XXX` tables, to identify the connection from which socket events come:

- For the Unix domain listener socket (`server_unix_socket`), the port is 0, and the IP is ''.
- For client connections via the Unix domain listener (`client_connection`), the port is 0, and the IP is ''.
- For the TCP/IP server listener socket (`server_tcpip_socket`), the port is always the master port (for example, 3306), and the IP is always `0.0.0.0`.
- For client connections via the TCP/IP listener (`client_connection`), the port is whatever the server assigns, but never 0. The IP is the IP of the originating host (`127.0.0.1` or `::1` for the local host)

The `socket_instances` table was added in MySQL 5.6.3.

## 20.9.3. Performance Schema Wait Event Tables

These tables store wait events:

- `events_waits_current`: Current wait events
- `events_waits_history`: The most recent wait events for each thread
- `events_waits_history_long`: The most recent wait events overall

### 20.9.3.1. The `events_waits_current` Table

The `events_waits_current` table contains current wait events, one row per thread showing the current status of the thread's most recent monitored wait event.

The `events_waits_current` table can be truncated with `TRUNCATE TABLE`.

Of the tables that contain wait event rows, `events_waits_current` is the most fundamental. Other tables that contain wait event rows are logically derived from the current events. For example, the `events_waits_history` and `events_waits_history_long` tables are collections of the most recent wait events, up to a fixed number of rows.

The `events_waits_current` table has these columns:

- `THREAD_ID`

The thread associated with the event. The `THREAD_ID` and `EVENT_ID` values taken together form a primary key that uniquely identifies the row. No two rows will have the same pair of values.

- `EVENT_ID`

The thread current event number when the event starts.

- `END_EVENT_ID`

This column is set to `NULL` when the event starts, and updated to the thread current event number when the event ends. This column was added in MySQL 5.6.4.

- `EVENT_NAME`

The name of the instrument that produced the event. This is a `setup_instruments.NAME` value. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

- `SOURCE`

The name of the source file containing the instrumented code that produced the event and the line number in the file at which the instrumentation occurs. This enables you to check the source to determine exactly what code is involved. For example, if a mutex or lock is being blocked, you can check the context in which this occurs.

- `TIMER_START`, `TIMER_END`, `TIMER_WAIT`

Timing information for the event. The unit for these values is picoseconds (trillionths of a second). The `TIMER_START` and `TIMER_END` values indicate when event timing started and ended. `TIMER_WAIT` is the event elapsed time (duration).

If an event has not finished, `TIMER_END` and `TIMER_WAIT` are `NULL`.

If an event is produced from an instrument that has `TIMED = NO`, timing information is not collected, and `TIMER_START`, `TIMER_END`, and `TIMER_WAIT` are all `NULL`.

For discussion of picoseconds as the unit for event times and factors that affect time values, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

- `SPINS`

For a mutex, the number of spin rounds. If the value is `NULL`, the code does not use spin rounds or spinning is not instrumented.

- `OBJECT_SCHEMA`, `OBJECT_NAME`, `OBJECT_TYPE`, `OBJECT_INSTANCE_BEGIN`

These columns identify the object “being acted on.” What that means depends on the object type.

For a synchronization object (`cond`, `mutex`, `rwlock`):

- `OBJECT_SCHEMA`, `OBJECT_NAME`, and `OBJECT_TYPE` are `NULL`.
- `OBJECT_INSTANCE_BEGIN` is the address of the synchronization object in memory.

For a file I/O object:

- `OBJECT_SCHEMA` is `NULL`.
- `OBJECT_NAME` is the file name.
- `OBJECT_TYPE` is `FILE`.
- `OBJECT_INSTANCE_BEGIN` is an address in memory.

For a socket object:

- `OBJECT_NAME` is the `IP:PORT` value for the socket.
- `OBJECT_INSTANCE_BEGIN` is an address in memory.

For a table I/O object:

- `OBJECT_SCHEMA` is the name of the schema that contains the table.
- `OBJECT_NAME` is the table name.

- `OBJECT_TYPE` is `TABLE` for a persistent base table or `TEMPORARY TABLE` for a temporary table.
- `OBJECT_INSTANCE_BEGIN` is an address in memory.

An `OBJECT_INSTANCE_BEGIN` value itself has no meaning, except that different values indicate different objects. `OBJECT_INSTANCE_BEGIN` can be used for debugging. For example, it can be used with `GROUP BY OBJECT_INSTANCE_BEGIN` to see whether the load on 1,000 mutexes (that protect, say, 1,000 pages or blocks of data) is spread evenly or just hitting a few bottlenecks. This can help you correlate with other sources of information if you see the same object address in a log file or another debugging or performance tool.

- `INDEX_NAME`

The name of the index used. `PRIMARY` indicates the table primary index. `NULL` means that no index was used.

- `NESTING_EVENT_ID`

The `EVENT_ID` value of the event within which this event is nested. Before MySQL 5.6.3, this column is always `NULL`.

- `NESTING_EVENT_TYPE`

The nesting event type. The value is `statement`, `stage`, or `wait`. This column was added in MySQL 5.6.3.

- `OPERATION`

The type of operation performed, such as `lock`, `read`, or `write`.

- `NUMBER_OF_BYTES`

The number of bytes read or written by the operation. For table I/O waits, `NUMBER_OF_BYTES` is `NULL`.

- `FLAGS`

Reserved for future use.

### 20.9.3.2. The `events_waits_history` Table

The `events_waits_history` table contains the most recent 10 wait events per thread. The table size may be changed by modifying the `performance_schema_events_waits_history_size` [2310] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_waits_history` table has the same structure as `events_waits_current`. See [Section 20.9.3.1, “The `events\_waits\_current` Table”](#).

The `events_waits_history` table can be truncated with `TRUNCATE TABLE`.

### 20.9.3.3. The `events_waits_history_long` Table

The `events_waits_history_long` table contains the most recent 10,000 wait events. The table size may be changed by modifying the `performance_schema_events_waits_history_long_size` [2310] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_waits_history_long` table has the same structure as `events_waits_current`. See [Section 20.9.3.1, “The `events\_waits\_current` Table”](#).

The `events_waits_history_long` table can be truncated with `TRUNCATE TABLE`.

## 20.9.4. Performance Schema Stage Event Tables

As of MySQL 5.6.3, the Performance Schema instruments stages, which are steps during the statement-execution process, such as parsing a statement, opening a table, or performing a `filesort` operation. Stages correspond to the thread states displayed by `SHOW PROCESSLIST` or that are visible in the `INFORMATION_SCHEMA.PROCESSLIST` table. Stages begin and end when state values change.

Within the event hierarchy, wait events nest within stage events, which nest within statement events.

These tables store stage events:

- `events_stages_current`: Current stage events
- `events_stages_history`: The most recent stage events for each thread
- `events_stages_history_long`: The most recent stage events overall

### 20.9.4.1. The `events_stages_current` Table

The `events_stages_current` table contains current stage events, one row per thread showing the current status of the thread's most recent monitored stage event.

The `events_stages_current` table can be truncated with `TRUNCATE TABLE`.

Of the tables that contain stage event rows, `events_stages_current` is the most fundamental. Other tables that contain stage event rows are logically derived from the current events. For example, the `events_stages_history` and `events_stages_history_long` tables are collections of the most recent stage events, up to a fixed number of rows.

The `events_stages_current` table has these columns:

- `THREAD_ID`

The thread associated with the event. The `THREAD_ID` and `EVENT_ID` values taken together form a primary key that uniquely identifies the row. No two rows will have the same pair of values.

- `EVENT_ID`

The thread current event number when the event starts.

- `END_EVENT_ID`

This column is set to `NULL` when the event starts, and updated to the thread current event number when the event ends. This column was added in MySQL 5.6.4.

- `EVENT_NAME`

The name of the instrument that produced the event. This is a `setup_instruments.NAME` value. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, "Performance Schema Instrument Naming Conventions"](#).

- `SOURCE`

The name of the source file containing the instrumented code that produced the event and the line number in the file at which the instrumentation occurs. This enables you to check the source to determine exactly what code is involved.

- `TIMER_START`, `TIMER_END`, `TIMER_WAIT`

Timing information for the event. The unit for these values is picoseconds (trillionths of a second). The `TIMER_START` and `TIMER_END` values indicate when event timing started and ended. `TIMER_WAIT` is the event elapsed time (duration).



If an event has not finished, `TIMER_END` and `TIMER_WAIT` are `NULL`.

If an event is produced from an instrument that has `TIMED = NO`, timing information is not collected, and `TIMER_START`, `TIMER_END`, and `TIMER_WAIT` are all `NULL`.

For discussion of picoseconds as the unit for event times and factors that affect time values, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

- `NESTING_EVENT_ID`

The `EVENT_ID` value of the event within which this event is nested. The nesting event for a stage event is usually a statement event.

- `NESTING_EVENT_TYPE`

The nesting event type. The value is `statement`, `stage`, or `wait`.

The `events_stages_current` table was added in MySQL 5.6.3.

#### 20.9.4.2. The `events_stages_history` Table

The `events_stages_history` table contains the most recent 10 stage events per thread. The table size may be changed by modifying the `performance_schema_events_stages_history_size` [2309] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_stages_history` table has the same structure as `events_stages_current`. See [Section 20.9.4.1, “The `events\_stages\_current` Table”](#).

The `events_stages_history` table can be truncated with `TRUNCATE TABLE`.

The `events_stages_history` table was added in MySQL 5.6.3.

#### 20.9.4.3. The `events_stages_history_long` Table

The `events_stages_history_long` table contains the most recent 10,000 stage events. The table size may be changed by modifying the `performance_schema_events_stages_history_long_size` [2308] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_stages_history_long` table has the same structure as `events_stages_current`. See [Section 20.9.4.1, “The `events\_stages\_current` Table”](#).

The `events_stages_history_long` table can be truncated with `TRUNCATE TABLE`.

The `events_stages_history_long` table was added in MySQL 5.6.3.

### 20.9.5. Performance Schema Statement Event Tables

As of MySQL 5.6.3, the Performance Schema instruments statement execution. Statement events occur at a high level of the event hierarchy. Wait events nest within stage events, which nest within statement events.

These tables store statement events:

- `events_statements_current`: Current statement events
- `events_statements_history`: The most recent statement events for each thread

- `events_statements_history_long`: The most recent statement events overall

Statement monitoring begins from the moment when the server sees that activity is requested on a thread, to the moment when all activity has ceased. Typically, this means from the time the server gets the first packet from the client to the time the server has finished sending the response. Monitoring occurs only for top-level statements. Statements within stored programs and subqueries are not seen separately.

A request from a client can be either a command or a SQL statement:

- Server commands correspond to the `COM_xxx` codes defined in the `mysql_com.h` header file and processed in `sql/sql_parse.cc`. Examples are `COM_PING`, `COM_QUERY`, and `COM_QUIT`. Instruments for commands have names that begin with `statement/com`, such as `statement/com/Ping`, `statement/com/Query`, and `statement/com/Quit`.
- SQL statements are expressed as text, such as `DELETE FROM t1` or `SELECT * FROM t2`. Instruments for SQL statements have names that begin with `statement/sql`, such as `statement/sql/delete` and `statement/sql/select`.

There are also special error-handling instruments:

- `statement/com/Error` accounts for messages received by the server that are out of band. It can be used to detect commands sent by clients that the server does not understand. This may be helpful for purposes such as identifying clients that are misconfigured or using a version of MySQL more recent than that of the server, or clients that are attempting to attack the server.
- `statement/sql/error` accounts for SQL statements that fail to parse. It can be used to detect malformed queries sent by clients. A query that fails to parse differs from a query that parses but fails due to an error during execution. For example, `SELECT * FROM` is malformed, and the `statement/sql/error` instrument is used. By contrast, `SELECT *` parses but fails with a `No tables used` error. In this case, `statement/sql/select` is used and the statement event contains information to indicate the nature of the error.

Details for a request are not initially known. For example, because a SQL statement is sent as text inside a `COM_QUERY` packet, the proper `statement/sql/*` instrument is not known at the moment the request is received. The Performance Schema at first instruments such a request using an event with an `EVENT_NAME` of `statement/com/Query`. Then it changes the `EVENT_NAME` value to either a valid `statement/sql/*` name if the statement parses, or to `statement/sql/error` if it does not.

### 20.9.5.1. The `events_statements_current` Table

The `events_statements_current` table contains current statement events, one row per thread showing the current status of the thread's most recent monitored statement event.

The `events_statements_current` table can be truncated with `TRUNCATE TABLE`.

Of the tables that contain statement event rows, `events_statements_current` is the most fundamental. Other tables that contain statement event rows are logically derived from the current events. For example, the `events_statements_history` and `events_statements_history_long` tables are collections of the most recent statement events, up to a fixed number of rows.

The `events_statements_current` table has these columns:

- `THREAD_ID`

The thread associated with the event. The `THREAD_ID` and `EVENT_ID` values taken together form a primary key that uniquely identifies the row. No two rows will have the same pair of values.

- `EVENT_ID`

The thread current event number when the event starts.

- `END_EVENT_ID`

This column is set to `NULL` when the event starts, and updated to the thread current event number when the event ends. This column was added in MySQL 5.6.4.

- `EVENT_NAME`

The name of the instrument from which the event was collected. This is a `setup_instruments.NAME` value. Instrument names have multiple parts and form a hierarchy, as discussed in [Section 20.4, “Performance Schema Instrument Naming Conventions”](#).

For SQL statements, the `EVENT_NAME` value initially is `statement/com/Query` until the statement is parsed, then changes to a more appropriate value, as described in [Section 20.9.5, “Performance Schema Statement Event Tables”](#).

- `SOURCE`

The name of the source file containing the instrumented code that produced the event and the line number in the file at which the instrumentation occurs. This enables you to check the source to determine exactly what code is involved.

- `TIMER_START`, `TIMER_END`, `TIMER_WAIT`

Timing information for the event. The unit for these values is picoseconds (trillionths of a second). The `TIMER_START` and `TIMER_END` values indicate when event timing started and ended. `TIMER_WAIT` is the event elapsed time (duration).

If an event has not finished, `TIMER_END` and `TIMER_WAIT` are `NULL`.

If an event is produced from an instrument that has `TIMED = NO`, timing information is not collected, and `TIMER_START`, `TIMER_END`, and `TIMER_WAIT` are all `NULL`.

For discussion of picoseconds as the unit for event times and factors that affect time values, see [Section 20.2.3.1, “Performance Schema Event Timing”](#).

- `LOCK_TIME`

The time spent waiting for table locks. This value is computed in microseconds but normalized to picoseconds for easier comparison with other Performance Schema timers.

- `SQL_TEXT`

The text of the SQL statement. For a command not associated with a SQL statement, the value is `NULL`.

- `DIGEST`

The statement digest MD5 value as a string of 32 hexadecimal characters, or `NULL` if the `statement_digest` consumer is `no`. For more information about statement digesting, see [Section 20.7, “Performance Schema Statement Digests”](#). This column was added in MySQL 5.6.5.

- `DIGEST_TEXT`

The normalized statement digest text, or `NULL` if the `statement_digest` consumer is `no`. For more information about statement digesting, see [Section 20.7, “Performance Schema Statement Digests”](#). This column was added in MySQL 5.6.5.

- `CURRENT_SCHEMA`

The default database for the statement, `NULL` if there is none.

- `OBJECT_SCHEMA`, `OBJECT_NAME`, `OBJECT_TYPE`

Reserved. Currently `NULL`.

- `OBJECT_INSTANCE_BEGIN`

This column identifies the statement. The value is the address of an object in memory.

- `MYSQL_ERRNO`

The statement error number, from the statement diagnostics area.

- `RETURNED_SQLSTATE`

The statement `SQLSTATE` value, from the statement diagnostics area.

- `MESSAGE_TEXT`

The statement error message, from the statement diagnostics area.

- `ERRORS`

Whether an error occurred for the statement. The value is 0 if the `SQLSTATE` value begins with 00 (completion) or 01 (warning). The value is 1 if the `SQLSTATE` value is anything else.

- `WARNINGS`

The number of warnings, from the statement diagnostics area.

- `ROWS_AFFECTED`

The number of rows affected by the statement. For a description of the meaning of “affected,” see [Section 21.8.6.1](#), “`mysql_affected_rows()`”.

- `ROWS_SENT`

The number of rows returned by the statement.

- `ROWS_EXAMINED`

The number of rows read from storage engines during statement execution.

- `CREATED_TMP_DISK_TABLES`

Like the `Created_tmp_disk_tables` [\[587\]](#) status variable, but specific to the statement.

- `CREATED_TMP_TABLES`

Like the `Created_tmp_tables` [\[588\]](#) status variable, but specific to the statement.

- `SELECT_FULL_JOIN`

Like the `Select_full_join` [\[598\]](#) status variable, but specific to the statement.

- `SELECT_FULL_RANGE_JOIN`

Like the `Select_full_range_join` [\[598\]](#) status variable, but specific to the statement.

- `SELECT_RANGE`

Like the `Select_range` [\[598\]](#) status variable, but specific to the statement.

- `SELECT_RANGE_CHECK`

Like the `Select_range_check` [\[598\]](#) status variable, but specific to the statement.

- `SELECT_SCAN`  
Like the `Select_scan` [598] status variable, but specific to the statement.
- `SORT_MERGE_PASSES`  
Like the `Sort_merge_passes` [599] status variable, but specific to the statement.
- `SORT_RANGE`  
Like the `Sort_range` [599] status variable, but specific to the statement.
- `SORT_ROWS`  
Like the `Sort_rows` [599] status variable, but specific to the statement.
- `SORT_SCAN`  
Like the `Sort_scan` [599] status variable, but specific to the statement.
- `NO_INDEX_USED`  
1 if the statement performed a table scan without using an index, 0 otherwise.
- `NO_GOOD_INDEX_USED`  
1 if the server found no good index to use for the statement, 0 otherwise. For additional information, see the description of the `Extra` column from `EXPLAIN` output for the `Range checked for each record` value in [Section 8.8.2, “EXPLAIN Output Format”](#).
- `NESTING_EVENT_ID`, `NESTING_EVENT_TYPE`  
Reserved. Currently `NULL`.

The `events_statements_current` table was added in MySQL 5.6.3.

### 20.9.5.2. The `events_statements_history` Table

The `events_statements_history` table contains the most recent 10 events per thread. The table size may be changed by modifying the `performance_schema_events_statements_history_size` [2309] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_statements_history` table has the same structure as `events_statements_current`. See [Section 20.9.5.1, “The `events\_statements\_current` Table”](#).

The `events_statements_history` table can be truncated with `TRUNCATE TABLE`.

The `events_statements_history` table was added in MySQL 5.6.3.

### 20.9.5.3. The `events_statements_history_long` Table

The `events_statements_history_long` table contains the most recent 10,000 events. The table size may be changed by modifying the `performance_schema_events_statements_history_long_size` [2309] system variable at server startup. As new events are added to the table, older events are discarded if the table is full. Events are not added to the table until they have ended.

The `events_statements_history_long` table has the same structure as `events_statements_current`. See [Section 20.9.5.1, “The `events\_statements\_current` Table”](#).

The `events_statements_history_long` table can be truncated with `TRUNCATE TABLE`.

The `events_statements_history_long` table was added in MySQL 5.6.3.

## 20.9.6. Performance Schema Connection Tables

As of MySQL 5.6.3, the Performance Schema provides statistics about connections to the server. When a client connects, it does so under a particular user name and from a particular host. The Performance Schema tracks connections per account (user name plus host name) and separately per user name and per host name, using these tables:

- `accounts`: Connection statistics per client account
- `hosts`: Connection statistics per client host name
- `users`: Connection statistics per client user name

The meaning of “account” in the connection tables is similar to its meaning in the MySQL grant tables in the `mysql` database, in the sense that the term refers to a combination of user and host values. Where they differ is that in the grant tables, the host part of an account can be a pattern, whereas in the connection tables the host value is always a specific nonpattern host name.

The connection tables all have `CURRENT_CONNECTIONS` and `TOTAL_CONNECTIONS` columns to track the current and total number of connections per “tracking value” on which statistics are based. The tables differ in what they use for the tracking value. The `accounts` table has `USER` and `HOST` columns to track connections per user name plus host name combination. The `users` and `hosts` tables have a `USER` and `HOST` column, respectively, to track connections per user name and per host name.

Suppose that clients named `user1` and `user2` each connect one time from `hosta` and `hostb`. The Performance Schema tracks the connections as follows:

- The `accounts` table will have four rows, for the `user1/hosta`, `user1/hostb`, `user2/hosta`, and `user2/hostb` account values, each row counting one connection per account.
- The `users` table will have two rows, for `user1` and `user2`, each row counting two connections per user name.
- The `hosts` table will have two rows, for `hosta` and `hostb`, each row counting two connections per host name.

When a client connects, the Performance Schema determines which row in each connection table applies to the connection, using the tracking value appropriate to each table. If there is no such row, one is added. Then the Performance Schema increments by one the `CURRENT_CONNECTIONS` and `TOTAL_CONNECTIONS` columns in that row.

When a client disconnects, the Performance Schema decrements by one the `CURRENT_CONNECTIONS` column in the row and leaves the `TOTAL_CONNECTIONS` column unchanged.

Each connection table can be truncated with `TRUNCATE TABLE`, which has this effect:

- Rows with `CURRENT_CONNECTIONS = 0` are deleted.
- For rows with `CURRENT_CONNECTIONS > 0`, `TOTAL_CONNECTIONS` is reset to `CURRENT_CONNECTIONS`.
- Connection summary tables that depend on the connection table are truncated implicitly (summary values are set to 0). For more information about implicit truncation, see [Section 20.9.8.7, “Connection Summary Tables”](#).

### 20.9.6.1. The `accounts` Table

The `accounts` table contains a row for each account that has connected to the MySQL server. For each account, the table counts the current and total number of connections. The table size may be changed by modifying the `performance_schema_accounts_size` [2308] system variable at server startup. To disable account statistics, set this variable to 0.

The `accounts` table has the following columns. For a description of how the Performance Schema maintains rows in this table, including the effect of `TRUNCATE TABLE`, see [Section 20.9.6, “Performance Schema Connection Tables”](#).

- `USER`

The client user name for the connection.

- `HOST`

The host name from which the client connected.

- `CURRENT_CONNECTIONS`

The current number of connections for the account.

- `TOTAL_CONNECTIONS`

The total number of connections for the account.

The `accounts` table was added in MySQL 5.6.3.

### 20.9.6.2. The `hosts` Table

The `hosts` table contains a row for each host from which clients have connected to the MySQL server. For each host name, the table counts the current and total number of connections. The table size may be changed by modifying the `performance_schema_hosts_size` [2310] system variable at server startup. To disable host statistics, set this variable to 0.

The `hosts` table has the following columns. For a description of how the Performance Schema maintains rows in this table, including the effect of `TRUNCATE TABLE`, see [Section 20.9.6, “Performance Schema Connection Tables”](#).

- `HOST`

The host name from which the client connected.

- `CURRENT_CONNECTIONS`

The current number of connections for the host.

- `TOTAL_CONNECTIONS`

The total number of connections for the host.

The `hosts` table was added in MySQL 5.6.3.

### 20.9.6.3. The `users` Table

The `users` table contains a row for each user who has connected to the MySQL server. For each user name, the table counts the current and total number of connections. The table size may be changed by modifying the `performance_schema_users_size` [2317] system variable at server startup. To disable user statistics, set this variable to 0.

The `users` table has the following columns. For a description of how the Performance Schema maintains rows in this table, including the effect of `TRUNCATE TABLE`, see [Section 20.9.6, “Performance Schema Connection Tables”](#).



- `USER`  
The client user name for the connection.
- `CURRENT_CONNECTIONS`  
The current number of connections for the user.
- `TOTAL_CONNECTIONS`  
The total number of connections for the user.

The `users` table was added in MySQL 5.6.3.

## 20.9.7. Performance Schema Connection Attribute Tables

As of MySQL 5.6.6, application programs can provide key/value connection attributes to be passed to the server at connect time, using the `mysql_options()` and `mysql_options4()` C API functions. The Performance Schema provides tables that expose this information through SQL statements:

- `session_account_connect_attrs`: Connection attributes per for the current session
- `session_connect_attrs`: Connection attributes for all sessions

Both tables have the same columns. The difference between them is that `session_connect_attrs` displays connection attributes for all sessions, whereas `session_account_connect_attrs` displays connections only for your own account.

- `PROCESSLIST_ID`  
The connection identifier for the session.
- `ATTR_NAME`  
The attribute name.
- `ATTR_VALUE`  
The attribute value.
- `ORDINAL_POSITION`  
The order in which the attribute was added to the set of connection attributes.

## 20.9.8. Performance Schema Summary Tables

Summary tables provide aggregated information for terminated events over time. The tables in this group summarize event data in different ways.

### Event Wait Summaries:

- `events_waits_summary_global_by_event_name`: Wait events summarized per event name
- `events_waits_summary_by_instance`: Wait events summarized per instance
- `events_waits_summary_by_thread_by_event_name`: Wait events summarized per thread and event name

### Stage Summaries:

- `events_stages_summary_by_thread_by_event_name`: Stage waits summarized per thread and event name
- `events_stages_summary_global_by_event_name`: Stage waits summarized per event name

**Statement Summaries:**

- [events\\_statements\\_summary\\_by\\_thread\\_by\\_event\\_name](#): Statement events summarized per thread and event name
- [events\\_statements\\_summary\\_global\\_by\\_event\\_name](#): Statement events summarized per event name

**Object Wait Summaries:**

- [objects\\_summary\\_global\\_by\\_type](#): Object summaries

**File I/O Summaries:**

- [file\\_summary\\_by\\_event\\_name](#): File events summarized per event name
- [file\\_summary\\_by\\_instance](#): File events summarized per file instance

**Table I/O and Lock Wait Summaries:**

- [table\\_io\\_waits\\_summary\\_by\\_index\\_usage](#): Table I/O waits per index
- [table\\_io\\_waits\\_summary\\_by\\_table](#): Table I/O waits per table
- [table\\_lock\\_waits\\_summary\\_by\\_table](#): Table lock waits per table

**Connection Summaries:**

- [events\\_waits\\_summary\\_by\\_account\\_by\\_event\\_name](#): Wait events summarized per account and event name
- [events\\_waits\\_summary\\_by\\_user\\_by\\_event\\_name](#): Wait events summarized per user name and event name
- [events\\_waits\\_summary\\_by\\_host\\_by\\_event\\_name](#): Wait events summarized per host name and event name
- [events\\_stages\\_summary\\_by\\_account\\_by\\_event\\_name](#): Stage events summarized per account and event name
- [events\\_stages\\_summary\\_by\\_user\\_by\\_event\\_name](#): Stage events summarized per user name and event name
- [events\\_stages\\_summary\\_by\\_host\\_by\\_event\\_name](#): Stage events summarized per host name and event name
- [events\\_statements\\_summary\\_by\\_account\\_by\\_event\\_name](#): Statement events summarized per account and event name
- [events\\_statements\\_summary\\_by\\_user\\_by\\_event\\_name](#): Statement events summarized per user name and event name
- [events\\_statements\\_summary\\_by\\_host\\_by\\_event\\_name](#): Statement events summarized per host name and event name

**Socket Summaries:**

- [socket\\_summary\\_by\\_instance](#): Socket waits and I/O summarized per instance
- [socket\\_summary\\_by\\_event\\_name](#): Socket waits and I/O summarized per event name

Each summary table has grouping columns that determine how to group the data to be aggregated, and summary columns that contain the aggregated values. Tables that summarize events in similar

ways often have similar sets of summary columns and differ only in the grouping columns used to determine how events are aggregated.

Summary tables can be truncated with `TRUNCATE TABLE`. Except for `events_statements_summary_by_digest`, the effect is to reset the summary columns to 0 or `NULL`, not to remove rows. This enables you to clear collected values and restart aggregation. That might be useful, for example, after you have made a runtime configuration change.

### 20.9.8.1. Event Wait Summary Tables

The event waits summary tables aggregate general information about event waits:

- `events_waits_summary_global_by_event_name`: Wait events summarized per event name
- `events_waits_summary_by_instance`: Wait events summarized per instance
- `events_waits_summary_by_thread_by_event_name`: Wait events summarized per thread and event name

For example:

```
mysql> SELECT * FROM events_waits_summary_global_by_event_name\G
...
***** 6. row *****
EVENT_NAME: wait/synch/mutex/sql/BINARY_LOG::LOCK_index
COUNT_STAR: 8
SUM_TIMER_WAIT: 2119302
MIN_TIMER_WAIT: 196092
AVG_TIMER_WAIT: 264912
MAX_TIMER_WAIT: 569421
...
***** 9. row *****
EVENT_NAME: wait/synch/mutex/sql/hash_filo::lock
COUNT_STAR: 69
SUM_TIMER_WAIT: 16848828
MIN_TIMER_WAIT: 0
AVG_TIMER_WAIT: 244185
MAX_TIMER_WAIT: 735345
...
```

`TRUNCATE TABLE` is permitted for wait summary tables. It resets the counters to zero rather than removing rows.

The event wait summary tables have these grouping columns to indicate how events are aggregated:

- `events_waits_summary_global_by_event_name` has an `EVENT_NAME` column. Each row summarizes events for a given instrument. An instrument might be used to create multiple instances of the instrumented object. For example, if there is an instrument for a mutex that is created for each connection, there are as many instances as there are connections. The summary row for the instrument summarizes over all these instances.
- `events_waits_summary_by_instance` has `EVENT_NAME` and `OBJECT_INSTANCE_BEGIN` columns. Each row summarizes events for a given instrument instance. If an instrument is used to create multiple instances, each instance has a unique `OBJECT_INSTANCE_BEGIN` value, so these instances are summarized separately in this table.
- `events_waits_summary_by_thread_by_event_name` has `THREAD_ID` and `EVENT_NAME` columns. Each row summarizes events for a given thread and instrument.

The event waits summary tables have these summary columns containing aggregated values:

- `COUNT_STAR`

The number of summarized events. This value includes all events, whether timed or not.

- `SUM_TIMER_WAIT`

The total wait time of the summarized timed events. This value is calculated only for timed events because nontimed events have a wait time of `NULL`. The same is true for the other `xxx_TIMER_WAIT` values.

- `MIN_TIMER_WAIT`

The minimum wait time of the summarized timed events.

- `AVG_TIMER_WAIT`

The average wait time of the summarized timed events.

- `MAX_TIMER_WAIT`

The maximum wait time of the summarized timed events.

## 20.9.8.2. Stage Summary Tables

The stage summary tables provide aggregated information about stage events:

- `events_stages_summary_by_thread_by_event_name`: Stage waits summarized per thread and event name
- `events_stages_summary_global_by_event_name`: Stage waits summarized per event name

For example:

```
mysql> SELECT * FROM events_stages_summary_global_by_event_name\G
...
***** 5. row *****
EVENT_NAME: stage/sql/checking permissions
COUNT_STAR: 57
SUM_TIMER_WAIT: 26501888880
MIN_TIMER_WAIT: 7317456
AVG_TIMER_WAIT: 464945295
MAX_TIMER_WAIT: 12858936792
...
***** 9. row *****
EVENT_NAME: stage/sql/closing tables
COUNT_STAR: 37
SUM_TIMER_WAIT: 662606568
MIN_TIMER_WAIT: 1593864
AVG_TIMER_WAIT: 17907891
MAX_TIMER_WAIT: 437977248
...
```

`TRUNCATE TABLE` is permitted for stage summary tables. It resets the counters to zero rather than removing rows.

The stage summary tables have these grouping columns to indicate how events are aggregated:

- `events_stages_summary_global_by_event_name` has an `EVENT_NAME` column. Each row summarizes events for a given instrument.
- `events_stages_summary_by_thread_by_event_name` has `THREAD_ID` and `EVENT_NAME` columns. Each row summarizes events for a given thread instrument instance.

The stage summary tables have these summary columns containing aggregated values: `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, and `MAX_TIMER_WAIT`. These columns are analogous to the columns of the same names in the `events_waits_summary_global_by_event_name` and `events_waits_summary_by_thread_by_event_name` tables, except that the stage summary tables aggregate waits from `events_stages_current` rather than `events_waits_current`.

These tables were added in MySQL 5.6.3.

### 20.9.8.3. Statement Summary Tables

The statement summary tables aggregate information about statement events:

- `events_statements_summary_by_digest`: Statement events summarized per schema and digest value
- `events_statements_summary_by_thread_by_event_name`: Statement events summarized per thread and event name
- `events_statements_summary_global_by_event_name`: Statement events summarized per event name

For example:

```
mysql> SELECT * FROM events_statements_summary_global_by_event_name\G
***** 1. row *****
      EVENT_NAME: statement/sql/select
      COUNT_STAR: 25
      SUM_TIMER_WAIT: 1535983999000
      MIN_TIMER_WAIT: 209823000
      AVG_TIMER_WAIT: 61439359000
      MAX_TIMER_WAIT: 1363397650000
      SUM_LOCK_TIME: 20186000000
      SUM_ERRORS: 0
      SUM_WARNINGS: 0
      SUM_ROWS_AFFECTED: 0
      SUM_ROWS_SENT: 388
      SUM_ROWS_EXAMINED: 370
SUM_CREATED_TMP_DISK_TABLES: 0
      SUM_CREATED_TMP_TABLES: 0
      SUM_SELECT_FULL_JOIN: 0
      SUM_SELECT_FULL_RANGE_JOIN: 0
      SUM_SELECT_RANGE: 0
      SUM_SELECT_RANGE_CHECK: 0
      SUM_SELECT_SCAN: 6
      SUM_SORT_MERGE_PASSES: 0
      SUM_SORT_RANGE: 0
      SUM_SORT_ROWS: 0
      SUM_SORT_SCAN: 0
      SUM_NO_INDEX_USED: 6
      SUM_NO_GOOD_INDEX_USED: 0
      ...
```

`TRUNCATE TABLE` is permitted for statement summary tables. For `events_statements_summary_by_digest`, it empties the table. For the other statement summary tables, it resets the counters to zero rather than removing rows.

The statement summary tables have these grouping columns to indicate how events are aggregated:

- `events_statements_summary_by_digest` has `SCHEMA_NAME` and `DIGEST` column. Each row summarizes events for given schema/digest values. (The `DIGEST_TEXT` column contains the corresponding normalized statement digest text, but is neither a grouping nor summary column.)

Before MySQL 5.6.9, there is no `SCHEMA_NAME` column and grouping is based on `DIGEST` values only.

- `events_statements_summary_global_by_event_name` has an `EVENT_NAME` column. Each row summarizes events for a given instrument.
- `events_statements_summary_by_thread_by_event_name` has `THREAD_ID` and `EVENT_NAME` columns. Each row summarizes events for a given thread instrument instance.

The statement summary tables have these summary columns containing aggregated values:

- `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, `MAX_TIMER_WAIT`

These columns are analogous to the columns of the same names in the `events_waits_summary_global_by_event_name` and `events_waits_summary_by_thread_by_event_name` tables, except that the statement summary tables aggregate waits from `events_statements_current` rather than `events_waits_current`.

- `SUM_xxx`

The aggregate of the corresponding `xxx` column in the `events_statements_current` table. For example, the `SUM_LOCK_TIME` and `SUM_ERRORS` columns in statement summary tables are the aggregates of the `LOCK_TIME` and `ERRORS` columns in `events_statements_current` table.

The `events_statements_summary_by_digest` table has these additional summary columns:

- `FIRST_SEEN_TIMESTAMP`, `LAST_SEEN_TIMESTAMP`

The times at which a statement with the given digest value were first seen and most recently seen.

These tables were added in MySQL 5.6.3, except that `events_statements_summary_by_digest` was added in 5.6.5.

## Statement Digest Aggregation Rules

If the `statement_digest` consumer is enabled, aggregation into `events_statements_summary_by_digest` occurs as follows when a statement completes. Aggregation is based on the `DIGEST` value computed for the statement.

- If a `events_statements_summary_by_digest` row already exists with the digest value for the statement that just completed, statistics for the statement are aggregated to that row. The `LAST_SEEN` column is updated to the current time.
- If no row has the digest value for the statement that just completed, and the table is not full, a new row is created for the statement. The `FIRST_SEEN` and `LAST_SEEN` columns are initialized with the current time.
- If no row has the statement digest value for the statement that just completed, and the table is full, the statistics for the statement that just completed are added to a special “catch-all” row with `DIGEST = NULL`, which is created if necessary. If the row is created, the `FIRST_SEEN` and `LAST_SEEN` columns are initialized with the current time. Otherwise, the `LAST_SEEN` column is updated with the current time.

The row with `DIGEST = NULL` is maintained because Performance Schema tables have a maximum size due to memory constraints. The `DIGEST = NULL` row permits digests that do not match other rows to be counted even if the summary table is full, using a common “other” bucket. This row helps you estimate whether the digest summary is representative:

- A `DIGEST = NULL` row that has a `COUNT_STAR` value that represents 5% of all digests shows that the digest summary table is very representative; the other rows cover 95% of the statements seen.
- A `DIGEST = NULL` row that has a `COUNT_STAR` value that represents 50% of all digests shows that the digest summary table is not very representative; the other rows cover only half the statements seen. Most likely the DBA should increase the maximum table size so that more of the rows counted in the `DIGEST = NULL` row would be counted using more specific rows instead. To do this, set the `performance_schema_digests_size` [2308] system variable to a larger value at server startup. The default size is 200.

### 20.9.8.4. Object Wait Summary Table

The `objects_summary_global_by_type` aggregates object wait events. For example:

```
mysql> SELECT * FROM objects_summary_global_by_type\G
...
```

```

***** 3. row *****
OBJECT_TYPE: TABLE
OBJECT_SCHEMA: test
OBJECT_NAME: t
COUNT_STAR: 3
SUM_TIMER_WAIT: 263126976
MIN_TIMER_WAIT: 1522272
AVG_TIMER_WAIT: 87708678
MAX_TIMER_WAIT: 258428280
...
***** 10. row *****
OBJECT_TYPE: TABLE
OBJECT_SCHEMA: mysql
OBJECT_NAME: user
COUNT_STAR: 14
SUM_TIMER_WAIT: 365567592
MIN_TIMER_WAIT: 1141704
AVG_TIMER_WAIT: 26111769
MAX_TIMER_WAIT: 334783032
...

```

`TRUNCATE TABLE` is permitted for object summary tables. It resets the counters to zero rather than removing rows.

The `objects_summary_global_by_type` table has these grouping columns to indicate how events are aggregated: `OBJECT_TYPE`, `OBJECT_SCHEMA`, and `OBJECT_NAME`. Each row summarizes events for the given object.

`objects_summary_global_by_type` has the same summary columns as the `events_waits_summary_by_xxx` tables. See [Section 20.9.8.1, “Event Wait Summary Tables”](#).

### 20.9.8.5. File I/O Summary Tables

The file I/O summary tables aggregate information about I/O operations:

- `file_summary_by_event_name`: File events summarized per event name
- `file_summary_by_instance`: File events summarized per file instance

For example:

```

mysql> SELECT * FROM file_summary_by_event_name\G
...
***** 2. row *****
EVENT_NAME: wait/io/file/sql/binlog
COUNT_STAR: 31
SUM_TIMER_WAIT: 8243784888
MIN_TIMER_WAIT: 0
AVG_TIMER_WAIT: 265928484
MAX_TIMER_WAIT: 6490658832
...
mysql> SELECT * FROM file_summary_by_instance\G
...
***** 2. row *****
FILE_NAME: /var/mysql/share/english/errmsg.sys
EVENT_NAME: wait/io/file/sql/ERRMSG
EVENT_NAME: wait/io/file/sql/ERRMSG
OBJECT_INSTANCE_BEGIN: 4686193384
COUNT_STAR: 5
SUM_TIMER_WAIT: 13990154448
MIN_TIMER_WAIT: 26349624
AVG_TIMER_WAIT: 2798030607
MAX_TIMER_WAIT: 8150662536
...

```

`TRUNCATE TABLE` is permitted for file I/O summary tables. It resets the counters to zero rather than removing rows.

The file I/O summary tables have these grouping columns to indicate how events are aggregated:



- `file_summary_by_event_name` has an `EVENT_NAME` column. Each row summarizes events for a given instrument.
- `file_summary_by_instance` has `FILE_NAME`, `EVENT_NAME`, and (as of MySQL 5.6.4) `OBJECT_INSTANCE_BEGIN` columns. Each row summarizes events for a given file instrument instance.

The file I/O summary tables have the following summary columns containing aggregated values. (Before MySQL 5.6.4, the tables contain only the Only `COUNT_READ` `COUNT_WRITE` `SUM_NUMBER_OF_BYTES_READ`, and `SUM_NUMBER_OF_BYTES_WRITE` aggregation columns.) Some columns are more general and have values that are the same as the sum of the values of more fine-grained columns. In this way, aggregations at higher levels are available directly without the need for user-defined views that sum lower-level columns.

- `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, `MAX_TIMER_WAIT`

These columns aggregate all I/O operations.

- `COUNT_READ`, `SUM_TIMER_READ`, `MIN_TIMER_READ`, `AVG_TIMER_READ`, `MAX_TIMER_READ`, `SUM_NUMBER_OF_BYTES_READ`

These columns aggregate all read operations, including `FGETS`, `FGETC`, `FREAD`, and `READ`.

- `COUNT_WRITE`, `SUM_TIMER_WRITE`, `MIN_TIMER_WRITE`, `AVG_TIMER_WRITE`, `MAX_TIMER_WRITE`, `SUM_NUMBER_OF_BYTES_WRITE`

These columns aggregate all write operations, including `FPUTS`, `Fputc`, `FPRINTF`, `VFPRTF`, `FWRITE`, and `PWRITE`.

- `COUNT_MISC`, `SUM_TIMER_MISC`, `MIN_TIMER_MISC`, `AVG_TIMER_MISC`, `MAX_TIMER_MISC`

These columns aggregate all other I/O operations, including `CREATE`, `DELETE`, `OPEN`, `CLOSE`, `STREAM_OPEN`, `STREAM_CLOSE`, `SEEK`, `TELL`, `FLUSH`, `STAT`, `FSTAT`, `CHSIZE`, `RENAME`, and `SYNC`. There are no byte counts for these operations.

The MySQL server uses several techniques to avoid I/O operations by caching information read from files, so it is possible that statements you might expect to result in I/O events will not. You may be able to ensure that I/O does occur by flushing caches or restarting the server to reset its state.

### 20.9.8.6. Table I/O and Lock Wait Summary Tables

The following sections describe the table I/O and lock wait summary tables:

- `table_io_waits_summary_by_index_usage`: Table I/O waits per index
- `table_io_waits_summary_by_table`: Table I/O waits per table
- `table_lock_waits_summary_by_table`: Table lock waits per table

#### 20.9.8.6.1. The `table_io_waits_summary_by_table` Table

The `table_io_waits_summary_by_table` table aggregates all table I/O wait events, as generated by the `wait/io/table/sql/handler` instrument. The grouping is by table.

`TRUNCATE TABLE` is permitted for table I/O summary tables. It resets the counters to zero rather than removing rows. Truncating this table also truncates the `table_io_waits_summary_by_index_usage` table.

The `table_io_waits_summary_by_table` table has these grouping columns to indicate how events are aggregated: `OBJECT_TYPE`, `OBJECT_SCHEMA`, and `OBJECT_NAME`. These columns have the same meaning as in the `events_waits_current` table. They identify the table to which the row applies.

`table_io_waits_summary_by_table` has the following summary columns containing aggregated values. As indicated in the column descriptions, some columns are more general and have values that are the same as the sum of the values of more fine-grained columns. For example, columns that aggregate all writes hold the sum of the corresponding columns that aggregate inserts, updates, and deletes. In this way, aggregations at higher levels are available directly without the need for user-defined views that sum lower-level columns.

- `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, `MAX_TIMER_WAIT`

These columns aggregate all I/O operations. They are the same as the sum of the corresponding `xxx_READ` and `xxx_WRITE` columns.

- `COUNT_READ`, `SUM_TIMER_READ`, `MIN_TIMER_READ`, `AVG_TIMER_READ`, `MAX_TIMER_READ`

These columns aggregate all read operations. They are the same as the sum of the corresponding `xxx_FETCH` columns.

- `COUNT_WRITE`, `SUM_TIMER_WRITE`, `MIN_TIMER_WRITE`, `AVG_TIMER_WRITE`, `MAX_TIMER_WRITE`

These columns aggregate all write operations. They are the same as the sum of the corresponding `xxx_INSERT`, `xxx_UPDATE`, and `xxx_DELETE` columns.

- `COUNT_FETCH`, `SUM_TIMER_FETCH`, `MIN_TIMER_FETCH`, `AVG_TIMER_FETCH`, `MAX_TIMER_FETCH`

These columns aggregate all fetch operations.

- `COUNT_INSERT`, `SUM_TIMER_INSERT`, `MIN_TIMER_INSERT`, `AVG_TIMER_INSERT`, `MAX_TIMER_INSERT`

These columns aggregate all insert operations.

- `COUNT_UPDATE`, `SUM_TIMER_UPDATE`, `MIN_TIMER_UPDATE`, `AVG_TIMER_UPDATE`, `MAX_TIMER_UPDATE`

These columns aggregate all update operations.

- `COUNT_DELETE`, `SUM_TIMER_DELETE`, `MIN_TIMER_DELETE`, `AVG_TIMER_DELETE`, `MAX_TIMER_DELETE`

These columns aggregate all delete operations.

#### 20.9.8.6.2. The `table_io_waits_summary_by_index_usage` Table

The `table_io_waits_summary_by_index_usage` table aggregates all table index I/O wait events, as generated by the `wait/io/table/sql/handler` instrument. The grouping is by table index.

`TRUNCATE TABLE` is permitted for table I/O summary tables. It resets the counters to zero rather than removing rows. This table is also truncated by truncation of the `table_io_waits_summary_by_table` table. A DDL operation that changes the index structure of a table may cause the per-index statistics to be reset.

The structure of `table_io_waits_summary_by_index_usage` is nearly identical to `table_io_waits_summary_by_table`. The only difference is the additional group column, `INDEX_NAME`, which corresponds to the name of the index that was used when the table I/O wait event was recorded:

- A value of `PRIMARY` indicates that table I/O used the primary index.
- A value of `NULL` means that table I/O used no index.
- Inserts are counted against `INDEX_NAME = NULL`.

### 20.9.8.6.3. The `table_lock_waits_summary_by_table` Table

The `table_lock_waits_summary_by_table` table aggregates all table lock wait events, as generated by the `wait/lock/table/sql/handler` instrument. The grouping is by table.

`TRUNCATE TABLE` is permitted for table lock summary tables. It resets the counters to zero rather than removing rows.

This table contains information about internal and external locks:

- An internal lock corresponds to a lock in the SQL layer. This is currently implemented by a call to `thr_lock()`. In event rows, these locks are distinguished by the `OPERATION` column, which will have one of these values:

```
read normal
read with shared locks
read high priority
read no insert
write allow write
write concurrent insert
write delayed
write low priority
write normal
```

- An external lock corresponds to a lock in the storage engine layer. This is currently implemented by a call to `handler::external_lock()`. In event rows, these locks are distinguished by the `OPERATION` column, which will have one of these values:

```
read external
write external
```

The `table_lock_waits_summary_by_table` table has these grouping columns to indicate how events are aggregated: `OBJECT_TYPE`, `OBJECT_SCHEMA`, and `OBJECT_NAME`. These columns have the same meaning as in the `events_waits_current` table. They identify the table to which the row applies.

`table_lock_waits_summary_by_table` has the following summary columns containing aggregated values. As indicated in the column descriptions, some columns are more general and have values that are the same as the sum of the values of more fine-grained columns. For example, columns that aggregate all writes hold the sum of the corresponding columns that aggregate inserts, updates, and deletes. In this way, aggregations at higher levels are available directly without the need for user-defined views that sum lower-level columns.

- `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, `MAX_TIMER_WAIT`

These columns aggregate all lock operations. They are the same as the sum of the corresponding `xxx_READ` and `xxx_WRITE` columns.

- `COUNT_READ`, `SUM_TIMER_READ`, `MIN_TIMER_READ`, `AVG_TIMER_READ`, `MAX_TIMER_READ`

These columns aggregate all read-lock operations. They are the same as the sum of the corresponding `xxx_READ_NORMAL`, `xxx_READ_WITH_SHARED_LOCKS`, `xxx_READ_HIGH_PRIORITY`, and `xxx_READ_NO_INSERT` columns.

- `COUNT_WRITE`, `SUM_TIMER_WRITE`, `MIN_TIMER_WRITE`, `AVG_TIMER_WRITE`, `MAX_TIMER_WRITE`

These columns aggregate all write-lock operations. They are the same as the sum of the corresponding `xxx_WRITE_ALLOW_WRITE`, `xxx_WRITE_CONCURRENT_INSERT`, `xxx_WRITE_DELAYED`, `xxx_WRITE_LOW_PRIORITY`, and `xxx_WRITE_NORMAL` columns.

- `COUNT_READ_NORMAL`, `SUM_TIMER_READ_NORMAL`, `MIN_TIMER_READ_NORMAL`, `AVG_TIMER_READ_NORMAL`, `MAX_TIMER_READ_NORMAL`

These columns aggregate internal read locks.

- `COUNT_READ_WITH_SHARED_LOCKS`, `SUM_TIMER_READ_WITH_SHARED_LOCKS`, `MIN_TIMER_READ_WITH_SHARED_LOCKS`, `AVG_TIMER_READ_WITH_SHARED_LOCKS`, `MAX_TIMER_READ_WITH_SHARED_LOCKS`

These columns aggregate internal read locks.

- `COUNT_READ_HIGH_PRIORITY`, `SUM_TIMER_READ_HIGH_PRIORITY`, `MIN_TIMER_READ_HIGH_PRIORITY`, `AVG_TIMER_READ_HIGH_PRIORITY`, `MAX_TIMER_READ_HIGH_PRIORITY`

These columns aggregate internal read locks.

- `COUNT_READ_NO_INSERT`, `SUM_TIMER_READ_NO_INSERT`, `MIN_TIMER_READ_NO_INSERT`, `AVG_TIMER_READ_NO_INSERT`, `MAX_TIMER_READ_NO_INSERT`

These columns aggregate internal read locks.

- `COUNT_READ_EXTERNAL`, `SUM_TIMER_READ_EXTERNAL`, `MIN_TIMER_READ_EXTERNAL`, `AVG_TIMER_READ_EXTERNAL`, `MAX_TIMER_READ_EXTERNAL`

These columns aggregate external read locks.

- `COUNT_WRITE_ALLOW_WRITE`, `SUM_TIMER_WRITE_ALLOW_WRITE`, `MIN_TIMER_WRITE_ALLOW_WRITE`, `AVG_TIMER_WRITE_ALLOW_WRITE`, `MAX_TIMER_WRITE_ALLOW_WRITE`

These columns aggregate internal write locks.

- `COUNT_WRITE_CONCURRENT_INSERT`, `SUM_TIMER_WRITE_CONCURRENT_INSERT`, `MIN_TIMER_WRITE_CONCURRENT_INSERT`, `AVG_TIMER_WRITE_CONCURRENT_INSERT`, `MAX_TIMER_WRITE_CONCURRENT_INSERT`

These columns aggregate internal write locks.

- `COUNT_WRITE_DELAYED`, `SUM_TIMER_WRITE_DELAYED`, `MIN_TIMER_WRITE_DELAYED`, `AVG_TIMER_WRITE_DELAYED`, `MAX_TIMER_WRITE_DELAYED`

These columns aggregate internal write locks.

As of MySQL 5.6.6, `DELAYED` inserts are deprecated, so these columns will be removed in a future release.

- `COUNT_WRITE_LOW_PRIORITY`, `SUM_TIMER_WRITE_LOW_PRIORITY`, `MIN_TIMER_WRITE_LOW_PRIORITY`, `AVG_TIMER_WRITE_LOW_PRIORITY`, `MAX_TIMER_WRITE_LOW_PRIORITY`

These columns aggregate internal write locks.

- `COUNT_WRITE_NORMAL`, `SUM_TIMER_WRITE_NORMAL`, `MIN_TIMER_WRITE_NORMAL`, `AVG_TIMER_WRITE_NORMAL`, `MAX_TIMER_WRITE_NORMAL`

These columns aggregate internal write locks.

- `COUNT_WRITE_EXTERNAL`, `SUM_TIMER_WRITE_EXTERNAL`, `MIN_TIMER_WRITE_EXTERNAL`, `AVG_TIMER_WRITE_EXTERNAL`, `MAX_TIMER_WRITE_EXTERNAL`

These columns aggregate external write locks.

### 20.9.8.7. Connection Summary Tables

The connection summary tables are similar to the corresponding [events\\_xxx\\_summary\\_by\\_thread\\_by\\_event\\_name](#) tables, except that aggregation occurs per account, user, or host, rather than by thread.

The Performance Schema maintains summary tables that aggregate connection statistics by event name and account, user, or host. Separate groups of tables are available that aggregate wait, stage, and statement events, which results in this set of connection summary tables:

- [events\\_waits\\_summary\\_by\\_account\\_by\\_event\\_name](#): Wait events summarized per account and event name
- [events\\_waits\\_summary\\_by\\_user\\_by\\_event\\_name](#): Wait events summarized per user name and event name
- [events\\_waits\\_summary\\_by\\_host\\_by\\_event\\_name](#): Wait events summarized per host name and event name
- [events\\_stages\\_summary\\_by\\_account\\_by\\_event\\_name](#): Stage events summarized per account and event name
- [events\\_stages\\_summary\\_by\\_user\\_by\\_event\\_name](#): Stage events summarized per user name and event name
- [events\\_stages\\_summary\\_by\\_host\\_by\\_event\\_name](#): Stage events summarized per host name and event name
- [events\\_statements\\_summary\\_by\\_account\\_by\\_event\\_name](#): Statement events summarized per account and event name
- [events\\_statements\\_summary\\_by\\_user\\_by\\_event\\_name](#): Statement events summarized per user name and event name
- [events\\_statements\\_summary\\_by\\_host\\_by\\_event\\_name](#): Statement events summarized per host name and event name

In other words, the connection summary tables have names of the form [events\\_xxx\\_summary\\_yyy\\_by\\_event\\_name](#), where *xxx* is *waits*, *stages*, or *statements*, and *yyy* is *account*, *user*, or *host*.

The connection summary tables provide an intermediate aggregation level:

- [xxx\\_summary\\_by\\_thread\\_by\\_event\\_name](#) tables are more detailed than connection summary tables
- [xxx\\_summary\\_global\\_by\\_event\\_name](#) tables are less detailed than connection summary tables

`TRUNCATE TABLE` is permitted for connection summary tables. It resets the counters to zero rather than removing rows. In addition, connection summary tables are implicitly truncated if a connection table on which they depend is truncated. [Table 20.1, “Effect of Implicit Table Truncation”](#), describes the relationship between connection table truncation and implicitly truncated tables.

**Table 20.1. Effect of Implicit Table Truncation**

Truncated Table	Implicitly Truncated Summary Tables
<code>accounts</code>	Tables with names matching <code>_%_by_account%</code> , <code>_%_by_thread%</code>
<code>hosts</code>	Tables with names matching <code>_%_by_account%</code> , <code>_%_by_host%</code> , <code>_%_by_thread%</code>
<code>users</code>	Tables with names matching <code>_%_by_account%</code> , <code>_%_by_user%</code> , <code>_%_by_thread%</code>

The connection summary tables have these grouping columns to indicate how events are aggregated:

- For tables with `_by_account` in the name, the `USER`, `HOST`, and `EVENT_NAME` columns group events per account and event name.
- For tables with `_by_host` in the name, the `HOST` and `EVENT_NAME` columns group events per host name and event name.
- For tables with `_by_user` in the name, the `USER` and `EVENT_NAME` columns group events per user name and event name.

The connection summary tables have these summary columns containing aggregated values: `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, and `MAX_TIMER_WAIT`. These are similar to the columns of the same names in the `events_waits_summary_by_instance` table.

The connection summary tables were added in MySQL 5.6.3.

### 20.9.8.8. Socket Summary Tables

The socket summary tables aggregate timer and byte count information for socket operations:

- `socket_summary_by_instance`: Aggregate timer and byte count statistics generated by the `wait/io/socket/*` instruments for all socket I/O operations, per socket instance. When a connection terminates, the row in `socket_summary_by_instance` corresponding to it is deleted.
- `socket_summary_by_event_name`: Aggregate timer and byte count statistics generated by the `wait/io/socket/*` instruments for all socket I/O operations, per socket instrument.

The socket summary tables do not aggregate waits generated by `idle` events while sockets are waiting for the next request from the client. For `idle` event aggregations, use the wait-event summary tables; see [Section 20.9.8.1, “Event Wait Summary Tables”](#).

`TRUNCATE TABLE` is permitted for socket summary tables. Except for `events_statements_summary_by_digest`, it resets the counters to zero rather than removing rows.

The socket summary tables have these grouping columns to indicate how events are aggregated:

- `socket_summary_by_instance` has an `OBJECT_INSTANCE_BEGIN` column. Each row summarizes events for a given object.
- `socket_summary_by_event_name` has an `EVENT_NAME` column. Each row summarizes events for a given instrument.

The socket summary tables have these summary columns containing aggregated values:

- `COUNT_STAR`, `SUM_TIMER_WAIT`, `MIN_TIMER_WAIT`, `AVG_TIMER_WAIT`, `MAX_TIMER_WAIT`

These columns aggregate all operations.

- `COUNT_READ`, `SUM_TIMER_READ`, `MIN_TIMER_READ`, `AVG_TIMER_READ`, `MAX_TIMER_READ`, `SUM_NUMBER_OF_BYTES_READ`

These columns aggregate all receive operations (`RECV`, `RECVFROM`, and `RECVMSG`).

- `COUNT_WRITE`, `SUM_TIMER_WRITE`, `MIN_TIMER_WRITE`, `AVG_TIMER_WRITE`, `MAX_TIMER_WRITE`, `SUM_NUMBER_OF_BYTES_WRITE`

These columns aggregate all send operations (`SEND`, `SENDTO`, and `SENDMSG`).

- `COUNT_MISC`, `SUM_TIMER_MISC`, `MIN_TIMER_MISC`, `AVG_TIMER_MISC`, `MAX_TIMER_MISC`

These columns aggregate all other socket operations, such as `CONNECT`, `LISTEN`, `ACCEPT`, `CLOSE`, and `SHUTDOWN`. There are no byte counts for these operations.

The `socket_summary_by_instance` table also has an `EVENT_NAME` column that indicates the class of the socket: `client_connection`, `server_tcpip_socket`, `server_unix_socket`. This column can be grouped on to isolate, for example, client activity from that of the server listening sockets.

These tables were added in MySQL 5.6.3.

## 20.9.9. Performance Schema Miscellaneous Tables

The following sections describe tables that do not fall into the table categories discussed in the preceding sections:

- `host_cache`: Information from the internal host cache.
- `performance_timers`: Which event timers are available.
- `threads`: Information about server threads.

### 20.9.9.1. The `host_cache` Table

The `host_cache` table provides access to the contents of the host cache, which contains client host name and IP address information and is used to avoid DNS lookups. (See [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#).) The `host_cache` table exposes the contents of the host cache so that it can be examined using `SELECT` statements. The Performance Schema must be enabled or this table is empty.

`FLUSH HOSTS` and `TRUNCATE TABLE host_cache` have the same effect: They clear the host cache. This also empties the `host_cache` table (because it is the visible representation of the cache) and unblocks any blocked hosts (see [Section C.5.2.6, “Host ‘host\\_name’ is blocked”](#).) `FLUSH HOSTS` requires the `RELOAD [755]` privilege. `TRUNCATE TABLE` requires the `DROP [754]` privilege for the `host_cache` table.

The `host_cache` table has these columns:

- `IP`

The IP address of the client that connected to the server, expressed as a string.

- `HOST`

The resolved DNS host name for that client IP, or `NULL` if the name is unknown.

- `HOST_VALIDATED`

Whether the IP-to-host name-to-IP DNS resolution was performed successfully for the client IP. If `HOST_VALIDATED` is `YES`, the `HOST` column is used as the host name corresponding to the IP so that calls to DNS can be avoided. While `HOST_VALIDATED` is `NO`, DNS resolution is attempted again for each connect, until it eventually completes with either a valid result or a permanent error. This information enables the server to avoid caching bad or missing host names during temporary DNS failures, which would affect clients forever.

- `SUM_CONNECT_ERRORS`

The number of connection errors that are deemed “blocking” (assessed against the `max_connect_errors [502]` system variable). Currently, only protocol handshake errors are counted, and only for hosts that passed validation (`HOST_VALIDATED = YES`).

- `COUNT_HOST_BLOCKED_ERRORS`

The number of connections that were blocked because `SUM_CONNECT_ERRORS` exceeded the value of the `max_connect_errors [502]` system variable.

- `COUNT_NAMEINFO_TRANSIENT_ERRORS`



The number of transient errors during IP-to-host name DNS resolution.

- [COUNT\\_NAMEINFO\\_PERMANENT\\_ERRORS](#)

The number of permanent errors during IP-to-host name DNS resolution.

- [COUNT\\_FORMAT\\_ERRORS](#)

The number of host name format errors. MySQL does not perform matching of `Host` column values in the `mysql.user` table against host names for which one or more of the initial components of the name are entirely numeric, such as `1.2.example.com`. The client IP address is used instead. For the rationale why this type of matching does not occur, see [Section 6.2.3, “Specifying Account Names”](#).

- [COUNT\\_ADDRINFO\\_TRANSIENT\\_ERRORS](#)

The number of transient errors during host name-to-IP reverse DNS resolution.

- [COUNT\\_ADDRINFO\\_PERMANENT\\_ERRORS](#)

The number of permanent errors during host name-to-IP reverse DNS resolution.

- [COUNT\\_FCRDNS\\_ERRORS](#)

The number of forward-confirmed reverse DNS errors. These errors occur when IP-to-host name-to-IP DNS resolution produces an IP address that does not match the client originating IP address.

- [COUNT\\_HOST\\_ACL\\_ERRORS](#)

The number of errors that occur because no user from the client host can possibly log in. In such cases, the server returns `ER_HOST_NOT_PRIVILEGED [3669]` and does not even ask for a user name or password.

- [COUNT\\_NO\\_AUTH\\_PLUGIN\\_ERRORS](#)

The number of errors due to requests for an unavailable authentication plugin. A plugin can be unavailable if, for example, it was never loaded or a load attempt failed.

- [COUNT\\_AUTH\\_PLUGIN\\_ERRORS](#)

The number of errors reported by authentication plugins.

An authentication plugin can report different error codes to indicate the root cause of a failure. Depending on the type of error, one of these columns is incremented: [COUNT\\_AUTHENTICATION\\_ERRORS](#), [COUNT\\_AUTH\\_PLUGIN\\_ERRORS](#), [COUNT\\_HANDSHAKE\\_ERRORS](#). New return codes are an optional extension to the existing plugin API. Unknown or unexpected plugin errors are counted in the [COUNT\\_AUTH\\_PLUGIN\\_ERRORS](#) column.

- [COUNT\\_HANDSHAKE\\_ERRORS](#)

The number of errors detected at the wire protocol level.

- [COUNT\\_PROXY\\_USER\\_ERRORS](#)

The number of errors detected when a proxy user A is proxied to another user B who does not exist.

- [COUNT\\_PROXY\\_USER\\_ACL\\_ERRORS](#)

The number of errors detected when a proxy user A is proxied to another user B who does exist but for whom A does not have the `PROXY [755]` privilege.

- [COUNT\\_AUTHENTICATION\\_ERRORS](#)

The number of errors caused by failed authentication.

- `COUNT_SSL_ERRORS`

The number of errors due to SSL problems.

- `COUNT_MAX_USER_CONNECTIONS_ERRORS`

The number of errors caused by exceeding per-user connection quotas. See [Section 6.3.4, “Setting Account Resource Limits”](#).

- `COUNT_MAX_USER_CONNECTIONS_PER_HOUR_ERRORS`

The number of errors caused by exceeding per-user connections-per-hour quotas. See [Section 6.3.4, “Setting Account Resource Limits”](#).

- `COUNT_DEFAULT_DATABASE_ERRORS`

The number of errors related to the default database. For example, the database did not exist or the user had no privileges for accessing it.

- `COUNT_INIT_CONNECT_ERRORS`

The number of errors caused by execution failures of statements in the `init_connect` [489] system variable value.

- `COUNT_LOCAL_ERRORS`

The number of errors local to the server implementation and not related to the network, authentication, or authorization. For example, out-of-memory conditions fall into this category.

- `COUNT_UNKNOWN_ERRORS`

The number of other, unknown errors not accounted for by other columns in this table. This column is reserved for future use, in case new error conditions must be reported, and if preserving the backward compatibility and table structure of the `host_cache` table is required.

- `FIRST_SEEN`

The timestamp of the first connection attempt seen from the client in the `IP` column.

- `LAST_SEEN`

The timestamp of the last connection attempt seen from the client in the `IP` column.

- `FIRST_ERROR_SEEN`

The timestamp of the first error seen from the client in the `IP` column.

- `LAST_ERROR_SEEN`

The timestamp of the last error seen from the client in the `IP` column.

The `host_cache` table was added in MySQL 5.6.5.

### 20.9.9.2. The `performance_timers` Table

The `performance_timers` table shows which event timers are available:

```
mysql> SELECT * FROM performance_timers;
```

TIMER_NAME	TIMER_FREQUENCY	TIMER_RESOLUTION	TIMER_OVERHEAD
CYCLE	2389029850	1	72

NANOSECOND	NULL	NULL	NULL
MICROSECOND	1000000	1	585
MILLISECOND	1035	1	738
TICK	101	1	630

If the values associated with a given timer name are `NULL`, that timer is not supported on your platform. The rows that do not contain `NULL` indicate which timers you can use in `setup_timers`.

The `performance_timers` table has these columns:

- `TIMER_NAME`

The name by which to refer to the timer when configuring the `setup_timers` table.

- `TIMER_FREQUENCY`

The number of timer units per second. For a cycle timer, the frequency is generally related to the CPU speed. For example, on a system with a 2.4GHz processor, the `CYCLE` may be close to 2400000000.

- `TIMER_RESOLUTION`

Indicates the number of timer units by which timer values increase. If a timer has a resolution of 10, its value increases by 10 each time.

- `TIMER_OVERHEAD`

The minimal number of cycles of overhead to obtain one timing with the given timer. The Performance Schema determines this value by invoking the timer 20 times during initialization and picking the smallest value. The total overhead really is twice this amount because the instrumentation invokes the timer at the start and end of each event. The timer code is called only for timed events, so this overhead does not apply for nontimed events.

### 20.9.9.3. The `threads` Table

The `threads` table contains a row for each server thread. Each row contains information about a thread and indicates whether monitoring is enabled for it:

```
mysql> SELECT * FROM threads\G
***** 1. row *****
      THREAD_ID: 1
        NAME: thread/sql/main
         TYPE: BACKGROUND
    PROCESSLIST_ID: NULL
    PROCESSLIST_USER: NULL
    PROCESSLIST_HOST: NULL
    PROCESSLIST_DB: NULL
    PROCESSLIST_COMMAND: NULL
    PROCESSLIST_TIME: 80284
    PROCESSLIST_STATE: NULL
    PROCESSLIST_INFO: NULL
    PARENT_THREAD_ID: NULL
         ROLE: NULL
    INSTRUMENTED: YES
...
***** 4. row *****
      THREAD_ID: 51
        NAME: thread/sql/one_connection
         TYPE: FOREGROUND
    PROCESSLIST_ID: 34
    PROCESSLIST_USER: paul
    PROCESSLIST_HOST: localhost
    PROCESSLIST_DB: performance_schema
    PROCESSLIST_COMMAND: Query
    PROCESSLIST_TIME: 0
    PROCESSLIST_STATE: Sending data
    PROCESSLIST_INFO: SELECT * FROM threads
```

```
PARENT_THREAD_ID: 1
          ROLE: NULL
          INSTRUMENTED: YES
...

```

If you have the [PROCESS \[755\]](#) privilege, you can see all threads. Otherwise, you can see only your own threads (that is, threads associated with the MySQL account that you are using).

The `threads` table has these columns:

- `THREAD_ID`

A unique thread identifier.

- `NAME`

The name associated with the thread instrumentation code in the server. For example, `thread/sql/one_connection` corresponds to the thread function in the code responsible for handling a user connection, and `thread/sql/main` stands for the `main()` function of the server.

- `TYPE`

The thread type, either `FOREGROUND` or `BACKGROUND`. User connection threads are foreground threads. Threads associated with internal server activity are background threads. Examples are internal `InnoDB` threads, “binlog dump” threads sending information to slaves, and slave I/O and SQL threads.

- `PROCESSLIST_ID`

For threads that are displayed in the `INFORMATION_SCHEMA.PROCESSLIST` table, this is the `PROCESSLIST.ID` value, which is also the value that `CONNECTION_ID()` [1267] would return within that thread. For background threads (threads not associated with a user connection), `PROCESSLIST_ID` is 0, so the values are not unique.

- `PROCESSLIST_USER`

The user associated with a foreground thread, `NULL` for a background thread.

- `PROCESSLIST_HOST`

The host name of the client associated with a foreground thread, `NULL` for a background thread.

- `PROCESSLIST_DB`

The default database for the thread, or `NULL` if there is none.

- `PROCESSLIST_COMMAND`

The type of command the thread is executing. For descriptions for thread commands, see [Section 8.12.5, “Examining Thread Information”](#). The value of this column corresponds to the `COM_xxx` commands of the client/server protocol and `Com_xxx` status variables. See [Section 5.1.6, “Server Status Variables”](#)

- `PROCESSLIST_TIME`

The time in seconds that the thread has been in its current state.

- `PROCESSLIST_STATE`

An action, event, or state that indicates what the thread is doing. For descriptions of `PROCESSLIST_STATE` values, see [Section 8.12.5, “Examining Thread Information”](#).

Most states correspond to very quick operations. If a thread stays in a given state for many seconds, there might be a problem that needs to be investigated.

- `PROCESSLIST_INFO`

The statement the thread is executing, or `NULL` if it is not executing any statement. The statement might be the one sent to the server, or an innermost statement if the statement executes other statements. For example, if a `CALL` statement executes a stored procedure that is executing a `SELECT` statement, the `PROCESSLIST_INFO` value shows the `SELECT` statement.

- `PARENT_THREAD_ID`

If this thread is a subthread (spawned by another thread), this is the `THREAD_ID` value of the spawning thread. Thread spawning occurs, for example, to handle insertion of rows from `INSERT DELAYED` statements.

- `ROLE`

Unused.

- `INSTRUMENTED`

Whether the thread is instrumented:

- For foreground threads, the initial `INSTRUMENTED` value is determined by whether the user account associated with the thread matches any row in the `setup_actors` table. Matching is based on the values of the `PROCESSLIST_USER` and `PROCESSLIST_HOST` columns.

If the thread spawns a subthread, matching occurs again for the subthread.

- For background threads, `INSTRUMENTED` is `YES` by default. `setup_actors` is not consulted because there is no associated user for background threads.
- For any thread, its `INSTRUMENTED` value can be changed during the lifetime of the thread. This is the only `threads` table column that can be modified.

For thread monitoring to occur, these things must be true:

- The `thread_instrumentation` consumer in the `setup_consumers` table must be `YES`.
- The `thread.INSTRUMENTED` column must be `YES`.
- Monitoring occurs only for those thread events produced from instruments that are enabled, as specified in the `setup_instruments` table.

The initial contents of the `threads` table are based on the threads in existence when Performance Schema initialization occurs. Thereafter, a new row is added each time the server creates a thread.

Removal of rows from the `threads` table occurs when threads end. For a thread associated with a client session, removal occurs when the session ends. If a client has auto-reconnect enabled and the session reconnects after a disconnect, the session will be associated with a new row in the `threads` table (one that has a different `PROCESSLIST_ID` value). The initial `INSTRUMENTED` value for the new thread may be different from that of the original thread: The `setup_actors` table may have changed in the meantime, and if the `INSTRUMENTED` value for the original thread was changed after it was initialized, that change does not carry over to the new thread.

The `threads` table columns with names having a prefix of `PROCESS_` provide information similar to that available from the `INFORMATION_SCHEMA.PROCESSLIST` table or the `SHOW PROCESSLIST` statement. Thus, all three sources provide thread-monitoring information. Use of `threads` differs from use of the other two sources in these ways:

- Access to `threads` does not require a mutex and has minimal impact on server performance. `INFORMATION_SCHEMA.PROCESSLIST` and `SHOW PROCESSLIST` have negative performance consequences because they require a mutex.

- `threads` provides additional information for each thread, such as whether it is a foreground or background thread, and the location within the server associated with the thread.
- `threads` provides information about background threads, so it can be used to monitor activity the other thread information sources cannot.
- You can turn thread monitoring on or off. To control monitoring of existing threads, set the `INSTRUMENTED` column of the `threads` table. To control the initial `INSTRUMENTED` value for new foreground threads, use the `setup_actors` table.

For these reasons, DBAs who perform server monitoring using `INFORMATION_SCHEMA.PROCESSLIST` or `SHOW PROCESSLIST` may wish to monitor using `threads` instead.

## 20.10. Performance Schema Option and Variable Reference

Table 20.2. Performance Schema Variable Reference

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">performance_schema</a> [2307]	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_accounts_lost</a> [2318]				Yes	Global	No
<a href="#">performance_schema_accounts_lost</a> [2308]	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_cond_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_cond_instances_lost</a> [2318]				Yes	Global	No
<a href="#">performance_schema_events_statements_history_size</a> [2308]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_size</a> [2309]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_long_size</a> [2309]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_statements_history_size</a> [2309]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_waits_history_long_size</a> [2310]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_events_waits_history_size</a> [2310]	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_file_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_file_handles_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_file_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_hosts_lost</a> [2318]				Yes	Global	No
<a href="#">performance_schema_hosts_size</a> [2310]	Yes	Yes	Yes		Global	No
<a href="#">performance-schema-instrument</a> [2306]	Yes	Yes				
<a href="#">Performance_schema_locker_lost</a> [2319]				Yes	Global	No
<a href="#">performance_schema_max_cond_classes</a> [2311]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_cond_instances</a> [2311]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_classes</a> [2311]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_handles</a> [2312]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_file_instances</a> [2312]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_mutex_classes</a> [2313]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_mutex_instances</a> [2313]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_rwlock_classes</a> [2313]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_rwlock_instances</a> [2313]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_socket_classes</a> [2313]	Yes	Yes	Yes		Global	No

Name	Cmd-Line	Option file	System Var	Status Var	Var Scope	Dynamic
<a href="#">performance_schema_max_socket_instances</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_stage_classes</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_statement_classes</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_table_handles</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_table_instances</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_thread_classes</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_max_thread_instances</a> [2314]	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_mutex_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_mutex_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_rwlock_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_rwlock_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_session_connect_attrs_lost</a> [2319]				Yes	Global	No
<a href="#">performance_schema_session_connect_attrs</a> [2316]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_setup_activity_size</a> [2317]	Yes	Yes	Yes		Global	No
<a href="#">performance_schema_setup_object_size</a> [2317]	Yes	Yes	Yes		Global	No
<a href="#">Performance_schema_socket_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_socket_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_stage_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_statement_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_table_handles_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_table_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_thread_classes_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_thread_instances_lost</a> [2318]				Yes	Global	No
<a href="#">Performance_schema_users_lost</a> [2319]				Yes	Global	No
<a href="#">performance_schema_users_size</a> [2317]	Yes	Yes	Yes		Global	No

## 20.11. Performance Schema Command Options

Performance Schema parameters can be specified at server startup on the command line or in option files to configure Performance Schema instruments and consumers. Runtime configuration is also possible in many cases (see [Section 20.2.3, “Performance Schema Runtime Configuration”](#)), but startup configuration must be used when runtime configuration is too late to affect instruments that have already been initialized during the startup process.

Performance Schema consumers and instruments can be configured at startup using the following syntax. For additional details, see [Section 20.2.2, “Performance Schema Startup Configuration”](#).

- `--performance-schema-consumer=consumer_name=value` [2306]

Configure a Performance Schema consumer. Consumer names in the `setup_consumers` table use underscores, but for consumers set at startup, dashes and underscores within the name are equivalent.

- `--performance-schema-instrument=instrument_name=value` [2306]

Configure a Performance Schema instrument. The name may be given as a pattern to configure instruments that match the pattern.



## 20.12. Performance Schema System Variables

The Performance Schema implements several system variables that provide configuration information:

```
mysql> SHOW VARIABLES LIKE 'perf%';
```

Variable_name	Value
performance_schema	ON
performance_schema_accounts_size	100
performance_schema_digests_size	200
performance_schema_events_stages_history_long_size	10000
performance_schema_events_stages_history_size	10
performance_schema_events_statements_history_long_size	10000
performance_schema_events_statements_history_size	10
performance_schema_events_waits_history_long_size	10000
performance_schema_events_waits_history_size	10
performance_schema_hosts_size	100
performance_schema_max_cond_classes	80
performance_schema_max_cond_instances	1000
performance_schema_max_file_classes	50
performance_schema_max_file_handles	32768
performance_schema_max_file_instances	10000
performance_schema_max_mutex_classes	200
performance_schema_max_mutex_instances	1000000
performance_schema_max_rwlock_classes	30
performance_schema_max_rwlock_instances	1000000
performance_schema_max_socket_classes	10
performance_schema_max_socket_instances	1000
performance_schema_max_stage_classes	150
performance_schema_max_statement_classes	165
performance_schema_max_table_handles	10000
performance_schema_max_table_instances	1000
performance_schema_max_thread_classes	50
performance_schema_max_thread_instances	1000
performance_schema_session_connect_attrs_size	512
performance_schema_setup_actors_size	100
performance_schema_setup_objects_size	100
performance_schema_users_size	100

Performance Schema system variables can be set at server startup on the command line or in option files, and many can be set at runtime. See [Section 20.10, “Performance Schema Option and Variable Reference”](#).

As of MySQL 5.6.6, the Performance Schema automatically sizes the values of several of its parameters at server startup if they are not set explicitly. For more information, see [Section 20.2.2, “Performance Schema Startup Configuration”](#).

Performance Schema variables have the following meanings:

- [performance\\_schema \[2307\]](#)

<b>Command-Line Format</b>	<code>--performance_schema=#</code>
<b>Option-File Format</b>	<code>performance_schema=#</code>
<b>Variable Name</b>	<code>performance_schema</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values (&lt;= 5.6.5)</b>
<b>Type</b>	<code>boolean</code>
<b>Default</b>	<code>OFF</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>
<b>Type</b>	<code>boolean</code>

	<b>Default</b>	ON
--	----------------	----

The value of this variable is `ON` or `OFF` to indicate whether the Performance Schema is enabled. By default, the value is `ON` by default as of MySQL 5.6.6 and `OFF` before that. At server startup, you can specify this variable with no value or a value of `ON` or `1` to enable it, or with a value of `OFF` or `0` to disable it.

- [performance\\_schema\\_accounts\\_size \[2308\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--performance_schema_accounts_size=#</code>	
<b>Option-File Format</b>	<code>performance_schema_accounts_size=#</code>	
<b>Variable Name</b>	<code>performance_schema_accounts_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10
	<b>Range</b>	0 .. 1048576
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)
	<b>Range</b>	-1 .. 1048576

The number of rows in the `accounts` table. If this variable is 0, the Performance Schema does not maintain connection statistics in the `accounts` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_digests\\_size \[2308\]](#)

The maximum number of rows in the `events_statements_summary_by_digest` table. This variable was added in MySQL 5.6.5.

- [performance\\_schema\\_events\\_stages\\_history\\_long\\_size \[2308\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--performance_schema_events_stages_history_long_size=#</code>	
<b>Option-File Format</b>	<code>performance_schema_events_stages_history_long_size=#</code>	
<b>Variable Name</b>	<code>performance_schema_events_stages_history_long_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows in the `events_stages_history_long` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_events\\_stages\\_history\\_size \[2309\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	--performance_schema_events_stages_history_size=#	
<b>Option-File Format</b>	performance_schema_events_stages_history_size=#	
<b>Variable Name</b>	performance_schema_events_stages_history_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows per thread in the `events_stages_history` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_events\\_statements\\_history\\_long\\_size \[2309\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	--performance_schema_events_statements_history_long_size=#	
<b>Option-File Format</b>	performance_schema_events_statements_history_long_size=#	
<b>Variable Name</b>	performance_schema_events_statements_history_long_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows in the `events_statements_history_long` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_events\\_statements\\_history\\_size \[2309\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	--performance_schema_events_statements_history_size=#	
<b>Option-File Format</b>	performance_schema_events_statements_history_size=#	
<b>Variable Name</b>	performance_schema_events_statements_history_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric

	<b>Default</b>	10
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows per thread in the `events_statements_history` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_events\\_waits\\_history\\_long\\_size \[2310\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_events_waits_history_long_size=#</code>	
<b>Option-File Format</b>	<code>performance_schema_events_waits_history_long_size=#</code>	
<b>Variable Name</b>	<code>performance_schema_events_waits_history_long_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows in the `events_waits_history_long` table.

- [performance\\_schema\\_events\\_waits\\_history\\_size \[2310\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_events_waits_history_size=#</code>	
<b>Option-File Format</b>	<code>performance_schema_events_waits_history_size=#</code>	
<b>Variable Name</b>	<code>performance_schema_events_waits_history_size</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The number of rows per thread in the `events_waits_history` table.

- [performance\\_schema\\_hosts\\_size \[2310\]](#)

<b>Version Introduced</b>	5.6.3
<b>Command-Line Format</b>	<code>--performance_schema_hosts_size=#</code>
<b>Option-File Format</b>	<code>performance_schema_hosts_size=#</code>
<b>Variable Name</b>	<code>performance_schema_hosts_size</code>
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No

	<b>Permitted Values (&lt;= 5.6.5)</b>	
<b>Type</b>	numeric	
<b>Default</b>	10	
<b>Range</b>	0 .. 1048576	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
<b>Type</b>	numeric	
<b>Default</b>	-1 (autosized)	
<b>Range</b>	-1 .. 1048576	

The number of rows in the `hosts` table. If this variable is 0, the Performance Schema does not maintain connection statistics in the `hosts` table. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_max\\_cond\\_classes \[2311\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_cond_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_cond_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_cond_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	80	

The maximum number of condition instruments.

- [performance\\_schema\\_max\\_cond\\_instances \[2311\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_cond_instances=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_cond_instances=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_cond_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
<b>Type</b>	numeric	
<b>Default</b>	1000	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
<b>Type</b>	numeric	
<b>Default</b>	-1 (autosized)	

The maximum number of instrumented condition objects.

- [performance\\_schema\\_max\\_file\\_classes \[2311\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_file_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_file_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_file_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	

	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	50	

The maximum number of file instruments.

- [performance\\_schema\\_max\\_file\\_handles \[2312\]](#)

<b>Command-Line Format</b>	--performance_schema_max_file_handles=#	
<b>Option-File Format</b>	performance_schema_max_file_handles=#	
<b>Variable Name</b>	performance_schema_max_file_handles	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	
<b>Default</b>	32768	

The maximum number of opened file objects.

The value of [performance\\_schema\\_max\\_file\\_handles \[2312\]](#) should be greater than the value of [open\\_files\\_limit \[518\]](#); [open\\_files\\_limit \[518\]](#) affects the maximum number of open file handles the server can support and [performance\\_schema\\_max\\_file\\_handles \[2312\]](#) affects how many of these file handles can be instrumented.

- [performance\\_schema\\_max\\_file\\_instances \[2312\]](#)

<b>Command-Line Format</b>	--performance_schema_max_file_instances=#	
<b>Option-File Format</b>	performance_schema_max_file_instances=#	
<b>Variable Name</b>	performance_schema_max_file_instances	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
<b>Type</b>	numeric	
<b>Default</b>	10000	
	<b>Permitted Values (&gt;= 5.6.6)</b>	
<b>Type</b>	numeric	
<b>Default</b>	-1 (autosized)	

The maximum number of instrumented file objects.

- [performance\\_schema\\_max\\_mutex\\_classes \[2312\]](#)

<b>Command-Line Format</b>	--performance_schema_max_mutex_classes=#	
<b>Option-File Format</b>	performance_schema_max_mutex_classes=#	
<b>Variable Name</b>	performance_schema_max_mutex_classes	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
<b>Type</b>	numeric	

	<b>Default</b>	200
--	----------------	-----

The maximum number of mutex instruments.

- [performance\\_schema\\_max\\_mutex\\_instances \[2313\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_mutex_instances=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_mutex_instances=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_mutex_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The maximum number of instrumented mutex objects.

- [performance\\_schema\\_max\\_rwlock\\_classes \[2313\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_rwlock_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_rwlock_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_rwlock_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	20

The maximum number of rwlock instruments.

- [performance\\_schema\\_max\\_rwlock\\_instances \[2313\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_rwlock_instances=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_rwlock_instances=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_rwlock_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The maximum number of instrumented rwlock objects.



- [performance\\_schema\\_max\\_socket\\_classes](#) [2314]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--performance_schema_max_socket_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_socket_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_socket_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>10</code>

The maximum number of socket instruments. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_max\\_socket\\_instances](#) [2314]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--performance_schema_max_socket_instances=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_socket_instances=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_socket_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>1000</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>

The maximum number of instrumented socket objects. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_max\\_stage\\_classes](#) [2314]

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	<code>--performance_schema_max_stage_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_stage_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_stage_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>100</code>

The maximum number of stage instruments. This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_max\\_statement\\_classes](#) [2314]

<b>Version Introduced</b>	5.6.3	2314
---------------------------	-------	------

<b>Command-Line Format</b>	<code>--performance_schema_max_statement_classes=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_statement_classes=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_statement_classes</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>autosized</code>

The maximum number of statement instruments. The default value is calculated at server build time based on the number of commands in the client/server protocol and the number of SQL statement types supported by the server.

This variable should not be changed, unless to set it to 0 to disable all statement instrumentation and save all memory associated with it. Setting the variable to nonzero values other than the default has no benefit; in particular, values larger than the default cause more memory to be allocated than is needed.

This variable was added in MySQL 5.6.3.

- [performance\\_schema\\_max\\_table\\_handles \[2315\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_table_handles=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_table_handles=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_table_handles</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>100000</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>-1 (autosized)</code>

The maximum number of opened table objects.

- [performance\\_schema\\_max\\_table\\_instances \[2315\]](#)

<b>Command-Line Format</b>	<code>--performance_schema_max_table_instances=#</code>	
<b>Option-File Format</b>	<code>performance_schema_max_table_instances=#</code>	
<b>Variable Name</b>	<code>performance_schema_max_table_instances</code>	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	<code>numeric</code>
	<b>Default</b>	<code>50000</code>
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	<code>numeric</code>

	<b>Default</b>	-1 (autosized)
--	----------------	----------------

The maximum number of instrumented table objects.

- [performance\\_schema\\_max\\_thread\\_classes](#) [2316]

<b>Command-Line Format</b>	--performance_schema_max_thread_classes=#	
<b>Option-File Format</b>	performance_schema_max_thread_classes=#	
<b>Variable Name</b>	performance_schema_max_thread_classes	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	50

The maximum number of thread instruments.

- [performance\\_schema\\_max\\_thread\\_instances](#) [2316]

<b>Command-Line Format</b>	--performance_schema_max_thread_instances=#	
<b>Option-File Format</b>	performance_schema_max_thread_instances=#	
<b>Variable Name</b>	performance_schema_max_thread_instances	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	1000
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)

The maximum number of instrumented thread objects.

The [max\\_connections](#) [503] and [max\\_delayed\\_threads](#) [504] system variables affect how many threads are run in the server. [performance\\_schema\\_max\\_thread\\_instances](#) [2316] affects how many of these running threads can be instrumented. If you increase [max\\_connections](#) [503] or [max\\_delayed\\_threads](#) [504], you should consider increasing [performance\\_schema\\_max\\_thread\\_instances](#) [2316] so that [performance\\_schema\\_max\\_thread\\_instances](#) [2316] is greater than the sum of [max\\_connections](#) [503] and [max\\_delayed\\_threads](#) [504].

- [performance\\_schema\\_session\\_connect\\_attrs\\_size](#) [2316]

<b>Version Introduced</b>	5.6.6
<b>Command-Line Format</b>	--performance_schema_session_connect_attrs_size=#
<b>Option-File Format</b>	performance_schema_session_connect_attrs_size=#
<b>Variable Name</b>	performance_schema_session_connect_attrs_size
<b>Variable Scope</b>	Global
<b>Dynamic Variable</b>	No
	<b>Permitted Values</b>

	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)
	<b>Range</b>	-1 .. 1048576

The amount of preallocated memory per thread used to hold connection attribute strings. If the connection attribute strings are larger than the reserved storage, the [Performance\\_schema\\_session\\_connect\\_attrs\\_lost \[2319\]](#) status variable is incremented. This variable was added in MySQL 5.6.7.

- [performance\\_schema\\_setup\\_actors\\_size \[2317\]](#)

<b>Version Introduced</b>	5.6.1	
<b>Command-Line Format</b>	--performance_schema_setup_actors_size=#	
<b>Option-File Format</b>	performance_schema_setup_actors_size=#	
<b>Variable Name</b>	performance_schema_setup_actors_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	100

The number of rows in the [setup\\_actors](#) table.

- [performance\\_schema\\_setup\\_objects\\_size \[2317\]](#)

<b>Version Introduced</b>	5.6.1	
<b>Command-Line Format</b>	--performance_schema_setup_objects_size=#	
<b>Option-File Format</b>	performance_schema_setup_objects_size=#	
<b>Variable Name</b>	performance_schema_setup_objects_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values</b>	
	<b>Type</b>	numeric
	<b>Default</b>	100

The number of rows in the [setup\\_objects](#) table.

- [performance\\_schema\\_users\\_size \[2317\]](#)

<b>Version Introduced</b>	5.6.3	
<b>Command-Line Format</b>	--performance_schema_users_size=#	
<b>Option-File Format</b>	performance_schema_users_size=#	
<b>Variable Name</b>	performance_schema_users_size	
<b>Variable Scope</b>	Global	
<b>Dynamic Variable</b>	No	
	<b>Permitted Values (&lt;= 5.6.5)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	10

	<b>Range</b>	0 .. 1048576
	<b>Permitted Values (&gt;= 5.6.6)</b>	
	<b>Type</b>	numeric
	<b>Default</b>	-1 (autosized)
	<b>Range</b>	-1 .. 1048576

The number of rows in the `users` table. If this variable is 0, the Performance Schema does not maintain connection statistics in the `users` table. This variable was added in MySQL 5.6.3.

## 20.13. Performance Schema Status Variables

The Performance Schema implements several status variables that provide information about instrumentation that could not be loaded or created due to memory constraints:

```
mysql> SHOW STATUS LIKE 'perf%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Performance_schema_accounts_lost | 0 |
| Performance_schema_cond_classes_lost | 0 |
| Performance_schema_cond_instances_lost | 0 |
| Performance_schema_file_classes_lost | 0 |
| Performance_schema_file_handles_lost | 0 |
| Performance_schema_file_instances_lost | 0 |
| Performance_schema_hosts_lost | 0 |
| Performance_schema_locker_lost | 0 |
| Performance_schema_mutex_classes_lost | 0 |
| Performance_schema_mutex_instances_lost | 0 |
| Performance_schema_rwlock_classes_lost | 0 |
| Performance_schema_rwlock_instances_lost | 0 |
| Performance_schema_socket_classes_lost | 0 |
| Performance_schema_socket_instances_lost | 0 |
| Performance_schema_stage_classes_lost | 0 |
| Performance_schema_statement_classes_lost | 0 |
| Performance_schema_table_handles_lost | 0 |
| Performance_schema_table_instances_lost | 0 |
| Performance_schema_thread_classes_lost | 0 |
| Performance_schema_thread_instances_lost | 0 |
| Performance_schema_users_lost | 0 |
+-----+-----+
```

Names for these variables have several forms:

- `Performance_schema_accounts_lost`

The number of times a row could not be added to the `accounts` table because it was full. This variable was added in MySQL 5.6.3.

- `Performance_schema_xxx_classes_lost`

How many instruments of type `xxx` could not be loaded.

- `Performance_schema_hosts_lost`

The number of times a row could not be added to the `hosts` table because it was full. This variable was added in MySQL 5.6.3.

- `Performance_schema_xxx_instances_lost`

How many instances of object type `xxx` could not be created.

- `Performance_schema_xxx_handles_lost`

How many instances of object type `xxx` could not be opened.

- `Performance_schema_locker_lost`

How many events are “lost” or not recorded, due to the following conditions:

- Events are recursive (for example, waiting for A caused a wait on B, which caused a wait on C).
- The depth of the nested events stack is greater than the limit imposed by the implementation.

Currently, events recorded by the Performance Schema are not recursive, so this variable should always be 0.

- `Performance_schema_session_connect_attrs_lost`

The number of times a connection attribute string was larger than the reserved storage. This variable was added in MySQL 5.6.7.

- `Performance_schema_users_lost`

The number of times a row could not be added to the `users` table because it was full. This variable was added in MySQL 5.6.3.

For information on using these variables to check Performance Schema status, see [Section 20.5, “Performance Schema Status Monitoring”](#).

## 20.14. Performance Schema and Plugins

Removing a plugin with `UNINSTALL PLUGIN` does not affect information already collected for code in that plugin. Time spent executing the code while the plugin was loaded was still spent even if the plugin is unloaded later. The associated event information, including aggregate information, remains readable in `performance_schema` database tables. For additional information about the effect of plugin installation and removal, see [Section 20.5, “Performance Schema Status Monitoring”](#).

A plugin implementor who instruments plugin code should document its instrumentation characteristics to enable those who load the plugin to account for its requirements. For example, a third-party storage engine should include in its documentation how much memory the engine needs for mutex and other instruments.

## 20.15. Using the Performance Schema to Diagnose Problems

The Performance Schema is a tool to help a DBA do performance tuning by taking real measurements instead of “wild guesses.” This section demonstrates some ways to use the Performance Schema for this purpose. The discussion here relies on the use of event filtering, which is described in [Section 20.2.3.2, “Performance Schema Event Filtering”](#).

The following example provides one methodology that you can use to analyze a repeatable problem, such as investigating a performance bottleneck. To begin, you should have a repeatable use case where performance is deemed “too slow” and needs optimization, and you should enable all instrumentation (no pre-filtering at all).

1. Run the use case.
2. Using the Performance Schema tables, analyze the root cause of the performance problem. This analysis will rely heavily on post-filtering.
3. For problem areas that are ruled out, disable the corresponding instruments. For example, if analysis shows that the issue is not related to file I/O in a particular storage engine, disable the file I/O instruments for that engine. Then truncate the history and summary tables to remove previously collected events.
4. Repeat the process at step 1.

At each iteration, the Performance Schema output, particularly the `events_waits_history_long` table, will contain less and less “noise” caused by nonsignificant instruments, and given that this table has a fixed size, will contain more and more data relevant to the analysis of the problem at hand.

At each iteration, investigation should lead closer and closer to the root cause of the problem, as the “signal/noise” ratio will improve, making analysis easier.

5. Once a root cause of performance bottleneck is identified, take the appropriate corrective action, such as:
  - Tune the server parameters (cache sizes, memory, and so forth).
  - Tune a query by writing it differently,
  - Tune the database schema (tables, indexes, and so forth).
  - Tune the code (this applies to storage engine or server developers only).
6. Start again at step 1, to see the effects of the changes on performance.

The `mutex_instances.LOCKED_BY_THREAD_ID` and `rwlock_instances.WRITE_LOCKED_BY_THREAD_ID` columns are extremely important for investigating performance bottlenecks or deadlocks. This is made possible by Performance Schema instrumentation as follows:

1. Suppose that thread 1 is stuck waiting for a mutex.
2. You can determine what the thread is waiting for:

```
SELECT * FROM events_waits_current WHERE THREAD_ID = thread_1;
```

Say the query result identifies that the thread is waiting for mutex A, found in `events_waits_current.OBJECT_INSTANCE_BEGIN`.

3. You can determine which thread is holding mutex A:

```
SELECT * FROM mutex_instances WHERE OBJECT_INSTANCE_BEGIN = mutex_A;
```

Say the query result identifies that it is thread 2 holding mutex A, as found in `mutex_instances.LOCKED_BY_THREAD_ID`.

4. You can see what thread 2 is doing:

```
SELECT * FROM events_waits_current WHERE THREAD_ID = thread_2;
```



---

# Chapter 21. Connectors and APIs

## Table of Contents

21.1. MySQL Connector/ODBC .....	2326
21.1.1. Connector/ODBC Versions .....	2327
21.1.2. Connector/ODBC Introduction .....	2327
21.1.3. Connector/ODBC Installation .....	2330
21.1.4. Configuring Connector/ODBC .....	2341
21.1.5. Connector/ODBC Examples .....	2360
21.1.6. Connector/ODBC Reference .....	2386
21.1.7. Connector/ODBC Notes and Tips .....	2392
21.1.8. Connector/ODBC Support .....	2402
21.2. MySQL Connector/Net .....	2403
21.2.1. Connector/Net Versions .....	2404
21.2.2. Connector/Net Installation .....	2406
21.2.3. Connector/Net Visual Studio Integration .....	2414
21.2.4. Connector/Net Tutorials .....	2444
21.2.5. Connector/Net Programming .....	2484
21.2.6. Connector/Net Connection String Options Reference .....	2523
21.2.7. Connector/Net API Reference .....	2530
21.2.8. Connector/Net Support .....	2633
21.2.9. Connector/Net FAQ .....	2634
21.3. MySQL Connector/J .....	2635
21.3.1. Overview of MySQL Connector/J .....	2635
21.3.2. Connector/J Versions .....	2635
21.3.3. Connector/J Installation .....	2637
21.3.4. Connector/J Examples .....	2641
21.3.5. Connector/J (JDBC) Reference .....	2642
21.3.6. JDBC Concepts .....	2675
21.3.7. Connection Pooling with Connector/J .....	2684
21.3.8. Load Balancing with Connector/J .....	2687
21.3.9. Failover with Connector/J .....	2689
21.3.10. Using the Connector/J Interceptor Classes .....	2691
21.3.11. Using Connector/J with Tomcat .....	2691
21.3.12. Using Connector/J with JBoss .....	2693
21.3.13. Using Connector/J with Spring .....	2694
21.3.14. Using Connector/J with GlassFish .....	2698
21.3.15. Troubleshooting Connector/J Applications .....	2704
21.3.16. Connector/J Support .....	2712
21.4. MySQL Connector/C++ .....	2713
21.4.1. How to Get MySQL Connector/C++ .....	2715
21.4.2. Installing MySQL Connector/C++ from a Binary Distribution .....	2716
21.4.3. Installing MySQL Connector/C++ from Source .....	2718
21.4.4. Building MySQL Connector/C++ Windows Applications with Microsoft Visual Studio .....	2723
21.4.5. Building MySQL Connector/C++ Linux Applications with NetBeans .....	2731
21.4.6. MySQL Connector/C++ Getting Started: Usage Examples .....	2735
21.4.7. MySQL Connector/C++ Tutorials .....	2741
21.4.8. MySQL Connector/C++ Debug Tracing .....	2749
21.4.9. MySQL Connector/C++ Usage Notes .....	2750
21.4.10. MySQL Connector/C++ Known Bugs and Issues .....	2754
21.4.11. MySQL Connector/C++ Support .....	2755
21.5. MySQL Connector/C .....	2755
21.5.1. MySQL Connector/C Supported Platforms .....	2755
21.5.2. MySQL Connector/C Distribution Contents .....	2756
21.5.3. Obtaining MySQL Connector/C .....	2756

21.5.4. Installing MySQL Connector/C from a Binary Distribution .....	2756
21.5.5. Installing MySQL Connector/C from Source .....	2756
21.5.6. Testing MySQL Connector/C .....	2758
21.5.7. Building MySQL Connector/C Applications .....	2758
21.6. MySQL Connector/Python .....	2759
21.6.1. Guidelines for Python Developers .....	2759
21.6.2. Connector/Python Versions .....	2761
21.6.3. Connector/Python Installation .....	2761
21.6.4. Connector/Python Coding Examples .....	2762
21.6.5. Connector/Python Tutorials .....	2768
21.6.6. Connector/Python Connection Arguments .....	2769
21.6.7. Connector/Python API Reference .....	2772
21.6.8. Connector/Python FAQ .....	2790
21.7. libmysqld, the Embedded MySQL Server Library .....	2790
21.7.1. Compiling Programs with <code>libmysqld</code> .....	2790
21.7.2. Restrictions When Using the Embedded MySQL Server .....	2791
21.7.3. Options with the Embedded Server .....	2791
21.7.4. Embedded Server Examples .....	2792
21.8. MySQL C API .....	2795
21.8.1. MySQL C API Implementations .....	2796
21.8.2. Example C API Client Programs .....	2797
21.8.3. Building and Running C API Client Programs .....	2797
21.8.4. C API Data Structures .....	2800
21.8.5. C API Function Overview .....	2805
21.8.6. C API Function Descriptions .....	2809
21.8.7. C API Prepared Statements .....	2859
21.8.8. C API Prepared Statement Data Structures .....	2859
21.8.9. C API Prepared Statement Function Overview .....	2865
21.8.10. C API Prepared Statement Function Descriptions .....	2868
21.8.11. C API Threaded Function Descriptions .....	2890
21.8.12. C API Embedded Server Function Descriptions .....	2892
21.8.13. C API Client Plugin Functions .....	2892
21.8.14. Common Questions and Problems When Using the C API .....	2895
21.8.15. Controlling Automatic Reconnection Behavior .....	2897
21.8.16. C API Support for Multiple Statement Execution .....	2898
21.8.17. C API Prepared Statement Problems .....	2900
21.8.18. C API Prepared Statement Handling of Date and Time Values .....	2900
21.8.19. C API Support for Prepared <code>CALL</code> Statements .....	2902
21.9. MySQL PHP API .....	2905
21.9.1. Overview of the MySQL PHP drivers .....	2906
21.9.2. Original MySQL API ( <code>Mysql</code> ) .....	2912
21.9.3. MySQL Improved Extension ( <code>Mysqli</code> ) .....	2982
21.9.4. MySQL Functions (PDO_MYSQL) ( <code>MySQL (PDO)</code> ) .....	3196
21.9.5. MySQL Native Driver ( <code>Mysqlnd</code> ) .....	3200
21.9.6. Mysqlnd replication and load balancing plugin ( <code>mysqlnd_ms</code> ) .....	3232
21.9.7. Mysqlnd query result cache plugin ( <code>mysqlnd_qc</code> ) .....	3333
21.9.8. Mysqlnd user handler plugin ( <code>mysqlnd_uh</code> ) .....	3384
21.9.9. Mysqlnd connection multiplexing plugin ( <code>mysqlnd_mux</code> ) .....	3467
21.9.10. Mysqlnd Memcache plugin ( <code>mysqlnd_memcache</code> ) .....	3471
21.9.11. Connector/PHP .....	3480
21.9.12. Common Problems with MySQL and PHP .....	3480
21.9.13. Enabling Both <code>mysql</code> and <code>mysqli</code> in PHP .....	3480
21.10. MySQL Perl API .....	3481
21.11. MySQL Python API .....	3482
21.12. MySQL Ruby APIs .....	3482
21.12.1. The MySQL/Ruby API .....	3482
21.12.2. The Ruby/MySQL API .....	3482
21.13. MySQL Tcl API .....	3482

MySQL Connectors provide connectivity to the MySQL server for client programs. APIs provide low-level access to the MySQL protocol and MySQL resources. Both Connectors and the APIs enable you to connect and execute MySQL statements from another language or environment, including ODBC, Java (JDBC), Perl, Python, PHP, Ruby, and native C and embedded MySQL instances.



#### Note

Connector version numbers do not correlate with MySQL Server version numbers. See [Table 21.2, “MySQL Connector Versions and MySQL Server Versions”](#).

## MySQL Connectors

Oracle develops a number of connectors:

- [Connector/ODBC](#) provides driver support for connecting to MySQL using the Open Database Connectivity (ODBC) API. Support is available for ODBC connectivity from Windows, Unix, and Mac OS X platforms.
- [Connector/Net](#) enables developers to create .NET applications that connect to MySQL. Connector/Net implements a fully functional ADO.NET interface and provides support for use with ADO.NET aware tools. Applications that use Connector/Net can be written in any supported .NET language.

The MySQL [Visual Studio Plugin](#) works with Connector/Net and Visual Studio 2005. The plugin is a MySQL DDEX Provider, which means that you can use the schema and data manipulation tools available in Visual Studio to create and edit objects within a MySQL database.

- [Connector/J](#) provides driver support for connecting to MySQL from Java applications using the standard Java Database Connectivity (JDBC) API.
- [Connector/Python](#) provides driver support for connecting to MySQL from Python applications using an API that is compliant with the [Python DB API version 2.0](#). No additional Python modules or MySQL client libraries are required.
- [Connector/C++](#) enables C++ applications to connect to MySQL.
- [Connector/C](#) is a standalone replacement for the MySQL Client Library (`libmysqlclient`), to be used for C applications.

## The MySQL C API

For direct access to using MySQL natively within a C application, there are two methods:

- The [C API](#) provides low-level access to the MySQL client/server protocol through the `libmysqlclient` client library. This is the primary method used to connect to an instance of the MySQL server, and is used both by MySQL command-line clients and many of the MySQL Connectors and third-party APIs detailed here.

`libmysqlclient` is included in MySQL distributions and in MySQL Connector/C distributions.

- `libmysqld` is an embedded MySQL server library that enables you to embed an instance of the MySQL server into your C applications.

`libmysqld` is included in MySQL distributions, but not in MySQL Connector/C distributions.

See also [Section 21.8.1, “MySQL C API Implementations”](#).

To access MySQL from a C application, or to build an interface to MySQL for a language not supported by the Connectors or APIs in this chapter, the [C API](#) is where to start. A number of programmer's

utilities are available to help with the process; see [Section 4.7, “MySQL Program Development Utilities”](#).

## Third-Party MySQL APIs

The remaining APIs described in this chapter provide an interface to MySQL from specific application languages. These third-party solutions are not developed or supported by Oracle. Basic information on their usage and abilities is provided here for reference purposes only.

All the third-party language APIs are developed using one of two methods, using `libmysqlclient` or by implementing a *native driver*. The two solutions offer different benefits:

- Using `libmysqlclient` offers complete compatibility with MySQL because it uses the same libraries as the MySQL client applications. However, the feature set is limited to the implementation and interfaces exposed through `libmysqlclient` and the performance may be lower as data is copied between the native language, and the MySQL API components.
- *Native drivers* are an implementation of the MySQL network protocol entirely within the host language or environment. Native drivers are fast, as there is less copying of data between components, and they can offer advanced functionality not available through the standard MySQL API. Native drivers are also easier for end users to build and deploy because no copy of the MySQL client libraries is needed to build the native driver components.

[Table 21.1, “MySQL APIs and Interfaces”](#) lists many of the libraries and interfaces available for MySQL. [Table 21.2, “MySQL Connector Versions and MySQL Server Versions”](#) shows which MySQL Server versions each connector supports.

**Table 21.1. MySQL APIs and Interfaces**

Environment	API	Type	Notes
Ada	GNU Ada MySQL Bindings	<code>libmysqlclient</code>	See <a href="#">MySQL Bindings for GNU Ada</a>
C	C API	<code>libmysqlclient</code>	See <a href="#">Section 21.8, “MySQL C API”</a> .
C	Connector/C	Replacement for <code>libmysqlclient</code>	See <a href="#">Section 21.5, “MySQL Connector/C”</a> .
C++	Connector/C++	<code>libmysqlclient</code>	See <a href="#">Section 21.4, “MySQL Connector/C++”</a> .
	MySQL++	<code>libmysqlclient</code>	See <a href="#">MySQL++ Web site</a> .
	MySQL wrapped	<code>libmysqlclient</code>	See <a href="#">MySQL wrapped</a> .
Cocoa	MySQL-Cocoa	<code>libmysqlclient</code>	Compatible with the Objective-C Cocoa environment. See <a href="http://mysql-cocoa.sourceforge.net/">http://mysql-cocoa.sourceforge.net/</a>
D	MySQL for D	<code>libmysqlclient</code>	See <a href="#">MySQL for D</a> .
Eiffel	Eiffel MySQL	<code>libmysqlclient</code>	See <a href="#">Section 21.14, “MySQL Eiffel Wrapper”</a> .
Erlang	<code>erlang-mysql-driver</code>	<code>libmysqlclient</code>	See <a href="#">erlang-mysql-driver</a> .
Haskell	Haskell MySQL Bindings	Native Driver	See <a href="#">Brian O’Sullivan’s pure Haskell MySQL bindings</a> .
	<code>hsqldb-mysql</code>	<code>libmysqlclient</code>	See <a href="#">MySQL driver for Haskell</a> .
Java/JDBC	Connector/J	Native Driver	See <a href="#">Section 21.3, “MySQL Connector/J”</a> .
Kaya	MyDB	<code>libmysqlclient</code>	See <a href="#">MyDB</a> .
Lua	LuaSQL	<code>libmysqlclient</code>	See <a href="#">LuaSQL</a> .

Environment	API	Type	Notes
.NET/Mono	Connector/Net	Native Driver	See <a href="#">Section 21.2, “MySQL Connector/Net”</a> .
Objective Caml	Objective Caml MySQL Bindings	<code>libmysqlclient</code>	See <a href="#">MySQL Bindings for Objective Caml</a> .
Octave	Database bindings for GNU Octave	<code>libmysqlclient</code>	See <a href="#">Database bindings for GNU Octave</a> .
ODBC	Connector/ODBC	<code>libmysqlclient</code>	See <a href="#">Section 21.1, “MySQL Connector/ODBC”</a> .
Perl	<code>DBI/DBD::mysql</code>	<code>libmysqlclient</code>	See <a href="#">Section 21.10, “MySQL Perl API”</a> .
	<code>Net::MySQL</code>	Native Driver	See <a href="#">Net::MySQL at CPAN</a>
PHP	<code>mysql</code> , <code>ext/mysql</code> interface (deprecated)	<code>libmysqlclient</code>	See <a href="#">Section 21.9.2, “Original MySQL API (mysql)”</a> .
	<code>mysqli</code> , <code>ext/mysqli</code> interface	<code>libmysqlclient</code>	See <a href="#">Section 21.9.3, “MySQL Improved Extension (mysqli)”</a> .
	<code>PDO_MYSQL</code>	<code>libmysqlclient</code>	See <a href="#">Section 21.9.4, “MySQL Functions (PDO_MYSQL) (MySQL (PDO))”</a> .
	PDO <code>mysqlnd</code>	Native Driver	
Python	Connector/Python	Native Driver	See <a href="#">Section 21.6, “MySQL Connector/Python”</a> .
	MySQLdb	<code>libmysqlclient</code>	See <a href="#">Section 21.11, “MySQL Python API”</a> .
Ruby	MySQL/Ruby	<code>libmysqlclient</code>	Uses <code>libmysqlclient</code> . See <a href="#">Section 21.12.1, “The MySQL/Ruby API”</a> .
	Ruby/MySQL	Native Driver	See <a href="#">Section 21.12.2, “The Ruby/MySQL API”</a> .
Scheme	<code>Myscsh</code>	<code>libmysqlclient</code>	See <a href="#">Myscsh</a> .
SPL	<code>sql_mysql</code>	<code>libmysqlclient</code>	See <a href="#">sql_mysql</a> for SPL.
Tcl	MySQLtcl	<code>libmysqlclient</code>	See <a href="#">Section 21.13, “MySQL Tcl API”</a> .

Table 21.2. MySQL Connector Versions and MySQL Server Versions

Connector	Connector version	MySQL Server version
Connector/C++	1.0.5 GA	5.6, 5.5, 5.4, 5.1
Connector/J	5.1.8	5.6, 5.5, 5.4, 5.1, 5.0, 4.1
Connector/Net	6.5	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	6.4	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	6.3	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	6.2 (No longer supported)	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	6.1 (No longer supported)	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	6.0 (No longer supported)	5.6, 5.5, 5.4, 5.1, 5.0
Connector/Net	5.2 (No longer supported)	5.6, 5.5, 5.4, 5.1, 5.0

Connector	Connector version	MySQL Server version
Connector/Net	1.0 (No longer supported)	5.0, 4.0
Connector/ODBC	5.1	5.6, 5.5, 5.4, 5.1, 5.0, 4.1.1+
Connector/ODBC	3.51 (Unicode not supported)	5.6, 5.5, 5.4, 5.1, 5.0, 4.1

## 21.1. MySQL Connector/ODBC

The MySQL Connector/ODBC is the name for the family of MySQL ODBC drivers (previously called MyODBC drivers) that provide access to a MySQL database using the industry standard Open Database Connectivity (ODBC) API. This reference covers Connector/ODBC 5.2, which includes the functionality of the Unicode driver and the ANSI driver, which formerly were split between Connector/ODBC 5.1 and Connector/ODBC 3.51.

MySQL Connector/ODBC provides both driver-manager based and native interfaces to the MySQL database, with full support for MySQL functionality, including stored procedures, [transactions](#) and, with Connector/ODBC 5.1 and higher, full Unicode compliance.

For more information on the ODBC API standard and how to use it, refer to <http://support.microsoft.com/kb/110093>.

The application development section of the ODBC API reference assumes a good working knowledge of C, general DBMS, and a familiarity with MySQL. For more information about MySQL functionality and its syntax, refer to <http://dev.mysql.com/doc/>.

Typically, you need to install Connector/ODBC only on Windows machines. For Unix and Mac OS X, you can use the native MySQL network or named pipes to communicate with your MySQL database. You may need Connector/ODBC for Unix or Mac OS X if you have an application that requires an ODBC interface to communicate with the database. Applications that require ODBC to communicate with MySQL include ColdFusion, Microsoft Office, and Filemaker Pro.

For release notes detailing the changes in each release of Connector/ODBC, see [MySQL Connector/ODBC Release Notes](#).

### Key connector/ODBC topics include:

- Installing Connector/ODBC: [Section 21.1.3, “Connector/ODBC Installation”](#).
- The configuration options: [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#).
- An example that connects to a MySQL database from a Windows host: [Section 21.1.5.2, “Step-by-step Guide to Connecting to a MySQL Database through Connector/ODBC”](#).
- An example that uses Microsoft Access as an interface to a MySQL database: [Section 21.1.5.4, “Using Connector/ODBC with Microsoft Access”](#).
- General tips and notes, including how to obtain the last auto-increment ID: [Section 21.1.7.1, “Connector/ODBC General Functionality”](#).
- Application-specific usage tips and notes: [Section 21.1.7.2, “Connector/ODBC Application-Specific Tips”](#).
- A FAQ (Frequently Asked Questions) list: [Section 21.1.7.3, “Connector/ODBC Errors and Resolutions \(FAQ\)”](#).
- Additional Connector/ODBC support options: [Section 21.1.8, “Connector/ODBC Support”](#).

## 21.1.1. Connector/ODBC Versions

These are the versions of Connector/ODBC that are currently available:

- Connector/ODBC 5.2, currently in GA status, upgrades the ANSI driver of Connector/ODBC 3.51 to the 5.x code base. It also includes new features, such as enabling server-side prepared statements by default. At installation time, you can choose the Unicode driver for the broadest compatibility with data sources using various character sets, or the ANSI driver for optimal performance with a more limited range of character sets. It works with MySQL versions 4.1.1 and higher.
- Connector/ODBC 5.1, now superseded by 5.2, is a partial rewrite of the of the 3.51 code base, and is designed to work with MySQL versions 4.1.1 and newer.

Connector/ODBC 5.1 also includes the following changes and improvements over the 3.51 release:

- Improved support on Windows 64-bit platforms.
- Full Unicode support at the driver level. This includes support for the `SQL_WCHAR` data type, and support for Unicode login, password and DSN configurations. For more information, see [Microsoft Knowledgebase Article #716246](#).
- Support for the `SQL_NUMERIC_STRUCT` data type, which provides easier access to the precise definition of numeric values. For more information, see [Microsoft Knowledgebase Article #714556](#).
- Native Windows setup library. This replaces the Qt library based interface for configuring DSN information within the ODBC Data Sources application.
- Support for the ODBC descriptor, which improves the handling and metadata of columns and parameter data. For more information, see [Microsoft Knowledgebase Article #716339](#).
- Connector/ODBC 3.51 is the current release of the 32-bit ODBC driver, also known as the MySQL ODBC 3.51 driver. Connector/ODBC 3.51 has support for ODBC 3.5x specification level 1 (complete core API + level 2 features) to continue to provide all functionality of ODBC for accessing MySQL.

With MySQL Connector/ODBC 5.2 now offering both ANSI and Unicode drivers, you should expect to upgrade relatively soon.

The manual for versions of Connector/ODBC older than 5.2 can be located in the corresponding binary or source distribution. Please note that versions of Connector/ODBC earlier than the 3.51 revision were not fully compliant with the ODBC specification.



### Note

From this section onward, the primary focus of this guide is the Connector/ODBC 5.2 driver.



### Note

Version numbers for MySQL products are formatted as X.X.X. However, Windows tools (Control Panel, properties display) may show the version numbers as XX.XX.XX. For example, the official MySQL formatted version number 5.0.9 may be displayed by Windows tools as 5.00.09. The two versions are the same; only the number display format is different.

## 21.1.2. Connector/ODBC Introduction

ODBC (Open Database Connectivity) provides a way for client programs to access a wide range of databases or data sources. ODBC is a standardized API that enables connections to SQL database servers. It was developed according to the specifications of the SQL Access Group and defines a set of function calls, error codes, and data types that can be used to develop database-independent



applications. ODBC usually is used when database independence or simultaneous access to different data sources is required.

For more information about ODBC, refer to <http://support.microsoft.com/kb/110093>.

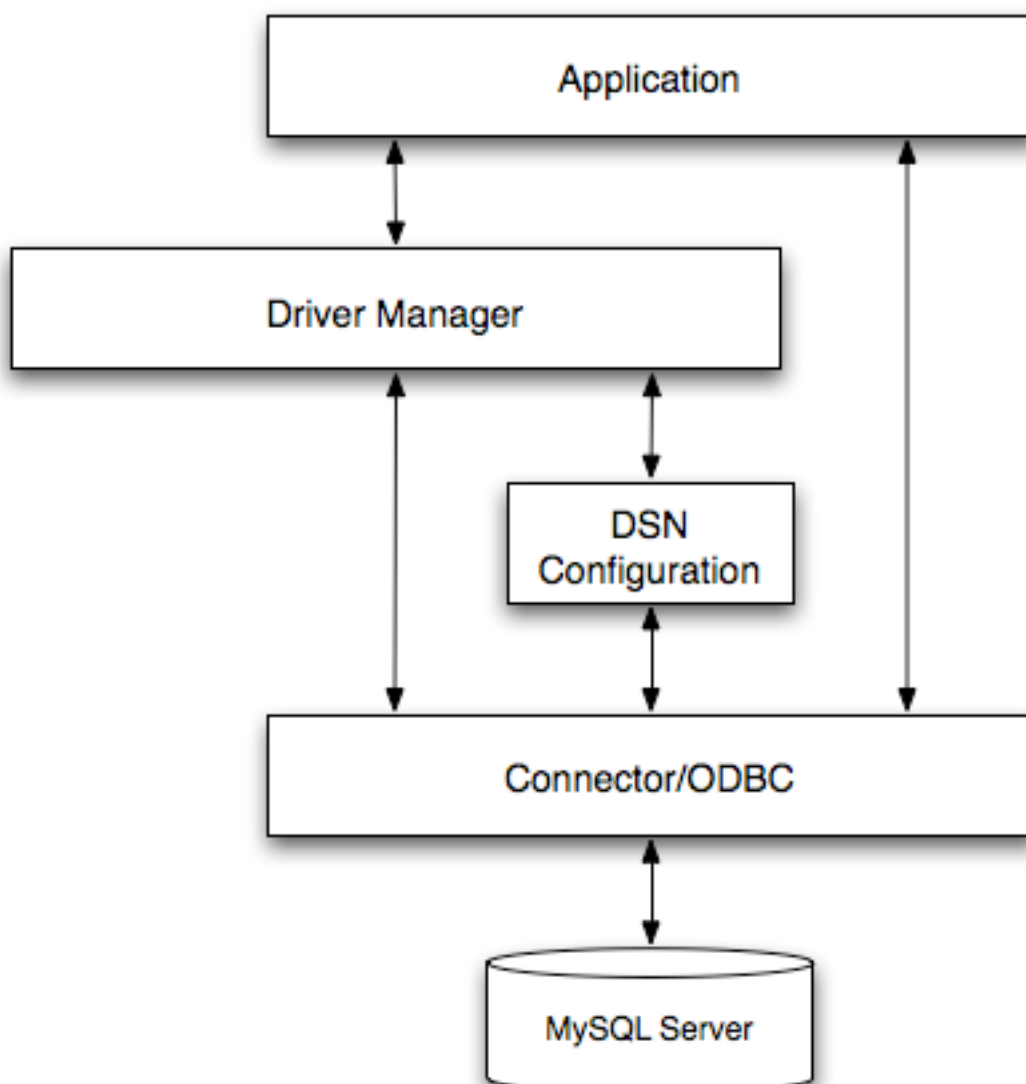
### 21.1.2.1. General Information About ODBC and Connector/ODBC

Open Database Connectivity (ODBC) is a widely accepted application-programming interface (API) for database access. It is based on the Call-Level Interface (CLI) specifications from X/Open and ISO/IEC for database APIs and uses Structured Query Language (SQL) as its database access language.

A survey of ODBC functions supported by Connector/ODBC is given at [Section 21.1.6.1, "Connector/ODBC API Reference"](#). For general information about ODBC, see <http://support.microsoft.com/kb/110093>.

#### 21.1.2.1.1. Connector/ODBC Architecture

The Connector/ODBC architecture is based on five components, as shown in the following diagram:



- **Application:**

The Application uses the ODBC API to access the data from the MySQL server. The ODBC API in turn communicates with the Driver Manager. The Application communicates with the Driver Manager

using the standard ODBC calls. The Application does not care where the data is stored, how it is stored, or even how the system is configured to access the data. It needs to know only the Data Source Name (DSN).

A number of tasks are common to all applications, no matter how they use ODBC. These tasks are:

- Selecting the MySQL server and connecting to it.
- Submitting SQL statements for execution.
- Retrieving results (if any).
- Processing errors.
- [Committing](#) or [rolling back](#) the [transaction](#) enclosing the SQL statement.
- Disconnecting from the MySQL server.

Because most data access work is done with SQL, the primary tasks for applications that use ODBC are submitting SQL statements and retrieving any results generated by those statements.

- **Driver manager:**

The Driver Manager is a library that manages communication between application and driver or drivers. It performs the following tasks:

- Resolves Data Source Names (DSN). The DSN is a configuration string that identifies a given database driver, database, database host and optionally authentication information that enables an ODBC application to connect to a database using a standardized reference.

Because the database connectivity information is identified by the DSN, any ODBC compliant application can connect to the data source using the same DSN reference. This eliminates the need to separately configure each application that needs access to a given database; instead you instruct the application to use a pre-configured DSN.

- Loading and unloading of the driver required to access a specific database as defined within the DSN. For example, if you have configured a DSN that connects to a MySQL database then the driver manager will load the Connector/ODBC driver to enable the ODBC API to communicate with the MySQL host.
- Processes ODBC function calls or passes them to the driver for processing.

- **Connector/ODBC Driver:**

The Connector/ODBC driver is a library that implements the functions supported by the ODBC API. It processes ODBC function calls, submits SQL requests to MySQL server, and returns results back to the application. If necessary, the driver modifies an application's request so that the request conforms to syntax supported by MySQL.

- **DSN Configuration:**

The ODBC configuration file stores the driver and database information required to connect to the server. It is used by the Driver Manager to determine which driver to be loaded according to the definition in the DSN. The driver uses this to read connection parameters based on the DSN specified. For more information, [Section 21.1.4, "Configuring Connector/ODBC"](#).

- **MySQL Server:**

The MySQL database where the information is stored. The database is used as the source of the data (during queries) and the destination for data (during inserts and updates).

### 21.1.2.1.2. ODBC Driver Managers

An ODBC Driver Manager is a library that manages communication between the ODBC-aware application and any drivers. Its main functionality includes:

- Resolving Data Source Names (DSN).
- Driver loading and unloading.
- Processing ODBC function calls or passing them to the driver.

Both Windows and Mac OS X include ODBC driver managers with the operating system. Most ODBC Driver Manager implementations also include an administration application that makes the configuration of DSN and drivers easier. Examples and information on these managers, including Unix ODBC driver managers, are listed below:

- Microsoft Windows ODBC Driver Manager (`odbc32.dll`), <http://support.microsoft.com/kb/110093>.
- Mac OS X includes `ODBC Administrator`, a GUI application that provides a simpler configuration mechanism for the Unix iODBC Driver Manager. You can configure DSN and driver information either through ODBC Administrator or through the iODBC configuration files. This also means that you can test ODBC Administrator configurations using the `iodbctest` command. <http://www.apple.com>.
- `unixODBC` Driver Manager for Unix (`libodbc.so`). See <http://www.unixodbc.org>, for more information. The `unixODBC` Driver Manager includes the Connector/ODBC driver 3.51 in the installation package, starting with version `unixODBC 2.1.2`.
- `iODBC` ODBC Driver Manager for Unix (`libiodbc.so`), see <http://www.iodbc.org>, for more information.

### 21.1.3. Connector/ODBC Installation

This section explains where to download Connector/ODBC, and how to run the installer, copy the files manually, or build from source.

#### Where to Get Connector/ODBC

You can get a copy of the latest version of Connector/ODBC binaries and sources from our Web site at <http://dev.mysql.com/downloads/connector/odbc/>.

For more information about Connector/ODBC, visit <http://www.mysql.com/products/myodbc/>.

For more information about licensing, visit <http://www.mysql.com/company/legal/licensing/>.

#### Choosing Unicode or ANSI Driver

Connector/ODBC offers the flexibility to handle data using any character set through its **Unicode-enabled** driver, or the maximum raw speed for a more limited range of character sets through its **ANSI** driver. Some users postponed their upgrade to Connector/ODBC 5.1, remaining with the older 3.51 version to keep this performance edge. As of Connector/ODBC 5.2, both kinds of drivers are available based on the 5.x code base: you can choose either a Unicode-enabled driver or an ANSI driver on the download page. The Unicode-enabled driver, recommended for most users, has no special qualifier in the download filename. The ANSI driver includes `-ansi-` in the download filename.



#### Note

You can install either the Unicode driver on a Windows system, or the ANSI driver, or both. The drivers are distinguished in the list of installed software and

in the names of libraries and directories by a `w` (for “wide characters”) for the Unicode driver, and `a` in the ANSI driver.

## Choosing Binary or Source Installation Method

You can install the Connector/ODBC drivers using two different methods:

- The **binary installation** is the easiest and most straightforward method of installation. You receive all the necessary libraries and other files pre-built, with an installer program or batch script to perform all necessary copying and configuration.
- The **source installation** method is intended for platforms where a binary installation package is not available, or in situations where you want to customize or modify the installation process or Connector/ODBC drivers before installation.

If a binary distribution is not available for a particular platform, and you build the driver from the original source code, you can contribute the binaries you create to MySQL by sending a mail message to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com), so that it becomes available for other users.

## Supported Platforms

Connector/ODBC can be used on all major platforms supported by MySQL. You can install it on:

- Windows 95, 98, Me, NT, 2000, XP, 2003, Vista and 7.
- All Unix-like Operating Systems, including: AIX, Amiga, BSDI, DEC, FreeBSD, HP-UX 10/11, Linux, NetBSD, OpenBSD, OS/2, SGI Irix, Solaris, SunOS, SCO OpenServer, SCO UnixWare, Tru64 Unix.
- Mac OS X and Mac OS X Server.



### Note

On all non-Windows platforms except Mac OS X, the driver is built against `unixODBC` and is expecting a 2-byte `SQLWCHAR`, not 4 bytes as `iODBC` is using. For this reason, the binaries are **only** compatible with `unixODBC`; recompile the driver against `iODBC` to use them together. For further information, see [Section 21.1.2.1.2, “ODBC Driver Managers”](#).

For further instructions, consult the documentation corresponding to the platform where you are installing and whether you are running a binary installer or building from source:

Platform	Binary Installer	Build from Source
Windows	<a href="#">Installation Instructions</a>	<a href="#">Build Instructions</a>
Unix/Linux	<a href="#">Installation Instructions</a>	<a href="#">Build Instructions</a>
Mac OS X	<a href="#">Installation Instructions</a>	

### 21.1.3.1. Installing Connector/ODBC on Windows

Before installing the Connector/ODBC drivers on Windows, ensure that your Microsoft Data Access Components (MDAC) are up to date. You can obtain the latest version from the [Microsoft Data Access and Storage](#) Web site.

There are different distribution types to use when installing for Windows. The software that is installed is identical in each case, only the installation method is different.

- The **zipped installer** consists of a zipped package containing a standalone installation application. To install from this package, unzip the installer, and then run the installation application.

The **MSI installer** is an installation file that can be used with the installer included in Windows 2000, Windows XP and Windows Server 2003.

See [Section 21.1.3.1.1, “Installing the Windows Connector/ODBC Driver Using an Installer”](#) for the steps to complete the installation with either of these installers.

- The **zipped DLL** package contains DLL files that must be manually installed. See [Section 21.1.3.1.2, “Installing the Windows Connector/ODBC Driver using the Zipped DLL Package”](#) for the steps to complete this type of installation.

**Note**

An OLEDB/ODBC driver for Windows 64-bit is available from [Microsoft Downloads](#).

### 21.1.3.1.1. Installing the Windows Connector/ODBC Driver Using an Installer

The zipped or MSI installer packages offer a very simple method for installing the Connector/ODBC drivers. If you have downloaded the zipped installer, extract the installer application. The basic installation process is identical for both installers.

Follow these steps to complete the installation:

1. Double-click the standalone installer that you extracted, or the MSI file you downloaded.
2. The MySQL Connector/ODBC Setup Wizard starts. Click the **Next** button to begin the installation process.
3. Choose the installation type. The **Typical** installation provides the standard files needed to connect to a MySQL database using ODBC. The **Complete** option installs all the available files, including debug and utility components. Oracle recommends choosing one of these two options to complete the installation. If you choose one of these methods, click **Next**, then proceed to step 5.

You can also choose a **Custom** installation, where you select the individual components to install. If you choose this method, click **Next**, then proceed to step 4.

4. If you have chosen a custom installation, use the pop-ups to select which components to install, then click **Next** to install the necessary files.
5. Once the files are copied to their final locations, the installation is complete. Click **Finish** to exit the installer.

Now that the installation is complete, configure your ODBC connections using [Section 21.1.4, “Configuring Connector/ODBC”](#).

### 21.1.3.1.2. Installing the Windows Connector/ODBC Driver using the Zipped DLL Package

If you have downloaded the zipped DLL package:

1. Unzip the installation files.
2. Run the included batch file to perform an installation to the default locations.
3. Alternatively, install the individual files required for Connector/ODBC operation manually.

**Note**

The following instructions only work for 32-bit Windows systems. If you have a 64-bit Windows system, use the MSI installer, which installs both the 32-bit and 64-bit drivers to the correct locations.

To install using the **batch file**:

1. Unzip the Connector/ODBC zipped DLL package.
2. Open a command prompt.
3. Change to the directory created when you unzipped the Connector/ODBC zipped DLL package.
4. Run `Install.bat`:

```
C:\> Install.bat
```

This copies the necessary files into the default location, and then registers the Connector/ODBC driver with the Windows ODBC manager.



#### Note

Changing or adding a new DSN (data source name) may be accomplished using either the GUI, or from the command-line using `myodbc-installer.exe` with Connector/ODBC 5.x, or `myodbc3i.exe` with Connector/ODBC 3.51,

Although Oracle recommends installing these files in the standard location, you can also copy the files by hand to an alternative location - for example, to run or test different versions of the Connector/ODBC driver on the same machine. To **copy the files** to a location of your choice, use the following steps:

1. Unzip the Connector/ODBC zipped DLL package.
2. Open a command prompt.
3. Change to the directory created when you unzipped the Connector/ODBC zipped DLL package.
4. Copy the library files to a suitable directory. The default location is the default Windows system directory `\Windows\System32`:

```
C:\> copy lib\myodbc5S.dll \Windows\System32
C:\> copy lib\myodbc5S.lib \Windows\System32

If installing the Unicode-enabled driver:
C:\> copy lib\myodbc5w.dll \Windows\System32
C:\> copy lib\myodbc5w.lib \Windows\System32

If installing the ANSI driver:
C:\> copy lib\myodbc5a.dll \Windows\System32
C:\> copy lib\myodbc5a.lib \Windows\System32
```

5. Copy the Connector/ODBC tools. These must be placed in a directory that is in the system `%PATH%`. The default is to install these into the Windows system directory `\Windows\System32`:

```
C:\> copy bin\myodbc3i.exe \Windows\System32
C:\> copy bin\myodbc3m.exe \Windows\System32
C:\> copy bin\myodbc3c.exe \Windows\System32
```

6. Optionally, copy the help files. For these files to be accessible through the help system, they must be installed in the Windows system directory:

```
C:\> copy doc\*.hlp \Windows\System32
```

7. Finally, register the Connector/ODBC driver with the ODBC manager:

```
For Unicode-enabled driver:
C:\> myodbc3i -a -d -t"MySQL ODBC 5.2 Driver;\
    DRIVER=myodbc5w.dll;SETUP=myodbc5S.dll"

For ANSI driver:
C:\> myodbc3i -a -d -t"MySQL ODBC 5.2 Driver;\
```

```
DRIVER=myodbc5a.dll;SETUP=myodbc5S.dll"
```

If you installed these files into a non-default location, change the references to the DLL files and command location in the above statement

### 21.1.3.2. Installing Connector/ODBC on Unix

There are two methods available for installing Connector/ODBC on Unix from a binary distribution. For most Unix environments, you will use the **tarball** distribution. For Linux systems, there is also an **RPM** distribution available.



#### Note

To install Connector/ODBC 5.x on Unix requires unixODBC 2.2.12 or later to be installed.

#### 21.1.3.2.1. Installing Connector/ODBC from a Binary Tarball Distribution

To install the driver from a tarball distribution (`.tar.gz` file), download the latest version of the driver for your operating system and follow these steps that demonstrate the process using the Linux version of the tarball, substituting the appropriate file and directory names based on the package you download:

```
shell> su root
shell> gunzip mysql-connector-odbc-5.2.2-i686-pc-linux.tar.gz
shell> tar xvf mysql-connector-odbc-5.2.2-i686-pc-linux.tar
shell> cd mysql-connector-odbc-5.2.2-i686-pc-linux
```

Read the installation instructions in the `INSTALL` file and execute these commands.

Then see [Section 21.1.4.5, "Configuring a Connector/ODBC DSN on Unix"](#) to configure the DSN for Connector/ODBC. For more information, refer to the `INSTALL` file that comes with your distribution.

#### 21.1.3.2.2. Installing Connector/ODBC from an RPM Distribution

To install or upgrade Connector/ODBC from an RPM distribution on Linux, simply download the RPM distribution of the latest version of Connector/ODBC and follow the instructions below. Use `su root` to become `root`, then install the RPM file.

If you are installing for the first time:

```
shell> su root
shell> rpm -ivh mysql-connector-odbc-5.2.2.i386.rpm
```

If the driver exists, upgrade it like this:

```
shell> su root
shell> rpm -Uvh mysql-connector-odbc-5.2.2.i386.rpm
```

If there is any dependency error for MySQL client library, `libmysqlclient`, simply ignore it by supplying the `--nodeps` option, and then make sure the MySQL client shared library is in the path or set through `LD_LIBRARY_PATH`.

This installs the driver libraries and related documents to `/usr/local/lib` and `/usr/share/doc/MyODBC`, respectively. See [Section 21.1.4.5, "Configuring a Connector/ODBC DSN on Unix"](#) for the post-installation configuration steps.

To **uninstall** the driver, become `root` and execute an `rpm` command:

```
shell> su root
shell> rpm -e mysql-connector-odbc
```



### 21.1.3.3. Installing Connector/ODBC on Mac OS X

Mac OS X is based on the FreeBSD operating system, and you can normally use the MySQL network port for connecting to MySQL servers on other hosts. Installing the Connector/ODBC driver lets you connect to MySQL databases on any platform through the ODBC interface. If your application requires an ODBC interface, install the Connector/ODBC driver. Applications that require or can use ODBC (and therefore the Connector/ODBC driver) include ColdFusion, Filemaker Pro, 4th Dimension and many other applications.

Mac OS X includes its own ODBC manager, based on the [iODBC](#) manager. Mac OS X includes an administration tool that provides easier administration of ODBC drivers and configuration, updating the underlying [iODBC](#) configuration files.

For Connector/ODBC 5.x, the package is provided as a compressed tar archive that you manually install.

In either case, the driver is designed to work with the iODBC driver manager included with Mac OS X.

To install Connector/ODBC 5.x:

1. Download the installation file for the appropriate level of Mac OS X (10.5 or higher).
2. Extract the archive:

```
shell> tar xvzf mysql-connector-odbc-5.x.y-osx10.z-x86-(32/64)bit.tar.gz
```

3. The directory created contains the subdirectories, [lib](#), [bin](#), and [test](#). Copy the contents of [lib](#) and [bin](#) to a suitable location such as [/usr/local](#):

```
shell> sudo cp bin/* /usr/local/bin
shell> sudo cp lib/* /usr/local/lib
```

4. Finally, register the driver with iODBC using the [myodbc3i](#) tool you just installed:

```
shell> myodbc3i -a -d -t"MySQL ODBC 3.51 Driver;Driver=/usr/local/lib/libmyodbc3.so;Setup=/usr/local/
```

To verify the installed drivers, either use the ODBC Administrator application or the [myodbc3i](#) utility:

```
shell> myodbc3i -q -d
```

### 21.1.3.4. Building Connector/ODBC from a Source Distribution on Windows

You only need to build Connector/ODBC from source on Windows to modify the source or installation location. If you are unsure whether to install from source, please use the binary installation detailed in [Section 21.1.3.1, "Installing Connector/ODBC on Windows"](#).

Building Connector/ODBC from source on Windows requires a number of different tools and packages:

- MDAC, Microsoft Data Access SDK from <http://support.microsoft.com/kb/110093>.
- Suitable C compiler, such as Microsoft Visual C++ or the C compiler included with Microsoft Visual Studio.

Microsoft Visual Studio 7 and 8 are preferred, and well-tested.

- Connector/ODBC 5.2: [cmake](#).
- MySQL client libraries and include files from MySQL 4.0.0 or higher. (Preferably MySQL 4.0.16 or higher). This is required because Connector/ODBC uses calls and structures that exist only starting from this version of the library. To get the client libraries and include files, visit <http://dev.mysql.com/downloads/>.

## Build Steps

Set the environment variables for the Visual Studio toolchain. Visual Studio includes a batch file to set these for you, and installs a **Start** menu shortcut that opens a command prompt with these variables set.

Set `MYSQL_DIR` to the MySQL server installation path, while using the short-style file names. For example:

```
C:\> set MYSQL_DIR=C:\PROGRA~1\MySQL\MYSQLS~1.0
```

Build Connector/ODBC using the `cmake` command-line tool by executing the following from the source root directory (in a command prompt window):

```
C:\> cmake -G "Visual Studio 8 2005"
```

This produces a project file that you can open with Visual Studio, or build from the command line with either of the following commands:

```
C:\> devenv.com MySQL_Connector_ODBC.sln /build Release
C:\> devenv.com MySQL_Connector_ODBC.sln /build RelWithDebInfo
```

To compile a debug build, set the `cmake` build type so that the correct versions of the MySQL client libraries are used:

```
C:\> cmake -G "Visual Studio 8 2005" -DCMAKE_BUILD_TYPE=Debug
C:\> devenv.com MySQL_Connector_ODBC.sln /build Debug
```

Upon completion, the executables are in the `bin/` and `lib/` subdirectories.

See [Section 21.1.3.1.2, "Installing the Windows Connector/ODBC Driver using the Zipped DLL Package"](#) for the `copy` commands to complete the installation.

### 21.1.3.5. Building Connector/ODBC from a Source Distribution on Unix



#### Note

Connector/5.2 uses `cmake` for build, rather than `autoconf` and its `configure` command.

You need the following tools to build MySQL from source on Unix:

- A working ANSI C++ compiler. GCC 4.2.1 or later, Sun Studio 10 or later, Visual Studio 2008 or later, and many current vendor-supplied compilers are known to work.
- `cmake`.
- MySQL client libraries and include files from MySQL 4.0.0 or higher. (Preferably MySQL 4.0.16 or higher). This is required because Connector/ODBC uses calls and structures that exist only starting from this version of the library. To get the client libraries and include files, visit <http://dev.mysql.com/downloads/>.

If you have built your own MySQL server or client libraries from source using the GNU autotools, use the `--enable-thread-safe-client` option to `configure` when the libraries were built. No special option is needed if you configure with `cmake`.

Also, ensure that the `libmysqlclient` library was built and installed as a shared library.

- A compatible ODBC manager must be installed. Connector/ODBC is known to work with the `iODBC` and `unixODBC` managers. See [Section 21.1.2.1.2, "ODBC Driver Managers"](#) for more information.
- If you are using a character set that is not compiled into the MySQL client library, install the MySQL character definitions from the `charsets` directory into `SHAREDIR` (by default, `/usr/local/`

`mysql/share/mysql/charsets`). These should be in place if you have installed the MySQL server on the same machine. See [Section 10.1, “Character Set Support”](#) for more information on character set support.

Once you have all the required files, unpack the source files to a separate directory, then run `configure` and build the library using `make`.

## Typical `cmake` Options

`iODBC` is the default ODBC library used by Connector/ODBC. Alternatively, `unixODBC` may be used by passing in the appropriate option to `cmake`. For example:

```
shell> cmake -G "Unix Makefiles" -DWITH_UNIXODBC=1
```

When you run `cmake`, you might add options to the command line. Here are some examples:

- `-DODBC_INCLUDES=dir_name [133]`: Used when the ODBC include directory is not found within the system `$PATH`.
- `-DODBC_LIB_DIR=dir_name [133]`: Used when the ODBC library directory is not found within the system `$PATH`.
- `-DWITH_UNIXODBC=1 [137]`: Enables `unixODBC` support.

## Additional `configure` Options

There are a number of other options that you need, or want, to set when configuring the Connector/ODBC driver before it is built.

- To link the driver with MySQL thread safe client libraries `libmysqlclient_r.so` or `libmysqlclient_r.a`, specify the following `configure` option:

```
--enable-thread-safe
```

and can be disabled (default) using

```
--disable-thread-safe
```

This option enables the building of the driver thread-safe library `libmyodbc3_r.so` from by linking with MySQL thread-safe client library `libmysqlclient_r.so`. (The extensions are OS-dependent.)

If the compilation with the thread-safe option fails, it may be because the correct thread-libraries on the system could not be located. Set the value of `LIBS` to point to the correct thread library for your system.

```
LIBS="-lpthread" ./configure ..
```

- You can enable or disable the shared and static versions of Connector/ODBC using these options:

```
--enable-shared[=yes|no]
--disable-shared
--enable-static[=yes|no]
--disable-static
```

- By default, all the binary distributions are built as non-debug versions (configured with `--without-debug`).

To enable debugging information, build the driver from a source distribution with the proper configuration option to enable debugging support. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

- This option is available only for source trees that have been obtained from the Subversion repository. This option does not apply to the packaged source distributions.

By default, the driver is built with the `--without-docs` option. If you would like the documentation to be built, then execute `configure` with:

```
--with-docs
```

## Build Steps for Unix

To build the driver libraries, execute `make`:

```
shell> make
```

If any errors occur, correct them and continue the build process. If you are not able to build, then send a detailed email to [myodbc@lists.mysql.com](mailto:myodbc@lists.mysql.com) for further assistance.

## Building Shared Libraries

On most platforms, MySQL does not build or support `.so` (shared) client libraries by default. This is based on our experience of problems when building shared libraries.

In cases like this, you have to download the MySQL distribution and configure it with these options:

```
--without-server --enable-shared
```

To build shared driver libraries, specify the `--enable-shared` option for `configure`. By default, `configure` does not enable this option.

If you have configured with the `--disable-shared` option, you can build the `.so` file from the static libraries using the following commands:

```
shell> cd mysql-connector-odbc-3.51.01
shell> make
shell> cd driver
shell> CC=/usr/bin/gcc \
    $CC -bundle -flat_namespace -undefined error \
    -o .libs/libmyodbc3-3.51.01.so \
    catalog.o connect.o cursor.o dll.o error.o execute.o \
    handle.o info.o misc.o myodbc3.o options.o prepare.o \
    results.o transact.o utility.o \
    -L/usr/local/mysql/lib/mysql/ \
    -L/usr/local/iodbc/lib/ \
    -lz -lc -lmysqlclient -liodbcinst
```

Make sure to change `-liodbcinst` to `-lodbcinst` if you are using `unixODBC` instead of `iODBC`, and configure the library paths accordingly.

This builds and places the `libmyodbc3-3.51.01.so` file in the `.libs` directory. Copy this file to the Connector/ODBC library installation directory (`/usr/local/lib`, or the `lib` directory under the installation directory that you supplied with the `--prefix` option).

```
shell> cd .libs
shell> cp libmyodbc3-3.51.01.so /usr/local/lib
shell> cd /usr/local/lib
shell> ln -s libmyodbc3-3.51.01.so libmyodbc3.so
```

To build the thread-safe driver library:

```
shell> CC=/usr/bin/gcc \
    $CC -bundle -flat_namespace -undefined error
    -o .libs/libmyodbc3_r-3.51.01.so
    catalog.o connect.o cursor.o dll.o error.o execute.o
    handle.o info.o misc.o myodbc3.o options.o prepare.o
    results.o transact.o utility.o
    -L/usr/local/mysql/lib/mysql/
    -L/usr/local/iodbc/lib/
    -lz -lc -lmysqlclient_r -liodbcinst
```

## Installing Driver Libraries

To install the driver libraries, execute the following command:

```
shell> make install
```

That command installs one of the following sets of libraries:

For Connector/ODBC 3.51:

- `libmyodbc3.so`
- `libmyodbc3-3.51.01.so`, where 3.51.01 is the version of the driver
- `libmyodbc3.a`

For thread-safe Connector/ODBC 3.51:

- `libmyodbc3_r.so`
- `libmyodbc3-3_r.51.01.so`
- `libmyodbc3_r.a`

For more information on build process, refer to the `INSTALL` file that comes with the source distribution. Note that if you are trying to use the `make` from Sun, you may end up with errors. On the other hand, GNU `gmake` should work fine on all platforms.

## Testing Connector/ODBC on Unix

To run the basic samples provided in the distribution with the libraries that you built, use the following command:

```
shell> make test
```

Before running the tests, create the DSN 'myodbc3' in `odbc.ini` and set the environment variable `ODBCINI` to the correct `odbc.ini` file; and MySQL server is running. You can find a sample `odbc.ini` with the driver distribution.

You can even modify the `samples/run-samples` script to pass the desired DSN, UID, and PASSWORD values as the command-line arguments to each sample.

### 21.1.3.6. Building Connector/ODBC from a Source Distribution on Mac OS X

To build the driver on Mac OS X (Darwin), make use of the following `configure` example:

```
shell> ./configure --prefix=/usr/local
--with-unixODBC=/usr/local
--with-mysql-path=/usr/local/mysql
--disable-shared
--enable-gui=no
--host=powerpc-apple
```

The command assumes that the `unixODBC` and MySQL are installed in the default locations. If not, configure accordingly.

On Mac OS X, `--enable-shared` builds `.dylib` files by default. You can build `.so` files like this:

```
shell> make
shell> cd driver
shell> CC=/usr/bin/gcc \
  $CC -bundle -flat_namespace -undefined error
-o .libs/libmyodbc3-3.51.01.so *.o
-L/usr/local/mysql/lib/
-L/usr/local/iodbc/lib
-liodbcinst -lmysqlclient -lz -lc
```

To build the thread-safe driver library:

```
shell> CC=/usr/bin/gcc \
      $CC -bundle -flat_namespace -undefined error
      -o .libs/libmyodbc3-3.51.01.so *.o
      -L/usr/local/mysql/lib/
      -L/usr/local/iodbc/lib
      -liodbcinst -lmysqlclient_r -lz -lc -lpthread
```

Make sure to change the `-liodbcinst` to `-lodbcinst` in case of using `unixODBC` instead of `iODBC` and configure the libraries path accordingly.

In Apple's version of GCC, both `cc` and `gcc` are actually symbolic links to `gcc3`.

Copy this library to the `$prefix/lib` directory and symlink to `libmyodbc3.so`.

You can cross-check the output shared-library properties using this command:

```
shell> otool -LD .libs/libmyodbc3-3.51.01.so
```

### 21.1.3.7. Installing Connector/ODBC from the Development Source Tree



#### Caution

Read this section only if you are interested in helping us test our new code. To just get MySQL Connector/ODBC up and running on your system, use a standard release distribution.

To obtain the most recent development source tree, first download and install Bazaar from the [Bazaar VCS Web site](#). Bazaar is supported by any platform that supports Python, and is therefore compatible with any Linux, Unix, Windows or Mac OS X host. Instructions for downloading and installing Bazaar on the different platforms are available on the Bazaar Web site.

Building from the source trees requires the following tools:

- `autoconf` 2.52 (or newer).
- `automake` 1.4 (or newer).
- `libtool` 1.4 (or newer).
- `m4`.

The most recent development source tree is available from our public Subversion trees at <http://dev.mysql.com/tech-resources/sources.html>.

To check out the Connector/ODBC sources, change to the directory where you want the copy of the Connector/ODBC tree to be stored, then use the following command:

```
shell> bzr branch lp:mysqlbc
```

You should now have a copy of the entire Connector/ODBC source tree in the directory `connector-odbc3`. To build from this source tree on Unix or Linux, follow these steps:

```
shell> cd myodbc
shell> aclocal
shell> autoheader
shell> libtoolize -c -f
shell> autoconf
shell> automake;
shell> ./configure # Add your favorite options here
shell> make
```

When the build is done, run `make install` to install the Connector/ODBC driver on your system.

If you have gotten to the `make` stage and the distribution does not compile, please report it to [myodbc@lists.mysql.com](mailto:myodbc@lists.mysql.com).

On Windows, make use of Windows Makefiles [WIN-Makefile](#) and [WIN-Makefile\\_debug](#) in building the driver. For more information, see [Section 21.1.3.4, “Building Connector/ODBC from a Source Distribution on Windows”](#).

After the initial checkout operation to get the source tree, run `bzr pull` periodically to update your source according to the latest version.

## 21.1.4. Configuring Connector/ODBC

Before you connect to a MySQL database using the Connector/ODBC driver, you configure an ODBC Data Source Name (DSN). The DSN associates the various configuration parameters required to communicate with a database to a specific name. You use the DSN in an application to communicate with the database, rather than specifying individual parameters within the application itself. DSN information can be user-specific, system-specific, or provided in a special file. ODBC data source names are configured in different ways, depending on your platform and ODBC driver.

### 21.1.4.1. Overview of Connector/ODBC Data Source Names

A Data Source Name associates the configuration parameters for communicating with a specific database. Generally, a DSN consists of the following parameters:

- Name
- Host Name
- Database Name
- Login
- Password

In addition, different ODBC drivers, including Connector/ODBC, may accept additional driver-specific options and parameters.

There are three types of DSN:

- A *System DSN* is a global DSN definition that is available to any user and application on a particular system. A System DSN can normally only be configured by a systems administrator, or by a user who has specific permissions that let them create System DSNs.
- A *User DSN* is specific to an individual user, and can be used to store database connectivity information that the user regularly uses.
- A *File DSN* uses a simple file to define the DSN configuration. File DSNs can be shared between users and machines and are therefore more practical when installing or deploying DSN information as part of an application across many machines.

DSN information is stored in different locations depending on your platform and environment.

### 21.1.4.2. Connector/ODBC Connection Parameters

You can specify the parameters in the following tables for Connector/ODBC when configuring a DSN:

- [Table 21.3, “Connector/ODBC DSN Configuration Options”](#)
- [Table 21.4, “Connector/ODBC Option Flags”](#)

Users on Windows can use the Options and Advanced panels when configuring a DSN to set these parameters; see the table for information on which options relate to which fields and check boxes. On Unix and Mac OS X, use the parameter name and value as the keyword/value pair in the DSN configuration. Alternatively, you can set these parameters within the `InConnectionString` argument in the `SQLDriverConnect()` call.



**Table 21.3. Connector/ODBC DSN Configuration Options**

Parameter	Default Value	Comment
<code>user</code>	ODBC	The user name used to connect to MySQL.
<code>uid</code>	ODBC	Synonymous with <code>user</code> . Added in 3.51.16.
<code>server</code>	<code>localhost</code>	The host name of the MySQL server.
<code>database</code>		The default database.
<code>option</code>	0	Options that specify how Connector/ODBC works. See <a href="#">Table 21.4, “Connector/ODBC Option Flags”</a> and <a href="#">Table 21.5, “Recommended Connector/ODBC option Values for Different Configurations”</a> .
<code>port</code>	3306	The TCP/IP port to use if <code>server</code> is not <code>localhost</code> .
<code>initstmt</code>		Initial statement. A statement to execute when connecting to MySQL. In version 3.51 the parameter is called <code>stmt</code> . Note, the driver supports the initial statement being executed only at the time of the initial connection.
<code>password</code>		The password for the <code>user</code> account on <code>server</code> .
<code>pwd</code>		Synonymous with <code>password</code> . Added in 3.51.16.
<code>socket</code> [542]		The Unix socket file or Windows named pipe to connect to if <code>server</code> is <code>localhost</code> .
<code>sslca</code>		The path to a file with a list of trust SSL CAs. Added in 3.51.16.
<code>sslcapath</code>		The path to a directory that contains trusted SSL CA certificates in PEM format. Added in 3.51.16.
<code>sslcert</code>		The name of the SSL certificate file to use for establishing a secure connection. Added in 3.51.16.
<code>sslcipher</code>		A list of permissible ciphers to use for SSL encryption. The cipher list has the same format as the <code>openssl ciphers</code> command. Added in 3.51.16.
<code>sslkey</code>		The name of the SSL key file to use for establishing a secure connection. Added in 3.51.16.
<code>charset</code>		The character set to use for the connection. Added in 3.51.17.
<code>sslverify</code>		If set to 1, the SSL certificate will be verified when used with the MySQL connection. If not set, then the default behavior is to ignore SSL certificate verification.
<code>readtimeout</code>		The timeout in seconds for attempts to read from the server. Each attempt uses this timeout value and there are retries if necessary, so the total effective timeout value is three times the option value. You can set the value so that a lost connection can be detected earlier than the TCP/IP <code>Close_Wait_Timeout</code> value of 10 minutes. This option works only for TCP/IP connections, and only for Windows prior to MySQL 5.1.12. Corresponds to the <code>MYSQL_OPT_READ_TIMEOUT</code> option of the MySQL Client Library. Added in 3.51.27.
<code>writetimeout</code>		The timeout in seconds for attempts to write to the server. Each attempt uses this timeout value and there are <code>net_retry_count</code> retries if necessary, so the total effective timeout value is <code>net_retry_count</code> times the option value. This option works only for TCP/IP connections, and only for Windows prior to MySQL 5.1.12. Corresponds to the <code>MYSQL_OPT_WRITE_TIMEOUT</code> option of the MySQL Client Library. Added in 3.51.27.

Parameter	Default Value	Comment
<code>interactive</code>		Enables the <code>CLIENT_INTERACTIVE</code> connection option of <code>mysql_real_connect</code> .
<code>prefetch</code>	0	When set to a non-zero value <i>N</i> , causes all queries in the connection to return <i>N</i> rows at a time rather than the entire result set. Useful for queries against very large tables where it is not practical to retrieve the whole result set at once. You can scroll through the result set, <i>N</i> records at a time.  This option works only with forward-only cursors. It does not work when the <code>FLAG_MULTI_STATEMENTS</code> flag is set. It can be used in combination with the <code>FLAG_NO_CACHE</code> option. Its behavior in ADO applications is undefined: the prefetching might or might not occur.
<code>no_ssps</code>	0	In Connector/ODBC 5.2, by default, server-side prepared statements are used. When this option is set to a non-zero value, prepared statements are emulated on the client side, which is the same behavior as in 5.1 and 3.51. Added in 5.2.
<code>can_handle_expired_pwd</code>		Indicates that the application can deal with an expired password, which is signalled by an SQL state of <code>08004</code> ("Server rejected the connection") and a native error code <code>ER_MUST_CHANGE_PASSWORD_LOGIN</code> (1862). The connection is "sandboxed", and can do nothing other than issue a <code>SET PASSWORD</code> statement. To establish a connection in this case, your application must either use the <code>initstmt</code> connection option to set a new password at the start, or issue a <code>SET PASSWORD</code> statement immediately after connecting. Once the expired password is reset, the restrictions on the connection are lifted. See <a href="#">Section 13.7.1.1, "ALTER USER Syntax"</a> for details about password expiration for MySQL server accounts. Added in 5.2.4.



**Note**

The SSL configuration parameters can also be automatically loaded from a `my.ini` or `my.cnf` file. See [Section 4.2.3.3, "Using Option Files"](#).

The `option` argument is used to tell Connector/ODBC that the client is not 100% ODBC-compliant. On Windows, you normally select options by toggling the check boxes in the connection screen, but you can also select them in the `option` argument. The following options are listed in the order in which they appear in the Connector/ODBC connect screen.

**Table 21.4. Connector/ODBC Option Flags**

Flag Name	GUI Option	Constant Value	Description
<code>FLAG_FIELD_LENGTH</code>	Do not Optimize Column Width		The client cannot handle that Connector/ODBC returns the real width of a column. This option was removed in 3.51.18.
<code>FLAG_FOUND_ROWS</code>	Return Matching Rows	2	The client cannot handle that MySQL returns the true value of affected rows. If this flag is set, MySQL returns "found rows" instead. You must have MySQL 3.21.14 or newer for this to work.
<code>FLAG_DEBUG</code>	Trace Driver Calls To myodbc.log		Make a debug log in <code>C:\myodbc.log</code> on Windows,

## Configuring Connector/ODBC

Flag Name	GUI Option	Constant Value	Description
			or <code>/tmp/myodbc.log</code> on Unix variants. This option was removed in Connector/ODBC 3.51.18.
<code>FLAG_BIG_PACKETS</code>	Allow Big Results	8	Do not set any packet limit for results and bind parameters. Without this option, parameter binding will be truncated to 255 characters.
<code>FLAG_NO_PROMPT</code>	Do not Prompt Upon Connect	16	Do not prompt for questions even if driver would like to prompt.
<code>FLAG_DYNAMIC_CURSOR</code>	Enable Dynamic Cursor	32	Enable or disable the dynamic cursor support.
<code>FLAG_NO_SCHEMA</code>	Ignore # in Table Name	64	Ignore use of database name in <code>db_name.tbl_name.col_name</code> .
<code>FLAG_NO_DEFAULT_CURSOR</code>	User Manager Cursors	128	Force use of ODBC manager cursors (experimental).
<code>FLAG_NO_LOCALE</code>	Do not Use Set Locale	256	Disable the use of extended fetch (experimental).
<code>FLAG_PAD_SPACE</code>	Pad Char To Full Length	512	Pad <code>CHAR</code> columns to full column length.
<code>FLAG_FULL_COLUMN_NAMES</code>	Return Table Names for SQLDescribeCol	1024	<code>SQLDescribeCol()</code> returns fully qualified column names.
<code>FLAG_COMPRESSED_PROTO</code>	Use Compressed Protocol	2048	Use the compressed client/server protocol.
<code>FLAG_IGNORE_SPACE</code>	Ignore Space After Function Names	4096	Tell server to ignore space after function name and before "(" (needed by PowerBuilder). This makes all function names keywords.
<code>FLAG_NAMED_PIPE</code>	Force Use of Named Pipes	8192	Connect with named pipes to a <code>mysqld</code> server running on NT.
<code>FLAG_NO_BIGINT</code>	Change BIGINT Columns to Int	16384	Change <code>BIGINT</code> columns to <code>INT</code> columns (some applications cannot handle <code>BIGINT</code> ).
<code>FLAG_NO_CATALOG</code>	No Catalog	32768	Forces results from the catalog functions, such as <code>SQLTables</code> , to always return <code>NULL</code> and the driver to report that catalogs are not supported.
<code>FLAG_USE_MYCNF</code>	Read Options From <code>my.cnf</code>	65536	Read parameters from the <code>[client]</code> and <code>[odbc]</code> groups from <code>my.cnf</code> .
<code>FLAG_SAFE</code>	Safe	131072	Add some extra safety checks.
<code>FLAG_NO_TRANSACTIONS</code>	Disable transactions	262144	Disable transactions.
<code>FLAG_LOG_QUERY</code>	Save queries to <code>myodbc.sql</code>	524288	Enable query logging to <code>c:\myodbc.sql(/tmp/</code>

Flag Name	GUI Option	Constant Value	Description
			<code>myodbc.sql</code> ) file. (Enabled only in debug mode.)
<code>FLAG_NO_CACHE</code>	Do not Cache Result (forward only cursors)	1048576	Do not cache the results locally in the driver, instead read from server ( <code>mysql_use_result()</code> ). This works only for forward-only cursors. This option is very important in dealing with large tables when you do not want the driver to cache the entire result set.
<code>FLAG_FORWARD_CURSOR</code>	Force Use Of Forward Only Cursors	2097152	Force the use of <code>Forward-only</code> cursor type. In case of applications setting the default static/dynamic cursor type, and one wants the driver to use noncache result sets, then this option ensures the forward-only cursor behavior.
<code>FLAG_AUTO_RECONNECT</code>	Enable auto-reconnect.	4194304	Enables auto-reconnection functionality. Do not use this option with <code>transactions</code> , since an auto-reconnection during a incomplete transaction may cause corruption. Note that an auto-reconnected connection will not inherit the same settings and environment as the original. Added in 3.51.13.
<code>FLAG_AUTO_IS_NULL</code>	Flag Auto Is Null	8388608	<p>When <code>FLAG_AUTO_IS_NULL</code> is set, the driver does not change the default value of <code>sql_auto_is_null</code> [543], leaving it at 1, so you get the MySQL default, not the SQL standard behavior.</p> <p>When <code>FLAG_AUTO_IS_NULL</code> is not set, the driver changes the default value of <code>SQL_AUTO_IS_NULL</code> to 0 after connecting, so you get the SQL standard, not the MySQL default behavior.</p> <p>Thus, omitting the flag disables the compatibility option and forces SQL standard behavior.</p> <p>See <code>IS NULL</code> [1168]. Added in 3.51.13.</p>

Flag Name	GUI Option	Constant Value	Description
<a href="#">FLAG_ZERO_DATE_TO_MIN</a>	Return SQL_NULL_DATA for zero date	16777216	Translates zero dates ( <a href="#">XXXX-00-00</a> ) into the minimum date values supported by ODBC, <a href="#">XXXX-01-01</a> . This resolves an issue where some statements will not work because the date returned and the minimum ODBC date value are incompatible. Added in 3.51.17.
<a href="#">FLAG_MIN_DATE_TO_ZERO</a>	Bind minimal date as zero date	33554432	Translates the minimum ODBC date value ( <a href="#">XXXX-01-01</a> ) to the zero date format supported by MySQL ( <a href="#">XXXX-00-00</a> ). This resolves an issue where some statements will not work because the date returned and the minimum ODBC date value are incompatible. Added in 3.51.17.
<a href="#">FLAG_MULTI_STATEMENTS</a>	Allow multiple statements	67108864	Enables support for batched statements. Added in 3.51.18.
<a href="#">FLAG_COLUMN_SIZE_S32</a>	Limit column size to 32-bit value	134217728	Limits the column size to a signed 32-bit value to prevent problems with larger column sizes in applications that do not support them. This option is automatically enabled when working with ADO applications. Added in 3.51.22.
<a href="#">FLAG_NO_BINARY_RESULT</a>	Always handle binary function results as character data	268435456	When set this option disables charset 63 for columns with an empty <a href="#">org_table</a> . Added in 3.51.26.
<a href="#">FLAG_DFLT_BIGINT_BIND_STR</a>		536870912	Causes <a href="#">BIGINT</a> parameters to be bound as strings. Microsoft Access treats <a href="#">BIGINT</a> as a string on linked tables. The value is read correctly, but bound as a string. This option is used automatically if the driver is used by Microsoft Access.
<a href="#">FLAG_NO_INFORMATION_SCHEMA</a>		1073741824	Tells catalog functions not to use <a href="#">INFORMATION_SCHEMA</a> , but rather use legacy algorithms. The trade-off here is usually speed for information quality. Using <a href="#">INFORMATION_SCHEMA</a> is often slow, but the information obtained is more complete.

To select multiple options, add together their values, using the numbers from the [Constant Value](#) column in the table.

**Note**

In MySQL Connector/ODBC 5.2, you can use the flag name directly as a parameter in the connection string, by specifying the flag name without the FLAG\_ prefix. So, in addition to using the `options` parameter with various flags set, you can use the flags directly as parameters. For example, `FIELD_LENGTH`, `FOUND_ROWS` and `DEBUG` could all be used as parameters.

The following table shows some recommended `option` values for various configurations:

**Table 21.5. Recommended Connector/ODBC option Values for Different Configurations**

Configuration	Option Value
Microsoft Access, Visual Basic	3
Driver trace generation (Debug mode)	4
Microsoft Access (with improved DELETE queries)	35
Large tables with too many rows	2049
Sybase PowerBuilder	135168
Query log generation (Debug mode)	524288
Generate driver trace as well as query log (Debug mode)	524292
Large tables with no-cache results	3145731

### 21.1.4.3. Configuring a Connector/ODBC DSN on Windows

The `ODBC Data Source Administrator` within Windows lets you create DSNs, check driver installation and configure ODBC systems such as tracing (used for debugging) and connection pooling.

Different editions and versions of Windows store the `ODBC Data Source Administrator` in different locations depending on the version of Windows that you are using.

To open the `ODBC Data Source Administrator` in Windows Server 2003:

**Tip**

To identify whether a DSN was created using the 32-bit or the 64-bit driver, include the driver being used within the DSN identifier. This will help you to identify the right DSN to use with applications such as Excel that are only compatible with the 32-bit driver. For example, you might add `Using32bitCODBC` to the DSN identifier for the 32-bit interface and `Using64bitCODBC` for those using the 64-bit Connector/ODBC driver.

1. On the `Start` menu, choose `Administrative Tools`, and then click `Data Sources (ODBC)`.

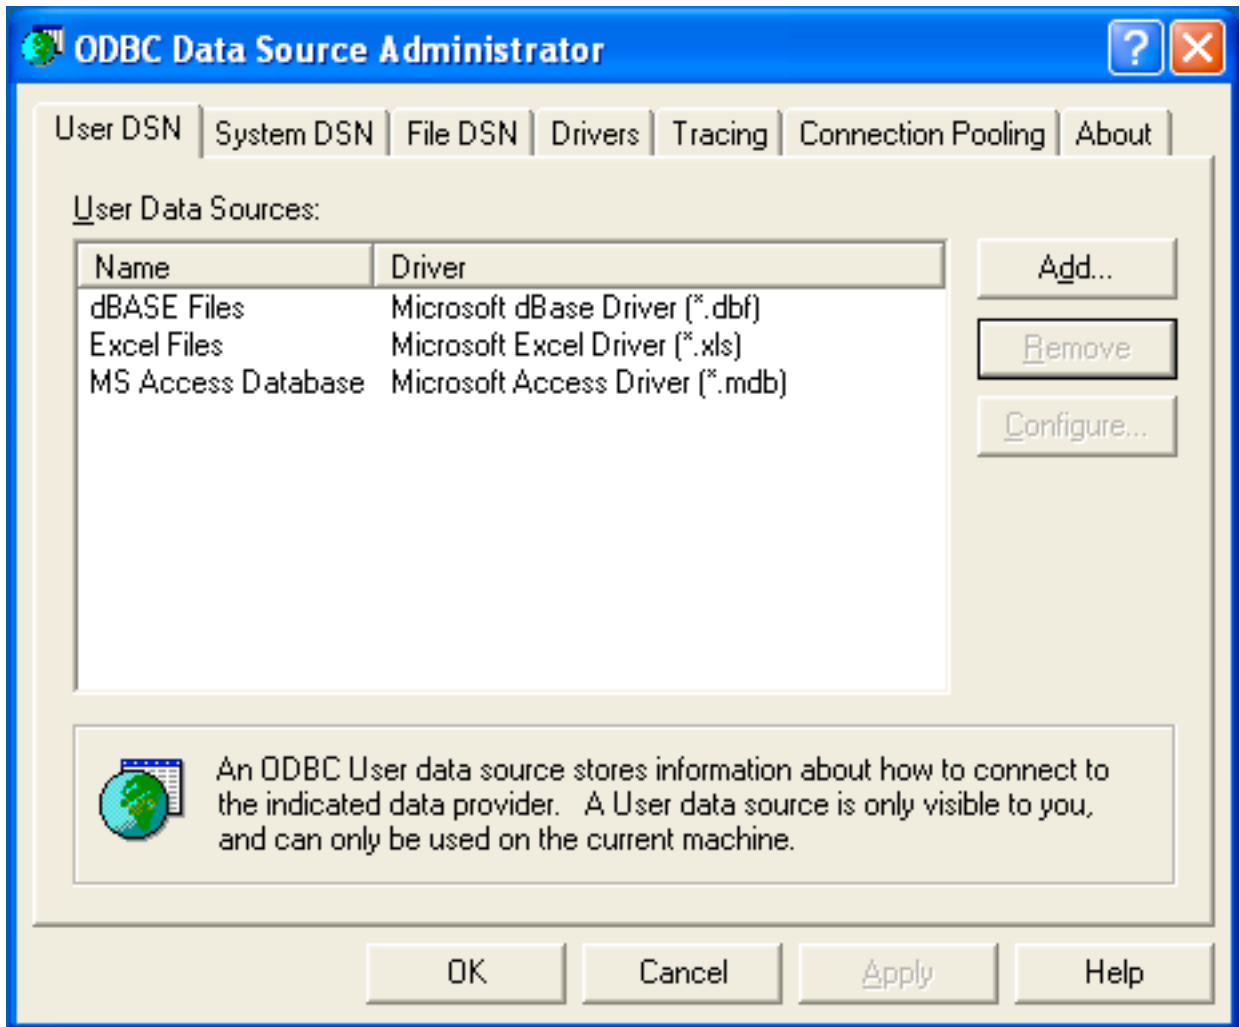
To open the `ODBC Data Source Administrator` in Windows 2000 Server or Windows 2000 Professional:

1. On the `Start` menu, choose `Settings`, and then click `Control Panel`.
2. In `Control Panel`, click `Administrative Tools`.
3. In `Administrative Tools`, click `Data Sources (ODBC)`.

To open the `ODBC Data Source Administrator` on Windows XP:

1. On the `Start` menu, click `Control Panel`.
2. In the `Control Panel` when in `Category View` click `Performance and Maintenance` and then click `Administrative Tools..` If you are viewing the `Control Panel` in `Classic View`, click `Administrative Tools`.
3. In `Administrative Tools`, click `Data Sources (ODBC)`.

Irrespective of your Windows version, you should be presented the [ODBC Data Source Administrator](#) window:



Within Windows XP, you can add the [Administrative Tools](#) folder to your **Start** menu to make it easier to locate the ODBC Data Source Administrator. To do this:

1. Right-click the **Start** menu.
2. Select [Properties](#).
3. Click [Customize...](#)
4. Select the [Advanced](#) tab.
5. Within [Start menu items](#), within the [System Administrative Tools](#) section, select [Display on the All Programs menu](#).

Within both Windows Server 2003 and Windows XP, consider permanently adding the [ODBC Data Source Administrator](#) to your **Start** menu. To do this, locate the [Data Sources \(ODBC\)](#) icon using the methods shown, then right-click on the icon and then choose [Pin to Start Menu](#).

The interfaces for the 3.51 and 5.x versions of the Connector/ODBC driver are different, although the fields and information that you need to enter remain the same.

To configure a DSN using Connector/ODBC 5.2, see [Section 21.1.4.3.1, "Configuring a Connector/ODBC 5.x DSN on Windows"](#).

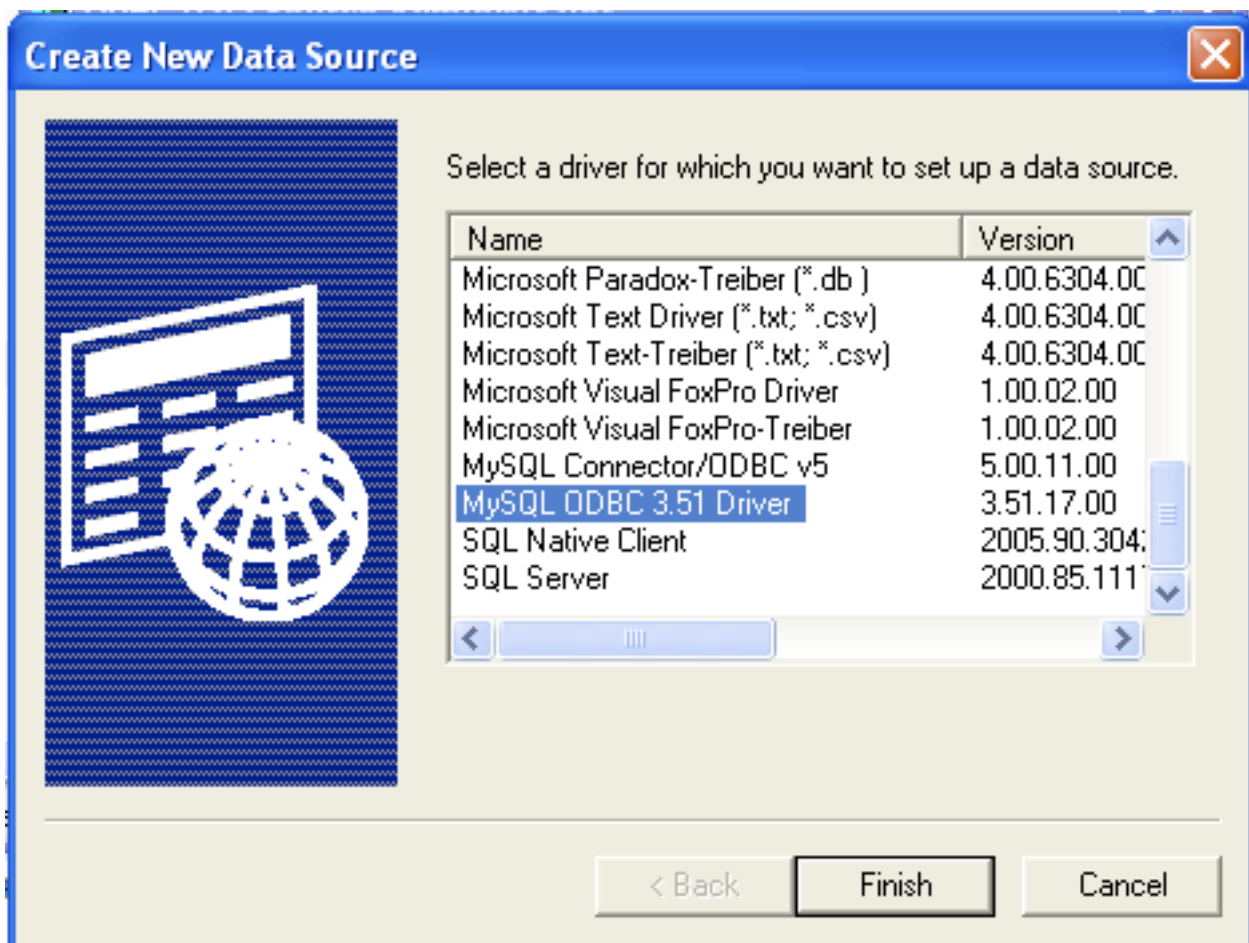


### 21.1.4.3.1. Configuring a Connector/ODBC 5.x DSN on Windows

Due to the native Unicode support within Connector/ODBC, you do not need to specify the initial character set to be used with your connection.


To add or configure a Connector/ODBC 5.x DSN on Windows, use either the command-line, or the [ODBC Data Source Administrator](#) GUI.

1. Open the [ODBC Data Source Administrator](#).
2. To create a System DSN (which will be available to all users), select the **System DSN** tab. To create a User DSN, which will be unique only to the current user, click the **Add...** button.
3. Select the ODBC driver for this DSN.



Select [MySQL ODBC 5.x Driver](#) for the appropriate level of Connector/ODBC, then click **Finish**.

4. You now need to configure the specific fields for the DSN you are creating through the [Connection Parameters](#) dialog.



In the **Data Source Name** box, enter the name of the data source to access. It can be any valid name that you choose.

5. In the **Description** box, enter some text to help identify the connection.
6. In the **Server** field, enter the name of the MySQL server host to access. By default, it is `localhost`.
7. In the **User** field, enter the user name to use for this connection.
8. In the **Password** field, enter the corresponding password for this connection.
9. The **Database** pop-up should automatically populate with the list of databases that the user has permissions to access.
10. To communicate over a different TCP/IP port than the default (3306), change the value of the **Port**.
11. Click **OK** to save the DSN.

To verify the connection using the parameters you have entered, click the **Test** button. If the connection could be made successfully, you will be notified with a `Success; connection was made!` dialog.

You can configure a number of options for a specific DSN by using the **Details** button.

**MySQL Connector/ODBC Data Source Configuration**

MySQL Connector/ODBC

**Connection Parameters**

Data Source Name:

Description:

Server:  Port:

User:

Password:

Database:

Flags 1 | Flags 2 | Flags 3 | Debug | SSL Settings

Return matched rows instead of affected rows

Allow big result sets

Use compression

Treat BIGINT columns as INT columns

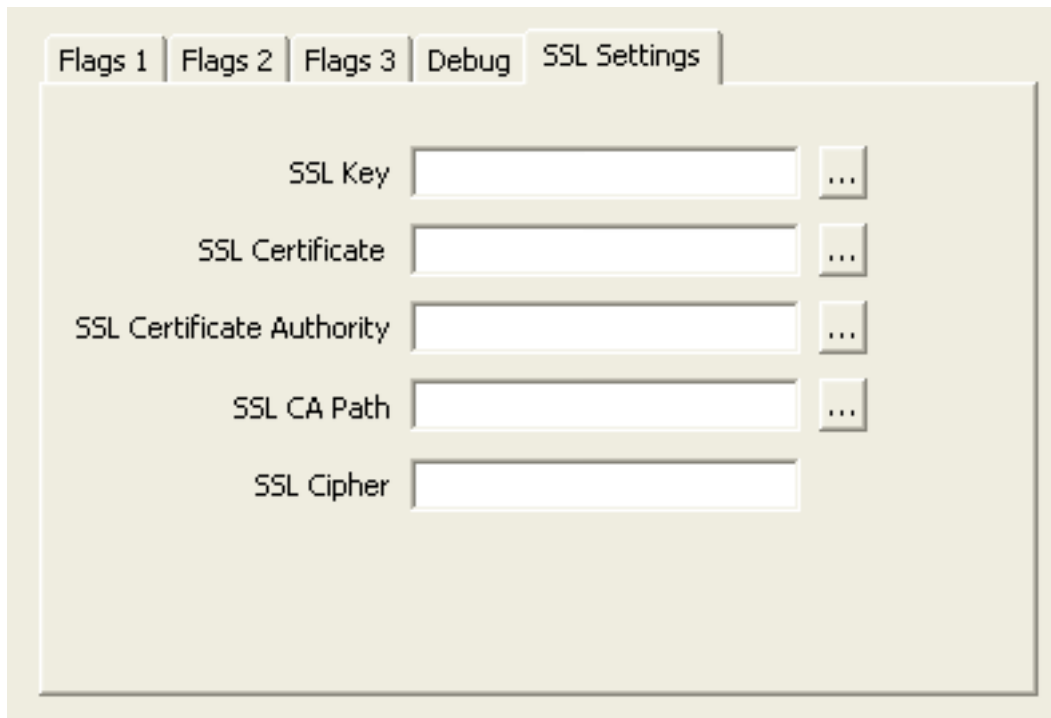
Enable safe options (see documentation)

Enable automatic reconnect

Enable SQL\_AUTO\_IS\_NULL

The **Details** button opens a tabbed display where you set additional options:

- **Flags 1**, **Flags 2**, and **Flags 3** enable you to select the additional flags for the DSN connection. For more information on these flags, see [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#).
- **Debug** lets you turn on ODBC debugging to record the queries you execute through the DSN to the `myodbc.sql` file. For more information, see [Section 21.1.4.8, “Getting an ODBC Trace File”](#).
- **SSL Settings** configures the additional options required for using the Secure Sockets Layer (SSL) when communicating with MySQL server. Note that you must also enable SSL and configure the MySQL server with suitable certificates to communicate over SSL.



The **Advanced** tab lets you configure Connector/ODBC connection parameters. Refer to [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#), for information about the meaning of these options.

#### 21.1.4.3.2. Configuring a Connector/ODBC 5.x DSN on Windows, Using the Command Line

Use `myodbc-installer.exe` when configuring Connector/ODBC 5.1 or later from the command-line.

Execute `myodbc-installer.exe` without arguments to view a list of available options.

#### 21.1.4.3.3. Troubleshooting ODBC Connection Problems

This section answers Connector/ODBC connection-related questions.

- **While configuring a Connector/ODBC DSN, a `Could Not Load Translator or Setup Library` error occurs**

For more information, refer to [MS KnowledgeBase Article\(Q260558\)](#). Also, make sure you have the latest valid `ct13d32.dll` in your system directory.

- On Windows, the default `myodbc5w.dll` (Unicode) or `myodbc5a.dll` (ANSI) is compiled for optimal performance. To debug Connector/ODBC (for example, to enable tracing), instead use `myodbc5d.dll`. To install this file, copy `myodbc5d.dll` over the installed `myodbc5w.dll` or `myodbc5a.dll` file. Make sure to revert back to the release version of the driver DLL once you are done with the debugging, because the debug version may cause performance issues.

#### 21.1.4.4. Configuring a Connector/ODBC DSN on Mac OS X

To configure a DSN on Mac OS X, you can either use the command-line utility (`myodbc-installer` with connector/ODBC 5.1, or `myodbc3i` with connector/OCBC 3.51), edit the `odbc.ini` file within the `Library/ODBC` directory of the user, or use the ODBC Administrator GUI. If you have Mac OS X 10.2 or earlier, refer to [Section 21.1.4.5, “Configuring a Connector/ODBC DSN on Unix”](#). Select whether to create a User DSN or a System DSN. When adding a System DSN, you might need to authenticate with the system. Click the padlock and enter a user and password with administrator privileges.

For correct operation of ODBC Administrator, ensure that the `/Library/ODBC/odbc.ini` file used to set up ODBC connectivity and DSNs are writable by the `admin` group. If this file is not writable by this group, then the ODBC Administrator may fail, or may appear to work but not generate the correct entry.



#### Warning

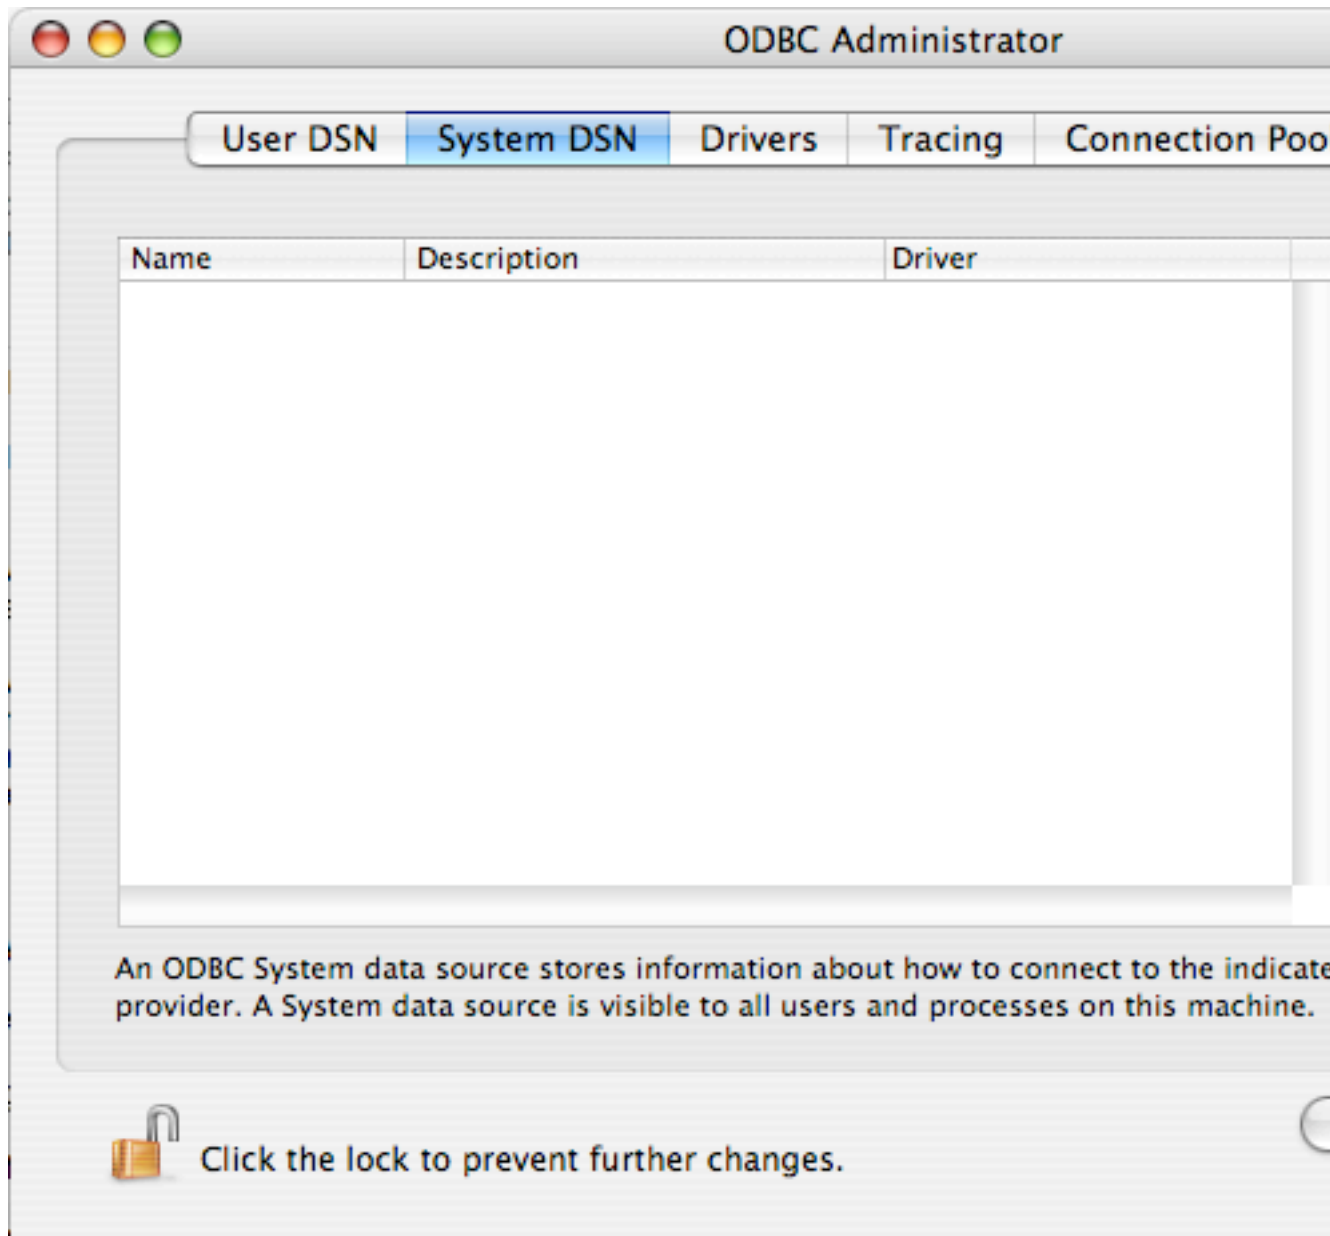
There are known issues with the Mac OS X ODBC Administrator and Connector/ODBC that may prevent you from creating a DSN using this method. In this case, use the command line or edit the `odbc.ini` file directly. Note that existing DSNs or those that you create using the `myodbc3i` or `myodbc-installer` tool can still be checked and edited using ODBC Administrator.

To create a DSN using the `myodbc3i` utility, you need only specify the DSN type and the DSN connection string. For example:

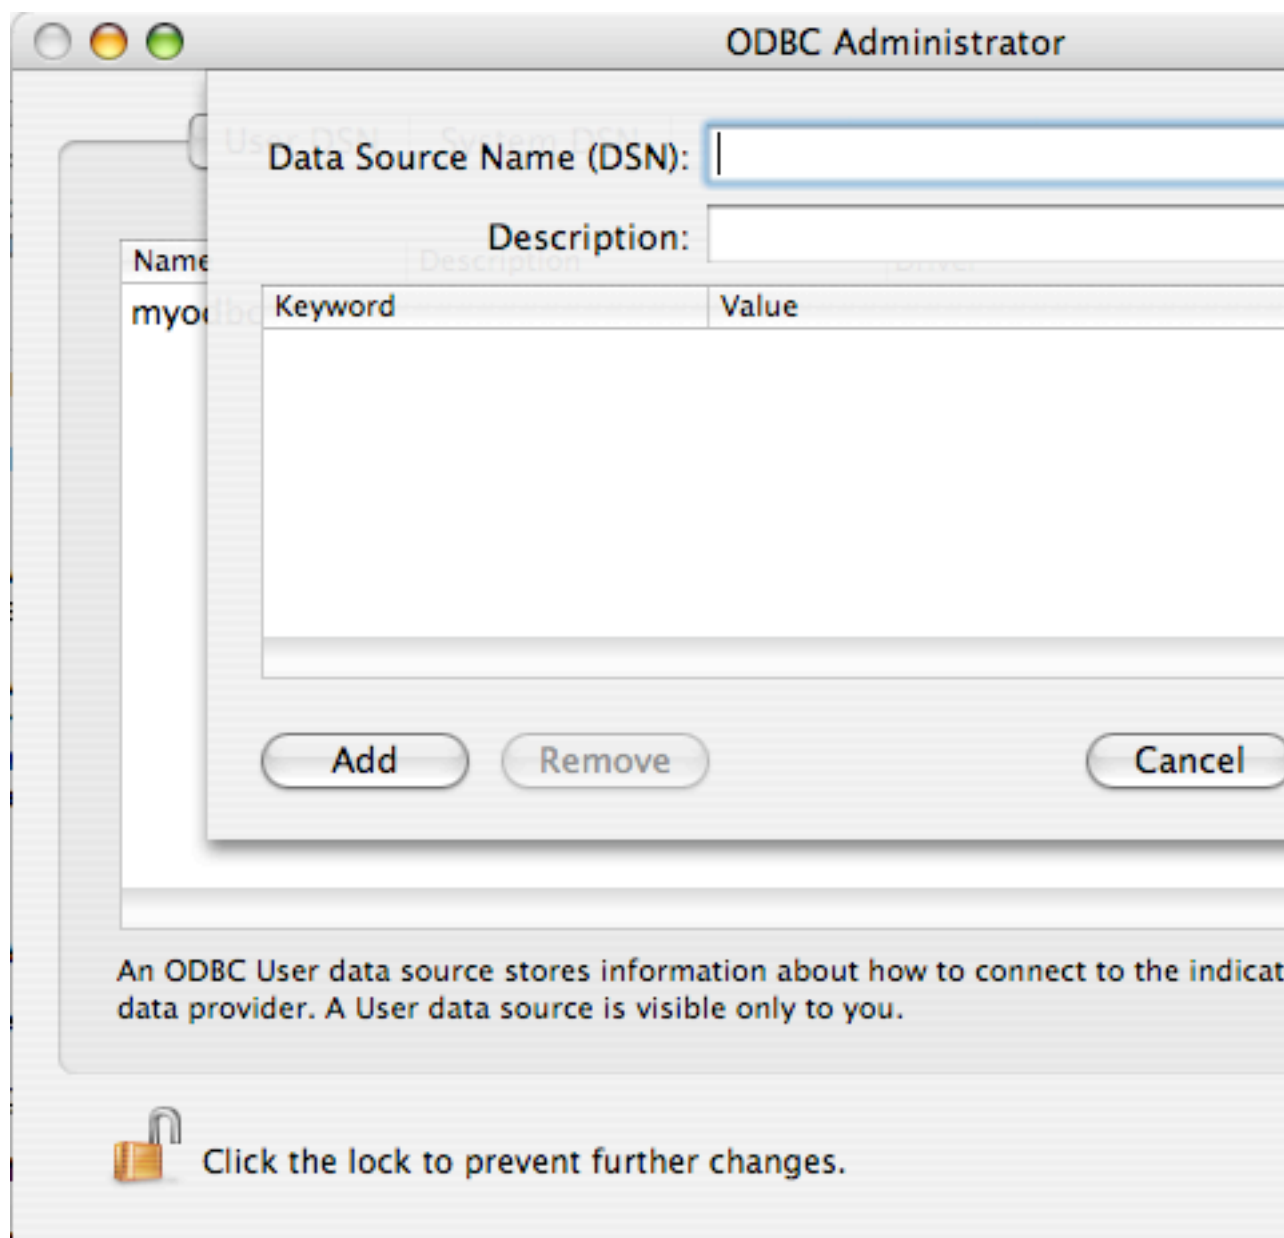
```
shell> myodbc3i -a -s -t"DSN=mydb;DRIVER=MySQL ODBC 3.51 Driver;SERVER=mysql;USER=username;PASSWORD=pa
```

To use ODBC Administrator:

1. Open the ODBC Administrator from the [Utilities](#) folder in the [Applications](#) folder.



2. On the User DSN or System DSN panel, click **Add**.
3. Select the Connector/ODBC driver and click **OK**.
4. You will be presented with the **Data Source Name** dialog. Enter the **Data Source Name** and an optional **Description** for the DSN.



5. Click **Add** to add a new keyword/value pair to the panel. Configure at least four pairs to specify the [server](#), [username](#), [password](#) and [database](#) connection parameters. See [Section 21.1.4.2, "Connector/ODBC Connection Parameters"](#).
6. Click **OK** to add the DSN to the list of configured data source names.

A completed DSN configuration may look like this:



Data Source Name (DSN): **WorldSample**

Description: **Connection to sample World database**

Keyword	Value
server	mysql
user	sakila
password	Sample
database	test_world

Buttons: Add, Remove, Cancel, OK

You can configure other ODBC options in your DSN by adding further keyword/value pairs and setting the corresponding values. See [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#).

#### 21.1.4.5. Configuring a Connector/ODBC DSN on Unix

On [Unix](#), you configure DSN entries directly in the `odbc.ini` file. Here is a typical `odbc.ini` file that configures `myodbc3` as the DSN name for Connector/ODBC 3.51:

```

;
; odbc.ini configuration for Connector/ODBC and Connector/ODBC 3.51 drivers
;

[ODBC Data Sources]
myodbc3      = MyODBC 3.51 Driver DSN

[myodbc3]
Driver       = /usr/local/lib/libmyodbc3.so
Description  = Connector/ODBC 3.51 Driver DSN
SERVER      = localhost
PORT        =
USER        = root
Password    =
Database    = test
OPTION      = 3
SOCKET      =

[Default]
Driver       = /usr/local/lib/libmyodbc3.so
Description  = Connector/ODBC 3.51 Driver DSN
SERVER      = localhost
PORT        =
USER        = root
Password    =
Database    = test
OPTION      = 3
SOCKET      =

```

Refer to the [Section 21.1.4.2, “Connector/ODBC Connection Parameters”](#), for the list of connection parameters that can be supplied.

**Note**

If you are using `unixODBC`, you can use the following tools to set up the DSN:

- `ODBCConfig` GUI tool ([HOWTO: ODBCConfig](#))
- `odbcinst`

In some cases when using `unixODBC`, you might get this error:

```
Data source name not found and no default driver specified
```

If this happens, make sure the `ODBCINI` and `ODBCSYSINI` environment variables are pointing to the right `odbc.ini` file. For example, if your `odbc.ini` file is located in `/usr/local/etc`, set the environment variables like this:

```
export ODBCINI=/usr/local/etc/odbc.ini
export ODBCSYSINI=/usr/local/etc
```

### 21.1.4.6. Connecting Without a Predefined DSN

You can connect to the MySQL server using `SQLDriverConnect`, by specifying the `DRIVER` name field. Here are the connection strings for Connector/ODBC using DSN-Less connections:

#### For Connector/ODBC 3.51:

```
ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};\
SERVER=localhost;\
DATABASE=test;\
USER=venu;\
PASSWORD=venu;\
OPTION=3;"
```

If your programming language converts backslash followed by whitespace to a space, it is preferable to specify the connection string as a single long string, or to use a concatenation of multiple strings that does not add spaces in between. For example:

```
ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};"
"SERVER=localhost;"
"DATABASE=test;"
"USER=venu;"
"PASSWORD=venu;"
"OPTION=3;"
```

**Note.** On Mac OS X, you might need to specify the full path to the Connector/ODBC driver library.

Refer to [Section 21.1.4.2, "Connector/ODBC Connection Parameters"](#) for the list of connection parameters that can be supplied.

### 21.1.4.7. ODBC Connection Pooling

Connection pooling enables the ODBC driver to re-use existing connections to a given database from a pool of connections, instead of opening a new connection each time the database is accessed. By enabling connection pooling you can improve the overall performance of your application by lowering the time taken to open a connection to a database in the connection pool.

For more information about connection pooling: <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q169470>.

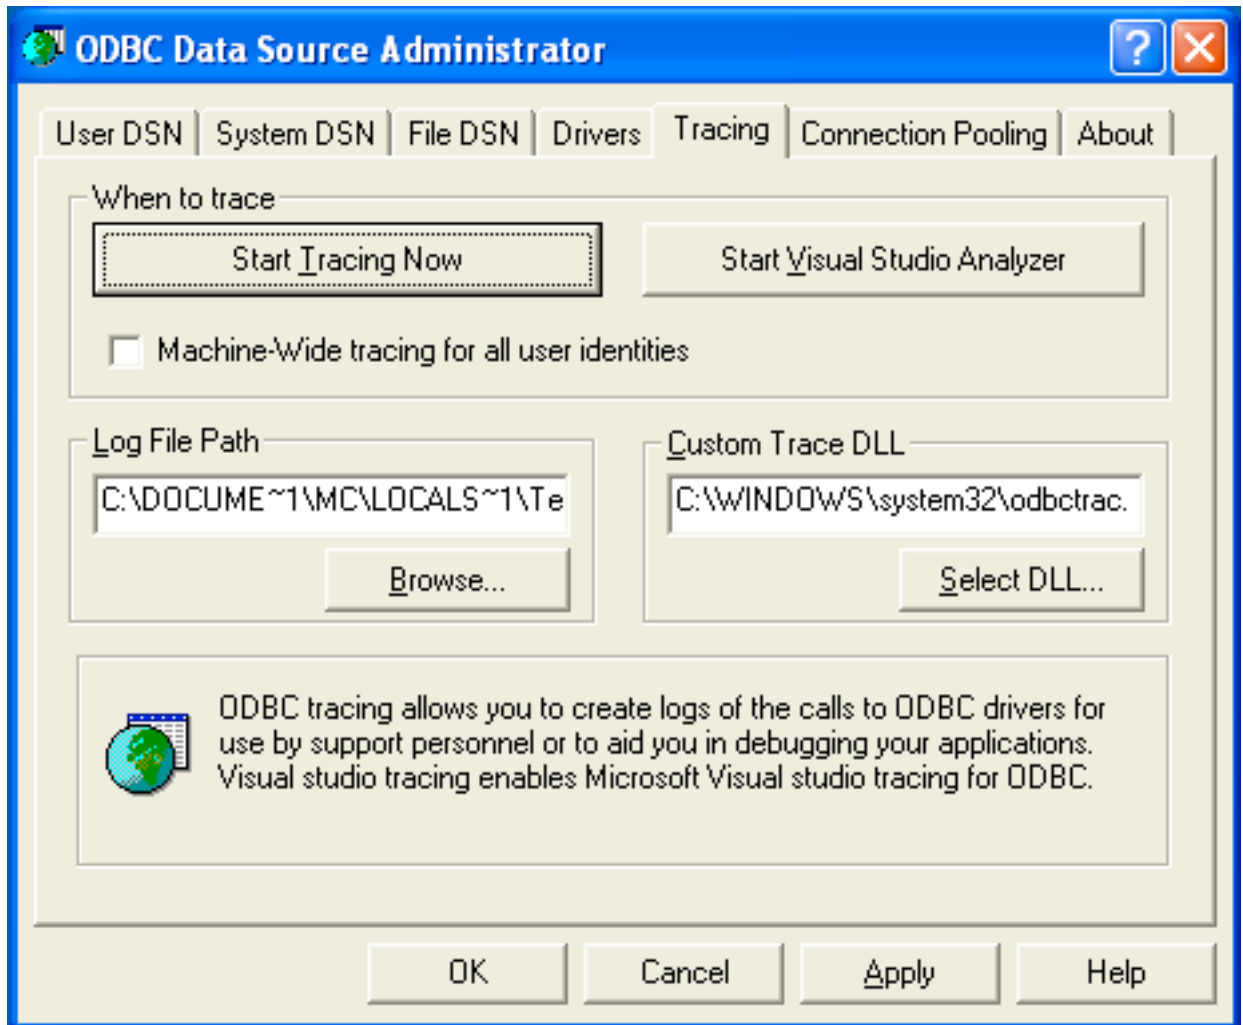
### 21.1.4.8. Getting an ODBC Trace File

If you encounter difficulties or problems with Connector/ODBC, start by making a log file from the `ODBC Manager` and Connector/ODBC. This is called *tracing*, and is enabled through the ODBC Manager. The procedure for this differs for Windows, Mac OS X and Unix.

### 21.1.4.8.1. Enabling ODBC Tracing on Windows

To enable the trace option on Windows:

1. The [Tracing](#) tab of the ODBC Data Source Administrator dialog box lets you configure the way ODBC function calls are traced.

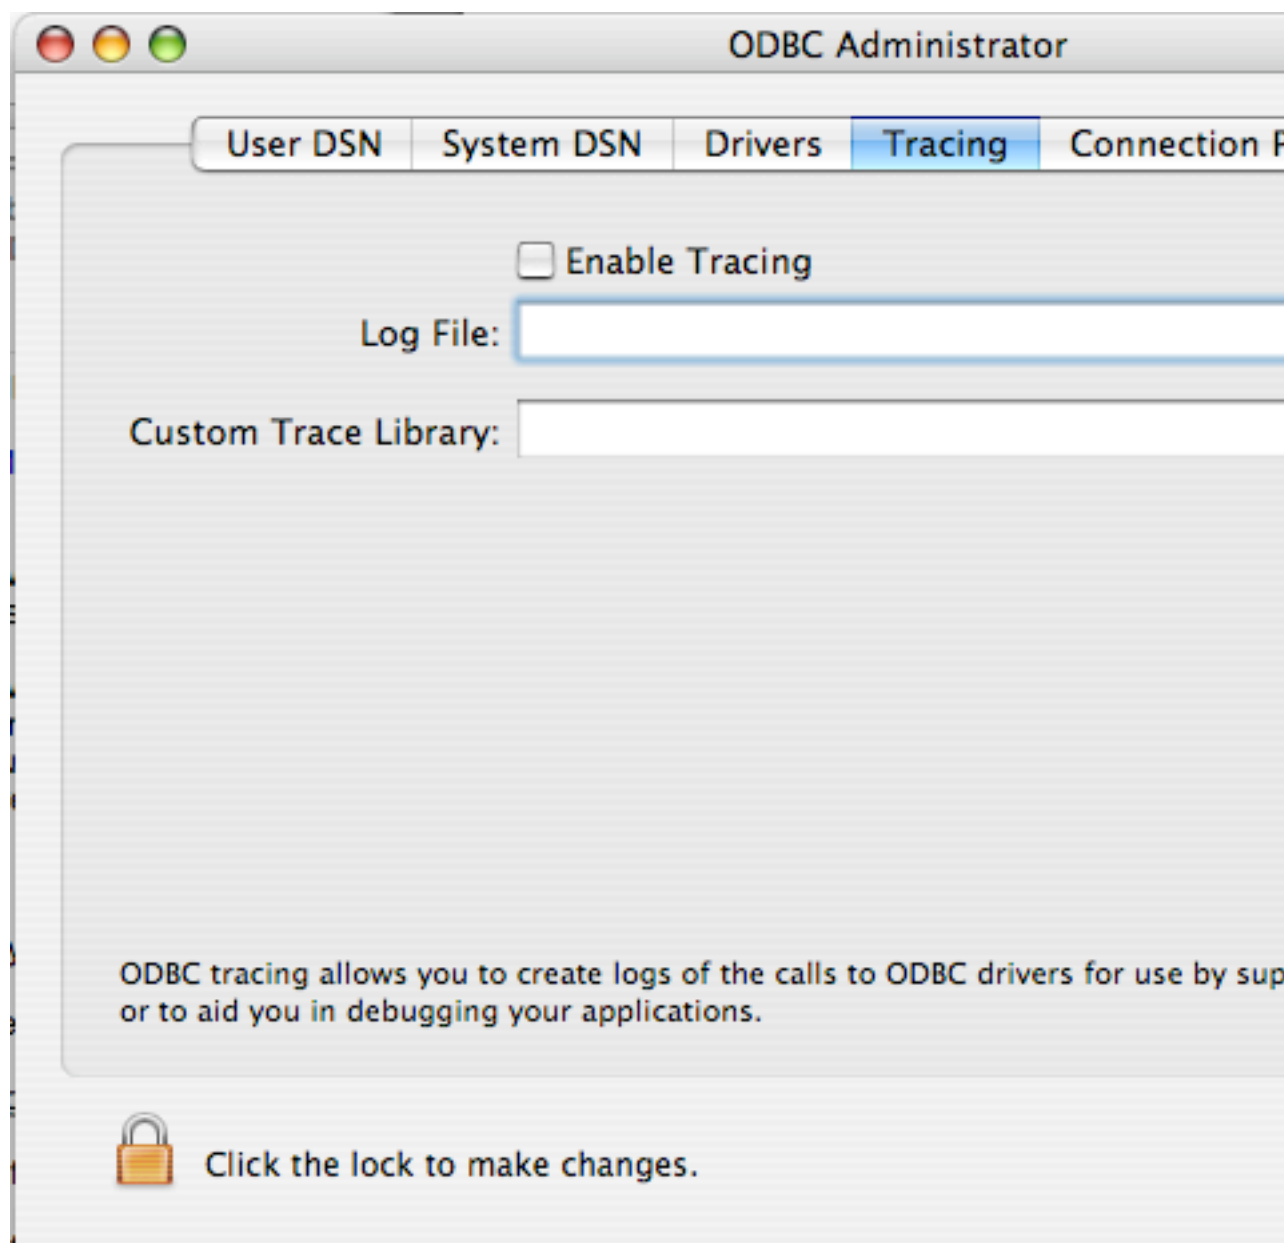


2. When you activate tracing from the [Tracing](#) tab, the [Driver Manager](#) logs all ODBC function calls for all subsequently run applications.
3. ODBC function calls from applications running before tracing is activated are not logged. ODBC function calls are recorded in a log file you specify.
4. Tracing ceases only after you click [Stop Tracing Now](#). Remember that while tracing is on, the log file continues to increase in size and that tracing affects the performance of all your ODBC applications.

### 21.1.4.8.2. Enabling ODBC Tracing on Mac OS X

To enable the trace option on Mac OS X 10.3 or later, use the [Tracing](#) tab within ODBC Administrator .

1. Open the ODBC Administrator.
2. Select the [Tracing](#) tab.



3. Select the `Enable Tracing` check box.
4. Enter the location to save the Tracing log. To append information to an existing log file, click the `Choose...` button.

#### 21.1.4.8.3. Enabling ODBC Tracing on Unix

To enable the trace option on Mac OS X 10.2 (or earlier) or Unix, add the `trace` option to the ODBC configuration:

1. On Unix, explicitly set the `Trace` option in the `ODBC.INI` file.

Set the tracing `ON` or `OFF` by using `TraceFile` and `Trace` parameters in `odbc.ini` as shown below:

```
TraceFile = /tmp/odbc.trace
Trace     = 1
```

`TraceFile` specifies the name and full path of the trace file and `Trace` is set to `ON` or `OFF`. You can also use `1` or `YES` for `ON` and `0` or `NO` for `OFF`. If you are using `ODBCConfig` from `unixODBC`, then follow the instructions for tracing `unixODBC` calls at [HOWTO-ODBCConfig](#).

#### 21.1.4.8.4. Enabling a Connector/ODBC Log

To generate a Connector/ODBC log, do the following:

1. Within Windows, enable the `Trace Connector/ODBC` option flag in the Connector/ODBC connect/configure screen. The log is written to file `C:\myodbc.log`. If the trace option is not remembered when you are going back to the above screen, it means that you are not using the `myodbcd.dll` driver, see [Section 21.1.4.3.3, "Troubleshooting ODBC Connection Problems"](#).

On Mac OS X, Unix, or if you are using a DSN-less connection, either supply `OPTION=4` in the connection string, or set the corresponding keyword/value pair in the DSN.

2. Start your application and try to get it to fail. Then check the Connector/ODBC trace file to find out what could be wrong.

If you need help determining what is wrong, see [Section 21.1.8.1, "Connector/ODBC Community Support"](#).

### 21.1.5. Connector/ODBC Examples

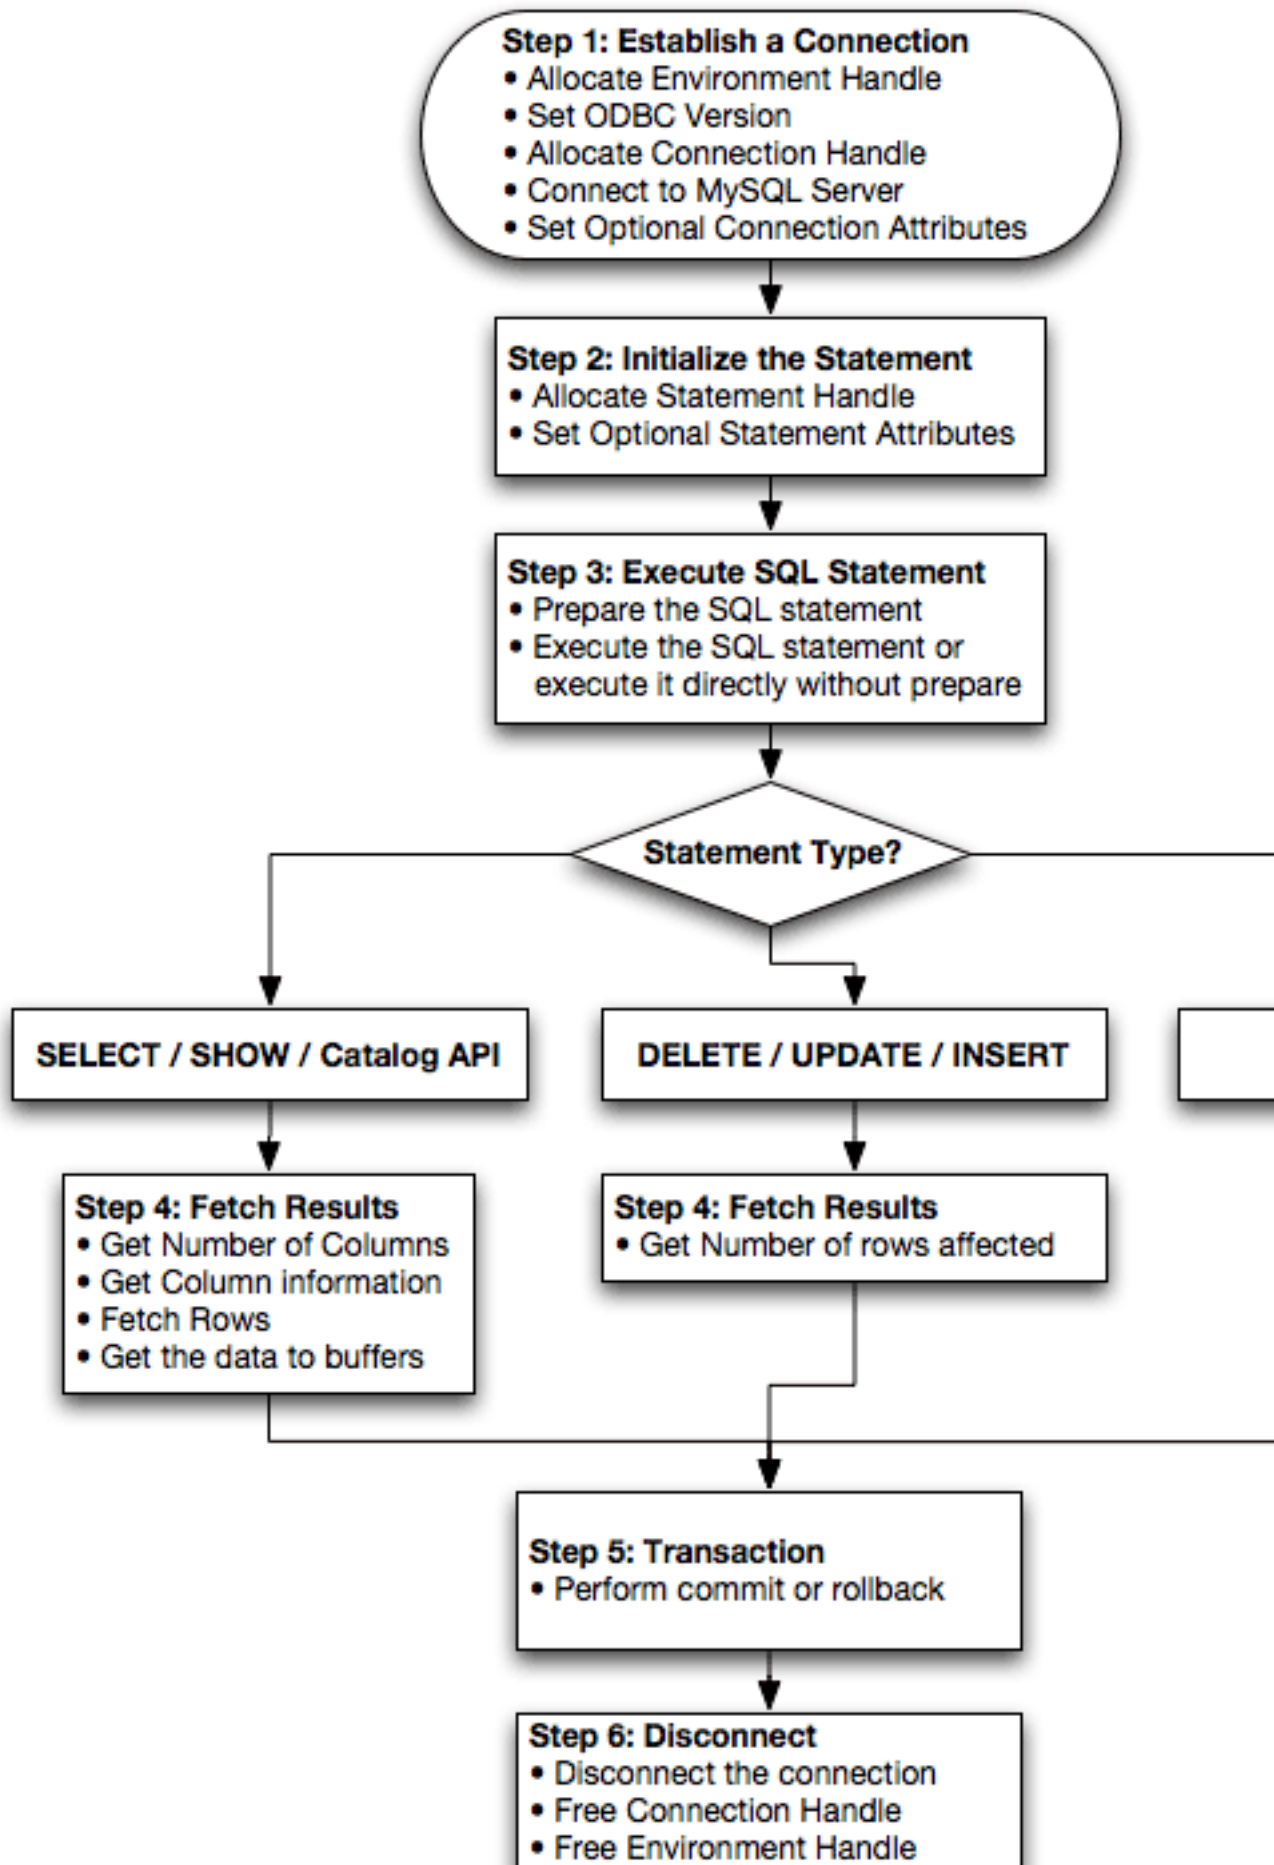
Once you have configured a DSN to provide access to a database, how you access and use that connection is dependent on the application or programming language. As ODBC is a standardized interface, any application or language that supports ODBC can use the DSN and connect to the configured database.

#### 21.1.5.1. Basic Connector/ODBC Application Steps

Interacting with a MySQL server from an applications using the Connector/ODBC typically involves the following operations:

- Configure the Connector/ODBC DSN.
- Connect to MySQL server.
- Initialization operations.
- Execute SQL statements.
- Retrieve results.
- Perform [transactions](#).
- Disconnect from the server.

Most applications use some variation of these steps. The basic application steps are shown in the following diagram:



### 21.1.5.2. Step-by-step Guide to Connecting to a MySQL Database through Connector/ODBC

A typical situation where you would install Connector/ODBC is to access a database on a Linux or Unix host from a Windows machine.

As an example of the process required to set up access between two machines, the steps below take you through the basic steps. These instructions assume that you connect to system ALPHA from system BETA with a user name and password of `myuser` and `mypassword`.

On system ALPHA (the MySQL server) follow these steps:

1. Start the MySQL server.
2. Use `GRANT` to set up an account with a user name of `myuser` that can connect from system BETA using a password of `myuser` to the database `test`:

```
GRANT ALL ON test.* to 'myuser'@'BETA' IDENTIFIED BY 'mypassword';
```

For more information about MySQL privileges, refer to [Section 6.3, "MySQL User Account Management"](#).

On system BETA (the Connector/ODBC client), follow these steps:

1. Configure a Connector/ODBC DSN using parameters that match the server, database and authentication information that you have just configured on system ALPHA.

Parameter	Value	Comment
DSN	remote_test	A name to identify the connection.
SERVER	ALPHA	The address of the remote server.
DATABASE	test	The name of the default database.
USER	myuser	The user name configured for access to this database.
PASSWORD	mypassword	The password for <code>myuser</code> .

2. Using an ODBC-capable application, such as Microsoft Office, connect to the MySQL server using the DSN you have just created. If the connection fails, use tracing to examine the connection process. See [Section 21.1.4.8, "Getting an ODBC Trace File"](#), for more information.

### 21.1.5.3. Connector/ODBC and Third-Party ODBC Tools

Once you have configured your Connector/ODBC DSN, you can access your MySQL database through any application that supports the ODBC interface, including programming languages and third-party applications. This section contains guides and help on using Connector/ODBC with various ODBC-compatible tools and applications, including Microsoft Word, Microsoft Excel and Adobe/Macromedia ColdFusion.

Connector/ODBC has been tested with the following applications:

Publisher	Application	Notes
Adobe	ColdFusion	Formerly Macromedia ColdFusion
Borland	C++ Builder	
	Builder 4	
	Delphi	
Business Objects	Crystal Reports	
Claris	Filemaker Pro	



Publisher	Application	Notes
Corel	Paradox	
Computer Associates	Visual Objects	Also known as CAVO
	AllFusion ERwin Data Modeler	
Gupta	Team Developer	Previously known as Centura Team Developer; Gupta SQL/Windows
Gensym	G2-ODBC Bridge	
Inline	iHTML	
Lotus	Notes	Versions 4.5 and 4.6
Microsoft	Access	
	Excel	
	Visio Enterprise	
	Visual C++	
	Visual Basic	
	ODBC.NET	Using C#, Visual Basic, C++
	FoxPro	
	Visual Interdev	
OpenOffice.org	OpenOffice.org	
Perl	DBD::ODBC	
Pervasive Software	DataJunction	
Sambar Technologies	Sambar Server	
SPSS	SPSS	
SoftVelocity	Clarion	
SQLExpress	SQLExpress for Xbase+ +	
Sun	StarOffice	
SunSystems	Vision	
Sybase	PowerBuilder	
	PowerDesigner	
theKompany.com	Data Architect	

If you know of any other applications that work with Connector/ODBC, please send mail to [myodbc@lists.mysql.com](mailto:myodbc@lists.mysql.com) about them.

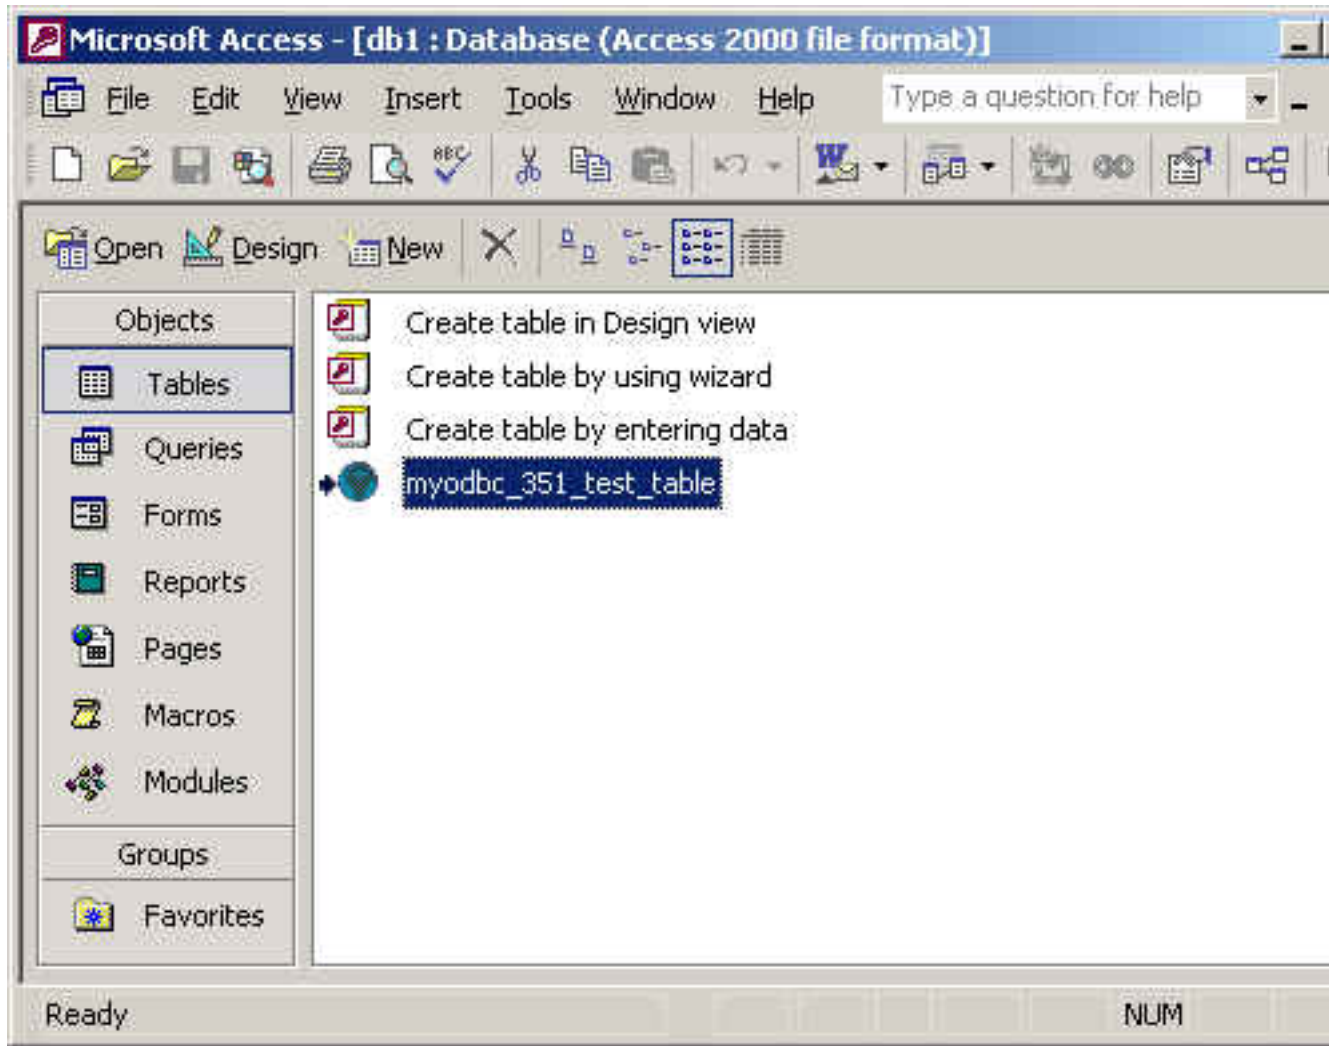
#### 21.1.5.4. Using Connector/ODBC with Microsoft Access

You can use MySQL database with Microsoft Access using Connector/ODBC. The MySQL database can be used as an import source, an export source, or as a linked table for direct use within an Access application, so you can use Access as the front-end interface to a MySQL database.

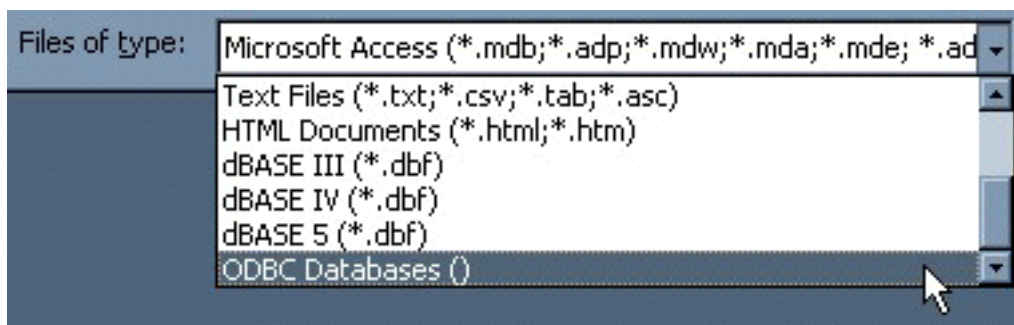
##### 21.1.5.4.1. Exporting Access Data to MySQL

To export a table of data from an Access database to MySQL, follow these instructions:

1. When you open an Access database or an Access project, a Database window appears. It displays shortcuts for creating new database objects and opening existing objects.



2. Click the name of the [table](#) or [query](#) to export, and then in the [File](#) menu, select [Export](#).
3. In the [Export Object Type Object name To](#) dialog box, in the [Save As Type](#) box, select [ODBC Databases \(\)](#) as shown here:



4. In the [Export](#) dialog box, enter a name for the file (or use the suggested name), and then select [OK](#).
5. The [Select Data Source](#) dialog box is displayed; it lists the defined data sources for any ODBC drivers installed on your computer. Click either the [File Data Source](#) or [Machine Data Source](#) tab, and then double-click the [Connector/ODBC](#) or [Connector/ODBC 3.51](#) data source to export to. To define a new data source for [Connector/ODBC](#), please [Section 21.1.4.3, "Configuring a Connector/ODBC DSN on Windows"](#).

**Note**

Ensure that the information that you are exporting to the MySQL table is valid for the corresponding MySQL data types. Values that are outside of the supported range of the MySQL data type but valid within Access may trigger an “overflow” error during the export.

Microsoft Access connects to the MySQL Server through this data source and exports new tables and or data.

#### 21.1.5.4.2. Importing MySQL Data to Access

To import a table or tables from MySQL to Access, follow these instructions:

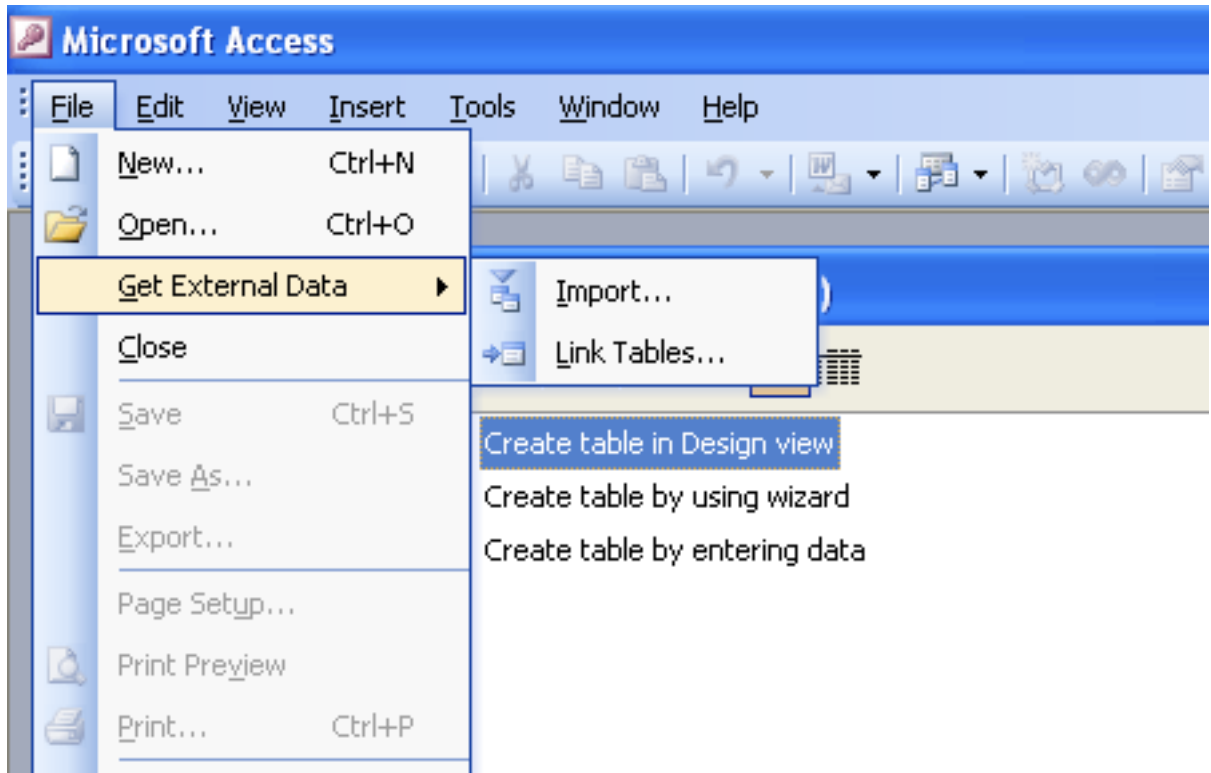
1. Open a database, or switch to the Database window for the open database.
2. To import tables, on the **File** menu, point to **Get External Data**, and then click **Import**.
3. In the **Import** dialog box, in the Files Of Type box, select **ODBC Databases (J)**. The Select Data Source dialog box lists the defined data sources **The Select Data Source** dialog box is displayed; it lists the defined data source names.
4. If the ODBC data source that you selected requires you to log on, enter your login ID and password (additional information might also be required), and then click **OK**.
5. Microsoft Access connects to the MySQL server through **ODBC data source** and displays the list of tables that you can **import**.
6. Click each table to **import**, and then click **OK**.

#### 21.1.5.4.3. Using Microsoft Access as a Front-end to MySQL

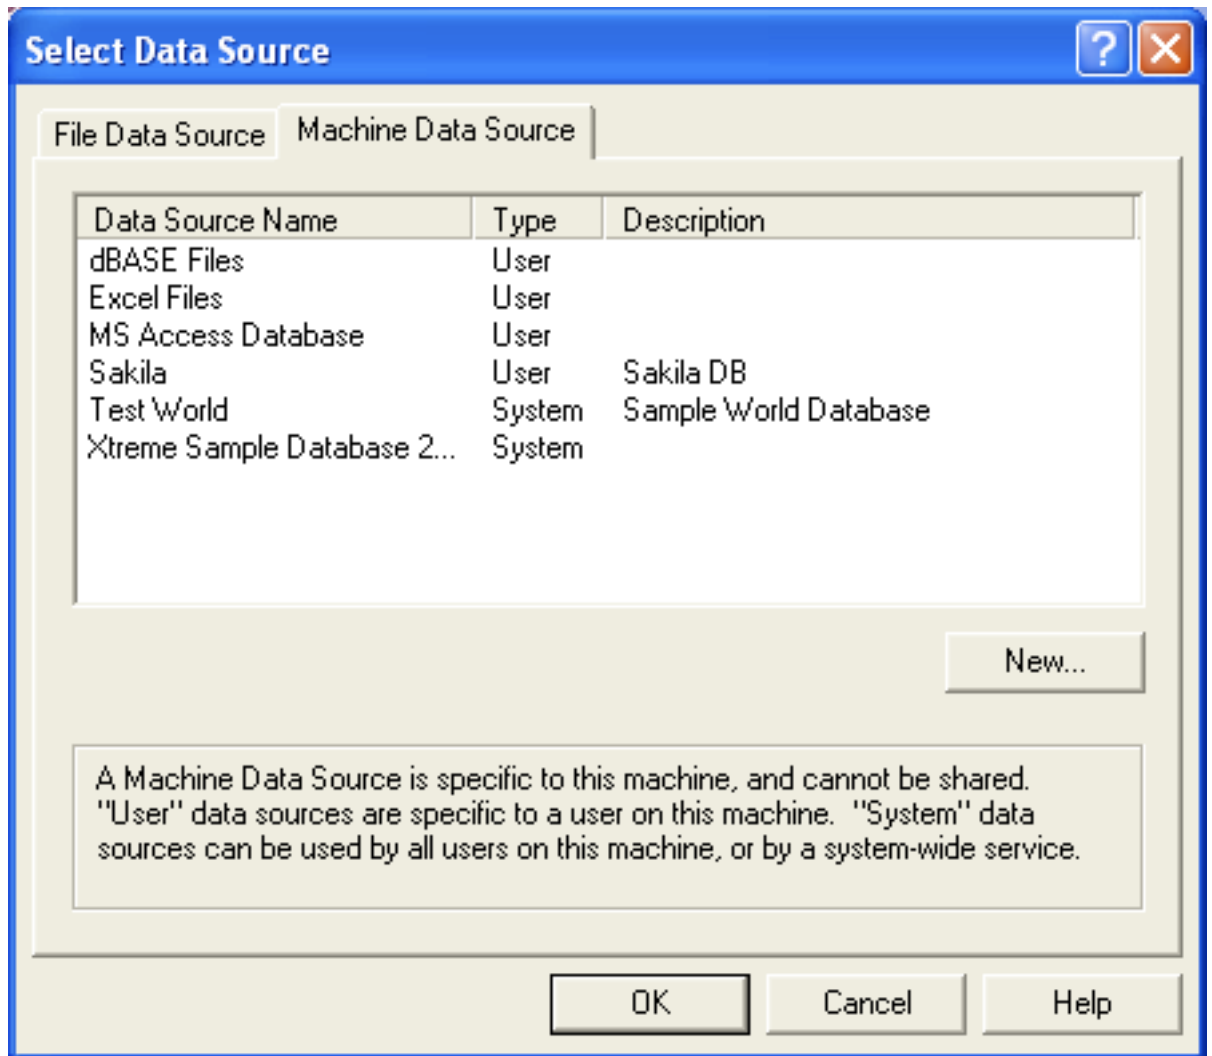
You can use Microsoft Access as a front end to a MySQL database by linking tables within your Microsoft Access database to tables that exist within your MySQL database. When a query is requested on a table within Access, ODBC is used to execute the queries on the MySQL database instead.

**To create a linked table:**

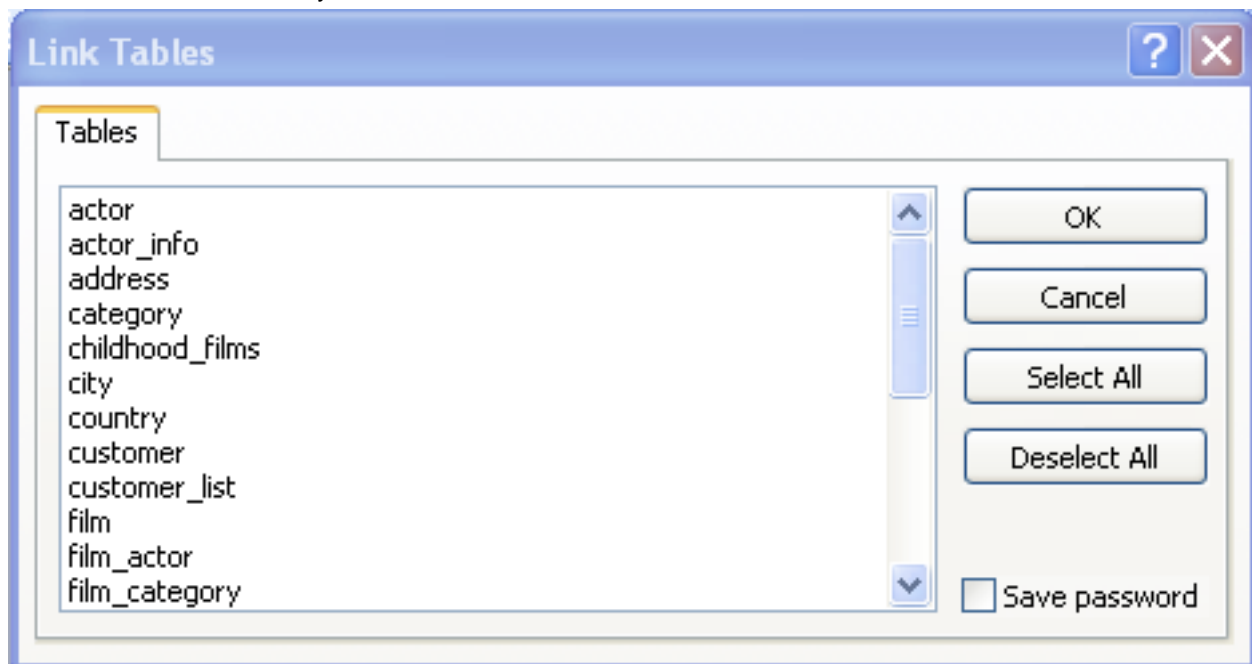
1. Open the Access database that you want to link to MySQL.
2. From the **File**, choose **Get External Data->Link Tables**.



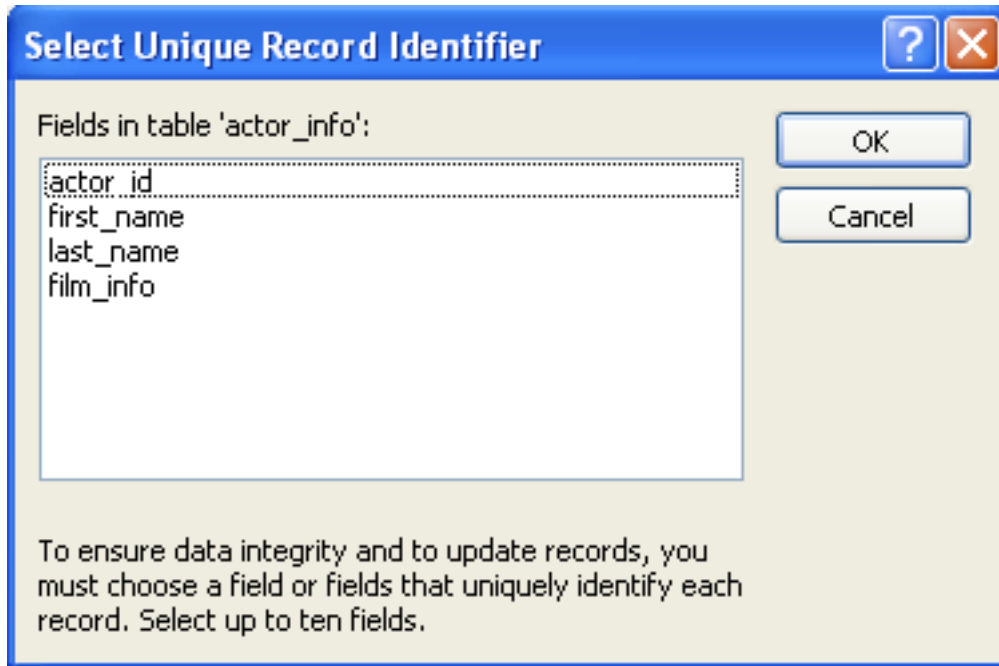
3. From the browser, choose **ODBC Databases ()** from the **Files of type** pop-up.
4. In the **Select Data Source** window, choose an existing DSN, either from a **File Data Source** or **Machine Data Source**. You can also create a new DSN using the **New...** button. For more information on creating a DSN see [Section 21.1.4.3, "Configuring a Connector/ODBC DSN on Windows"](#).



- In the **Link Tables** dialog, select one or more tables from the MySQL database. A link will be created to each table that you select from this list.



6. If Microsoft Access is unable to determine the unique record identifier for a table automatically then it may ask you to confirm the column, or combination of columns, to be used to uniquely identify each row from the source table. Select the columns to use and click **OK**.



Once the process has been completed, you can now build interfaces and queries to the linked tables just as you would for any Access database.

Use the following procedure to view or to refresh links when the structure or location of a linked table has changed. The Linked Table Manager lists the paths to all currently linked tables.

**To view or refresh links:**

1. Open the database that contains links to MySQL tables.
2. On the **Tools** menu, point to **Add-ins** (**Database Utilities** in Access 2000 or newer), and then click **Linked Table Manager**.
3. Select the check box for the tables whose links you want to refresh.
4. Click **OK** to refresh the links.

Microsoft Access confirms a successful refresh or, if the table wasn't found, displays the **Select New Location of <table name>** dialog box in which you can specify its the table's new location. If several selected tables have moved to the new location that you specify, the Linked Table Manager searches that location for all selected tables, and updates all links in one step.

**To change the path for a set of linked tables:**

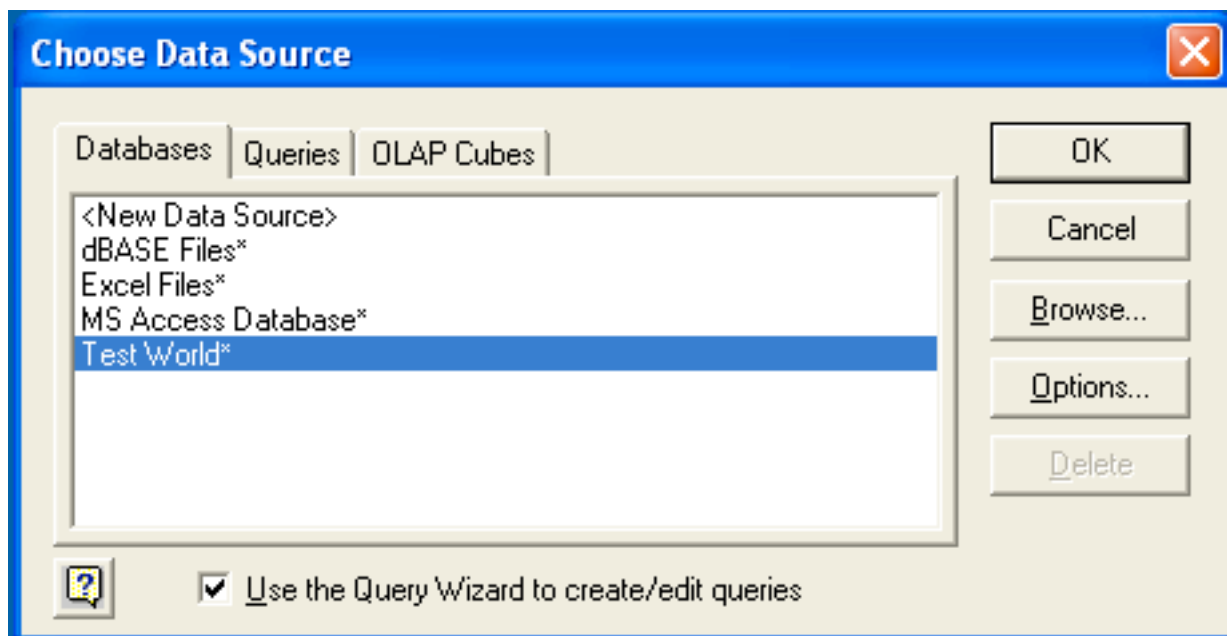
1. Open the database that contains links to tables.
2. On the **Tools** menu, point to **Add-ins** (**Database Utilities** in Access 2000 or newer), and then click **Linked Table Manager**.
3. Select the **Always Prompt For A New Location** check box.
4. Select the check box for the tables whose links you want to change, and then click **OK**.
5. In the **Select New Location of <table name>** dialog box, specify the new location, click **Open**, and then click **OK**.

### 21.1.5.5. Using Connector/ODBC with Microsoft Word or Excel

You can use Microsoft Word and Microsoft Excel to access information from a MySQL database using Connector/ODBC. Within Microsoft Word, this facility is most useful when importing data for mailmerge, or for tables and data to be included in reports. Within Microsoft Excel, you can execute queries on your MySQL server and import the data directly into an Excel Worksheet, presenting the data as a series of rows and columns.

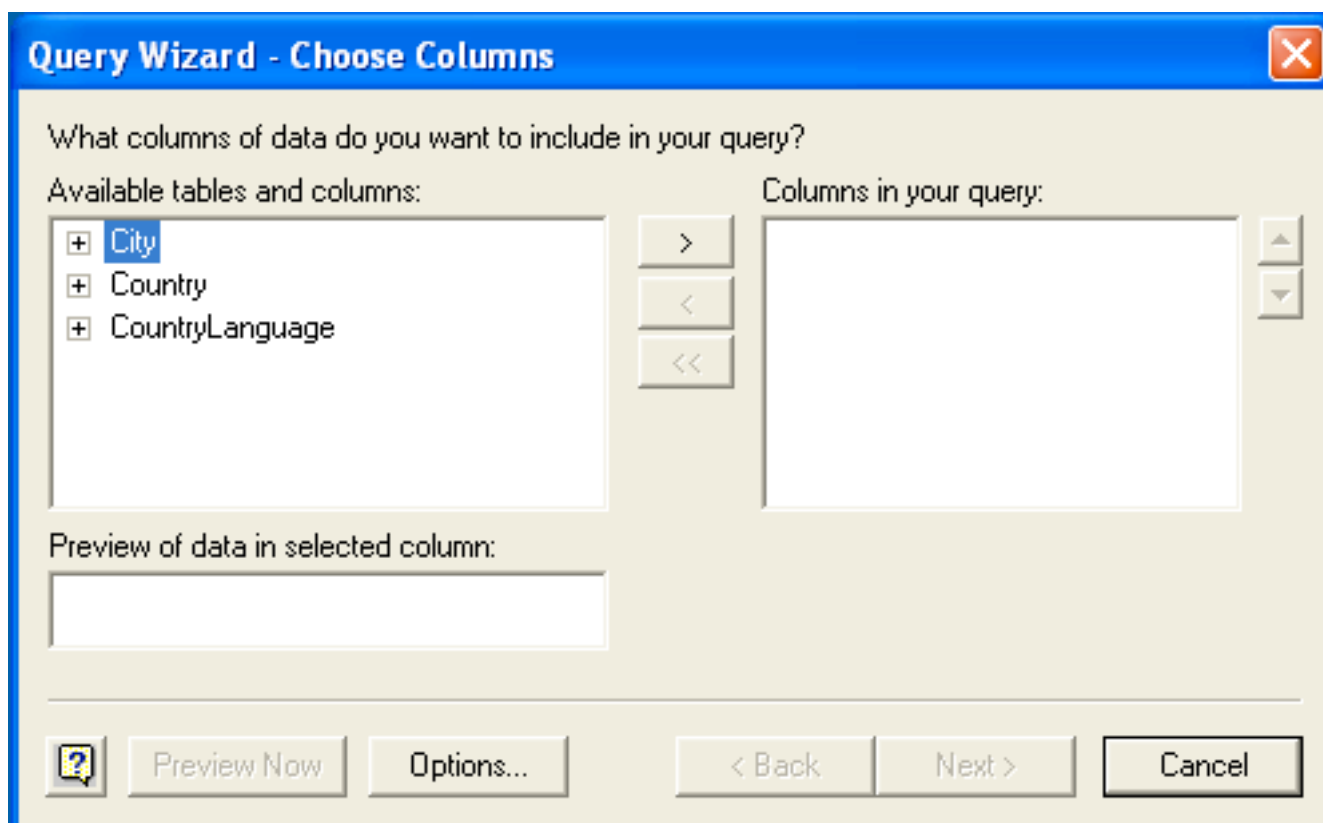
With both applications, data is accessed and imported into the application using Microsoft Query, which lets you execute a query through an ODBC source. You use Microsoft Query to build the SQL statement to be executed, selecting the tables, fields, selection criteria and sort order. For example, to insert information from a table in the World test database into an Excel spreadsheet, using the DSN samples shown in [Section 21.1.4, "Configuring Connector/ODBC"](#):

1. Create a new Worksheet.
2. From the **Data** menu, choose **Import External Data**, and then select **New Database Query**.
3. Microsoft Query will start. First, you need to choose the data source, by selecting an existing Data Source Name.

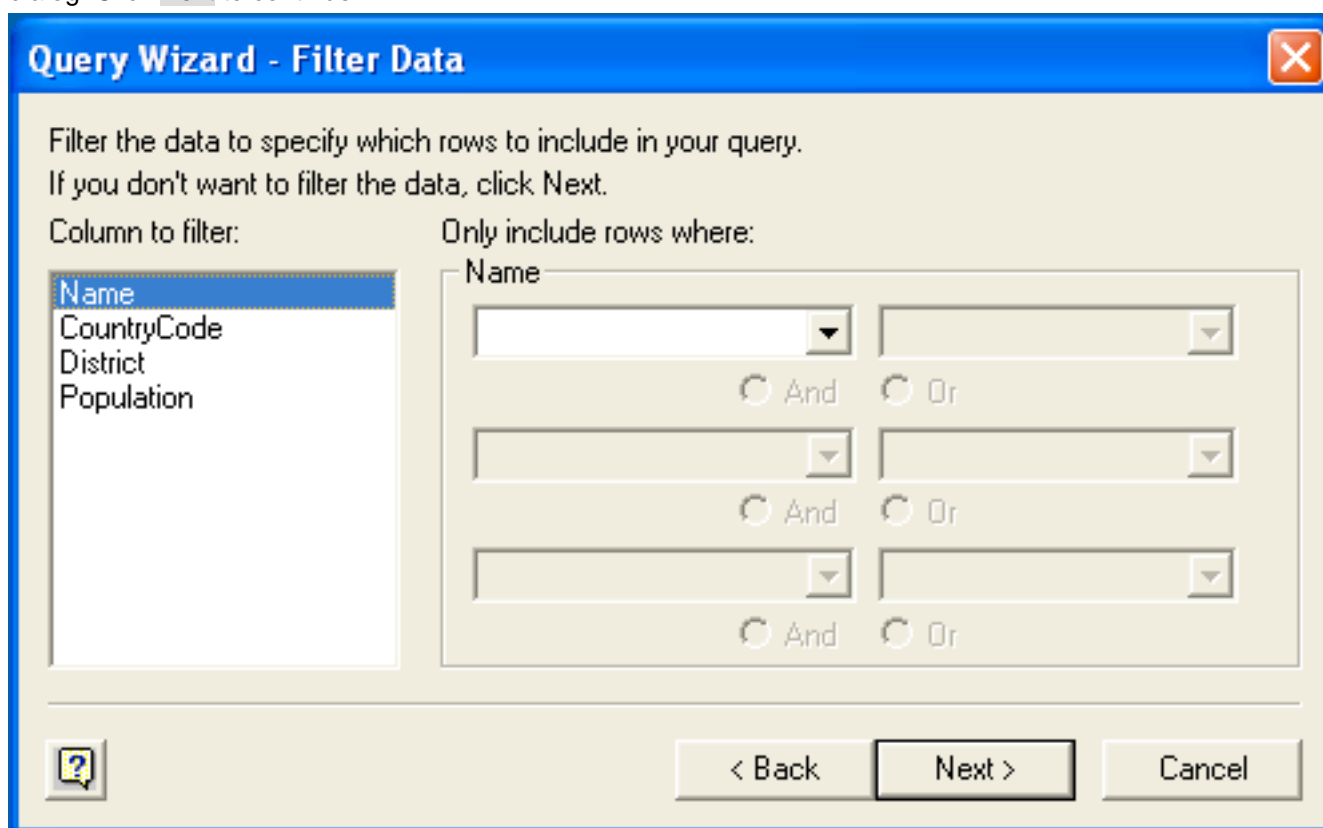


4. Within the **Query Wizard**, choose the columns to import. The list of tables available to the user configured through the DSN is shown on the left, the columns that will be added to your query are shown on the right. The columns you choose are equivalent to those in the first section of a **SELECT** query. Click **Next** to continue.





5. You can filter rows from the query (the equivalent of a `WHERE` clause) using the [Filter Data](#) dialog. Click **Next** to continue.



6. Select an (optional) sort order for the data. This is equivalent to using a `ORDER BY` clause in your SQL query. You can select up to three fields for sorting the information returned by the query. Click **Next** to continue.


**Query Wizard - Sort Order**

Specify how you want your data sorted.  
If you don't want to sort the data, click **Next**.

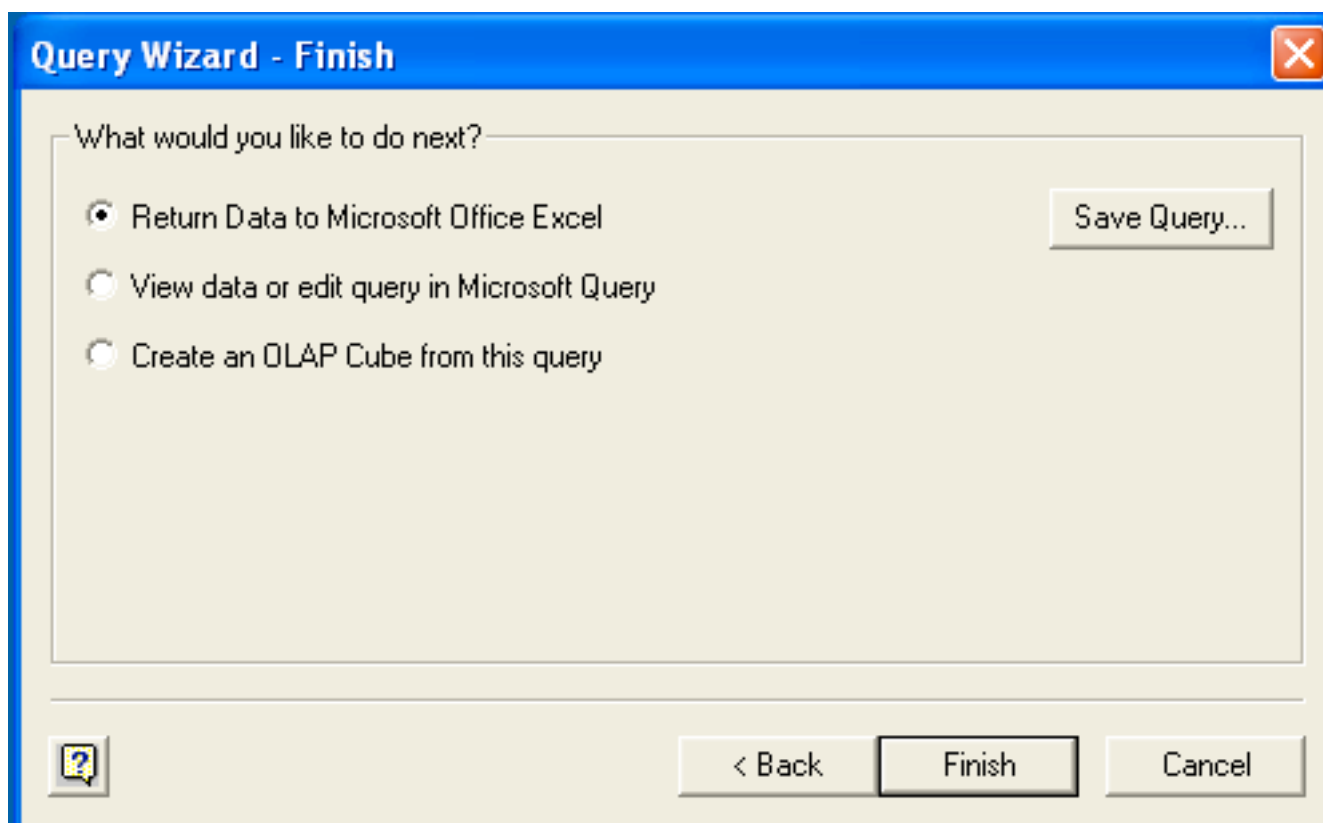
Sort by  
Name  Ascending  Descending

Then by  
 Ascending  Descending

Then by  
 Ascending  Descending

 < Back Next > Cancel

7. Select the destination for your query. You can select to return the data Microsoft Excel, where you can choose a worksheet and cell where the data will be inserted; you can continue to view the query and results within Microsoft Query, where you can edit the SQL query and further filter and sort the information returned; or you can create an OLAP Cube from the query, which can then be used directly within Microsoft Excel. Click **Finish**.



The same process can be used to import data into a Word document, where the data will be inserted as a table. This can be used for mail merge purposes (where the field data is read from a Word table), or where you want to include data and reports within a report or other document.

### 21.1.5.6. Using Connector/ODBC with Crystal Reports

Crystal Reports can use an ODBC DSN to connect to a database from which you to extract data and information for reporting purposes.



#### Note

There is a known issue with certain versions of Crystal Reports where the application is unable to open and browse tables and fields through an ODBC connection. Before using Crystal Reports with MySQL, please ensure that you have update to the latest version, including any outstanding service packs and hotfixes. For more information on this issue, see the [Business\) Objects Knowledgebase](#) for more information.

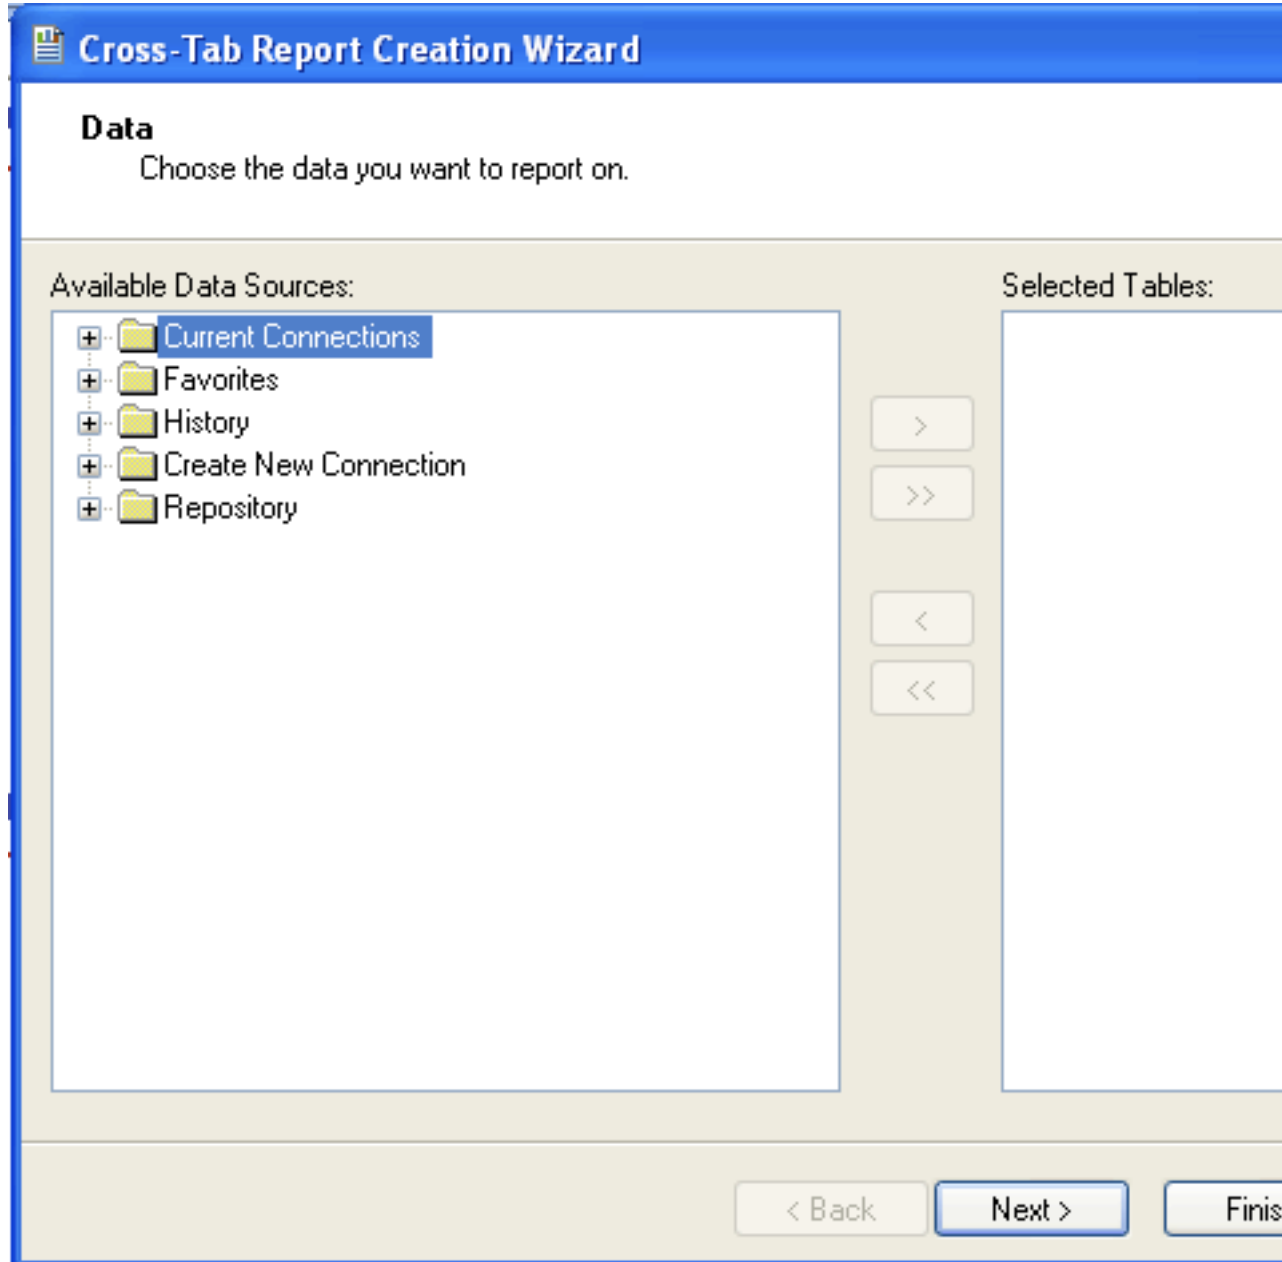
For example, to create a simple crosstab report within Crystal Reports XI, follow these steps:

1. Create a DSN using the [Data Sources \(ODBC\)](#) tool. You can either specify a complete database, including user name and password, or you can build a basic DSN and use Crystal Reports to set the user name and password.

For the purposes of this example, a DSN that provides a connection to an instance of the MySQL Sakila sample database has been created.

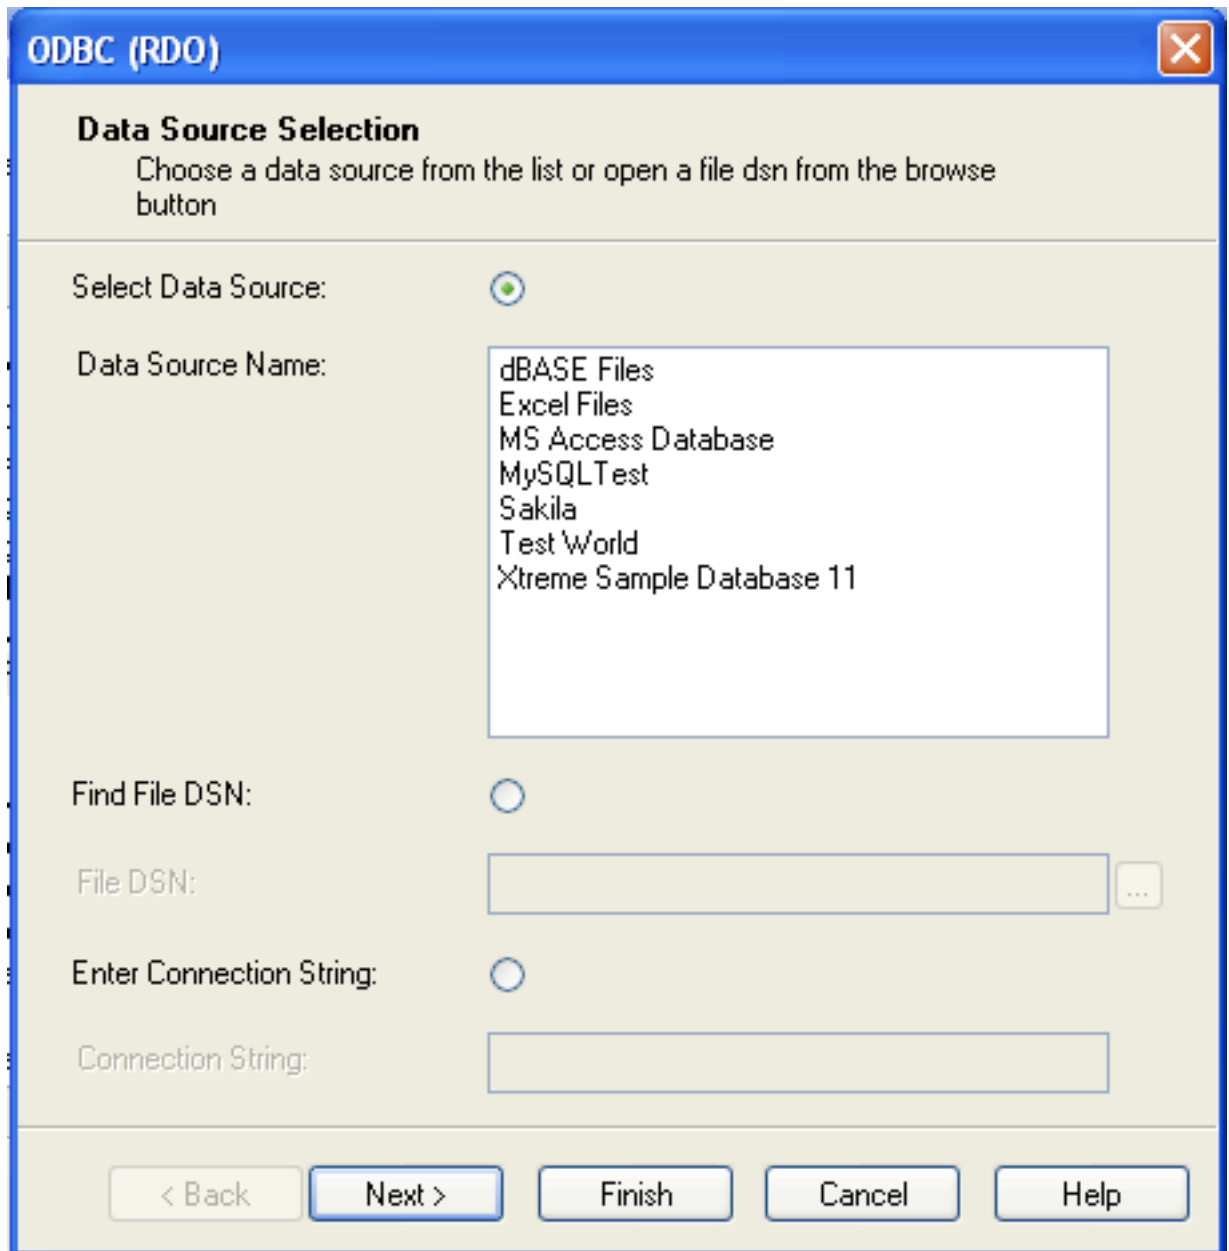
2. Open Crystal Reports and create a new project, or an open an existing reporting project into which you want to insert data from your MySQL data source.
3. Start the Cross-Tab Report Wizard, either by clicking the option on the Start Page. Expand the **Create New Connection** folder, then expand the **ODBC (RDO)** folder to obtain a list of ODBC data sources.

You will be asked to select a data source.



4. When you first expand the **ODBC (RDO)** folder you will be presented the Data Source Selection screen. From here you can select either a pre-configured DSN, open a file-based DSN or enter and manual connection string. For this example, the **Sakila** DSN will be used.

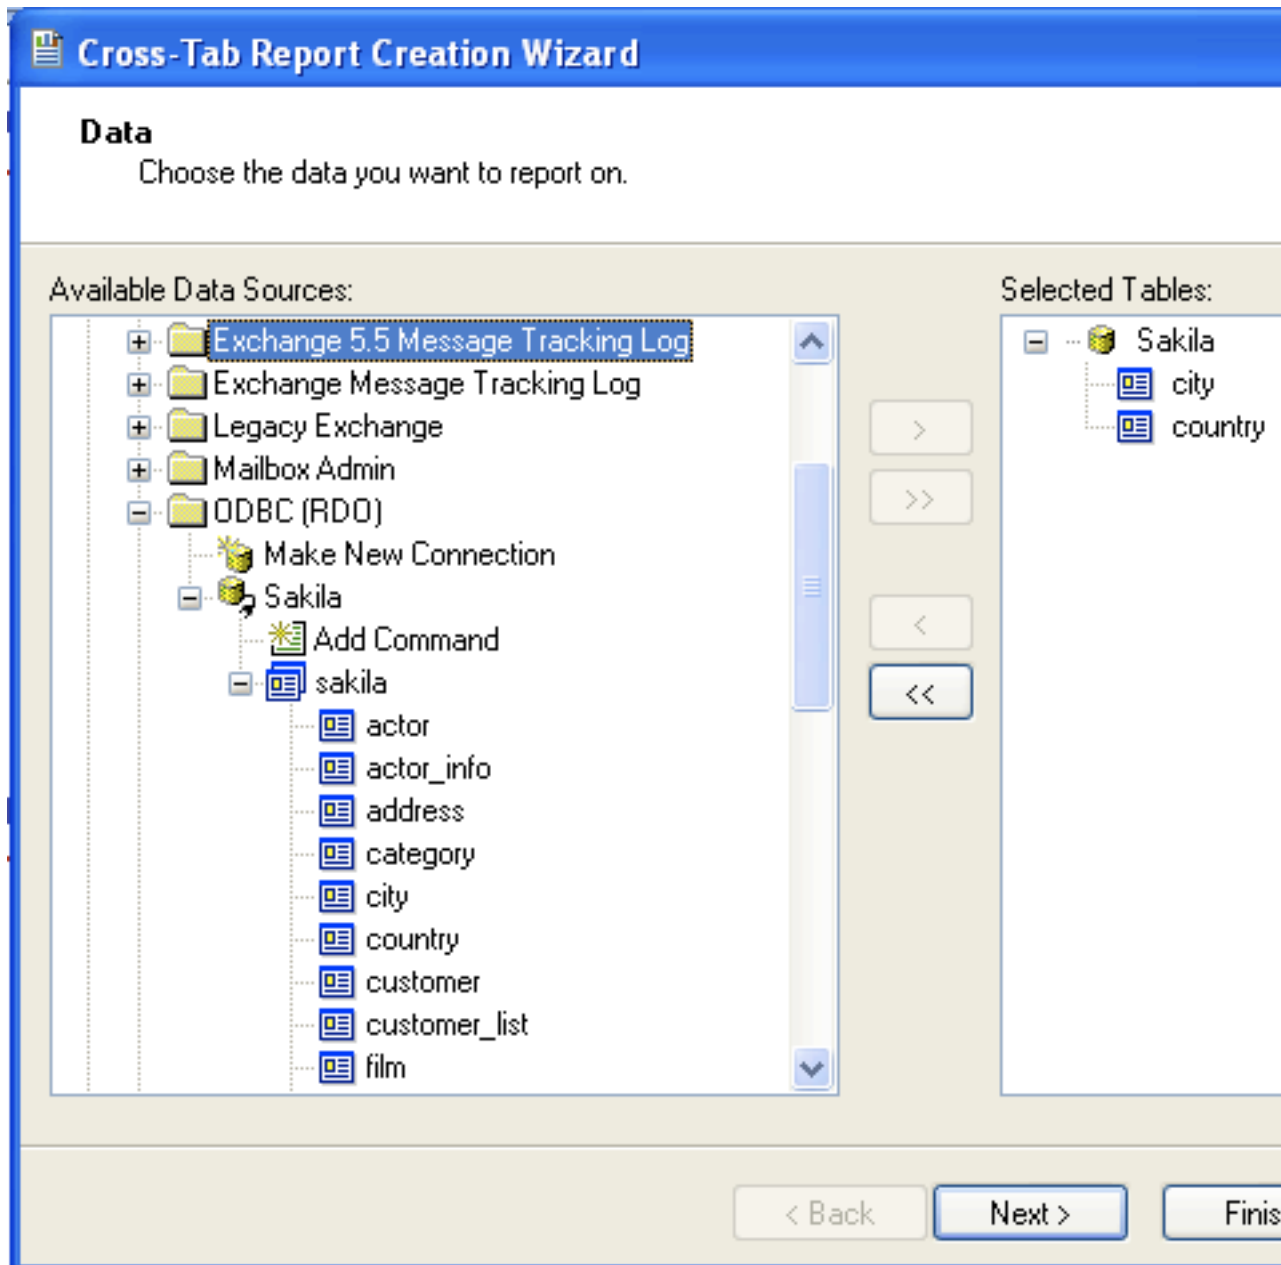
If the DSN contains a user name/password combination, or you want to use different authentication credentials, click **Next** to enter the user name and password that you want to use. Otherwise, click **Finish** to continue the data source selection wizard.



5. You will be returned the Cross-Tab Report Creation Wizard. You now need to select the database and tables that you want to include in your report. For our example, we will expand the selected Sakila database. Click the `city` table and use the `>` button to add the table to the report. Then repeat the action with the `country` table. Alternatively you can select multiple tables and add them to the report.

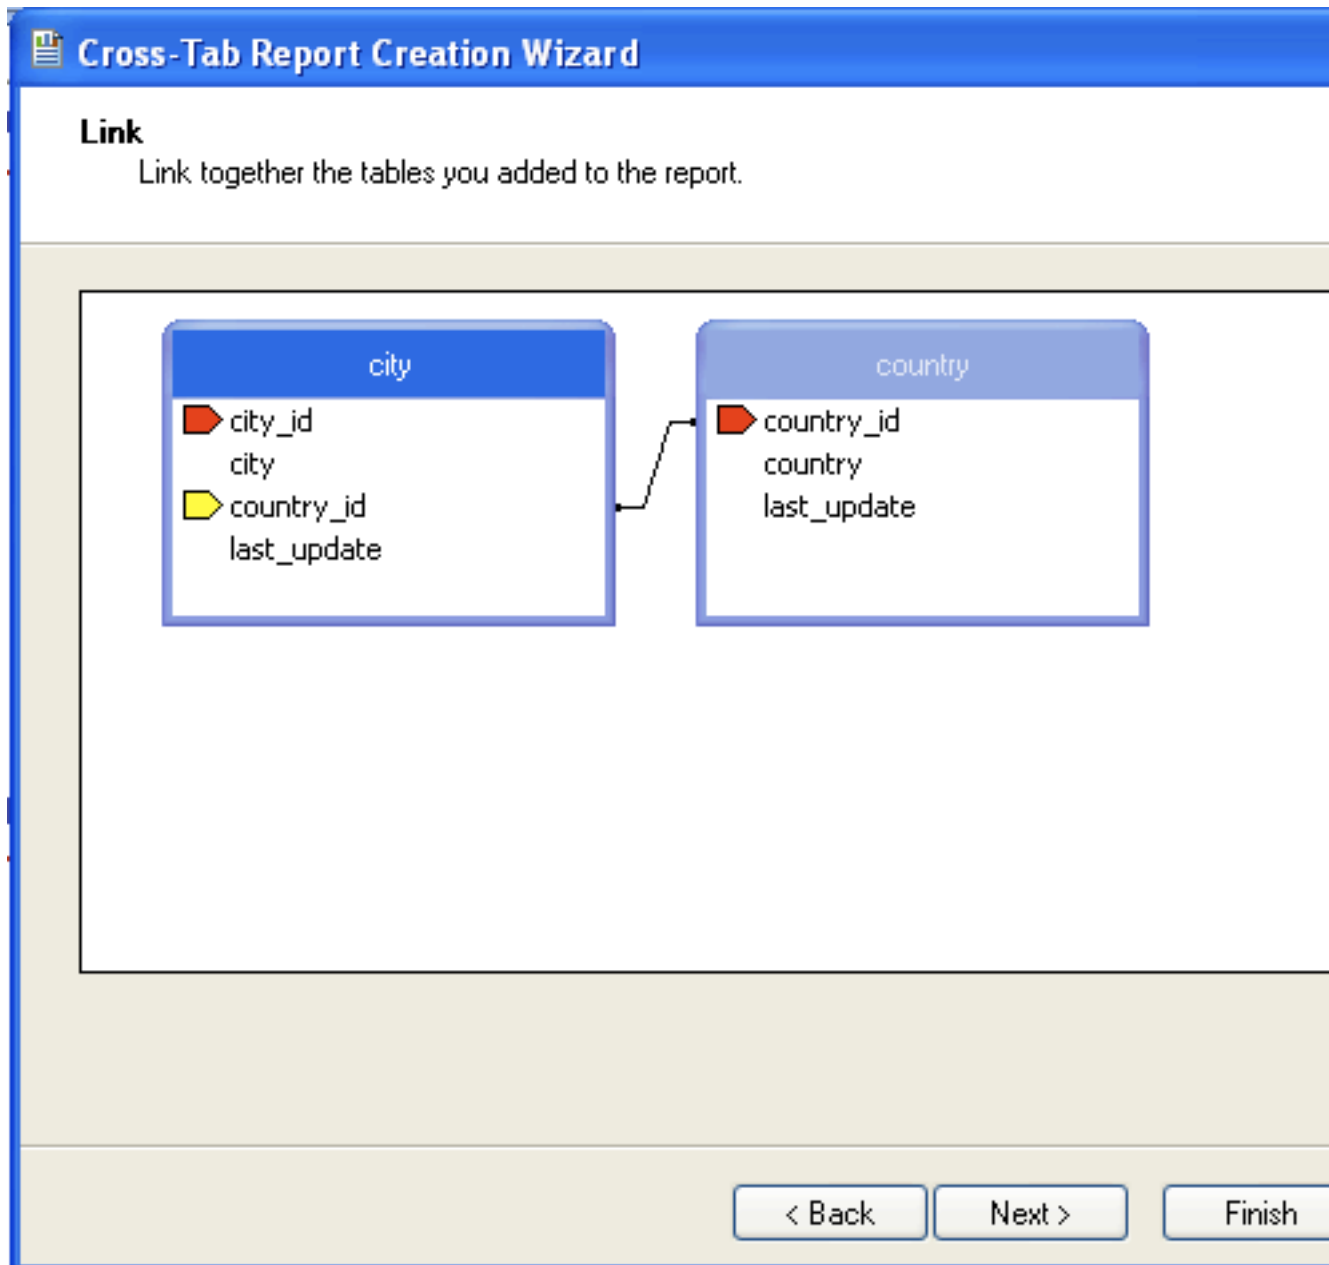
Finally, you can select the parent **Sakila** resource and add of the tables to the report.

Once you have selected the tables you want to include, click **Next** to continue.



6. Crystal Reports will now read the table definitions and automatically identify the links between the tables. The identification of links between tables enables Crystal Reports to automatically lookup and summarize information based on all the tables in the database according to your query. If Crystal Reports is unable to perform the linking itself, you can manually create the links between fields in the tables you have selected.

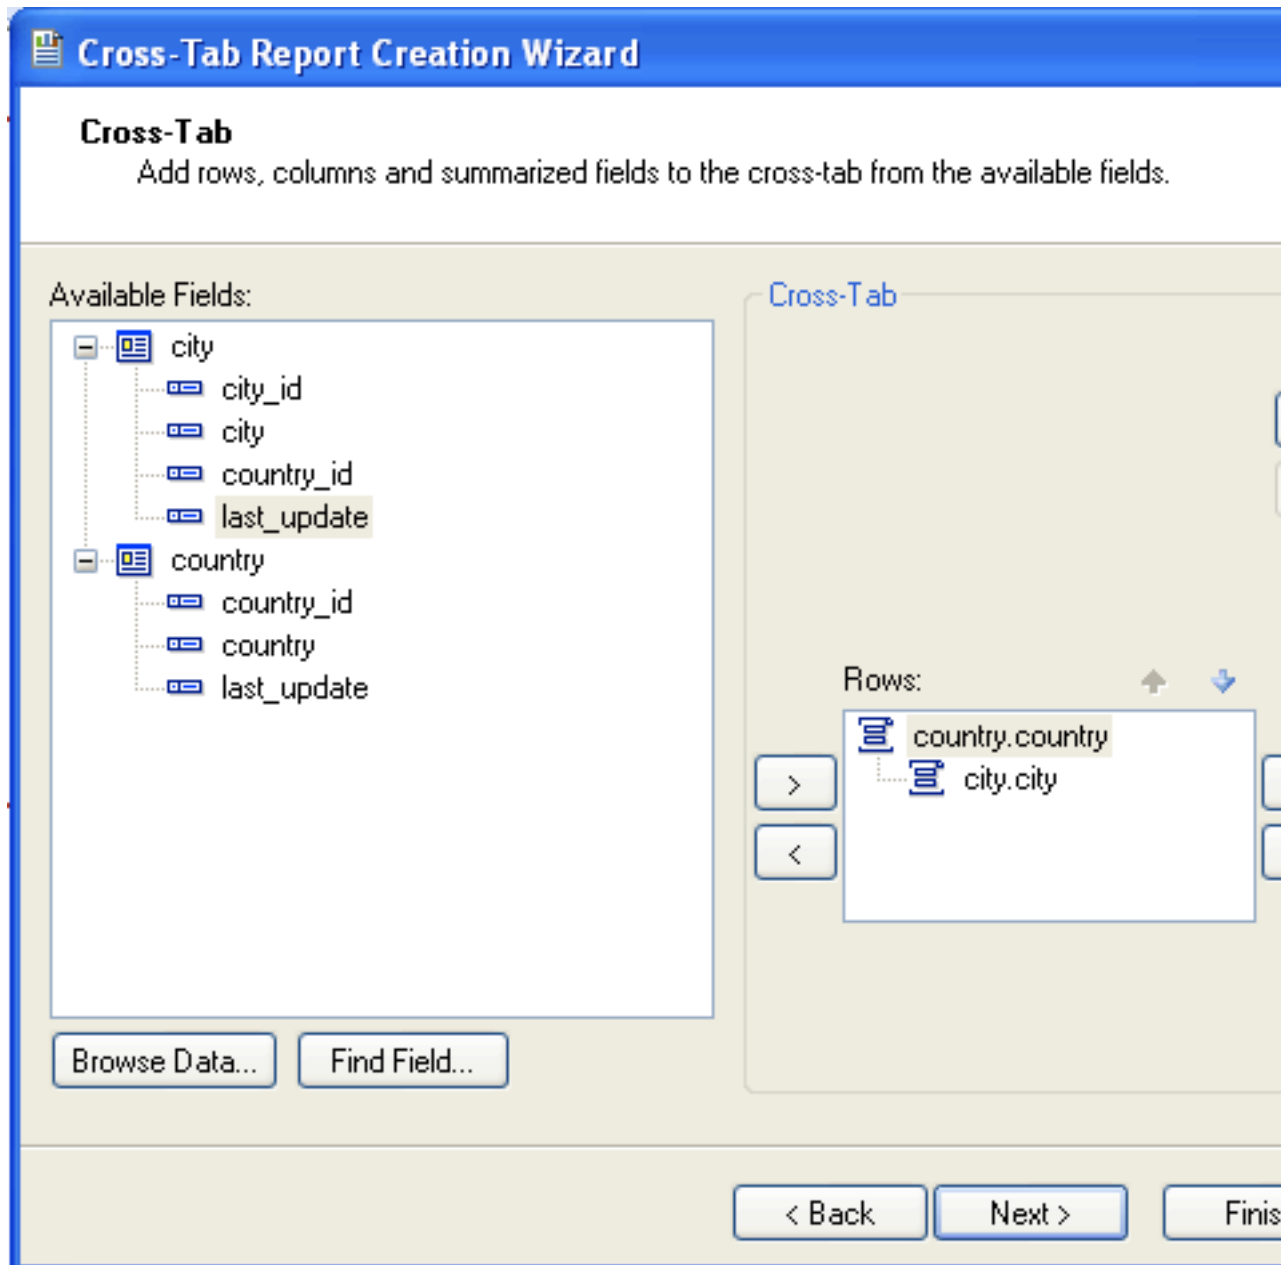
Click **Next** to continue the process.



7. You can now select the columns and rows that to include within the Cross-Tab report. Drag and drop or use the **>** buttons to add fields to each area of the report. In the example shown, we will report on cities, organized by country, incorporating a count of the number of cities within each country. If you want to browse the data, select a field and click the **Browse Data...** button.

Click **Next** to create a graph of the results. Since we are not creating a graph from this data, click **Finish** to generate the report.





8. The finished report will be shown, a sample of the output from the Sakila sample database is shown below.

		Total
<b>Total</b>		600
<b>Afghanistan</b>	<b>Total</b>	1
	<b>Kabul</b>	1
<b>Algeria</b>	<b>Total</b>	3
	<b>Batna</b>	1
	<b>Bchar</b>	1
	<b>Skikda</b>	1
<b>American Samoa</b>	<b>Total</b>	1
	<b>Tafuna</b>	1
<b>Angola</b>	<b>Total</b>	2
	<b>Benguela</b>	1
	<b>Namibe</b>	1
<b>Anguilla</b>	<b>Total</b>	1
	<b>South Hill</b>	1
<b>Argentina</b>	<b>Total</b>	13
	<b>Almirante Brow</b>	1

Once the ODBC connection has been opened within Crystal Reports, you can browse and add any fields within the available tables into your reports.

### 21.1.5.7. Connector/ODBC Programming

With a suitable ODBC Manager and the Connector/ODBC driver installed, any programming language or environment that can support ODBC can connect to a MySQL database through Connector/ODBC.

This includes, but is not limited to, Microsoft support languages (including Visual Basic, C# and interfaces such as ODBC.NET), Perl (through the DBI module, and the DBD::ODBC driver).

#### 21.1.5.7.1. Using Connector/ODBC with Visual Basic Using ADO, DAO and RDO

This section contains simple examples of the use of Connector/ODBC with ADO, DAO and RDO.

### 21.1.5.7.1.1. ADO: `rs.addNew`, `rs.delete`, and `rs.update`

The following ADO (ActiveX Data Objects) example creates a table `my_ado` and demonstrates the use of `rs.addNew`, `rs.delete`, and `rs.update`.

```
Private Sub myodbc_ado_Click()

Dim conn As ADODB.Connection
Dim rs As ADODB.Recordset
Dim fld As ADODB.Field
Dim sql As String

'connect to MySQL server using Connector/ODBC
Set conn = New ADODB.Connection
conn.ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"

conn.Open

'create table
conn.Execute "DROP TABLE IF EXISTS my_ado"
conn.Execute "CREATE TABLE my_ado(id int not null primary key, name varchar(20)," _
& "txt text, dt date, tm time, ts timestamp)"

'direct insert
conn.Execute "INSERT INTO my_ado(id,name,txt) values(1,100,'venu')"
conn.Execute "INSERT INTO my_ado(id,name,txt) values(2,200,'MySQL')"
conn.Execute "INSERT INTO my_ado(id,name,txt) values(3,300,'Delete')"

Set rs = New ADODB.Recordset
rs.CursorLocation = adUseServer

'fetch the initial table ..
rs.Open "SELECT * FROM my_ado", conn
Debug.Print rs.RecordCount
rs.MoveFirst
Debug.Print String(50, "-") & "Initial my_ado Result Set " & String(50, "-")
For Each fld In rs.Fields
Debug.Print fld.Name,
Next
Debug.Print

Do Until rs.EOF
For Each fld In rs.Fields
Debug.Print fld.Value,
Next
rs.MoveNext
Debug.Print
Loop
rs.Close

'rs insert
rs.Open "select * from my_ado", conn, adOpenDynamic, adLockOptimistic
rs.AddNew
rs!Name = "Monty"
rs!txt = "Insert row"
rs.Update
rs.Close

'rs update
rs.Open "SELECT * FROM my_ado"
rs!Name = "update"
rs!txt = "updated-row"
rs.Update
rs.Close

'rs update second time..
rs.Open "SELECT * FROM my_ado"
rs!Name = "update"
rs!txt = "updated-second-time"
```

```

rs.Update
rs.Close

'rs delete
rs.Open "SELECT * FROM my_ado"
rs.MoveNext
rs.MoveNext
rs.Delete
rs.Close

'fetch the updated table ..
rs.Open "SELECT * FROM my_ado", conn
Debug.Print rs.RecordCount
rs.MoveFirst
Debug.Print String(50, "-") & "Updated my_ado Result Set " & String(50, "-")
For Each fld In rs.Fields
Debug.Print fld.Name,
Next
Debug.Print

Do Until rs.EOF
For Each fld In rs.Fields
Debug.Print fld.Value,
Next
rs.MoveNext
Debug.Print
Loop
rs.Close
conn.Close
End Sub

```

#### 21.1.5.7.1.2. DAO: `rs.addNew`, `rs.update`, and Scrolling

The following DAO (Data Access Objects) example creates a table `my_dao` and demonstrates the use of `rs.addNew`, `rs.update`, and result set scrolling.

```

Private Sub myodbc_dao_Click()

Dim ws As Workspace
Dim conn As Connection
Dim queryDef As queryDef
Dim str As String

'connect to MySQL using MySQL ODBC 3.51 Driver
Set ws = DBEngine.CreateWorkspace("", "venu", "venu", dbUseODBC)
str = "odbc;DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"
Set conn = ws.OpenConnection("test", dbDriverNoPrompt, False, str)

'Create table my_dao
Set queryDef = conn.CreateQueryDef("", "drop table if exists my_dao")
queryDef.Execute

Set queryDef = conn.CreateQueryDef("", "create table my_dao(Id INT AUTO_INCREMENT PRIMARY KEY, " _
& "Ts TIMESTAMP(14) NOT NULL, Name varchar(20), Id2 INT)")
queryDef.Execute

'Insert new records using rs.addNew
Set rs = conn.OpenRecordset("my_dao")
Dim i As Integer

For i = 10 To 15
rs.AddNew
rs!Name = "insert record" & i
rs!Id2 = i
rs.Update
Next i
rs.Close

'rs update..

```

```

Set rs = conn.OpenRecordset("my_dao")
rs.Edit
rs!Name = "updated-string"
rs.Update
rs.Close

'fetch the table back...
Set rs = conn.OpenRecordset("my_dao", dbOpenDynamic)
str = "Results:"
rs.MoveFirst
While Not rs.EOF
str = " " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print "DATA:" & str
rs.MoveNext
Wend

'rs Scrolling
rs.MoveFirst
str = " FIRST ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

rs.MoveLast
str = " LAST ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

rs.MovePrevious
str = " LAST-1 ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

'free all resources
rs.Close
queryDef.Close
conn.Close
ws.Close

End Sub

```

### 21.1.5.7.1.3. RDO: `rs.addNew` and `rs.update`

The following RDO (Remote Data Objects) example creates a table `my_rdo` and demonstrates the use of `rs.addNew` and `rs.update`.

```

Dim rs As rdoResultset
Dim cn As New rdoConnection
Dim cl As rdoColumn
Dim SQL As String

'cn.Connect = "DSN=test;"
cn.Connect = "DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"

cn.CursorDriver = rdUseOdbc
cn.EstablishConnection rdDriverPrompt

'drop table my_rdo
SQL = "drop table if exists my_rdo"
cn.Execute SQL, rdExecDirect

'create table my_rdo
SQL = "create table my_rdo(id int, name varchar(20))"
cn.Execute SQL, rdExecDirect

'insert - direct
SQL = "insert into my_rdo values (100,'venu')"
cn.Execute SQL, rdExecDirect

SQL = "insert into my_rdo values (200,'MySQL')"
cn.Execute SQL, rdExecDirect

```

```

'rs insert
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.AddNew
rs!id = 300
rs!Name = "Insert1"
rs.Update
rs.Close

'rs insert
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.AddNew
rs!id = 400
rs!Name = "Insert 2"
rs.Update
rs.Close

'rs update
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.Edit
rs!id = 999
rs!Name = "updated"
rs.Update
rs.Close

'fetch back...
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
Do Until rs.EOF
For Each cl In rs.rdoColumns
Debug.Print cl.Value,
Next
rs.MoveNext
Debug.Print
Loop
Debug.Print "Row count="; rs.RowCount

'close
rs.Close
cn.Close

End Sub

```

### 21.1.5.7.2. Using Connector/ODBC with .NET

This section contains simple examples that demonstrate the use of Connector/ODBC drivers with ODBC.NET.

#### 21.1.5.7.2.1. Using Connector/ODBC with ODBC.NET and C# (C sharp)

The following sample creates a table `my_odbc_net` and demonstrates its use in C#.

```

/**
 * @sample      : mycon.cs
 * @purpose     : Demo sample for ODBC.NET using Connector/ODBC
 * @author      : Venu, <myodbc@lists.mysql.com>
 *
 * (C) Copyright MySQL AB, 1995-2006
 *
 **/

/* build command
 *
 * csc /t:exe
 *      /out:mycon.exe mycon.cs
 *      /r:Microsoft.Data.Odbc.dll
 */

using Console = System.Console;
using Microsoft.Data.Odbc;

```

```
namespace myodbc3
{
    class mycon
    {
        static void Main(string[] args)
        {
            try
            {
                //Connection string for Connector/ODBC 3.51
                string MyConString = "DRIVER={MySQL ODBC 3.51 Driver};" +
                    "SERVER=localhost;" +
                    "DATABASE=test;" +
                    "UID=venu;" +
                    "PASSWORD=venu;" +
                    "OPTION=3";

                //Connect to MySQL using Connector/ODBC
                OdbcConnection MyConnection = new OdbcConnection(MyConString);
                MyConnection.Open();

                Console.WriteLine("\n !!! success, connected successfully !!!\n");

                //Display connection information
                Console.WriteLine("Connection Information:");
                Console.WriteLine("\tConnection String:" +
                    MyConnection.ConnectionString);
                Console.WriteLine("\tConnection Timeout:" +
                    MyConnection.ConnectionTimeout);
                Console.WriteLine("\tDatabase:" +
                    MyConnection.Database);
                Console.WriteLine("\tDataSource:" +
                    MyConnection.DataSource);
                Console.WriteLine("\tDriver:" +
                    MyConnection.Driver);
                Console.WriteLine("\tServerVersion:" +
                    MyConnection.ServerVersion);

                //Create a sample table
                OdbcCommand MyCommand =
                    new OdbcCommand("DROP TABLE IF EXISTS my_odbc_net",
                        MyConnection);
                MyCommand.ExecuteNonQuery();
                MyCommand.CommandText =
                    "CREATE TABLE my_odbc_net(id int, name varchar(20), idb bigint)";
                MyCommand.ExecuteNonQuery();

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(10,'venu', 300)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(20,'mysql',400)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(20,'mysql',500)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Update
                MyCommand.CommandText =
                    "UPDATE my_odbc_net SET id=999 WHERE id=20";
                Console.WriteLine("Update, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //COUNT(*)
            }
        }
    }
}
```





```

Module myvb
Sub Main()
Try

'Connector/ODBC 3.51 connection string
Dim MyConString As String = "DRIVER={MySQL ODBC 3.51 Driver};" & _
"SERVER=localhost;" & _
"DATABASE=test;" & _
"UID=venu;" & _
"PASSWORD=venu;" & _
"OPTION=3;"

'Connection
Dim MyConnection As New OdbcConnection(MyConString)
MyConnection.Open()

Console.WriteLine("Connection State::" & MyConnection.State.ToString)

'Drop
Console.WriteLine("Dropping table")
Dim MyCommand As New OdbcCommand()
MyCommand.Connection = MyConnection
MyCommand.CommandText = "DROP TABLE IF EXISTS my_vb_net"
MyCommand.ExecuteNonQuery()

'Create
Console.WriteLine("Creating...")
MyCommand.CommandText = "CREATE TABLE my_vb_net(id int, name varchar(30))"
MyCommand.ExecuteNonQuery()

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(10,'venu')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(20,'mysql')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(20,'mysql')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net(id) VALUES(30)"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Update
MyCommand.CommandText = "UPDATE my_vb_net SET id=999 WHERE id=20"
Console.WriteLine("Update, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'COUNT(*)
MyCommand.CommandText = "SELECT COUNT(*) as TRows FROM my_vb_net"
Console.WriteLine("Total Rows:" & MyCommand.ExecuteScalar())

'Select
Console.WriteLine("Select * FROM my_vb_net")
MyCommand.CommandText = "SELECT * FROM my_vb_net"
Dim MyDataReader As OdbcDataReader
MyDataReader = MyCommand.ExecuteReader
While MyDataReader.Read
If MyDataReader("name") Is DBNull.Value Then
Console.WriteLine("id = " & _
CStr(MyDataReader("id")) & " name = " & _
"NULL")
Else
Console.WriteLine("id = " & _

```

```

        CStr(MyDataReader("id")) & " name = " & _
        CStr(MyDataReader("name")))
    End If
End While

'Catch ODBC Exception
Catch MyOdbcException As OdbcException
    Dim i As Integer
    Console.WriteLine(MyOdbcException.ToString)

'Catch program exception
Catch MyException As Exception
    Console.WriteLine(MyException.ToString)
End Try
End Sub

```

## 21.1.6. Connector/ODBC Reference

This section provides reference material for the Connector/ODBC API, showing supported functions and methods, supported MySQL column types and the corresponding native type in Connector/ODBC, and the error codes returned by Connector/ODBC when a fault occurs.

### 21.1.6.1. Connector/ODBC API Reference

This section summarizes ODBC routines, categorized by functionality.

For the complete ODBC API reference, please refer to the ODBC Programmer's Reference at <http://msdn.microsoft.com/en-us/library/ms714177.aspx>.

An application can call [SQLGetInfo](#) function to obtain conformance information about Connector/ODBC. To obtain information about support for a specific function in the driver, an application can call [SQLGetFunctions](#).



#### Note

For backward compatibility, the Connector/ODBC driver supports all deprecated functions.

The following tables list Connector/ODBC API calls grouped by task:

**Table 21.6. ODBC API Calls for Connecting to a Data Source**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<a href="#">SQLAllocHandle</a>	Yes	ISO 92	Obtains an environment, connection, statement, or descriptor handle.
<a href="#">SQLConnect</a>	Yes	ISO 92	Connects to a specific driver by data source name, user ID, and password.
<a href="#">SQLDriverConnect</a>	Yes	ODBC	Connects to a specific driver by connection string or requests that the Driver Manager and driver display connection dialog boxes for the user.
<a href="#">SQLAllocEnv</a>	Yes	Deprecated	Obtains an environment handle allocated from driver.
<a href="#">SQLAllocConnect</a>	Yes	Deprecated	Obtains a connection handle

**Table 21.7. ODBC API Calls for Obtaining Information about a Driver and Data Source**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<a href="#">SQLDataSources</a>	No	ISO 92	Returns the list of available data sources, handled by the Driver Manager
<a href="#">SQLDrivers</a>	No	ODBC	Returns the list of installed drivers and their attributes, handles by Driver Manager
<a href="#">SQLGetInfo</a>	Yes	ISO 92	Returns information about a specific driver and data source.
<a href="#">SQLGetFunctions</a>	Yes	ISO 92	Returns supported driver functions.
<a href="#">SQLGetTypeInfo</a>	Yes	ISO 92	Returns information about supported data types.

**Table 21.8. ODBC API Calls for Setting and Retrieving Driver Attributes**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<a href="#">SQLSetConnectAttr</a>	Yes	ISO 92	Sets a connection attribute.
<a href="#">SQLGetConnectAttr</a>	Yes	ISO 92	Returns the value of a connection attribute.
<a href="#">SQLSetConnectOption</a>	Yes	Deprecated	Sets a connection option
<a href="#">SQLGetConnectOption</a>	Yes	Deprecated	Returns the value of a connection option
<a href="#">SQLSetEnvAttr</a>	Yes	ISO 92	Sets an environment attribute.
<a href="#">SQLGetEnvAttr</a>	Yes	ISO 92	Returns the value of an environment attribute.
<a href="#">SQLSetStmtAttr</a>	Yes	ISO 92	Sets a statement attribute.
<a href="#">SQLGetStmtAttr</a>	Yes	ISO 92	Returns the value of a statement attribute.
<a href="#">SQLSetStmtOption</a>	Yes	Deprecated	Sets a statement option
<a href="#">SQLGetStmtOption</a>	Yes	Deprecated	Returns the value of a statement option

**Table 21.9. ODBC API Calls for Preparing SQL Requests**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<a href="#">SQLAllocStmt</a>	Yes	Deprecated	Allocates a statement handle
<a href="#">SQLPrepare</a>	Yes	ISO 92	Prepares an SQL statement for later execution.
<a href="#">SQLBindParameter</a>	Yes	ODBC	Assigns storage for a parameter in an SQL statement. Connector/ODBC 5.2 adds support for “out” and “inout” parameters, through the <a href="#">SQL_PARAM_OUTPUT</a> or <a href="#">SQL_PARAM_INPUT_OUTPUT</a> type specifiers. (“Out” and “inout” parameters are not supported for <a href="#">LONGTEXT</a> and <a href="#">LONGBLOB</a> columns.)
<a href="#">SQLGetCursorName</a>	Yes	ISO 92	Returns the cursor name associated with a statement handle.
<a href="#">SQLSetCursorName</a>	Yes	ISO 92	Specifies a cursor name.
<a href="#">SQLSetScrollOptions</a>	Yes	ODBC	Sets options that control cursor behavior.

**Table 21.10. ODBC API Calls for Submitting Requests**

Function Name	Connector/ ODBC Supports?	Standard	Purpose
<code>SQLExecute</code>	Yes	ISO 92	Executes a prepared statement.
<code>SQLExecDirect</code>	Yes	ISO 92	Executes a statement
<code>SQLNativeSql</code>	Yes	ODBC	Returns the text of an SQL statement as translated by the driver.
<code>SQLDescribeParam</code>	Yes	ODBC	Returns the description for a specific parameter in a statement.
<code>SQLNumParams</code>	Yes	ISO 92	Returns the number of parameters in a statement.
<code>SQLParamData</code>	Yes	ISO 92	Used in conjunction with <code>SQLPutData</code> to supply parameter data at execution time. (Useful for long data values.)
<code>SQLPutData</code>	Yes	ISO 92	Sends part or all of a data value for a parameter. (Useful for long data values.)

**Table 21.11. ODBC API Calls for Retrieving Results and Information about Results**

Function Name	Connector/ ODBC Supports?	Standard	Purpose
<code>SQLRowCount</code>	Yes	ISO 92	Returns the number of rows affected by an insert, update, or delete request.
<code>SQLNumResultCols</code>	Yes	ISO 92	Returns the number of columns in the result set.
<code>SQLDescribeCol</code>	Yes	ISO 92	Describes a column in the result set.
<code>SQLColAttribute</code>	Yes	ISO 92	Describes attributes of a column in the result set.
<code>SQLColAttributes</code>	Yes	Deprecated	Describes attributes of a column in the result set.
<code>SQLFetch</code>	Yes	ISO 92	Returns multiple result rows.
<code>SQLFetchScroll</code>	Yes	ISO 92	Returns scrollable result rows.
<code>SQLExtendedFetch</code>	Yes	Deprecated	Returns scrollable result rows.
<code>SQLSetPos</code>	Yes	ODBC	Positions a cursor within a fetched block of data and enables an application to refresh data in the rowset or to update or delete data in the result set.
<code>SQLBulkOperations</code>	Yes	ODBC	Performs bulk insertions and bulk bookmark operations, including update, delete, and fetch by bookmark.

**Table 21.12. ODBC API Calls for Retrieving Error or Diagnostic Information**

Function Name	Connector/ ODBC Supports?	Standard	Purpose
<code>SQLError</code>	Yes	Deprecated	Returns additional error or status information
<code>SQLGetDiagField</code>	Yes	ISO 92	Returns additional diagnostic information (a single field of the diagnostic data structure).
<code>SQLGetDiagRec</code>	Yes	ISO 92	Returns additional diagnostic information (multiple fields of the diagnostic data structure).

**Table 21.13. ODBC API Calls for Obtaining Information about the Data Source's System Tables (Catalog Functions) Item**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<code>SQLColumnPrivileges</code>	Yes	ODBC	Returns a list of columns and associated privileges for one or more tables.
<code>SQLColumns</code>	Yes	X/Open	Returns the list of column names in specified tables.
<code>SQLForeignKeys</code>	Yes	ODBC	Returns a list of column names that make up foreign keys, if they exist for a specified table.
<code>SQLPrimaryKeys</code>	Yes	ODBC	Returns the list of column names that make up the primary key for a table.
<code>SQLSpecialColumns</code>	Yes	X/Open	Returns information about the optimal set of columns that uniquely identifies a row in a specified table, or the columns that are automatically updated when any value in the row is updated by a transaction.
<code>SQLStatistics</code>	Yes	ISO 92	Returns statistics about a single table and the list of indexes associated with the table.
<code>SQLTablePrivileges</code>	Yes	ODBC	Returns a list of tables and the privileges associated with each table.
<code>SQLTables</code>	Yes	X/Open	Returns the list of table names stored in a specific data source.

**Table 21.14. ODBC API Calls for Performing Transactions**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<code>SQLTransact</code>	Yes	Deprecated	Commits or rolls back a transaction
<code>SQLEndTran</code>	Yes	ISO 92	<a href="#">Commits</a> or <a href="#">rolls back</a> a <a href="#">transaction</a> .

**Table 21.15. ODBC API Calls for Terminating a Statement**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<code>SQLFreeStmt</code>	Yes	ISO 92	Ends statement processing, discards pending results, and, optionally, frees all resources associated with the statement handle.
<code>SQLCloseCursor</code>	Yes	ISO 92	Closes a cursor that has been opened on a statement handle.
<code>SQLCancel</code>	Yes	ISO 92	Cancels an SQL statement.

**Table 21.16. ODBC API Calls for Terminating a Connection**

Function Name	Connector/ODBC Supports?	Standard	Purpose
<code>SQLDisconnect</code>	Yes	ISO 92	Closes the connection.
<code>SQLFreeHandle</code>	Yes	ISO 92	Releases an environment, connection, statement, or descriptor handle.

Function Name	Connector/ODBC Supports?	Standard	Purpose
<code>SQLFreeConnect</code>	Yes	Deprecated	Releases connection handle.
<code>SQLFreeEnv</code>	Yes	Deprecated	Releases an environment handle.

### 21.1.6.2. Connector/ODBC Data Types

The following table illustrates how Connector/ODBC maps the server data types to default SQL and C data types.

**Table 21.17. How Connector/ODBC Maps MySQL Data Types to SQL and C Data Types**

Native Value	SQL Type	C Type
<code>bigint unsigned</code>	<code>SQL_BIGINT</code>	<code>SQL_C_UBIGINT</code>
<code>bigint</code>	<code>SQL_BIGINT</code>	<code>SQL_C_SBIGINT</code>
<code>bit</code>	<code>SQL_BIT</code>	<code>SQL_C_BIT</code>
<code>bit</code>	<code>SQL_CHAR</code>	<code>SQL_C_CHAR</code>
<code>blob</code>	<code>SQL_LONGVARBINARY</code>	<code>SQL_C_BINARY</code>
<code>bool</code>	<code>SQL_CHAR</code>	<code>SQL_C_CHAR</code>
<code>char</code>	<code>SQL_CHAR</code>	<code>SQL_C_CHAR</code>
<code>date</code>	<code>SQL_DATE</code>	<code>SQL_C_DATE</code>
<code>datetime</code>	<code>SQL_TIMESTAMP</code>	<code>SQL_C_TIMESTAMP</code>
<code>decimal</code>	<code>SQL_DECIMAL</code>	<code>SQL_C_CHAR</code>
<code>double precision</code>	<code>SQL_DOUBLE</code>	<code>SQL_C_DOUBLE</code>
<code>double</code>	<code>SQL_FLOAT</code>	<code>SQL_C_DOUBLE</code>
<code>enum</code>	<code>SQL_VARCHAR</code>	<code>SQL_C_CHAR</code>
<code>float</code>	<code>SQL_REAL</code>	<code>SQL_C_FLOAT</code>
<code>int unsigned</code>	<code>SQL_INTEGER</code>	<code>SQL_C_ULONG</code>
<code>int</code>	<code>SQL_INTEGER</code>	<code>SQL_C_SLONG</code>
<code>integer unsigned</code>	<code>SQL_INTEGER</code>	<code>SQL_C_ULONG</code>
<code>integer</code>	<code>SQL_INTEGER</code>	<code>SQL_C_SLONG</code>
<code>long varbinary</code>	<code>SQL_LONGVARBINARY</code>	<code>SQL_C_BINARY</code>
<code>long varchar</code>	<code>SQL_LONGVARCHAR</code>	<code>SQL_C_CHAR</code>
<code>longblob</code>	<code>SQL_LONGVARBINARY</code>	<code>SQL_C_BINARY</code>
<code>longtext</code>	<code>SQL_LONGVARCHAR</code>	<code>SQL_C_CHAR</code>
<code>mediumblob</code>	<code>SQL_LONGVARBINARY</code>	<code>SQL_C_BINARY</code>
<code>mediumint unsigned</code>	<code>SQL_INTEGER</code>	<code>SQL_C_ULONG</code>
<code>mediumint</code>	<code>SQL_INTEGER</code>	<code>SQL_C_SLONG</code>
<code>mediumtext</code>	<code>SQL_LONGVARCHAR</code>	<code>SQL_C_CHAR</code>
<code>numeric</code>	<code>SQL_NUMERIC</code>	<code>SQL_C_CHAR</code>
<code>real</code>	<code>SQL_FLOAT</code>	<code>SQL_C_DOUBLE</code>
<code>set</code>	<code>SQL_VARCHAR</code>	<code>SQL_C_CHAR</code>
<code>smallint unsigned</code>	<code>SQL_SMALLINT</code>	<code>SQL_C_USHORT</code>
<code>smallint</code>	<code>SQL_SMALLINT</code>	<code>SQL_C_SSHORT</code>
<code>text</code>	<code>SQL_LONGVARCHAR</code>	<code>SQL_C_CHAR</code>



Native Value	SQL Type	C Type
time	SQL_TIME	SQL_C_TIME
timestamp [558]	SQL_TIMESTAMP	SQL_C_TIMESTAMP
tinyblob	SQL_LONGVARBINARY	SQL_C_BINARY
tinyint unsigned	SQL_TINYINT	SQL_C_UTINYINT
tinyint	SQL_TINYINT	SQL_C_STINYINT
tinytext	SQL_LONGVARCHAR	SQL_C_CHAR
varchar	SQL_VARCHAR	SQL_C_CHAR
year	SQL_SMALLINT	SQL_C_SHORT

### 21.1.6.3. Connector/ODBC Error Codes

The following tables lists the error codes returned by Connector/ODBC apart from the server errors.

**Table 21.18. Special Error Codes Returned by Connector/ODBC**

Native Code	SQLSTATE 2	SQLSTATE 3	Error Message
500	01000	01000	General warning
501	01004	01004	String data, right truncated
502	01S02	01S02	Option value changed
503	01S03	01S03	No rows updated/deleted
504	01S04	01S04	More than one row updated/deleted
505	01S06	01S06	Attempt to fetch before the result set returned the first row set
506	07001	07002	<code>SQLBindParameter</code> not used for all parameters
507	07005	07005	Prepared statement not a cursor-specification
508	07009	07009	Invalid descriptor index
509	08002	08002	Connection name in use
510	08003	08003	Connection does not exist
511	24000	24000	Invalid cursor state
512	25000	25000	Invalid transaction state
513	25S01	25S01	Transaction state unknown
514	34000	34000	Invalid cursor name
515	S1000	HY000	General driver defined error
516	S1001	HY001	Memory allocation error
517	S1002	HY002	Invalid column number
518	S1003	HY003	Invalid application buffer type
519	S1004	HY004	Invalid SQL data type
520	S1009	HY009	Invalid use of null pointer
521	S1010	HY010	Function sequence error
522	S1011	HY011	Attribute can not be set now
523	S1012	HY012	Invalid transaction operation code
524	S1013	HY013	Memory management error
525	S1015	HY015	No cursor name available

Native Code	SQLSTATE 2	SQLSTATE 3	Error Message
526	S1024	HY024	Invalid attribute value
527	S1090	HY090	Invalid string or buffer length
528	S1091	HY091	Invalid descriptor field identifier
529	S1092	HY092	Invalid attribute/option identifier
530	S1093	HY093	Invalid parameter number
531	S1095	HY095	Function type out of range
532	S1106	HY106	Fetch type out of range
533	S1117	HY117	Row value out of range
534	S1109	HY109	Invalid cursor position
535	S1C00	HYC00	Optional feature not implemented
0	21S01	21S01	Column count does not match value count
0	23000	23000	Integrity constraint violation
0	42000	42000	Syntax error or access violation
0	42S02	42S02	Base table or view not found
0	42S12	42S12	Index not found
0	42S21	42S21	Column already exists
0	42S22	42S22	Column not found
0	08S01	08S01	Communication link failure

## 21.1.7. Connector/ODBC Notes and Tips

Here are some common notes and tips for using Connector/ODBC within different environments, applications and tools. The notes provided here are based on the experiences of Connector/ODBC developers and users.

### 21.1.7.1. Connector/ODBC General Functionality

This section provides help with common queries and areas of functionality in MySQL and how to use them with Connector/ODBC.

#### 21.1.7.1.1. Obtaining Auto-Increment Values

Obtaining the value of column that uses `AUTO_INCREMENT` after an `INSERT` statement can be achieved in a number of different ways. To obtain the value immediately after an `INSERT`, use a `SELECT` query with the `LAST_INSERT_ID()` [1269] function.

For example, using Connector/ODBC you would execute two separate statements, the `INSERT` statement and the `SELECT` query to obtain the auto-increment value.

```
INSERT INTO tbl (auto,text) VALUES(NULL,'text');
SELECT LAST_INSERT_ID();
```

If you do not require the value within your application, but do require the value as part of another `INSERT`, the entire process can be handled by executing the following statements:

```
INSERT INTO tbl (auto,text) VALUES(NULL,'text');
INSERT INTO tbl2 (id,text) VALUES(LAST_INSERT_ID(),'text');
```

Certain ODBC applications (including Delphi and Access) may have trouble obtaining the auto-increment value using the previous examples. In this case, try the following statement as an alternative:

```
SELECT * FROM tbl WHERE auto IS NULL;
```

This alternative method requires that `sql_auto_is_null` [543] variable is not set to 0. See [Section 5.1.4, “Server System Variables”](#).

See also [Section 21.8.14.3, “How to Get the Unique ID for the Last Inserted Row”](#).

#### 21.1.7.1.2. Dynamic Cursor Support

Support for the `dynamic cursor` is provided in Connector/ODBC 3.51, but dynamic cursors are not enabled by default. You can enable this function within Windows by selecting the `Enable Dynamic Cursor` check box within the ODBC Data Source Administrator.

On other platforms, you can enable the dynamic cursor by adding `32` to the `OPTION` value when creating the DSN.

#### 21.1.7.1.3. Connector/ODBC Performance

The Connector/ODBC driver has been optimized to provide very fast performance. If you experience problems with the performance of Connector/ODBC, or notice a large amount of disk activity for simple queries, there are a number of aspects to check:

- Ensure that `ODBC Tracing` is not enabled. With tracing enabled, a lot of information is recorded in the tracing file by the ODBC Manager. You can check, and disable, tracing within Windows using the `Tracing` panel of the ODBC Data Source Administrator. Within Mac OS X, check the `Tracing` panel of ODBC Administrator. See [Section 21.1.4.8, “Getting an ODBC Trace File”](#).
- Make sure you are using the standard version of the driver, and not the debug version. The debug version includes additional checks and reporting measures.
- Disable the Connector/ODBC driver trace and query logs. These options are enabled for each DSN, so make sure to examine only the DSN that you are using in your application. Within Windows, you can disable the Connector/ODBC and query logs by modifying the DSN configuration. Within Mac OS X and Unix, ensure that the driver trace (option value 4) and query logging (option value 524288) are not enabled.

#### 21.1.7.1.4. Setting ODBC Query Timeout in Windows

For more information on how to set the query timeout on Microsoft Windows when executing queries through an ODBC connection, read the Microsoft knowledgebase document at <http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B153756>.

### 21.1.7.2. Connector/ODBC Application-Specific Tips

Most programs should work with Connector/ODBC, but for each of those listed here, there are specific notes and tips to improve or enhance the way you work with Connector/ODBC and these applications.

With all applications, ensure that you are using the latest Connector/ODBC drivers, ODBC Manager and any supporting libraries and interfaces used by your application. For example, on Windows, using the latest version of Microsoft Data Access Components (MDAC) will improve the compatibility with ODBC in general, and with the Connector/ODBC driver.

#### 21.1.7.2.1. Using Connector/ODBC with Microsoft Applications

The majority of Microsoft applications have been tested with Connector/ODBC, including Microsoft Office, Microsoft Access and the various programming languages supported within ASP and Microsoft Visual Studio.

##### 21.1.7.2.1.1. Microsoft Access

To improve the integration between Microsoft Access and MySQL through Connector/ODBC:

- For all versions of Access, enable the Connector/ODBC `Return matching rows` option. For Access 2.0, also enable the `Simulate ODBC 1.0` option.

- Include a `TIMESTAMP` column in all tables that you want to be able to update. For maximum portability, do not use a length specification in the column declaration (which is unsupported within MySQL in versions earlier than 4.1).
- Include a `primary key` in each MySQL table you want to use with Access. If not, new or updated rows may show up as `#DELETED#`.
- Use only `DOUBLE` float fields. Access fails when comparing with single-precision floats. The symptom usually is that new or updated rows may show up as `#DELETED#` or that you cannot find or update rows.
- If you are using Connector/ODBC to link to a table that has a `BIGINT` column, the results are displayed as `#DELETED#`. The work around solution is:
  - Have one more dummy column with `TIMESTAMP` as the data type.
  - Select the `Change BIGINT columns to INT` option in the connection dialog in ODBC DSN Administrator.
  - Delete the table link from Access and re-create it.

Old records may still display as `#DELETED#`, but newly added/updated records are displayed properly.

- If you still get the error `Another user has changed your data` after adding a `TIMESTAMP` column, the following trick may help you:

Do not use a `table` data sheet view. Instead, create a form with the fields you want, and use that `form` data sheet view. Set the `DefaultValue` property for the `TIMESTAMP` column to `NOW() [1220]`. Consider hiding the `TIMESTAMP` column from view so your users are not confused.

- In some cases, Access may generate SQL statements that MySQL cannot understand. You can fix this by selecting `"Query | SQLSpecific | Pass-Through"` from the Access menu.
- On Windows NT, Access reports `BLOB` columns as `OLE OBJECTS`. If you want to have `MEMO` columns instead, change `BLOB` columns to `TEXT` with `ALTER TABLE`.
- Access cannot always handle the MySQL `DATE` column properly. If you have a problem with these, change the columns to `DATETIME`.
- If you have in Access a column defined as `BYTE`, Access tries to export this as `TINYINT` instead of `TINYINT UNSIGNED`. This gives you problems if you have values larger than 127 in the column.
- If you have very large (long) tables in Access, it might take a very long time to open them. Or you might run low on virtual memory and eventually get an `ODBC Query Failed` error and the table cannot open. To deal with this, select the following options:
  - Return Matching Rows (2)
  - Allow BIG Results (8).

These add up to a value of 10 (`OPTION=10`).

Some external articles and tips that may be useful when using Access, ODBC and Connector/ODBC:

- Read [How to Trap ODBC Login Error Messages in Access](#)
- Optimizing Access ODBC Applications
  - [Optimizing for Client/Server Performance](#)
  - [Tips for Converting Applications to Using ODBCDirect](#)

- [Tips for Optimizing Queries on Attached SQL Tables](#)
- For a list of tools that can be used with Access and ODBC data sources, refer to <http://www.mysql.com/portal/software/convertors/> section for list of available tools.

#### 21.1.7.2.1.2. Microsoft Excel and Column Types

If you have problems importing data into Microsoft Excel, particularly numeric, date, and time values, this is probably because of a bug in Excel, where the column type of the source data is used to determine the data type when that data is inserted into a cell within the worksheet. The result is that Excel incorrectly identifies the content and this affects both the display format and the data when it is used within calculations.

To address this issue, use the `CONCAT()` [1178] function in your queries. The use of `CONCAT()` [1178] forces Excel to treat the value as a string, which Excel will then parse and usually correctly identify the embedded information.

However, even with this option, some data may be incorrectly formatted, even though the source data remains unchanged. Use the `Format Cells` option within Excel to change the format of the displayed information.

#### 21.1.7.2.1.3. Microsoft Visual Basic

To be able to update a table, you must define a `primary key` for the table.

Visual Basic with ADO cannot handle big integers. This means that some queries like `SHOW PROCESSLIST` do not work properly. The fix is to use `OPTION=16384` in the ODBC connect string or to select the `Change BIGINT columns to INT` option in the Connector/ODBC connect screen. You may also want to select the `Return matching rows` option.

#### 21.1.7.2.1.4. Microsoft Visual InterDev

If you have a `BIGINT` in your result, you may get the error `[Microsoft][ODBC Driver Manager] Driver does not support this parameter`. Try selecting the `Change BIGINT columns to INT` option in the Connector/ODBC connect screen.

#### 21.1.7.2.1.5. Visual Objects

Select the `Don't optimize column widths` option.

#### 21.1.7.2.1.6. Microsoft ADO

When you are coding with the ADO API and Connector/ODBC, you need to pay attention to some default properties that aren't supported by the MySQL server. For example, using the `CursorLocation Property` as `adUseServer` returns a result of `-1` for the `RecordCount Property`. To have the right value, you need to set this property to `adUseClient`, as shown in the VB code here:

```
Dim myconn As New ADODB.Connection
Dim myrs As New Recordset
Dim mySQL As String
Dim myrows As Long

myconn.Open "DSN=MyODBCsample"
mySQL = "SELECT * from user"
myrs.Source = mySQL
Set myrs.ActiveConnection = myconn
myrs.CursorLocation = adUseClient
myrs.Open
myrows = myrs.RecordCount

myrs.Close
myconn.Close
```

Another workaround is to use a `SELECT COUNT(*)` statement for a similar query to get the correct row count.

To find the number of rows affected by a specific SQL statement in ADO, use the `RecordsAffected` property in the ADO execute method. For more information on the usage of execute method, refer to <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/ado270/hlm/mdmthcnnextecute.asp>.

For information, see [ActiveX Data Objects\(ADO\) Frequently Asked Questions](#).

#### 21.1.7.2.1.7. Using Connector/ODBC with Active Server Pages (ASP)

Select the `Return matching rows` option in the DSN.

For more information about how to access MySQL through ASP using Connector/ODBC, refer to the following articles:

- [Using MyODBC To Access Your MySQL Database Via ASP](#)
- [ASP and MySQL at DWAM.NT](#)

A Frequently Asked Questions list for ASP can be found at <http://support.microsoft.com/default.aspx?scid=/Support/ActiveServer/faq/data/adofaq.asp>.

#### 21.1.7.2.1.8. Using Connector/ODBC with Visual Basic (ADO, DAO and RDO) and ASP

Some articles that may help with Visual Basic and ASP:

- [MySQL BLOB columns and Visual Basic 6](#) by Mike Hillyer (<mike@openwin.org>).
- [How to map Visual basic data type to MySQL types](#) by Mike Hillyer (<mike@openwin.org>).

#### 21.1.7.2.2. Using Connector/ODBC with Borland Applications

With all Borland applications where the Borland Database Engine (BDE) is used, follow these steps to improve compatibility:

- Update to BDE 3.2 or newer.
- Enable the `Don't optimize column widths` option in the DSN.
- Enabled the `Return matching rows` option in the DSN.

#### 21.1.7.2.2.1. Using Connector/ODBC with Borland Builder 4

When you start a query, you can use the `Active` property or the `Open` method. Note that `Active` starts by automatically issuing a `SELECT * FROM ...` query. That may not be a good thing if your tables are large.

#### 21.1.7.2.2.2. Using Connector/ODBC with Delphi

Also, here is some potentially useful Delphi code that sets up both an ODBC entry and a BDE entry for Connector/ODBC. The BDE entry requires a BDE Alias Editor that is free at a Delphi Super Page near you. (Thanks to Bryan Brunton <bryan@flesherfab.com> for this):

```
fReg:= TRegistry.Create;
fReg.OpenKey('\Software\ODBC\ODBC.INI\DocumentsFab', True);
fReg.WriteString('Database', 'Documents');
fReg.WriteString('Description', '');
fReg.WriteString('Driver', 'C:\WINNT\System32\myodbc.dll');
fReg.WriteString('Flag', '1');
fReg.WriteString('Password', '');
fReg.WriteString('Port', '');
fReg.WriteString('Server', 'xmark');
fReg.WriteString('User', 'winuser');
fReg.OpenKey('\Software\ODBC\ODBC.INI\ODBC Data Sources', True);
```

```
fReg.WriteString('DocumentsFab', 'MySQL');
fReg.CloseKey;
fReg.Free;

Memo1.Lines.Add('DATABASE NAME=');
Memo1.Lines.Add('USER NAME=');
Memo1.Lines.Add('ODBC DSN=DocumentsFab');
Memo1.Lines.Add('OPEN MODE=READ/WRITE');
Memo1.Lines.Add('BATCH COUNT=200');
Memo1.Lines.Add('LANGDRIVER=');
Memo1.Lines.Add('MAX ROWS=-1');
Memo1.Lines.Add('SCHEMA CACHE DIR=');
Memo1.Lines.Add('SCHEMA CACHE SIZE=8');
Memo1.Lines.Add('SCHEMA CACHE TIME=-1');
Memo1.Lines.Add('SQLPASSTHRU MODE=SHARED AUTOCOMMIT');
Memo1.Lines.Add('SQLQRYMODE=');
Memo1.Lines.Add('ENABLE SCHEMA CACHE=FALSE');
Memo1.Lines.Add('ENABLE BCD=FALSE');
Memo1.Lines.Add('ROWSET SIZE=20');
Memo1.Lines.Add('BLOBS TO CACHE=64');
Memo1.Lines.Add('BLOB SIZE=32');

AliasEditor.Add('DocumentsFab', 'MySQL', Memo1.Lines);
```

#### 21.1.7.2.2.3. Using Connector/ODBC with C++ Builder

Tested with BDE 3.0. The only known problem is that when the table schema changes, query fields are not updated. BDE, however, does not seem to recognize primary keys, only the index named `PRIMARY`, although this has not been a problem.

#### 21.1.7.2.3. Using Connector/ODBC with ColdFusion

The following information is taken from the ColdFusion documentation:

Use the following information to configure ColdFusion Server for Linux to use the `unixODBC` driver with Connector/ODBC for MySQL data sources. You can download Connector/ODBC at <http://dev.mysql.com/downloads/connector/odbc/>.

ColdFusion version 4.5.1 lets you use the ColdFusion Administrator to add the MySQL data source. However, the driver is not included with ColdFusion version 4.5.1. Before the MySQL driver appears in the ODBC data sources drop-down list, build and copy the Connector/ODBC driver to `/opt/coldfusion/lib/libmyodbc.so`.

The Contrib directory contains the program `mysdn-xxx.zip` which lets you build and remove the DSN registry file for the Connector/ODBC driver on ColdFusion applications.

For more information and guides on using ColdFusion and Connector/ODBC, see the following external sites:

- [Troubleshooting Data Sources and Database Connectivity for Unix Platforms](#).

#### 21.1.7.2.4. Using Connector/ODBC with OpenOffice.org

Open Office (<http://www.openoffice.org>) [How-to: MySQL + OpenOffice](#). [How-to: OpenOffice + MyODBC + unixODBC](#).

#### 21.1.7.2.5. Using Connector/ODBC with Sambar Server

Sambar Server (<http://www.sambarserver.info>) [How-to: MyODBC + SambarServer + MySQL](#).

#### 21.1.7.2.6. Using Connector/ODBC with Pervasive Software DataJunction

You have to change it to output `VARCHAR` rather than `ENUM`, as it exports the latter in a manner that causes MySQL problems.

#### 21.1.7.2.7. Using Connector/ODBC with SunSystems Vision



---

Select the [Return matching rows](#) option.

### 21.1.7.3. Connector/ODBC Errors and Resolutions (FAQ)

The following section details some common errors and their suggested fix or alternative solution. If you are still experiencing problems, use the Connector/ODBC mailing list; see [Section 21.1.8.1, "Connector/ODBC Community Support"](#).

Many problems can be resolved by upgrading your Connector/ODBC drivers to the latest available release. On Windows, make sure that you have the latest versions of the Microsoft Data Access Components (MDAC) installed.

#### 64-Bit Windows and ODBC Data Source Administrator

I have installed Connector/ODBC on Windows XP x64 Edition or Windows Server 2003 R2 x64. The installation completed successfully, but the Connector/ODBC driver does not appear in [ODBC Data Source Administrator](#).

This is not a bug, but is related to the way Windows x64 editions operate with the ODBC driver. On Windows x64 editions, the Connector/ODBC driver is installed in the `%SystemRoot%\SysWOW64` folder. However, the default [ODBC Data Source Administrator](#) that is available through the [Administrative Tools](#) or [Control Panel](#) in Windows x64 Editions is located in the `%SystemRoot%\system32` folder, and only searches this folder for ODBC drivers.

On Windows x64 editions, use the ODBC administration tool located at `%SystemRoot%\SysWOW64\odbcad32.exe`, this will correctly locate the installed Connector/ODBC drivers and enable you to create a Connector/ODBC DSN.

This issue was originally reported as Bug #20301.

#### Error 10061 (Cannot connect to server)

When connecting or using the [Test](#) button in [ODBC Data Source Administrator](#) I get error 10061 (Cannot connect to server)

This error can be raised by a number of different issues, including server problems, network problems, and firewall and port blocking problems. For more information, see [Section C.5.2.2, "Can't connect to \[local\] MySQL server"](#).

#### "Transactions are not enabled" Error

The following error is reported when using transactions: [Transactions are not enabled](#)

This error indicates that you are trying to use [transactions](#) with a MySQL table that does not support transactions. Transactions are supported within MySQL when using the [InnoDB](#) database engine, which is the default storage engine in MySQL 5.5 and higher. In versions of MySQL before MySQL 5.1, you may also use the [BDB](#) engine.

Check the following before continuing:

- Verify that your MySQL server supports a transactional database engine. Use [SHOW ENGINES](#) to obtain a list of the available engine types.
- Verify that the tables you are updating use a transactional database engine.
- Ensure that you have not enabled the [disable transactions](#) option in your DSN.

#### #DELETED# Records Reported by Access

Access reports records as [#DELETED#](#) when inserting or updating records in linked tables.

If the inserted or updated records are shown as [#DELETED#](#) in Access, then:

- If you are using Access 2000, get and install the newest (version 2.6 or higher) Microsoft MDAC (Microsoft Data Access Components) from <http://support.microsoft.com/kb/110093>. This fixes a bug in Access that when you export data to MySQL, the table and column names aren't specified.

Also, get and apply the Microsoft Jet 4.0 Service Pack 5 (SP5), which can be found at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q239114>. This fixes some cases where columns are marked as #DELETED# in Access.

- For all versions of Access, enable the Connector/ODBC `Return matching rows` option. For Access 2.0, also enable the `Simulate ODBC 1.0` option.
- Include a `TIMESTAMP` in all tables that you want to be able to update.
- Include a `primary key` in the table. If not, new or updated rows may show up as #DELETED#.
- Use only `DOUBLE` float fields. Access fails when comparing with single-precision floats. The symptom usually is that new or updated rows may show up as #DELETED# or that you cannot find or update rows.
- If you are using Connector/ODBC to link to a table that has a `BIGINT` column, the results are displayed as #DELETED#. The work around solution is:
  - Have one more dummy column with `TIMESTAMP` as the data type.
  - Select the `Change BIGINT columns to INT` option in the connection dialog in ODBC DSN Administrator.
  - Delete the table link from Access and re-create it.

Old records still display as #DELETED#, but newly added/updated records are displayed properly.

## Write Conflicts or Row Location Errors

How do I handle Write Conflicts or Row Location errors?

If you see the following errors, select the `Return Matching Rows` option in the DSN configuration dialog, or specify `OPTION=2`, as the connection parameter:

```
Write Conflict. Another user has changed your data.
```

```
Row cannot be located for updating. Some values may have been changed since it was last read.
```

## Importing from Access 97

Exporting data from Access 97 to MySQL reports a `Syntax Error`.

This error is specific to Access 97 and versions of Connector/ODBC earlier than 3.51.02. Update to the latest version of the Connector/ODBC driver to resolve this problem.

## Importing from Microsoft DTS

Exporting data from Microsoft DTS to MySQL reports a `Syntax Error`.

This error occurs only with MySQL tables using the `TEXT` or `VARCHAR` data types. You can fix this error by upgrading your Connector/ODBC driver to version 3.51.02 or higher.

## SQL\_NO\_DATA Exception from ODBC.NET

Using ODBC.NET with Connector/ODBC, while fetching empty string (0 length), it starts giving the `SQL_NO_DATA` exception.

You can get the patch that addresses this problem from <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q319243>.

## Error with SELECT COUNT(\*)

Using `SELECT COUNT(*) FROM tbl_name` within Visual Basic and ASP returns an error.

This error occurs because the `COUNT(*)` [1283] expression is returning a `BIGINT`, and ADO cannot make sense of a number this big. Select the `Change BIGINT columns to INT` option (option value 16384).

## Multiple-Step Operation Error

Using the `AppendChunk()` or `GetChunk()` ADO methods, the `Multiple-step operation generated errors. Check each status value` error is returned.

The `GetChunk()` and `AppendChunk()` methods from ADO do not work as expected when the cursor location is specified as `adUseServer`. On the other hand, you can overcome this error by using `adUseClient`.

A simple example can be found from [http://www.dwam.net/iishelp/ado/docs/adomth02\\_4.htm](http://www.dwam.net/iishelp/ado/docs/adomth02_4.htm)

## Modified Record Error

Access returns `Another user had modified the record that you have modified` while editing records on a Linked Table.

In most cases, this can be solved by doing one of the following things:

- Add a `primary key` for the table if one doesn't exist.
- Add a timestamp column if one doesn't exist.
- Only use double-precision float fields. Some programs may fail when they compare single-precision floats.

If these strategies do not help, start by making a log file from the ODBC manager (the log you get when requesting logs from ODBCADMIN) and a Connector/ODBC log to help you figure out why things go wrong. For instructions, see [Section 21.1.4.8, "Getting an ODBC Trace File"](#).

## Direct Application Linking Under Unix or Linux

When linking an application directly to the Connector/ODBC library under Unix or Linux, the application crashes.

Connector/ODBC under Unix or Linux is not compatible with direct application linking. To connect to an ODBC source, use a driver manager, such as `iODBC` or `unixODBC`.

## Microsoft Office and DATE or TIMESTAMP Columns

Applications in the Microsoft Office suite cannot update tables that have `DATE` or `TIMESTAMP` columns.

This is a known issue with Connector/ODBC. Ensure that the field has a default value (rather than `NULL`) and that the default value is nonzero (that is, something other than `0000-00-00 00:00:00`).

## INFORMATION\_SCHEMA Database

When connecting Connector/ODBC 5.x to a MySQL 4.x server, the error `1044 Access denied for user 'xxx'@'%' to database 'information_schema'` is returned.

Connector/ODBC 5.x is designed to work with MySQL 5.0 or later, taking advantage of the `INFORMATION_SCHEMA` database to determine data definition information. Support for MySQL 4.1 is planned for the final release.

## S1T00 Error

When calling `SQLTables`, the error `S1T00` is returned, but I cannot find this in the list of error numbers for Connector/ODBC.

The `S1T00` error indicates that a general timeout has occurred within the ODBC system and is not a MySQL error. Typically it indicates that the connection you are using is stale, the server is too busy to accept your request or that the server has gone away.

### "Table does not exist" Error in Access 2000

When linking to tables in Access 2000 and generating links to tables programmatically, rather than through the table designer interface, you may get errors about tables not existing.

There is a known issue with a specific version of the `msjet40.dll` that exhibits this issue. The version affected is 4.0.9025.0. Reverting to an older version will enable you to create the links. If you have recently updated your version, check your `WINDOWS` directory for the older version of the file and copy it to the drivers directory.

### Batched Statements

When I try to use batched statements, the execution of the batched statements fails.

Batched statement support was added in 3.51.18. Support for batched statements is not enabled by default. Enable option `FLAG_MULTI_STATEMENTS`, value 67108864, or select the **Allow multiple statements** flag within a GUI configuration.

### Packet Errors with ADODB and Excel

When connecting to a MySQL server using ADODB and Excel, occasionally the application fails to communicate with the server and the error `Got an error reading communication packets` appears in the error log.

This error may be related to Keyboard Logger 1.1 from PanteraSoft.com, which is known to interfere with the network communication between MySQL Connector/ODBC and MySQL.

### Outer Join Error

When using some applications to access a MySQL server using Connector/ODBC and outer joins, an error is reported regarding the Outer Join Escape Sequence.

This is a known issue with MySQL Connector/ODBC which is not correctly parsing the "Outer Join Escape Sequence", as per the specs at [Microsoft ODBC Specs](#). Currently, Connector/ODBC will return a value `> 0` when asked for `SQL_OJ_CAPABILITIES` even though no parsing takes place in the driver to handle the outer join escape sequence.

### Hebrew/CJK Characters

I can correctly store extended characters in the database (Hebrew/CJK) using Connector/ODBC 5.1, but when I retrieve the data, the text is not formatted correctly and I get garbled characters.

When using ASP and UTF8 characters, add the following to your ASP files to ensure that the data returned is correctly encoded:

```
Response.CodePage = 65001
Response.CharSet = "utf-8"
```

### Duplicate Entry in Installed Programs List

I have a duplicate MySQL Connector/ODBC entry within my **Installed Programs** list, but I cannot delete one of them.

This problem can occur when you upgrade an existing Connector/ODBC installation, rather than removing and then installing the updated version.



#### Warning

To fix the problem, use any working uninstallers to remove existing installations; then may have to edit the contents of the registry. Make sure you have a

backup of your registry information before attempting any editing of the registry contents.

## Values Truncated to 255 Characters

When submitting queries with parameter binding using `UPDATE`, my field values are being truncated to 255 characters.

Ensure that the `FLAG_BIG_PACKETS` option is set for your connection. This removes the 255 character limitation on bound parameters.

## Disabling Data-At-Execution

Is it possible to disable data-at-execution using a flag?

If you do not want to use data-at-execution, remove the corresponding calls. For example:

```
SQLLEN ylen = SQL_LEN_DATA_AT_EXEC(10);
SQLBindCol(hstmt,2,SQL_C_BINARY, buf, 10, &ylen);
```

Would become:

```
SQLBindCol(hstmt,2,SQL_C_BINARY, buf, 10, NULL);
```

Note that in the call to `SQLBindCol()`, `&ylen` has been replaced by `NULL`.

For further information, refer to the [MSDN documentation](#) for `SQLBindCol()`.

## NULLABLE Attribute for AUTO\_INCREMENT Columns

When you call `SQLColumns()` for a table column that is `AUTO_INCREMENT`, the `NULLABLE` column of the result set is always `SQL_NULLABLE (1)`.

This is because MySQL reports the `DEFAULT` value for such a column as `NULL`. It means, if you insert a `NULL` value into the column, you will get the next integer value for the table's `auto_increment` counter.

## 21.1.8. Connector/ODBC Support

There are many different places where you can get support for using Connector/ODBC. Always try the Connector/ODBC Mailing List or Connector/ODBC Forum. See [Section 21.1.8.1, "Connector/ODBC Community Support"](#), for help before reporting a specific bug or issue to MySQL.

### 21.1.8.1. Connector/ODBC Community Support

Oracle provides assistance to the user community by means of its mailing lists. For Connector/ODBC-related issues, you can get help from experienced users by using the [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com) mailing list. Archives are available online at <http://lists.mysql.com/myodbc>.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See [Section 1.6.1, "MySQL Mailing Lists"](#).

Community support from experienced users is also available through the [ODBC Forum](#). You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See [Section 1.6.2, "MySQL Community Support at the MySQL Forums"](#).

### 21.1.8.2. How to Report Connector/ODBC Problems or Bugs

If you encounter difficulties or problems with Connector/ODBC, start by making a log file from the [ODBC Manager](#) (the log you get when requesting logs from `ODBC ADMIN`) and Connector/ODBC. The procedure for doing this is described in [Section 21.1.4.8, "Getting an ODBC Trace File"](#).

Check the Connector/ODBC trace file to find out what could be wrong. Determine what statements were issued by searching for the string `>mysql_real_query` in the `myodbc.log` file.

Also, try issuing the statements from the `mysql` client program or from `admindemo`. This helps you determine whether the error is in Connector/ODBC or MySQL.

If you find out something is wrong, please only send the relevant rows (maximum 40 rows) to the `myodbc` mailing list. See [Section 1.6.1, "MySQL Mailing Lists"](#). Please never send the whole Connector/ODBC or ODBC log file!

Ideally, include the following information with the email:

- Operating system and version
- Connector/ODBC version
- ODBC Driver Manager type and version
- MySQL server version
- ODBC trace from Driver Manager
- Connector/ODBC log file from Connector/ODBC driver
- Simple reproducible sample

Remember that the more information you can supply to us, the more likely it is that we can fix the problem!

Also, before posting the bug, check the MyODBC mailing list archive at <http://lists.mysql.com/myodbc>.

If you are unable to find out what is wrong, the last option is to create an archive in `tar` or `zip` format that contains a Connector/ODBC trace file, the ODBC log file, and a `README` file that explains the problem. Initiate a bug report for our bugs database at <http://bugs.mysql.com/>, then click the Files tab in the bug report for instructions on uploading the archive to the bugs database. Only MySQL engineers have access to the files you upload, and we are very discreet with the data.

If you can create a program that also demonstrates the problem, please include it in the archive as well.

If the program works with another SQL server, include an ODBC log file where you perform exactly the same SQL statements so that we can compare the results between the two systems.

Remember that the more information you can supply to us, the more likely it is that we can fix the problem.

### 21.1.8.3. How to Submit a Connector/ODBC Patch

You can send a patch or suggest a better solution for any existing code or problems by sending a mail message to `<myodbc@lists.mysql.com>`.

## 21.2. MySQL Connector/Net

Connector/Net lets you easily develop .NET applications that require secure, high-performance data connectivity with MySQL. It implements the required ADO.NET interfaces and integrates into ADO.NET-aware tools. Developers can build applications using their choice of .NET languages. Connector/Net is a fully managed ADO.NET driver written in 100% pure C#.

For release notes detailing the changes in each release of Connector/Net, see [MySQL Connector/Net Release Notes](#).

Connector/Net includes full support for:



- Features provided by MySQL Server up to and including MySQL Server version 5.5.
- Large-packet support for sending and receiving rows and BLOBs up to 2 gigabytes in size.
- Protocol compression, which enables compressing the data stream between the client and server.
- Connections using TCP/IP sockets, named pipes, or shared memory on Windows.
- Connections using TCP/IP sockets or Unix sockets on Unix.
- The Open Source Mono framework developed by Novell.
- Fully managed, does not utilize the MySQL client library.

This document is intended as a user's guide to Connector/Net and includes a full syntax reference. Syntax information is also included within the [Documentation.chm](#) file included with the Connector/Net distribution.

If you are using MySQL 5.0 or later, and Visual Studio as your development environment, you can also use the MySQL Visual Studio Plugin. The plugin acts as a DDEX (Data Designer Extensibility) provider: you can use the data design tools within Visual Studio to manipulate the schema and objects within a MySQL database. For more information, see [Section 21.2.3, "Connector/Net Visual Studio Integration"](#).



**Note**

Connector/Net 5.1.2 and later include the Visual Studio Plugin by default.

MySQL Connector/Net supports full versions of Visual Studio 2005, 2008, and 2010, although certain features are only available in Visual Studio 2010 when using MySQL Connector/Net version 6.3.2 and later. Note that MySQL Connector/Net does not currently support Express versions of Microsoft products, including Microsoft Visual Web Developer.

**Key topics:**

- For connection string properties when using the `MySqlConnection` class, see [Section 21.2.6, "Connector/Net Connection String Options Reference"](#).

As of Connector/Net 5.1.2, the Visual Studio Plugin is part of the main Connector/Net package. However, even for versions of Connector/Net older than 5.1.2, the Visual Studio Plugin release notes can be found at [MySQL Connector/Net Release Notes](#).

## 21.2.1. Connector/Net Versions

There are several versions of Connector/Net available:

- Connector/Net 6.6 includes support for MySQL Server 5.6, 5.5, 5.1, and 5.0. Important new features include stored procedure debugging in Microsoft Visual Studio, support for pluggable authentication including the ability to write your own authentication plugin, Entity Framework 4.3 Code First support, and enhancements to partial trust support to allow hosting services to deploy applications without installing the Connector/Net library in the GAC.
- Connector/Net 6.5 includes support for MySQL Server 5.6, 5.5, 5.1, and 5.0. Important new features include interceptor classes for exceptions and commands, support for the MySQL 5.6 fractional seconds feature, better partial-trust support, and better IntelliSense, including auto-completion when editing stored procedures or `.mysql` files.
- Connector/Net 6.4 includes support for MySQL Server 5.5, 5.1, and 5.0. Important new features include support for Windows authentication (when connecting to MySQL Server 5.5+), table caching on the client side, simple connection fail-over support, and improved SQL generation from the Entity Framework provider.



- Connector/Net 6.3 includes support for MySQL Server 5.5, 5.1, and 5.0. Important new features include integration with Visual Studio 2010, such as availability of DDL T4 template for Entity Framework, and a custom MySQL SQL Editor. Other features include refactored transaction scope: Connector/Net now supports nested transactions in a scope where they use the same connection string.
- Connector/Net 6.2 includes support for MySQL Server 5.5, 5.1, 5.0, and 4.1. Important new features include a new logging system and client SSL certificates.
- Connector/Net 6.1 includes support for MySQL Server 5.5, 5.1, 5.0, and 4.1. Important new features include the MySQL Website Configuration Tool and a Session State Provider.
- Connector/Net 6.0 includes support for MySQL Server 5.5, 5.1, 5.0, and 4.1.

This version of Connector/Net is no longer supported.

- Connector/Net 5.2 includes support for MySQL Server 5.5, 5.1, 5.0, and 4.1 features. Connector/Net 5.2 also includes support for a new membership/role provider, Compact Framework 2.0, a new stored procedure parser and improvements to [GetSchema](#). Connector/Net 5.2 also includes the Visual Studio Plugin as a standard installable component.

This version of Connector/Net is no longer supported.

- Connector/Net 5.1 includes support for MySQL Server 5.5, 5.1, 5.0, 4.1, and 4.0 features. Connector/Net 5.1 also includes support for a new membership/role provider, Compact Framework 2.0, a new stored procedure parser and improvements to [GetSchema](#). Connector/Net 5.1 also includes the Visual Studio Plugin as a standard installable component.

This version of Connector/Net is no longer supported.

- Connector/Net 5.0 includes support for MySQL Server 5.1, 5.0, 4.1 and 4.0 features. Connector/Net 5.0 also includes full support for the ADO.Net 2.0 interfaces and subclasses, includes support for the usage advisor and performance monitor (PerfMon) hooks.

This version of Connector/Net is no longer supported.

- Connector/Net 1.0 includes support for MySQL Server 5.0, 4.1, and 4.0 features, and full compatibility with the ADO.NET driver interface.

This version of Connector/Net is no longer supported.

The latest source code for Connector/Net can be downloaded from the MySQL public Subversion server. For further details, see [Section 21.2.2.3, "Installing Connector/Net from the Source Code"](#).

The following table shows the .NET Framework version required, and the MySQL Server version supported by Connector/Net:

**Table 21.19. Connector/Net Requirements for Related Products**

Connector/Net version	ADO.NET version supported	.NET Framework version required	MySQL Server version supported	Currently supported
6.6	2.x+	2.x+, 4.x+ for VS 2010 support	5.6, 5.5, 5.1, 5.0	Yes
6.5	2.x+	2.x+, 4.x+ for VS 2010 support	5.6, 5.5, 5.1, 5.0	Yes
6.4	2.x+	2.x+, 4.x+ for VS 2010 support	5.6, 5.5, 5.1, 5.0	Yes
6.3	2.x+	2.x+, 4.x+ for VS 2010 support	5.6, 5.5, 5.1, 5.0	Yes

Connector/Net version	ADO.NET version supported	.NET Framework version required	MySQL Server version supported	Currently supported
6.2	2.x+	2.x+	5.6, 5.5, 5.1, 5.0, 4.1	Yes
6.1	2.x+	2.x+	5.6, 5.5, 5.1, 5.0, 4.1	Yes
6.0	2.x+	2.x+	5.5, 5.1, 5.0, 4.1	Critical issues only
5.2	2.x+	2.x+	5.5, 5.1, 5.0, 4.1	No
5.1	2.x+	2.x+	5.5, 5.1, 5.0, 4.1, 4.0	No
5.0	2.x+	2.x+	5.0, 4.1, 4.0	No
1.0	1.x	1.x	5.0, 4.1, 4.0	No



#### Note

Version numbers for MySQL products are formatted as X.Y.Z, where Z=0 indicates alpha, Z=1 indicates beta, and Z>=2 indicates GA. However, Windows tools (Control Panel, properties display) may show the version numbers as XX.YY.ZZ. For example, the official MySQL formatted version number 5.0.9 may be displayed by Windows tools as 5.00.09. The two versions are the same; only the number display format is different.

## 21.2.2. Connector/Net Installation

Connector/Net runs on any platform that supports the .NET framework. The .NET framework is primarily supported on recent versions of Microsoft Windows, and is supported on Linux through the Open Source Mono framework (see <http://www.mono-project.com>).

Connector/Net is available for download from <http://dev.mysql.com/downloads/connector/net/>.

### 21.2.2.1. Installing Connector/Net on Windows

On Windows, you can install either through a binary installation process or by downloading a zip file with the Connector/Net components.

Before installing, ensure that your system is up to date, including installing the latest version of the .NET Framework.

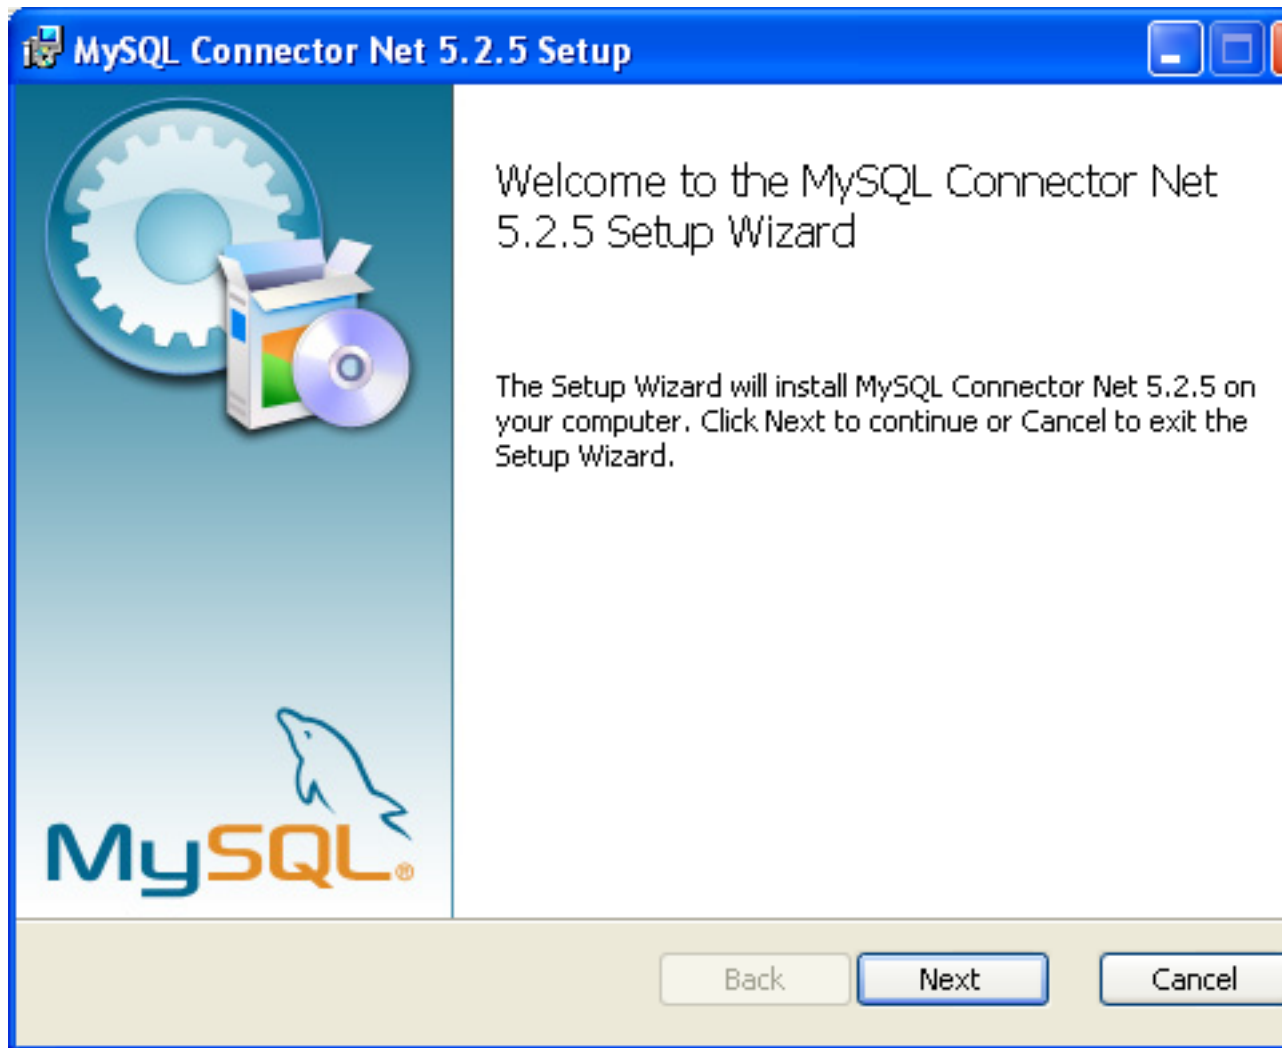
#### 21.2.2.1.1. Installing Connector/Net using the Installer

Using the installer is the most straightforward method of installing Connector/Net on Windows and the installed components include the source code, test code and full reference documentation.

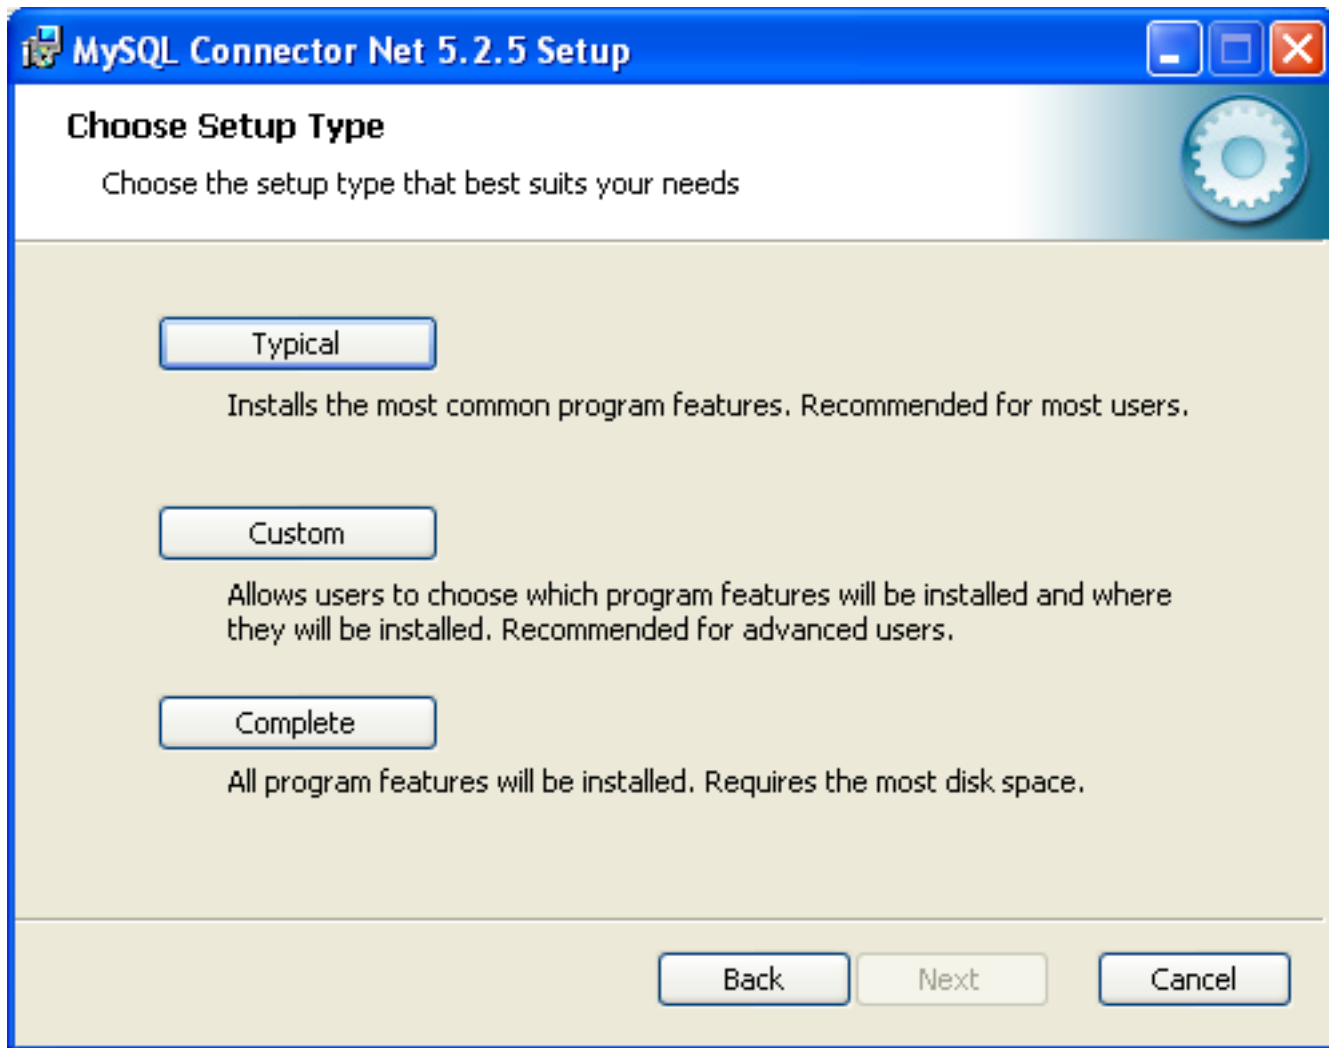
You install Connector/Net through a Windows Installer (.msi) installation package, which can install Connector/Net on all Windows operating systems. The MSI package is contained within a zip archive named `mysql-connector-net-version.zip`, where *version* indicates the Connector/Net version.

To install Connector/Net:

1. Double-click the MSI installer file extracted from the zip you downloaded. Click **Next** to start the installation.



2. You must choose the type of installation to perform.



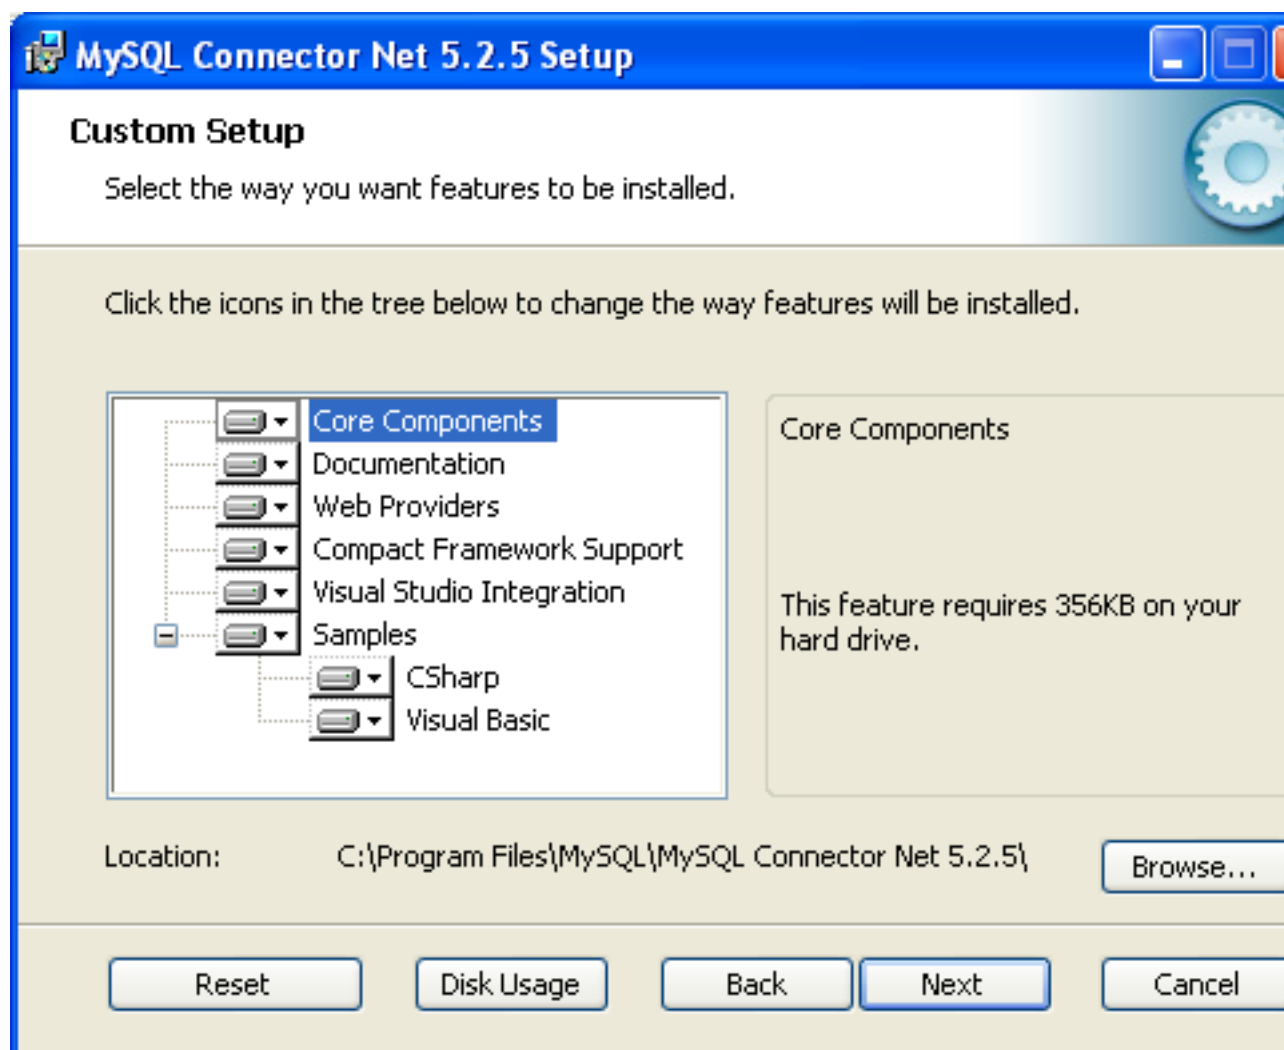
For most situations, the Typical installation is suitable. Click the **Typical** button and proceed to Step 5. A Complete installation installs all the available files. To conduct a Complete installation, click the **Complete** button and proceed to step 5. To customize your installation, including choosing the components to install and some installation options, click the **Custom** button and proceed to Step 3.

The Connector/Net installer will register the connector within the Global Assembly Cache (GAC) - this will make the Connector/Net component available to all applications, not just those where you explicitly reference the Connector/Net component. The installer will also create the necessary links in the Start menu to the documentation and release notes.

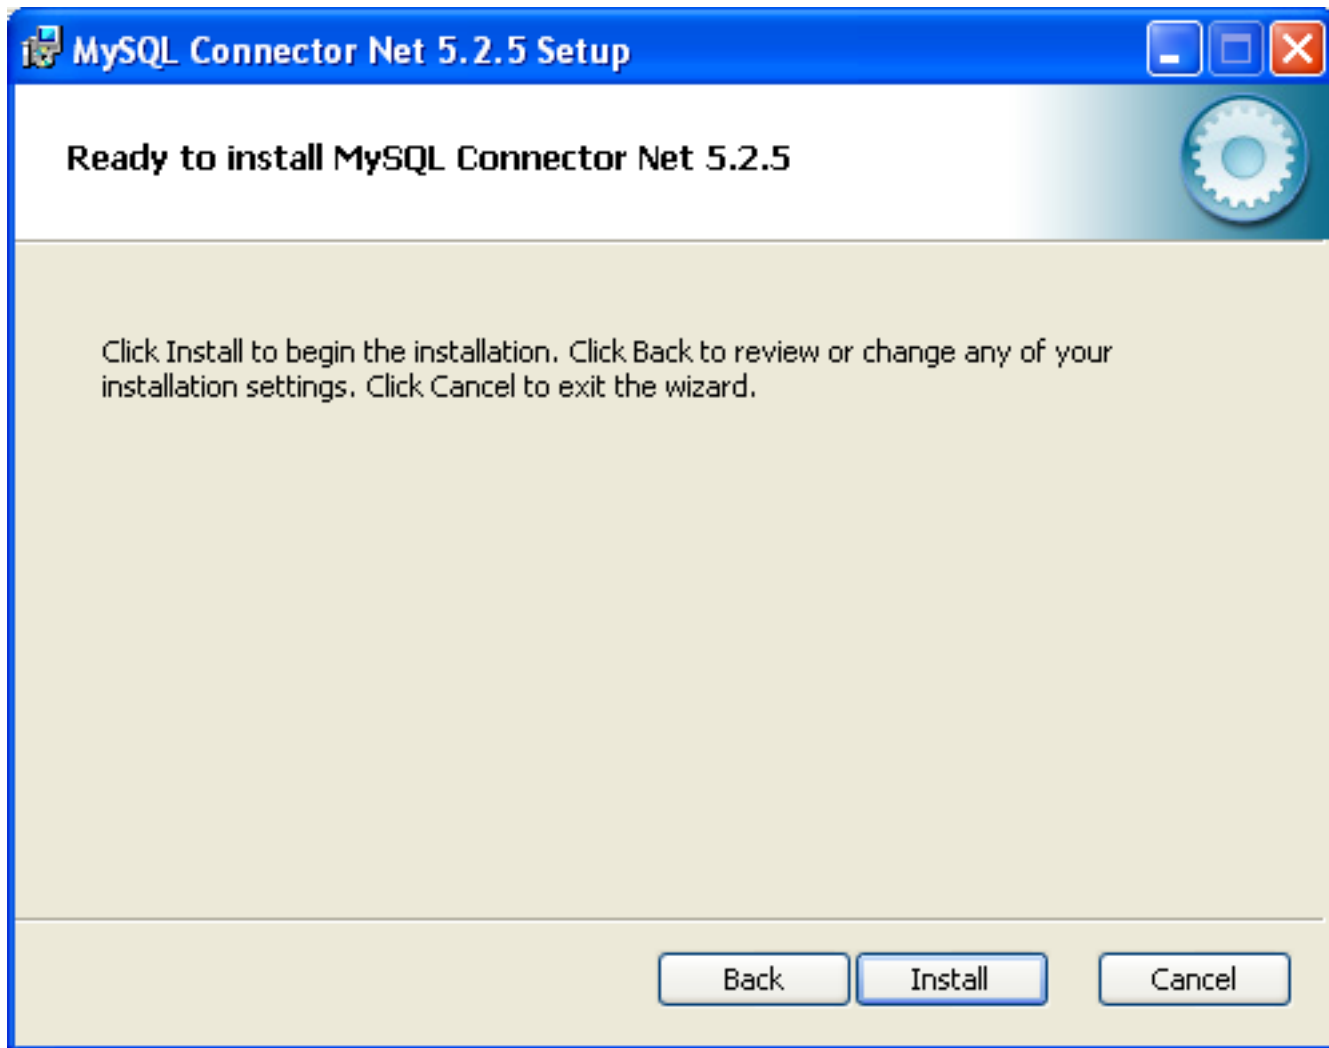
3. If you have chosen a custom installation, you can select the individual components to install, including the core interface component, supporting documentation (a CHM file) samples and examples, and the source code. Select the items, and their installation level, and then click **Next** to continue the installation.

**Note**


For Connector/Net 1.0.8 or lower and Connector 5.0.4 and lower the installer will attempt to install binaries for both 1.x and 2.x of the .NET Framework. If you only have one version of the framework installed, the connector installation may fail. If this happens, you can choose the framework version to be installed through the custom installation step.



4. You will be given a final opportunity to confirm the installation. Click **Install** to copy and install the files onto your machine.



5. Once the installation has been completed, click **Finish** to exit the installer.



Unless you choose otherwise, Connector/Net is installed in `C:\Program Files\MySQL\MySQL Connector Net X.X.X`, where `X.X.X` is replaced with the version of Connector/Net you are installing. New installations do not overwrite existing versions of Connector/Net.

Depending on your installation type, the installed components will include some or all of the following components:

- `bin`: Connector/Net MySQL libraries for different versions of the .NET environment.
- `docs`: Connector/Net documentation in CHM format.
- `samples`: Sample code and applications that use the Connector/Net component.
- `src`: The source code for the Connector/Net component.

You may also use the `/quiet` or `/q` command-line option with the `msiexec` tool to install the Connector/Net package automatically (using the default options) with no notification to the user. Using this method the user cannot select options. Additionally, no prompts, messages or dialog boxes will be displayed.

```
C:\> msiexec /package connector-net.msi /quiet
```

To provide a progress bar to the user during automatic installation, use the `/passive` option.

#### 21.2.2.1.2. Installing Connector/Net Using the Zip Packages



If you have problems running the installer, you can download a zip file without an installer as an alternative. That file is called `mysql-connector-net-version-noinstall.zip`. After downloading the zip archive, extract the files to a location of your choice.

The file contains the following directories:

- `bin`: Connector/Net MySQL libraries for different versions of the .NET environment.
- `Docs`: Connector/Net documentation in CHM format.
- `Samples`: Sample code and applications that use the Connector/Net component.

Connector/Net 6.0.x has a different directory structure:

- `Assemblies`: A collection of DLLs that make up the connector functionality.
- `Documentation`: Connector/Net documentation in CHM format.
- `Samples`: sample code and applications that use the Connector/Net component.

Another zip file available for download contains the source code distribution. This file is named `mysql-connector-net-version-src.zip`.

The file contains the following directories:

- `Documentation`: Source files to build the documentation into the compiled HTML (CHM) format.
- `Installer`: Source files to build the Connector/Net installer program.
- `MySql.Data`: Source files for the core data provider.
- `MySql.VisualStudio`: Source files for the Microsoft Visual Studio extensions.
- `MySql.Web`: Source files for the web providers. This includes code for the membership provider, role provider and profile provider. These are used in ASP.NET web sites.
- `Samples`: Source files for several example applications.
- `Tests`: A spreadsheet listing test cases.
- `VisualStudio`: Resources used by the Visual Studio plugin.

Finally, ensure that `MySql.Data.dll` is accessible to your program at build time (and runtime). If using Microsoft Visual Studio, add `MySql.Data` as a Reference to your project.



#### Note

If using MySQL Connector/Net 6.3.5 and above, the `MySql.Data.dll` file provided will work with both .NET Framework 2.x and 4.x.

### 21.2.2.2. Installing Connector/Net on Unix with Mono

There is no installer available for installing the Connector/Net component on your Unix installation. Before installing, ensure that you have a working Mono project installation. To test whether your system has Mono installed, enter:

```
shell> mono --version
```

The version of the Mono JIT compiler is displayed.

To compile C# source code, make sure a Mono C# compiler is installed. Note that there are two Mono C# compilers available, `mcs`, which accesses the 1.0-profile libraries, and `gmcs`, which accesses the 2.0-profile libraries.

To install Connector/Net on Unix/Mono:

1. Download the `mysql-connector-net-version-noinstall.zip` and extract the contents to a directory of your choice, for example: `~/connector-net/`.
2. In the directory where you unzipped the connector to, change into the `bin` subdirectory. Ensure the file `MySQL.Data.dll` is present. This filename is case-sensitive.
3. You must register the Connector/Net component, `MySQL.Data`, in the Global Assembly Cache (GAC). In the current directory enter the `gacutil` command:

```
root-shell> gacutil /i MySQL.Data.dll
```

This will register `MySQL.Data` into the GAC. You can check this by listing the contents of `/usr/lib/mono/gac`, where you will find `MySQL.Data` if the registration has been successful.

You are now ready to compile your application. You must ensure that when you compile your application you include the Connector/Net component using the `-r:` command-line option. For example:

```
shell> gmcs -r:System.dll -r:System.Data.dll -r:MySQL.Data.dll HelloWorld.cs
```

Note, the assemblies that are referenced depend on the requirements of the application, but applications using Connector/Net must provide `-r:MySQL.Data` as a minimum.

You can further check your installation by running the compiled program, for example:

```
shell> mono HelloWorld.exe
```

### 21.2.2.3. Installing Connector/Net from the Source Code



#### Caution

Read this section only if you are interested in helping us test our new code. If you just want to get Connector/Net up and running on your system, use a standard release distribution.

#### Obtaining the Source Code

To obtain the most recent development source tree, first download and install Bazaar. You can obtain Bazaar from the [Bazaar VCS Website](#). Bazaar is supported by any platform that supports Python, and is therefore compatible with any Linux, Unix, Windows or Mac OS X host. Instructions for downloading and installing Bazaar on the different platforms are available on the Bazaar Web site.

The most recent development source tree is available from our public Subversion trees at <http://dev.mysql.com/tech-resources/sources.html>.

To check out the Connector/Net sources, change to the directory where you want the copy of the Connector/Net tree to be stored, then use the following command:

```
shell> bzip branch lp:connectornet/trunk
```

To download a specific version of Connector/Net, specify the version number instead of `trunk`. For example, to obtain a copy of the 6.0 version of the source tree:

```
shell> bzip branch lp:connectornet/6.0
```

Source packages are also available on the downloads page.

#### Building the Source Code on Windows

The following procedure can be used to build the connector on Microsoft Windows.

- Obtain the source code, either from the Subversion server, or through one of the prepared source code packages.
- Navigate to the root of the source code tree.
- A Microsoft Visual Studio 2005 solution file is available to build the connector, this is called [MySQL-VS2005.sln](#). Click this file to load the solution into Visual Studio.
- Select **Build, Build Solution** from the main menu to build the solution.

## Building the Source Code on Unix

Support for building Connector/Net on Mono/Unix is currently not available.

## 21.2.3. Connector/Net Visual Studio Integration

When MySQL Connector/Net is installed on Microsoft Windows, Visual Studio integration components are also installed and initialized. This enables the developer to work seamlessly with MySQL Connector/Net in the familiar Visual Studio environment, as described in the following sections of the manual.

MySQL Connector/Net supports Visual Studio versions 2005, 2008, and 2010. However, only MySQL Connector/Net version 6.3 and higher fully integrate with Visual Studio 2010, although applications using earlier versions of the connector can be built with the Visual Studio 2010 environment using .NET 2.x frameworks.

Visual Studio 2010 support was introduced with MySQL Connector/Net 6.3.2. From version 6.3.2 the connector ships with both .NET 2.x and .NET 4.x versions of the entity framework support files, [mysql.data.ef.dll](#) and [mysql.visualstudio.dll](#). The .NET 4.x versions are required to enable new integration features supported in Visual Studio 2010, including:

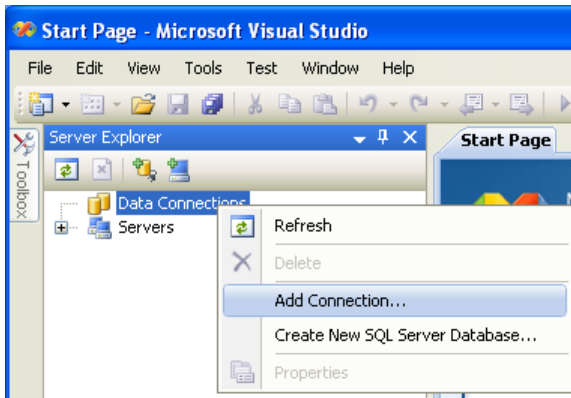
- New DDL T4 template for the Entity Framework (EF) - This enables developers to design an EF model from scratch and use the native Visual Studio 2010 facility to generate MySQL DDL from that model. This is done by creating the model, and with the model open, choosing the SSDLToMySQL template in the properties window. The correct DDL is then generated. The developer can then save this code as a [.mysql](#) file in their project and execute it against the MySQL server.
- New SQL Editor - A new SQL editor lets you connect to a MySQL server to execute SQL. This is activated by creating a new file with a [.mysql](#) extension. A new template lets you create files with this file type using the Visual Studio 2010 main menu item **File, New**. Note that the MySQL SQL Editor is also available in Visual Studio 2005 and 2008.

As of Connector/Net 5.1.2, the Visual Studio Plugin is part of the main Connector/Net package. However, even for versions of Connector/Net older than 5.1.2, the Visual Studio Plugin release notes can be found at [MySQL Connector/Net Release Notes](#).

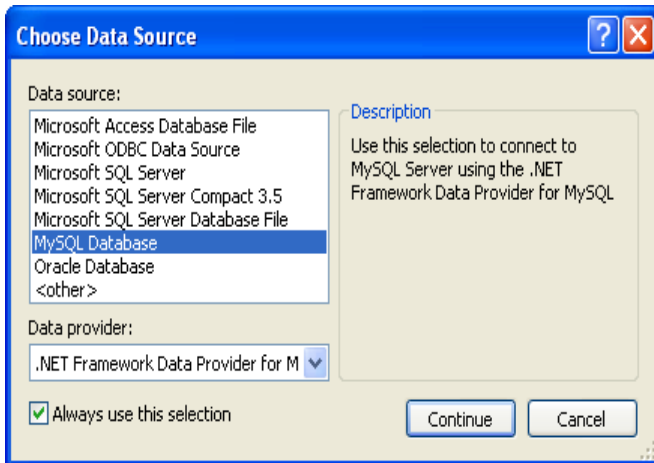
### 21.2.3.1. Making a Connection

Once the connector is installed, you can use it to create, modify, and delete connections to MySQL databases. To create a connection with a MySQL database, perform the following steps:

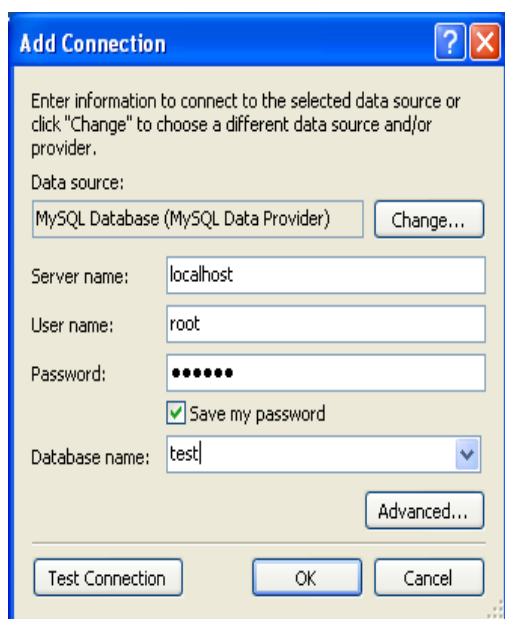
- Start Visual Studio, and open the Server Explorer window (**View, Server Explorer** option in the main Visual Studio menu, or **Control+W, L** keyboard shortcuts).
- Right-click the Data Connections node, and choose the **Add Connection...** menu item.
- Add Connection dialog opens. Press the **Change** button to choose MySQL Database as a data source.

**Figure 21.1. Add Connection Context Menu**

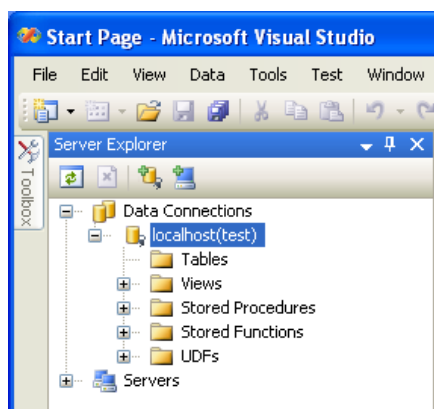
- **Change Data Source** dialog opens. Choose **MySQL Database** in the list of data sources (or the **<other>** option, if MySQL Database is absent), and then choose **.NET Framework Data Provider for MySQL** in the combo box of data providers.

**Figure 21.2. Choose Data Source**

- Input the connection settings: the server host name (for example, localhost if the MySQL server is installed on the local machine), the user name, the password, and the default schema name. Note that you must specify the default schema name to open the connection.

**Figure 21.3. Add Connection Dialog**

- You can also set the port to connect with the MySQL server by pressing the **Advanced** button. To test connection with the MySQL server, set the server host name, the user name, and the password, and press the **Test Connection** button. If the test succeeds, the success confirmation dialog opens.
- After you set all settings and test the connection, press **OK**. The newly created connection is displayed in Server Explorer. Now you can work with the MySQL server through standard Server Explorer GUI.

**Figure 21.4. New Data Connection**

After the connection is successfully established, all settings are saved for future use. When you start Visual Studio for the next time, open the connection node in Server Explorer to establish a connection to the MySQL server again.

To modify and delete a connection, use the Server Explorer context menu for the corresponding node. You can modify any of the settings by overwriting the existing values with new ones. Note that the connection may be modified or deleted only if no active editor for its objects is opened: otherwise, you may lose your data.

### 21.2.3.2. Using IntelliSense in the SQL Editor

IntelliSense support is available starting in Connector/Net 6.5. Once you have established a connection, for example, using the **Connect to MySQL** toolbar button, you can get autocompletion as

you type, or by pressing **Control+J**. Depending on the context, the autocompletion dialog can show the list of available tables, table columns, or stored procedures (with the routine's signature as a tooltip). Typing some characters before pressing **Control+J** filters the choices to those items starting with that prefix.

### 21.2.3.3. Editing Tables

Connector/Net contains a table editor, which enables the visual creation and modification of tables.

The Table Designer can be accessed through a mouse action on table-type node of Server Explorer. To create a new table, right-click the **Tables** node (under the connection node) and choose the **Create Table** command from the context menu.

To modify an existing table, double-click the node of the table to modify, or right-click this node and choose the **Design** item from the context menu. Either of the commands opens the Table Designer.

The table editor is implemented in the manner of the well-known Query Browser Table Editor, but with minor differences.

**Figure 21.5. Editing New Table**

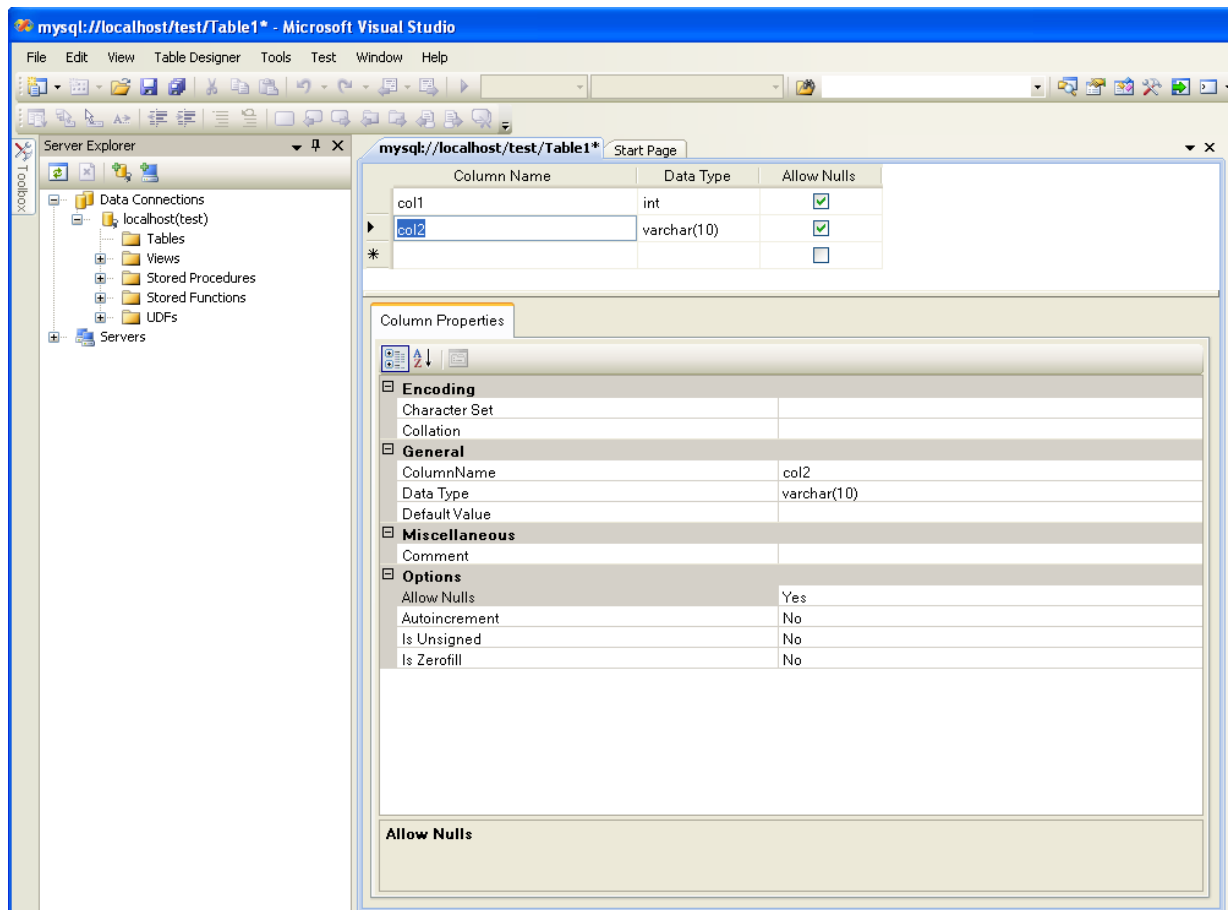


Table Designer consists of the following parts:


- Columns Editor - a data grid on top of the Table Designer. Use the Columns grid for column creation, modification, and deletion.
- Indexes tab - a tab on bottom of the Table Designer. Use the Indexes tab for indexes management.
- Foreign Keys tab - a tab on bottom of the Table Designer. Use the Foreign Keys tab for foreign keys management.

- Column Details tab - a tab on bottom of the Table Designer. Use the Column Details tab to set advanced column options.
- Properties window - a standard Visual Studio Properties window, where the properties of the edited table are displayed. Use the Properties window to set the table properties.

Each of these areas is discussed in more detail in subsequent sections.

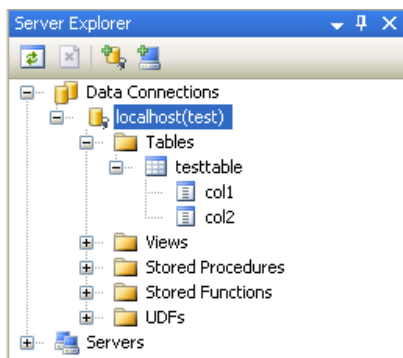
To save changes you have made in the Table Designer, use either **Save** or **Save All** button of the Visual Studio main toolbar, or press **Control+S**. If you have not already named the table, you will be prompted to do so.

**Figure 21.6. Choose Table Name**




Once the table is created, you can view it in the Server Explorer.

**Figure 21.7. Newly Created Table**



The Table Designer main menu lets you set a [primary key](#) column, edit relationships such as [foreign keys](#), and create [indexes](#).

**Figure 21.8. Table Designer Main Menu**



### 21.2.3.3.1. Column Editor

You can use the Column Editor to set or change the name, data type, default value, and other properties of a table column. To set the focus to a needed cell of a grid, use the mouse click. Also you can move through the grid using **Tab** and **Shift+Tab** keys.



To set or change the name, data type, default value and comment of a column, activate the appropriate cell and type the desired value.

To set or unset flag-type column properties (`NOT NULL`, auto incremented, flags), check or uncheck the corresponding check boxes. Note that the set of column flags depends on its data type.

To reorder columns, index columns or foreign key columns in the Column Editor, select the whole column to reorder by clicking the selector column on the left of the column grid. Then move the column by using **Control+Up** (to move the column up) or **Control+Down** (to move the column down) keys.

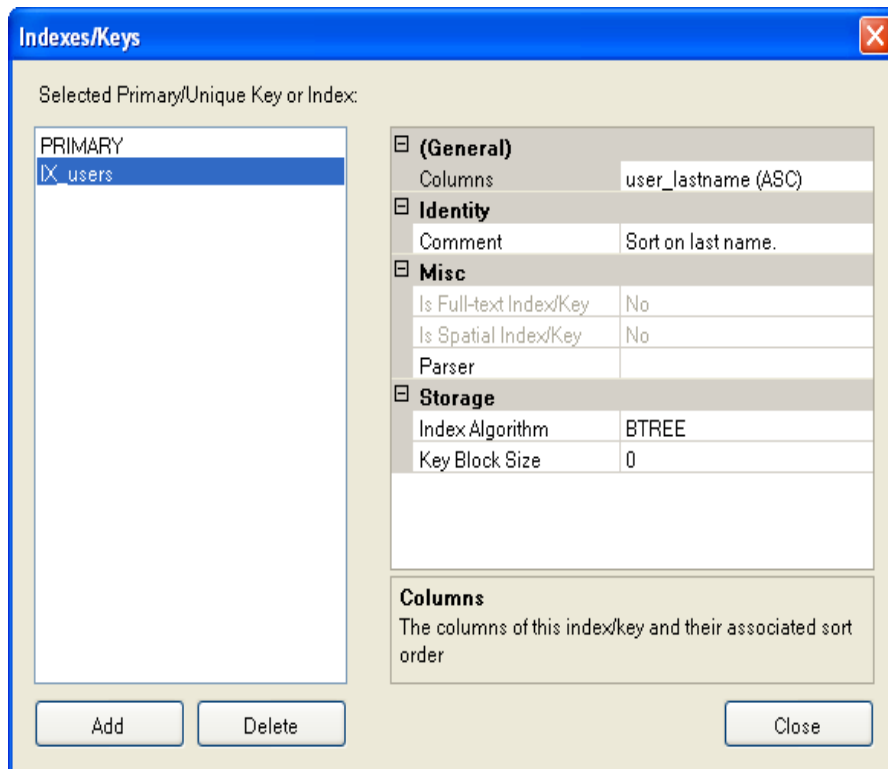
To delete a column, select it by clicking the selector column on the left of the column grid, then press the **Delete** button on a keyboard.

### 21.2.3.3.2. Editing Indexes

Indexes management is performed using the **Indexes/Keys** dialog.

To add an index, select **Table Designer, Indexes/Keys...** from the main menu, and click **Add** to add a new index. You can then set the index name, index kind, index type, and a set of index columns.

**Figure 21.9. Indexes Dialog**



To remove an index, select it in the list box on the left, and click the **Delete** button.

To change index settings, select the needed index in the list box on the left. The detailed information about the index is displayed in the panel on the right hand side. Change the desired values.

### 21.2.3.3.3. Editing Foreign Keys

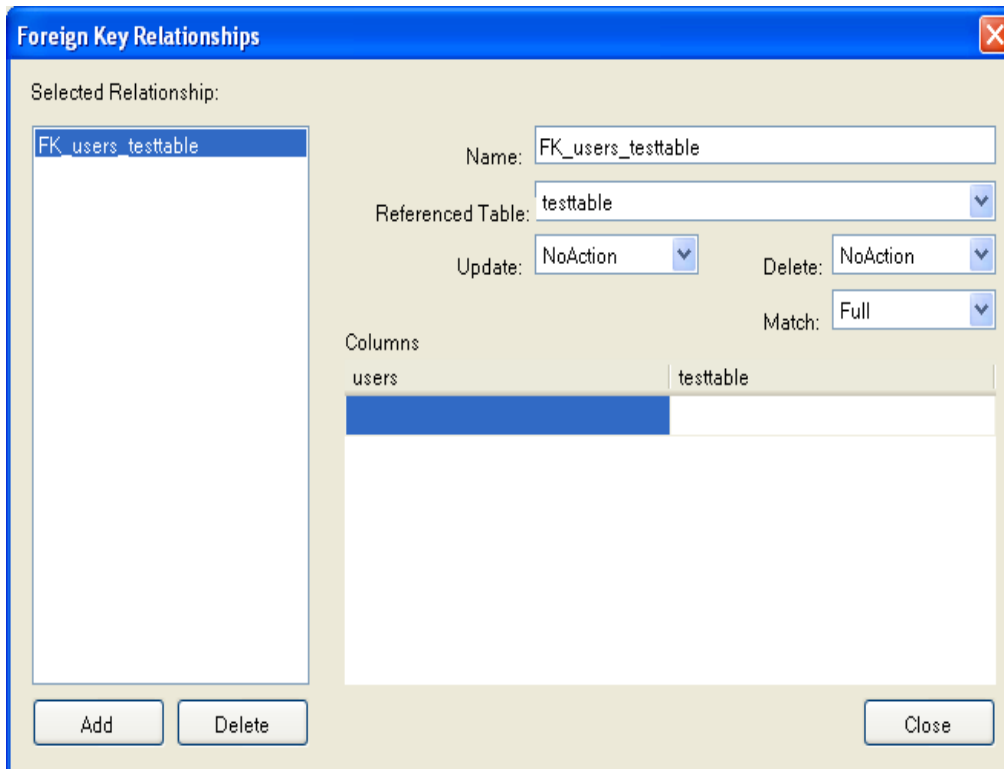
You manage **foreign keys** for **InnoDB** tables using the **Foreign Key Relationships** dialog.

To add a foreign key, select **Table Designer, Relationships...** from the main menu. This displays the **Foreign Key Relationship** dialog. Click **Add**. You can then set the foreign key name, referenced table name, foreign key columns, and actions upon update and delete.

To remove a foreign key, select it in the list box on the left, and click the **Delete** button.

To change foreign key settings, select the required foreign key in the list box on the left. The detailed information about the foreign key is displayed in the right hand panel. Change the desired values.

**Figure 21.10. Foreign Key Relationships Dialog**



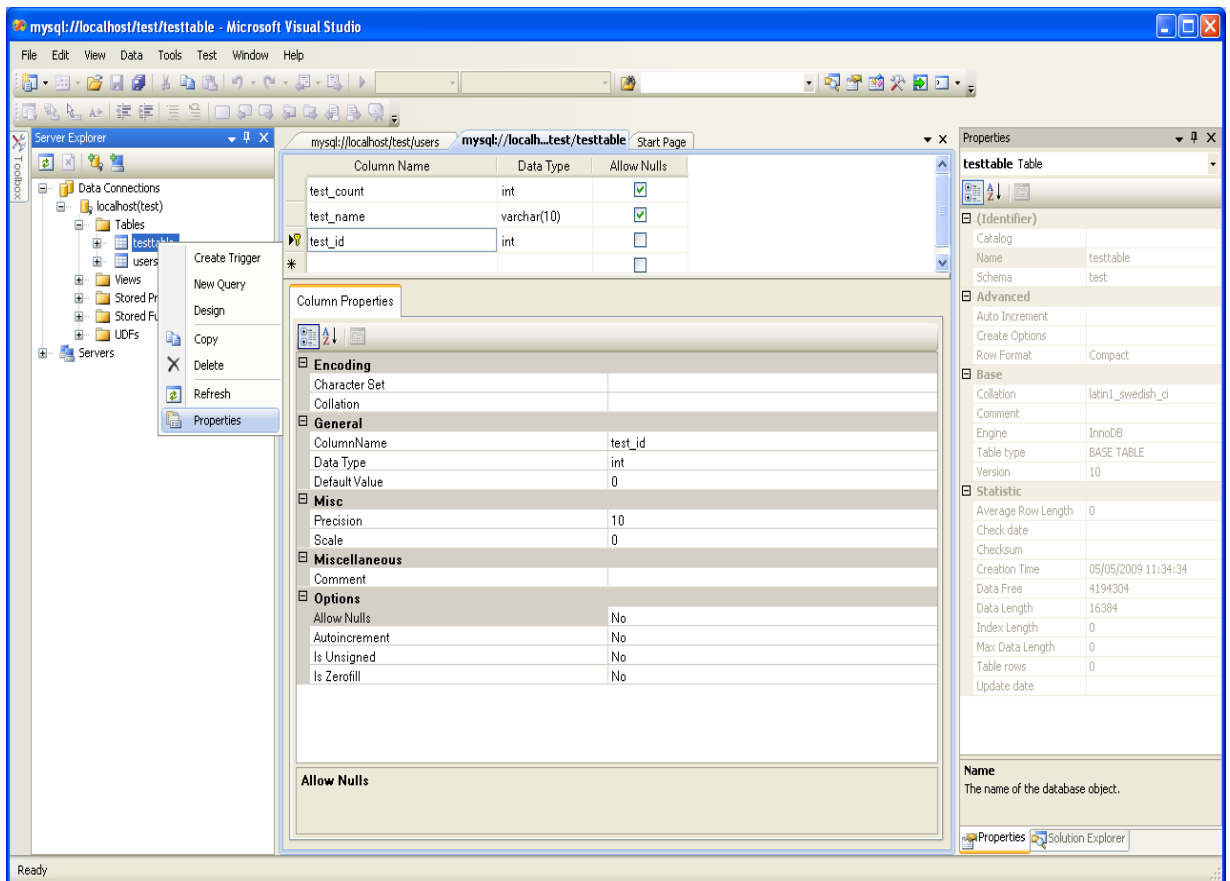
#### 21.2.3.3.4. Column Properties

The **Column Properties** tab can be used to set column options. In addition to the general column properties presented in the Column Editor, in the **Column Properties** tab you can set additional properties such as Character Set, Collation and Precision.

#### 21.2.3.3.5. Table Properties

To bring up Table Properties select the table and right-click to activate the context menu. Select **Properties**. The **Table Properties** dockable window will be displayed.

Figure 21.11. Table Properties Menu Item

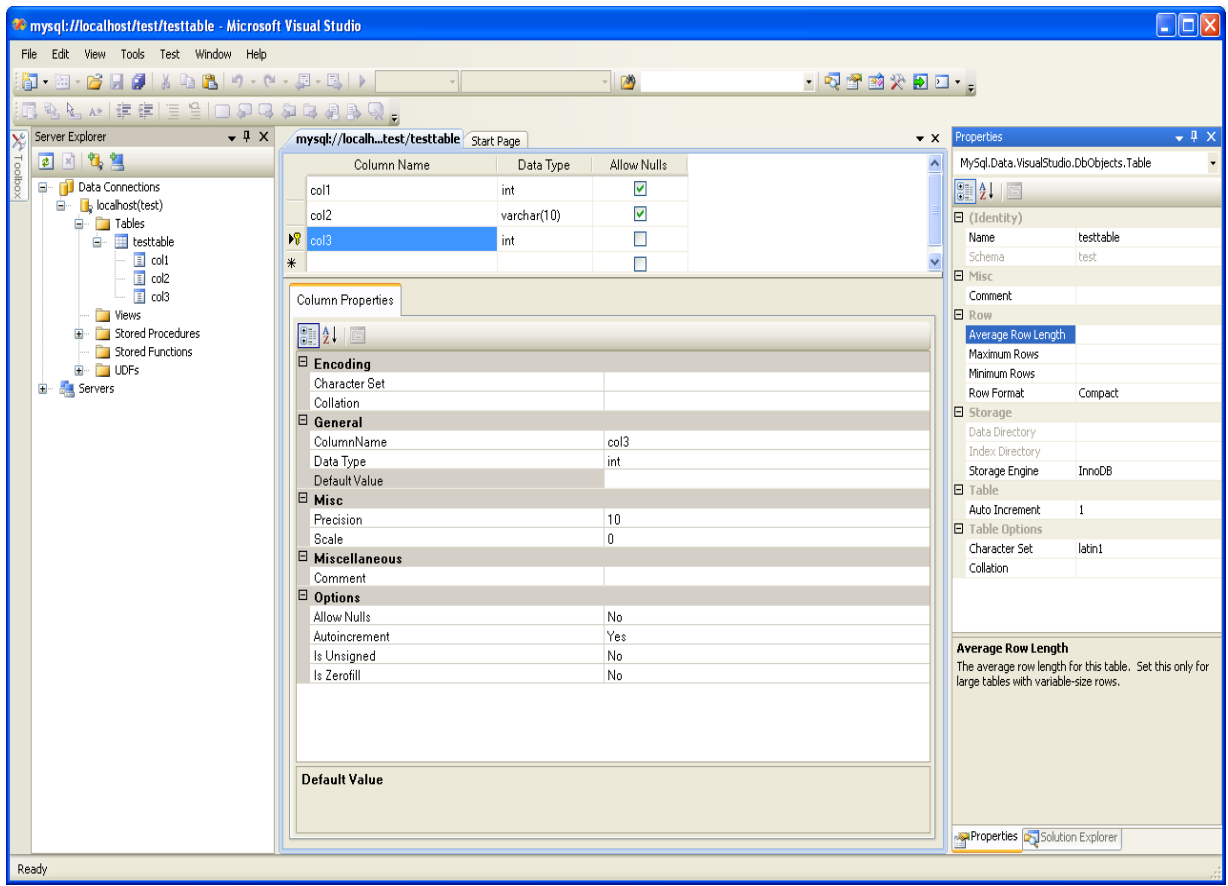


The following table properties can be set:

- Auto Increment.
- Average Row Length.
- Character Set.
- Collation.
- Comment.
- Data Directory.
- Index Directory.
- Maximum Rows.
- Minimum Rows.
- Name.
- Row Format.
- Schema.
- [Storage Engine](#). Note that in MySQL 5.5 and higher, the default storage engine for new tables is [InnoDB](#). See [Section 14.2.1.1, “InnoDB as the Default MySQL Storage Engine”](#) for more information about the choice of storage engine, and considerations if you convert existing tables to [InnoDB](#).

The property [Schema](#) is read-only.


Figure 21.12. Table Properties



### 21.2.3.4. Editing Views


To create a new view, right-click the Views node under the connection node in Server Explorer. From the node's context menu, choose the **Create View** command. This command opens the SQL Editor.

Figure 21.13. Editing View SQL



You can then enter the SQL for your view.

Figure 21.14. View SQL Added



To modify an existing view, double-click a node of the view to modify, or right-click this node and choose the **Alter View** command from a context menu. Either of the commands opens the SQL Editor.

All other view properties can be set in the Properties window. These properties are:

- Catalog

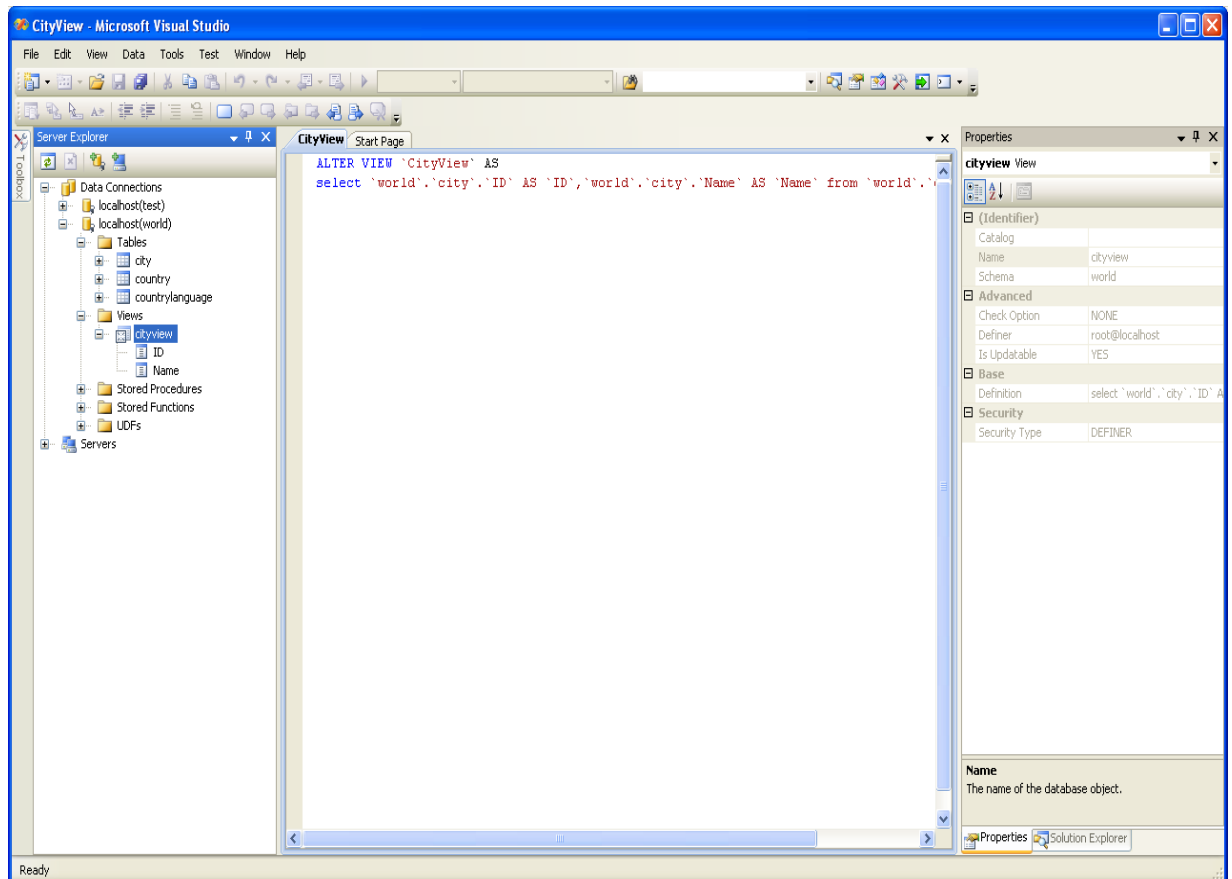
- Check Option
- Definer
- Definition
- Definer
- Is Updatable
- Name
- Schema
- Security Type

Some of these properties can have arbitrary text values, others accept values from a predefined set. In the latter case, set the desired value with an embedded combobox.

The properties `Is Updatable` and `Schema` are readonly.

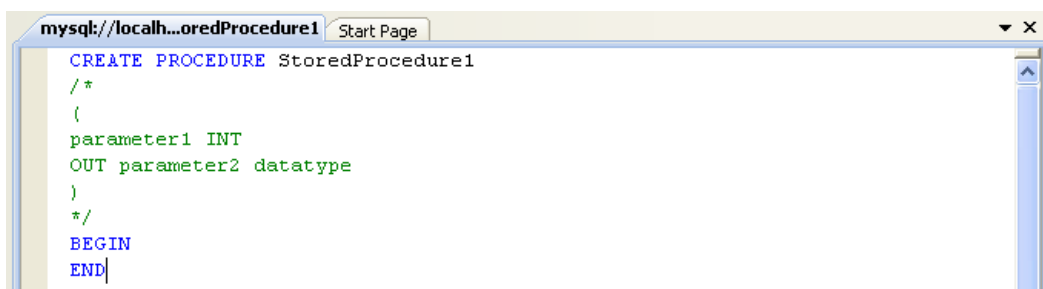
To save changes you have made, use either `Save` or `Save All` buttons of the Visual Studio main toolbar, or press **Control+S**.

**Figure 21.15. View SQL Saved**



### 21.2.3.5. Editing Stored Procedures and Functions

To create a new stored procedure, right-click the **Stored Procedures** node under the connection node in Server Explorer. From the node's context menu, choose the **Create Routine** command. This command opens the SQL Editor.

**Figure 21.16. Edit Stored Procedure SQL**

To create a new stored function, right-click the **Functions** node under the connection node in Server Explorer. From the node's context menu, choose the **Create Routine** command.

To modify an existing stored routine (procedure or function), double-click the node of the routine to modify, or right-click this node and choose the **Alter Routine** command from the context menu. Either of the commands opens the SQL Editor.

To create or alter the routine definition using SQL Editor, type this definition in the SQL Editor using standard SQL. All other routine properties can be set in the Properties window. These properties are:

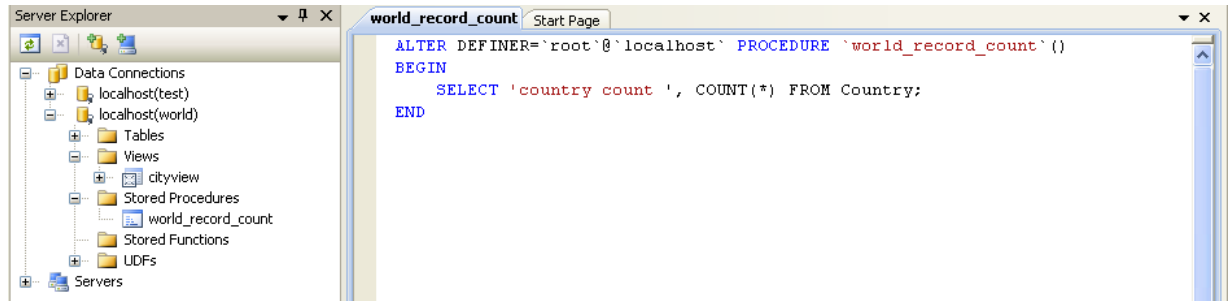
- Body
- Catalog
- Comment
- Creation Time
- Data Access
- Definer
- Definition
- External Name
- External Language
- Is Deterministic
- Last Modified
- Name
- Parameter Style
- Returns
- Schema
- Security Type
- Specific Name
- SQL Mode
- SQL Path
- Type

Some of these properties can have arbitrary text values, others accept values from a predefined set. In the latter case, set the desired value using the embedded combo box.

You can also set all the options directly in the SQL Editor, using the standard `CREATE PROCEDURE` or `CREATE FUNCTION` statement. However, it is recommended to use the Properties window instead.

To save changes you have made, use either **Save** or **Save All** buttons of the Visual Studio main toolbar, or press **Control+S**.

**Figure 21.17. Stored Procedure SQL Saved**



To observe the runtime behavior of a stored routine and debug any problems, use the Stored Procedure Debugger (available in Connector/Net 6.6 and higher). See [Section 21.2.3.8, “Debugging Stored Procedures and Functions”](#) for details.

### 21.2.3.6. Editing Triggers

To create a new trigger, right-click the node of the table in which to add the trigger. From the node's context menu, choose the **Create Trigger** command. This command opens the SQL Editor.

To modify an existing trigger, double-click the node of the trigger to modify, or right-click this node and choose the **Alter Trigger** command from the context menu. Either of the commands opens the SQL Editor.

To create or alter the trigger definition using SQL Editor, type the trigger statement in the SQL Editor using standard SQL.



#### Note

Enter only the trigger statement, that is, the part of the `CREATE TRIGGER` query that is placed after the `FOR EACH ROW` clause.

All other trigger properties are set in the Properties window. These properties are:

- Definer
- Event Manipulation
- Name
- Timing

Some of these properties can have arbitrary text values, others accept values from a predefined set. In the latter case, set the desired value using the embedded combo box.

The properties `Event Table`, `Schema`, and `Server` in the Properties window are read-only.

To save changes you have made, use either **Save** or **Save All** buttons of the Visual Studio main toolbar, or press **Control+S**. Before changes are saved, you will be asked to confirm the execution of the corresponding SQL query in a confirmation dialog.

To observe the runtime behavior of a stored routine and debug any problems, use the Stored Procedure Debugger (available in Connector/Net 6.6 and higher). See [Section 21.2.3.8, “Debugging Stored Procedures and Functions”](#) for details.



### 21.2.3.7. Editing User Defined Functions (UDF)

To create a new User Defined Function (UDF), right-click the **UDFs** node under the connection node in Server Explorer. From the node's context menu, choose the [Create UDF](#) command. This command opens the UDF Editor.

To modify an existing UDF, double-click the node of the UDF to modify, or right-click this node and choose the [Alter UDF](#) command from the context menu. Either of these commands opens the UDF Editor.

The UDF editor enables you to set the following properties:

- Name
- So-name (DLL name)
- Return type
- Is Aggregate

There are text fields for both names, a combo box for the return type, and a check box to indicate if the UDF is aggregate. All these options are also accessible using the Properties window.

The property `Server` in the Properties window is read-only.

To save changes you have made, use either `Save` or `Save All` buttons of the Visual Studio main toolbar, or press **Control+S**. Before changes are saved, you will be asked to confirm the execution of the corresponding SQL query in a confirmation dialog.

To observe the runtime behavior of a stored routine and debug any problems, use the Stored Procedure Debugger (available in Connector/Net 6.6 and higher). See [Section 21.2.3.8, "Debugging Stored Procedures and Functions"](#) for details.

### 21.2.3.8. Debugging Stored Procedures and Functions

The stored procedure debugger, new in Connector/Net 6.6, provides facilities for setting breakpoints, stepping into individual statements (Step Into, Step Out, Step Over), evaluating and changing local variable values, evaluating breakpoints, and other typical debugging tasks.

#### Installing the Debugger

To enable the stored procedure debugger, install Connector/Net 6.6 or higher and choose the **Complete** option.

#### Privileges

The debugger recreates at the start of each debug session a `serversidedebugger` database in your server. This database helps to track the instrumented code and implement observability logic in the debugged routine. Your current connection needs to have privileges to create that database, and its associated stored routines, functions, and tables.

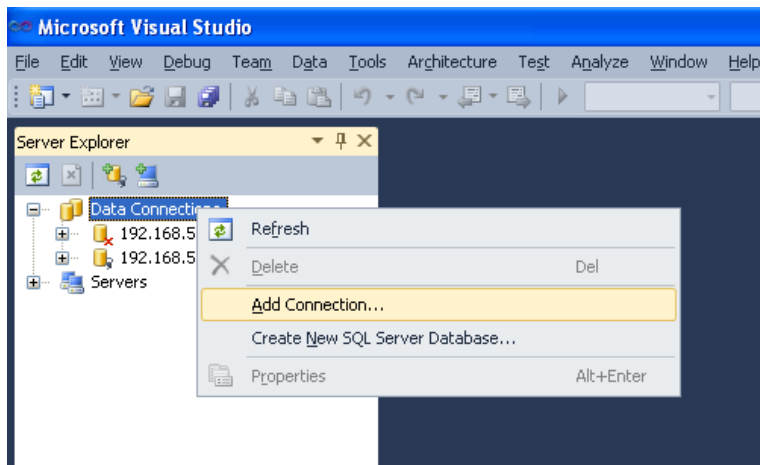
The debugger makes changes behind the scenes to temporarily add instrumentation code to the stored routines that you debug. You must have the `ALTER ROUTINE` privilege for each stored procedure, function, or trigger that you debug. (Including procedures and functions that are called, and triggers that are fired, by a procedure that you are debugging.)

#### Starting the Debugger

To start the debugger, follow these steps:

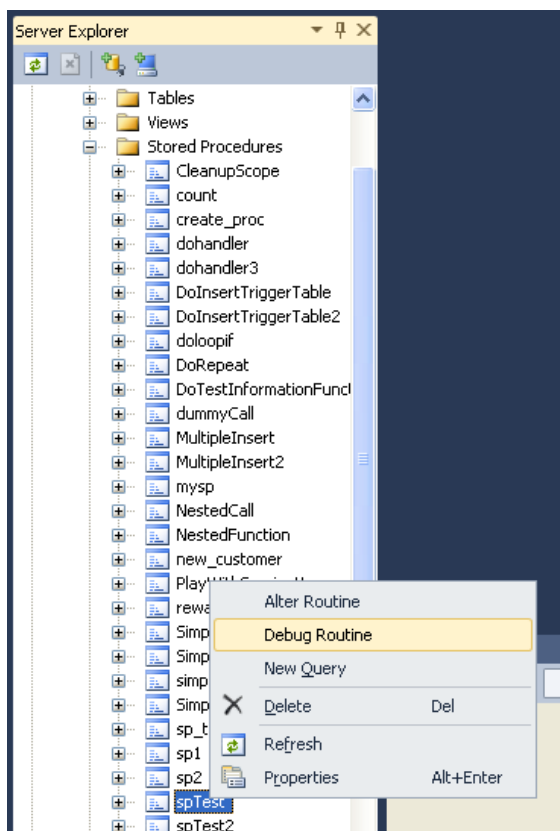
- Choose a connection in the Visual Studio Server Explorer.

**Figure 21.18. Connection Dialog**



- Expand the **Stored Procedures** folder. Only stored procedures can be debugged directly. To debug a user-defined function, create a stored procedure that calls the function.
- Click on a stored procedure node, then right-click and from the context menu choose **Debug Routine**.

**Figure 21.19. Choose a Stored Routine to Debug**

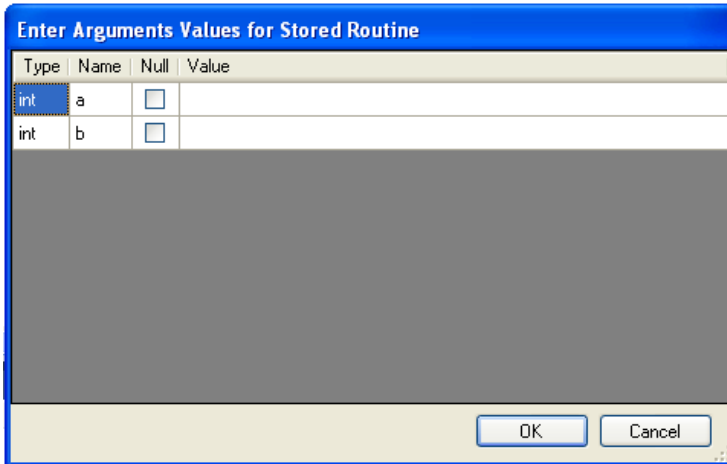


At this point, Visual Studio switches to debug mode, opening the source code of the routine being debugged in step mode, positioned on the first statement.

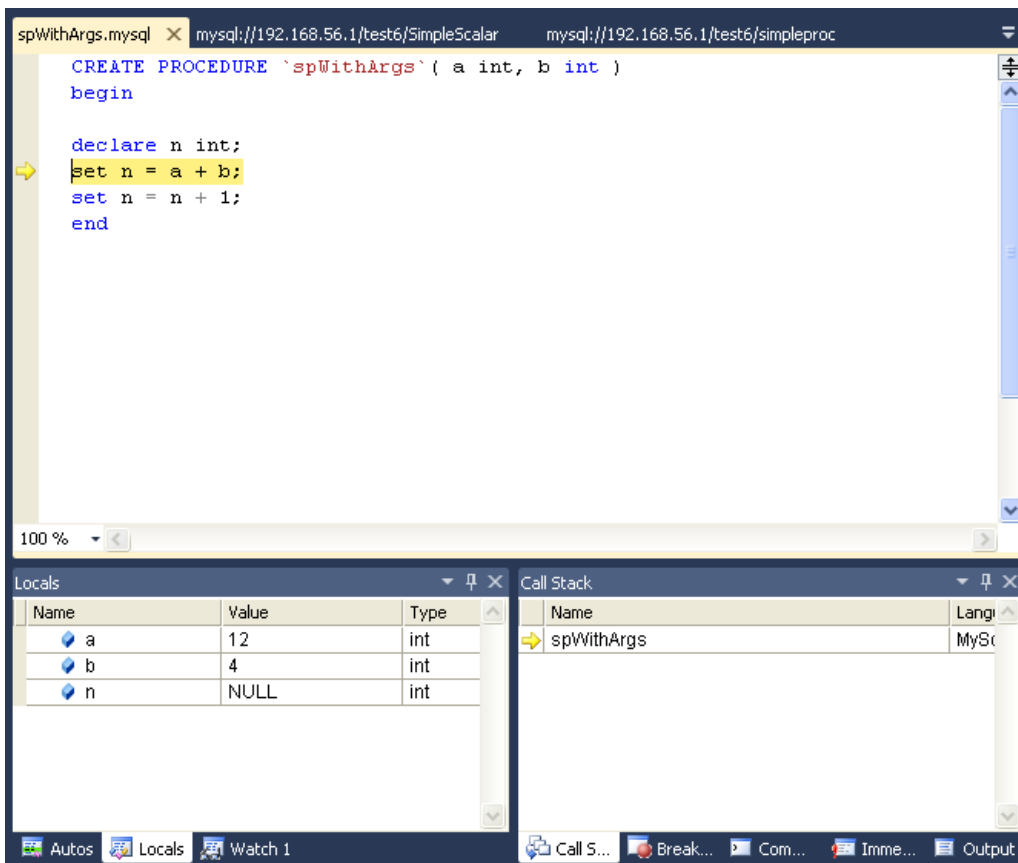
If the initial routine you debug has one or more arguments, a popup will show up with a grid (a row per each argument and three columns: one for the argument, one for the argument value (this is editable)

and one for nullifying that argument value (a checkbox). After setting up all the argument values, you can press **OK** to start the debug session, or **Cancel** to cancel the debug session.

**Figure 21.20. Setting Arguments (1 of 2)**



**Figure 21.21. Setting Arguments (2 of 2)**



## How the Debugger Works

To have visibility into the internal workings of a stored routine, the debugger prepares a special version of the procedure, function, or trigger being debugged, instrumented with extra code to keep track of the current line being stepped into and the values of all the local variables. Any other stored procedures, functions, or triggers called from the routine being debugged are instrumented the same way. The debug versions of the routines are prepared for you automatically, and when the debug session ends (by either pressing **F5** or **Shift+F5**), the original versions of the routines are automatically restored.

A copy of the original version of each instrumented routine (the version without instrumentation) is stored in the `AppData\Roaming\MySQLDebuggerCache` folder for the current Windows user (the path returned by calling `System.Environment.GetFolderPath(Environment.SpecialFolder.ApplicationData)` in .NET, plus appending `MySQLDebuggerCache`). There is one file for each instrumented routine, named `routine_name.mysql`. For example, in Windows 7, for a user named `fergs`, the path is `C:\Users\fergs\AppData\Roaming\MySQLDebuggerCache`.

Two threads are used, one for the debugger and one for the routine being debugged. The threads run in strict alternation, switching between the debugger and the routine as each statement is executed in the stored routine.

## Basic Debugging Operations

The debugger has the same look and feel as the standard Visual Studio debuggers for C#, VB.NET or C++. In particular, the following are true:

### Locals and Watches

- To show the Locals tab, choose the menu item **Debug -> Windows -> Locals**.

The **Locals** tab lists all the variables available in the current scope: variables defined with `DECLARE` at any point in the routine, argument parameters, and session variables that are referenced.

- If the last step operation changes the value of a local, its value will be highlighted in red (until another statement is executed or stepped).
- You can change the value of any local.
- To show the **Watch** tab, choose the menu item **Debug -> Windows -> Watch**.

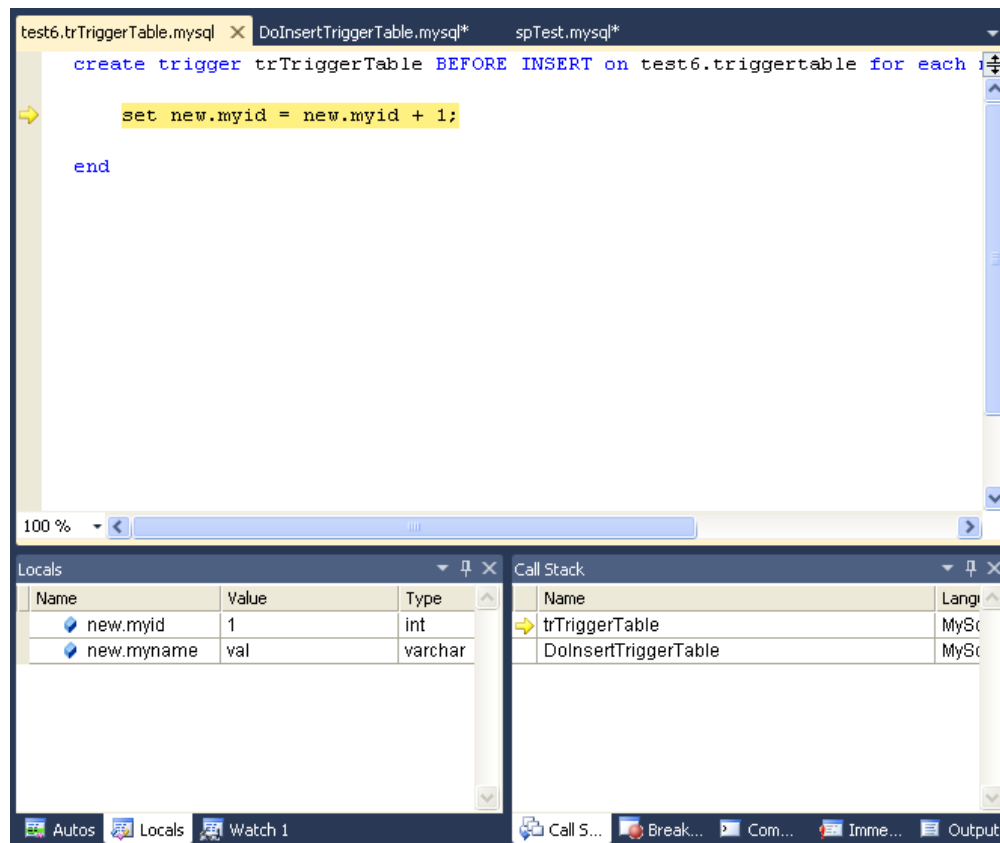
To define a watch, type any valid MySQL expression, optionally including function calls. If the watch evaluation makes sense in the current context (current stack frame), it will show its value, otherwise it will show an error message in the same row the watch was defined.

- When debugging a trigger, in addition to any locals declared or session variables referenced, the new and old object (when applicable) will be listed. For example in a trigger for `INSERT`, for a table defined like:

```
create table t1( id int, myname varchar( 50 ));
```

the locals will list the extra variables `new.id` and `new.myname`. For an `UPDATE` trigger, you will also get the extra variables `old.id` and `old.myname`. These variables from the new and old objects can be manipulated the same way as any ordinary local variable.

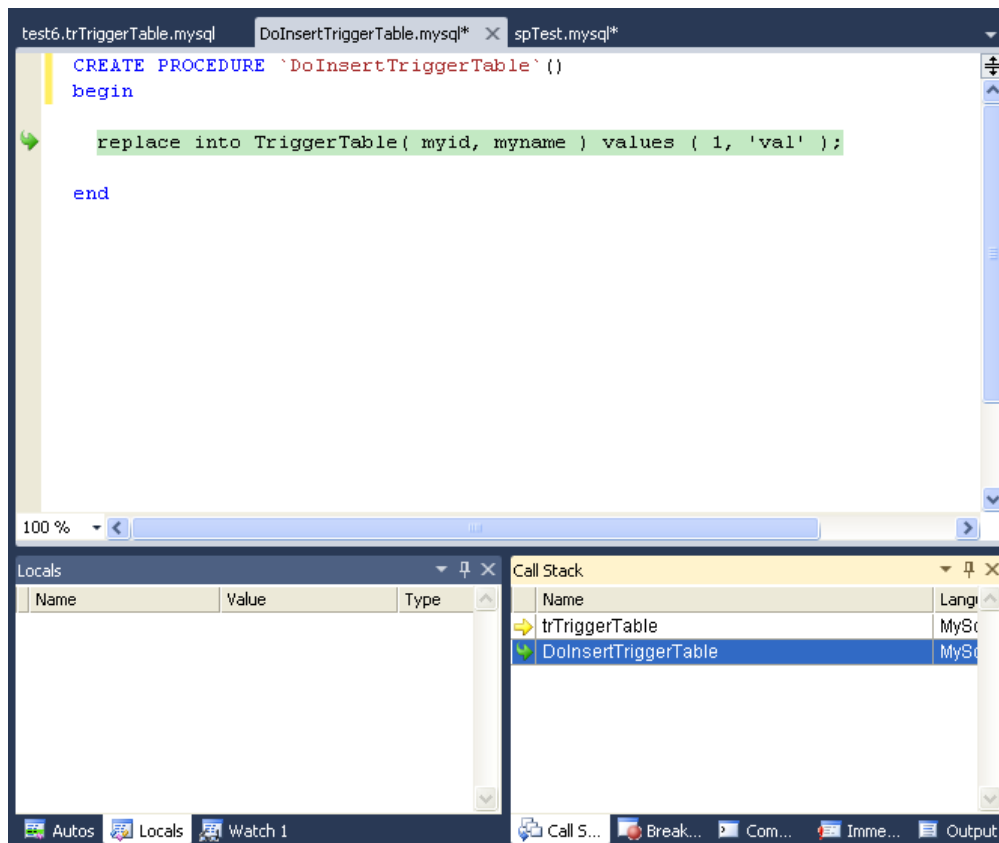
Figure 21.22. Debugging a Trigger



### Call Stack

- To show the **Call Stack** tab, choose the menu item **Debug -> Windows -> Call Stack**.
- The stack trace (in the **Call Stack** tab) will list all the stack traces, one for each routine invocation. The one with a yellow mark is the current stepping point. Clicking in another will activate in the editor the tab for that routine source, highlighting in green the last statement stepped.

Figure 21.23. Call Stack



### Stepping

- Stepping of a new routine starts in the first executable instruction (excluding declares, handlers, cursor declarations, and so on).

Figure 21.24. Debug Stepping

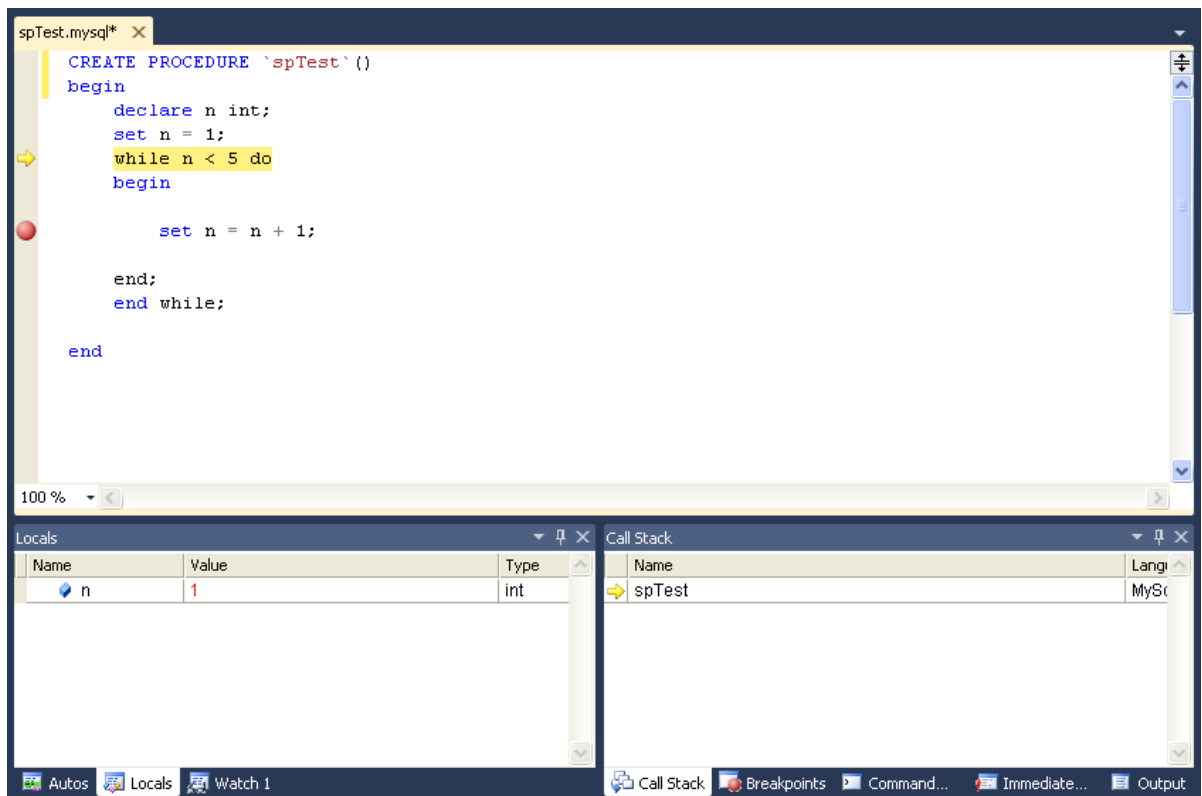


Figure 21.25. Function Stepping (1 of 2)

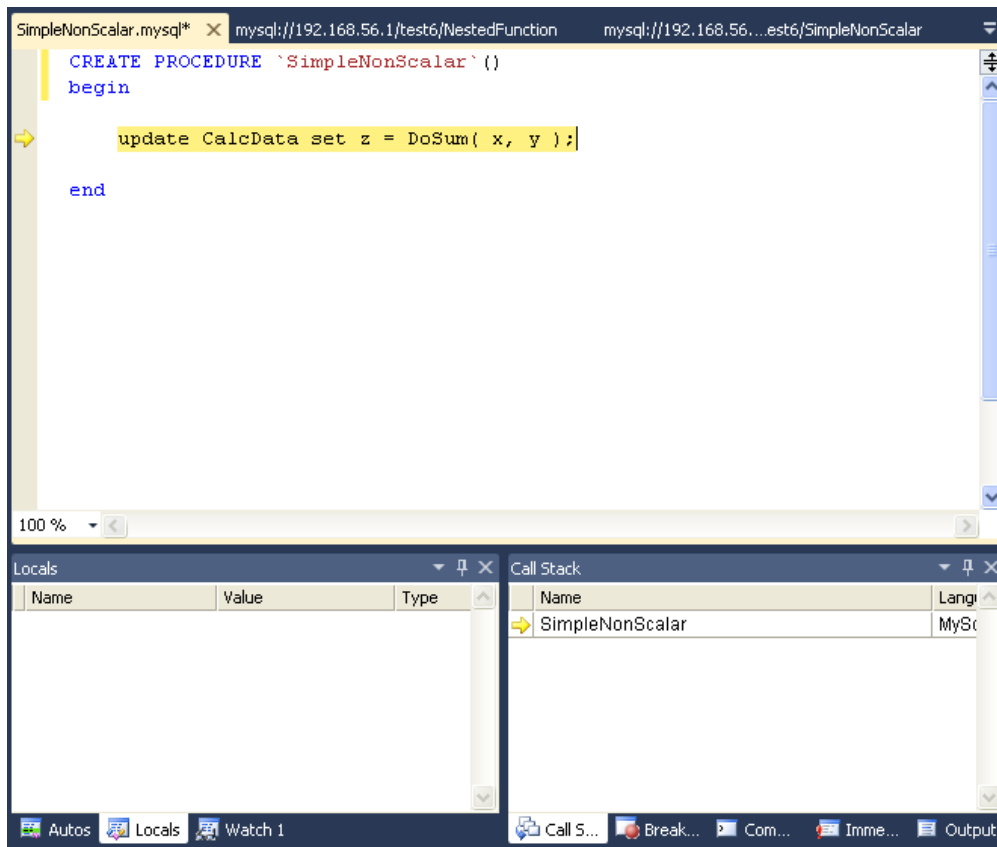
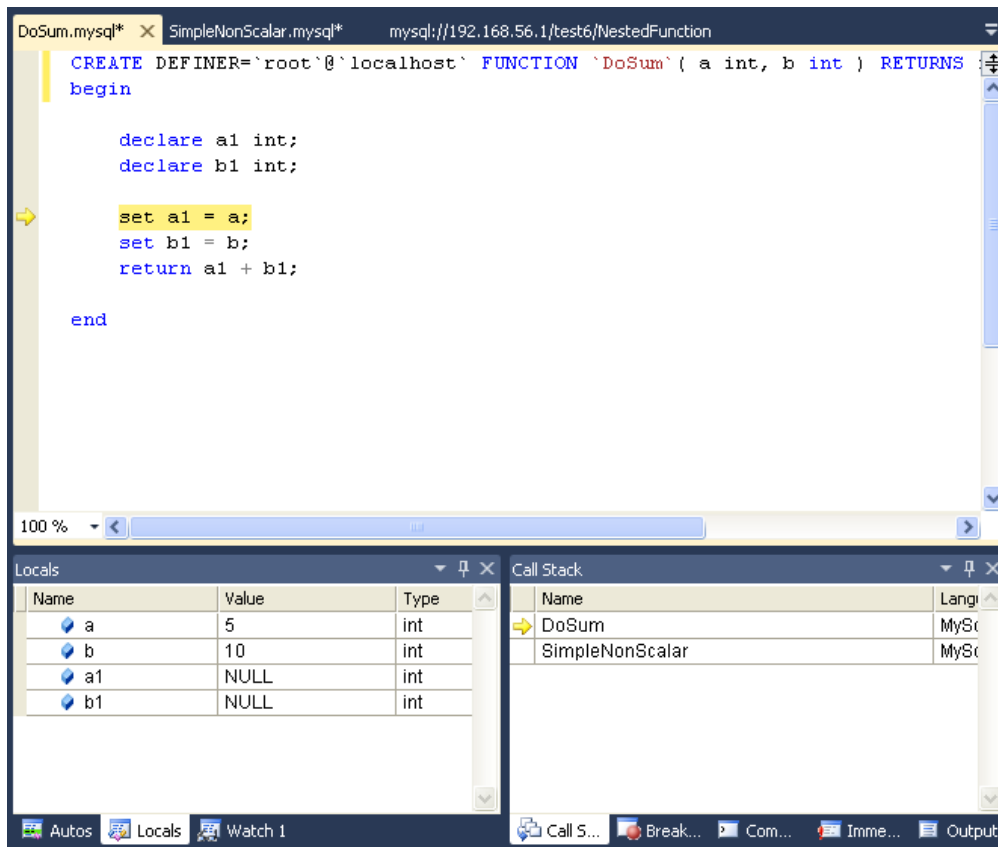




Figure 21.26. Function Stepping (2 of 2)



- To step into the code of a condition handler, the condition must be triggered in the rest of the MySQL routine.
- The next statement to be executed is highlighted in yellow.
- To continue stepping, you can choose between **Step Into** (by pressing **F11**), **Step Out** (by pressing **F10**) or **Step Over** (by pressing **Shift+F11**).
- You can step out of any of functions, triggers or stored procedures. If you step from the main routine, it will run that routine to completion and finish the debug session.
- You can step over stored procedure calls, stored functions, and triggers. (To step over a trigger, step over the statement that would cause the trigger to fire.)
- When stepping into a single statement, the debugger will step into each individual function invoked by that statement and each trigger fired by that statement. The order in which they are debugged is the same order in which the MySQL server executes them.
- You can step into triggers triggered from [INSERT](#), [DELETE](#), [UPDATE](#), and [REPLACE](#) statements.
- Also, the number of times you enter into a stored function or trigger depends on how many rows are evaluated by the function or affected by the trigger. For example, if you press **F11 (Step Into)** into an [UPDATE](#) statement that modifies three rows (calling a function for a column in the [SET](#) clause, thus invoking the function for each of the three rows), you will step into into that function three times in succession, once for each of the rows. You can accelerate this debug session by disabling any breakpoints defined in the given stored function and pressing **Shift+F11** to step out. In this example, the order in which the different instances of the stored function are debugged is server-specific: the same order used by the current MySQL server instance to evaluate the three function invocations.

### Breakpoints

- To show the **Breakpoints** tab, choose the menu item **Debug -> Windows -> Breakpoints**.
- The **Breakpoints** tab will show all the breakpoints defined. From here, you can enable and disable breakpoints one by one or all at once (using the toolbar on top of the **Breakpoints** tab).
- You can define new breakpoints only in the middle of a debug session. Click in the left gray border of any MySQL editor, or click anywhere in a MySQL editor and press **F9**. In the familiar Visual Studio way, you press **F9** once to create a breakpoint in that line, and press it again to remove that breakpoint.
- Once a breakpoint is defined, it will appear enabled (as filled red circle left to the current row if that line is a valid statement to put a breakpoint) or disabled (as a non-filled red circle left to the current row if that row is not valid to put a breakpoint).
- To define conditional breakpoints, after creating the breakpoint, right click in the red dot and choose **Condition...**. There you can put any valid MySQL expression and state if the condition is **Is True** or **Has changed**. The former will trigger the breakpoint every time the condition is true, the latter every time the condition value has changed. (If you define a conditional breakpoint, it is not enough to step into the line with the breakpoint defined to trigger such a breakpoint.)

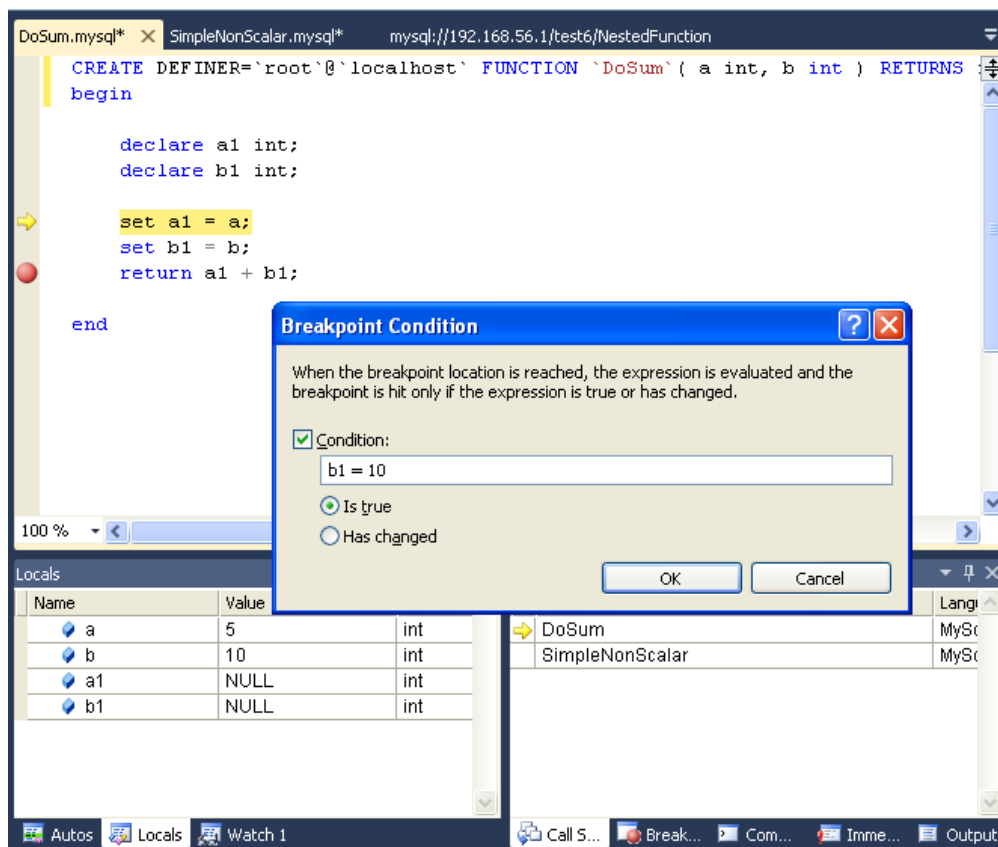
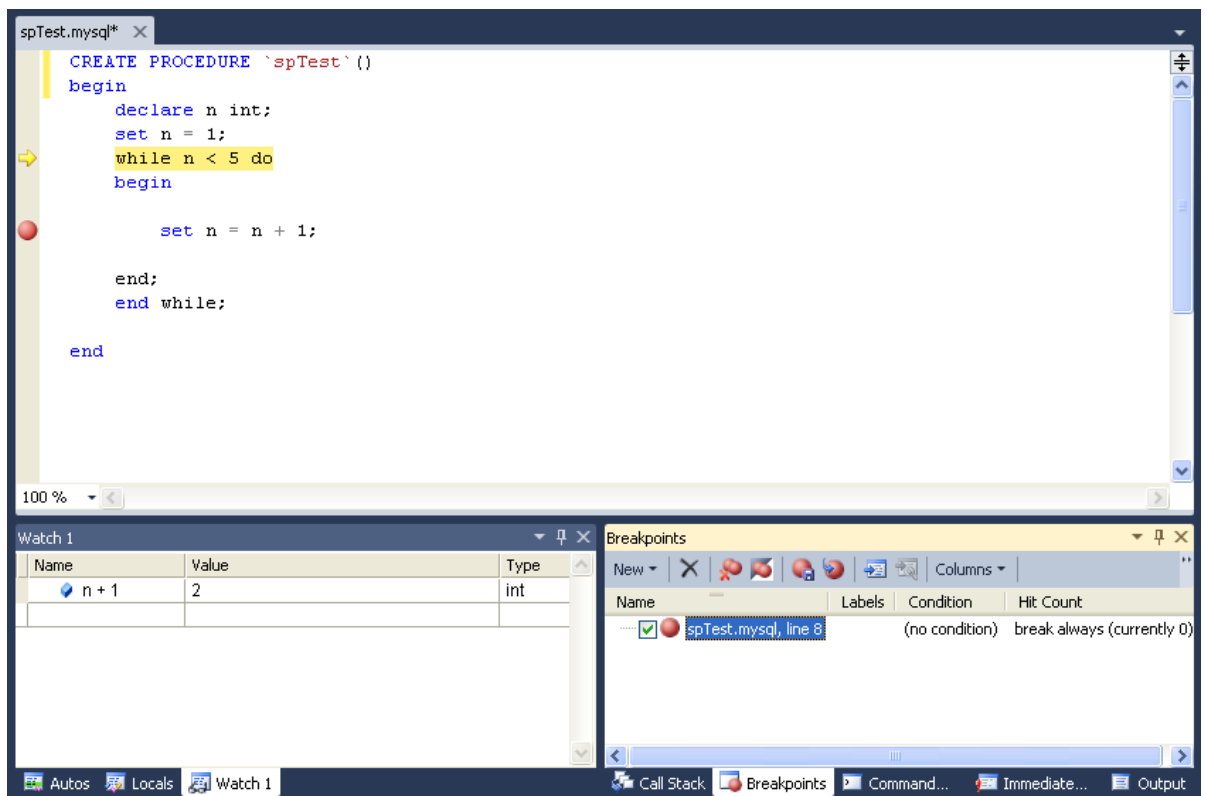
**Figure 21.27. Conditional Breakpoints**

Figure 21.28. Expressions and Breakpoints



- To define pass count breakpoints, after creating the breakpoint, right click in the red dot and choose **Hit Count...**. In the popup dialog, define the specific condition to set. For example, **break when the hit count is equal to** and a value 3 will trigger the breakpoint the third time it is hit.

## Other Features

- To abort the debug session (and the execution of the current call stack of routines), press Shift+F5.
- To run the routine to completion (or until next breakpoint hit), press F5.
- For all functionality you can use (in addition to the shortcuts documented), see the options in the **Debug** menu of Visual Studio.

## Limitations

- Code being debugged must not use `get_lock` or `release_lock` MySQL functions, since they are used internally by the debugger infrastructure to synchronize the debugger and the debugged routine.
- Code being debugged must avoid using any transaction code (`START TRANSACTION`, `COMMIT`, `ROLLBACK`) due to the possibility of wiping out the contents of the debugger tables. (This limitation may be removed in the future).
- You cannot debug the routines in the `serversidedebugger` database.
- The MySQL server running the routine being debugged can be any version between 5.0 and 5.6, running on Windows, Linux, or any other supported platform.
- We recommend always running debug sessions on test and development servers, rather than against a MySQL production server, because debugging can cause temporary performance issues or even deadlocks. The instrumented versions of the routines being debugged use locks, that the rest of the production code may not be aware of.

## Keyboard Shortcuts

The following list summarizes the keyboard shortcuts for debugging:

- **F9** Toggles breakpoints
- **F11**: Step into once
- **F10**: Step over once
- **Shift+F11**: Step out once
- **F5**: Run
- **Shift+F5**: Abort current debug session

### 21.2.3.9. Cloning Database Objects

Tables, views, stored procedures, and functions can be cloned using the appropriate Clone command from the context menu: [Clone Table](#), [Clone View](#), [Clone Routine](#). The clone commands open the corresponding editor for a new object: the **Table Editor** for cloning a table, and the **SQL Editor** for cloning a view or a routine.

The editor is filled with values of the original object. You can modify these values in the usual manner.

To save the cloned object, use either **Save** or **Save All** buttons of the Visual Studio main toolbar, or press **Control+S**. Before changes are saved, you will be asked to confirm the execution of the corresponding SQL query in a confirmation dialog.

### 21.2.3.10. Dropping Database Objects

Tables, views, stored routines, triggers, and UDFs can be dropped with the appropriate Drop command selected from its context menu: [Drop Table](#), [Drop View](#), [Drop Routine](#), [Drop Trigger](#), [Drop UDF](#).

You will be asked to confirm the execution of the corresponding drop query in a confirmation dialog.

You can only drop a single object at a time.

### 21.2.3.11. Using the ADO.NET Entity Framework

Connector/Net 6.0 introduced support for the ADO.NET Entity Framework. ADO.NET Entity Framework was included with .NET Framework 3.5 Service Pack 1, and Visual Studio 2008 Service Pack 1. ADO.NET Entity Framework was released on 11th August 2008.

ADO.NET Entity Framework provides an Object Relational Mapping (ORM) service, mapping the relational database schema to objects. The ADO.NET Entity Framework defines several layers, these can be summarized as:

- **Logical** - this layer defines the relational data and is defined by the Store Schema Definition Language (SSDL).
- **Conceptual** - this layer defines the .NET classes and is defined by the Conceptual Schema Definition Language (CSDL)
- **Mapping** - this layer defines the mapping from .NET classes to relational tables and associations, and is defined by Mapping Specification Language (MSL).

Connector/Net integrates with Visual Studio 2008 to provide a range of helpful tools to assist the developer.

A full treatment of ADO.NET Entity Framework is beyond the scope of this manual. If you are unfamiliar with ADO.NET, review the [Microsoft ADO.NET Entity Framework documentation](#).


Tutorials on getting started with ADO.NET Entity Framework are available. See [Section 21.2.4.5, "Tutorial: Using an Entity Framework Entity as a Windows Forms Data Source"](#) and [Section 21.2.4.6, "Tutorial: Databinding in ASP.NET using LINQ on Entities"](#).

### 21.2.3.12. MySQL Website Configuration Tool

MySQL Connector/Net 6.1 introduced the MySQL Website Configuration Tool. This is a facility available in Visual Studio that enables you to configure the Membership, Role, Session State and Profile Provider, without editing configuration files. You set your configuration options within the tool, and the tool modifies your `web.config` file accordingly.

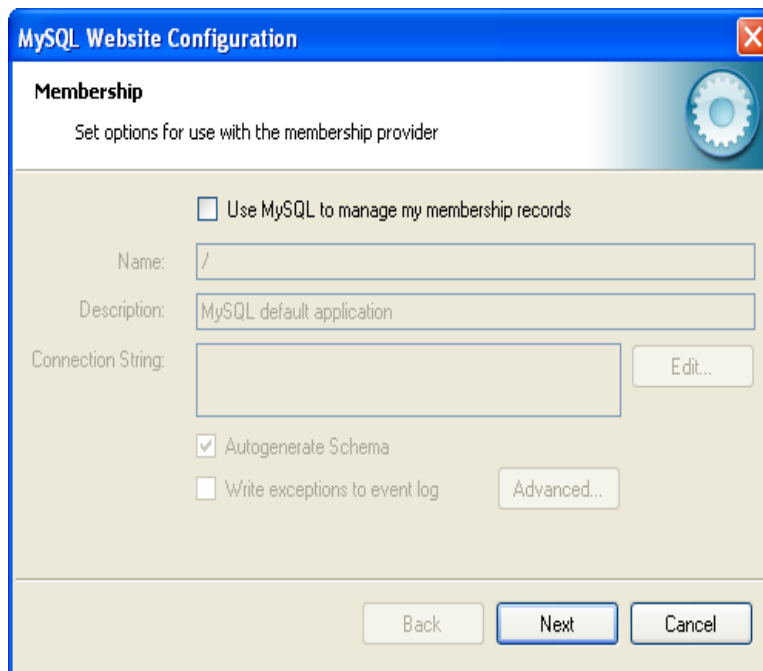
The MySQL Website Configuration Tool appears as a small icon on the Solution Explorer toolbar in Visual Studio, as show by the following screenshot:

**Figure 21.29. MySQL Website Configuration Tool**



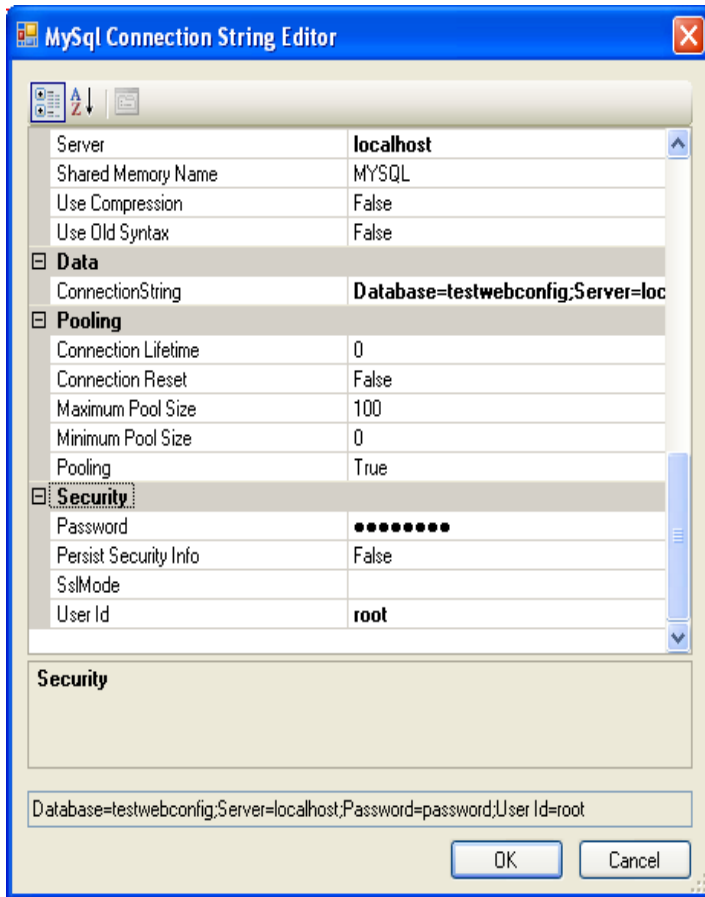
Clicking the Website Configuration Tool icon launches the wizard and displays the first screen:

**Figure 21.30. MySQL Website Configuration Tool - Membership**



This allows you to enable use of the MySQL Membership Provider. Click the check box to enable this. You can now enter the name of the application that you are creating the configuration for. You can also enter a description for the application.

You can then click the `Edit...` button to launch the Connection String Editor:

**Figure 21.31. MySQL Website Configuration Tool - Connection String Editor**

Note that if you have already defined a connection string for the providers manually in `web.config`, or previously using the tool, this will be automatically loaded and displayed, and can then be modified in this dialog.

You can also ensure that the necessary schemas are created automatically for you by selecting the Autogenerate Schema check box. These schemas are used to store membership information. The database used to storage is the one specified in the connection string.

You can also ensure that exceptions generated by the application will be written to the Windows event log by selecting the **Write exceptions to event log** check box.

Clicking the **Advanced...** button launches a dialog that enables you to set Membership Options. These options dictate such variables as password length required when a user signs up, whether the password is encrypted and whether the user can reset their password or not.

**Figure 21.32. MySQL Website Configuration Tool - Advanced Options**

The screenshot shows the 'Membership Options' dialog box with the following settings:

- Password Options:**
  - Minimum Required Password Length: 7
  - Min Required Non-alphanumeric Characters: 1
  - Password strength regular expression: (empty text box)
  - Password Format: Clear
- User Options:**
  - Require Question/Answer
  - Enable Password Reset
  - Requires Unique Email
  - Enable Password Retrieval
- Access Options:**
  - Maximum Invalid Password Attempts: 5
  - Password Attempt Window: 10

Buttons: OK, Cancel

Once information has been set up as required for configuration of the Membership Provider, the **Next** button can be clicked to display the Roles Provider screen:

**Figure 21.33. MySQL Website Configuration Tool - Roles**


The screenshot shows the 'Roles' configuration dialog box with the following settings:

- Roles:** Set options for use with the role provider
- Use MySQL to manage my roles
- Name: /
- Description: (empty text box)
- Connection String: (empty text box)
- Autogenerate Schema
- Write exceptions to event log

Buttons: Back, Next, Cancel

Again the connection string can be edited, a description added and Autogenerate Schema can be enabled before clicking **Next** to go to the Profiles Provider screen:

**Figure 21.34. MySQL Website Configuration Tool - Profiles**



The screenshot shows the 'MySQL Website Configuration' dialog box with the 'Profiles' tab selected. The title bar reads 'MySQL Website Configuration' and the subtitle is 'Profiles'. Below the subtitle, it says 'Set options for use with the profile provider'. There is a gear icon in the top right corner. The main area contains the following options:


- Use MySQL to manage my profiles
- Name: /
- Description: Profiles
- Connection String: [Empty text box] with an 'Edit...' button to its right.
- Autogenerate Schema

At the bottom, there are three buttons: 'Back', 'Finish', and 'Cancel'.

This screen display similar options to the previous screens.

Click **Next** to proceed to the Session State configuration page:

**Figure 21.35. MySQL Website Configuration Tool - Session State**



The screenshot shows the 'MySQL Website Configuration' dialog box with the 'Session State' tab selected. The title bar reads 'MySQL Website Configuration' and the subtitle is 'Session State'. Below the subtitle, it says 'Set options for use with the session state provider'. There is a gear icon in the top right corner. The main area contains the following options:

- Use MySQL to manage my ASP.Net session date
- Name: /
- Description: [Empty text box]
- Connection String: [Empty text box] with an 'Edit...' button to its right.
- Autogenerate Schema
- Write exceptions to event log

At the bottom, there are three buttons: 'Back', 'Finish', and 'Cancel'.


Once you have set up the Session State Provider as required, click **Finish** to exit the wizard.

At this point, select the **Authentication Type** to **From Internet**. Launch the ASP.NET Configuration Tool and select the Security tab. Click the **Select authentication type** link and ensure that the



From the **internet** radio button is selected. You can now examine the database you created to store membership information. All the necessary tables will have been created for you:

Figure 21.36. MySQL Website Configuration Tool - Tables



```
mysql> show tables;
Empty set (0.01 sec)

mysql> show tables;
Empty set (0.00 sec)

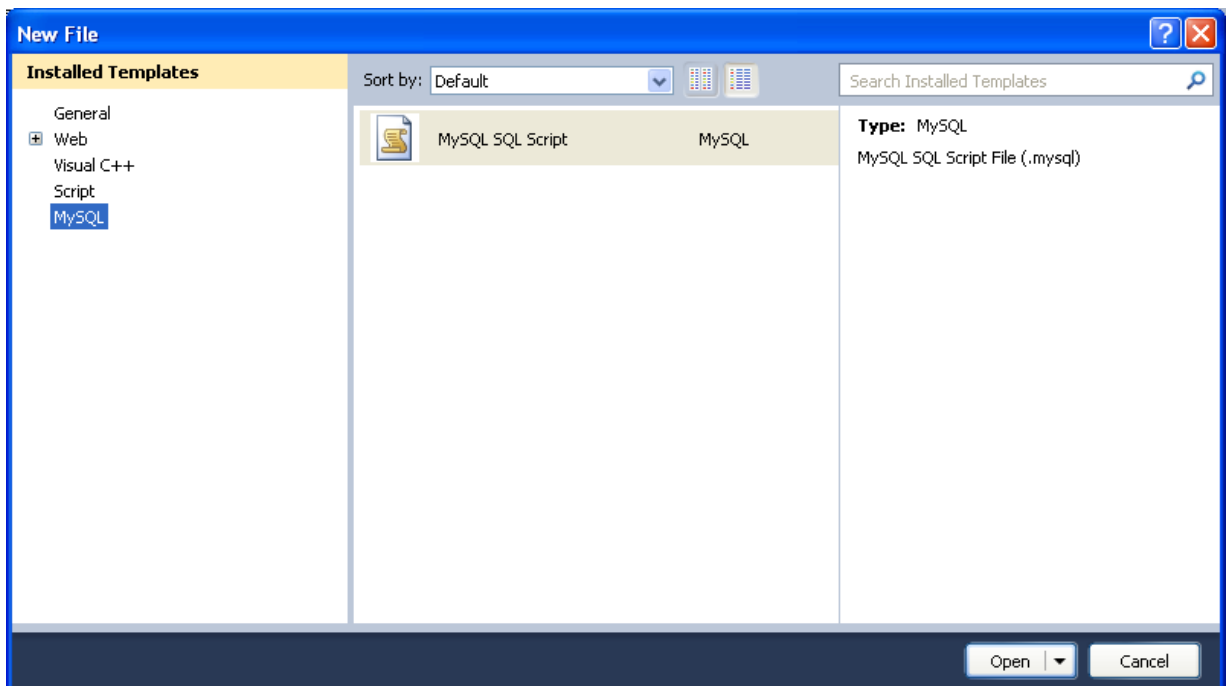
mysql> show tables;
+-----+
| Tables_in_testwebconfig |
+-----+
| my_aspnet_applications  |
| my_aspnet_membership    |
| my_aspnet_profiles      |
| my_aspnet_roles        |
| my_aspnet_schemaversion |
| my_aspnet_sessioncleanup |
| my_aspnet_sessions     |
| my_aspnet_users        |
| my_aspnet_usersinroles |
+-----+
9 rows in set (0.01 sec)

mysql>
```

### 21.2.3.13. MySQL SQL Editor

MySQL Connector/Net 6.3.2 introduced a new MySQL SQL Editor. The easiest way to invoke the editor is by selecting the **New, File** menu item from the Visual Studio main menu. This displays the **New File** dialog:

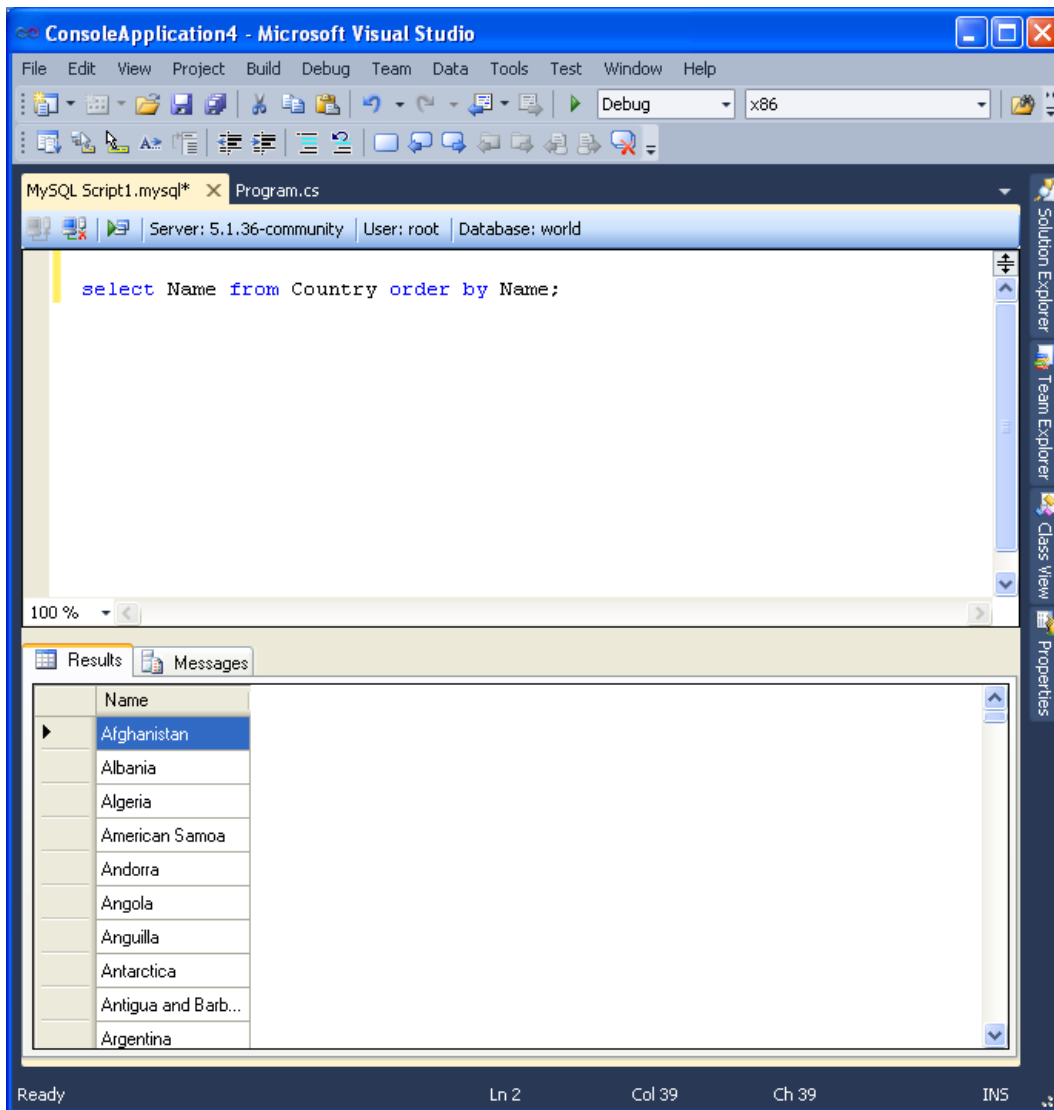
Figure 21.37. MySQL SQL Editor - New File



From the **New File** dialog, select the MySQL template, and then double-click the **MySQL SQL Script** document, or click the **Open** button.

The MySQL SQL Editor will be displayed. You can now enter SQL code as required, or connect to a MySQL server. Click the **Connect to MySQL** button in the MySQL SQL Editor toolbar. You can enter the connection details into the **Connect to MySQL** dialog that is displayed. You can enter the server name, user ID, password and database to connect to, or click the **Advanced** button to select other connection string options. Click the **Connect** button to connect to the MySQL server. To execute your SQL code against the server, click the **Run SQL** button on the toolbar.

**Figure 21.38. MySQL SQL Editor - Query**

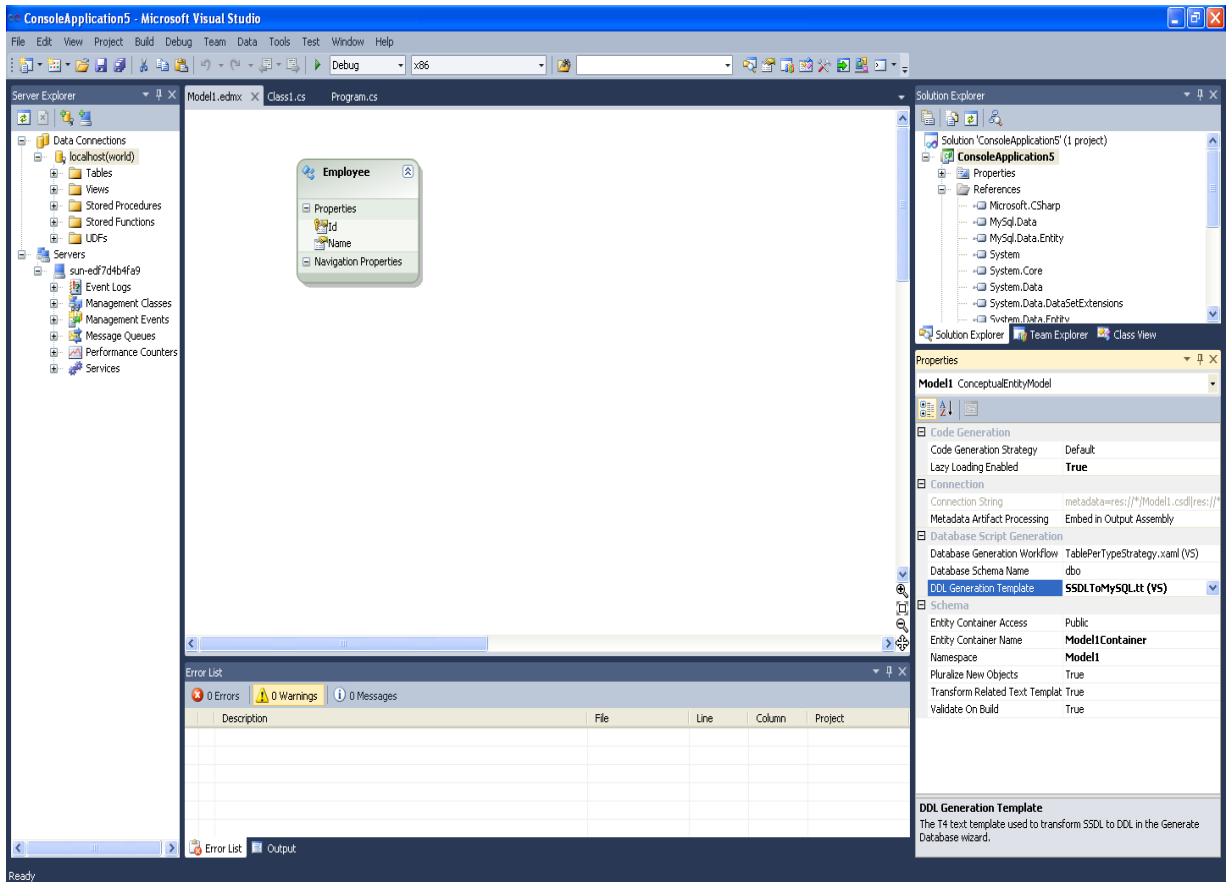


The results from any queries are displayed on the **Results** tab. Any errors are displayed on the **Messages** tab.

#### 21.2.3.14. DDL T4 Template Macro

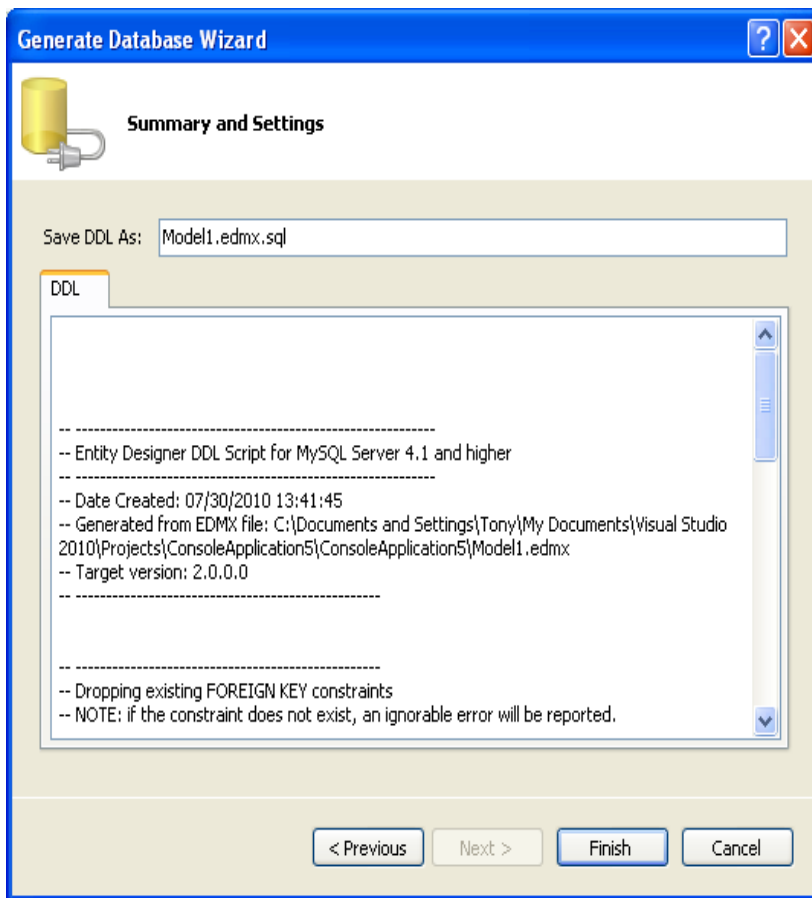
MySQL Connector/Net 6.3 introduced the ability to convert an Entity Framework model to MySQL **DDL** code. Starting with a blank model, you can develop an entity model in Visual Studio's designer. Once the model is created, you can select the model's properties, and in the Database Script Generation category of the model's properties, the property **DDL Generation** can be found. Select the value **SSDLToMySQL.tt(VS)** from the drop-down listbox.

Figure 21.39. DDL T4 Template Macro - Model Properties



Right-clicking the model design area displays a context-sensitive menu. Selecting **Generate Database from Model** from the menu displays the **Generate Database Wizard**. The wizard can then be used to generate MySQL DDL code.

Figure 21.40. DDL T4 Template Macro - Generate Database Wizard



## 21.2.4. Connector/Net Tutorials

The following tutorials illustrate how to develop MySQL programs using technologies such as Visual Studio, C#, ASP.NET, and the .NET and Mono frameworks. Work through the first tutorial to verify that you have the right software components installed and configured, then choose other tutorials to try depending on the features you intend to use in your applications.

### 21.2.4.1. Tutorial: An Introduction to Connector/Net Programming

This section provides a gentle introduction to programming with Connector/Net. The example code is written in C#, and is designed to work on both Microsoft .NET Framework and Mono.

This tutorial is designed to get you up and running with Connector/Net as quickly as possible, it does not go into detail on any particular topic. However, the following sections of this manual describe each of the topics introduced in this tutorial in more detail. In this tutorial you are encouraged to type in and run the code, modifying it as required for your setup.

This tutorial assumes you have MySQL and Connector/Net already installed. It also assumes that you have installed the World example database, which can be downloaded from the [MySQL Documentation page](#). You can also find details on how to install the database on the same page.



#### Note

Before compiling the example code, make sure that you have added References to your project as required. The References required are [System](#), [System.Data](#) and [MySql.Data](#).

### 21.2.4.1.1. The MySqlConnection Object

For your Connector/Net application to connect to a MySQL database, it must establish a connection by using a [MySqlConnection](#) object.

The [MySqlConnection](#) constructor takes a connection string as one of its parameters. The connection string provides necessary information to make the connection to the MySQL database. The connection string is discussed more fully in [Section 21.2.5.1, "Connecting to MySQL Using Connector/Net"](#). For a list of supported connection string options, see [Section 21.2.6, "Connector/Net Connection String Options Reference"](#).

The following code shows how to create a connection object:

```
using System;
using System.Data;

using MySql.Data;
using MySql.Data.MySqlClient;

public class Tutorial1
{
    public static void Main()
    {
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
        MySqlConnection conn = new MySqlConnection(connStr);
        try
        {
            Console.WriteLine("Connecting to MySQL...");
            conn.Open();
            // Perform database operations
        }
        catch (Exception ex)
        {
            Console.WriteLine(ex.ToString());
        }
        conn.Close();
        Console.WriteLine("Done.");
    }
}
```

When the [MySqlConnection](#) constructor is invoked, it returns a connection object, which is used for subsequent database operations. Open the connection before any other operations take place. Before the application exits, close the connection to the database by calling [Close](#) on the connection object.

Sometimes an attempt to perform an [Open](#) on a connection object can fail, generating an exception that can be handled using standard exception handling code.

In this section you have learned how to create a connection to a MySQL database, and open and close the corresponding connection object.

### 21.2.4.1.2. The MySqlCommand Object

Once a connection has been established with the MySQL database, the next step is do carry out the desired database operations. This can be achieved through the use of the [MySqlCommand](#) object.

You will see how to create a [MySqlCommand](#) object. Once it has been created, there are three main methods of interest that you can call:

- **ExecuteReader** - used to query the database. Results are usually returned in a [MySqlDataReader](#) object, created by [ExecuteReader](#).
- **ExecuteNonQuery** - used to insert and delete data.
- **ExecuteScalar** - used to return a single value.

Once a [MySqlCommand](#) object has been created, you will call one of the above methods on it to carry out a database operation, such as perform a query. The results are usually returned into a

[MySqlDataReader](#) object, and then processed, for example the results might be displayed. The following code demonstrates how this could be done.

```
using System;
using System.Data;

using MySql.Data;
using MySql.Data.MySqlClient;

public class Tutorial2
{
    public static void Main()
    {
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
        MySqlConnection conn = new MySqlConnection(connStr);
        try
        {
            Console.WriteLine("Connecting to MySQL...");
            conn.Open();

            string sql = "SELECT Name, HeadOfState FROM Country WHERE Continent='Oceania'";
            MySqlCommand cmd = new MySqlCommand(sql, conn);
            MySqlDataReader rdr = cmd.ExecuteReader();

            while (rdr.Read())
            {
                Console.WriteLine(rdr[0]+" -- "+rdr[1]);
            }
            rdr.Close();
        }
        catch (Exception ex)
        {
            Console.WriteLine(ex.ToString());
        }

        conn.Close();
        Console.WriteLine("Done.");
    }
}
```

When a connection has been created and opened, the code then creates a [MySqlCommand](#) object. Note that the SQL query to be executed is passed to the [MySqlCommand](#) constructor. The [ExecuteReader](#) method is then used to generate a [MySqlReader](#) object. The [MySqlReader](#) object contains the results generated by the SQL executed on the command object. Once the results have been obtained in a [MySqlReader](#) object, the results can be processed. In this case, the information is printed out by a `while` loop. Finally, the [MySqlReader](#) object is disposed of by running its `Close` method on it.

In the next example, you will see how to use the [ExecuteNonQuery](#) method.

The procedure for performing an [ExecuteNonQuery](#) method call is simpler, as there is no need to create an object to store results. This is because [ExecuteNonQuery](#) is only used for inserting, updating and deleting data. The following example illustrates a simple update to the [Country](#) table:

```
using System;
using System.Data;

using MySql.Data;
using MySql.Data.MySqlClient;

public class Tutorial3
{
    public static void Main()
    {
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
        MySqlConnection conn = new MySqlConnection(connStr);
        try
        {
            Console.WriteLine("Connecting to MySQL...");
            conn.Open();
```

```
        string sql = "INSERT INTO Country (Name, HeadOfState, Continent) VALUES ('Disneyland', 'Mick  
        MySqlCommand cmd = new MySqlCommand(sql, conn);  
        cmd.ExecuteNonQuery();  
    }  
    catch (Exception ex)  
    {  
        Console.WriteLine(ex.ToString());  
    }  
  
    conn.Close();  
    Console.WriteLine("Done.");  
}  
}
```

The query is constructed, the command object created and the [ExecuteNonQuery](#) method called on the command object. You can access your MySQL database with the `mysql` command interpreter and verify that the update was carried out correctly.

Finally, you will see how the [ExecuteScalar](#) method can be used to return a single value. Again, this is straightforward, as a [MySqlCommandReader](#) object is not required to store results, a simple variable will do. The following code illustrates how to use [ExecuteScalar](#):

```
using System;  
using System.Data;  
  
using MySql.Data;  
using MySql.Data.MySqlClient;  
  
public class Tutorial4  
{  
    public static void Main()  
    {  
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";  
        MySqlConnection conn = new MySqlConnection(connStr);  
        try  
        {  
            Console.WriteLine("Connecting to MySQL...");  
            conn.Open();  
  
            string sql = "SELECT COUNT(*) FROM Country";  
            MySqlCommand cmd = new MySqlCommand(sql, conn);  
            object result = cmd.ExecuteScalar();  
            if (result != null)  
            {  
                int r = Convert.ToInt32(result);  
                Console.WriteLine("Number of countries in the World database is: " + r);  
            }  
        }  
        catch (Exception ex)  
        {  
            Console.WriteLine(ex.ToString());  
        }  
  
        conn.Close();  
        Console.WriteLine("Done.");  
    }  
}
```

This example uses a simple query to count the rows in the [Country](#) table. The result is obtained by calling [ExecuteScalar](#) on the command object.

### 21.2.4.1.3. Working with Decoupled Data

Previously, when using [MySqlCommandReader](#), the connection to the database was continually maintained, unless explicitly closed. It is also possible to work in a manner where a connection is only established when needed. For example, in this mode, a connection could be established to read a chunk of data, the data could then be modified by the application as required. A connection could then

be reestablished only if and when the application writes data back to the database. This decouples the working data set from the database.

This decoupled mode of working with data is supported by Connector/Net. There are several parts involved in allowing this method to work:

- **Data Set** - The Data Set is the area in which data is loaded to read or modify it. A `DataSet` object is instantiated, which can store multiple tables of data.
- **Data Adapter** - The Data Adapter is the interface between the Data Set and the database itself. The Data Adapter is responsible for efficiently managing connections to the database, opening and closing them as required. The Data Adapter is created by instantiating an object of the `MySqlDataAdapter` class. The `MySqlDataAdapter` object has two main methods: `Fill` which reads data into the Data Set, and `Update`, which writes data from the Data Set to the database.
- **Command Builder** - The Command Builder is a support object. The Command Builder works in conjunction with the Data Adapter. When a `MySqlDataAdapter` object is created, it is typically given an initial `SELECT` statement. From this `SELECT` statement the Command Builder can work out the corresponding `INSERT`, `UPDATE` and `DELETE` statements that would be required to update the database. To create the Command Builder, an object of the class `MySqlCommandBuilder` is created.

Each of these classes will now be discussed in more detail.

### Instantiating a DataSet object

A `DataSet` object can be created simply, as shown in the following example code snippet:

```
DataSet dsCountry;  
...  
dsCountry = new DataSet();
```

Although this creates the `DataSet` object, it has not yet filled it with data. For that, a Data Adapter is required.

### Instantiating a MySqlDataAdapter object

The `MySqlDataAdapter` can be created as illustrated by the following example:

```
MySqlDataAdapter daCountry;  
...  
string sql = "SELECT Code, Name, HeadOfState FROM Country WHERE Continent='North America'";  
daCountry = new MySqlDataAdapter (sql, conn);
```

Note, the `MySqlDataAdapter` is given the SQL specifying the data to work with.

### Instantiating a MySqlCommandBuilder object

Once the `MySqlDataAdapter` has been created, it is necessary to generate the additional statements required for inserting, updating and deleting data. There are several ways to do this, but in this tutorial you will see how this can most easily be done with `MySqlCommandBuilder`. The following code snippet illustrates how this is done:

```
MySqlCommandBuilder cb = new MySqlCommandBuilder(daCountry);
```

Note that the `MySqlDataAdapter` object is passed as a parameter to the command builder.

### Filling the Data Set

To do anything useful with the data from your database, you need to load it into a Data Set. This is one of the jobs of the `MySqlDataAdapter` object, and is carried out with its `Fill` method. The following example code illustrates this:

```
DataSet dsCountry;  
...  
dsCountry = new DataSet();  
...
```



```
daCountry.Fill(dsCountry, "Country");
```

Note the `Fill` method is a `MySqlDataAdapter` method, the Data Adapter knows how to establish a connection with the database and retrieve the required data, and then populates the Data Set when the `Fill` method is called. The second parameter "Country" is the table in the Data Set to update.

### Updating the Data Set

The data in the Data Set can now be manipulated by the application as required. At some point, changes to data will need to be written back to the database. This is achieved through a `MySqlDataAdapter` method, the `Update` method.

```
daCountry.Update(dsCountry, "Country");
```

Again, the Data Set and the table within the Data Set to update are specified.

### Working Example

The interactions between the `DataSet`, `MySqlDataAdapter` and `MySqlCommandBuilder` classes can be a little confusing, so their operation can perhaps be best illustrated by working code.

In this example, data from the World database is read into a Data Grid View control. Here, the data can be viewed and changed before clicking an update button. The update button then activates code to write changes back to the database. The code uses the principles explained above. The application was built using the Microsoft Visual Studio to place and create the user interface controls, but the main code that uses the key classes described above is shown below, and is portable.

```
using System;
using System.Collections.Generic;
using System.ComponentModel;
using System.Data;
using System.Drawing;
using System.Linq;
using System.Text;
using System.Windows.Forms;

using MySql.Data;
using MySql.Data.MySqlClient;

namespace WindowsFormsApplication5
{
    public partial class Form1 : Form
    {
        MySqlDataAdapter daCountry;
        DataSet dsCountry;

        public Form1()
        {
            InitializeComponent();
        }

        private void Form1_Load(object sender, EventArgs e)
        {
            string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);
            try
            {
                label2.Text = "Connecting to MySQL...";

                string sql = "SELECT Code, Name, HeadOfState FROM Country WHERE Continent='North America'";
                daCountry = new MySqlDataAdapter(sql, conn);
                MySqlCommandBuilder cb = new MySqlCommandBuilder(daCountry);

                dsCountry = new DataSet();
                daCountry.Fill(dsCountry, "Country");
                dataGridView1.DataSource = dsCountry;
                dataGridView1.DataMember = "Country";
            }
        }
    }
}
```

```

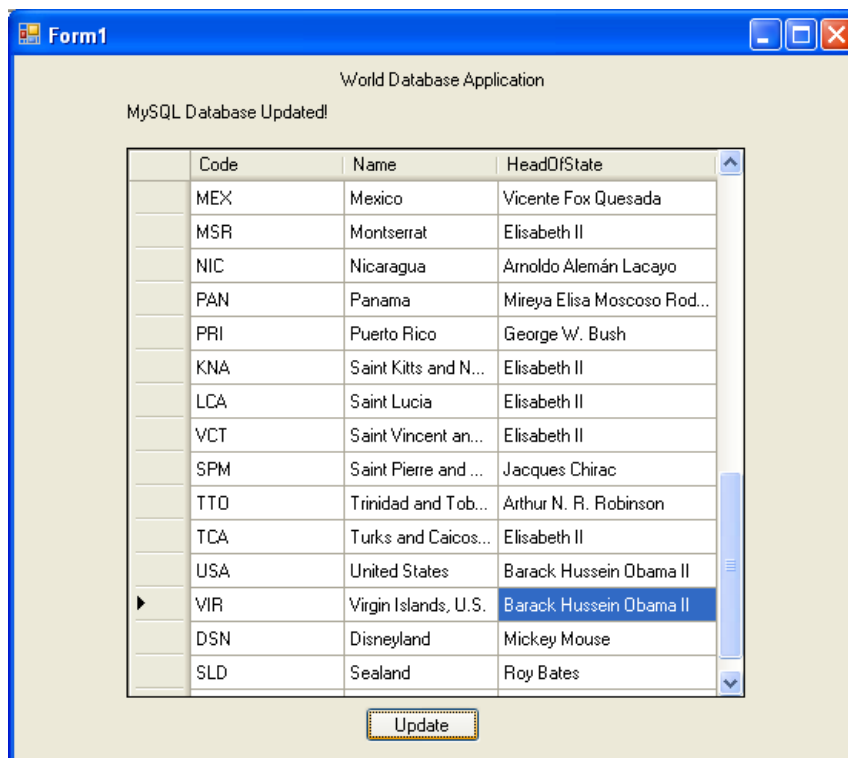
        catch (Exception ex)
        {
            label2.Text = ex.ToString();
        }
    }

    private void button1_Click(object sender, EventArgs e)
    {
        daCountry.Update(dsCountry, "Country");
        label2.Text = "MySQL Database Updated!";
    }
}

```

The application running is shown below:

**Figure 21.41. World Database Application**



#### 21.2.4.1.4. Working with Parameters

This part of the tutorial shows you how to use parameters in your Connector/Net application.

Although it is possible to build SQL query strings directly from user input, this is not advisable as it does not prevent erroneous or malicious information being entered. It is safer to use parameters as they will be processed as field data only. For example, imagine the following query was constructed from user input:

```
string sql = "SELECT Name, HeadOfState FROM Country WHERE Continent = "+user_continent;
```

If the string `user_continent` came from a Text Box control, there would potentially be no control over the string entered by the user. The user could enter a string that generates a runtime error, or in the worst case actually harms the system. When using parameters it is not possible to do this because a parameter is only ever treated as a field parameter, rather than an arbitrary piece of SQL code.

The same query written using a parameter for user input would be:

```
string sql = "SELECT Name, HeadOfState FROM Country WHERE Continent = @Continent";
```

Note that the parameter is preceded by an '@' symbol to indicate it is to be treated as a parameter.

As well as marking the position of the parameter in the query string, it is necessary to add a parameter to the Command object. This is illustrated by the following code snippet:

```
cmd.Parameters.AddWithValue("@Continent", "North America");
```

In this example the string "North America" is supplied as the parameter value statically, but in a more practical example it would come from a user input control.

A further example illustrates the complete process:

```
using System;
using System.Data;

using MySql.Data;
using MySql.Data.MySqlClient;

public class Tutorial5
{
    public static void Main()
    {
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
        MySqlConnection conn = new MySqlConnection(connStr);
        try
        {
            Console.WriteLine("Connecting to MySQL...");
            conn.Open();

            string sql = "SELECT Name, HeadOfState FROM Country WHERE Continent=@Continent";
            MySqlCommand cmd = new MySqlCommand(sql, conn);

            Console.WriteLine("Enter a continent e.g. 'North America', 'Europe': ");
            string user_input = Console.ReadLine();

            cmd.Parameters.AddWithValue("@Continent", user_input);

            MySqlDataReader rdr = cmd.ExecuteReader();

            while (rdr.Read())
            {
                Console.WriteLine(rdr["Name"]+" --- "+rdr["HeadOfState"]);
            }
            rdr.Close();
        }
        catch (Exception ex)
        {
            Console.WriteLine(ex.ToString());
        }

        conn.Close();
        Console.WriteLine("Done.");
    }
}
```

In this part of the tutorial you have see how to use parameters to make your code more secure.

#### 21.2.4.1.5. Working with Stored Procedures

This section illustrates how to work with stored procedures. Putting database-intensive operations into stored procedures lets you define an API for your database application. You can reuse this API across multiple applications and multiple programming languages. This technique avoids duplicating database code, saving time and effort when you make updates due to schema changes, tune the performance of queries, or add new database operations for logging, security, and so on. Before working through this tutorial, familiarize yourself with the [CREATE PROCEDURE](#) and [CREATE FUNCTION](#) statements that create different kinds of stored routines.

For the purposes of this tutorial, you will create a simple stored procedure to see how it can be called from Connector/Net. In the MySQL Client program, connect to the World database and enter the following stored procedure:

```
DELIMITER //
CREATE PROCEDURE country_hos
(IN con CHAR(20))
BEGIN
    SELECT Name, HeadOfState FROM Country
    WHERE Continent = con;
END //
DELIMITER ;
```

Test that the stored procedure works as expected by typing the following into the `mysql` command interpreter:

```
CALL country_hos('Europe');
```

Note that The stored routine takes a single parameter, which is the continent to restrict your search to.

Having confirmed that the stored procedure is present and correct, you can see how to access it from Connector/Net.

Calling a stored procedure from your Connector/Net application is similar to techniques you have seen earlier in this tutorial. A `MySqlCommand` object is created, but rather than taking an SQL query as a parameter, it takes the name of the stored procedure to call. Set the `MySqlCommand` object to the type of stored procedure, as shown by the following code snippet:

```
string rtn = "country_hos";
MySqlCommand cmd = new MySqlCommand(rtn, conn);
cmd.CommandType = CommandType.StoredProcedure;
```

In this case, the stored procedure requires you to pass a parameter. This can be achieved using the techniques seen in the previous section on parameters, [Section 21.2.4.1.4, "Working with Parameters"](#), as shown in the following code snippet:

```
cmd.Parameters.AddWithValue("@con", "Europe");
```

The value of the parameter `@con` could more realistically have come from a user input control, but for simplicity it is set as a static string in this example.

At this point, everything is set up and you can call the routine using techniques also learned in earlier sections. In this case, the `ExecuteReader` method of the `MySqlCommand` object is used.

Complete working code for the stored procedure example is shown below:

```
using System;
using System.Data;

using MySql.Data;
using MySql.Data.MySqlClient;

public class Tutorial6
{
    public static void Main()
    {
        string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
        MySqlConnection conn = new MySqlConnection(connStr);
        try
        {
            Console.WriteLine("Connecting to MySQL...");
            conn.Open();

            string rtn = "country_hos";
            MySqlCommand cmd = new MySqlCommand(rtn, conn);
            cmd.CommandType = CommandType.StoredProcedure;

            cmd.Parameters.AddWithValue("@con", "Europe");

            MySqlDataReader rdr = cmd.ExecuteReader();
            while (rdr.Read())
            {
                Console.WriteLine(rdr[0] + " --- " + rdr[1]);
            }
        }
    }
}
```

```

        rdr.Close();
    }
    catch (Exception ex)
    {
        Console.WriteLine(ex.ToString());
    }

    conn.Close();
    Console.WriteLine("Done.");
}
}

```

In this section, you have seen how to call a stored procedure from Connector/Net. For the moment, this concludes our introductory tutorial on programming with Connector/Net.

#### 21.2.4.2. Tutorial: MySQL Connector/Net ASP.NET Membership and Role Provider

Many web sites feature the facility for the user to create a user account. They can then log into the web site and enjoy a personalized experience. This requires that the developer creates database tables to store user information, along with code to gather and process this data. This represents a burden on the developer, and there is the possibility for security issues to creep into the developed code. However, ASP.NET 2.0 introduced the Membership system. This system is designed around the concept of Membership, Profile and Role Providers, which together provide all of the functionality to implement a user system, that previously would have to have been created by the developer from scratch.

Currently, MySQL Connector/Net provides Membership, Role, Profile and Session State Providers.

This tutorial shows you how to set up your ASP.NET web application to use the MySQL Connector/Net Membership and Role Providers. It assumes that you have MySQL Server installed, along with MySQL Connector/Net and Microsoft Visual Studio. This tutorial was tested with MySQL Connector/Net 6.0.4 and Microsoft Visual Studio 2008 Professional Edition. It is recommended you use 6.0.4 or above for this tutorial.

1. Create a new database in the MySQL Server using the MySQL Command-Line Client program ([mysql](#)), or other suitable tool. It does not matter what name is used for the database, but record it. You specify it in the connection string constructed later in this tutorial. This database contains the tables, automatically created for you later, used to store data about users and roles.
2. Create a new ASP.NET Web Site in Visual Studio. If you are not sure how to do this, refer to [Section 21.2.4.6, "Tutorial: Databinding in ASP.NET using LINQ on Entities"](#), which demonstrates how to create a simple ASP.NET web site.
3. Add References to [MySQL.Data](#) and [MySQL.Web](#) to the web site project.
4. Locate the [machine.config](#) file on your system, which is the configuration file for the .NET Framework.
5. Search the [machine.config](#) file to find the membership provider [MySQLMembershipProvider](#).
6. Add the attribute [autogenerateschema="true"](#). The appropriate section should now resemble the following (note: for the sake of brevity some information has been excluded):

```

<membership>
  <providers>
    <add name="AspNetSqlMembershipProvider"
      type="System.Web.Security.SqlMembershipProvider"
      ...
      connectionStringName="LocalSqlServer"
      ... />
    <add name="MySQLMembershipProvider"
      autogenerateschema="true"
      type="MySQL.Web.Security.MySQLMembershipProvider, MySQL.Web, Version=6.0.4.0, Culture=neutral,
      connectionStringName="LocalMySQLServer"
      ... />
  </providers>

```

```
</membership>
```

Note that the name for the connection string to be used to connect to the server that contains the membership database is `LocalMySqlServer`.

The `autogenerateschema="true"` attribute will cause MySQL Connector/Net to silently create, or upgrade, the schema on the database server, to contain the required tables for storing membership information.

7. It is now necessary to create the connection string referenced in the previous step. Load the web site's `web.config` file into Visual Studio.
8. Locate the section marked `<connectionStrings>`. Add the following connection string information:

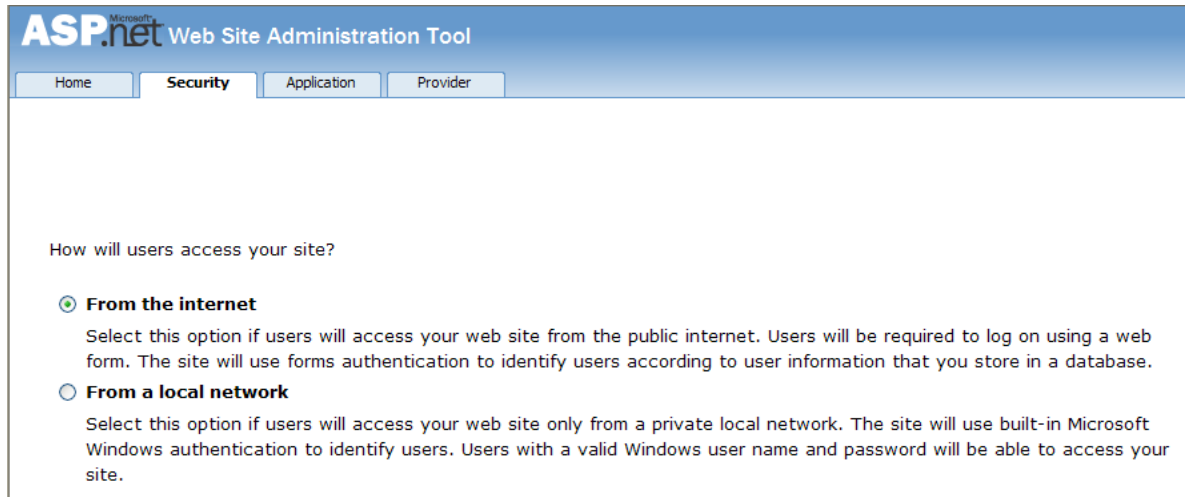
```
<connectionStrings>
  <remove name="LocalMySqlServer" />
  <add name="LocalMySqlServer"
        connectionString="Datasource=localhost;Database=users;uid=root;pwd=password;"
        providerName="MySql.Data.MySqlClient" />
</connectionStrings>
```

The database specified is the one created in the first step. You could alternatively have used an existing database.

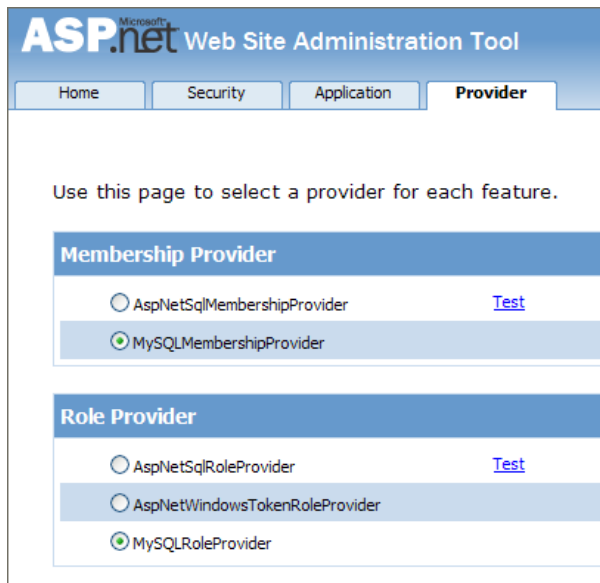
9. At this point build the solution to ensure no errors are present. This can be done by selecting **Build**, **Build Solution** from the main menu, or pressing **F6**.
10. ASP.NET supports the concept of locally and remotely authenticated users. With local authentication the user is validated using their Windows credentials when they attempt to access the web site. This can be useful in an Intranet environment. With remote authentication, a user is prompted for their login details when accessing the web site, and these credentials are checked against the membership information stored in a database server such as MySQL Server. You will now see how to choose this form of authentication.

Start the ASP.NET Web Site Administration Tool. This can be done quickly by clicking the small hammer/Earth icon in the Solution Explorer. You can also launch this tool by selecting **Website**, **ASP.NET Configuration** from the main menu.

11. In the ASP.NET Web Site Administration Tool click the **Security** tab.
12. Now click the **User Authentication Type** link.
13. Select the **From the internet** radio button. The web site will now need to provide a form to allow the user to enter their login details. These will be checked against membership information stored in the MySQL database.


**Figure 21.42. Authentication Type**

14. You now need to specify the Role and Membership Provider to be used. Click the **Provider** tab.
15. Click the **Select a different provider for each feature (advanced)** link.
16. Now select the **MySQLMembershipProvider** and the **MySQLRoleProvider** radio buttons.

**Figure 21.43. Select Membership and Role Provider**

17. In Visual Studio, rebuild the solution by selecting **Build, Rebuild Solution** from the main menu.
18. Check that the necessary schema has been created. This can be achieved using the `mysql` command interpreter.

Figure 21.44. Membership and Role Provider Tables



```

MySQL Command Line Client
+-----+
| Database |
+-----+
| information_schema |
| mysql |
| test |
| users |
| world |
+-----+
5 rows in set (0.01 sec)


mysql> show tables;
+-----+
| Tables_in_users |
+-----+
| my_aspnet_applications |
| my_aspnet_membership |
| my_aspnet_profiles |
| my_aspnet_roles |
| my_aspnet_schemaversion |
| my_aspnet_users |
| my_aspnet_usersinroles |
+-----+
7 rows in set (0.00 sec)

mysql>

```

19. Assuming all is present and correct, you can now create users and roles for your web application. The easiest way to do this is with the ASP.NET Web Site Administration Tool. However, many web applications contain their own modules for creating roles and users. For simplicity, the ASP.NET Web Site Administration Tool will be used in this tutorial.
20. In the ASP.NET Web Site Administration Tool, click the **Security** tab. Now that both the Membership and Role Provider are enabled, you will see links for creating roles and users. Click the **Create or Manage Roles** link.

Figure 21.45. Security Tab



**ASP.NET** Web Site Administration Tool

Home Security Application Provider

You can use the Web Site Administration Tool to manage all the security settings for your application. You can set up users and passwords (authentication), create roles (groups of users), and create permissions (rules for controlling access to parts of your application).

By default, user information is stored in a Microsoft SQL Server Express database in the Data folder of your Web site. If you want to store user information in a different database, use the Provider tab to select a different provider.

[Use the security Setup Wizard to configure security step by step.](#)

Click the links in the table to manage the settings for your application.

Users	Roles	Access Rules
Existing users: <b>1</b> <a href="#">Create user</a> <a href="#">Manage users</a> <a href="#">Select authentication type</a>	Existing roles: <b>2</b> <a href="#">Disable Roles</a> <a href="#">Create or Manage roles</a>	<a href="#">Create access rules</a> <a href="#">Manage access rules</a>

21. You can now enter the name of a new Role and click **Add Role** to create the new Role. Create new Roles as required.



22. Click the **Back** button.
23. Click the **Create User** link. You can now fill in information about the user to be created, and also allocate that user to one or more Roles.

**Figure 21.46. Create User**

ASP.NET Web Site Administration Tool

Home Security Application Provider

Add a user by entering the user's ID, password, and e-mail address on this page.

**Create User**

Sign Up for Your New Account

User Name: johndoe

Password: ●●●●●●●●

Confirm Password: ●●●●●●●●

E-mail: johndoe@example.com

Security Question: Statler and ...

Security Answer: Waldorf

Create User

Active User

**Roles**

Select roles for this user:

Visitors

Admins

24. Using the `mysql` command interpreter, you can check that your database has been correctly populated with the Membership and Role data.

**Figure 21.47. Membership and Roles Table Contents**

```

MySQL Command Line Client
mysql> show tables;
+-----+
| my_aspnet_roles |
| my_aspnet_schemaversion |
| my_aspnet_users |
| my_aspnet_usersinroles |
+-----+
7 rows in set (0.00 sec)

mysql> select * from my_aspnet_users;
+----+-----+-----+-----+-----+
| id | applicationId | name | isAnonymous | lastActivityDate |
+----+-----+-----+-----+-----+
| 1 | 1 | johndoe | 0 | 2009-07-14 17:17:25 |
+----+-----+-----+-----+-----+
1 row in set (0.01 sec)

mysql> select * from my_aspnet_roles;
+----+-----+-----+
| id | applicationId | name |
+----+-----+-----+
| 1 | 1 | Visitors |
| 2 | 1 | Admins |
+----+-----+-----+
2 rows in set (0.00 sec)

mysql>

```

In this tutorial, you have seen how to set up the MySQL Connector/Net Membership and Role Providers for use in your ASP.NET web application.

### 21.2.4.3. Tutorial: MySQL Connector/Net ASP.NET Session State Provider

MySQL Connector/Net from version 6.1 has included a MySQL Session State Provider. This provider enables you to store session state in a MySQL database. The following tutorial shows you how to prepare to use the MySQL Session State Provider, and then store session data into the MySQL database. This tutorial uses Microsoft Visual Studio 2008 Professional Edition, MySQL Connector/Net 6.1.1 and MySQL Server 5.1. This tutorial also assumes you have created an empty database, for example `test`, where you will store session data. You could do this using the `mysql` command interpreter.

1. In Visual Studio, create a new ASP.NET web site. If you are not sure how to do this, refer to the tutorial [Section 21.2.4.6, "Tutorial: Databinding in ASP.NET using LINQ on Entities"](#), which demonstrates how to do this.
2. Launch the MySQL Website Configuration tool. Due to a bug in 6.1.1, this may not appear unless you are connected to a server in the Server Explorer. If you are unfamiliar with the MySQL Website Configuration tool, consider first working through the tutorial in [Section 21.2.3.12, "MySQL Website Configuration Tool"](#).
3. Navigate through the wizard to the Session State page. Make sure the check box **Use MySQL to manage my ASP.NET session data** is selected.
4. On the same page, configure the connection string to the database that will contain your session data. If this database is empty, MySQL Connector/Net will create the schema required to store session data.
5. Ensure that the check box **Autogenerate Schema** is selected so that MySQL Connector/Net will create the schema in your database to store the session data correctly.
6. Enter the name of your application.
7. Click **Finish**. The MySQL Website Configuration tool will now update your application's `web.config` file with information about the connection string and default providers to be used. In this case, we have selected the MySQL Session State Provider.

At this point, you are ready to use the MySQL database to store session data. To test that the set up has worked, you can write a simple program that uses session variables.

1. Open `Default.aspx.cs`. In the `Page_Load` method, add the following code:

```
Session["SessionVariable1"] = "Test string";
```

2. Build your solution.
3. Run the solution (without debugging). When the application runs, the provider will autogenerate tables required in the database you chose when setting up the application.
4. Check that the schema was in fact created. Using the MySQL Command-Line Client use the target database and then type `SHOW TABLES;`. You will see that MySQL Connector/Net has created the required schema automatically, as we selected this to happen in the MySQL Website Configuration tool.
5. Now view the contents of these tables by typing `SELECT * FROM my_aspnet_sessions;` in the `mysql` command interpreter. This will display the session data our application used. Note that this is stored in binary format so some data may not display as expected.

At this point, you have installed the Session State Provider and carried out a preliminary test of the installation. You will now work a bit more with the Session State Provider.

In this part of the tutorial, you will set and retrieve a session variable. You can work with your existing project.

1. Select the `Default.aspx` and switch to Design View. Add a text box and three buttons. Change the text property for the buttons to “Store Session Variable”, “Clear Textbox”, and “Show Session Variable”. These will be `Button1`, `Button2` and `Button3` respectively. Build your solution to ensure that no errors have been introduced.
2. Still in the Design View, double-click `Button1`. Now to the `Button1_Click` event handler add code some the handler resembles the following:

```
protected void Button1_Click(object sender, EventArgs e)
{
    Session["SessionString"] = TextBox1.Text;
}
```

You have created a new Session variable accessed using the key “SessionString”. This will be set to the text that was entered into the text box when `Button1` is clicked.

3. In Design View, double-click `Button2` to add its click event handler. This button needs to clear text from the text box. The code to do this is as follows:

```
protected void Button2_Click(object sender, EventArgs e)
{
    TextBox1.Text = "";
}
```

The code simply assigns an empty string to the `Text` property of the text box.

4. In the Design View double-click `Button3` and modify the click handler as follows:

```
protected void Button3_Click(object sender, EventArgs e)
{
    TextBox1.Text = (String)Session["SessionString"];
}
```

This will retrieve the session string and display it in the text box.

5. Now modify the `Page_Load` method as follows:

```
protected void Page_Load(object sender, EventArgs e)
{
    if (!IsPostBack)
    {
        TextBox1.Text = "Enter some text";
    }
}
```

This ensures that when the page loads the text box `Text` property is reset.

6. Ensure that the solution is saved and then rebuild the solution.
7. Run the solution without debugging.
8. The form will be displayed. Enter some text into the text box. Now click `Store Session Variable`. At this point you have stored the string in a session variable.
9. Now click `Clear Text` to clear the text box.
10. Now click `Show Session Variable` to retrieve and display the session variable.
11. Refresh the page to destroy the form and display a new form.
12. Click `Show Session Variable` the text box will display the stored session variable, demonstrating that the refreshing the page does not destroy the session variable.

This illustrates that the session state data is not destroyed when a page is reloaded.

#### 21.2.4.4. Tutorial: MySQL Connector/Net ASP.NET Profile Provider

This tutorial shows you how to use the MySQL Profile Provider to store user profile information in a MySQL database. The tutorial uses MySQL Connector/Net 6.1.1, MySQL Server 5.1 and Microsoft Visual Studio 2008 Professional Edition.

Many modern web sites allow the user to create a personal profile. This requires a significant amount of code, but ASP.NET reduces this considerably by including the functionality in its Profile classes. The Profile Provider provides an abstraction between these classes and a data source. The MySQL Profile Provider enables profile data to be stored in a MySQL database. This enables the profile properties to be written to a persistent store, and be retrieved when required. The Profile Provider also enables profile data to be managed effectively, for example it enables profiles that have not been accessed since a specific date to be deleted.

The following steps show you how you can select the MySQL Profile Provider.

1. Create a new ASP.NET web project.
2. Select the MySQL Website Configuration tool. Due to a bug in 6.1.1 you may have to first connect to a server in Server Explorer before the tool's icon will display in the toolbar of the Solution Explorer.
3. In the MySQL Website Configuration tool navigate through the tool to the Profiles page.
4. Select the **Use MySQL to manage my profiles** check box.
5. Select the **Autogenerate Schema** check box.
6. Click the **Edit...** button and configure a connection string for the database that will be used to store user profile information.
7. Navigate to the last page of the tool and click **Finish** to save your changes and exit the tool.

At this point you are now ready to start using the MySQL Profile Provider. With the following steps you can carry out a preliminary test of your installation.

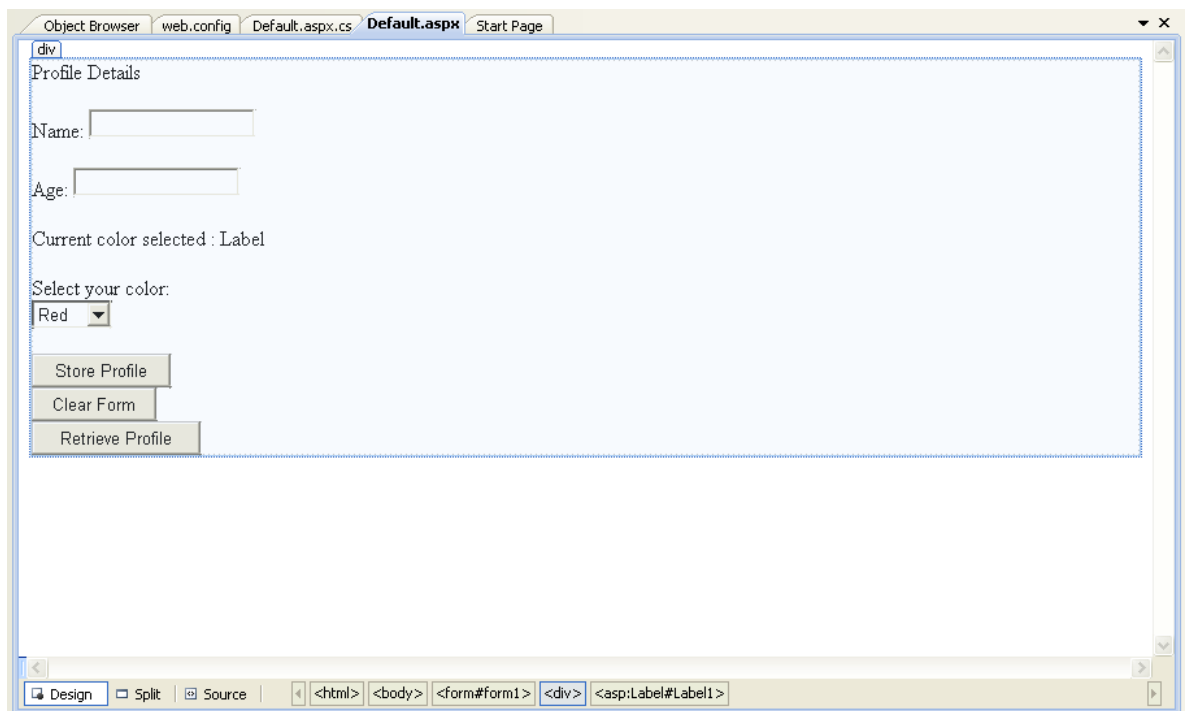
1. Open your `web.config` file.
2. Add a simple profile such as the following:

```
<system.web>
  <anonymousIdentification enabled="true" />
  <profile defaultProvider="MySQLProfileProvider">
    ...
    <properties>
      <add name="Name" allowAnonymous="true" />
      <add name="Age" allowAnonymous="true" type="System.UInt16" />
      <group name="UI">
        <add name="Color" allowAnonymous="true" defaultValue="Blue" />
        <add name="Style" allowAnonymous="true" defaultValue="Plain" />
      </group>
    </properties>
  </profile>
  ...
```

Note that `anonymousIdentification` has been set to true. This enables users who have not been authenticated to use profiles. They are identified by a GUID in a cookie rather than by user name.

Now that the simple profile has been defined in `web.config`, the next step is to write some code to test the profile.

1. In Design View design a simple page with the following controls:

**Figure 21.48. Simple Profile Application**

These will allow the user to enter some profile information. The user can also use the buttons to save their profile, clear the page, and restore their profile data.

2. In the Code View add code as follows:

```

...
protected void Page_Load(object sender, EventArgs e)
{
    if (!IsPostBack)
    {
        TextBox1.Text = Profile.Name;
        TextBox2.Text = Profile.Age.ToString();
        Label1.Text = Profile.UI.Color;
    }
}

// Store Profile
protected void Button1_Click(object sender, EventArgs e)
{
    Profile.Name = TextBox1.Text;
    Profile.Age = UInt16.Parse(TextBox2.Text);
}

// Clear Form
protected void Button2_Click(object sender, EventArgs e)
{
    TextBox1.Text = "";
    TextBox2.Text = "";
    Label1.Text = "";
}

// Retrieve Profile
protected void Button3_Click(object sender, EventArgs e)
{
    TextBox1.Text = Profile.Name;
    TextBox2.Text = Profile.Age.ToString();
    Label1.Text = Profile.UI.Color;
}

protected void DropDownList1_SelectedIndexChanged(object sender, EventArgs e)

```

```
{
    Profile.UI.Color = DropDownList1.SelectedValue;
}
...
```

3. Save all files and build the solution to check that no errors have been introduced.
4. Run the application.
5. Enter your name, age and select a color from the listbox. Now store this information in your profile by clicking **Store Profile**. Note that if you do not select a color from the listbox your profile will use the default color **Blue** that was specified in the `web.config` file.
6. Click **Clear Form** to clear text from the textboxes and the label that displays your chosen color.
7. Now click **Retrieve Profile** to restore your profile data from the MySQL database.
8. Now exit the browser to terminate the application.
9. Run the application again. Note that when the page loads your profile information is restored from the MySQL database.

In this tutorial you have seen how to using the MySQL Profile Provider with MySQL Connector/Net.

#### 21.2.4.5. Tutorial: Using an Entity Framework Entity as a Windows Forms Data Source

In this tutorial you will learn how to create a Windows Forms Data Source from an Entity in an Entity Data Model. This tutorial assumes that you have installed the World example database, which can be downloaded from the [MySQL Documentation page](#). You can also find details on how to install the database on the same page. It will also be convenient for you to create a connection to the World database after it is installed. For instructions on how to do this see [Section 21.2.3.1, "Making a Connection"](#).

##### Creating a new Windows Forms application

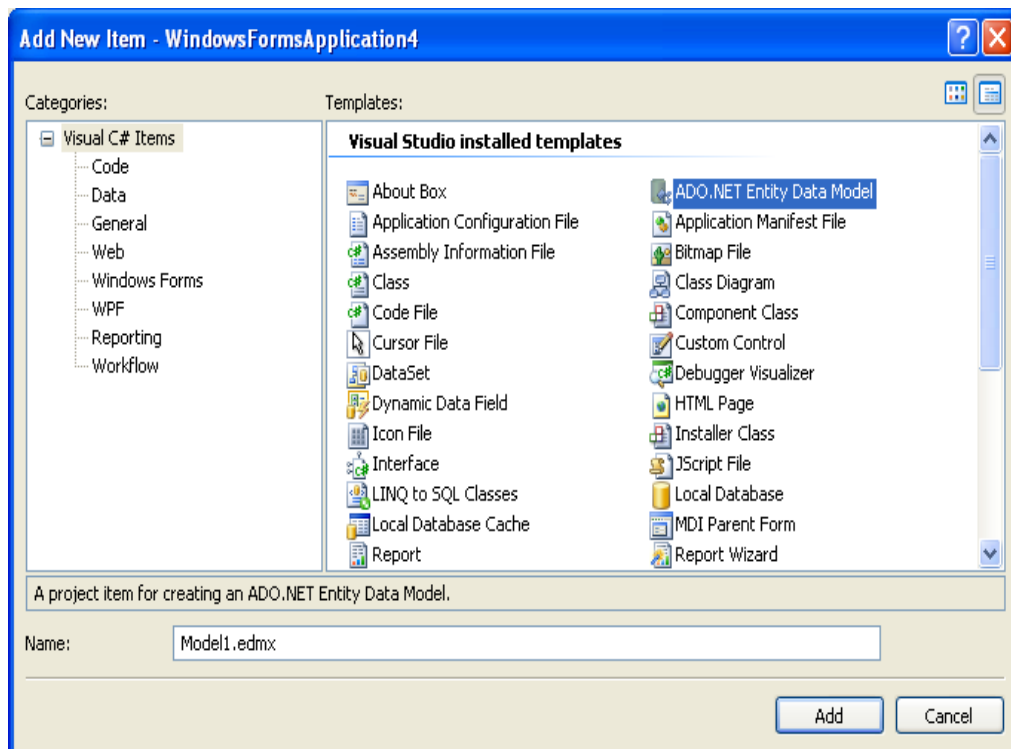
The first step is to create a new Windows Forms application.

1. In Visual Studio, select **File, New, Project** from the main menu.
2. Choose the **Windows Forms Application** installed template. Click **OK**. The solution is created.

##### Adding an Entity Data Model

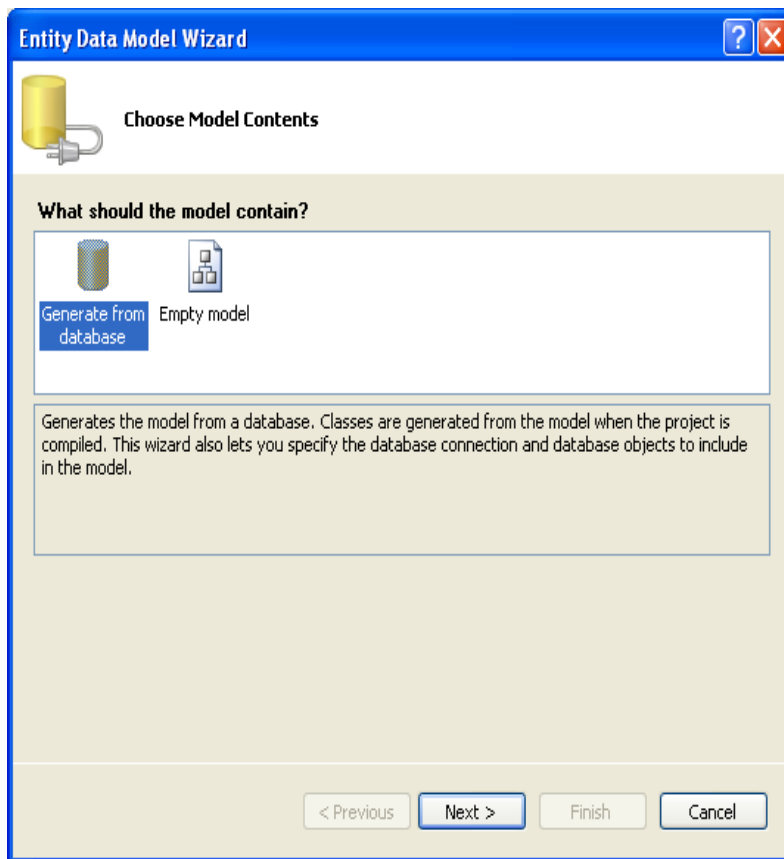
You will now add an Entity Data Model to your solution.

1. In the Solution Explorer, right-click your application and select **Add, New Item...**. From **Visual Studio installed templates** select **ADO.NET Entity Data Model**. Click **Add**.

**Figure 21.49. Add Entity Data Model**

2. You will now see the Entity Data Model Wizard. You will use the wizard to generate the Entity Data Model from the world example database. Select the icon **Generate from database**. Click **Next**.

Figure 21.50. Entity Data Model Wizard Screen 1



3. You can now select the connection you made earlier to the World database. If you have not already done so, you can create the new connection at this time by clicking **New Connection...**. For further instructions on creating a connection to a database see [Section 21.2.3.1, "Making a Connection"](#).



Figure 21.51. Entity Data Model Wizard Screen 2

**Entity Data Model Wizard**

**Choose Your Data Connection**

Which data connection should your application use to connect to the database?

localhost(world) [v] New Connection...

This connection string appears to contain sensitive data (for example, a password) that is required to connect to the database. Storing sensitive data in the connection string can be a security risk. Do you want to include this sensitive data in the connection string?

No, exclude sensitive data from the connection string. I will set it in my application code.

Yes, include the sensitive data in the connection string.

Entity connection string:

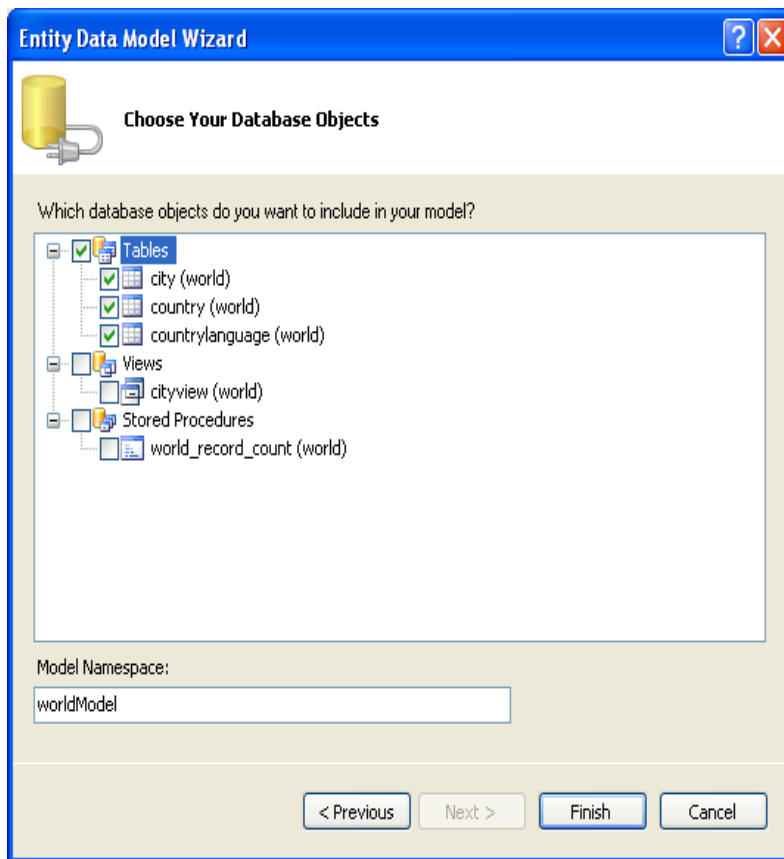
```
metadata=res://*/Model1.csdl|res://*/Model1.ssdl|res://*/Model1.msl;provider=MySql.Data.MySqlClient;provider connection string="server=localhost;user id=root;persist security info=True;database=world"
```

Save entity connection settings in App.Config as:

worldEntities

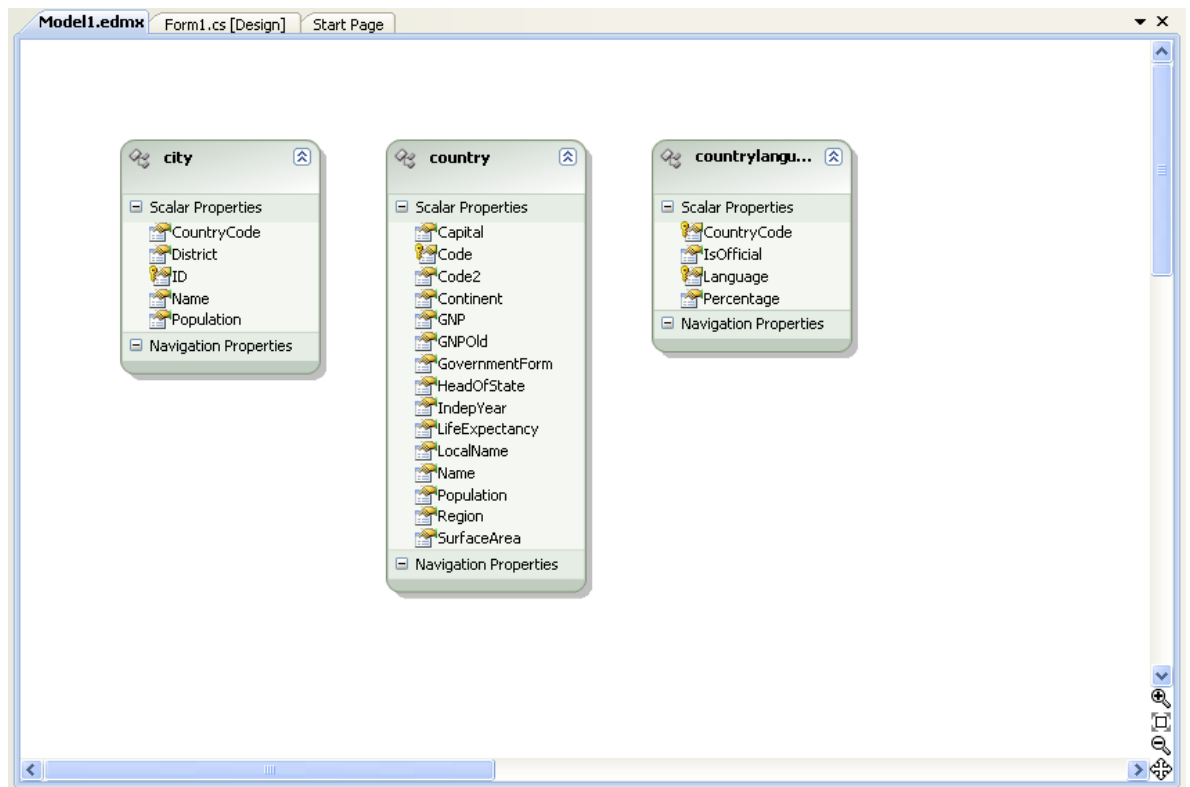
< Previous Next > Finish Cancel

4. Make a note of the entity connection settings to be used in App.Config, as these will be used later to write the necessary control code.
5. Click **Next**.
6. The Entity Data Model Wizard connects to the database. You are then presented with a tree structure of the database. From this you can select the object you would like to include in your model. If you had created Views and Stored Routines these will be displayed along with any tables. In this example you just need to select the tables. Click **Finish** to create the model and exit the wizard.

**Figure 21.52. Entity Data Model Wizard Screen 3**

7. Visual Studio will generate the model and then display it.

Figure 21.53. Entity Data Model Diagram



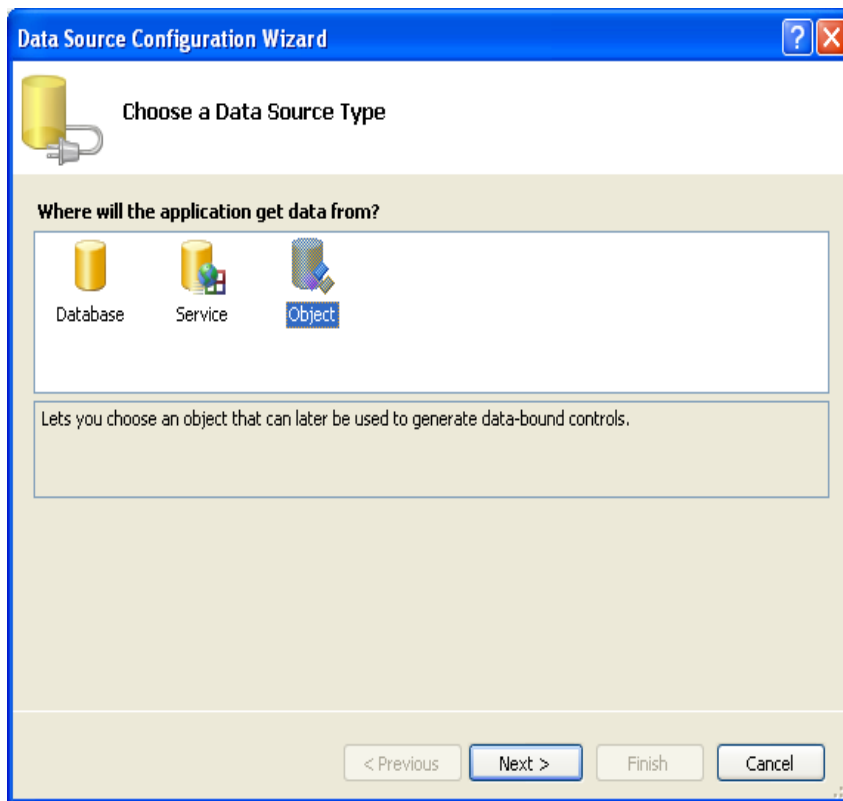
8. From the Visual Studio main menu select **Build, Build Solution**, to ensure that everything compiles correctly so far.

### Adding a new Data Source

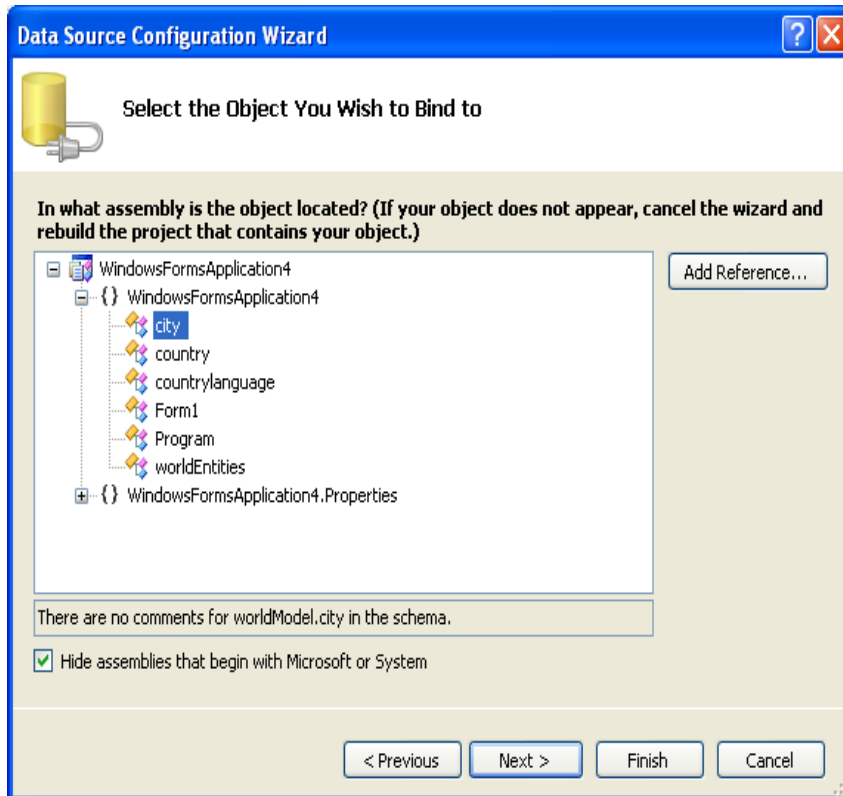
You will now add a new Data Source to your project and see how it can be used to read and write to the database.

1. From the Visual Studio main menu select **Data, Add New Data Source...**. You will be presented with the Data Source Configuration Wizard.

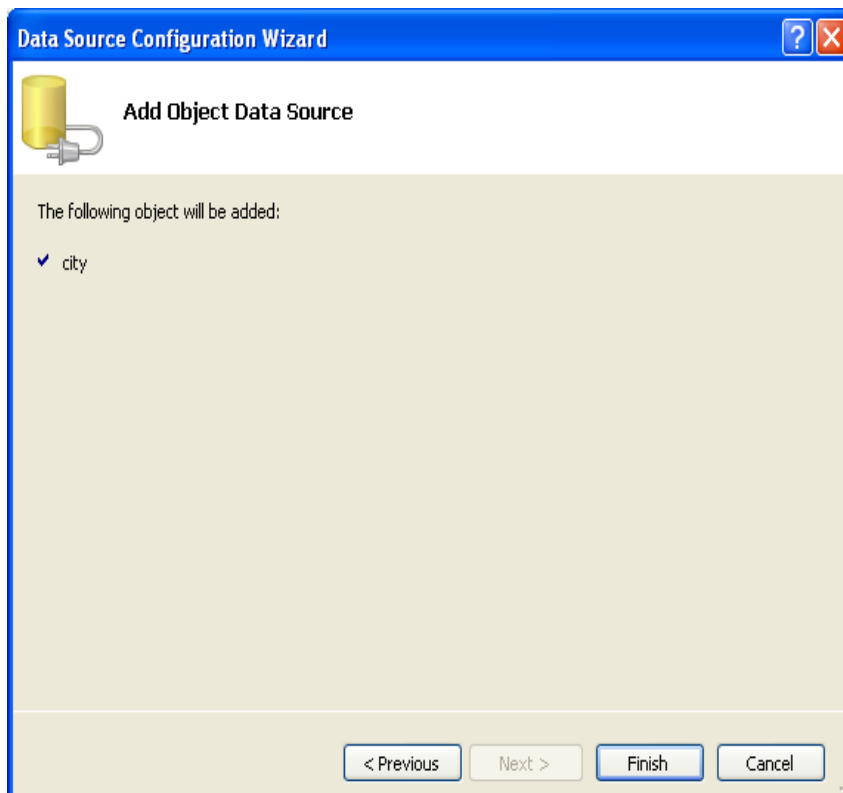
Figure 21.54. Entity Data Source Configuration Wizard Screen 1



2. Select the **Object** icon. Click **Next**.
3. You will now select the Object to bind to. Expand the tree. In this tutorial, you will select the city table. Once the city table has been selected click **Next**.

**Figure 21.55. Entity Data Source Configuration Wizard Screen 2**

4. The wizard will confirm that the city object is to be added. Click **Finish**.

**Figure 21.56. Entity Data Source Configuration Wizard Screen 3**

5. The city object will be display in the Data Sources panel. If the Data Sources panel is not displayed, select **Data**, **Show Data Sources** from the Visual Studio main menu. The docked panel will then be displayed.

**Figure 21.57. Data Sources**




### Using the Data Source in a Windows Form

You will now learn how to use the Data Source in a Windows Form.

1. In the Data Sources panel select the Data Source you just created and drag and drop it onto the Form Designer. By default the Data Source object will be added as a Data Grid View control. Note that the Data Grid View control is bound to the `cityBindingSource` and the Navigator control is bound to `cityBindingNavigator`.

Figure 21.58. Data Form Designer

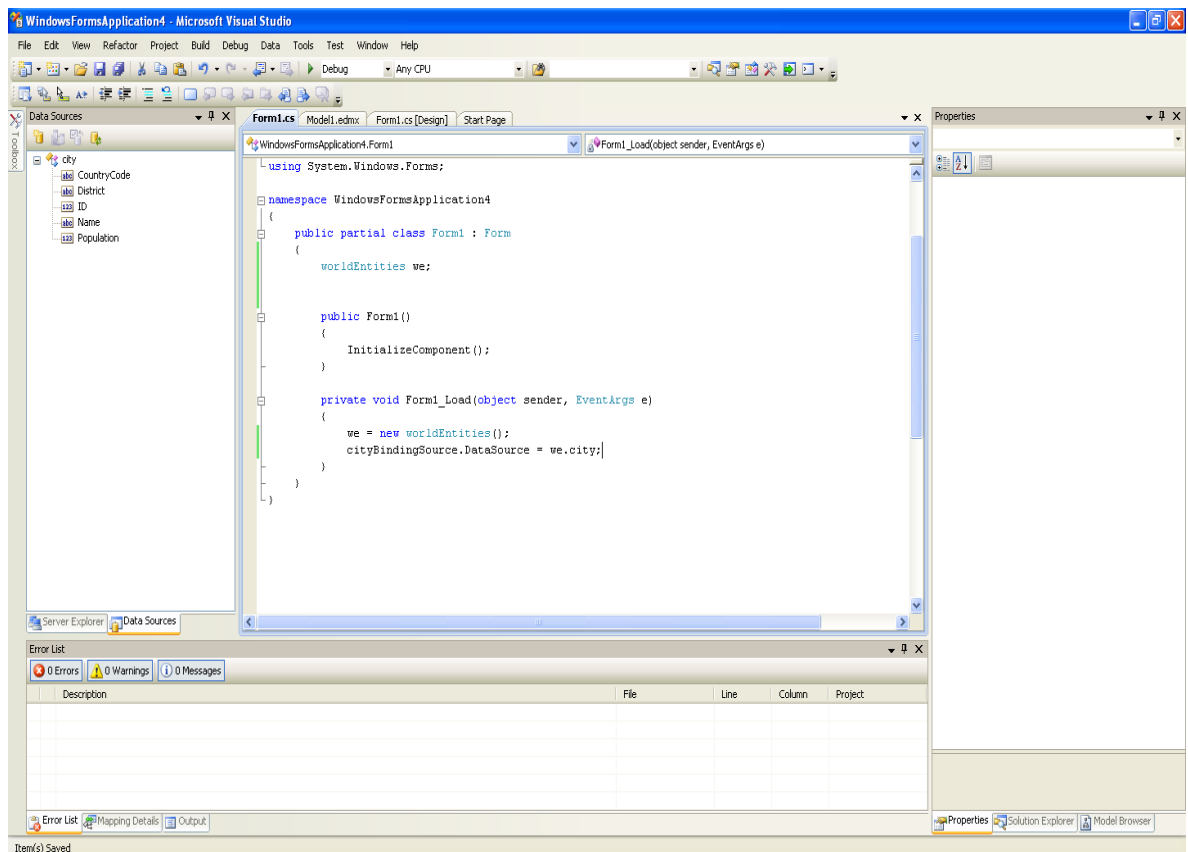


2. Save and rebuild the solution before continuing.

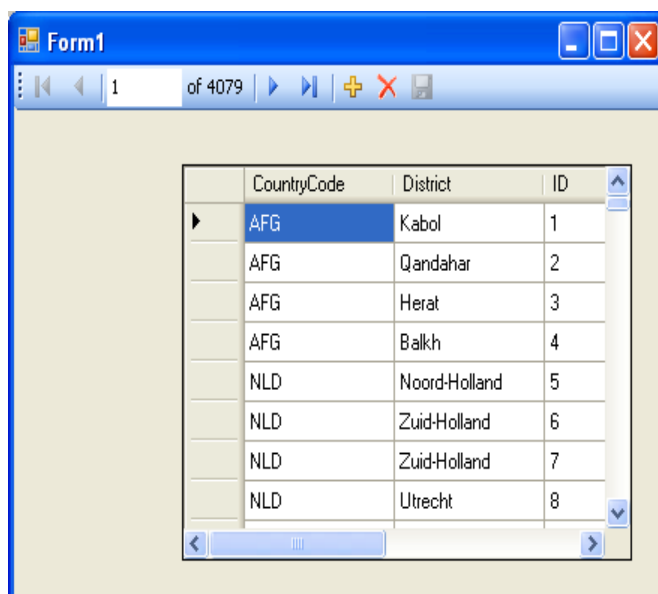
### Adding Code to Populate the Data Grid View

You are now ready to add code to ensure that the Data Grid View control will be populated with data from the City database table.

1. Double-click the form to access its code.
2. Add code to instantiate the Entity Data Model's EntityContainer object and retrieve data from the database to populate the control.

**Figure 21.59. Adding Code to the Form**

3. Save and rebuild the solution.
4. Run the solution. Ensure the grid is populated and you can navigate the database.

**Figure 21.60. The Populated Grid Control**

### Adding Code to Save Changes to the Database


You will now add code to enable you to save changes to the database.



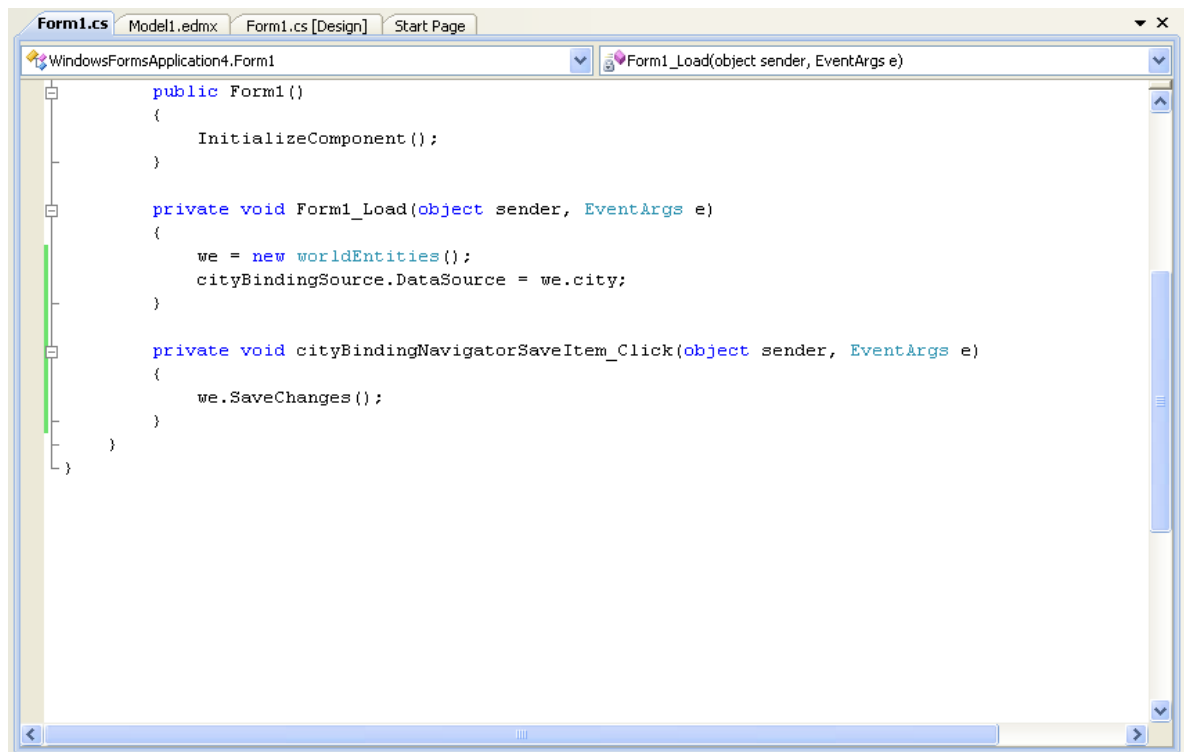
The Binding source component ensures that changes made in the Data Grid View control are also made to the Entity classes bound to it. However, that data needs to be saved back from the entities to the database itself. This can be achieved by the enabling of the Save button in the Navigator control, and the addition of some code.

1. In the Form Designer, click the Save icon in the Form toolbar and ensure that its Enabled property is set to True.

**Figure 21.61. Save Button Enabled**



2. Double-click the Save icon in the Form toolbar to display its code.
3. You now need to add code to ensure that data is saved to the database when the save button is clicked in the application.

**Figure 21.62. Adding Save Code to the Form**

4. Once the code has been added, save the solution and rebuild it. Run the application and verify that changes made in the grid are saved.

#### 21.2.4.6. Tutorial: Databinding in ASP.NET using LINQ on Entities

In this tutorial you create an ASP.NET web page that binds LINQ queries to entities using the Entity Framework mapping.

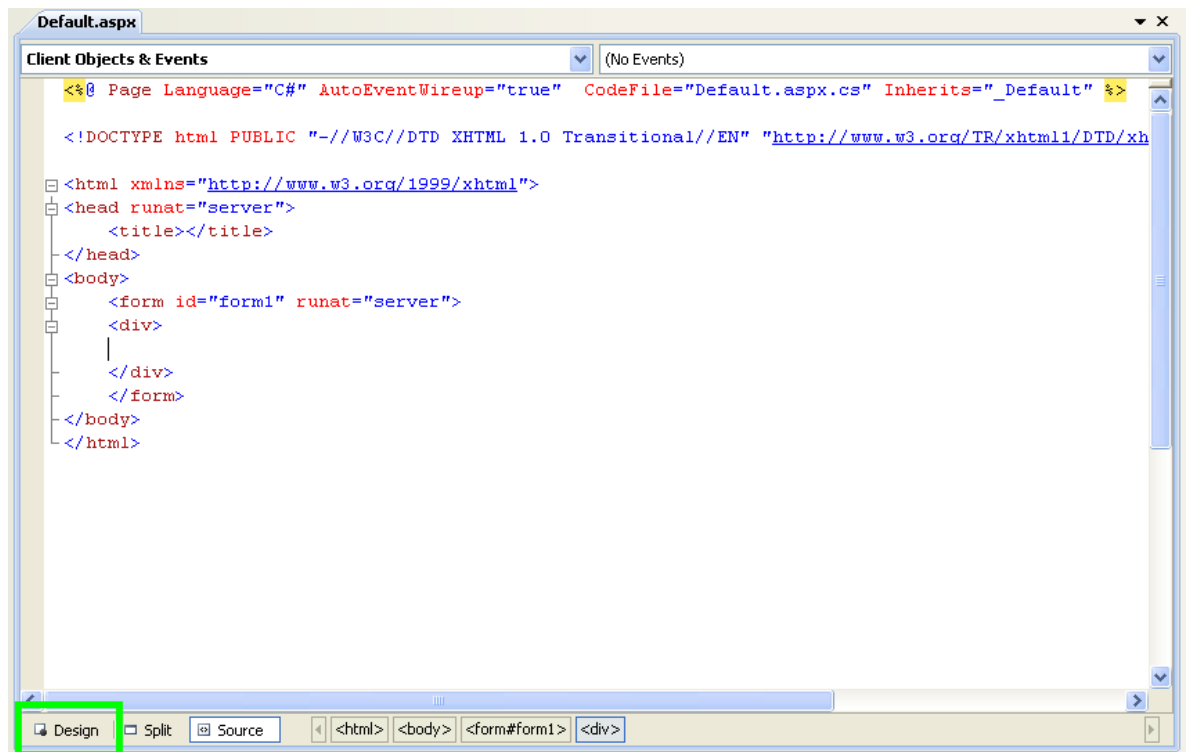
If you have not already done so, install the World example database prior to attempting this tutorial. See the tutorial [Section 21.2.4.5, "Tutorial: Using an Entity Framework Entity as a Windows Forms Data Source"](#) for instructions on downloading and installing this database.

##### Creating an ASP.NET web site

In this part of the tutorial, you create an ASP.NET web site. The web site uses the World database. The main web page features a drop down list from which you can select a country. Data about that country's cities is then displayed in a grid view control.


1. From the Visual Studio main menu select **File**, **New**, **Web Site...**.
2. From the Visual Studio installed templates select **ASP.NET Web Site**. Click **OK**. You will be presented with the Source view of your web page by default.
3. Click the Design view tab situated underneath the Source view panel.

Figure 21.63. The Design Tab




- In the Design view panel, enter some text to decorate the blank web page.
- Click Toolbox. From the list of controls select **DropDownList**. Drag and drop the control to a location beneath the text on your web page.

Figure 21.64. Drop Down List




- From the **DropDownList** control's context menu, ensure that the **Enable AutoPostBack** check box is enabled. This will ensure the control's event handler is called when an item is selected. The user's choice will in turn be used to populate the **GridView** control.

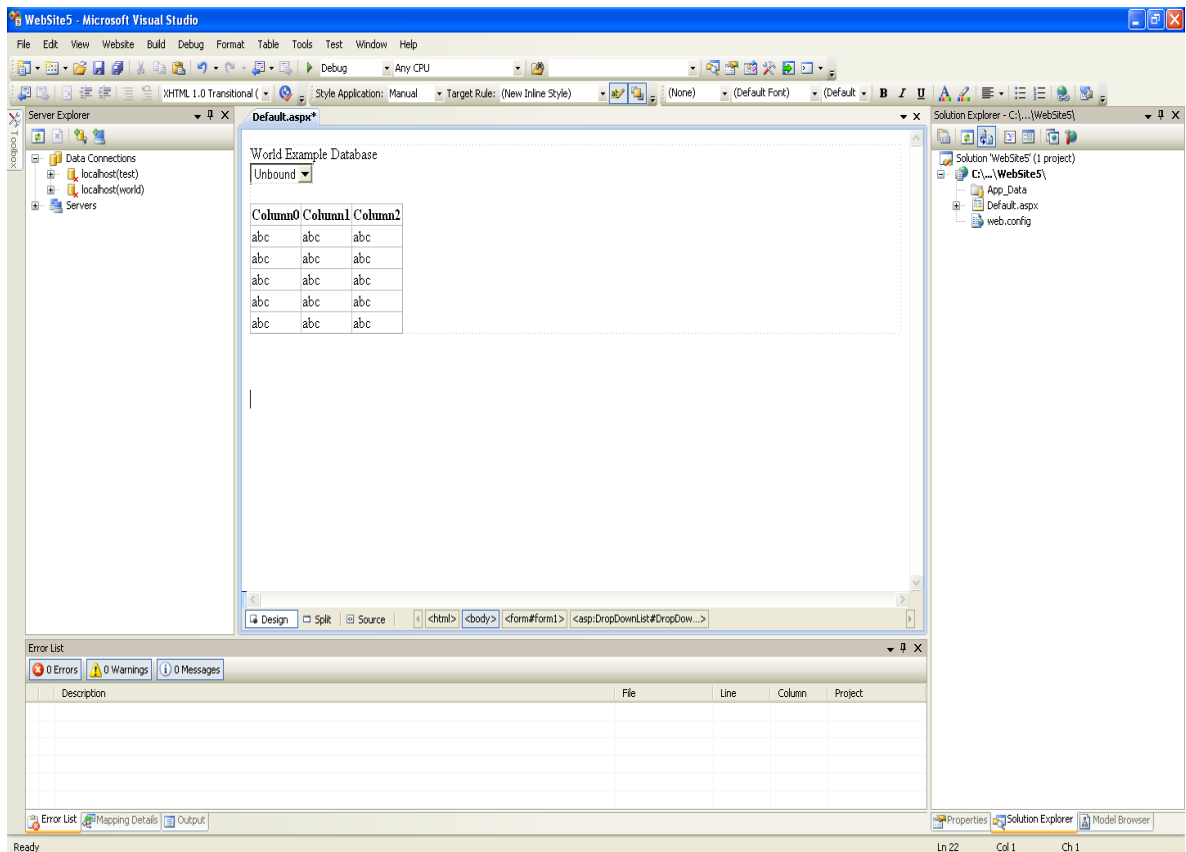
Figure 21.65. Enable AutoPostBack



- From the Toolbox select the **GridView** control.

**Figure 21.66. Grid View Control**

Drag and drop the Grid View control to a location just below the Drop Down List you already placed.

**Figure 21.67. Placed Grid View Control**

8. At this point it is recommended that you save your solution, and build the solution to ensure that there are no errors.
9. If you run the solution you will see that the text and drop down list are displayed, but the list is empty. Also, the grid view does not appear at all. Adding this functionality is described in the following sections.

At this stage you have a web site that will build, but further functionality is required. The next step will be to use the Entity Framework to create a mapping from the World database into entities that you can control programmatically.

### Creating an ADO.NET Entity Data Model

In this stage of the tutorial you will add an ADO.NET Entity Data Model to your project, using the World database at the storage level. The procedure for doing this is described in the tutorial [Section 21.2.4.5, "Tutorial: Using an Entity Framework Entity as a Windows Forms Data Source"](#), and so will not be repeated here.

## Populating a Drop Data List Box with using the results of a entity LINQ query

In this part of the tutorial you will write code to populate the DropDownList control. When the web page loads the data to populate the list will be achieved by using the results of a LINQ query on the model created previously.

1. In the Design view panel, double-click any blank area. This brings up the [Page\\_Load](#) method.
2. Modify the relevant section of code according to the following listing:

```
...
public partial class _Default : System.Web.UI.Page
{
    worldModel.worldEntities we;

    protected void Page_Load(object sender, EventArgs e)
    {
        we = new worldModel.worldEntities();

        if (!IsPostBack)
        {
            var countryQuery = from c in we.country
                               orderby c.Name
                               select new { c.Code, c.Name };
            DropDownList1.DataValueField = "Code";
            DropDownList1.DataTextField = "Name";
            DropDownList1.DataSource = countryQuery;
            DataBind();
        }
    }
}
...
```

Note that the list control only needs to be populated when the page first loads. The conditional code ensures that if the page is subsequently reloaded, the list control is not repopulated, which would cause the user selection to be lost.

3. Save the solution, build it and run it. You should see the list control has been populated. You can select an item, but as yet the grid view control does not appear.

At this point you have a working Drop Down List control, populated by a LINQ query on your entity data model.

## Populating a Grid View control using an entity LINQ query

In the last part of this tutorial you will populate the Grid View Control using a LINQ query on your entity data model.

1. In the Design view, double-click the **DropDownList** control. This causes its [SelectedIndexChanged](#) code to be displayed. This method is called when a user selects an item in the list control and thus fires an AutoPostBack event.
2. Modify the relevant section of code accordingly to the following listing:

```
...
protected void DropDownList1_SelectedIndexChanged(object sender, EventArgs e)
{
    var cityQuery = from c in we.city
                    where c.CountryCode == DropDownList1.SelectedValue
                    orderby c.Name
                    select new { c.Name, c.Population, c.CountryCode };
    GridView1.DataSource = cityQuery;
    DataBind();
}
...
```

The grid view control is populated from the result of the LINQ query on the entity data model.

3. As a check compare your code to that shown in the following screenshot:

Figure 21.68. Source Code

```

protected void Page_Load(object sender, EventArgs e)
{
    we = new worldModel.worldEntities();

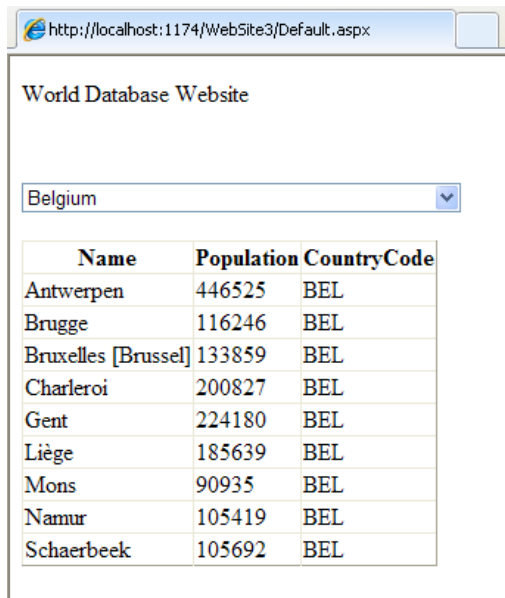
    if (!IsPostBack)
    {
        var countryQuery = from c in we.country
                           orderby c.Name
                           select new { c.Code, c.Name };
        DropDownList1.DataValueField = "Code";
        DropDownList1.DataTextField = "Name";
        DropDownList1.DataSource = countryQuery;
        DataBind();
    }
}

protected void DropDownList1_SelectedIndexChanged(object sender, EventArgs e)
{
    var cityQuery = from c in we.city
                    where c.CountryCode == DropDownList1.SelectedValue
                    orderby c.Name
                    select new { c.Name, c.Population, c.CountryCode };
    GridView1.DataSource = cityQuery;
    DataBind();
}

```

4. Save, build and run the solution. As you select a country you will see its cities are displayed in the grid view control.

Figure 21.69. The Working Web Site



In this tutorial you have seen how to create an ASP.NET web site, you have also seen how you can access a MySQL database using LINQ queries on an entity data model.

#### 21.2.4.7. Tutorial: Using SSL with MySQL Connector/Net

In this tutorial you will learn how you can use MySQL Connector/Net to connect to a MySQL server configured to use SSL. Support for SSL client certificates was added with MySQL Connector/Net 6.2.

MySQL Server uses the PEM format for certificates and private keys. This tutorial will use the test certificates from the server test suite by way of example. You can obtain the MySQL Server source code from [MySQL Downloads](#). The certificates can be found in the directory `./mysql-test/std_data`.

To carry out the steps in this tutorial, you must have Open SSL installed. This can be downloaded for Microsoft Windows at no charge from [Shining Light Productions](#).

Further details on the connection string options used in this tutorial can be found at [Section 21.2.6, "Connector/Net Connection String Options Reference"](#).

### Configuring the MySQL Server to use SSL

1. In the MySQL Server configuration file, set the SSL parameters as follows:

```
ssl-ca=path/to/repo/mysql-test/std_data/cacert.pem
ssl-cert=path/to/repo/mysql-test/std_data/server-cert.pem
ssl-key=path/to/repo/mysql-test/std_data/server-key.pem
```

Adjust the directories according to the location in which you installed the MySQL source code.

2. In this step you create a test user and set the user to require SSL.

Using the MySQL Command-Line Client, connect as `root` and create the user `sslclient`.

3. To set privileges and requirements, issue the following command:

```
GRANT ALL PRIVILEGES ON *.* TO sslclient@%' REQUIRE SSL;
```

### Creating a certificate file to use with the .NET client

1. The .NET client does not use the PEM file format, as .NET does not support this format natively. You will be using test client certificates from the same server repository, for the purposes of this example. Convert these to PFX format first. This format is also known as PKCS#12. An article describing this procedure can be found at the [Citrix website](#). From the directory `server-repository-root/mysql-test/std_data`, issue the following command:

```
openssl pkcs12 -export -in client-cert.pem -inkey client-key.pem -certfile cacert.pem -out client.pfx
```

2. When asked for an export password, enter the password "pass". The file `client.pfx` will be generated. This file is used in the remainder of the tutorial.

### Connecting to the server using a file-based certificate

1. You will use PFX file, `client.pfx` you created in the previous step to authenticate the client. The following example demonstrates how to connect using the `SSL Mode`, `CertificateFile` and `CertificatePassword` connection string options:

```
using (MySQLConnection connection = new MySQLConnection(
    "database=test;user=sslclient;" +
    "CertificateFile=H:\\bzt\\mysql-trunk\\mysql-test\\std_data\\client.pfx" +
    "CertificatePassword=pass;" +
    "SSL Mode=Required "))
{
    connection.Open();
}
```

The path to the certificate file will need to be changed to reflect your individual installation.

### Connecting to the server using a store-based certificate

1. The first step is to import the PFX file, `client.pfx`, into the Personal Store. Double-click the file in Windows explorer. This launches the Certificate Import Wizard.
2. Follow the steps dictated by the wizard, and when prompted for the password for the PFX file, enter "pass".

- Click **Finish** to close the wizard and import the certificate into the personal store.

### Examine certificates in the Personal Store

- Start the Microsoft Management Console by entering `mmc.exe` at a command prompt.
- Select **File**, **Add/Remove snap-in**. Click **Add**. Select **Certificates** from the list of available snap-ins in the dialog.
- Click **Add** button in the dialog, and select the **My user account** radio button. This is used for personal certificates.
- Click the **Finish** button.
- Click **OK** to close the Add/Remove Snap-in dialog.
- You will now have **Certificates – Current User** displayed in the left panel of the Microsoft Management Console. Expand the Certificates - Current User tree item and select **Personal, Certificates**. The right-hand panel will display a certificate issued to MySQL. This is the certificate that was previously imported. Double-click the certificate to display its details.
- After you have imported the certificate to the Personal Store, you can use a more succinct connection string to connect to the database, as illustrated by the following code:

```
using (MySQLConnection connection = new MySQLConnection(
    "database=test;user=sslclient;" +
    "Certificate Store Location=CurrentUser;" +
    "SSL Mode=Required"))
{
    connection.Open();
}
```

### Certificate Thumbprint Parameter

If you have a large number of certificates in your store, and many have the same Issuer, this can be a source of confusion and result in the wrong certificate being used. To alleviate this situation, there is an optional Certificate Thumbprint parameter that can additionally be specified as part of the connection string. As mentioned before, you can double-click a certificate in the Microsoft Management Console to display the certificate's details. When the Certificate dialog is displayed click the **Details** tab and scroll down to see the thumbprint. The thumbprint will typically be a number such as `#47 94 36 00 9a 40 f3 01 7a 14 5c f8 47 9e 76 94 d7 aa de f0`. This thumbprint can be used in the connection string, as the following code illustrates:

```
using (MySQLConnection connection = new MySQLConnection(
    "database=test;user=sslclient;" +
    "Certificate Store Location=CurrentUser;" +
    "Certificate Thumbprint=479436009a40f3017a145cf8479e7694d7aadeef0;" +
    "SSL Mode=Required"))
{
    connection.Open();
}
```

Spaces in the thumbprint parameter are optional and the value is case-insensitive.

### 21.2.4.8. Tutorial: Using MySqlCommand

This tutorial teaches you how to use the `MySqlCommand` class. This class enables you to execute a series of statements. Depending on the circumstances, this can be more convenient than using the `MySQLCommand` approach.

Further details of the `MySqlCommand` class can be found in the reference documentation supplied with MySQL Connector/Net.

To run the example programs in this tutorial, set up a simple test database and table using the `mysql` Command-Line Client or MySQL Workbench. Commands for the `mysql` Command-Line Client are given here:



```
CREATE DATABASE TestDB;
USE TestDB;
CREATE TABLE TestTable (id INT NOT NULL PRIMARY KEY
    AUTO_INCREMENT, name VARCHAR(100));
```

The main method of the `MySqlCommand` class is the `Execute` method. This method causes the script (sequence of statements) assigned to the `Query` property of the `MySqlCommand` object to be executed. Note the `Query` property can be set through the `MySqlCommand` constructor or using the `Query` property. `Execute` returns the number of statements executed.

The `MySqlCommand` object will execute the specified script on the connection set using the `Connection` property. Again, this property can be set directly or through the `MySqlCommand` constructor. The following code snippets illustrate this:

```
string sql = "SELECT * FROM TestTable";
...
MySqlCommand script = new MySqlCommand(conn, sql);
...
MySqlCommand script = new MySqlCommand();
script.Query = sql;
script.Connection = conn;
...
script.Execute();
```

The `MySqlCommand` class has several events associated with it. There are:

1. `Error` - generated if an error occurs.
2. `ScriptCompleted` - generated when the script successfully completes execution.
3. `StatementExecuted` - generated after each statement is executed.

It is possible to assign event handlers to each of these events. These user-provided routines are called back when the connected event occurs. The following code shows how the event handlers are set up.

```
script.Error += new MySqlCommandErrorHandler(script_Error);
script.ScriptCompleted += new EventHandler(script_ScriptCompleted);
script.StatementExecuted += new MySqlCommandStatementExecutedEventHandler(script_StatementExecuted);
```

In VisualStudio, you can save typing by using tab completion to fill out stub routines. Start by typing, for example, "`script.Error +=`". Then press **TAB**, and then press **TAB** again. The assignment is completed, and a stub event handler created. A complete working example is shown below:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

using System.Data;
using MySql.Data;
using MySql.Data.MySqlClient;

namespace MySqlCommandTest
{
    class Program
    {
        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=TestDB;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);

            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                string sql = "INSERT INTO TestTable(name) VALUES ('Superman');" +
                    "INSERT INTO TestTable(name) VALUES ('Batman');" +
                    "INSERT INTO TestTable(name) VALUES ('Wolverine');" +
                    "INSERT INTO TestTable(name) VALUES ('Storm');";
```

```

        MySqlScript script = new MySqlScript(conn, sql);

        script.Error += new MySqlScriptErrorHandler(script_Error);
        script.ScriptCompleted += new EventHandler(script_ScriptCompleted);
        script.StatementExecuted += new MySqlStatementExecutedEventHandler(script_StatementExecuted);

        int count = script.Execute();

        Console.WriteLine("Executed " + count + " statement(s).");
        Console.WriteLine("Delimiter: " + script.Delimiter);
    }
    catch (Exception ex)
    {
        Console.WriteLine(ex.ToString());
    }

    conn.Close();
    Console.WriteLine("Done.");
}

static void script_StatementExecuted(object sender, MySqlScriptEventArgs args)
{
    Console.WriteLine("script_StatementExecuted");
}

static void script_ScriptCompleted(object sender, EventArgs e)
{
    /// EventArgs e will be EventArgs.Empty for this method
    Console.WriteLine("script_ScriptCompleted!");
}

static void script_Error(Object sender, MySqlScriptErrorEventArgs args)
{
    Console.WriteLine("script_Error: " + args.Exception.ToString());
}
}
}
}

```

Note that in the `script_ScriptCompleted` event handler, the `EventArgs` parameter `e` will be `EventArgs.Empty`. In the case of the `ScriptCompleted` event there is no additional data to be obtained, which is why the event object is `EventArgs.Empty`.

#### 21.2.4.8.1. Using Delimiters with MySqlScript

Depending on the nature of the script, you may need control of the delimiter used to separate the statements that will make up a script. The most common example of this is where you have a multi-statement stored routine as part of your script. In this case if the default delimiter of “;” is used you will get an error when you attempt to execute the script. For example, consider the following stored routine:

```

CREATE PROCEDURE test_routine()
BEGIN
    SELECT name FROM TestTable ORDER BY name;
    SELECT COUNT(name) FROM TestTable;
END

```

This routine actually needs to be executed on the MySQL Server as a single statement. However, with the default delimiter of “;”, the `MySqlScript` class would interpret the above as two statements, the first being:

```

CREATE PROCEDURE test_routine()
BEGIN
    SELECT name FROM TestTable ORDER BY name;

```

Executing this as a statement would generate an error. To solve this problem `MySqlScript` supports the ability to set a different delimiter. This is achieved through the `Delimiter` property. For example, you could set the delimiter to “??”, in which case the above stored routine would no longer generate an error when executed. Multiple statements can be delimited in the script, so for example, you could have a three statement script such as:

```
string sql = "DROP PROCEDURE IF EXISTS test_routine??" +
    "CREATE PROCEDURE test_routine() " +
    "BEGIN " +
    "SELECT name FROM TestTable ORDER BY name;" +
    "SELECT COUNT(name) FROM TestTable;" +
    "END??" +
    "CALL test_routine()";
```

You can change the delimiter back at any point by setting the Delimiter property. The following code shows a complete working example:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

using MySql.Data;
using MySql.Data.MySqlClient;

namespace ConsoleApplication8
{
    class Program
    {
        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=TestDB;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);

            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                string sql = "DROP PROCEDURE IF EXISTS test_routine??" +
                    "CREATE PROCEDURE test_routine() " +
                    "BEGIN " +
                    "SELECT name FROM TestTable ORDER BY name;" +
                    "SELECT COUNT(name) FROM TestTable;" +
                    "END??" +
                    "CALL test_routine()";

                MySqlScript script = new MySqlScript(conn);

                script.Query = sql;
                script.Delimiter = "??";
                int count = script.Execute();
                Console.WriteLine("Executed " + count + " statement(s)");
                script.Delimiter = ";";
                Console.WriteLine("Delimiter: " + script.Delimiter);
                Console.WriteLine("Query: " + script.Query);
            }
            catch (Exception ex)
            {
                Console.WriteLine(ex.ToString());
            }

            conn.Close();
            Console.WriteLine("Done.");
        }
    }
}
```

#### 21.2.4.9. Tutorial: Generating MySQL DDL from an Entity Framework Model

In this tutorial, you will learn how to create MySQL [DDL](#) from an Entity Framework model. Use Visual Studio 2010 and MySQL Connector/Net 6.3 to carry out this tutorial.

1. Create a new console application in Visual Studio 2010.
2. Using the **Solution Explorer**, add a reference to `MySql.Data.Entity`.

3. From the **Solution Explorer** select **Add, New Item**. In the **Add New Item** dialog select **Online Templates**. Select **ADO.NET Entity Data Model** and click Add. The **Entity Data Model** dialog will be displayed.
4. In the **Entity Data Model** dialog select **Empty Model**. Click **Finish**. A blank model will be created.
5. Create a simple model. A single Entity will do for the purposes of this tutorial.
6. In the **Properties** panel select **ConceptualEntityModel** from the drop-down listbox.
7. In the **Properties** panel, locate the **DDL Generation Template** in the category **Database Script Generation**.
8. For the **DDL Generation** property select **SSDLToMySQL.tt(VS)** from the drop-down listbox.
9. Save the solution.
10. Right-click an empty space in the model design area. The context-sensitive menu will be displayed.
11. From the context-sensitive menu select **Generate Database from Model**. The **Generate Database Wizard** dialog will be displayed.
12. In the **Generate Database Wizard** dialog select an existing connection, or create a new connection to a server. Select an appropriate radio button to show or hide sensitive data. For the purposes of this tutorial you can select **Yes** (although you might skip this for commercial applications).
13. Click **Next**. MySQL compatible DDL code will be generated. Click **Finish** to exit the wizard.

You have seen how to create MySQL DDL code from an Entity Framework model.

## 21.2.5. Connector/Net Programming

Connector/Net comprises several classes that are used to connect to the database, execute queries and statements, and manage query results.

The following are the major classes of Connector/Net:

- **MySqlCommand**: Represents an SQL statement to execute against a MySQL database.
- **MySqlCommandBuilder**: Automatically generates single-table commands used to reconcile changes made to a DataSet with the associated MySQL database.
- **MySqlConnection**: Represents an open connection to a MySQL Server database.
- **MySqlDataAdapter**: Represents a set of data commands and a database connection that are used to fill a data set and update a MySQL database.
- **MySqlDataReader**: Provides a means of reading a forward-only stream of rows from a MySQL database.
- **MySqlException**: The exception that is thrown when MySQL returns an error.
- **MySqlHelper**: Helper class that makes it easier to work with the provider.
- **MySqlTransaction**: Represents an SQL **transaction** to be made in a MySQL database.

In the following sections, you will learn about some common use cases for Connector/Net, including BLOB handling, date handling, and using Connector/Net with common tools such as Crystal Reports.

### 21.2.5.1. Connecting to MySQL Using Connector/Net

All interaction between a .NET application and the MySQL server is routed through a **MySqlConnection** object. Before your application can interact with the server, it must instantiate, configure, and open a **MySqlConnection** object.

Even when using the `MySQLHelper` class, a `MySQLConnection` object is created by the helper class.

This section describes how to connect to MySQL using the `MySQLConnection` object.

### 21.2.5.2. Creating a Connector/Net Connection String

The `MySQLConnection` object is configured using a connection string. A connection string contains several key/value pairs, separated by semicolons. In each key/value pair, the option name and its corresponding value are joined by an equal sign. For the list of option names to use in the connection string, see [Section 21.2.6, "Connector/Net Connection String Options Reference"](#).

The following is a sample connection string:

```
Server=127.0.0.1;Uid=root;Pwd=12345;Database=test;
```

In this example, the `MySQLConnection` object is configured to connect to a MySQL server at `127.0.0.1`, with a user name of `root` and a password of `12345`. The default database for all statements will be the `test` database.



#### Note

Using the '@' symbol for parameters is now the preferred approach, although the old pattern of using '?' is still supported. To avoid conflicts when using the '@' symbol in combination with user variables, see the [Allow User Variables](#) connection string option in [Section 21.2.6, "Connector/Net Connection String Options Reference"](#). The [Old Syntax](#) connection string option has now been deprecated.

#### 21.2.5.2.1. Opening a Connection

Once you have created a connection string it can be used to open a connection to the MySQL server.

The following code is used to create a `MySQLConnection` object, assign the connection string, and open the connection.

Connector/Net can also connect using the native Windows authentication plugin. See [Section 21.2.5.5, "Using the Windows Native Authentication Plugin"](#) for details.

You can further extend the authentication mechanism by writing your own authentication plugin. See [Section 21.2.5.6, "Writing a Custom Authentication Plugin"](#) for details.

#### Visual Basic Example

```
Dim conn As New MySql.Data.MySqlClient.MySqlConnection
Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    conn.ConnectionString = myConnectionString
    conn.Open()

Catch ex As MySql.Data.MySqlClient.MySqlException
    MessageBox.Show(ex.Message)
End Try
```

#### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";
```

```

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection();
    conn.ConnectionString = myConnectionString;
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message);
}

```

You can also pass the connection string to the constructor of the [MySqlConnection](#) class:

### Visual Basic Example

```

Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    Dim conn As New MySql.Data.MySqlClient.MySqlConnection(myConnectionString)
    conn.Open()
Catch ex As MySql.Data.MySqlClient.MySqlException
    MessageBox.Show(ex.Message)
End Try

```

### C# Example

```

MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection(myConnectionString);
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message);
}

```

Once the connection is open it can be used by the other Connector/Net classes to communicate with the MySQL server.

#### 21.2.5.2.2. Handling Connection Errors

Because connecting to an external server is unpredictable, it is important to add error handling to your .NET application. When there is an error connecting, the [MySqlConnection](#) class will return a [MySqlException](#) object. This object has two properties that are of interest when handling errors:

- **Message:** A message that describes the current exception.
- **Number:** The MySQL error number.

When handling errors, you can your application's response based on the error number. The two most common error numbers when connecting are as follows:

- **0:** Cannot connect to server.
- **1045:** Invalid user name and/or password.

The following code shows how to adapt the application's response based on the actual error:

## Visual Basic Example

```

Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    Dim conn As New MySql.Data.MySqlClient.MySqlConnection(myConnectionString)
    conn.Open()
Catch ex As MySql.Data.MySqlClient.MySqlException
    Select Case ex.Number
        Case 0
            MessageBox.Show("Cannot connect to server. Contact administrator")
        Case 1045
            MessageBox.Show("Invalid username/password, please try again")
    End Select
End Try

```

## C# Example

```

MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection(myConnectionString);
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    switch (ex.Number)
    {
        case 0:
            MessageBox.Show("Cannot connect to server. Contact administrator");
        case 1045:
            MessageBox.Show("Invalid username/password, please try again");
    }
}

```



### Important

Note that if you are using multilanguage databases you must specify the character set in the connection string. If you do not specify the character set, the connection defaults to the `latin1` charset. You can specify the character set as part of the connection string, for example:

```

MySqlConnection myConnection = new MySqlConnection("server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;Charset=latin1;");

```

### 21.2.5.2.3. Using GetSchema on a Connection

The `GetSchema()` method of the connection object can be used to retrieve schema information about the database currently connected to. The schema information is returned in the form of a `DataTable`. The schema information is organized into a number of collections. Different forms of the `GetSchema()` method can be used depending on the information required. There are three forms of the `GetSchema()` method:

- `GetSchema()` - This call will return a list of available collections.
- `GetSchema(String)` - This call returns information about the collection named in the string parameter. If the string "MetaDataCollections" is used then a list of all available collections is returned. This is the same as calling `GetSchema()` without any parameters.

- `GetSchema(String, String[])` - In this call the first string parameter represents the collection name, and the second parameter represents a string array of restriction values. Restriction values limit the amount of data that will be returned. Restriction values are explained in more detail in the [Microsoft .NET documentation](#).

#### 21.2.5.2.3.1. Collections

The collections can be broadly grouped into two types: collections that are common to all data providers, and collections specific to a particular provider.

##### Common

The following collections are common to all data providers:

- `MetaDataCollections`
- `DataSourceInformation`
- `DataTypes`
- `Restrictions`
- `ReservedWords`

##### Provider-specific

The following are the collections currently provided by MySQL Connector/Net, in addition to the common collections above:

- `Databases`
- `Tables`
- `Columns`
- `Users`
- `Foreign Keys`
- `IndexColumns`
- `Indexes`
- `Foreign Key Columns`
- `UDF`
- `Views`
- `ViewColumns`
- `Procedure Parameters`
- `Procedures`
- `Triggers`

##### Example Code

A list of available collections can be obtained using the following code:

```
using System;
using System.Data;
using System.Text;
using MySql.Data;
```



```

using MySql.Data.MySqlClient;

namespace ConsoleApplication2
{
    class Program
    {
        private static void DisplayData(System.Data.DataTable table)
        {
            foreach (System.Data.DataRow row in table.Rows)
            {
                foreach (System.Data.DataColumn col in table.Columns)
                {
                    Console.WriteLine("{0} = {1}", col.ColumnName, row[col]);
                }
                Console.WriteLine("=====");
            }
        }

        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=world;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);

            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                DataTable table = conn.GetSchema("MetaDataCollections");
                //DataTable table = conn.GetSchema("UDF");
                DisplayData(table);

                conn.Close();
            }
            catch (Exception ex)
            {
                Console.WriteLine(ex.ToString());
            }
            Console.WriteLine("Done.");
        }
    }
}

```

Further information on the [GetSchema\(\)](#) method and schema collections can be found in the [Microsoft .NET documentation](#).

### 21.2.5.3. Using MySqlCommand

A [MySqlCommand](#) has the [CommandText](#) and [CommandType](#) properties associated with it. The [CommandText](#) will be handled differently depending on the setting of [CommandType](#). [CommandType](#) can be one of:

1. Text - A SQL text command (default)
2. StoredProcedure - The name of a Stored Procedure
3. TableDirect - The name of a table (new in Connector/Net 6.2)

The default [CommandType](#), [Text](#), is used for executing queries and other SQL commands. Some example of this can be found in the following section [Section 21.2.4.1.2, "The MySqlCommand Object"](#).

If [CommandType](#) is set to [StoredProcedure](#), set [CommandText](#) to the name of the Stored Procedure to access.

If [CommandType](#) is set to [TableDirect](#), all rows and columns of the named table will be returned when you call one of the Execute methods. In effect, this command performs a `SELECT *` on the table

specified. The `CommandText` property is set to the name of the table to query. This is illustrated by the following code snippet:

```
...
MySQLCommand cmd = new MySQLCommand();
cmd.CommandText = "mytable";
cmd.Connection = someConnection;
cmd.CommandType = CommandType.TableDirect;
MySQLDataReader reader = cmd.ExecuteReader();
while (reader.Read())
{
    Console.WriteLine(reader[0], reader[1]...);
}
...
```

Examples of using the `CommandType` of `StoredProcedure` can be found in the section [Section 21.2.5.9, "Accessing Stored Procedures with Connector/Net"](#).

Commands can have a timeout associated with them. This is useful as you may not want a situation where a command takes up an excessive amount of time. A timeout can be set using the `CommandTimeout` property. The following code snippet sets a timeout of one minute:

```
MySQLCommand cmd = new MySQLCommand();
cmd.CommandTimeout = 60;
```

The default value is 30 seconds. Avoid a value of 0, which indicates an indefinite wait. To change the default command timeout, use the connection string option `Default Command Timeout`.

Prior to MySQL Connector/Net 6.2, `MySQLCommand.CommandTimeout` included user processing time, that is processing time not related to direct use of the connector. Timeout was implemented through a .NET Timer, that triggered after `CommandTimeout` seconds. This timer consumed a thread.

MySQL Connector/Net 6.2 introduced timeouts that are aligned with how Microsoft handles `SqlCommand.CommandTimeout`. This property is the cumulative timeout for all network reads and writes during command execution or processing of the results. A timeout can still occur in the `MySQLReader.Read` method after the first row is returned, and does not include user processing time, only IO operations. The 6.2 implementation uses the underlying stream timeout facility, so is more efficient in that it does not require the additional timer thread as was the case with the previous implementation.

Further details on this can be found in the relevant [Microsoft documentation](#).

#### 21.2.5.4. Using Connector/Net with Connection Pooling

The Connector/Net supports connection pooling for better performance and scalability with database-intensive applications. This is enabled by default. You can turn it off or adjust its performance characteristics using the connection string options `Pooling`, `Connection Reset`, `Connection Lifetime`, `Cache Server Properties`, `Max Pool Size` and `Min Pool Size`. See [Section 21.2.5.2, "Creating a Connector/Net Connection String"](#) for further information.

Connection pooling works by keeping the native connection to the server live when the client disposes of a `MySQLConnection`. Subsequently, if a new `MySQLConnection` object is opened, it will be created from the connection pool, rather than creating a new native connection. This improves performance.

#### Guidelines

To work as designed, it is best to let the connection pooling system manage all connections. Do not create a globally accessible instance of `MySQLConnection` and then manually open and close it. This interferes with the way the pooling works and can lead to unpredictable results or even exceptions.

One approach that simplifies things is to avoid manually creating a `MySQLConnection` object. Instead use the overloaded methods that take a connection string as an argument. Using this approach,

Connector/Net will automatically create, open, close and destroy connections, using the connection pooling system for best performance.

Typed Datasets and the [MembershipProvider](#) and [RoleProvider](#) classes use this approach. Most classes that have methods that take a [MySQLConnection](#) as an argument, also have methods that take a connection string as an argument. This includes [MySQLDataAdapter](#).

Instead of manually creating [MySQLCommand](#) objects, you can use the static methods of the [MySQLHelper](#) class. These take a connection string as an argument, and they fully support connection pooling.

## Resource Usage

Starting with MySQL Connector/Net 6.2, there is a background job that runs every three minutes and removes connections from pool that have been idle (unused) for more than three minutes. The pool cleanup frees resources on both client and server side. This is because on the client side every connection uses a socket, and on the server side every connection uses a socket and a thread.

Prior to this change, connections were never removed from the pool, and the pool always contained the peak number of open connections. For example, a web application that peaked at 1000 concurrent database connections would consume 1000 threads and 1000 open sockets at the server, without ever freeing up those resources from the connection pool. Note, connections, no matter how old, will not be closed if the number of connections in the pool is less than or equal to the value set by the [Min Pool Size](#) connection string parameter.

### 21.2.5.5. Using the Windows Native Authentication Plugin

Connector/Net applications can authenticate to a MySQL server using the Windows Native Authentication Plugin as of Connector/Net 6.4.4 and MySQL 5.5.16/5.6.10. Users who have logged in to Windows can connect from MySQL client programs to the server based on the information in their environment without specifying an additional password. For background and usage information about the authentication plugin, see, [Section 6.3.7.4, “The Windows Native Authentication Plugin”](#).

The interface matches the [MySQL.Data.MySqlClient](#) object. To enable, pass in [Integrated Security](#) to the connection string with a value of `yes` or `sspi`.

Passing in a user ID is optional. When Windows authentication is set up, a MySQL user is created and configured to be used by Windows authentication. By default, this user ID is named `auth_windows`, but can be defined using a different name. If the default name is used, then passing the user ID to the connection string from Connector/Net is optional, because it will use the `auth_windows` user. Otherwise, the name must be passed to the [connection string](#) using the standard user ID element.

### 21.2.5.6. Writing a Custom Authentication Plugin

Advanced users with special security requirements can create their own authentication plugins for Connector/Net applications. You can extend the handshake protocol, adding custom logic. This capability requires Connector/Net 6.6.3 or higher, and MySQL 5.5.16 or higher. For background and usage information about MySQL authentication plugins, see, [Section 22.2.3.7, “Authentication Plugins”](#) and [Section 22.2.4.9, “Writing Authentication Plugins”](#).

To write a custom authentication plugin, you will need a reference to the assembly `MySQL.Data.dll`. The classes relevant for writing authentication plugins are available at the namespace `MySQL.Data.MySqlClient.Authentication`.

## How the Custom Authentication Plugin Works

At some point during handshake, the internal method

```
void Authenticate(bool reset)
```

of [MySQLAuthenticationPlugin](#) is called. This method in turns calls several overridable methods of the current plugin.

## Creating the Authentication Plugin Class

You put the authentication plugin logic inside a new class derived from `MySql.Data.MySqlClient.Authentication.MySqlAuthenticationPlugin`. The following methods are available to be overridden:

```
protected virtual void CheckConstraints()
protected virtual void AuthenticationFailed(Exception ex)
protected virtual void AuthenticationSuccessful()
protected virtual byte[] MoreData(byte[] data)
protected virtual void AuthenticationChange()
public abstract string PluginName { get; }
public virtual string GetUsername()
public virtual object GetPassword()
protected byte[] AuthData;
```

The following is a brief explanation of each one:

```
/// <summary>
/// This method must check authentication method specific constraints in the
environment and throw an Exception
/// if the conditions are not met. The default implementation does nothing.
/// </summary>
protected virtual void CheckConstraints()

/// <summary>
/// This method, called when the authentication failed, provides a chance to
plugins to manage the error
/// the way they consider decide (either showing a message, logging it, etc.).
/// The default implementation wraps the original exception in a MySqlConnection
with an standard message and rethrows it.
/// </summary>
/// <param name="ex">The exception with extra information on the error.</param>
protected virtual void AuthenticationFailed(Exception ex)

/// <summary>
/// This method is invoked when the authentication phase was successful accepted
by the server.
/// Derived classes must override this if they want to be notified of such
condition.
/// </summary>
/// <remarks>The default implementation does nothing.</remarks>
protected virtual void AuthenticationSuccessful()

/// <summary>
/// This method provides a chance for the plugin to send more data when the
server requests so during the
/// authentication phase. This method will be called at least once, and more
than one depending upon whether the
/// server response packets have the 0x01 prefix.
/// </summary>
/// <param name="data">The response data from the server, during the
authentication phase the first time is called is null, in
subsequent calls contains the server response.</param>
/// <returns>The data generated by the plugin for server consumption.</returns>
/// <remarks>The default implementation always returns null.</remarks>
protected virtual byte[] MoreData(byte[] data)

/// <summary>
/// The plugin name.
/// </summary>
public abstract string PluginName { get; }

/// <summary>
/// Gets the user name to send to the server in the authentication phase.
/// </summary>
/// <returns>An string with the user name</returns>
/// <remarks>Default implementation returns the UserId passed from the
connection string.</remarks>
public virtual string GetUsername()
```

```

/// <summary>
/// Gets the password to send to the server in the authentication phase. This
/// can be a string or a
/// </summary>
/// <returns>An object, can be byte[], string or null, with the password.
</returns>
/// <remarks>Default implementation returns null.</remarks>
public virtual object GetPassword()

/// <summary>
/// The authentication data passed when creating the plugin.
/// For example in mysql_native_password this is the seed to encrypt the
/// password.
/// </summary>
protected byte[] AuthData;

```

## Sample Authentication Plugin

Here is an example showing how to create the authentication plugin, then enable it by means of a configuration file. Follow these steps:

1. Create a console app, adding a reference to `MySql.Data.dll`.
2. Design the main program as follows:

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using MySql.Data.MySqlClient;

namespace AuthPluginTest
{
    class Program
    {
        static void Main(string[] args)
        {
            // Customize the connection string as necessary.
            MySqlConnection con = new MySqlConnection("server=localhost;
database=test; user id=myuser; password=mypass");
            con.Open();
            con.Close();
        }
    }
}

```

3. Create your plugin class. In this example, we add an “alternative” implementation of the Native password plugin by just using the same code from the original plugin. We name our class `MySqlNativePasswordPlugin2`:

```

using System.IO;
using System;
using System.Text;
using System.Security.Cryptography;
using MySql.Data.MySqlClient.Authentication;
using System.Diagnostics;

namespace AuthPluginTest
{
    public class MySqlNativePasswordPlugin2 : MySqlAuthenticationPlugin
    {
        public override string PluginName
        {
            get { return "mysql_native_password"; }
        }

        public override object GetPassword()
        {

```

```

    Debug.WriteLine("Calling MySqlNativePasswordPlugin2.GetPassword");
    return Get411Password(Settings.Password, AuthData);
}

/// <summary>
/// Returns a byte array containing the proper encryption of the
/// given password/seed according to the new 4.1.1 authentication scheme.
/// </summary>
/// <param name="password"></param>
/// <param name="seed"></param>
/// <returns></returns>
private byte[] Get411Password(string password, byte[] seedBytes)
{
    // if we have no password, then we just return 1 zero byte
    if (password.Length == 0) return new byte[1];

    SHA1 sha = new SHA1CryptoServiceProvider();

    byte[] firstHash = sha.ComputeHash(Encoding.Default.GetBytes(password));
    byte[] secondHash = sha.ComputeHash(firstHash);

    byte[] input = new byte[seedBytes.Length + secondHash.Length];
    Array.Copy(seedBytes, 0, input, 0, seedBytes.Length);
    Array.Copy(secondHash, 0, input, seedBytes.Length, secondHash.Length);
    byte[] thirdHash = sha.ComputeHash(input);

    byte[] finalHash = new byte[thirdHash.Length + 1];
    finalHash[0] = 0x14;
    Array.Copy(thirdHash, 0, finalHash, 1, thirdHash.Length);

    for (int i = 1; i < finalHash.Length; i++)
        finalHash[i] = (byte)(finalHash[i] ^ firstHash[i - 1]);
    return finalHash;
}
}
}

```

4. Notice that the plugin implementation just overrides `GetPassword`, and provides an implementation to encrypt the password using the 4.1 protocol. We also put the following line in the `GetPassword` body:

```
Debug.WriteLine("Calling MySqlNativePasswordPlugin2.GetPassword");
```

to provide confirmation that the plugin was effectively used. (You could also put a breakpoint on that method.)

5. Enable the new plugin in the configuration file:

```

<?xml version="1.0"?>
<configuration>
  <configSections>
    <section name="MySQL" type="MySql.Data.MySqlClient.MySqlConfiguration,
MySql.Data"/>
  </configSections>
  <MySQL>
    <AuthenticationPlugins>
      <add name="mysql_native_password"
type="AuthPluginTest.MySqlNativePasswordPlugin2, AuthPluginTest"></add>
    </AuthenticationPlugins>
  </MySQL>
</startup><supportedRuntime version="v4.0" sku=".NETFramework,Version=v4.0"/>
</startup></configuration>

```

6. Run the application. In Visual Studio, you will see the message `Calling MySqlNativePasswordPlugin2.GetPassword` in the debug window.
7. Continue enhancing the authentication logic, overriding more methods if you required.

### 21.2.5.7. Using Connector/Net with Table Caching

This feature exists with Connector/Net versions 6.4 and above.

Table caching is a feature that can be used to cache slow-changing datasets on the client side. This is useful for applications that are designed to use readers, but still want to minimize trips to the server for slow-changing tables.

This feature is transparent to the application, and is disabled by default.

## Configuration

- To enable table caching, add `'table cache = true'` to the connection string.
- Optionally, specify the `'Default Table Cache Age'` connection string option, which represents the number of seconds a table is cached before the cached data is discarded. The default value is `60`.
- You can turn caching on and off and set caching options at runtime, on a per-command basis.

### 21.2.5.8. Using the Connector/Net with Prepared Statements

As of MySQL 4.1, it is possible to use prepared statements with Connector/Net. Use of prepared statements can provide significant performance improvements on queries that are executed more than once.

Prepared execution is faster than direct execution for statements executed more than once, primarily because the query is parsed only once. In the case of direct execution, the query is parsed every time it is executed. Prepared execution also can provide a reduction of network traffic because for each execution of the prepared statement, it is necessary only to send the data for the parameters.

Another advantage of prepared statements is that it uses a binary protocol that makes data transfer between client and server more efficient.

#### 21.2.5.8.1. Preparing Statements in Connector/Net

To prepare a statement, create a command object and set the `.CommandText` property to your query.

After entering your statement, call the `.Prepare` method of the `MySQLCommand` object. After the statement is prepared, add parameters for each of the dynamic elements in the query.

After you enter your query and enter parameters, execute the statement using the `.ExecuteNonQuery()`, `.ExecuteScalar()`, or `.ExecuteReader` methods.

For subsequent executions, you need only modify the values of the parameters and call the execute method again, there is no need to set the `.CommandText` property or redefine the parameters.

#### Visual Basic Example

```
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

conn.ConnectionString = strConnection

Try
    conn.Open()
    cmd.Connection = conn

    cmd.CommandText = "INSERT INTO myTable VALUES(NULL, @number, @text)"
    cmd.Prepare()

    cmd.Parameters.AddWithValue("@number", 1)
    cmd.Parameters.AddWithValue("@text", "One")

    For i = 1 To 1000
        cmd.Parameters("@number").Value = i
        cmd.Parameters("@text").Value = "A string value"

        cmd.ExecuteNonQuery()
```

```

        Next
    Catch ex As MySqlException
        MessageBox.Show("Error " & ex.Number & " has occurred: " & ex.Message, "Error", MessageBoxButtons.OK, MessageBoxIcon.Error)
    End Try

```

## C# Example

```

MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();

conn.ConnectionString = strConnection;

try
{
    conn.Open();
    cmd.Connection = conn;

    cmd.CommandText = "INSERT INTO myTable VALUES(NULL, @number, @text)";
    cmd.Prepare();

    cmd.Parameters.AddWithValue("@number", 1);
    cmd.Parameters.AddWithValue("@text", "One");

    for (int i=1; i <= 1000; i++)
    {
        cmd.Parameters["@number"].Value = i;
        cmd.Parameters["@text"].Value = "A string value";

        cmd.ExecuteNonQuery();
    }
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

### 21.2.5.9. Accessing Stored Procedures with Connector/Net

MySQL server version 5 and up supports stored procedures with the SQL 2003 stored procedure syntax.

A stored procedure is a set of SQL statements that is stored in the server. Clients make a single call to the stored procedure, passing parameters that can influence the procedure logic and query conditions, rather than issuing individual hardcoded SQL statements.

Stored procedures can be particularly useful in situations such as the following:

- Stored procedures can act as an API or abstraction layer, allowing multiple client applications to perform the same database operations. The applications can be written in different languages and run on different platforms. The applications do not need to hardcode table and column names, complicated queries, and so on. When you extend and optimize the queries in a stored procedure, all the applications that call the procedure automatically receive the benefits.
- When security is paramount, stored procedures keep applications from directly manipulating tables, or even knowing details such as table and column names. Banks, for example, use stored procedures for all common operations. This provides a consistent and secure environment, and procedures can ensure that each operation is properly logged. In such a setup, applications and users would not get any access to the database tables directly, but can only execute specific stored procedures.

Connector/Net supports the calling of stored procedures through the [MySqlCommand](#) object. Data can be passed in and out of a MySQL stored procedure through use of the [MySqlCommand.Parameters](#) collection.



**Note**

When you call a stored procedure, the command object makes an additional `SELECT` call to determine the parameters of the stored procedure. You must ensure that the user calling the procedure has the `SELECT [755]` privilege on the `mysql.proc` table to enable them to verify the parameters. Failure to do this will result in an error when calling the procedure.

This section will not provide in-depth information on creating Stored Procedures. For such information, please refer to <http://dev.mysql.com/doc/mysql/en/stored-routines.html>.

A sample application demonstrating how to use stored procedures with Connector/Net can be found in the `Samples` directory of your Connector/Net installation.

### 21.2.5.9.1. Using Stored Routines from Connector/Net

Stored procedures in MySQL can be created using a variety of tools. First, stored procedures can be created using the `mysql` command-line client. Second, stored procedures can be created using MySQL Workbench. Finally, stored procedures can be created using the `.ExecuteNonQuery` method of the `MySqlCommand` object.

Unlike the command-line and GUI clients, you are not required to specify a special delimiter when creating stored procedures in Connector/Net.

To call a stored procedure using Connector/Net, you create a `MySqlCommand` object and pass the stored procedure name as the `.CommandText` property. You then set the `.CommandType` property to `CommandType.StoredProcedure`.

After the stored procedure is named, you create one `MySqlCommand` parameter for every parameter in the stored procedure. `IN` parameters are defined with the parameter name and the object containing the value, `OUT` parameters are defined with the parameter name and the data type that is expected to be returned. All parameters need the parameter direction defined.

After defining the parameters, you call the stored procedure by using the `MySqlCommand.ExecuteNonQuery()` method.

Once the stored procedure is called, the values of the output parameters can be retrieved by using the `.Value` property of the `MySqlConnection.Parameters` collection.

**Note**

When a stored procedure is called using `MySqlCommand.ExecuteReader`, and the stored procedure has output parameters, the output parameters are only set after the `MySqlDataReader` returned by `ExecuteReader` is closed.

The following C# example code demonstrates the use of stored procedures. It assumes the database 'employees' has already been created:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

using System.Data;
using MySql.Data;
using MySql.Data.MySqlClient;

namespace UsingStoredRoutines
{
    class Program
    {
        static void Main(string[] args)
        {
            MySqlConnection conn = new MySqlConnection();
            conn.ConnectionString = "server=localhost;user=root;database=employees;port=3306;password=";
```

```

MySQLCommand cmd = new MySQLCommand();

try
{
    Console.WriteLine("Connecting to MySQL...");
    conn.Open();
    cmd.Connection = conn;
    cmd.CommandText = "DROP PROCEDURE IF EXISTS add_emp";
    cmd.ExecuteNonQuery();
    cmd.CommandText = "DROP TABLE IF EXISTS emp";
    cmd.ExecuteNonQuery();
    cmd.CommandText = "CREATE TABLE emp (empno INT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY
    cmd.ExecuteNonQuery();

    cmd.CommandText = "CREATE PROCEDURE add_emp(" +
        "IN fname VARCHAR(20), IN lname VARCHAR(20), IN bday DATETIME, OUT empno
        "BEGIN INSERT INTO emp(first_name, last_name, birthdate) " +
        "VALUES(fname, lname, DATE(bday)); SET empno = LAST_INSERT_ID(); END";

    cmd.ExecuteNonQuery();
}
catch (MySQLException ex)
{
    Console.WriteLine ("Error " + ex.Number + " has occurred: " + ex.Message);
}
conn.Close();
Console.WriteLine("Connection closed.");
try
{
    Console.WriteLine("Connecting to MySQL...");
    conn.Open();
    cmd.Connection = conn;

    cmd.CommandText = "add_emp";
    cmd.CommandType = CommandType.StoredProcedure;

    cmd.Parameters.AddWithValue("@lname", "Jones");
    cmd.Parameters["@lname"].Direction = ParameterDirection.Input;

    cmd.Parameters.AddWithValue("@fname", "Tom");
    cmd.Parameters["@fname"].Direction = ParameterDirection.Input;

    cmd.Parameters.AddWithValue("@bday", "1940-06-07");
    cmd.Parameters["@bday"].Direction = ParameterDirection.Input;

    cmd.Parameters.AddWithValue("@empno", MySQLDbType.Int32);
    cmd.Parameters["@empno"].Direction = ParameterDirection.Output;

    cmd.ExecuteNonQuery();

    Console.WriteLine("Employee number: "+cmd.Parameters["@empno"].Value);
    Console.WriteLine("Birthday: " + cmd.Parameters["@bday"].Value);
}
catch (MySQL.Data.MySqlClient.MySQLException ex)
{
    Console.WriteLine("Error " + ex.Number + " has occurred: " + ex.Message);
}
conn.Close();
Console.WriteLine("Done.");
}
}
}

```

The following code shows the same application in Visual Basic:

```

Imports System
Imports System.Collections.Generic
Imports System.Linq
Imports System.Text

Imports System.Data

```

```

Imports MySql.Data
Imports MySql.Data.MySqlClient

Module Module1

    Sub Main()
        Dim conn As New MySqlConnection()
        conn.ConnectionString = "server=localhost;user=root;database=world;port=3306;password=*****;"
        Dim cmd As New MySqlCommand()

        Try
            Console.WriteLine("Connecting to MySQL...")
            conn.Open()
            cmd.Connection = conn
            cmd.CommandText = "DROP PROCEDURE IF EXISTS add_emp"
            cmd.ExecuteNonQuery()
            cmd.CommandText = "DROP TABLE IF EXISTS emp"
            cmd.ExecuteNonQuery()
            cmd.CommandText = "CREATE TABLE emp (empno INT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY

            cmd.CommandText = "CREATE PROCEDURE add_emp(" & "IN fname VARCHAR(20), IN lname VARCHAR(20)

            cmd.ExecuteNonQuery()
        Catch ex As MySqlException
            Console.WriteLine(("Error " & ex.Number & " has occurred: ") + ex.Message)
        End Try
        conn.Close()
        Console.WriteLine("Connection closed.")
        Try
            Console.WriteLine("Connecting to MySQL...")
            conn.Open()
            cmd.Connection = conn

            cmd.CommandText = "add_emp"
            cmd.CommandType = CommandType.StoredProcedure

            cmd.Parameters.AddWithValue("@lname", "Jones")
            cmd.Parameters("@lname").Direction = ParameterDirection.Input

            cmd.Parameters.AddWithValue("@fname", "Tom")
            cmd.Parameters("@fname").Direction = ParameterDirection.Input

            cmd.Parameters.AddWithValue("@bday", "1940-06-07")
            cmd.Parameters("@bday").Direction = ParameterDirection.Input

            cmd.Parameters.AddWithValue("@empno", MySqlDbType.Int32)
            cmd.Parameters("@empno").Direction = ParameterDirection.Output

            cmd.ExecuteNonQuery()

            Console.WriteLine("Employee number: " & cmd.Parameters("@empno").Value)
            Console.WriteLine("Birthday: " & cmd.Parameters("@bday").Value)
        Catch ex As MySql.Data.MySqlClient.MySqlException
            Console.WriteLine(("Error " & ex.Number & " has occurred: ") + ex.Message)
        End Try
        conn.Close()
        Console.WriteLine("Done.")

    End Sub

End Module

```

### 21.2.5.10. Handling BLOB Data With Connector/Net

One common use for MySQL is the storage of binary data in [BLOB](#) columns. MySQL supports four different BLOB data types: [TINYBLOB](#), [BLOB](#), [MEDIUMBLOB](#), and [LONGBLOB](#), all described in [Section 11.4.3, “The BLOB and TEXT Types”](#) and [Section 11.6, “Data Type Storage Requirements”](#).

Data stored in a [BLOB](#) column can be accessed using Connector/Net and manipulated using client-side code. There are no special requirements for using Connector/Net with [BLOB](#) data.

Simple code examples will be presented within this section, and a full sample application can be found in the [Samples](#) directory of the Connector/Net installation.

### 21.2.5.10.1. Preparing the MySQL Server

The first step is using MySQL with **BLOB** data is to configure the server. Let's start by creating a table to be accessed. In my file tables, I usually have four columns: an **AUTO\_INCREMENT** column of appropriate size (**UNSIGNED SMALLINT**) to serve as a primary key to identify the file, a **VARCHAR** column that stores the file name, an **UNSIGNED MEDIUMINT** column that stores the size of the file, and a **MEDIUMBLOB** column that stores the file itself. For this example, I will use the following table definition:

```
CREATE TABLE file(  
file_id SMALLINT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,  
file_name VARCHAR(64) NOT NULL,  
file_size MEDIUMINT UNSIGNED NOT NULL,  
file MEDIUMBLOB NOT NULL);
```

After creating a table, you might need to modify the `max_allowed_packet` system variable. This variable determines how large of a packet (that is, a single row) can be sent to the MySQL server. By default, the server only accepts a maximum size of 1MB from the client application. If you intend to exceed 1MB in your file transfers, increase this number.

The `max_allowed_packet` option can be modified using the MySQL Workbench **Server Administration** screen. Adjust the Maximum permitted option in the **Data / Memory size** section of the Networking tab to an appropriate setting. After adjusting the value, click the **Apply** button and restart the server using the **Startup / Shutdown** screen of MySQL Workbench. You can also adjust this value directly in the `my.cnf` file (add a line that reads `max_allowed_packet=xxM`), or use the `SET max_allowed_packet=xxM;` syntax from within MySQL.

Try to be conservative when setting `max_allowed_packet`, as transfers of BLOB data can take some time to complete. Try to set a value that will be adequate for your intended use and increase the value if necessary.

### 21.2.5.10.2. Writing a File to the Database

To write a file to a database, we need to convert the file to a byte array, then use the byte array as a parameter to an **INSERT** query.

The following code opens a file using a **FileStream** object, reads it into a byte array, and inserts it into the `file` table:

#### Visual Basic Example

```
Dim conn As New MySqlConnection  
Dim cmd As New MySqlCommand  
  
Dim SQL As String  
  
Dim fileSize As UInt32  
Dim rawData() As Byte  
Dim fs As FileStream  
  
conn.ConnectionString = "server=127.0.0.1;" _  
    & "uid=root;" _  
    & "pwd=12345;" _  
    & "database=test"  
  
Try  
    fs = New FileStream("c:\image.png", FileMode.Open, FileAccess.Read)  
    fileSize = fs.Length  
  
    rawData = New Byte(fileSize) {}  
    fs.Read(rawData, 0, fileSize)  
    fs.Close()
```

```

conn.Open()

SQL = "INSERT INTO file VALUES(NULL, @FileName, @FileSize, @File)"

cmd.Connection = conn
cmd.CommandText = SQL
cmd.Parameters.AddWithValue("@FileName", strFileName)
cmd.Parameters.AddWithValue("@FileSize", FileSize)
cmd.Parameters.AddWithValue("@File", rawData)

cmd.ExecuteNonQuery()

MessageBox.Show("File Inserted into database successfully!", _
"Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk)

conn.Close()
Catch ex As Exception
    MessageBox.Show("There was an error: " & ex.Message, "Error", _
        MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try

```

### C# Example

```

MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();

string SQL;
UInt32 FileSize;
byte[] rawData;
FileStream fs;

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    fs = new FileStream(@"c:\image.png", FileMode.Open, FileAccess.Read);
    FileSize = fs.Length;

    rawData = new byte[FileSize];
    fs.Read(rawData, 0, FileSize);
    fs.Close();

    conn.Open();

    SQL = "INSERT INTO file VALUES(NULL, @FileName, @FileSize, @File)";

    cmd.Connection = conn;
    cmd.CommandText = SQL;
    cmd.Parameters.AddWithValue("@FileName", strFileName);
    cmd.Parameters.AddWithValue("@FileSize", FileSize);
    cmd.Parameters.AddWithValue("@File", rawData);

    cmd.ExecuteNonQuery();

    MessageBox.Show("File Inserted into database successfully!",
        "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk);

    conn.Close();
}
catch (MySQL.Data.MySqlClient.MySQLException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

The [Read](#) method of the [FileStream](#) object is used to load the file into a byte array which is sized according to the [Length](#) property of the [FileStream](#) object.

After assigning the byte array as a parameter of the `MySQLCommand` object, the `ExecuteNonQuery` method is called and the `BLOB` is inserted into the `file` table.

### 21.2.5.10.3. Reading a BLOB from the Database to a File on Disk

Once a file is loaded into the `file` table, we can use the `MySQLDataReader` class to retrieve it.

The following code retrieves a row from the `file` table, then loads the data into a `FileStream` object to be written to disk:

#### Visual Basic Example

```
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myData As MySQLDataReader
Dim SQL As String
Dim rawData() As Byte
Dim fileSize As UInt32
Dim fs As FileStream

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

SQL = "SELECT file_name, file_size, file FROM file"

Try
    conn.Open()

    cmd.Connection = conn
    cmd.CommandText = SQL

    myData = cmd.ExecuteReader

    If Not myData.HasRows Then Throw New Exception("There are no BLOBs to save")

    myData.Read()

    fileSize = myData.GetUInt32(myData.GetOrdinal("file_size"))
    rawData = New Byte(fileSize) {}

    myData.GetBytes(myData.GetOrdinal("file"), 0, rawData, 0, fileSize)

    fs = New FileStream("C:\newfile.png", FileMode.OpenOrCreate, FileAccess.Write)
    fs.Write(rawData, 0, fileSize)
    fs.Close()

    MessageBox.Show("File successfully written to disk!", "Success!", MessageBoxButtons.OK, MessageBoxIcon.Information)

    myData.Close()
    conn.Close()
Catch ex As Exception
    MessageBox.Show("There was an error: " & ex.Message, "Error", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try
```

#### C# Example

```
MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;
MySQL.Data.MySqlClient.MySqlDataReader myData;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();

string SQL;
UInt32 fileSize;
byte[] rawData;
FileStream fs;
```

```

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

SQL = "SELECT file_name, file_size, file FROM file";

try
{
    conn.Open();

    cmd.Connection = conn;
    cmd.CommandText = SQL;

    myData = cmd.ExecuteReader();

    if (!myData.HasRows)
        throw new Exception("There are no BLOBs to save");

    myData.Read();

    FileSize = myData.GetUInt32(myData.GetOrdinal("file_size"));
    rawData = new byte[FileSize];

    myData.GetBytes(myData.GetOrdinal("file"), 0, rawData, 0, (int)FileSize);

    fs = new FileStream(@"C:\newfile.png", FileMode.OpenOrCreate, FileAccess.Write);
    fs.Write(rawData, 0, (int)FileSize);
    fs.Close();

    MessageBox.Show("File successfully written to disk!",
        "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk);

    myData.Close();
    conn.Close();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

After connecting, the contents of the `file` table are loaded into a `MySqlDataReader` object. The `GetBytes` method of the `MySqlDataReader` is used to load the `BLOB` into a byte array, which is then written to disk using a `FileStream` object.

The `GetOrdinal` method of the `MySqlDataReader` can be used to determine the integer index of a named column. Use of the `GetOrdinal` method prevents errors if the column order of the `SELECT` query is changed.

### 21.2.5.11. Using the Connector/Net Interceptor Classes

An interceptor is a software design pattern that provides a transparent way to extend or modify some aspect of a program, similar to a user exit. No recompiling is required. With Connector/Net, the interceptors are enabled and disabled by updating the connection string to refer to different sets of interceptor classes that you instantiate.

Connector/Net includes the following interceptor classes:

- The `BaseCommandInterceptor` lets you perform additional operations when a program issues a SQL command. For example, you can examine the SQL statement for logging or debugging purposes, substitute your own result set to implement a caching mechanism, and so on. Depending on the use case, your code can supplement the SQL command or replace it entirely.

The `BaseCommandInterceptor` class has these methods that you can override:

```

public virtual bool ExecuteScalar(string sql, ref object returnValue);
public virtual bool ExecuteNonQuery(string sql, ref int returnValue);
public virtual bool ExecuteReader(string sql, CommandBehavior behavior, ref MySqlDataReader returnValue);
public virtual void Init(MySqlConnection connection);

```

If your interceptor overrides one of the `Execute...` methods, set the `returnValue` output parameter and return `true` if you handled the event, or `false` if you did not handle the event. The SQL command is processed normally only when all command interceptors return `false`.

The connection passed to the `Init` method is the connection that is attached to this interceptor.

- The `BaseExceptionInterceptor` lets you perform additional operations when a program encounters a SQL exception. The exception interception mechanism is modeled after the Connector/J model. You can code an interceptor class and connect it to an existing program without recompiling, and intercept exceptions when they are created. You can then change the exception type and optionally attach information to it. This capability lets you turn on and off logging and debugging code without hardcoding anything in the application. This technique applies to exceptions raised at the SQL level, not to lower-level system or I/O errors.

You develop an exception interceptor first by creating a subclass of the `BaseExceptionInterceptor` class. You must override the `InterceptException()` method. You can also override the `Init()` method to do some one-time initialization.

Each exception interceptor has 2 methods:

```
public abstract Exception InterceptException(Exception exception,
    MySqlConnection connection);
public virtual void Init(MySqlConnection connection);
```

The connection passed to `Init()` is the connection that is attached to this interceptor.

Each interceptor is required to override `InterceptException` and return an exception. It can return the exception it is given, or it can wrap it in a new exception. We currently do not offer the ability to suppress the exception.

Here are examples of using the FQN (fully qualified name) on the connection string:

```
MySqlConnection c1 = new MySqlConnection(@"server=localhost;pooling=false;
commandinterceptors=CommandApp.MyCommandInterceptor,CommandApp");

MySqlConnection c2 = new MySqlConnection(@"server=localhost;pooling=false;
exceptioninterceptors=ExceptionStackTraceTest.MyExceptionInterceptor,ExceptionStackTraceTest");
```

In this example, the command interceptor is called `CommandApp.MyCommandInterceptor` and exists in the `CommandApp` assembly. The exception interceptor is called `ExceptionStackTraceTest.MyExceptionInterceptor` and exists in the `ExceptionStackTraceTest` assembly.

To shorten the connection string, you can register your exception interceptors in your `app.config` or `web.config` file like this:

```
<configSections>
<section name="MySQL" type="MySql.Data.MySqlClient.MySqlConfiguration,
MySql.Data" />
</configSections>
<MySQL>
<CommandInterceptors>
<add name="myC" type="CommandApp.MyCommandInterceptor,CommandApp" />
</CommandInterceptors>
</MySQL>

<configSections>
<section name="MySQL" type="MySql.Data.MySqlClient.MySqlConfiguration,
MySql.Data" />
</configSections>
<MySQL>
<ExceptionInterceptors>
<add name="myE"
type="ExceptionStackTraceTest.MyExceptionInterceptor,ExceptionStackTraceTest" />
</ExceptionInterceptors>
</MySQL>
```



```
</ExceptionInterceptors>  
</MySQL>
```

Once you have done that, your connection strings can look like these:

```
MySqlConnection c1 = new MySqlConnection(@"server=localhost;pooling=false;  
commandinterceptors=myC");  
  
MySqlConnection c2 = new MySqlConnection(@"server=localhost;pooling=false;  
exceptioninterceptors=myE");
```

### 21.2.5.12. Handling Date and Time Information in Connector/Net

MySQL and the .NET languages handle date and time information differently, with MySQL allowing dates that cannot be represented by a .NET data type, such as '0000-00-00 00:00:00'. These differences can cause problems if not properly handled.

The following sections demonstrate how to properly handle date and time information when using Connector/Net.

#### 21.2.5.12.1. Fractional Seconds

Connector/Net 6.5 and higher support the fractional seconds feature introduced in MySQL 5.6.4. Fractional seconds could always be specified in a date literal or passed back and forth as parameters and return values, but the fractional part was always stripped off when stored in a table column. In MySQL 5.6.4 and higher, the fractional part is now preserved in data stored and retrieved through SQL. For fractional second handling in MySQL 5.6.4 and higher, see [Section 11.3.6, "Fractional Seconds in Time Values"](#). For the behavior of fractional seconds prior to MySQL 5.6.4, see [Fractional Seconds in Time Values](#).

To use the more precise date and time types, specify a value from 1 to 6 when creating the table column, for example `TIME(3)` or `DATETIME(6)`, representing the number of digits of precision after the decimal point. Specifying a precision of 0 leaves the fractional part out entirely. In your C# or Visual Basic code, refer to the `Millisecond` member to retrieve the fractional second value from the `MySqlDateTime` object returned by the `GetMySqlDateTime` function. The `DateTime` object returned by the `GetDateTime` function also contains the fractional value, but only the first 3 digits.

For related code examples, see the following blog post: [https://blogs.oracle.com/MySqlOnWindows/entry/milliseconds\\_value\\_support\\_on\\_datetime](https://blogs.oracle.com/MySqlOnWindows/entry/milliseconds_value_support_on_datetime)

#### 21.2.5.12.2. Problems when Using Invalid Dates

The differences in date handling can cause problems for developers who use invalid dates. Invalid MySQL dates cannot be loaded into native .NET `DateTime` objects, including `NULL` dates.

Because of this issue, .NET `DataSet` objects cannot be populated by the `Fill` method of the `MySqlDataAdapter` class as invalid dates will cause a `System.ArgumentOutOfRangeException` exception to occur.

#### 21.2.5.12.3. Restricting Invalid Dates

The best solution to the date problem is to restrict users from entering invalid dates. This can be done on either the client or the server side.

Restricting invalid dates on the client side is as simple as always using the .NET `DateTime` class to handle dates. The `DateTime` class will only allow valid dates, ensuring that the values in your database are also valid. The disadvantage of this is that it is not useful in a mixed environment where .NET and non .NET code are used to manipulate the database, as each application must perform its own date validation.

Users of MySQL 5.0.2 and higher can use the new `traditional` SQL mode to restrict invalid date values. For information on using the `traditional` SQL mode, see [Section 5.1.7, "Server SQL Modes"](#).

#### 21.2.5.12.4. Handling Invalid Dates

Although it is strongly recommended that you avoid the use of invalid dates within your .NET application, it is possible to use invalid dates by means of the `MySqlDateTime` data type.

The `MySqlDateTime` data type supports the same date values that are supported by the MySQL server. The default behavior of Connector/Net is to return a .NET `DateTime` object for valid date values, and return an error for invalid dates. This default can be modified to cause Connector/Net to return `MySqlDateTime` objects for invalid dates.

To instruct Connector/Net to return a `MySqlDateTime` object for invalid dates, add the following line to your connection string:

```
Allow Zero Datetime=True
```

Please note that the use of the `MySqlDateTime` class can still be problematic. The following are some known issues:

1. Data binding for invalid dates can still cause errors (zero dates like 0000-00-00 do not seem to have this problem).
2. The `ToString` method return a date formatted in the standard MySQL format (for example, `2005-02-23 08:50:25`). This differs from the `ToString` behavior of the .NET `DateTime` class.
3. The `MySqlDateTime` class supports NULL dates, while the .NET `DateTime` class does not. This can cause errors when trying to convert a `MySQLDateTime` to a `DateTime` if you do not check for NULL first.

Because of the known issues, the best recommendation is still to use only valid dates in your application.

#### 21.2.5.12.5. Handling NULL Dates

The .NET `DateTime` data type cannot handle `NULL` values. As such, when assigning values from a query to a `DateTime` variable, you must first check whether the value is in fact `NULL`.

When using a `MySqlDataReader`, use the `.IsDBNull` method to check whether a value is `NULL` before making the assignment:

#### Visual Basic Example

```
If Not myReader.IsDBNull(myReader.GetOrdinal("mytime")) Then
    myTime = myReader.GetDateTime(myReader.GetOrdinal("mytime"))
Else
    myTime = DateTime.MinValue
End If
```

#### C# Example

```
if (! myReader.IsDBNull(myReader.GetOrdinal("mytime")))
    myTime = myReader.GetDateTime(myReader.GetOrdinal("mytime"));
else
    myTime = DateTime.MinValue;
```

`NULL` values will work in a data set and can be bound to form controls without special handling.

#### 21.2.5.13. Using the `MySqlBulkLoader` Class

MySQL Connector/Net features a bulk loader class that wraps the MySQL statement `LOAD DATA INFILE`. This gives MySQL Connector/Net the ability to load a data file from a local or remote host

to the server. The class concerned is [MySQLBulkLoader](#). This class has various methods, the main one being `load` to cause the specified file to be loaded to the server. Various parameters can be set to control how the data file is processed. This is achieved through setting various properties of the class. For example, the field separator used, such as comma or tab, can be specified, along with the record terminator, such as newline.

The following code shows a simple example of using the [MySQLBulkLoader](#) class. First an empty table needs to be created, in this case in the `test` database:

```
CREATE TABLE Career (
    Name VARCHAR(100) NOT NULL,
    Age INTEGER,
    Profession VARCHAR(200)
);
```

A simple tab-delimited data file is also created (it could use any other field delimiter such as comma):

```
Table Career in Test Database
Name Age Profession

Tony 47 Technical Writer
Ana 43 Nurse
Fred 21 IT Specialist
Simon 45 Hairy Biker
```

Note that with this test file the first three lines will need to be ignored, as they do not contain table data. This can be achieved using the [NumberOfLinesToSkip](#) property. This file can then be loaded and used to populate the `Career` table in the `test` database:

```
using System;
using System.Text;
using MySql.Data;
using MySql.Data.MySqlClient;

namespace ConsoleApplication1
{
    class Program
    {
        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=test;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);

            MySQLBulkLoader bl = new MySQLBulkLoader(conn);
            bl.TableName = "Career";
            bl.FieldTerminator = "\t";
            bl.LineTerminator = "\n";
            bl.FileName = "c:/career_data.txt";
            bl.NumberOfLinesToSkip = 3;

            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                // Upload data from file
                int count = bl.Load();
                Console.WriteLine(count + " lines uploaded.");

                string sql = "SELECT Name, Age, Profession FROM Career";
                MySqlCommand cmd = new MySqlCommand(sql, conn);
                MySqlDataReader rdr = cmd.ExecuteReader();

                while (rdr.Read())
                {
                    Console.WriteLine(rdr[0] + " -- " + rdr[1] + " -- " + rdr[2]);
                }

                rdr.Close();
            }
        }
    }
}
```

```

        conn.Close();
    }
    catch (Exception ex)
    {
        Console.WriteLine(ex.ToString());
    }
    Console.WriteLine("Done.");
}
}
}

```

Further information on `LOAD DATA INFILE` can be found in [Section 13.2.6, "LOAD DATA INFILE Syntax"](#). Further information on `MySQLBulkLoader` can be found in the reference documentation that was included with your connector.

#### 21.2.5.14. Using the MySQL Connector/Net Trace Source Object

MySQL Connector/Net 6.2 introduced support for .NET 2.0 compatible tracing, using `TraceSource` objects.

The .NET 2.0 tracing architecture consists of four main parts:

- *Source* - This is the originator of the trace information. The source is used to send trace messages. The name of the source provided by MySQL Connector/Net is `mysql`.
- *Switch* - This defines the level of trace information to emit. Typically, this is specified in the `app.config` file, so that it is not necessary to recompile an application to change the trace level.
- *Listener* - Trace listeners define where the trace information will be written to. Supported listeners include, for example, the Visual Studio Output window, the Windows Event Log, and the console.
- *Filter* - Filters can be attached to listeners. Filters determine the level of trace information that will be written. While a switch defines the level of information that will be written to all listeners, a filter can be applied on a per-listener basis, giving finer grained control of trace information.

To use tracing a `TraceSource` object first needs to be created. To create a `TraceSource` object in MySQL Connector/Net you would use code similar to the following:

```
TraceSource ts = new TraceSource("mysql");
```

To enable trace messages, configure a trace switch. There are three main switch classes, `BooleanSwitch`, `SourceSwitch`, and `TraceSwitch`. Trace switches also have associated with them a trace level enumeration, these are `Off`, `Error`, `Warning`, `Info`, and `Verbose`. The following code snippet illustrates creating a switch:

```
ts.Switch = new SourceSwitch("MySwitch", "Verbose");
```

This creates a `SourceSwitch`, called `MySwitch`, and sets the trace level to `Verbose`, meaning that all trace messages will be written.

It is convenient to be able to change the trace level without having to recompile the code. This is achieved by specifying the trace level in application configuration file, `app.config`. You then simply need to specify the desired trace level in the configuration file and restart the application. The trace source is configured within the `system.diagnostics` section of the file. The following XML snippet illustrates this:

```

<configuration>
  ...
  <system.diagnostics>
    <sources>
      <source name="mysql" switchName="MySwitch"
        switchType="System.Diagnostics.SourceSwitch" />
    </sources>
  </system.diagnostics>
</configuration>

```

```

...
</sources>
<switches>
  <add name="MySwitch" value="Verbose"/>
  ...
</switches>
</system.diagnostics>
...
</configuration>

```

By default, trace information is written to the Output window of Microsoft Visual Studio. There are a wide range of listeners that can be attached to the trace source, so that trace messages can be written out to various destinations. You can also create custom listeners to allow trace messages to be written to other destinations as mobile devices and web services. A commonly used example of a listener is [ConsoleTraceListener](#), which writes trace messages to the console.

To add a listener at runtime, use code such as the following:

```
ts.Listeners.Add(new ConsoleTraceListener());
```

Then, call methods on the trace source object to generate trace information. For example, the [TraceInformation\(\)](#), [TraceEvent\(\)](#), or [TraceData\(\)](#) methods can be used.

The [TraceInformation\(\)](#) method simply prints a string passed as a parameter. The [TraceEvent\(\)](#) method, as well as the optional informational string, requires a [TraceEventType](#) value to be passed to indicate the trace message type, and also an application specific ID. The [TraceEventType](#) can have a value of [Verbose](#), [Information](#), [Warning](#), [Error](#), and [Critical](#). Using the [TraceData\(\)](#) method you can pass any object, for example an exception object, instead of a message.

To ensure that these generated trace messages gets flushed from the trace source buffers to listeners, invoke the [Flush\(\)](#) method. When you are finished using a trace source, call the [Close\(\)](#) method. The [Close\(\)](#) method first calls [Flush\(\)](#), to ensure any remaining data is written out. It then frees up resources, and closes the listeners associated with the trace source.

```

ts.TraceInformation("Informational message");
ts.TraceEvent(TraceEventType.Error, 3, "Optional error message");
ts.TraceData(TraceEventType.Error, 3, ex); // pass exception object
ts.Flush();
...
ts.Close();

```

#### 21.2.5.14.1. Viewing MySQL Trace Information

This section describes how to set up your application to view MySQL trace information.

The first thing you need to do is create a suitable `app.config` file for your application. An example is shown in the following code:

```

<?xml version="1.0" encoding="utf-8" ?>
<configuration>
  <system.diagnostics>
    <sources>
      <source name="mysql" switchName="SourceSwitch"
        switchType="System.Diagnostics.SourceSwitch" >
        <listeners>
          <add name="console" />
          <remove name="Default" />
        </listeners>
      </source>
    </sources>
    <switches>
      <!-- You can set the level at which tracing is to occur -->
      <add name="SourceSwitch" value="Verbose" />
      <!-- You can turn tracing off -->

```

```

    <!--add name="SourceSwitch" value="Off" -->
  </switches>
  <sharedListeners>
    <add name="console"
      type="System.Diagnostics.ConsoleTraceListener"
      initializeData="false"/>
  </sharedListeners>
</system.diagnostics>
</configuration>

```

This ensures a suitable trace source is created, along with a switch. The switch level in this case is set to `Verbose` to display the maximum amount of information.

In the application the only other step required is to add `logging=true` to the connection string. An example application could be:

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Diagnostics;
using MySql.Data;
using MySql.Data.MySqlClient;
using MySql.Web;

namespace ConsoleApplication1
{
    class Program
    {
        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=world;port=3306;password=*****;logging=true";
            MySqlConnection conn = new MySqlConnection(connStr);
            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                string sql = "SELECT Name, HeadOfState FROM Country WHERE Continent='Oceania'";
                MySqlCommand cmd = new MySqlCommand(sql, conn);
                MySqlDataReader rdr = cmd.ExecuteReader();

                while (rdr.Read())
                {
                    Console.WriteLine(rdr[0] + " -- " + rdr[1]);
                }

                rdr.Close();

                conn.Close();
            }
            catch (Exception ex)
            {
                Console.WriteLine(ex.ToString());
            }
            Console.WriteLine("Done.");
        }
    }
}

```

This simple application will then generate the following output:

```

Connecting to MySQL...
mysql Information: 1 : 1: Connection Opened: connection string = 'server=localhost;User Id=root;database=world;password=*****;logging=True'
mysql Information: 3 : 1: Query Opened: SHOW VARIABLES
mysql Information: 4 : 1: Resultset Opened: field(s) = 2, affected rows = -1, inserted id = -1
mysql Information: 5 : 1: Resultset Closed. Total rows=272, skipped rows=0, size (bytes)=7058
mysql Information: 6 : 1: Query Closed

```

```
mysql Information: 3 : 1: Query Opened: SHOW COLLATION
mysql Information: 4 : 1: Resultset Opened: field(s) = 6, affected rows = -1, inserted id = -1
mysql Information: 5 : 1: Resultset Closed. Total rows=127, skipped rows=0, size (bytes)=4102
mysql Information: 6 : 1: Query Closed
mysql Information: 3 : 1: Query Opened: SET character_set_results=NULL
mysql Information: 4 : 1: Resultset Opened: field(s) = 0, affected rows = 0, inserted id = 0
mysql Information: 5 : 1: Resultset Closed. Total rows=0, skipped rows=0, size (bytes)=0
mysql Information: 6 : 1: Query Closed
mysql Information: 10 : 1: Set Database: world
mysql Information: 3 : 1: Query Opened: SELECT Name, HeadOfState FROM Country WHERE Continent='Oceania'
mysql Information: 4 : 1: Resultset Opened: field(s) = 2, affected rows = -1, inserted id = -1
American Samoa -- George W. Bush
Australia -- Elisabeth II
...
Wallis and Futuna -- Jacques Chirac
Vanuatu -- John Bani
United States Minor Outlying Islands -- George W. Bush
mysql Information: 5 : 1: Resultset Closed. Total rows=28, skipped rows=0, size (bytes)=788
mysql Information: 6 : 1: Query Closed
Done.
mysql Information: 2 : 1: Connection Closed
```

The first number displayed in the trace message corresponds to the MySQL event type:

Event	Description
1	ConnectionOpened: connection string
2	ConnectionClosed:
3	QueryOpened: mysql server thread id, query text
4	ResultOpened: field count, affected rows (-1 if select), inserted id (-1 if select)
5	ResultClosed: total rows read, rows skipped, size of resultset in bytes
6	QueryClosed:
7	StatementPrepared: prepared sql, statement id
8	StatementExecuted: statement id, mysql server thread id
9	StatementClosed: statement id
10	NonQuery: [varies]
11	UsageAdvisorWarning: usage advisor flag. NoIndex = 1, BadIndex = 2, SkippedRows = 3, SkippedColumns = 4, FieldConversion = 5.
12	Warning: level, code, message
13	Error: error number, error message

The second number displayed in the trace message is the connection count.

Although this example uses the `ConsoleTraceListener`, any of the other standard listeners could have been used. Another possibility is to create a custom listener that uses the information passed using the `TraceEvent` method. For example, a custom trace listener could be created to perform active monitoring of the MySQL event messages, rather than simply writing these to an output device.

It is also possible to add listeners to the MySQL Trace Source at runtime. This can be done with the following code:

```
MySqlTrace.Listeners.Add(new ConsoleTraceListener());
```

MySQL Connector/Net 6.3.2 introduced the ability to switch tracing on and off at runtime. This can be achieved using the calls `MySqlTrace.EnableQueryAnalyzer(string host, int postInterval)` and `MySqlTrace.DisableQueryAnalyzer()`. The parameter `host` is the URL of

the MySQL Enterprise Monitor server to monitor. The parameter `postInterval` is how often to post the data to MySQL Enterprise Monitor, in seconds.

#### 21.2.5.14.2. Building Custom Listeners

To build custom listeners that work with the MySQL Connector/Net Trace Source, it is necessary to understand the key methods used, and the event data formats used.

The main method involved in passing trace messages is the `TraceSource.TraceEvent` method. This has the prototype:

```
public void TraceEvent(  
    TraceEventType eventType,  
    int id,  
    string format,  
    params Object[] args  
)
```

This trace source method will process the list of attached listeners and call the listener's `TraceListener.TraceEvent` method. The prototype for the `TraceListener.TraceEvent` method is as follows:

```
public virtual void TraceEvent(  
    TraceEventCache eventCache,  
    string source,  
    TraceEventType eventType,  
    int id,  
    string format,  
    params Object[] args  
)
```

The first three parameters are used in the standard as [defined by Microsoft](#). The last three parameters contain MySQL-specific trace information. Each of these parameters is now discussed in more detail.

##### `int id`

This is a MySQL-specific identifier. It identifies the MySQL event type that has occurred, resulting in a trace message being generated. This value is defined by the `MySQLTraceEventType` public enum contained in the MySQL Connector/Net code:

```
public enum MySQLTraceEventType : int  
{  
    ConnectionOpened = 1,  
    ConnectionClosed,  
    QueryOpened,  
    ResultOpened,  
    ResultClosed,  
    QueryClosed,  
    StatementPrepared,  
    StatementExecuted,  
    StatementClosed,  
    NonQuery,  
    UsageAdvisorWarning,  
    Warning,  
    Error  
}
```

The MySQL event type also determines the contents passed using the parameter `params Object[] args`. The nature of the `args` parameters are described in further detail in the following material.

##### `string format`

This is the format string that contains zero or more format items, which correspond to objects in the `args` array. This would be used by a listener such as `ConsoleTraceListener` to write a message to the output device.



`params Object[] args`

This is a list of objects that depends on the MySQL event type, `id`. However, the first parameter passed using this list is always the driver id. The driver id is a unique number that is incremented each time the connector is opened. This enables groups of queries on the same connection to be identified. The parameters that follow driver id depend on the MySQL event id, and are as follows:

MySQL-specific event type	Arguments (params Object[] args)
ConnectionOpened	Connection string
ConnectionClosed	No additional parameters
QueryOpened	mysql server thread id, query text
ResultOpened	field count, affected rows (-1 if select), inserted id (-1 if select)
ResultClosed	total rows read, rows skipped, size of resultset in bytes
QueryClosed	No additional parameters
StatementPrepared	prepared sql, statement id
StatementExecuted	statement id, mysql server thread id
StatementClosed	statement id
NonQuery	Varies
UsageAdvisorWarning	usage advisor flag. NoIndex = 1, BadIndex = 2, SkippedRows = 3, SkippedColumns = 4, FieldConversion = 5.
Warning	level, code, message
Error	error number, error message

This information will allow you to create custom trace listeners that can actively monitor the MySQL-specific events.

### 21.2.5.15. Binary/Nonbinary Issues

There are certain situations where MySQL will return incorrect metadata about one or more columns. More specifically, the server will sometimes report that a column is binary when it is not and vice versa. In these situations, it becomes practically impossible for the connector to be able to correctly identify the correct metadata.

Some examples of situations that may return incorrect metadata are:

- Execution of `SHOW PROCESSLIST`. Some of the columns will be returned as binary even though they only hold string data.
- When a temporary table is used to process a resultset, some columns may be returned with incorrect binary flags.
- Some server functions such `DATE_FORMAT` will incorrectly return the column as binary.

With the availability of `BINARY` and `VARBINARY` data types, it is important that we respect the metadata returned by the server. However, we are aware that some existing applications may break with this change, so we are creating a connection string option to enable or disable it. By default, Connector/Net 5.1 respects the binary flags returned by the server. You might need to make small changes to your application to accommodate this change.

In the event that the changes required to your application would be too large, adding `'respect binary flags=false'` to your connection string causes the connector to use the prior behavior: any column that is marked as string, regardless of binary flags, will be returned as string. Only columns that are specifically marked as a `BLOB` will be returned as `BLOB`.

### 21.2.5.16. Character Set Considerations for Connector/Net

## Treating Binary Blobs As UTF8

MySQL doesn't currently support 4-byte UTF8 sequences. This makes it difficult to represent some multi-byte languages such as Japanese. To try and alleviate this, Connector/Net now supports a mode where binary blobs can be treated as strings.

To do this, you set the '[Treat Blobs As UTF8](#)' connection string keyword to [yes](#). This is all that needs to be done to enable conversion of all binary blobs to UTF8 strings. To convert only some of your BLOB columns, you can make use of the '[BlobAsUTF8IncludePattern](#)' and '[BlobAsUTF8ExcludePattern](#)' keywords. Set these to a regular expression pattern that matches the column names to include or exclude respectively.

When the regular expression patterns both match a single column, the include pattern is applied before the exclude pattern. The result, in this case, would be that the column would be excluded. Also, be aware that this mode does not apply to columns of type [BINARY](#) or [VARBINARY](#) and also do not apply to nonbinary [BLOB](#) columns.

Currently, this mode only applies to reading strings out of MySQL. To insert 4-byte UTF8 strings into blob columns, use the .NET [Encoding.GetBytes](#) function to convert your string to a series of bytes. You can then set this byte array as a parameter for a [BLOB](#) column.

## 21.2.5.17. Using Connector/Net with Crystal Reports

Crystal Reports is a common tool used by Windows application developers to perform reporting and document generation. In this section we will show how to use Crystal Reports XI with MySQL and Connector/Net.

### 21.2.5.17.1. Creating a Data Source

When creating a report in Crystal Reports there are two options for accessing the MySQL data while designing your report.

The first option is to use Connector/ODBC as an ADO data source when designing your report. You will be able to browse your database and choose tables and fields using drag and drop to build your report. The disadvantage of this approach is that additional work must be performed within your application to produce a data set that matches the one expected by your report.

The second option is to create a data set in VB.NET and save it as XML. This XML file can then be used to design a report. This works quite well when displaying the report in your application, but is less versatile at design time because you must choose all relevant columns when creating the data set. If you forget a column you must re-create the data set before the column can be added to the report.

The following code can be used to create a data set from a query and write it to disk:

### Visual Basic Example

```
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=world"

Try
    conn.Open()
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " _
        & "country.name, country.population, country.continent " _
        & "FROM country, city ORDER BY country.continent, country.name"
    cmd.Connection = conn
```

```

myAdapter.SelectCommand = cmd
myAdapter.Fill(myData)

myData.WriteXml("C:\dataset.xml", XmlWriteMode.WriteSchema)
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try

```

## C# Example

```

DataSet myData = new DataSet();
MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;
MySQL.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();
myAdapter = new MySQL.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " +
        "country.name, country.population, country.continent " +
        "FROM country, city ORDER BY country.continent, country.name";
    cmd.Connection = conn;

    myAdapter.SelectCommand = cmd;
    myAdapter.Fill(myData);

    myData.WriteXml(@"C:\dataset.xml", XmlWriteMode.WriteSchema);
}
catch (MySQL.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

The resulting XML file can be used as an ADO.NET XML datasource when designing your report.

If you choose to design your reports using Connector/ODBC, it can be downloaded from [dev.mysql.com](http://dev.mysql.com).

### 21.2.5.17.2. Creating the Report

For most purposes, the Standard Report wizard helps with the initial creation of a report. To start the wizard, open Crystal Reports and choose the **New > Standard Report** option from the File menu.

The wizard first prompts you for a data source. If you use Connector/ODBC as your data source, use the OLEDB provider for ODBC option from the OLE DB (ADO) tree instead of the ODBC (RDO) tree when choosing a data source. If using a saved data set, choose the ADO.NET (XML) option and browse to your saved data set.

The remainder of the report creation process is done automatically by the wizard.

After the report is created, choose the Report Options... entry of the File menu. Un-check the Save Data With Report option. This prevents saved data from interfering with the loading of data within our application.

### 21.2.5.17.3. Displaying the Report

To display a report we first populate a data set with the data needed for the report, then load the report and bind it to the data set. Finally we pass the report to the crViewer control for display to the user.

The following references are needed in a project that displays a report:

- CrystalDecisions.CrystalReports.Engine
- CrystalDecisions.ReportSource
- CrystalDecisions.Shared
- CrystalDecisions.Windows.Forms

The following code assumes that you created your report using a data set saved using the code shown in [Section 21.2.5.17.1, "Creating a Data Source"](#), and have a crViewer control on your form named `myViewer`.

### Visual Basic Example

```
Imports CrystalDecisions.CrystalReports.Engine
Imports System.Data
Imports MySql.Data.MySqlClient

Dim myReport As New ReportDocument
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = _
    "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

Try
    conn.Open()

    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " _
        & "country.name, country.population, country.continent " _
        & "FROM country, city ORDER BY country.continent, country.name"
    cmd.Connection = conn

    myAdapter.SelectCommand = cmd
    myAdapter.Fill(myData)

    myReport.Load(".\world_report.rpt")
    myReport.SetDataSource(myData)
    myViewer.ReportSource = myReport
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try
```

### C# Example

```
using CrystalDecisions.CrystalReports.Engine;
using System.Data;
using MySql.Data.MySqlClient;

ReportDocument myReport = new ReportDocument();
DataSet myData = new DataSet();
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;
MySql.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();
myAdapter = new MySql.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " +
        "country.name, country.population, country.continent " +
```

```

        "FROM country, city ORDER BY country.continent, country.name";
cmd.Connection = conn;

myAdapter.SelectCommand = cmd;
myAdapter.Fill(myData);

myReport.Load(@".\world_report.rpt");
myReport.SetDataSource(myData);
myViewer.ReportSource = myReport;
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

A new data set is generated using the same query used to generate the previously saved data set. Once the data set is filled, a ReportDocument is used to load the report file and bind it to the data set. The ReportDocument is then passed as the ReportSource of the crViewer.

This same approach is taken when a report is created from a single table using Connector/ODBC. The data set replaces the table used in the report and the report is displayed properly.

When a report is created from multiple tables using Connector/ODBC, a data set with multiple tables must be created in our application. This enables each table in the report data source to be replaced with a report in the data set.

We populate a data set with multiple tables by providing multiple [SELECT](#) statements in our MySqlCommand object. These [SELECT](#) statements are based on the SQL query shown in Crystal Reports in the Database menu's Show SQL Query option. Assume the following query:

```

SELECT `country`.`Name`, `country`.`Continent`, `country`.`Population`, `city`.`Name`, `city`.`Populati
FROM `world`.`country` `country` LEFT OUTER JOIN `world`.`city` `city` ON `country`.`Code`=`city`.`Coun
ORDER BY `country`.`Continent`, `country`.`Name`, `city`.`Name`

```

This query is converted to two [SELECT](#) queries and displayed with the following code:

### Visual Basic Example

```

Imports CrystalDecisions.CrystalReports.Engine
Imports System.Data
Imports MySql.Data.MySqlClient

Dim myReport As New ReportDocument
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=world"

Try
    conn.Open()
    cmd.CommandText = "SELECT name, population, countrycode FROM city ORDER BY countrycode, name;" _
        & "SELECT name, population, code, continent FROM country ORDER BY continent, name"
    cmd.Connection = conn

    myAdapter.SelectCommand = cmd
    myAdapter.Fill(myData)

    myReport.Load(@".\world_report.rpt")
    myReport.Database.Tables(0).SetDataSource(myData.Tables(0))
    myReport.Database.Tables(1).SetDataSource(myData.Tables(1))
    myViewer.ReportSource = myReport
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Err

```

```
End Try
```

## C# Example

```
using CrystalDecisions.CrystalReports.Engine;
using System.Data;
using MySql.Data.MySqlClient;

ReportDocument myReport = new ReportDocument();
DataSet myData = new DataSet();
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;
MySql.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();
myAdapter = new MySql.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT name, population, countrycode FROM city ORDER " +
        "BY countrycode, name; SELECT name, population, code, continent FROM " +
        "country ORDER BY continent, name";
    cmd.Connection = conn;

    myAdapter.SelectCommand = cmd;
    myAdapter.Fill(myData);

    myReport.Load(@".\world_report.rpt");
    myReport.Database.Tables(0).SetDataSource(myData.Tables(0));
    myReport.Database.Tables(1).SetDataSource(myData.Tables(1));
    myViewer.ReportSource = myReport;
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}
```

It is important to order the [SELECT](#) queries in alphabetic order, as this is the order the report will expect its source tables to be in. One `SetDataSource` statement is needed for each table in the report.

This approach can cause performance problems because Crystal Reports must bind the tables together on the client-side, which will be slower than using a pre-saved data set.

### 21.2.5.18. ASP.NET Provider Model

MySQL Connector/Net provides support for the ASP.NET 2.0 provider model. This model enables application developers to focus on the business logic of their application instead of having to recreate such boilerplate items as membership and roles support.

MySQL Connector/Net supplies the following providers:

- Membership Provider
- Role Provider
- Profile Provider
- Session State Provider (MySQL Connector/Net 6.1 and later)

The following tables show the supported providers, their default provider and the corresponding MySQL provider.

#### Membership Provider

Default Provider	MySQL Provider
System.Web.Security.SqlMembershipProvider	MySql.Web.Security.MySQLMembershipProvider

### Role Provider

Default Provider	MySQL Provider
System.Web.Security.SqlRoleProvider	MySql.Web.Security.MySQLRoleProvider

### Profile Provider

Default Provider	MySQL Provider
System.Web.Profile.SqlProfileProvider	MySql.Web.Profile.MySQLProfileProvider

### SessionState Provider

Default Provider	MySQL Provider
System.Web.SessionState.InProcSessionStateStore	MySql.Web.SessionState.MySqlSessionStateStore



#### Note

The MySQL Session State provider uses slightly different capitalization on the class name compared to the other MySQL providers.

### Installing The Providers

The installation of Connector/Net 5.1 or later will install the providers and register them in your machine's .NET configuration file, `machine.config`. The additional entries created will result in the `system.web` section appearing similar to the following code:

```
<system.web>
  <processModel autoConfig="true" />
  <httpHandlers />
  <membership>
    <providers>
      <add name="AspNetSqlMembershipProvider" type="System.Web.Security.SqlMembershipProvider, System.W
      <add name="MySQLMembershipProvider" type="MySql.Web.Security.MySQLMembershipProvider, MySql.Web,
    </providers>
  </membership>
  <profile>
    <providers>
      <add name="AspNetSqlProfileProvider" connectionStringName="LocalSqlServer" applicationName="/" ty
      <add name="MySQLProfileProvider" type="MySql.Web.Profile.MySQLProfileProvider, MySql.Web, Version
    </providers>
  </profile>
  <roleManager>
    <providers>
      <add name="AspNetSqlRoleProvider" connectionStringName="LocalSqlServer" applicationName="/" type=
      <add name="AspNetWindowsTokenRoleProvider" applicationName="/" type="System.Web.Security.WindowsT
      <add name="MySQLRoleProvider" type="MySql.Web.Security.MySQLRoleProvider, MySql.Web, Version=6.1.
    </providers>
  </roleManager>
</system.web>
```

Each provider type can have multiple provider implementations. The default provider can also be set here using the `defaultProvider` attribute, but usually this is set in the `web.config` file either manually or by using the ASP.NET configuration tool.

At time of writing, the `MySqlSessionStateStore` is not added to `machine.config` at install time, and so add the following:

```
<sessionState>
```

```
<providers>
  <add name="MySQLSessionStateStore" type="MySQL.Web.SessionState.MySQLSessionStateStore, MySQL.Web, Ver
</providers>
</sessionState>
```

The SessionState Provider uses the `customProvider` attribute, rather than `defaultProvider`, to set the provider as the default. A typical `web.config` file might contain:

```
<system.web>
  <membership defaultProvider="MySQLMembershipProvider" />
  <roleManager defaultProvider="MySQLRoleProvider" />
  <profile defaultProvider="MySQLProfileProvider" />
  <sessionState customProvider="MySQLSessionStateStore" />
  <compilation debug="false">
    ...
```

This sets the MySQL Providers as the defaults to be used in this web application.

The providers are implemented in the file `mysql.web.dll` and this file can be found in your MySQL Connector/Net installation folder. There is no need to run any type of SQL script to set up the database schema, as the providers create and maintain the proper schema automatically.

### Using The Providers

The easiest way to start using the providers is to use the ASP.NET configuration tool that is available on the Solution Explorer toolbar when you have a website project loaded.

In the web pages that open, you can select the MySQL membership and roles providers by picking a custom provider for each area.

When the provider is installed, it creates a dummy connection string named `LocalMySQLServer`. Although this has to be done so that the provider will work in the ASP.NET configuration tool, you override this connection string in your `web.config` file. You do this by first removing the dummy connection string and then adding in the proper one, as shown in the following example:

```
<connectionStrings>
  <remove name="LocalMySQLServer" />
  <add name="LocalMySQLServer" connectionString="server=xxx;uid=xxx;pwd=xxx;database=xxx;" />
</connectionStrings>
```

Note the database to connect to must be specified.

Rather than manually editing configuration files, consider using the MySQL Website Configuration tool to configure your desired provider setup. From MySQL Connector/Net 6.1.1 onwards, all providers can be selected and configured from this wizard. The tool modifies your `website.config` file to the desired configuration. A tutorial on doing this is available in the following section [Section 21.2.3.12, "MySQL Website Configuration Tool"](#).

A tutorial demonstrating how to use the Membership and Role Providers can be found in the following section [Section 21.2.4.2, "Tutorial: MySQL Connector/Net ASP.NET Membership and Role Provider"](#).

### Deployment

To use the providers on a production server, distribute the `MySQL.Data` and the `MySQL.Web` assemblies, and either register them in the remote systems Global Assembly Cache or keep them in your application's `bin/` directory.

#### 21.2.5.19. Working with Partial Trust / Medium Trust

.NET applications operate under a given trust level. Normal desktop applications operate under full trust, while web applications that are hosted in shared environments are normally run under the partial trust level (also known as "medium trust"). Some hosting providers host shared applications in their own app pools and allow the application to run under full trust, but this configuration is relatively rare.



The Connector/Net support for partial trust has improved over time to simplify the configuration and deployment process for hosting providers.

### 21.2.5.19.1. Evolution of Partial Trust Support Across Connector/Net Versions

The partial trust support for Connector/Net has improved rapidly throughout the 6.5.x and 6.6.x versions. The latest enhancements do require some configuration changes in existing deployments. Here is a summary of the changes for each version.

#### 6.6.4 and Above: Library Can Be Inside or Outside GAC

Now you can install the `MySql.Data.dll` library in the Global Assembly Cache (GAC) as explained in [Section 21.2.5.19.2, “Configuring Partial Trust with Connector/Net Library Installed in GAC”](#), or in a `bin` or `lib` folder inside the project or solution as explained in [Section 21.2.5.19.3, “Configuring Partial Trust with Connector/Net Library Not Installed in GAC”](#). If the library is not in the GAC, the only protocol supported is TCP/IP.

#### 6.5.1 and Above: Partial Trust Requires Library in the GAC

Connector/Net 6.5 fully enables our provider to run in a partial trust environment when the library is installed in the Global Assembly Cache (GAC). The new `MySQLClientPermission` class, derived from the `.NET DBDataPermission` class, helps to simplify the permission setup.

#### 5.0.8 / 5.1.3 and Above: Partial Trust Requires Socket Permissions

Starting with these versions, Connector/Net can be used under partial trust hosting that has been modified to allow the use of sockets for communication. By default, partial trust does not include `SocketPermission`. Connector/Net uses sockets to talk with the MySQL server, so the hosting provider must create a new trust level that is an exact clone of partial trust but that has the following permissions added:

- `System.Net.SocketPermission`
- `System.Security.Permissions.ReflectionPermission`
- `System.Net.DnsPermission`
- `System.Security.Permissions.SecurityPermission`

#### Prior to 5.0.8 / 5.1.3: Partial Trust Not Supported

Connector/Net versions prior to 5.0.8 and 5.1.3 were not compatible with partial trust hosting.

### 21.2.5.19.2. Configuring Partial Trust with Connector/Net Library Installed in GAC

If the library is installed in the GAC, you must include the connection option `includesecurityasserts=true` in your connection string. This is a new requirement as of Connector/Net 6.6.4.

The following list shows steps and code fragments needed to run a Connector/Net application in a partial trust environment. For illustration purposes, we use the Pipe Connections protocol in this example.

1. Install Connector/Net: version 6.6.1 or higher, or 6.5.4 or higher.
2. After installing the library, make the following configuration changes:

In the `SecurityClasses` section, add a definition for the `MySQLClientPermission` class, including the version to use.

```
<configuration>
  <mscorlib>
    <security>
```

```
<policy>
  <PolicyLevel version="1">
    <SecurityClasses>
      ....
    <SecurityClass Name="MySQLClientPermission" Description="MySQL.Data.MySQLClient.MySQLClientPermission" />
  </PolicyLevel>
</policy>
```

Scroll down to the [ASP.Net](#) section:

```
<PermissionSet class="NamedPermissionSet" version="1" Name="ASP.Net">
```

Add a new entry for the detailed configuration of the [MySQLClientPermission](#) class:

```
<IPermission class="MySQLClientPermission" version="1" Unrestricted="true"/>
```

Note: This configuration is the most generalized way that includes all keywords.

- Configure the MySQL server to accept pipe connections, by adding the `--enable-named-pipe` option on the command line. If you need more information about this, see [Section 2.3, "Installing MySQL on Microsoft Windows"](#).
- Confirm that the hosting provider has installed the Connector/Net library (`MySQL.Data.dll`) in the GAC.
- Optionally, the hosting provider can avoid granting permissions globally by using the new [MySQLClientPermission](#) class in the trust policies. (The alternative is to globally enable the permissions [System.Net.SocketPermission](#), [System.Security.Permissions.ReflectionPermission](#), [System.Net.DnsPermission](#), and [System.Security.Permissions.SecurityPermission](#).)
- Create a simple web application using Visual Studio 2010.
- Add the reference in your application for the [MySQL.Data.MySQLClient](#) library.
- Edit your `web.config` file so that your application runs using a Medium trust level:

```
<system.web>
  <trust level="Medium"/>
</system.web>
```

- Add the [MySQL.Data.MySQLClient](#) namespace to your server-code page.
- Define the connection string, in slightly different ways depending on the Connector/Net version.

**Only for 6.6.4 or later:** To use the connections inside any web application that will run in Medium trust, add the new `includesecurityasserts` option to the connection string. `includesecurityasserts=true` that makes the library request the following permissions when required: [SocketPermissions](#), [ReflectionPermissions](#), [DnsPermissions](#), [SecurityPermissions](#) among others that are not granted in Medium trust levels.

**For Connector/Net 6.6.3 or earlier:** No special setting for security is needed within the connection string.

```
MySQLConnectionStringBuilder myconnString = new MySQLConnectionStringBuilder("server=localhost;User Id=
myconnString.PipeName = "MySQL55";
myconnString.ConnectionProtocol = MySQLConnectionProtocol.Pipe;
// Following attribute is a new requirement when the library is in the GAC.
// Could also be done by adding includesecurityasserts=true; to the string literal
// in the constructor above.
// Not needed with Connector/Net 6.6.3 and earlier.myconnString.IncludeSecurityAsserts = true;
```

- Define the [MySQLConnection](#) to use:

```
MySQLConnection myconn = new MySQLConnection(myconnString.ConnectionString);
```

```
myconn.Open();
```

- Retrieve some data from your tables:

```
MySQLCommand cmd = new MySQLCommand("Select * from products", myconn);
MySQLDataAdapter da = new MySQLDataAdapter(cmd);
DataSet1 tds = new DataSet1();
da.Fill(tds, tds.Tables[0].TableName);
GridView1.DataSource = tds;
GridView1.DataBind();
myconn.Close();
```

- Run the program. It should execute successfully, without requiring any special code or encountering any security problems.

### 21.2.5.19.3. Configuring Partial Trust with Connector/Net Library Not Installed in GAC

When deploying a web application to a Shared Hosted environment, where this environment is configured to run all their .NET applications under a partial or medium trust level, you might not be able to install the Connector/Net library in the GAC. Instead, you put a reference to the library in the [bin](#) or [lib](#) folder inside the project or solution. In this case, you configure the security in a different way than when the library is in the GAC.

Connector/Net is commonly used by applications that run in Windows environments where the default communication for the protocol is used via sockets or by TCP/IP. For this protocol to operate is necessary have the required socket permissions in the web configuration file as follows:

- Open the medium trust policy web configuration file, which should be under this folder:

```
%windir%\Microsoft.NET\Framework\{version}\CONFIG\web_mediumtrust.config
```

Use [Framework64](#) in the path instead of [Framework](#) if you are using a 64-bit installation of the framework.

- Locate the [SecurityClasses](#) tag:

```
<SecurityClass Name="SocketPermission"
Description="System.Net.SocketPermission, System, Version=4.0.0.0,
Culture=neutral, PublicKeyToken=b77a5c561934e089" />
```

- Scroll down and look for the following [PermissionSet](#):

```
<PermissionSet version="1" Name="ASP.Net">
```

- Add the following inside this [PermissionSet](#):

```
<IPermission class="SocketPermission" version="1" Unrestricted="true" />
```

This configuration lets you use the driver with the default Windows protocol TCP/IP without having any security issues. This approach only supports the TCP/IP protocol, so you cannot use any other type of connection.

Also, since the [MySQLClientPermissions](#) class is not added to the medium trust policy, you cannot use it. This configuration is the minimum required in order to work with Connector/Net without the GAC.

## 21.2.6. Connector/Net Connection String Options Reference

For usage information about connection strings, see [Section 21.2.5.2, "Creating a Connector/Net Connection String"](#). The first table list options that apply generally to all server configurations. The options related to systems using a connection pool are split into a separate table.

## General Options

**Table 21.20. Connector/Net Connection String Options - General**

Name	Default	Description
<code>Allow Batch, AllowBatch</code>	true	When true, multiple SQL statements can be sent with one command execution. Note: starting with MySQL 4.1.1, batch statements should be separated by the server-defined separator character. Statements sent to earlier versions of MySQL should be separated by ';'. 
<code>Allow User Variables, AllowUserVariables</code>	false	Setting this to <code>true</code> indicates that the provider expects user variables in the SQL. This option was added in Connector/Net version 5.2.2. 
<code>Allow Zero Datetime, AllowZeroDateTime</code>	false	If set to <code>True</code> , <code>MySqlDataReader.GetValue()</code> returns a <code>MySqlDateTime</code> object for date or datetime columns that have disallowed values, such as zero datetime values, and a <code>System.DateTime</code> object for valid values. If set to <code>False</code> (the default setting) it causes a <code>System.DateTime</code> object to be returned for all valid values and an exception to be thrown for disallowed values, such as zero datetime values. 
<code>Auto Enlist, AutoEnlist</code>	true	If <code>AutoEnlist</code> is set to <code>true</code> , which is the default, a connection opened using <code>TransactionScope</code> participates in this scope, it commits when the scope commits and rolls back if <code>TransactionScope</code> does not commit. However, this feature is considered security sensitive and therefore cannot be used in a medium trust environment. 
<code>BlobAsUTF8ExcludePattern</code>	null	A POSIX-style regular expression that matches the names of BLOB columns that do not contain UTF-8 character data. See <a href="#">Section 21.2.5.16, "Character Set Considerations for Connector/Net"</a> for usage details. 
<code>BlobAsUTF8IncludePattern</code>	null	A POSIX-style regular expression that matches the names of BLOB columns containing UTF-8 character data. See <a href="#">Section 21.2.5.16, "Character Set Considerations for Connector/Net"</a> for usage details. 
<code>Certificate File, CertificateFile</code>	null	This option specifies the path to a certificate file in PKCS #12 format ( <code>.pfx</code> ). For an example of usage, see <a href="#">Section 21.2.4.7, "Tutorial: Using SSL with MySQL Connector/Net"</a> . Was introduced with 6.2.1. 
<code>Certificate Password, CertificatePassword</code>	null	Specifies a password that is used in conjunction with a certificate specified using the option <code>CertificateFile</code> . For an example of usage, see <a href="#">Section 21.2.4.7, "Tutorial: Using SSL with MySQL Connector/Net"</a> . Was introduced with 6.2.1. 
<code>Certificate Store Location, CertificateStoreLocation</code>	null	Enables you to access a certificate held in a personal store, rather than use a certificate file and password combination. For an example of usage, see <a href="#">Section 21.2.4.7, "Tutorial: Using SSL with MySQL Connector/Net"</a> . Was introduced with 6.2.1. 
<code>Certificate Thumbprint, CertificateThumbprint</code>	null	Specifies a certificate thumbprint to ensure correct identification of a certificate contained within a personal store. For an example of usage, see <a href="#">Section 21.2.4.7, "Tutorial: Using SSL with MySQL Connector/Net"</a> . Was introduced with 6.2.1. 

Name	Default	Description
<code>CharSet</code> , <code>Character Set</code> , <code>CharacterSet</code>		Specifies the character set that should be used to encode all queries sent to the server. Resultsets are still returned in the character set of the result data.
<code>Check Parameters</code> , <code>CheckParameters</code>	true	Indicates if stored routine parameters should be checked against the server.
<code>Command Interceptors</code> , <code>CommandInterceptors</code>		The list of interceptors that can intercept SQL command operations.
<code>Connect Timeout</code> , <code>Connection Timeout</code> , <code>ConnectionTimeout</code>	15	The length of time (in seconds) to wait for a connection to the server before terminating the attempt and generating an error.
<code>Connect_Attrs</code>		Passes a predefined set of key/value pairs containing connection-specific data items to the server, for application-defined purposes. connector/Net automatically transmits the following attributes: <code>_pid</code> , <code>_client_version</code> , <code>_platform</code> , <code>_program_name</code> , <code>_os</code> , <code>_os_details</code> , and <code>_thread</code> . See <a href="#">Section 20.9.7, "Performance Schema Connection Attribute Tables"</a> for the <a href="#">Performance Schema</a> tables to query on the server side to make use of these attributes.
<code>Convert Zero Datetime</code> , <code>ConvertZeroDateTime</code>	false	True to have <code>MySqlDataReader.GetValue()</code> and <code>MySqlDataReader.GetDateTime()</code> return <code>DateTime.MinValue</code> for date or datetime columns that have disallowed values.
<code>Default Command Timeout</code> , <code>DefaultCommandTimeout</code>	30	Sets the default value of the command timeout to be used. This does not supercede the individual command timeout property on an individual command object. If you set the command timeout property, that will be used. This option was added in Connector/Net 5.1.4
<code>Default Table Cache Age</code> , <code>DefaultTableCacheAge</code>	60	Specifies how long a <code>TableDirect</code> result should be cached, in seconds. For usage information about table caching, see <a href="#">Section 21.2.5.7, "Using Connector/Net with Table Caching"</a> . This option was added in Connector/Net 6.4.
<code>enableSessionExpireCallback</code>	false	When set to <code>true</code> , causes the session-expiry scanner to raise the <code>session_end</code> event before deleting the session data stored in the <code>my_aspnet_sessions</code> table, when a session times out. Enable this option to write additional application-specific cleanup code to handle the <code>session_end</code> event of the <code>global.asax</code> class, before the stored data of the session gets deleted. Within the <code>session_end</code> method, any other required cleanup can be done. This option was added in Connector/Net 6.4.5; it is not yet available in Connector/Net 6.5.x releases.
<code>Encrypt</code> , <code>UseSSL</code>	false	For Connector/Net 5.0.3 and later, when <code>true</code> , SSL encryption is used for all data sent between the client and server if the server has a certificate installed. Recognized values are <code>true</code> , <code>false</code> , <code>yes</code> , and <code>no</code> . In versions before 5.0.3, this option had no effect. From version 6.2.1, this option is deprecated and is replaced by <code>SSL Mode</code> . The option still works if used. If this

Name	Default	Description
		option is set to true, it is equivalent to <code>SSL Mode = Preferred</code> .
<code>Exception Interceptors</code> , <code>ExceptionInterceptors</code>		The list of interceptors that can triage thrown <code>MySqlConnection</code> exceptions.
<code>Functions Return String</code> , <code>FunctionsReturnString</code>	false	Causes the connector to return binary/varbinary values as strings, if they do not have a tablename in the metadata.
<code>Host</code> , <code>Server</code> , <code>Data Source</code> , <code>DataSource</code> , <code>Address</code> , <code>Addr</code> , <code>Network Address</code>	localhost	The name or network address of the instance of MySQL to which to connect. Multiple hosts can be specified separated by commas. This can be useful where multiple MySQL servers are configured for replication and you are not concerned about the precise server you are connecting to. No attempt is made by the provider to synchronize writes to the database, so take care when using this option. In Unix environment with Mono, this can be a fully qualified path to a MySQL socket file. With this configuration, the Unix socket is used instead of the TCP/IP socket. Currently, only a single socket name can be given, so accessing MySQL in a replicated environment using Unix sockets is not currently supported.
<code>Ignore Prepare</code> , <code>IgnorePrepare</code>	true	When true, instructs the provider to ignore any calls to <code>MySqlCommand.Prepare()</code> . This option is provided to prevent issues with corruption of the statements when used with server-side prepared statements. If you use server-side prepare statements, set this option to false. This option was added in Connector/Net 5.0.3 and Connector/Net 1.0.9.
<code>includesecurityasserts</code> , <code>include security asserts</code>	false	Must be set to <code>true</code> when using the <code>MySQLClientPermissions</code> class in a partial trust environment, with the library installed in the GAC of the hosting environment. This requirement is new for partial-trust applications in Connector/Net 6.6.4 and higher. See <a href="#">Section 21.2.5.19, "Working with Partial Trust / Medium Trust"</a> for details.
<code>Initial Catalog</code> , <code>Database</code>	mysql	The case-sensitive name of the database to use initially.
<code>Interactive</code> , <code>Interactive Session</code> , <code>InteractiveSession</code>	false	If set to true, the client is interactive. An interactive client is one where the server variable <code>CLIENT_INTERACTIVE</code> is set. If an interactive client is set, the <code>wait_timeout</code> variable is set to the value of <code>interactive_timeout</code> . The client will then time out after this period of inactivity. For more details, see <a href="#">Section 5.1.4, "Server System Variables"</a> in the MySQL Reference Manual.
<code>Integrated Security</code> , <code>IntegratedSecurity</code>	no	Use Windows authentication when connecting to server. By default, it is turned off. To enable, specify a value of <code>yes</code> . (You can also use the value <code>sspi</code> as an alternative to <code>yes</code> .) For details, see <a href="#">Section 21.2.5.5, "Using the Windows Native Authentication Plugin"</a> . This option was introduced in Connector/Net 6.4.4.



Name	Default	Description
<code>Keep Alive, Keepalive</code>	0	For TCP connections, idle connection time measured in seconds, before the first keepalive packet is sent. A value of 0 indicates that keepalive is not used.
<code>Logging</code>	false	When true, various pieces of information is output to any configured TraceListeners. See <a href="#">Section 21.2.5.14, "Using the MySQL Connector/Net Trace Source Object"</a> for further details.
<code>Old Guids, OldGuids</code>	false	This option was introduced in Connector/Net 6.1.1. The backend representation of a GUID type was changed from <code>BINARY(16)</code> to <code>CHAR(36)</code> . This was done to allow developers to use the server function <code>UUID() [1280]</code> to populate a GUID table - <code>UUID()</code> generates a 36-character string. Developers of older applications can add <code>'Old Guids=true'</code> to the connection string to use a GUID of data type <code>BINARY(16)</code> .
<code>Old Syntax, OldSyntax, Use Old Syntax, UseOldSyntax</code>	false	This option was deprecated in Connector/Net 5.2.2. All code should now be written using the '@' symbol as the parameter marker.
<code>Password, pwd</code>		The password for the MySQL account being used.
<code>Persist Security Info, PersistSecurityInfo</code>	false	When set to <code>false</code> or <code>no</code> (strongly recommended), security-sensitive information, such as the password, is not returned as part of the connection if the connection is open or has ever been in an open state. Resetting the connection string resets all connection string values, including the password. Recognized values are <code>true</code> , <code>false</code> , <code>yes</code> , and <code>no</code> .
<code>Pipe Name, Pipe, PipeName</code>	mysql	When set to the name of a named pipe, the <code>MySqlConnection</code> attempts to connect to MySQL on that named pipe. This setting only applies to the Windows platform.
<code>Port</code>	3306	The port MySQL is using to listen for connections. This value is ignored if Unix socket is used.
<code>Procedure Cache Size, ProcedureCacheSize, procedure cache, procedurecache</code>	25	Sets the size of the stored procedure cache. By default, Connector/Net stores the metadata (input/output data types) about the last 25 stored procedures used. To disable the stored procedure cache, set the value to zero (0). This option was added in Connector/Net 5.0.2 and Connector/Net 1.0.9.
<code>Protocol, Connection Protocol, ConnectionProtocol</code>	socket	Specifies the type of connection to make to the server. Values can be: <code>socket</code> or <code>tcp</code> for a socket connection, <code>pipe</code> for a named pipe connection, <code>unix</code> for a Unix socket connection, <code>memory</code> to use MySQL shared memory.
<code>Replication</code>	false	Indicates if this connection is to use replicated servers.
<code>Respect Binary Flags, RespectBinaryFlags</code>	true	Setting this option to <code>false</code> means that Connector/Net ignores a column's binary flags as set by the server. This option was added in Connector/Net version 5.1.3.
<code>Shared Memory Name, SharedMemoryName</code>	MYSQL	The name of the shared memory object to use for communication if the connection protocol is set to <code>memory</code> .

Name	Default	Description
Sql Server Mode, sqlservermode	false	Allow SQL Server syntax. When set to <code>true</code> , enables Connector/Net to support square brackets around symbols instead of backticks. This enables Visual Studio wizards that bracket symbols with <code>[]</code> to work with Connector/Net. This option incurs a performance hit, so should only be used if necessary. This option was added in version 6.3.1.
SSL Mode, SslMode	None	<p>This option has the following values:</p> <ul style="list-style-type: none"> <li>• <b>None</b> - do not use SSL.</li> <li>• <b>Preferred</b> - use SSL if the server supports it, but allow connection in all cases.</li> <li>• <b>Required</b> - Always use SSL. Deny connection if server does not support SSL.</li> <li>• <b>VerifyCA</b> - Always use SSL. Validate the CA but tolerate name mismatch.</li> <li>• <b>VerifyFull</b> - Always use SSL. Fail if the host name is not correct.</li> </ul> <p>This option was introduced in MySQL Connector/Net 6.2.1.</p>
Table Cache, tablecache, tablecaching	false	Enables or disables caching of <code>TableDirect</code> commands. A value of <code>true</code> enables the cache while <code>false</code> disables it. For usage information about table caching, see <a href="#">Section 21.2.5.7, “Using Connector/Net with Table Caching”</a> . This option was added in Connector/Net 6.4.
Treat BLOBs as UTF8, TreatBlobsAsUTF8	false	
Treat Tiny As Boolean, TreatTinyAsBoolean	true	Setting this value to <code>false</code> causes <code>TINYINT(1)</code> to be treated as an <code>INT</code> . See <a href="#">Section 11.1.1, “Numeric Type Overview”</a> for a further explanation of the <code>TINYINT</code> and <code>BOOL</code> data types.
Use Affected Rows, UseAffectedRows	false	When <code>true</code> , the connection reports changed rows instead of found rows. This option was added in Connector/Net version 5.2.6.
Use Procedure Bodies, UseProcedureBodies, procedure bodies	true	When set to <code>true</code> , the default value, MySQL Connector/Net expects the body of the procedure to be viewable. This enables it to determine the parameter types and order. Set the option to <code>false</code> when the user connecting to the database does not have the <code>SELECT [755]</code> privileges for the <code>mysql.proc</code> (stored procedures) table, or cannot view <code>INFORMATION_SCHEMA.ROUTINES</code> . In this case, MySQL Connector/Net cannot determine the types and order of the parameters, and must be alerted to this fact by setting this option to <code>false</code> . When set to <code>false</code> , MySQL Connector/Net does not rely on this information being available when the procedure is called. Because MySQL Connector/Net will not be able to determine this information, explicitly set the types of all the parameters



Name	Default	Description
		before the call and add the parameters to the command in the same order as they appear in the procedure definition. This option was added in MySQL Connector/Net 5.0.4 and MySQL Connector/Net 1.0.10.
User Id, UserID, Username, Uid, User name, User		The MySQL login account being used.
Compress, Use Compression, UseCompression	false	<p>Setting this option to <code>true</code> enables compression of packets exchanged between the client and the server. This exchange is defined by the MySQL client/server protocol.</p> <p>Compression is used if both client and server support ZLIB compression, and the client has requested compression using this option.</p> <p>A compressed packet header is: packet length (3 bytes), packet number (1 byte), and Uncompressed Packet Length (3 bytes). The Uncompressed Packet Length is the number of bytes in the original, uncompressed packet. If this is zero, the data in this packet has not been compressed. When the compression protocol is in use, either the client or the server may compress packets. However, compression will not occur if the compressed length is greater than the original length. Thus, some packets will contain compressed data while other packets will not.</p>
Use Usage Advisor, Usage Advisor, UseUsageAdvisor	false	Logs inefficient database operations.
Use Performance Monitor, UsePerformanceMonitor, userperfmon, perfmon	false	Indicates that performance counters should be updated during execution.

## Connection Pooling Options

The following table lists the valid names for options related to connection pooling within the `ConnectionString`. For more information about connection pooling, see [Section 21.2.5.4, "Using Connector/Net with Connection Pooling"](#).

**Table 21.21. Connector/Net Connection String Options - Connection Pooling**

Name	Default	Description
Cache Server Properties, CacheServerProperties	false	Specifies whether server variable settings are updated by a <code>SHOW VARIABLES</code> command each time a pooled connection is returned. Enabling this setting speeds up connections in a connection pool environment. Your application is not informed of any changes to configuration variables made by other connections. This option was added in Connector/Net 6.3.
Connection Lifetime, ConnectionLifeTime	0	When a connection is returned to the pool, its creation time is compared with the current time, and the connection is destroyed if that time span (in seconds) exceeds the value specified by <code>Connection Lifetime</code> . This is useful in clustered configurations to force load balancing between a running server and a server just brought online. A value of zero (0) causes

Name	Default	Description
		pooled connections to have the maximum connection timeout.
<a href="#">Connection Reset</a> , <a href="#">ConnectionReset</a>	false	If true, the connection state is reset when it is retrieved from the pool. The default value of false avoids making an additional server round trip when obtaining a connection, but the connection state is not reset.
<a href="#">Maximum Pool Size</a> , <a href="#">Max Pool Size</a> , <a href="#">maximumpoolsize</a>	100	The maximum number of connections allowed in the pool.
<a href="#">Minimum Pool Size</a> , <a href="#">Min Pool Size</a> , <a href="#">MinimumPoolSize</a>	0	The minimum number of connections allowed in the pool.
<a href="#">Pooling</a>	true	When <a href="#">true</a> , the <a href="#">MySQLConnection</a> object is drawn from the appropriate pool, or if necessary, is created and added to the appropriate pool. Recognized values are <a href="#">true</a> , <a href="#">false</a> , <a href="#">yes</a> , and <a href="#">no</a> .

## 21.2.7. Connector/Net API Reference

This section of the manual contains a complete reference to the Connector/Net ADO.NET component, automatically generated from the embedded documentation.

### 21.2.7.1. [MySQL.Data.MySqlClient](#) Namespace

[Namespace hierarchy](#)

#### Classes

Class	Description
<a href="#">BaseCommandInterceptor</a>	Provides a means of enhancing or replacing SQL commands through the connection string rather than recompiling.
<a href="#">BaseExceptionInterceptor</a>	Provides a means of enabling and disabling exception handling through the connection string rather than recompiling.
<a href="#">MySQLClientPermission</a>	Derived from the .NET <a href="#">DBDataPermission</a> class. For usage information, see <a href="#">Section 21.2.5.19</a> , “ <a href="#">Working with Partial Trust / Medium Trust</a> ”.
<a href="#">MySQLCommand</a>	
<a href="#">MySQLCommandBuilder</a>	
<a href="#">MySQLConnection</a>	
<a href="#">MySQLDataAdapter</a>	
<a href="#">MySQLDataReader</a>	Provides a means of reading a forward-only stream of rows from a MySQL database. This class cannot be inherited.
<a href="#">MySQLError</a>	Collection of error codes that can be returned by the server
<a href="#">MySQLException</a>	The exception that is thrown when MySQL returns an error. This class cannot be inherited.
<a href="#">MySQLHelper</a>	Helper class that makes it easier to work with the provider.
<a href="#">MySQLInfoMessageEventArgs</a>	Provides data for the InfoMessage event. This class cannot be inherited.

Class	Description
<a href="#">MySqlParameter</a>	Represents a parameter to a <a href="#">MySqlCommand</a> , and optionally, its mapping to DataSet columns. This class cannot be inherited.
<a href="#">MySqlParameterCollection</a>	Represents a collection of parameters relevant to a <a href="#">MySqlCommand</a> as well as their respective mappings to columns in a DataSet. This class cannot be inherited.
<a href="#">MySqlRowUpdatedEventArgs</a>	Provides data for the RowUpdated event. This class cannot be inherited.
<a href="#">MySqlRowUpdatingEventArgs</a>	Provides data for the RowUpdating event. This class cannot be inherited.
<a href="#">MySqlTransaction</a>	

### Delegates

Delegate	Description
<a href="#">MySqlInfoMessageEventHandler</a>	Represents the method that will handle the <a href="#">InfoMessage</a> event of a <a href="#">MySqlConnection</a> .
<a href="#">MySqlRowUpdatedEventHandler</a>	Represents the method that will handle the RowUpdated event of a <a href="#">MySqlDataAdapter</a> .
<a href="#">MySqlRowUpdatingEventHandler</a>	Represents the method that will handle the RowUpdating event of a <a href="#">MySqlDataAdapter</a> .

### Enumerations

Enumeration	Description
<a href="#">MySqlDbType</a>	Specifies MySQL-specific data type of a field, property, for use in a <a href="#">MySqlParameter</a> .
<a href="#">MySqlErrorCode</a>	

#### 21.2.7.1.1. [MySql.Data.MySqlClientHierarchy](#)

##### See Also

[MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.2. [BaseCommandInterceptor](#)

The [BaseCommandInterceptor](#) class has these methods that you can override:

```
public virtual bool ExecuteScalar(string sql, ref object returnValue);
public virtual bool ExecuteNonQuery(string sql, ref int returnValue);
public virtual bool ExecuteReader(string sql, CommandBehavior behavior, ref MySqlDataReader returnValue);
public virtual void Init(MySqlConnection connection);
```

If your interceptor overrides one of the [Execute...](#) methods, set the [returnValue](#) output parameter and return [true](#) if you handled the event, or [false](#) if you did not handle the event. The SQL command is processed normally only when all command interceptors return [false](#).

The connection passed to the [Init](#) method is the connection that is attached to this interceptor.

For full usage and examples, see [Section 21.2.5.11, "Using the Connector/Net Interceptor Classes"](#).

#### 21.2.7.1.3. [BaseExceptionInterceptor](#)

You develop an exception interceptor first by creating a subclass of the [BaseExceptionInterceptor](#) class. You must override the [InterceptException\(\)](#) method. You can also override the [Init\(\)](#) method to do some one-time initialization.

Each exception interceptor has 2 methods:

```
public abstract Exception InterceptException(Exception exception,
    MySqlConnection connection);
public virtual void Init(MySqlConnection connection);
```

The connection passed to [Init\(\)](#) is the connection that is attached to this interceptor.

Each interceptor is required to override [InterceptException](#) and return an exception. It can return the exception it is given, or it can wrap it in a new exception. We currently do not offer the ability to suppress the exception.

For full usage and examples, see [Section 21.2.5.11, "Using the Connector/Net Interceptor Classes"](#).

#### 21.2.7.1.4. [MySqlCommand](#) Class

For a list of all members of this type, see [MySqlCommand Members](#).

##### Syntax: Visual Basic

```
NotInheritable Public Class MySqlCommand_
    Inherits Component_
    Implements IDbCommand, ICloneable
```

##### Syntax: C#

```
public sealed class MySqlCommand : Component, IDbCommand, ICloneable
```

##### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

##### Requirements

Namespace: [MySql.Data.MySqlClient](#)

Assembly: [MySql.Data](#) (in [MySql.Data.dll](#))

##### See Also

[MySqlCommand Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1. [MySqlCommand](#) Members

[MySqlCommand overview](#)

##### Public Instance Constructors

<a href="#">MySqlCommand</a>	Overloaded. Initializes a new instance of the <a href="#">MySqlCommand</a> class.
------------------------------	-----------------------------------------------------------------------------------

##### Public Instance Properties

<a href="#">CommandText</a>	
<a href="#">CommandTimeout</a>	
<a href="#">CommandType</a>	

<a href="#">Connection</a>	
Container (inherited from Component)	Gets the IContainer that contains the Component.
<a href="#">IsPrepared</a>	
<a href="#">Parameters</a>	
Site (inherited from Component)	Gets or sets the ISite of the Component.
<a href="#">Transaction</a>	
<a href="#">UpdatedRowSource</a>	

**Public Instance Methods**

<a href="#">Cancel</a>	Attempts to cancel the execution of a MySqlCommand. This operation is not supported.
CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<a href="#">CreateParameter</a>	Creates a new instance of a <a href="#">MySqlParameter</a> object.
Dispose (inherited from Component)	Releases all resources used by the Component.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">ExecuteNonQuery</a>	
<a href="#">ExecuteReader</a>	Overloaded.
<a href="#">ExecuteScalar</a>	
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType (inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">Prepare</a>	
ToString (inherited from Component)	Returns a String containing the name of the Component, if any. This method should not be overridden.

**Public Instance Events**

Disposed (inherited from Component)	Adds an event handler to listen to the Disposed event on the component.
-------------------------------------	-------------------------------------------------------------------------

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1. [MySqlCommand](#) Constructor

Initializes a new instance of the [MySqlCommand](#) class.

**Overload List**

Initializes a new instance of the [MySqlCommand](#) class.

- [public MySqlCommand\(\);](#)
- [public MySqlCommand\(string\);](#)
- [public MySqlCommand\(string, MySqlConnection\);](#)
- [public MySqlCommand\(string, MySqlConnection, MySqlTransaction\);](#)

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.1. [MySqlCommand](#) Constructor ()

Initializes a new instance of the [MySqlCommand](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlCommand();
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlCommand Constructor Overload List](#)

#### 21.2.7.1.4.1.1.2. [MySqlCommand](#) Constructor (String)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal cmdText As String _  
)
```

**Syntax: C#**

```
public MySqlCommand(  
    stringcmdText  
) ;
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlCommand Constructor Overload List](#)

#### 21.2.7.1.4.1.1.3. [MySqlCommand](#) Constructor

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal cmdText As String, _  
    ByVal connection As MySqlConnection _  
)
```

**Syntax: C#**

```
public MySqlCommand(  
    stringcmdText,  
    MySqlConnectionconnection  
) ;
```

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLCommand Constructor Overload List](#)

#### 21.2.7.1.4.1.1.3.1. [MySQLConnection](#) Class

For a list of all members of this type, see [MySQLConnection Members](#) .

##### Syntax: Visual Basic

```
NotInheritable Public Class MySqlConnection_
    Inherits Component_
    Implements IDbConnection, ICloneable
```

##### Syntax: C#

```
public sealed class MySqlConnection : Component, IDbConnection, ICloneable
```

##### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

##### Requirements

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

##### See Also

[MySQLConnection Members](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1. [MySQLConnection](#) Members

[MySQLConnection overview](#)

##### Public Instance Constructors

<a href="#">MySQLConnection</a>	Overloaded. Initializes a new instance of the MySqlConnection class.
---------------------------------	----------------------------------------------------------------------

##### Public Instance Properties

<a href="#">ConnectionString</a>	
<a href="#">ConnectionTimeout</a>	
Container (inherited from Component)	Gets the IContainer that contains the Component.
<a href="#">Database</a>	
<a href="#">DataSource</a>	Gets the name of the MySQL server to which to connect.
<a href="#">ServerThread</a>	Returns the ID of the server thread this connection is executing on.
<a href="#">ServerVersion</a>	
Site (inherited from Component)	Gets or sets the ISite of the Component.
<a href="#">State</a>	
<a href="#">UseCompression</a>	Indicates if this connection should use compression when communicating with the server.

##### Public Instance Methods

<a href="#">BeginTransaction</a>	Overloaded.
<a href="#">ChangeDatabase</a>	
<a href="#">Close</a>	
<a href="#">CreateCommand</a>	
CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Dispose (inherited from Component)	Releases all resources used by the Component.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType (inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">Open</a>	
<a href="#">Ping</a>	Ping
ToString (inherited from Component)	Returns a String containing the name of the Component, if any. This method should not be overridden.

### Public Instance Events

Disposed (inherited from Component)	Adds an event handler to listen to the Disposed event on the component.
<a href="#">InfoMessage</a>	
<a href="#">StateChange</a>	

### See Also

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.1. [MySQLConnection](#) Constructor

Initializes a new instance of the [MySQLConnection](#) class.

#### Overload List

Initializes a new instance of the [MySQLConnection](#) class.

- [public MySQLConnection\(\);](#)
- [public MySQLConnection\(string\);](#)

### See Also

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.1.1. [MySQLConnection](#) Constructor

Initializes a new instance of the [MySQLConnection](#) class.



**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlConnection();
```

**See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlConnection Constructor Overload List](#)

21.2.7.1.4.1.1.3.1.1.2. [MySqlConnection](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal connectionString As String _  
)
```

**Syntax: C#**

```
public MySqlConnection(  
    string connectionString  
) ;
```

**See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlConnection Constructor Overload List](#)

21.2.7.1.4.1.1.3.1.1.2. [ConnectionString](#) Property**Syntax: Visual Basic**

```
NotOverridable Public Property ConnectionString As String _  
    Implements IDbConnection.ConnectionString
```

**Syntax: C#**

```
public string ConnectionString {get; set;}
```

**Implements**

[IDbConnection.ConnectionString](#)

**See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.3. [ConnectionTimeout](#) Property**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property ConnectionTimeout As Integer _  
    Implements IDbConnection.ConnectionTimeout
```

**Syntax: C#**

```
public int ConnectionTimeout {get;}
```

**Implements**

[IDbConnection.ConnectionTimeout](#)

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.4. Database Property

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Database As String _  
-  
    Implements IDbConnection.Database
```

**Syntax: C#**

```
public string Database {get;}
```

**Implements**

IDbConnection.Database

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.5. DataSource Property

Gets the name of the MySQL server to which to connect.

**Syntax: Visual Basic**

```
Public ReadOnly Property DataSource As String
```

**Syntax: C#**

```
public string DataSource {get;}
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.6. ServerThread Property

Returns the ID of the server thread this connection is executing on

**Syntax: Visual Basic**

```
Public ReadOnly Property ServerThread As Integer
```

**Syntax: C#**

```
public int ServerThread {get;}
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.7. ServerVersion Property

**Syntax: Visual Basic**

```
Public ReadOnly Property ServerVersion As String
```

**Syntax: C#**

```
public string ServerVersion {get;}
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.8. State Property

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property State As ConnectionState _  
-  
    Implements IDbConnection.State
```

**Syntax: C#**

```
public System.Data.ConnectionState State {get;}
```

**Implements**

IDbConnection.State

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.9. UseCompression Property

Indicates if this connection should use compression when communicating with the server.

**Syntax: Visual Basic**

```
Public ReadOnly Property UseCompression As Boolean
```

**Syntax: C#**

```
public bool UseCompression {get;}
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.10. BeginTransaction Method

**Overload List**

- [public MySqlTransaction BeginTransaction\(\);](#)
- [public MySqlTransaction BeginTransaction\(IsolationLevel\);](#)

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.10.1. [MySQLConnection.BeginTransaction](#) Method**Syntax: Visual Basic**

```
Overloads Public Function BeginTransaction() As MySqlTransaction
```

**Syntax: C#**

```
public MySqlTransaction BeginTransaction();
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLConnection.BeginTransaction Overload List](#)

21.2.7.1.4.1.1.3.1.1.10.1.1.1. [MySQLTransaction](#) Class

For a list of all members of this type, see [MySQLTransaction Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLTransaction_
    Implements IDbTransaction, IDisposable
```

**Syntax: C#**

```
public sealed class MySQLTransaction : IDbTransaction, IDisposable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySQLTransaction Members](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.10.1.1.1.1. [MySQLTransaction](#) Members

[MySQLTransaction overview](#)

**Public Instance Properties**

<a href="#">Connection</a>	Gets the <a href="#">MySQLConnection</a> object associated with the transaction, or a null reference (Nothing in Visual Basic) if the transaction is no longer valid.
<a href="#">IsolationLevel</a>	Specifies the IsolationLevel for this transaction.

**Public Instance Methods**

<a href="#">Commit</a>	
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType (inherited from Object)	Gets the Type of the current instance.
<a href="#">Rollback</a>	
ToString (inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySQLTransaction Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.10.1.1.1.1.1.1. Connection Property

Gets the [MySQLConnection](#) object associated with the transaction, or a null reference (Nothing in Visual Basic) if the transaction is no longer valid.

**Syntax: Visual Basic**

```
Public ReadOnly Property Connection As MySqlConnection
```

**Syntax: C#**

```
public MySqlConnection Connection {get;}
```

**Property Value**

The [MySqlConnection](#) object associated with this transaction.

**Remarks**

A single application may have multiple database connections, each with zero or more transactions. This property enables you to determine the connection object associated with a particular transaction created by [BeginTransaction](#).

**See Also**

[MySqlTransaction Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.10.1.1.1.2. IsolationLevel Property

Specifies the IsolationLevel for this transaction.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property IsolationLevel As IsolationLevel _  
_ Implements IDbTransaction.IsolationLevel
```

**Syntax: C#**

```
public System.Data.IsolationLevel IsolationLevel {get;}
```

**Property Value**

The IsolationLevel for this transaction. The default is ReadCommitted.

**Implements**

IDbTransaction.IsolationLevel

**Remarks**

Parallel transactions are not supported. Therefore, the IsolationLevel applies to the entire transaction.

**See Also**

[MySqlTransaction Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.10.1.1.1.3. [MySqlTransaction.Commit](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Sub Commit() _  
_ Implements IDbTransaction.Commit
```

**Syntax: C#**

```
public void Commit();
```

**Implements**

---

IDbTransaction.Commit

**See Also**

[MySQLTransaction Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.10.1.1.1.4. [MySQLTransaction.Rollback](#) Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub Rollback() _  
-  
    Implements IDbTransaction.Rollback
```

**Syntax: C#**

```
public void Rollback();
```

**Implements**

IDbTransaction.Rollback

**See Also**

[MySQLTransaction Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.10.2. [MySQLConnection.BeginTransaction](#) Method

**Syntax: Visual Basic**

```
Overloads Public Function BeginTransaction( _  
    ByVal iso As IsolationLevel _  
) As MySQLTransaction
```

**Syntax: C#**

```
public MySQLTransaction BeginTransaction(  
    IsolationLevel iso  
);
```

**See Also**

[MySQLConnection Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLConnection.BeginTransaction Overload List](#)

21.2.7.1.4.1.1.3.1.1.11. [MySQLConnection.ChangeDatabase](#) Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub ChangeDatabase( _  
    ByVal databaseName As String _  
) _  
-  
    Implements IDbConnection.ChangeDatabase
```

**Syntax: C#**

```
public void ChangeDatabase(  
    string databaseName  
);
```

**Implements**

IDbConnection.ChangeDatabase

**See Also**

[MySQLConnection Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.12. [MySQLConnection.Close](#) Method

##### Syntax: Visual Basic

```
NotOverridable Public Sub Close() _  
-  
    Implements IDbConnection.Close
```

##### Syntax: C#

```
public void Close();
```

##### Implements

IDbConnection.Close

##### See Also

[MySQLConnection Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.13. [MySQLConnection.CreateCommand](#) Method

##### Syntax: Visual Basic

```
Public Function CreateCommand() As MySqlCommand
```

##### Syntax: C#

```
public MySqlCommand CreateCommand();
```

##### See Also

[MySQLConnection Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.14. [MySQLConnection.Open](#) Method

##### Syntax: Visual Basic

```
NotOverridable Public Sub Open() _  
-  
    Implements IDbConnection.Open
```

##### Syntax: C#

```
public void Open();
```

##### Implements

IDbConnection.Open

##### See Also

[MySQLConnection Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.15. [MySQLConnection.Ping](#) Method

Ping

##### Syntax: Visual Basic

```
Public Function Ping() As Boolean
```

##### Syntax: C#

```
public bool Ping();
```

**Return Value****See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16. [MySqlConnection.InfoMessage](#) Event**Syntax: Visual Basic**

```
Public Event InfoMessage As MySqlInfoMessageEventHandler
```

**Syntax: C#**

```
public event MySqlInfoMessageEventHandler InfoMessage;
```

**See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1. [MySqlInfoMessageEventHandler](#) Delegate

Represents the method that will handle the [InfoMessage](#) event of a [MySqlConnection](#).

**Syntax: Visual Basic**

```
Public Delegate Sub MySqlInfoMessageEventHandler( _  
    ByVal sender As Object, _  
    ByVal args As MySqlInfoMessageEventArgs _  
)
```

**Syntax: C#**

```
public delegate void MySqlInfoMessageEventHandler(  
    object sender,  
    MySqlInfoMessageEventArgs args  
);
```

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1. [MySqlInfoMessageEventArgs](#) Class

Provides data for the [InfoMessage](#) event. This class cannot be inherited.

For a list of all members of this type, see [MySqlInfoMessageEventArgs Members](#) .

**Syntax: Visual Basic**

```
Public Class MySqlInfoMessageEventArgs_  
    Inherits EventArgs
```

**Syntax: C#**

```
public class MySqlInfoMessageEventArgs : EventArgs
```

**Thread Safety**



Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

### Requirements

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlInfoMessageEventArgs Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.1.3.1.1.16.1.1.1.1. [MySqlInfoMessageEventArgs](#) Members

[MySqlInfoMessageEventArgs](#) overview

### Public Instance Constructors

<a href="#">MySqlInfoMessageEventArgs</a> Constructor	Initializes a new instance of the <a href="#">MySqlInfoMessageEventArgs</a> class.
-------------------------------------------------------	------------------------------------------------------------------------------------

### Public Instance Fields

<a href="#">errors</a>	
------------------------	--

### Public Instance Methods

Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType (inherited from Object)	Gets the Type of the current instance.
ToString (inherited from Object)	Returns a String that represents the current Object.

### Protected Instance Methods

Finalize (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone (inherited from Object)	Creates a shallow copy of the current Object.

### See Also

[MySqlInfoMessageEventArgs](#) Class, [MySql.Data.MySqlClient](#) Namespace

#### 21.2.7.1.4.1.1.3.1.1.16.1.1.1.1.1. [MySqlInfoMessageEventArgs](#) Constructor

Initializes a new instance of the [MySqlInfoMessageEventArgs](#) class.

### Syntax: Visual Basic

```
Public Sub New()
```

### Syntax: C#

```
public MySqlInfoMessageEventArgs();
```

**See Also**

[MySqlInfoMessageEventArgs Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1.1.2. [MySqlInfoMessageEventArgs.errors](#) Field

**Syntax: Visual Basic**

```
Public errors As MySqlError()
```

**Syntax: C#**

```
public MySqlError[] errors;
```

**See Also**

[MySqlInfoMessageEventArgs Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1. [MySqlError](#) Class

Collection of error codes that can be returned by the server

For a list of all members of this type, see [MySqlError Members](#) .

**Syntax: Visual Basic**

```
Public Class MySqlError
```

**Syntax: C#**

```
public class MySqlError
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlError Members](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1.1. [MySqlError](#) Members

[MySqlError overview](#)

**Public Instance Constructors**

<a href="#">MySqlError Constructor</a>	
----------------------------------------	--

**Public Instance Properties**

<a href="#">Code</a>	Error code
<a href="#">Level</a>	Error level
<a href="#">Message</a>	Error message

**Public Instance Methods**

Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType (inherited from Object)	Gets the Type of the current instance.
ToString (inherited from Object)	Returns a String that represents the current Object.

### Protected Instance Methods

Finalize (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone (inherited from Object)	Creates a shallow copy of the current Object.

### See Also

[MySqlError Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1.1.1. [MySqlError](#) Constructor

### Syntax: Visual Basic

```
Public Sub New( _
    ByVal level As String, _
    ByVal code As Integer, _
    ByVal message As String _
)
```

### Syntax: C#

```
public MySqlError(
    string level,
    int code,
    string message
);
```

### Parameters

- `level`:
- `code`:
- `message`:

### See Also

[MySqlError Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1.1.2. Code Property

Error code

### Syntax: Visual Basic

```
Public ReadOnly Property Code As Integer
```

### Syntax: C#

```
public int Code {get;}
```

**See Also**

[MySqlError Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1.1.3. Level Property

Error level

**Syntax: Visual Basic**

```
Public ReadOnly Property Level As String
```

**Syntax: C#**

```
public string Level {get;}
```

**See Also**

[MySqlError Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.1.3.1.1.16.1.1.1.2.1.1.4. Message Property

Error message

**Syntax: Visual Basic**

```
Public ReadOnly Property Message As String
```

**Syntax: C#**

```
public string Message {get;}
```

**See Also**

[MySqlError Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.3.1.1.17. [MySqlConnection.StateChange](#) Event**Syntax: Visual Basic**

```
Public Event StateChange As StateChangeEventHandler
```

**Syntax: C#**

```
public event StateChangeEventHandler StateChange;
```

**See Also**

[MySqlConnection Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.1.4. [MySqlCommand](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal cmdText As String, _  
    ByVal connection As MySqlConnection, _  
    ByVal transaction As MySqlTransaction _  
)
```

**Syntax: C#**

```
public MySqlCommand(  
    stringcmdText,  
    MySqlConnectionconnection,  
    MySqlTransactiontransaction
```

```
);
```

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#),  [MySqlCommand Constructor Overload List](#)

## 21.2.7.1.4.1.2. CommandText Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandText As String _  
-  
    Implements IDbCommand.CommandText
```

**Syntax: C#**

```
public string CommandText {get; set;}
```

**Implements**

IDbCommand.CommandText

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.3. CommandTimeout Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandTimeout As Integer _  
-  
    Implements IDbCommand.CommandTimeout
```

**Syntax: C#**

```
public int CommandTimeout {get; set;}
```

**Implements**

IDbCommand.CommandTimeout

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.4. CommandType Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandType As CommandType _  
-  
    Implements IDbCommand.CommandType
```

**Syntax: C#**

```
public System.Data.CommandType CommandType {get; set;}
```

**Implements**

IDbCommand.CommandType

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

---

#### 21.2.7.1.4.1.5. Connection Property

**Syntax: Visual Basic**

```
Public Property Connection As MySqlConnection
```

**Syntax: C#**

```
public MySqlConnection Connection {get; set;}
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.6. IsPrepared Property

**Syntax: Visual Basic**

```
Public ReadOnly Property IsPrepared As Boolean
```

**Syntax: C#**

```
public bool IsPrepared {get;}
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7. Parameters Property

**Syntax: Visual Basic**

```
Public ReadOnly Property Parameters As MySqlParameterCollection
```

**Syntax: C#**

```
public MySqlParameterCollection Parameters {get;}
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1. [MySqlParameterCollection](#) Class

Represents a collection of parameters relevant to a [MySqlCommand](#) as well as their respective mappings to columns in a DataSet. This class cannot be inherited.

For a list of all members of this type, see [MySqlParameterCollection Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlParameterCollection_  
    Inherits MarshalByRefObject_  
    Implements IDataParameterCollection, IList, ICollection, IEnumerable
```

**Syntax: C#**

```
public sealed class MySqlParameterCollection : MarshalByRefObject, IDataParameterCollection, IList, ICollection
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlParameterCollection Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1. [MySqlParameterCollection](#) Members

[MySqlParameterCollection overview](#)

### Public Instance Constructors

<a href="#">MySqlParameterCollection Constructor</a>	Initializes a new instance of the <a href="#">MySqlParameterCollection</a> class.
------------------------------------------------------	-----------------------------------------------------------------------------------

### Public Instance Properties

<a href="#">Count</a>	Gets the number of <a href="#">MySqlParameter</a> objects in the collection.
<a href="#">Item</a>	Overloaded. Gets the <a href="#">MySqlParameter</a> with a specified attribute. In C#, this property is the indexer for the <a href="#">MySqlParameterCollection</a> class.

### Public Instance Methods

<a href="#">Add</a>	Overloaded. Adds the specified <a href="#">MySqlParameter</a> object to the <a href="#">MySqlParameterCollection</a> .
<a href="#">Clear</a>	Removes all items from the collection.
<a href="#">Contains</a>	Overloaded. Gets a value indicating whether a <a href="#">MySqlParameter</a> exists in the collection.
<a href="#">CopyTo</a>	Copies <a href="#">MySqlParameter</a> objects from the <a href="#">MySqlParameterCollection</a> to the specified array.
CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType (inherited from Object)	Gets the Type of the current instance.
<a href="#">IndexOf</a>	Overloaded. Gets the location of a <a href="#">MySqlParameter</a> in the collection.
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">Insert</a>	Inserts a <a href="#">MySqlParameter</a> into the collection at the specified index.
<a href="#">Remove</a>	Removes the specified <a href="#">MySqlParameter</a> from the collection.

<a href="#">RemoveAt</a>	Overloaded. Removes the specified <a href="#">MySQLParameter</a> from the collection.
ToString (inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.1. [MySQLParameterCollection](#) Constructor

Initializes a new instance of the [MySQLParameterCollection](#) class.

**Syntax: Visual Basic**

```
Public Sub New()
```

**Syntax: C#**

```
public MySQLParameterCollection();
```

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.2. Count Property

Gets the number of [MySQLParameter](#) objects in the collection.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Count As Integer _  
_ Implements ICollection.Count
```

**Syntax: C#**

```
public int Count {get;}
```

**Implements**

ICollection.Count

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3. Item Property

Gets the [MySQLParameter](#) with a specified attribute. In C#, this property is the indexer for the [MySQLParameterCollection](#) class.

**Overload List**

Gets the [MySQLParameter](#) at the specified index.

- `public MySQLParameter this[int] {get; set;}`

Gets the [MySQLParameter](#) with the specified name.

- `public MySQLParameter this[string] {get; set;}`

**See Also**



[MySQLParameterCollection Class, MySql.Data.MySqlClient Namespace](#)21.2.7.1.4.1.7.1.1.3.1. [MySQLParameter](#) Class

Represents a parameter to a [MySQLCommand](#), and optionally, its mapping to DataSet columns. This class cannot be inherited.

For a list of all members of this type, see [MySQLParameter Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLParameter_
    Inherits MarshalByRefObject_
    Implements IDataParameter, IDbDataParameter, ICloneable
```

**Syntax: C#**

```
public sealed class MySQLParameter : MarshalByRefObject, IDataParameter, IDbDataParameter, ICloneable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySQLParameter Members, MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.3.1.1. [MySQLParameter](#) Members

[MySQLParameter overview](#)

**Public Instance Constructors**

<a href="#">MySQLParameter</a>	Overloaded. Initializes a new instance of the MySQLParameter class.
--------------------------------	---------------------------------------------------------------------

**Public Instance Properties**

<a href="#">DbType</a>	Gets or sets the DbType of the parameter.
<a href="#">Direction</a>	Gets or sets a value indicating whether the parameter is input-only, output-only, bidirectional, or a stored procedure return value parameter. As of MySQL version 4.1 and earlier, input-only is the only valid choice.
<a href="#">IsNullable</a>	Gets or sets a value indicating whether the parameter accepts null values.
<a href="#">IsUnsigned</a>	
<a href="#">MySqlDbType</a>	Gets or sets the MySqlDbType of the parameter.
<a href="#">ParameterName</a>	Gets or sets the name of the MySQLParameter.
<a href="#">Precision</a>	Gets or sets the maximum number of digits used to represent the <a href="#">Value</a> property.
<a href="#">Scale</a>	Gets or sets the number of decimal places to which <a href="#">Value</a> is resolved.

<a href="#">Size</a>	Gets or sets the maximum size, in bytes, of the data within the column.
<a href="#">SourceColumn</a>	Gets or sets the name of the source column that is mapped to the DataSet and used for loading or returning the <a href="#">Value</a> .
<a href="#">SourceVersion</a>	Gets or sets the DataRowVersion to use when loading <a href="#">Value</a> .
<a href="#">Value</a>	Gets or sets the value of the parameter.

**Public Instance Methods**

CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType (inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">ToString</a>	Overridden. Gets a string containing the <a href="#">ParameterName</a> .

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.3.1.1.1. [MySQLParameter](#) Constructor

Initializes a new instance of the MySQLParameter class.

**Overload List**

Initializes a new instance of the MySQLParameter class.

- [public MySQLParameter\(\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and the data type.

- [public MySQLParameter\(string,MySQLDbType\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the [MySQLDbType](#), and the size.

- [public MySQLParameter\(string,MySQLDbType,int\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the type of the parameter, the size of the parameter, a ParameterDirection, the precision of the parameter, the scale of the parameter, the source column, a DataRowVersion to use, and the value of the parameter.

- [public MySQLParameter\(string,MySQLDbType,int,ParameterDirection,bool,byte,byte,string,DataRowVersion,object\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the [MySQLDbType](#), the size, and the source column name.

- `public MySQLParameter(string,MySQLDbType,int,string);`

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and a value of the new [MySQLParameter](#).

- `public MySQLParameter(string,object);`

#### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.1.1. [MySQLParameter](#) Constructor ()

Initializes a new instance of the [MySQLParameter](#) class.

#### Syntax: Visual Basic

```
Overloads Public Sub New()
```

#### Syntax: C#

```
public MySQLParameter();
```

#### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLParameter Constructor Overload List](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.1.2. [MySQLParameter](#) Constructor

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and the data type.

#### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySQLDbType _  
)
```

#### Syntax: C#

```
public MySQLParameter(  
    stringparameterName,  
    MySQLDbTypedbType  
);
```

#### Parameters

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySQLDbType](#) values.

#### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLParameter Constructor Overload List](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.1.2.1. [MySQLDbType](#) Enumeration

Specifies MySQL-specific data type of a field, property, for use in a [MySQLParameter](#).

#### Syntax: Visual Basic

```
Public Enum MySqlDbType
```

**Syntax: C#**

```
public enum MySqlDbType
```

**Members**

Member Name	Description
Newdate	Obsolete. Use Datetime or Date type.
Timestamp	A timestamp. The range is '1970-01-01 00:00:01.000000' to '2038-01-19 03:14:07.999999'. (Fractional seconds can only be stored with a MySQL 5.6.4 or higher database server.)
Time	The range is '-838:59:59.000000' to '838:59:59.000000'. (Fractional seconds can only be stored with a MySQL 5.6.4 or higher database server.)
Date	The supported range is '1000-01-01' to '9999-12-31'.
Datetime	The supported range is '1000-01-01 00:00:00.000000' to '9999-12-31 23:59:59.999999'. (Fractional seconds can only be stored with a MySQL 5.6.4 or higher database server.)
Year	A year in 2- or 4-digit format (default is 4-digit). The allowable values are 1901 to 2155, 0000 in the 4-digit year format, and 1970-2069 if you use the 2-digit format (70-69).
TinyBlob	A BLOB column with a maximum length of 255 (2 <sup>8</sup> - 1) characters.
Blob	A BLOB column with a maximum length of 65535 (2 <sup>16</sup> - 1) characters.
MediumBlob	A BLOB column with a maximum length of 16777215 (2 <sup>24</sup> - 1) characters.
LongBlob	A BLOB column with a maximum length of 4294967295 or 4G (2 <sup>32</sup> - 1) characters.
Int16	A 16-bit signed integer. The signed range is -32768 to 32767. The unsigned range is 0 to 65535.
Int24	Specifies a 24 (3 byte) signed or unsigned value.
Int32	A 32-bit signed integer.
Int64	A 64-bit signed integer.
Byte	The signed range is -128 to 127. The unsigned range is 0 to 255.
Float	A small (single-precision) floating-point number. Allowable values are -3.402823466E+38 to -1.175494351E-38, 0, and 1.175494351E-38 to 3.402823466E+38.
Double	A normal-size (double-precision) floating-point number. Allowable

	values are -1.7976931348623157E+308 to -2.2250738585072014E-308, 0, and 2.2250738585072014E-308 to 1.7976931348623157E+308.
UByte	An 8-bit unsigned value.
UInt16	A 16-bit unsigned value.
UInt24	A 24-bit unsigned value.
UInt32	A 32-bit unsigned value.
UInt64	A 64-bit unsigned value.
Decimal	A fixed precision and scale numeric value between $-10^{38-1}$ and $10^{38-1}$ .
NewDecimal	New Decimal
Set	A set. A string object that can have zero or more values, each of which must be chosen from the list of values 'value1', 'value2', ... A SET can have a maximum of 64 members.
String	Obsolete. Use VarChar type.
VarChar	A variable-length string containing 0 to 255 characters.
VarString	A variable-length string containing 0 to 65535 characters.
Enum	An enumeration. A string object that can have only one value, chosen from the list of values 'value1', 'value2', ..., NULL or the special "" error value. An ENUM can have a maximum of 65535 distinct values.
Geometry	
Bit	Bit-field data type.
TinyText	A nonbinary string column supporting a maximum length of 255 ( $2^8 - 1$ ) characters.
Text	A nonbinary string column supporting a maximum length of 65535 ( $2^{16} - 1$ ) characters.
MediumText	A nonbinary string column supporting a maximum length of 16777215 ( $2^{24} - 1$ ) characters.
LongText	A nonbinary string column supporting a maximum length of 4294967295 ( $2^{32} - 1$ ) characters.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.3. [MySqlParameter](#) Constructor (String, MySqlDbType, Int32)

Initializes a new instance of the [MySqlParameter](#) class with the parameter name, the [MySqlDbType](#), and the size.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType, _
    ByVal size As Integer _
)
```

**Syntax: C#**

```
public MySqlParameter(
    stringparameterName,
    MySqlDbTypedbType,
    intsize
);
```

**Parameters**

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySqlDbType](#) values.
- `size`: The length of the parameter.

**See Also**

[MySqlParameter Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlParameter Constructor Overload List](#)

21.2.7.1.4.1.7.1.1.3.1.1.4. [MySqlParameter](#) Constructor

Initializes a new instance of the [MySqlParameter](#) class with the parameter name, the type of the parameter, the size of the parameter, a [ParameterDirection](#), the precision of the parameter, the scale of the parameter, the source column, a [DataRowVersion](#) to use, and the value of the parameter.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType, _
    ByVal size As Integer, _
    ByVal direction As ParameterDirection, _
    ByVal isNullable As Boolean, _
    ByVal precision As Byte, _
    ByVal scale As Byte, _
    ByVal sourceColumn As String, _
    ByVal sourceVersion As DataRowVersion, _
    ByVal value As Object _
)
```

**Syntax: C#**

```
public MySqlParameter(
    stringparameterName,
    MySqlDbTypedbType,
    intsize,
    ParameterDirectiondirection,
    boolisNullable,
    byteprecision,
    bytescale,
    stringsourceColumn,
    DataRowVersionsourceVersion,
    objectvalue
);
```

**Parameters**

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySqlDbType](#) values.

- **size**: The length of the parameter.
- **direction**: One of the ParameterDirection values.
- **isNullable**: true if the value of the field can be null, otherwise false.
- **precision**: The total number of digits to the left and right of the decimal point to which **Value** is resolved.
- **scale**: The total number of decimal places to which **Value** is resolved.
- **sourceColumn**: The name of the source column.
- **sourceVersion**: One of the DataRowVersion values.
- **value**: An Object that is the value of the **MySqlParameter**.

### Exceptions

Exception Type	Condition
ArgumentException	

### See Also

[MySqlParameter Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlParameter Constructor Overload List](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.4.1. Value Property

Gets or sets the value of the parameter.

#### Syntax: Visual Basic

```
NotOverridable Public Property Value As Object _
-
    Implements IDataParameter.Value
```

#### Syntax: C#

```
public object Value {get; set;}
```

#### Implements

IDataParameter.Value

#### See Also

[MySqlParameter Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.5. [MySqlParameter](#) Constructor

Initializes a new instance of the **MySqlParameter** class with the parameter name, the **MySqlDbType**, the size, and the source column name.

#### Syntax: Visual Basic

```
Overloads Public Sub New( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType, _
    ByVal size As Integer, _
    ByVal sourceColumn As String _
)
```

#### Syntax: C#

```
public MySqlParameter(  
    stringparameterName,  
    MySqlDbTypeDbType,  
    intsize,  
    stringsourceColumn  
);
```

### Parameters

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySqlDbType](#) values.
- `size`: The length of the parameter.
- `sourceColumn`: The name of the source column.

### See Also

[MySqlParameter Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlParameter Constructor Overload List](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.6. [MySqlParameter](#) Constructor

Initializes a new instance of the [MySqlParameter](#) class with the parameter name and a value of the new [MySqlParameter](#).

### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal parameterName As String, _  
    ByVal value As Object _  
)
```

### Syntax: C#

```
public MySqlParameter(  
    stringparameterName,  
    objectvalue  
);
```

### Parameters

- `parameterName`: The name of the parameter to map.
- `value`: An Object that is the value of the [MySqlParameter](#).

### See Also

[MySqlParameter Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlParameter Constructor Overload List](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.2. DbType Property

Gets or sets the DbType of the parameter.

### Syntax: Visual Basic

```
NotOverridable Public Property DbType As DbType _  
-  
    Implements IDataParameter.DbType
```

### Syntax: C#

```
public System.Data.DbType DbType {get; set;}
```



**Implements**

IDataParameter.DbType

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.3. Direction Property

Gets or sets a value indicating whether the parameter is input-only, output-only, bidirectional, or a stored procedure return value parameter. As of MySQL version 4.1 and earlier, input-only is the only valid choice.

**Syntax: Visual Basic**

```
NotOverridable Public Property Direction As ParameterDirection _
-
    Implements IDataParameter.Direction
```

**Syntax: C#**

```
public System.Data.ParameterDirection Direction {get; set;}
```

**Implements**

IDataParameter.Direction

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.4. IsNullable Property

Gets or sets a value indicating whether the parameter accepts null values.

**Syntax: Visual Basic**

```
NotOverridable Public Property IsNullable As Boolean _
-
    Implements IDataParameter.IsNullable
```

**Syntax: C#**

```
public bool IsNullable {get; set;}
```

**Implements**

IDataParameter.IsNullable

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.5. IsUnsigned Property

**Syntax: Visual Basic**

```
Public Property IsUnsigned As Boolean
```

**Syntax: C#**

```
public bool IsUnsigned {get; set;}
```

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.6. MySQLDbType Property

Gets or sets the MySQLDbType of the parameter.

##### Syntax: Visual Basic

```
Public Property MySQLDbType As MySQLDbType
```

##### Syntax: C#

```
public MySQLDbType MySQLDbType {get; set;}
```

##### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.7. ParameterName Property

Gets or sets the name of the MySQLParameter.

##### Syntax: Visual Basic

```
NotOverridable Public Property ParameterName As String _  
-  
    Implements IDataParameter.ParameterName
```

##### Syntax: C#

```
public string ParameterName {get; set;}
```

##### Implements

IDataParameter.ParameterName

##### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.8. Precision Property

Gets or sets the maximum number of digits used to represent the [Value](#) property.

##### Syntax: Visual Basic

```
NotOverridable Public Property Precision As Byte _  
-  
    Implements IDbDataParameter.Precision
```

##### Syntax: C#

```
public byte Precision {get; set;}
```

##### Implements

IDbDataParameter.Precision

##### See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.3.1.1.9. Scale Property

Gets or sets the number of decimal places to which [Value](#) is resolved.

**Syntax: Visual Basic**

```
NotOverridable Public Property Scale As Byte _  
-  
    Implements IDbDataParameter.Scale
```

**Syntax: C#**

```
public byte Scale {get; set;}
```

**Implements**

IDbDataParameter.Scale

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.1.10. Size Property

Gets or sets the maximum size, in bytes, of the data within the column.

**Syntax: Visual Basic**

```
NotOverridable Public Property Size As Integer _  
-  
    Implements IDbDataParameter.Size
```

**Syntax: C#**

```
public int Size {get; set;}
```

**Implements**

IDbDataParameter.Size

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.1.11. SourceColumn Property

Gets or sets the name of the source column that is mapped to the DataSet and used for loading or returning the [Value](#).

**Syntax: Visual Basic**

```
NotOverridable Public Property SourceColumn As String _  
-  
    Implements IDataParameter.SourceColumn
```

**Syntax: C#**

```
public string SourceColumn {get; set;}
```

**Implements**

IDataParameter.SourceColumn

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.1.1.12. SourceVersion Property

Gets or sets the DataRowVersion to use when loading [Value](#).

**Syntax: Visual Basic**

```
NotOverridable Public Property SourceVersion As DataRowVersion _
    Implements IDataParameter.SourceVersion
```

**Syntax: C#**

```
public System.Data.DataRowVersion SourceVersion {get; set;}
```

**Implements**

IDataParameter.SourceVersion

**See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.3.1.13. [MySQLParameter.ToString](#) Method

Overridden. Gets a string containing the [ParameterName](#).

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**Return Value****See Also**

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.3.2. Item Property (Int32)

Gets the [MySQLParameter](#) at the specified index.

**Syntax: Visual Basic**

```
Overloads Public Default Property Item( _
    ByVal index As Integer _
) As MySQLParameter
```

**Syntax: C#**

```
public MySQLParameter this[
    int index
] {get; set;}
```

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Item Overload List](#)

## 21.2.7.1.4.1.7.1.1.3.3. Item Property (String)

Gets the [MySQLParameter](#) with the specified name.

**Syntax: Visual Basic**

```
Overloads Public Default Property Item( _
    ByVal name As String _
) As MySQLParameter
```

**Syntax: C#**

```
public MySqlParameter this[
    stringname
] {get; set;}
```

**See Also**

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#),  
[MySqlParameterCollection.Item Overload List](#)

## 21.2.7.1.4.1.7.1.1.4. Add Method

Adds the specified [MySqlParameter](#) object to the [MySqlParameterCollection](#).

**Overload List**

Adds the specified [MySqlParameter](#) object to the [MySqlParameterCollection](#).

- [public MySqlParameter Add\(MySqlParameter\);](#)

Adds the specified [MySqlParameter](#) object to the [MySqlParameterCollection](#).

- [public int Add\(object\);](#)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) given the parameter name and the data type.

- [public MySqlParameter Add\(string,MySqlDbType\);](#)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) with the parameter name, the data type, and the column length.

- [public MySqlParameter Add\(string,MySqlDbType,int\);](#)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) with the parameter name, the data type, the column length, and the source column name.

- [public MySqlParameter Add\(string,MySqlDbType,int,string\);](#)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) given the specified parameter name and value.

- [public MySqlParameter Add\(string,object\);](#)

**See Also**

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.4.1. [MySqlParameterCollection.Add](#) Method

Adds the specified [MySqlParameter](#) object to the [MySqlParameterCollection](#).

**Syntax: Visual Basic**

```
Overloads Public Function Add( _
    ByVal value As MySqlParameter _
) As MySqlParameter
```

**Syntax: C#**

```
public MySqlParameter Add(
    MySqlParameter value
);
```

**Parameters**

- **value**: The [MySQLParameter](#) to add to the collection.

**Return Value**

The newly added [MySQLParameter](#) object.

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Add Overload List](#)

21.2.7.1.4.1.7.1.1.4.2. [MySQLParameterCollection.Add](#) Method

Adds the specified [MySQLParameter](#) object to the [MySQLParameterCollection](#).

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Function Add( _
    ByVal value As Object _
) As Integer _
-
    Implements IList.Add
```

**Syntax: C#**

```
public int Add(
    objectvalue
);
```

**Parameters**

- **value**: The [MySQLParameter](#) to add to the collection.

**Return Value**

The index of the new [MySQLParameter](#) object.

**Implements**

IList.Add

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Add Overload List](#)

21.2.7.1.4.1.7.1.1.4.3. [MySQLParameterCollection.Add](#) Method

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) given the parameter name and the data type.

**Syntax: Visual Basic**

```
Overloads Public Function Add( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType _
) As MySQLParameter
```

**Syntax: C#**

```
public MySQLParameter Add(
    stringparameterName,
    MySqlDbTypedbType
);
```

## Parameters

- `parameterName`: The name of the parameter.
- `dbType`: One of the [MySQLDbType](#) values.

## Return Value

The newly added [MySQLParameter](#) object.

## See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Add Overload List](#)

### 21.2.7.1.4.1.7.1.1.4.4. [MySQLParameterCollection.Add](#) Method

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) with the parameter name, the data type, and the column length.

#### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySQLDbType, _  
    ByVal size As Integer _  
) As MySQLParameter
```

#### Syntax: C#

```
public MySQLParameter Add(  
    string parameterName,  
    MySQLDbType dbType,  
    int size  
);
```

## Parameters

- `parameterName`: The name of the parameter.
- `dbType`: One of the [MySQLDbType](#) values.
- `size`: The length of the column.

## Return Value

The newly added [MySQLParameter](#) object.

## See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Add Overload List](#)

### 21.2.7.1.4.1.7.1.1.4.5. [MySQLParameterCollection.Add](#) Method

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) with the parameter name, the data type, the column length, and the source column name.

#### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySQLDbType, _  
    ByVal size As Integer, _  
    ByVal sourceColumn As String _  
) As MySQLParameter
```

**Syntax: C#**

```
public MySqlParameter Add(  
stringparameterName,  
MySqlDbTypeDbType,  
intsize,  
stringsourceColumn  
);
```

**Parameters**

- `parameterName`: The name of the parameter.
- `dbType`: One of the [MySqlDbType](#) values.
- `size`: The length of the column.
- `sourceColumn`: The name of the source column.

**Return Value**

The newly added [MySqlParameter](#) object.

**See Also**

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#),  
[MySqlParameterCollection.Add Overload List](#)

**21.2.7.1.4.1.7.1.1.4.6. [MySqlParameterCollection.Add](#) Method**

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) given the specified parameter name and value.

**Syntax: Visual Basic**

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal value As Object _  
) As MySqlParameter
```

**Syntax: C#**

```
public MySqlParameter Add(  
stringparameterName,  
objectvalue  
);
```

**Parameters**

- `parameterName`: The name of the parameter.
- `value`: The [Value](#) of the [MySqlParameter](#) to add to the collection.

**Return Value**

The newly added [MySqlParameter](#) object.

**See Also**

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#),  
[MySqlParameterCollection.Add Overload List](#)

**21.2.7.1.4.1.7.1.1.5. [MySqlParameterCollection.Clear](#) Method**

Removes all items from the collection.



**Syntax: Visual Basic**

```
NotOverridable Public Sub Clear() _
-
  Implements IList.Clear
```

**Syntax: C#**

```
public void Clear();
```

**Implements**

IList.Clear

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.6. Contains Method

Gets a value indicating whether a [MySQLParameter](#) exists in the collection.

**Overload List**

Gets a value indicating whether a [MySQLParameter](#) exists in the collection.

- [public bool Contains\(object\);](#)

Gets a value indicating whether a [MySQLParameter](#) with the specified parameter name exists in the collection.

- [public bool Contains\(string\);](#)

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.6.1. MySQLParameterCollection.Contains Method

Gets a value indicating whether a [MySQLParameter](#) exists in the collection.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Function Contains( _
  ByVal value As Object _
) As Boolean _
-
  Implements IList.Contains
```

**Syntax: C#**

```
public bool Contains(
  objectvalue
);
```

**Parameters**

- [value](#): The value of the [MySQLParameter](#) object to find.

**Return Value**

true if the collection contains the [MySQLParameter](#) object; otherwise, false.

**Implements**

IList.Contains

## See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Contains Overload List](#)

### 21.2.7.1.4.1.7.1.1.6.2. [MySQLParameterCollection.Contains](#) Method

Gets a value indicating whether a [MySQLParameter](#) with the specified parameter name exists in the collection.

#### Syntax: Visual Basic

```
NotOverridable Overloads Public Function Contains( _  
    ByVal name As String _  
) As Boolean _  
-  
    Implements IDataParameterCollection.Contains
```

#### Syntax: C#

```
public bool Contains(  
    stringname  
);
```

#### Parameters

- **name**: The name of the [MySQLParameter](#) object to find.

#### Return Value

true if the collection contains the parameter; otherwise, false.

#### Implements

[IDataParameterCollection.Contains](#)

#### See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.Contains Overload List](#)

### 21.2.7.1.4.1.7.1.1.7. [MySQLParameterCollection.CopyTo](#) Method

Copies [MySQLParameter](#) objects from the [MySQLParameterCollection](#) to the specified array.

#### Syntax: Visual Basic

```
NotOverridable Public Sub CopyTo( _  
    ByVal array As Array, _  
    ByVal index As Integer _  
) _  
-  
    Implements ICollection.CopyTo
```

#### Syntax: C#

```
public void CopyTo(  
    Arrayarray,  
    intindex  
);
```

#### Parameters

- **array**:
- **index** [899]:

**Implements**

ICollection.CopyTo

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.1.8. IndexOf Method

Gets the location of a [MySQLParameter](#) in the collection.

**Overload List**

Gets the location of a [MySQLParameter](#) in the collection.

- `public int IndexOf(object);`

Gets the location of the [MySQLParameter](#) in the collection with a specific parameter name.

- `public int IndexOf(string);`

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.1.8.1. [MySQLParameterCollection.IndexOf](#) Method

Gets the location of a [MySQLParameter](#) in the collection.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Function IndexOf( _  
    ByVal value As Object _  
) As Integer _  
-  
    Implements IList.IndexOf
```

**Syntax: C#**

```
public int IndexOf(  
    objectvalue  
);
```

**Parameters**

- `value`: The [MySQLParameter](#) object to locate.

**Return Value**

The zero-based location of the [MySQLParameter](#) in the collection.

**Implements**

IList.IndexOf

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.IndexOf Overload List](#)

21.2.7.1.4.1.7.1.1.8.2. [MySQLParameterCollection.IndexOf](#) Method

Gets the location of the [MySQLParameter](#) in the collection with a specific parameter name.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Function IndexOf( _  
    ByVal parameterName As String _  
) As Integer _  
-  
    Implements IDataParameterCollection.IndexOf
```

### Syntax: C#

```
public int IndexOf(  
    string parameterName  
);
```

### Parameters

- `parameterName`: The name of the [MySqlParameter](#) object to retrieve.

### Return Value

The zero-based location of the [MySqlParameter](#) in the collection.

### Implements

[IDataParameterCollection.IndexOf](#)

### See Also

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#),  
[MySqlParameterCollection.IndexOf Overload List](#)

#### 21.2.7.1.4.1.7.1.1.9. [MySqlParameterCollection.Insert](#) Method

Inserts a [MySqlParameter](#) into the collection at the specified index.

### Syntax: Visual Basic

```
NotOverridable Public Sub Insert( _  
    ByVal index As Integer, _  
    ByVal value As Object _  
) _  
-  
    Implements IList.Insert
```

### Syntax: C#

```
public void Insert(  
    int index,  
    object value  
);
```

### Parameters

- `index` [899]:
- `value`:

### Implements

[IList.Insert](#)

### See Also

[MySqlParameterCollection Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.7.1.1.10. [MySqlParameterCollection.Remove](#) Method

Removes the specified [MySqlParameter](#) from the collection.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Remove( _
    ByVal value As Object _
) _
- Implements IList.Remove
```

**Syntax: C#**

```
public void Remove(
    objectvalue
);
```

**Parameters**

- `value`:

**Implements**

IList.Remove

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.7.1.11. RemoveAt Method

Removes the specified [MySQLParameter](#) from the collection.

**Overload List**

Removes the specified [MySQLParameter](#) from the collection using a specific index.

- `public void RemoveAt(int);`

Removes the specified [MySQLParameter](#) from the collection using the parameter name.

- `public void RemoveAt(string);`

**See Also**

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.7.1.11.1. [MySQLParameterCollection.RemoveAt](#) Method

Removes the specified [MySQLParameter](#) from the collection using a specific index.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Sub RemoveAt( _
    ByVal index As Integer _
) _
- Implements IList.RemoveAt
```

**Syntax: C#**

```
public void RemoveAt(
    intindex
);
```

**Parameters**

- `index` [899]: The zero-based index of the parameter.

**Implements**

IList.RemoveAt

### See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.RemoveAt Overload List](#)

#### 21.2.7.1.4.1.7.1.11.2. [MySQLParameterCollection.RemoveAt](#) Method

Removes the specified [MySQLParameter](#) from the collection using the parameter name.

### Syntax: Visual Basic

```
NotOverridable Overloads Public Sub RemoveAt( _  
    ByVal name As String _  
) _  
- Implements IDataParameterCollection.RemoveAt
```

### Syntax: C#

```
public void RemoveAt(  
    stringname  
);
```

### Parameters

- [name](#): The name of the [MySQLParameter](#) object to retrieve.

### Implements

[IDataParameterCollection.RemoveAt](#)

### See Also

[MySQLParameterCollection Class](#), [MySQL.Data.MySqlClient Namespace](#),  
[MySQLParameterCollection.RemoveAt Overload List](#)

#### 21.2.7.1.4.1.8. Transaction Property

### Syntax: Visual Basic

```
Public Property Transaction As MySQLTransaction
```

### Syntax: C#

```
public MySQLTransaction Transaction {get; set;}
```

### See Also

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.9. UpdatedRowSource Property

### Syntax: Visual Basic

```
NotOverridable Public Property UpdatedRowSource As UpdateRowSource _  
- Implements IDbCommand.UpdatedRowSource
```

### Syntax: C#

```
public System.Data.UpdateRowSource UpdatedRowSource {get; set;}
```

### Implements

IDbCommand.UpdatedRowSource

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.10.  [MySqlCommand.Cancel](#) Method

Attempts to cancel the execution of a MySqlCommand. This operation is not supported.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Cancel() _
-
    Implements IDbCommand.Cancel
```

**Syntax: C#**

```
public void Cancel();
```

**Implements**

IDbCommand.Cancel

**Remarks**

Cancelling an executing command is currently not supported on any version of MySQL.

**Exceptions**

Exception Type	Condition
NotSupportedException	This operation is not supported.

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.11.  [MySqlCommand.CreateParameter](#) Method

Creates a new instance of a  [MySqlParameter](#) object.

**Syntax: Visual Basic**

```
Public Function CreateParameter() As MySqlParameter
```

**Syntax: C#**

```
public MySqlParameter CreateParameter();
```

**Return Value**

A  [MySqlParameter](#) object.

**Remarks**

This method is a strongly-typed version of CreateParameter.

**See Also**

[MySqlCommand Class](#),  [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.12.  [MySqlCommand.ExecuteNonQuery](#) Method

**Syntax: Visual Basic**

```
NotOverridable Public Function ExecuteNonQuery() As Integer _
_
_ Implements IDbCommand.ExecuteNonQuery
```

**Syntax: C#**

```
public int ExecuteNonQuery();
```

**Implements**

IDbCommand.ExecuteNonQuery

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.13. ExecuteReader Method

**Overload List**

- [public MySqlDataReader ExecuteReader\(\);](#)
- [public MySqlDataReader ExecuteReader\(CommandBehavior\);](#)

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1. [MySQLCommand.ExecuteReader](#) Method**Syntax: Visual Basic**

```
Overloads Public Function ExecuteReader() As MySqlDataReader
```

**Syntax: C#**

```
public MySqlDataReader ExecuteReader();
```

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLCommand.ExecuteReader Overload List](#)

21.2.7.1.4.1.13.1.1. [MySqlDataReader](#) Class

Provides a means of reading a forward-only stream of rows from a MySQL database. This class cannot be inherited.

For a list of all members of this type, see [MySqlDataReader Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlDataReader_
_ Inherits MarshalByRefObject_
_ Implements IEnumerable, IDataReader, IDisposable, IDataRecord
```

**Syntax: C#**

```
public sealed class MySqlDataReader : MarshalByRefObject, IEnumerable, IDataReader, IDisposable, IDataRecord
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**



Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlDataReader Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.1. [MySqlDataReader](#) Members

[MySqlDataReader overview](#)

### Public Instance Properties

<a href="#">Depth</a>	Gets a value indicating the depth of nesting for the current row. This method is not supported currently and always returns 0.
<a href="#">FieldCount</a>	Gets the number of columns in the current row.
<a href="#">HasRows</a>	Gets a value indicating whether the <a href="#">MySqlDataReader</a> contains one or more rows.
<a href="#">IsClosed</a>	Gets a value indicating whether the data reader is closed.
<a href="#">Item</a>	Overloaded. Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the <a href="#">MySqlDataReader</a> class.
<a href="#">RecordsAffected</a>	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.

### Public Instance Methods

<a href="#">Close</a>	Closes the <a href="#">MySqlDataReader</a> object.
<a href="#">CreateObjRef</a> (inherited from <a href="#">MarshalByRefObject</a> )	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<a href="#">Equals</a> (inherited from <a href="#">Object</a> )	Determines whether the specified <a href="#">Object</a> is equal to the current <a href="#">Object</a> .
<a href="#">GetBoolean</a>	Gets the value of the specified column as a <a href="#">Boolean</a> .
<a href="#">GetByte</a>	Gets the value of the specified column as a byte.
<a href="#">GetBytes</a>	Reads a stream of bytes from the specified column offset into the buffer an array starting at the given buffer offset.
<a href="#">GetChar</a>	Gets the value of the specified column as a single character.
<a href="#">GetChars</a>	Reads a stream of characters from the specified column offset into the buffer as an array starting at the given buffer offset.
<a href="#">GetDataTypeName</a>	Gets the name of the source data type.
<a href="#">GetDateTime</a>	
<a href="#">GetDecimal</a>	
<a href="#">GetDouble</a>	
<a href="#">GetFieldType</a>	Gets the <a href="#">Type</a> that is the data type of the object.

<a href="#">GetFloat</a>	
<a href="#">GetGuid</a>	
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetInt16</a>	
<a href="#">GetInt32</a>	
<a href="#">GetInt64</a>	
<a href="#">GetLifetimeService</a> (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
<a href="#">GetMySqlDateTime</a>	
<a href="#">GetName</a>	Gets the name of the specified column.
<a href="#">GetOrdinal</a>	Gets the column ordinal, given the name of the column.
<a href="#">GetSchemaTable</a>	Returns a DataTable that describes the column metadata of the MySqlDataReader.
<a href="#">GetString</a>	
<a href="#">GetTimeSpan</a>	
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.
<a href="#">GetUInt16</a>	
<a href="#">GetUInt32</a>	
<a href="#">GetUInt64</a>	
<a href="#">GetValue</a>	Gets the value of the specified column in its native format.
<a href="#">GetValues</a>	Gets all attribute columns in the collection for the current row.
<a href="#">InitializeLifetimeService</a> (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">IsDBNull</a>	Gets a value indicating whether the column contains non-existent or missing values.
<a href="#">NextResult</a>	Advances the data reader to the next result, when reading the results of batch SQL statements.
<a href="#">Read</a>	Advances the MySqlDataReader to the next record.
<a href="#">ToString</a> (inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.13.1.1.1.1. Depth Property

Gets a value indicating the depth of nesting for the current row. This method is not supported currently and always returns 0.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Depth As Integer _
    Implements IDataReader.Depth
```

**Syntax: C#**

```
public int Depth {get;}
```

**Implements**

IDataReader.Depth

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.13.1.1.1.2. FieldCount Property

Gets the number of columns in the current row.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property FieldCount As Integer _  
- Implements IDataRecord.FieldCount
```

**Syntax: C#**

```
public int FieldCount {get;}
```

**Implements**

IDataRecord.FieldCount

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.13.1.1.1.3. HasRows Property

Gets a value indicating whether the MySqlDataReader contains one or more rows.

**Syntax: Visual Basic**

```
Public ReadOnly Property HasRows As Boolean
```

**Syntax: C#**

```
public bool HasRows {get;}
```

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.4.1.13.1.1.1.4. IsClosed Property

Gets a value indicating whether the data reader is closed.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property IsClosed As Boolean _  
- Implements IDataReader.IsClosed
```

**Syntax: C#**

```
public bool IsClosed {get;}
```

**Implements**

IDataReader.IsClosed

### See Also

[MySQLDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.1.5. Item Property

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

### Overload List

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

- [public object this\[int\] {get;}](#)

Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

- [public object this\[string\] {get;}](#)

### See Also

[MySQLDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.1.5.1. Item Property (Int32)

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

### Syntax: Visual Basic

```
NotOverridable Overloads Public Default ReadOnly Property Item( _
    ByVal i As Integer _
) _
-
Implements IDataRecord.Item As Object _
-
Implements IDataRecord.Item
```

### Syntax: C#

```
public object this[
    inti
] {get;}
```

### Implements

IDataRecord.Item

### See Also

[MySQLDataReader Class](#), [MySql.Data.MySqlClient Namespace](#), [MySQLDataReader.Item Overload List](#)

#### 21.2.7.1.4.1.13.1.1.1.5.2. Item Property (String)

Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

### Syntax: Visual Basic

```
NotOverridable Overloads Public Default ReadOnly Property Item( _
    ByVal name As String _
```

```
) _  
- Implements IDataRecord.Item As Object _  
- Implements IDataRecord.Item
```

**Syntax: C#**

```
public object this[  
stringname  
] {get;}
```

**Implements**

IDataRecord.Item

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLDataReader.Item Overload List](#)

## 21.2.7.1.4.1.13.1.1.1.6. RecordsAffected Property

Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property RecordsAffected As Integer _  
- Implements IDataReader.RecordsAffected
```

**Syntax: C#**

```
public int RecordsAffected {get;}
```

**Implements**

IDataReader.RecordsAffected

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.7. [MySQLDataReader.Close](#) Method

Closes the MySQLDataReader object.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Close() _  
- Implements IDataReader.Close
```

**Syntax: C#**

```
public void Close();
```

**Implements**

IDataReader.Close

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.8. [MySQLDataReader.GetBoolean](#) Method

Gets the value of the specified column as a Boolean.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetBoolean( _  
    ByVal i As Integer _  
) As Boolean _  
-  
    Implements IDataRecord.GetBoolean
```

**Syntax: C#**

```
public bool GetBoolean(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetBoolean

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

**21.2.7.1.4.1.13.1.1.1.9. MySQLDataReader.GetByte Method**

Gets the value of the specified column as a byte.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetByte( _  
    ByVal i As Integer _  
) As Byte _  
-  
    Implements IDataRecord.GetByte
```

**Syntax: C#**

```
public byte GetByte(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetByte

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

**21.2.7.1.4.1.13.1.1.1.10. MySQLDataReader.GetBytes Method**

Reads a stream of bytes from the specified column offset into the buffer an array starting at the given buffer offset.

**Syntax: Visual Basic**

```

NotOverridable Public Function GetBytes( _
    ByVal i As Integer, _
    ByVal dataIndex As Long, _
    ByVal buffer As Byte(), _
    ByVal bufferIndex As Integer, _
    ByVal length As Integer _
) As Long _
-
    Implements IDataRecord.GetBytes

```

**Syntax: C#**

```

public long GetBytes(
    inti,
    longdataIndex,
    byte[]buffer,
    intbufferIndex,
    intlength
);

```

**Parameters**

- **i**: The zero-based column ordinal.
- **dataIndex**: The index within the field from which to begin the read operation.
- **buffer**: The buffer into which to read the stream of bytes.
- **bufferIndex**: The index for buffer to begin the read operation.
- **length**: The maximum length to copy into the buffer.

**Return Value**

The actual number of bytes read.

**Implements**

IDataRecord.GetBytes

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.11. [MySqlDataReader.GetChar](#) Method

Gets the value of the specified column as a single character.

**Syntax: Visual Basic**

```

NotOverridable Public Function GetChar( _
    ByVal i As Integer _
) As Char _
-
    Implements IDataRecord.GetChar

```

**Syntax: C#**

```

public char GetChar(
    inti
);

```

**Parameters**

- **i**:

**Return Value**

**Implements**

IDataRecord.GetChar

**See Also**[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)21.2.7.1.4.1.13.1.1.12. [MySQLDataReader.GetChars](#) Method

Reads a stream of characters from the specified column offset into the buffer as an array starting at the given buffer offset.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetChars( _  
    ByVal i As Integer, _  
    ByVal fieldOffset As Long, _  
    ByVal buffer As Char(), _  
    ByVal bufferoffset As Integer, _  
    ByVal length As Integer _  
) As Long _  
-  
    Implements IDataRecord.GetChars
```

**Syntax: C#**

```
public long GetChars(  
    inti,  
    longfieldOffset,  
    char[]buffer,  
    intbufferoffset,  
    intlength  
);
```

**Parameters**

- `i`:
- `fieldOffset`:
- `buffer`:
- `bufferoffset`:
- `length`:

**Return Value****Implements**

IDataRecord.GetChars

**See Also**[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)21.2.7.1.4.1.13.1.1.13. [MySQLDataReader.GetDataTypeName](#) Method

Gets the name of the source data type.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDataTypeName( _  
    ByVal i As Integer _  
) As String _  
-
```



Implements `IDataRecord.GetDataTypeName`

**Syntax: C#**

```
public string GetDataTypeName(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

`IDataRecord.GetDataTypeName`

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.14. `MySqlDataReader.GetDateTime` Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDateTime( _  
    ByVal index As Integer _  
) As Date _  
—  
Implements IDataRecord.GetDateTime
```

**Syntax: C#**

```
public DateTime GetDateTime(  
    int index  
);
```

**Implements**

`IDataRecord.GetDateTime`

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.15. `MySqlDataReader.GetDecimal` Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDecimal( _  
    ByVal index As Integer _  
) As Decimal _  
—  
Implements IDataRecord.GetDecimal
```

**Syntax: C#**

```
public decimal GetDecimal(  
    int index  
);
```

**Implements**

`IDataRecord.GetDecimal`

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.16. [MySQLDataReader.GetDouble](#) Method

##### Syntax: Visual Basic

```
NotOverridable Public Function GetDouble( _  
    ByVal index As Integer _  
) As Double _  
-  
    Implements IDataRecord.GetDouble
```

##### Syntax: C#

```
public double GetDouble(  
    int index  
);
```

##### Implements

IDataRecord.GetDouble

##### See Also

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.17. [MySQLDataReader.GetFieldType](#) Method

Gets the Type that is the data type of the object.

##### Syntax: Visual Basic

```
NotOverridable Public Function GetFieldType( _  
    ByVal i As Integer _  
) As Type _  
-  
    Implements IDataRecord.GetFieldType
```

##### Syntax: C#

```
public Type GetFieldType(  
    inti  
);
```

##### Parameters

- `i`:

##### Return Value

##### Implements

IDataRecord.GetFieldType

##### See Also

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.18. [MySQLDataReader.GetFloat](#) Method

##### Syntax: Visual Basic

```
NotOverridable Public Function GetFloat( _  
    ByVal index As Integer _  
) As Single _  
-  
    Implements IDataRecord.GetFloat
```

**Syntax: C#**

```
public float GetFloat(  
    int index  
);
```

**Implements**

IDataRecord.GetFloat

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.19. [MySqlDataReader.GetGuid](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function GetGuid( _  
    ByVal index As Integer _  
) As Guid _  
—  
    Implements IDataRecord.GetGuid
```

**Syntax: C#**

```
public Guid GetGuid(  
    int index  
);
```

**Implements**

IDataRecord.GetGuid

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.20. [MySqlDataReader.GetInt16](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt16( _  
    ByVal index As Integer _  
) As Short _  
—  
    Implements IDataRecord.GetInt16
```

**Syntax: C#**

```
public short GetInt16(  
    int index  
);
```

**Implements**

IDataRecord.GetInt16

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.21. [MySqlDataReader.GetInt32](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt32( _
```

```
    ByVal index As Integer _  
) As Integer _  
-  
    Implements IDataRecord.GetInt32
```

**Syntax: C#**

```
public int GetInt32(  
    int index  
);
```

**Implements**

IDataRecord.GetInt32

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.22. [MySQLDataReader.GetInt64](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt64( _  
    ByVal index As Integer _  
) As Long _  
-  
    Implements IDataRecord.GetInt64
```

**Syntax: C#**

```
public long GetInt64(  
    int index  
);
```

**Implements**

IDataRecord.GetInt64

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.23. [MySQLDataReader.GetMySqlDateTime](#) Method**Syntax: Visual Basic**

```
Public Function GetMySqlDateTime( _  
    ByVal index As Integer _  
) As MySqlDateTime
```

**Syntax: C#**

```
public MySqlDateTime GetMySqlDateTime(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.24. [MySQLDataReader.GetName](#) Method

Gets the name of the specified column.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetName( _  
    ByVal i As Integer _  
    ) As String _  
-  
    Implements IDataRecord.GetName
```

**Syntax: C#**

```
public string GetName(  
    inti  
    );
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetName

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.1.25. [MySqlDataReader.GetOrdinal](#) Method

Gets the column ordinal, given the name of the column.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetOrdinal( _  
    ByVal name As String _  
    ) As Integer _  
-  
    Implements IDataRecord.GetOrdinal
```

**Syntax: C#**

```
public int GetOrdinal(  
    stringname  
    );
```

**Parameters**

- `name`:

**Return Value****Implements**

IDataRecord.GetOrdinal

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.4.1.13.1.1.1.26. [MySqlDataReader.GetSchemaTable](#) Method

Returns a DataTable that describes the column metadata of the MySqlDataReader.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetSchemaTable() As DataTable _  
-
```

---

```
Implements IDataReader.GetSchemaTable
```

**Syntax: C#**

```
public DataTable GetSchemaTable();
```

**Return Value****Implements**

IDataReader.GetSchemaTable

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.27. [MySQLDataReader.GetString](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function GetString( _  
    ByVal index As Integer _  
) As String _  
—  
Implements IDataRecord.GetString
```

**Syntax: C#**

```
public string GetString(  
    int index  
);
```

**Implements**

IDataRecord.GetString

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.28. [MySQLDataReader.GetTimeSpan](#) Method**Syntax: Visual Basic**

```
Public Function GetTimeSpan( _  
    ByVal index As Integer _  
) As TimeSpan
```

**Syntax: C#**

```
public TimeSpan GetTimeSpan(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.29. [MySQLDataReader.GetUInt16](#) Method**Syntax: Visual Basic**

```
Public Function GetUInt16( _  
    ByVal index As Integer _  
) As UInt16
```

**Syntax: C#**

```
public ushort GetUInt16(  
    int index  
);
```

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.30. [MySqlDataReader.GetUInt32](#) Method**Syntax: Visual Basic**

```
Public Function GetUInt32( _  
    ByVal index As Integer _  
    ) As UInt32
```

**Syntax: C#**

```
public uint GetUInt32(  
    int index  
);
```

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.31. [MySqlDataReader.GetUInt64](#) Method**Syntax: Visual Basic**

```
Public Function GetUInt64( _  
    ByVal index As Integer _  
    ) As UInt64
```

**Syntax: C#**

```
public ulong GetUInt64(  
    int index  
);
```

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.32. [MySqlDataReader.GetValue](#) Method

Gets the value of the specified column in its native format.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetValue( _  
    ByVal i As Integer _  
    ) As Object _  
    Implements IDataRecord.GetValue
```

**Syntax: C#**

```
public object GetValue(  
    int i  
);
```

**Parameters**

- *i*:

**Return Value**

**Implements**

IDataRecord.GetValue

**See Also**[MySQLDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)21.2.7.1.4.1.13.1.1.1.33. [MySQLDataReader.GetValue](#) Method

Gets all attribute columns in the collection for the current row.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetValue( _  
    ByVal values As Object() _  
) As Integer _  
-  
    Implements IDataRecord.GetValue
```

**Syntax: C#**

```
public int GetValue(  
    object[] values  
);
```

**Parameters**

- `values`:

**Return Value****Implements**

IDataRecord.GetValue

**See Also**[MySQLDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)21.2.7.1.4.1.13.1.1.1.34. [MySQLDataReader.IsDBNull](#) Method

Gets a value indicating whether the column contains non-existent or missing values.

**Syntax: Visual Basic**

```
NotOverridable Public Function IsDBNull( _  
    ByVal i As Integer _  
) As Boolean _  
-  
    Implements IDataRecord.IsDBNull
```

**Syntax: C#**

```
public bool IsDBNull(  
    int i  
);
```

**Parameters**

- `i`:

**Return Value****Implements**



IDataRecord.IsDBNull

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.35. [MySqlDataReader.NextResult](#) Method

Advances the data reader to the next result, when reading the results of batch SQL statements.

**Syntax: Visual Basic**

```
NotOverridable Public Function NextResult() As Boolean _  
-  
    Implements IDataReader.NextResult
```

**Syntax: C#**

```
public bool NextResult();
```

**Return Value**

**Implements**

IDataReader.NextResult

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.1.1.1.36. [MySqlDataReader.Read](#) Method

Advances the MySqlDataReader to the next record.

**Syntax: Visual Basic**

```
NotOverridable Public Function Read() As Boolean _  
-  
    Implements IDataReader.Read
```

**Syntax: C#**

```
public bool Read();
```

**Return Value**

**Implements**

IDataReader.Read

**See Also**

[MySqlDataReader Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.13.2. [MySqlCommand.ExecuteReader](#) Method

**Syntax: Visual Basic**

```
Overloads Public Function ExecuteReader( _  
    ByVal behavior As CommandBehavior _  
) As MySqlDataReader
```

**Syntax: C#**

```
public MySqlDataReader ExecuteReader(  
    CommandBehaviorbehavior
```

```
);
```

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLCommand.ExecuteReader Overload List](#)

21.2.7.1.4.1.14. [MySQLCommand.ExecuteScalar](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Function ExecuteScalar() As Object _  
- Implements IDbCommand.ExecuteScalar
```

**Syntax: C#**

```
public object ExecuteScalar();
```

**Implements**

IDbCommand.ExecuteScalar

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.4.1.15. [MySQLCommand.Prepare](#) Method**Syntax: Visual Basic**

```
NotOverridable Public Sub Prepare() _  
- Implements IDbCommand.Prepare
```

**Syntax: C#**

```
public void Prepare();
```

**Implements**

IDbCommand.Prepare

**See Also**

[MySQLCommand Class](#), [MySQL.Data.MySqlClient Namespace](#)

21.2.7.1.5. [MySQLCommandBuilder](#) Class

For a list of all members of this type, see [MySQLCommandBuilder Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLCommandBuilder_  
- Inherits Component
```

**Syntax: C#**

```
public sealed class MySQLCommandBuilder : Component
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlCommandBuilder Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1. MySqlCommandBuilder Members

[MySqlCommandBuilder overview](#)

#### Public Static (Shared) Methods

<a href="#">DeriveParameters</a>	Overloaded. Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.
----------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

#### Public Instance Constructors

<a href="#">MySqlCommandBuilder</a>	Overloaded. Initializes a new instance of the MySqlCommandBuilder class.
-------------------------------------	--------------------------------------------------------------------------

#### Public Instance Properties

Container (inherited from Component)	Gets the IContainer that contains the Component.
<a href="#">DataAdapter</a>	
<a href="#">QuotePrefix</a>	
<a href="#">QuoteSuffix</a>	
Site (inherited from Component)	Gets or sets the ISite of the Component.

#### Public Instance Methods

CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Dispose (inherited from Component)	Releases all resources used by the Component.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetDeleteCommand</a>	
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetInsertCommand</a>	
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType (inherited from Object)	Gets the Type of the current instance.
<a href="#">GetUpdateCommand</a>	
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.

<a href="#">RefreshSchema</a>	
ToString (inherited from Component)	Returns a String containing the name of the Component, if any. This method should not be overridden.

### Public Instance Events

Disposed (inherited from Component)	Adds an event handler to listen to the Disposed event on the component.
-------------------------------------	-------------------------------------------------------------------------

### See Also

[MySQLCommandBuilder Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.1. DeriveParameters Method

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

### Overload List

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

- [public static void DeriveParameters\(MySqlCommand\);](#)
- [public static void DeriveParameters\(MySqlCommand,bool\);](#)

### See Also

[MySQLCommandBuilder Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.1.1. [MySQLCommandBuilder.DeriveParameters](#) Method

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

### Syntax: Visual Basic

```
Overloads Public Shared Sub DeriveParameters( _
    ByVal command As MySqlCommand _
)
```

### Syntax: C#

```
public static void DeriveParameters(
    MySqlCommand command
);
```

### Parameters

- [command](#): The MySqlCommand referencing the stored procedure from which the parameter information is to be derived. The derived parameters are added to the Parameters collection of the MySqlCommand.

### Exceptions

Exception Type	Condition
----------------	-----------

InvalidOperationException	The command text is not a valid stored procedure name.
---------------------------	--------------------------------------------------------

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#),  
 [MySqlCommandBuilder.DeriveParameters Overload List](#)

21.2.7.1.5.1.1.2.  [MySqlCommandBuilder.DeriveParameters](#) Method**Syntax: Visual Basic**

```
Overloads Public Shared Sub DeriveParameters( _  
    ByVal command As MySqlCommand, _  
    ByVal useProc As Boolean _  
)
```

**Syntax: C#**

```
public static void DeriveParameters(  
    MySqlCommand command,  
    bool useProc  
);
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#),  
 [MySqlCommandBuilder.DeriveParameters Overload List](#)

21.2.7.1.5.1.2.  [MySqlCommandBuilder](#) Constructor

Initializes a new instance of the  [MySqlCommandBuilder](#) class.

**Overload List**

Initializes a new instance of the  [MySqlCommandBuilder](#) class.

- [public MySqlCommandBuilder\(\);](#)
- [public MySqlCommandBuilder\(MySqlDataAdapter\);](#)
- [public MySqlCommandBuilder\(MySqlDataAdapter,bool\);](#)
- [public MySqlCommandBuilder\(bool\);](#)

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.1.  [MySqlCommandBuilder](#) Constructor

Initializes a new instance of the  [MySqlCommandBuilder](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlCommandBuilder();
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#),  [MySqlCommandBuilder Constructor Overload List](#)

21.2.7.1.5.1.2.2. [MySQLCommandBuilder](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal adapter As MySQLDataAdapter _
)
```

**Syntax: C#**

```
public MySQLCommandBuilder(
    MySQLDataAdapter adapter
);
```

**See Also**

[MySQLCommandBuilder Class](#), [MySQL.Data.MySQLClient Namespace](#), [MySQLCommandBuilder Constructor Overload List](#)

21.2.7.1.5.1.2.2.1. [MySQLDataAdapter](#) Class

For a list of all members of this type, see [MySQLDataAdapter Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLDataAdapter_
    Inherits DbDataAdapter
```

**Syntax: C#**

```
public sealed class MySQLDataAdapter : DbDataAdapter
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySQLClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySQLDataAdapter Members](#), [MySQL.Data.MySQLClient Namespace](#)

21.2.7.1.5.1.2.2.1.1. [MySQLDataAdapter](#) Members

[MySQLDataAdapter overview](#)

**Public Instance Constructors**

<a href="#">MySQLDataAdapter</a>	Overloaded. Initializes a new instance of the <a href="#">MySQLDataAdapter</a> class.
----------------------------------	---------------------------------------------------------------------------------------

**Public Instance Properties**

<a href="#">AcceptChangesDuringFill</a> (inherited from <a href="#">DataAdapter</a> )	Gets or sets a value indicating whether <a href="#">AcceptChanges</a> is called on a <a href="#">DataRow</a> after it is added to the <a href="#">DataTable</a> during any of the Fill operations.
---------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

AcceptChangesDuringUpdate (inherited from DataAdapter)	Gets or sets whether AcceptChanges is called during a Update.
Container (inherited from Component)	Gets the IContainer that contains the Component.
ContinueUpdateOnError (inherited from DataAdapter)	Gets or sets a value that specifies whether to generate an exception when an error is encountered during a row update.
<a href="#">DeleteCommand</a>	Overloaded.
FillLoadOption (inherited from DataAdapter)	Gets or sets the LoadOption that determines how the adapter fills the DataTable from the DbDataReader.
<a href="#">InsertCommand</a>	Overloaded.
MissingMappingAction (inherited from DataAdapter)	Determines the action to take when incoming data does not have a matching table or column.
MissingSchemaAction (inherited from DataAdapter)	Determines the action to take when existing DataSet schema does not match incoming data.
ReturnProviderSpecificTypes (inherited from DataAdapter)	Gets or sets whether the Fill method should return provider-specific values or common CLS-compliant values.
<a href="#">SelectCommand</a>	Overloaded.
Site (inherited from Component)	Gets or sets the ISite of the Component.
TableMappings (inherited from DataAdapter)	Gets a collection that provides the master mapping between a source table and a DataTable.
UpdateBatchSize (inherited from DbDataAdapter)	Gets or sets a value that enables or disables batch processing support, and specifies the number of commands that can be executed in a batch.
<a href="#">UpdateCommand</a>	Overloaded.

### Public Instance Methods

CreateObjRef (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Dispose (inherited from Component)	Releases all resources used by the Component.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
Fill (inherited from DbDataAdapter)	Overloaded. Adds or refreshes rows in the DataSet to match those in the data source using the DataSet name, and creates a DataTable named "Table."
FillSchema (inherited from DbDataAdapter)	Overloaded. Configures the schema of the specified DataTable based on the specified SchemaType.
GetFillParameters (inherited from DbDataAdapter)	Gets the parameters set by the user when executing an SQL <code>SELECT</code> statement.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService (inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.

GetType (inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService (inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
ResetFillLoadOption (inherited from DataAdapter)	Resets FillLoadOption to its default state and causes Fill to honor AcceptChangesDuringFill.
ShouldSerializeAcceptChangesDuringFill (inherited from DataAdapter)	Determines whether the AcceptChangesDuringFill property should be persisted.
ShouldSerializeFillLoadOption (inherited from DataAdapter)	Determines whether the FillLoadOption property should be persisted.
ToString (inherited from Component)	Returns a String containing the name of the Component, if any. This method should not be overridden.
Update (inherited from DbDataAdapter)	Overloaded. Calls the respective INSERT, UPDATE, or DELETE statements for each inserted, updated, or deleted row in the specified DataSet.

### Public Instance Events

Disposed (inherited from Component)	Adds an event handler to listen to the Disposed event on the component.
FillError (inherited from DataAdapter)	Returned when an error occurs during a fill operation.
<a href="#">RowUpdated</a>	Occurs during Update after a command is executed against the data source. The attempt to update is made, so the event fires.
<a href="#">RowUpdating</a>	Occurs during Update before a command is executed against the data source. The attempt to update is made, so the event fires.

### Protected Internal Instance Properties

FillCommandBehavior (inherited from DbDataAdapter)	Gets or sets the behavior of the command used to fill the data adapter.
----------------------------------------------------	-------------------------------------------------------------------------

### See Also

[MySQLDataAdapter Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.2.2.1.1.1. [MySQLDataAdapter](#) Constructor

Initializes a new instance of the [MySQLDataAdapter](#) class.

#### Overload List

Initializes a new instance of the [MySQLDataAdapter](#) class.

- [public MySQLDataAdapter\(\);](#)
- [public MySQLDataAdapter\(MySqlCommand\);](#)
- [public MySQLDataAdapter\(string, MySqlConnection\);](#)
- [public MySQLDataAdapter\(string, string\);](#)

### See Also



[MySqlDataAdapter Class, MySql.Data.MySqlClient Namespace](#)21.2.7.1.5.1.2.2.1.1.1.1. [MySqlDataAdapter](#) Constructor

Initializes a new instance of the [MySqlDataAdapter](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlDataAdapter();
```

**See Also**

[MySqlDataAdapter Class, MySql.Data.MySqlClient Namespace, MySqlDataAdapter Constructor Overload List](#)

21.2.7.1.5.1.2.2.1.1.1.2. [MySqlDataAdapter](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal selectCommand As MySqlCommand _  
)
```

**Syntax: C#**

```
public MySqlDataAdapter(  
    MySqlCommandselectCommand  
);
```

**See Also**

[MySqlDataAdapter Class, MySql.Data.MySqlClient Namespace, MySqlDataAdapter Constructor Overload List](#)

21.2.7.1.5.1.2.2.1.1.1.3. [MySqlDataAdapter](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal selectCommandText As String, _  
    ByVal connection As MySqlConnection _  
)
```

**Syntax: C#**

```
public MySqlDataAdapter(  
    stringselectCommandText,  
    MySqlConnectionconnection  
);
```

**See Also**

[MySqlDataAdapter Class, MySql.Data.MySqlClient Namespace, MySqlDataAdapter Constructor Overload List](#)

21.2.7.1.5.1.2.2.1.1.1.4. [MySqlDataAdapter](#) Constructor**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal selectCommandText As String, _  
    ByVal selectConnString As String _  
)
```

```
)
```

**Syntax: C#**

```
public MySqlDataAdapter(  
    stringselectCommandText,  
    stringselectConnString  
);
```

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlDataAdapter Constructor Overload List](#)

## 21.2.7.1.5.1.2.2.1.1.2. DeleteCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property DeleteCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand DeleteCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.2.2.1.1.3. InsertCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property InsertCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand InsertCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.2.2.1.1.4. SelectCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property SelectCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand SelectCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.2.2.1.1.5. UpdateCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property UpdateCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand UpdateCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.6. [MySqlDataAdapter.RowUpdated](#) Event

Occurs during Update after a command is executed against the data source. The attempt to update is made, so the event fires.

**Syntax: Visual Basic**

```
Public Event RowUpdated As MySqlRowUpdatedEventHandler
```

**Syntax: C#**

```
public event MySqlRowUpdatedEventHandler RowUpdated;
```

**Event Data**

The event handler receives an argument of type [MySqlRowUpdatedEventArgs](#) containing data related to this event. The following [MySqlRowUpdatedEventArgsproperties](#) provide information specific to this event.

Property	Description
<a href="#">Command</a>	Gets or sets the MySqlCommand executed when Update is called.
Errors	Gets any errors generated by the .NET Framework data provider when the Commandwas executed.
RecordsAffected	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.
Row	Gets the DataRow sent through an Update.
RowCount	Gets the number of rows processed in a batch of updated records.
StatementType	Gets the type of SQL statement executed.
Status	Gets the UpdateStatus of the Commandproperty.
TableMapping	Gets the DataTableMapping sent through an Update.

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.6.1. [MySqlRowUpdatedEventHandler](#) Delegate

Represents the method that will handle the RowUpdatedevent of a [MySqlDataAdapter](#).

**Syntax: Visual Basic**

```
Public Delegate Sub MySqlRowUpdatedEventHandler( _
    ByVal sender As Object, _
    ByVal e As MySqlRowUpdatedEventArgs _
)
```

**Syntax: C#**

```
public delegate void MySqlRowUpdatedEventHandler(
    object sender,
    MySqlRowUpdatedEventArgs
```

);

**Requirements**Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**[MySql.Data.MySqlClient Namespace](#)21.2.7.1.5.1.2.2.1.1.6.1.1. [MySqlRowUpdatedEventArgs](#) Class

Provides data for the RowUpdated event. This class cannot be inherited.

For a list of all members of this type, see [MySqlRowUpdatedEventArgs Members](#) .**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlRowUpdatedEventArgs_
    Inherits RowUpdatedEventArgs
```

**Syntax: C#**

```
public sealed class MySqlRowUpdatedEventArgs : RowUpdatedEventArgs
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**[MySqlRowUpdatedEventArgs Members](#), [MySql.Data.MySqlClient Namespace](#)21.2.7.1.5.1.2.2.1.1.6.1.1.1. [MySqlRowUpdatedEventArgs](#) Members[MySqlRowUpdatedEventArgs overview](#)**Public Instance Constructors**

<a href="#">MySqlRowUpdatedEventArgs Constructor</a>	Initializes a new instance of the <a href="#">MySqlRowUpdatedEventArgs</a> class.
------------------------------------------------------	-----------------------------------------------------------------------------------

**Public Instance Properties**

<a href="#">Command</a>	Overloaded. Gets or sets the <a href="#">MySqlCommand</a> executed when Update is called.
Errors (inherited from <a href="#">RowUpdatedEventArgs</a> )	Gets any errors generated by the .NET Framework data provider when the <a href="#">Command</a> was executed.
RecordsAffected (inherited from <a href="#">RowUpdatedEventArgs</a> )	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.
Row (inherited from <a href="#">RowUpdatedEventArgs</a> )	Gets the <a href="#">DataRow</a> sent through an Update.

RowCount (inherited from RowUpdatedEventArgs)	Gets the number of rows processed in a batch of updated records.
StatementType (inherited from RowUpdatedEventArgs)	Gets the type of SQL statement executed.
Status (inherited from RowUpdatedEventArgs)	Gets the UpdateStatus of the Commandproperty.
TableMapping (inherited from RowUpdatedEventArgs)	Gets the DataTableMappingsent through an Update.

### Public Instance Methods

CopyToRows (inherited from RowUpdatedEventArgs)	Overloaded. Copies references to the modified rows into the provided array.
Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType (inherited from Object)	Gets the Type of the current instance.
ToString (inherited from Object)	Returns a String that represents the current Object.

### See Also

[MySQLRowUpdatedEventArgs Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.2.2.1.1.6.1.1.1.1. MySQLRowUpdatedEventArgs Constructor

Initializes a new instance of the MySQLRowUpdatedEventArgs class.

#### Syntax: Visual Basic

```
Public Sub New( _
    ByVal row As DataRow, _
    ByVal command As IDbCommand, _
    ByVal statementType As StatementType, _
    ByVal tableMapping As DataTableMapping _
)
```

#### Syntax: C#

```
public MySQLRowUpdatedEventArgs(
    DataRow row,
    IDbCommand command,
    StatementType statementType,
    DataTableMapping tableMapping
);
```

#### Parameters

- `row`: The DataRow sent through an Update.
- `command`: The IDbCommand executed when Update is called.
- `statementType`: One of the StatementType values that specifies the type of query executed.
- `tableMapping`: The DataTableMapping sent through an Update.

### See Also

[MySQLRowUpdatedEventArgs Class, MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.2.2.1.1.6.1.1.2. Command Property

Gets or sets the MySqlCommand executed when Update is called.

**Syntax: Visual Basic**

```
Overloads Public ReadOnly Property Command As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand Command {get;}
```

**See Also**

[MySqlRowUpdatedEventArgs Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.7. [MySqlDataAdapter.RowUpdating](#) Event

Occurs during Update before a command is executed against the data source. The attempt to update is made, so the event fires.

**Syntax: Visual Basic**

```
Public Event RowUpdating As MySqlRowUpdatingEventHandler
```

**Syntax: C#**

```
public event MySqlRowUpdatingEventHandler RowUpdating;
```

**Event Data**

The event handler receives an argument of type [MySqlRowUpdatingEventArgs](#) containing data related to this event. The following [MySqlRowUpdatingEventArgsproperties](#) provide information specific to this event.

Property	Description
<a href="#">Command</a>	Gets or sets the MySqlCommand to execute when performing the Update.
Errors	Gets any errors generated by the .NET Framework data provider when the Command executes.
Row	Gets the DataRow that will be sent to the server as part of an insert, update, or delete operation.
StatementType	Gets the type of SQL statement to execute.
Status	Gets or sets the UpdateStatus of the Commandproperty.
TableMapping	Gets the DataTableMapping to send through the Update.

**See Also**

[MySqlDataAdapter Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.7.1. [MySqlRowUpdatingEventHandler](#) Delegate

Represents the method that will handle the RowUpdatingevent of a [MySqlDataAdapter](#).

**Syntax: Visual Basic**

```
Public Delegate Sub MySqlRowUpdatingEventHandler( _
```

```

    ByVal sender As Object, _
    ByVal e As MySqlRowUpdatingEventArgs _
)

```

**Syntax: C#**

```

public delegate void MySqlRowUpdatingEventHandler(
    object sender,
    MySqlRowUpdatingEventArgs
);

```

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.7.1.1. [MySqlRowUpdatingEventArgs](#) Class

Provides data for the RowUpdating event. This class cannot be inherited.

For a list of all members of this type, see [MySqlRowUpdatingEventArgs Members](#) .

**Syntax: Visual Basic**

```

NotInheritable Public Class MySqlRowUpdatingEventArgs_
    Inherits RowUpdatingEventArgs

```

**Syntax: C#**

```

public sealed class MySqlRowUpdatingEventArgs : RowUpdatingEventArgs

```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlRowUpdatingEventArgs Members](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.2.2.1.1.7.1.1.1. [MySqlRowUpdatingEventArgs](#) Members

[MySqlRowUpdatingEventArgs overview](#)

**Public Instance Constructors**

<a href="#">MySqlRowUpdatingEventArgs Constructor</a>	Initializes a new instance of the <a href="#">MySqlRowUpdatingEventArgs</a> class.
-------------------------------------------------------	------------------------------------------------------------------------------------

**Public Instance Properties**

<a href="#">Command</a>	Overloaded. Gets or sets the <a href="#">MySqlCommand</a> to execute when performing the Update.
-------------------------	--------------------------------------------------------------------------------------------------

Errors (inherited from RowUpdatingEventArgs)	Gets any errors generated by the .NET Framework data provider when the Command executes.
Row (inherited from RowUpdatingEventArgs)	Gets the DataRow that will be sent to the server as part of an insert, update, or delete operation.
StatementType (inherited from RowUpdatingEventArgs)	Gets the type of SQL statement to execute.
Status (inherited from RowUpdatingEventArgs)	Gets or sets the UpdateStatus of the Commandproperty.
TableMapping (inherited from RowUpdatingEventArgs)	Gets the DataTableMapping to send through the Update.

### Public Instance Methods

Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType (inherited from Object)	Gets the Type of the current instance.
ToString (inherited from Object)	Returns a String that represents the current Object.

### See Also

[MySQLRowUpdatingEventArgs Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.2.2.1.1.7.1.1.1.1. MySQLRowUpdatingEventArgs Constructor

Initializes a new instance of the MySQLRowUpdatingEventArgs class.

### Syntax: Visual Basic

```
Public Sub New( _
    ByVal row As DataRow, _
    ByVal command As IDbCommand, _
    ByVal statementType As StatementType, _
    ByVal tableMapping As DataTableMapping _
)
```

### Syntax: C#

```
public MySQLRowUpdatingEventArgs(
    DataRowrow,
    IDbCommandcommand,
    StatementTypestatementType,
    DataTableMappingtableMapping
);
```

### Parameters

- `row`: The DataRowto Update.
- `command`: The IDbCommandto execute during Update.
- `statementType`: One of the StatementTypevalues that specifies the type of query executed.
- `tableMapping`: The DataTableMappingssent through an Update.

### See Also



[MySqlCommandUpdatingEventArgs Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.2.2.1.1.7.1.1.1.2. Command Property

Gets or sets the MySqlCommand to execute when performing the Update.

##### Syntax: Visual Basic

```
Overloads Public Property Command As MySqlCommand
```

##### Syntax: C#

```
new public MySqlCommand Command {get; set;}
```

##### See Also

[MySqlCommandUpdatingEventArgs Class, MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.5.1.2.3. MySqlCommandBuilder Constructor

##### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal adapter As MySqlDataAdapter, _  
    ByVal lastOneWins As Boolean _  
)
```

##### Syntax: C#

```
public MySqlCommandBuilder(  
    MySqlDataAdapteradapter,  
    boollastOneWins  
);
```

##### See Also

[MySqlCommandBuilder Class, MySql.Data.MySqlClient Namespace](#), [MySqlCommandBuilder Constructor Overload List](#)

#### 21.2.7.1.5.1.2.4. MySqlCommandBuilder Constructor

##### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal lastOneWins As Boolean _  
)
```

##### Syntax: C#

```
public MySqlCommandBuilder(  
    boollastOneWins  
);
```

##### See Also

[MySqlCommandBuilder Class, MySql.Data.MySqlClient Namespace](#), [MySqlCommandBuilder Constructor Overload List](#)

#### 21.2.7.1.5.1.3. DataAdapter Property

##### Syntax: Visual Basic

```
Public Property DataAdapter As MySqlDataAdapter
```

##### Syntax: C#

```
public MySqlDataAdapter DataAdapter {get; set;}
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.4. QuotePrefix Property

**Syntax: Visual Basic**

```
Public Property QuotePrefix As String
```

**Syntax: C#**

```
public string QuotePrefix {get; set;}
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.5. QuoteSuffix Property

**Syntax: Visual Basic**

```
Public Property QuoteSuffix As String
```

**Syntax: C#**

```
public string QuoteSuffix {get; set;}
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.6. MySqlCommandBuilder.GetDeleteCommand Method

**Syntax: Visual Basic**

```
Public Function GetDeleteCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetDeleteCommand();
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.7. MySqlCommandBuilder.GetInsertCommand Method

**Syntax: Visual Basic**

```
Public Function GetInsertCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetInsertCommand();
```

**See Also**

[MySqlCommandBuilder Class](#),  [MySql.Data.MySqlClient Namespace](#)

## 21.2.7.1.5.1.8. MySqlCommandBuilder.GetUpdateCommand Method

**Syntax: Visual Basic**

```
Public Function GetUpdateCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetUpdateCommand();
```

**See Also**

[MySqlCommandBuilder Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.5.1.9. [MySqlCommandBuilder.RefreshSchema](#) Method**Syntax: Visual Basic**

```
Public Sub RefreshSchema()
```

**Syntax: C#**

```
public void RefreshSchema();
```

**See Also**

[MySqlCommandBuilder Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.6. [MySqlException](#) Class

The exception that is thrown when MySQL returns an error. This class cannot be inherited.

For a list of all members of this type, see [MySqlException Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlException_  
    Inherits SystemException
```

**Syntax: C#**

```
public sealed class MySqlException : SystemException
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlException Members](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.6.1. [MySqlException](#) Members

[MySqlException overview](#)

**Public Instance Properties**

Data (inherited from Exception)	Gets a collection of key/value pairs that provide additional, user-defined information about the exception.
HelpLink (inherited from Exception)	Gets or sets a link to the help file associated with this exception.

InnerException (inherited from Exception)	Gets the <a href="#">Exceptioninstance</a> that caused the current exception.
Message (inherited from Exception)	Gets a message that describes the current exception.
<a href="#">Number</a>	Gets a number that identifies the type of error. This number corresponds to the error numbers given in <a href="#">Section C.3, "Server Error Codes and Messages"</a> .
Source (inherited from Exception)	Gets or sets the name of the application or the object that causes the error.
StackTrace (inherited from Exception)	Gets a string representation of the frames on the call stack at the time the current exception was thrown.
TargetSite (inherited from Exception)	Gets the method that throws the current exception.

### Public Instance Methods

Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetBaseException (inherited from Exception)	When overridden in a derived class, returns the Exception that is the root cause of one or more subsequent exceptions.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetObjectData (inherited from Exception)	When overridden in a derived class, sets the SerializationInfo with information about the exception.
GetType (inherited from Exception)	Gets the runtime type of the current instance.
ToString (inherited from Exception)	Creates and returns a string representation of the current exception.

### See Also

[MySQLException Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.6.1.1. Number Property

Gets a number that identifies the type of error.

#### Syntax: Visual Basic

```
Public ReadOnly Property Number As Integer
```

#### Syntax: C#

```
public int Number {get;}
```

### See Also

[MySQLException Class](#), [MySQL.Data.MySqlClient Namespace](#)

#### 21.2.7.1.7. MySQLHelper Class

Helper class that makes it easier to work with the provider.

For a list of all members of this type, see [MySqlHelper Members](#) .

### Syntax: Visual Basic

```
NotInheritable Public Class MySqlHelper
```

### Syntax: C#

```
public sealed class MySqlHelper
```

### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

### Requirements

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlHelper Members](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.7.1. [MySqlHelper](#) Members

[MySqlHelper overview](#)

#### Public Static (Shared) Methods

<a href="#">ExecuteDataRow</a>	Executes a single SQL statement and returns the first row of the resultset. A new <a href="#">MySqlConnection</a> object is created, opened, and closed during this method.
<a href="#">ExecuteDataset</a>	Overloaded. Executes a single SQL statement and returns the resultset in a <a href="#">DataSet</a> . A new <a href="#">MySqlConnection</a> object is created, opened, and closed during this method.
<a href="#">ExecuteNonQuery</a>	Overloaded. Executes a single command against a MySQL database. The <a href="#">MySqlConnection</a> is assumed to be open when the method is called and remains open after the method completes.
<a href="#">ExecuteReader</a>	Overloaded. Executes a single command against a MySQL database.
<a href="#">ExecuteScalar</a>	Overloaded. Execute a single command against a MySQL database.
<a href="#">UpdateDataSet</a>	Updates the given table with data from the given <a href="#">DataSet</a>

#### Public Instance Methods

<a href="#">Equals</a> (inherited from <a href="#">Object</a> )	Determines whether the specified <a href="#">Object</a> is equal to the current <a href="#">Object</a> .
<a href="#">GetHashCode</a> (inherited from <a href="#">Object</a> )	Serves as a hash function for a particular type. <a href="#">GetHashCode</a> is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from <a href="#">Object</a> )	Gets the <a href="#">Type</a> of the current instance.

ToString (inherited from Object)

Returns a String that represents the current Object.

**See Also**[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#)21.2.7.1.7.1.1. [MySQLHelper.ExecuteDataRow](#) Method

Executes a single SQL statement and returns the first row of the resultset. A new [MySQLConnection](#) object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Public Shared Function ExecuteDataRow( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray parms As MySqlParameter() _
) As DataRow
```

**Syntax: C#**

```
public static DataRow ExecuteDataRow(
    string connectionString,
    string commandText,
    params MySqlParameter[] parms
);
```

**Parameters**

- [connectionString](#): Settings to be used for the connection
- [commandText](#): Command to execute
- [parms](#): Parameters to use for the command

**Return Value**

DataRow containing the first row of the resultset

**See Also**[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#)21.2.7.1.7.1.2. [ExecuteDataset](#) Method

Executes a single SQL statement and returns the resultset in a DataSet. The state of the [MySQLConnection](#) object remains unchanged after execution of this method.

**Overload List**

Executes a single SQL statement and returns the resultset in a DataSet. The state of the [MySQLConnection](#) object remains unchanged after execution of this method.

- [public static DataSet ExecuteDataset\(MySqlConnection,string\);](#)

Executes a single SQL statement and returns the resultset in a DataSet. The state of the [MySQLConnection](#) object remains unchanged after execution of this method.

- [public static DataSet ExecuteDataset\(MySqlConnection,string,params MySqlParameter\[\]\);](#)

Executes a single SQL statement and returns the resultset in a DataSet. A new [MySQLConnection](#) object is created, opened, and closed during this method.

- [public static DataSet ExecuteDataset\(string,string\);](#)

Executes a single SQL statement and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

- [public static DataSet ExecuteDataset\(string, string, params MySqlParameter\[\]\)](#);

#### See Also

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#)

#### 21.2.7.1.7.1.2.1. [MySqlHelper.ExecuteDataset](#) Method

Executes a single SQL statement and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String _  
) As DataSet
```

#### Syntax: C#

```
public static DataSet ExecuteDataset(  
    MySqlConnection connection,  
    string commandText  
);
```

#### Parameters

- `connection`: [MySqlConnection](#) object to use
- `commandText`: Command to execute

#### Return Value

DataSet containing the resultset

#### See Also

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteDataset Overload List](#)

#### 21.2.7.1.7.1.2.2. [MySqlHelper.ExecuteDataset](#) Method

Executes a single SQL statement and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As DataSet
```

#### Syntax: C#

```
public static DataSet ExecuteDataset(  
    MySqlConnection connection,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```

#### Parameters

- `connection`: [MySqlConnection](#) object to use

- `commandText`: Command to execute
- `commandParameters`: Parameters to use for the command

**Return Value**

DataSet containing the resultset

**See Also**

[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLHelper.ExecuteDataset Overload List](#)

**21.2.7.1.7.1.2.3. MySQLHelper.ExecuteDataset Method**

Executes a single SQL statement and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connectionString As String, _  
    ByVal commandText As String _  
) As DataSet
```

**Syntax: C#**

```
public static DataSet ExecuteDataset(  
    string connectionString,  
    string commandText  
);
```

**Parameters**

- `connectionString`: Settings to be used for the connection
- `commandText`: Command to execute

**Return Value**

DataSet containing the resultset

**See Also**

[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLHelper.ExecuteDataset Overload List](#)

**21.2.7.1.7.1.2.4. MySQLHelper.ExecuteDataset Method**

Executes a single SQL statement and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As DataSet
```

**Syntax: C#**

```
public static DataSet ExecuteDataset(  
    string connectionString,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```



**Parameters**

- `connectionString`: Settings to be used for the connection
- `commandText`: Command to execute
- `commandParameters`: Parameters to use for the command

**Return Value**

DataSet containing the resultset

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteDataset Overload List](#)

## 21.2.7.1.7.1.3. ExecuteNonQuery Method

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

**Overload List**

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

- `public static int ExecuteNonQuery(MySqlConnection, string, params MySqlParameter[]);`

Executes a single command against a MySQL database. A new [MySqlConnection](#) is created using the [ConnectionString](#) given.

- `public static int ExecuteNonQuery(string, string, params MySqlParameter[]);`

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#)

21.2.7.1.7.1.3.1. [MySqlHelper.ExecuteNonQuery](#) Method

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteNonQuery( _
    ByVal connection As MySqlConnection, _
    ByVal commandText As String, _
    ParamArray commandParameters As MySqlParameter() _
) As Integer
```

**Syntax: C#**

```
public static int ExecuteNonQuery(
    MySqlConnection connection,
    string commandText,
    params MySqlParameter[] commandParameters
);
```

**Parameters**

- `connection`: [MySqlConnection](#) object to use
- `commandText`: SQL statement to be executed
- `commandParameters`: Array of [MySqlParameter](#) objects to use with the command.

## Return Value

## See Also

[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLHelper.ExecuteNonQuery Overload List](#)

### 21.2.7.1.7.1.3.2. [MySQLHelper.ExecuteNonQuery](#) Method

Executes a single command against a MySQL database. A new [MySQLConnection](#) is created using the [ConnectionString](#) given.

## Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteNonQuery( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ParamArray parms As MySqlParameter() _  
    ) As Integer
```

## Syntax: C#

```
public static int ExecuteNonQuery(  
    string connectionString,  
    string commandText,  
    params MySqlParameter[] parms  
);
```

## Parameters

- `connectionString`: [ConnectionString](#) to use
- `commandText`: SQL statement to be executed
- `parms`: Array of [MySqlParameter](#) objects to use with the command.

## Return Value

## See Also

[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLHelper.ExecuteNonQuery Overload List](#)

### 21.2.7.1.7.1.4. [ExecuteReader](#) Method

Executes a single command against a MySQL database.

## Overload List

Executes a single command against a MySQL database.

- [public static MySqlDataReader ExecuteReader\(string,string\);](#)

Executes a single command against a MySQL database.

- [public static MySqlDataReader ExecuteReader\(string,string,params MySqlParameter\[\]\);](#)

## See Also

[MySQLHelper Class](#), [MySQL.Data.MySqlClient Namespace](#)

### 21.2.7.1.7.1.4.1. [MySQLHelper.ExecuteReader](#) Method

Executes a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteReader( _
    ByVal connectionString As String, _
    ByVal commandText As String _
) As MySqlDataReader
```

**Syntax: C#**

```
public static MySqlDataReader ExecuteReader(
    string connectionString,
    string commandText
);
```

**Parameters**

- `connectionString`: Settings to use for this command
- `commandText`: Command text to use

**Return Value**

[MySqlDataReader](#) object ready to read the results of the command

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteReader Overload List](#)

21.2.7.1.7.1.4.2. [MySqlHelper.ExecuteReader](#) Method

Executes a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteReader( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray commandParameters As MySqlParameter() _
) As MySqlDataReader
```

**Syntax: C#**

```
public static MySqlDataReader ExecuteReader(
    string connectionString,
    string commandText,
    params MySqlParameter[] commandParameters
);
```

**Parameters**

- `connectionString`: Settings to use for this command
- `commandText`: Command text to use
- `commandParameters`: Array of [MySqlParameter](#) objects to use with the command

**Return Value**

[MySqlDataReader](#) object ready to read the results of the command

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteReader Overload List](#)

21.2.7.1.7.1.5. [ExecuteScalar](#) Method

Execute a single command against a MySQL database.

## Overload List

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(MySqlConnection,string\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(MySqlConnection,string,params MySqlParameter\[\]\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(string,string\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(string,string,params MySqlParameter\[\]\);](#)

## See Also

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#)

### 21.2.7.1.7.1.5.1. [MySqlHelper.ExecuteScalar](#) Method

Execute a single command against a MySQL database.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteScalar( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String _  
) As Object
```

#### Syntax: C#

```
public static object ExecuteScalar(  
    MySqlConnectionconnection,  
    stringcommandText  
);
```

#### Parameters

- `connection`: [MySqlConnection](#) object to use
- `commandText`: Command text to use for the command

#### Return Value

The first column of the first row in the result set, or a null reference if the result set is empty.

## See Also

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteScalar Overload List](#)

### 21.2.7.1.7.1.5.2. [MySqlHelper.ExecuteScalar](#) Method

Execute a single command against a MySQL database.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteScalar( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As Object
```

**Syntax: C#**

```
public static object ExecuteScalar(
    MySqlConnection connection,
    string commandText,
    params MySqlParameter[] commandParameters
);
```

**Parameters**

- `connection`: [MySqlConnection](#) object to use
- `commandText`: Command text to use for the command
- `commandParameters`: Parameters to use for the command

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteScalar Overload List](#)

21.2.7.1.7.1.5.3. [MySqlHelper.ExecuteScalar](#) Method

Execute a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connectionString As String, _
    ByVal commandText As String _
) As Object
```

**Syntax: C#**

```
public static object ExecuteScalar(
    string connectionString,
    string commandText
);
```

**Parameters**

- `connectionString`: Settings to use for the update
- `commandText`: Command text to use for the update

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteScalar Overload List](#)

21.2.7.1.7.1.5.4. [MySqlHelper.ExecuteScalar](#) Method

Execute a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray commandParameters As MySqlParameter() _
) As Object
```

**Syntax: C#**

```
public static object ExecuteScalar(  
    string connectionString,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```

**Parameters**

- `connectionString`: Settings to use for the command
- `commandText`: Command text to use for the command
- `commandParameters`: Parameters to use for the command

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlHelper.ExecuteScalar Overload List](#)

**21.2.7.1.7.1.6. `MySqlHelper.UpdateDataSet` Method**

Updates the given table with data from the given DataSet

**Syntax: Visual Basic**

```
Public Shared Sub UpdateDataSet( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ByVal ds As DataSet, _  
    ByVal tablename As String _  
)
```

**Syntax: C#**

```
public static void UpdateDataSet(  
    string connectionString,  
    string commandText,  
    DataSet ds,  
    string tablename  
);
```

**Parameters**

- `connectionString`: Settings to use for the update
- `commandText`: Command text to use for the update
- `ds`: DataSet containing the new data to use in the update
- `tablename`: Tablename in the data set to update

**See Also**

[MySqlHelper Class](#), [MySql.Data.MySqlClient Namespace](#)

**21.2.7.1.8. `MySqlErrorCode` Enumeration****Syntax: Visual Basic**

```
Public Enum MySqlErrorCode
```

**Syntax: C#**

```
public enum MySQLErrorCode
```

### Members

Member Name	Description
PacketTooLarge	
PasswordNotAllowed	
DuplicateKeyEntry	
HostNotPrivileged	
PasswordNoMatch	
AnonymousUser	
DuplicateKey	
KeyNotFound	
DuplicateKeyName	

### Requirements

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

### See Also

[MySQL.Data.MySqlClient Namespace](#)

## 21.2.7.2. [MySQL.Data.Types](#) Namespace

[Namespace hierarchy](#)

### Classes

Class	Description
<a href="#">MySQLConversionException</a>	Summary description for MySQLConversionException.
<a href="#">MySQLDateTime</a>	Summary description for MySQLDateTime.
<a href="#">MySQLValue</a>	

### 21.2.7.2.1. [MySQL.Data.TypesHierarchy](#)

#### See Also

[MySQL.Data.Types Namespace](#)

### 21.2.7.2.2. [MySQLConversionException](#) Class

Summary description for MySQLConversionException.

For a list of all members of this type, see [MySQLConversionException Members](#) .

#### Syntax: Visual Basic

```
Public Class MySQLConversionException_  
    Inherits ApplicationException
```

#### Syntax: C#

```
public class MySQLConversionException : ApplicationException
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.Types](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlConversionException Members](#), [MySql.Data.Types Namespace](#)

**21.2.7.2.2.1. [MySqlConversionException](#) Members**

[MySqlConversionException overview](#)

**Public Instance Constructors**

<a href="#">MySqlConversionException Constructor</a>	Ctor
------------------------------------------------------	------

**Public Instance Properties**

Data (inherited from Exception)	Gets a collection of key/value pairs that provide additional, user-defined information about the exception.
HelpLink (inherited from Exception)	Gets or sets a link to the help file associated with this exception.
InnerException (inherited from Exception)	Gets the <a href="#">Exceptioninstance</a> that caused the current exception.
Message (inherited from Exception)	Gets a message that describes the current exception.
Source (inherited from Exception)	Gets or sets the name of the application or the object that causes the error.
StackTrace (inherited from Exception)	Gets a string representation of the frames on the call stack at the time the current exception was thrown.
TargetSite (inherited from Exception)	Gets the method that throws the current exception.

**Public Instance Methods**

Equals (inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetBaseException (inherited from Exception)	When overridden in a derived class, returns the Exception that is the root cause of one or more subsequent exceptions.
GetHashCode (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetObjectData (inherited from Exception)	When overridden in a derived class, sets the SerializationInfo with information about the exception.
GetType (inherited from Exception)	Gets the runtime type of the current instance.



ToString (inherited from Exception)	Creates and returns a string representation of the current exception.
-------------------------------------	-----------------------------------------------------------------------

**Protected Instance Properties**

HResult (inherited from Exception)	Gets or sets <a href="#">HRESULT</a> , a coded numeric value that is assigned to a specific exception.
------------------------------------	--------------------------------------------------------------------------------------------------------

**Protected Instance Methods**

Finalize (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone (inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySQLConversionException Class](#), [MySQL.Data.Types Namespace](#)

**21.2.7.2.2.1.1. MySQLConversionException Constructor****Syntax: Visual Basic**

```
Public Sub New( _
    ByVal msg As String _
)
```

**Syntax: C#**

```
public MySQLConversionException(
    string msg
);
```

**See Also**

[MySQLConversionException Class](#), [MySQL.Data.Types Namespace](#)

**21.2.7.2.3. MySQLDateTime Class**

Summary description for MySQLDateTime.

For a list of all members of this type, see [MySQLDateTime Members](#) .

**Syntax: Visual Basic**

```
Public Class MySQLDateTime_
    Inherits MySQLValue_
    Implements IConvertible, IComparable
```

**Syntax: C#**

```
public class MySQLDateTime : MySQLValue, IConvertible, IComparable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.Types](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySqlDateTime Members](#), [MySql.Data.Types Namespace](#)

**21.2.7.2.3.1. MySqlDateTime Members**

[MySqlDateTime overview](#)

**Public Static (Shared) Type Conversions**

<a href="#">Explicit MySqlDateTime to DateTime Conversion</a>	
---------------------------------------------------------------	--

**Public Instance Properties**

<a href="#">Day</a>	Returns the day portion of this datetime
<a href="#">Hour</a>	Returns the hour portion of this datetime
<a href="#">IsNull</a> (inherited from MySqlValue)	
<a href="#">IsValidDateTime</a>	Indicates if this object contains a value that can be represented as a DateTime
<a href="#">Minute</a>	Returns the minute portion of this datetime
<a href="#">Month</a>	Returns the month portion of this datetime
<a href="#">Second</a>	Returns the second portion of this datetime
<a href="#">Millisecond</a>	Returns the millisecond portion of this datetime
<a href="#">ValueAsObject</a> (inherited from MySqlValue)	Returns the value of this field as an object
<a href="#">Year</a>	Returns the year portion of this datetime

**Public Instance Methods**

<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetDateTime</a>	Returns this value as a DateTime
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.
<a href="#">ToString</a>	Returns a MySQL-specific string representation of this value

**Protected Instance Fields**

<a href="#">classType</a> (inherited from MySqlValue)	The system type represented by this value
<a href="#">dbType</a> (inherited from MySqlValue)	The generic dbtype of this value
<a href="#">isNull</a> (inherited from MySqlValue)	Is this value null
<a href="#">mysqlDbType</a> (inherited from MySqlValue)	The specific MySQL db type
<a href="#">mysqlTypeName</a> (inherited from MySqlValue)	The MySQL-specific typename of this value
<a href="#">objectValue</a> (inherited from MySqlValue)	

**Protected Instance Methods**

<a href="#">Finalize</a> (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
--------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------

MemberwiseClone (inherited from Object)	Creates a shallow copy of the current Object.
-----------------------------------------	-----------------------------------------------

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

21.2.7.2.3.1.1. [MySQLDateTime](#) Explicit [MySQLDateTime](#) to [DateTime](#) Conversion**Syntax: Visual Basic**

```
MySQLDateTime.op_Explicit(val)
```

**Syntax: C#**

```
public static explicit operator DateTime(  
MySQLDateTime val  
)  
;
```

**Parameters**

- `val`:

**Return Value****See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

## 21.2.7.2.3.1.2. Day Property

Returns the day portion of this datetime

**Syntax: Visual Basic**

```
Public Property Day As Integer
```

**Syntax: C#**

```
public int Day {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

## 21.2.7.2.3.1.3. Hour Property

Returns the hour portion of this datetime

**Syntax: Visual Basic**

```
Public Property Hour As Integer
```

**Syntax: C#**

```
public int Hour {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

## 21.2.7.2.3.1.4. IsNull Property

**Syntax: Visual Basic**

```
Public Property IsNull As Boolean
```

**Syntax: C#**

```
public bool IsNull {get; set;}
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1. [MySqlValue](#) Class

For a list of all members of this type, see [MySqlValue Members](#) .

**Syntax: Visual Basic**

```
MustInherit Public Class MySqlValue
```

**Syntax: C#**

```
public abstract class MySqlValue
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.Types](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlValue Members](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1. [MySqlValue](#) Members

[MySqlValue overview](#)

**Protected Static (Shared) Fields**

<a href="#">numberFormat</a>	
------------------------------	--

**Public Instance Constructors**

<a href="#">MySqlValue Constructor</a>	Initializes a new instance of the <a href="#">MySqlValue</a> class.
----------------------------------------	---------------------------------------------------------------------

**Public Instance Properties**

<a href="#">IsNull</a>	
<a href="#">ValueAsObject</a>	Returns the value of this field as an object

**Public Instance Methods**

<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.

<a href="#">ToString</a>	Returns a string representation of this value
--------------------------	-----------------------------------------------

**Protected Instance Fields**

<a href="#">classType</a>	The system type represented by this value
<a href="#">dbType</a>	The generic dbtype of this value
<a href="#">isNull</a>	Is this value null
<a href="#">mysqlDbType</a>	The specific MySQL db type
<a href="#">mysqlTypeName</a>	The MySQL-specific typename of this value
<a href="#">objectValue</a>	

**Protected Instance Methods**

Finalize (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone (inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.1. [MySqlValue.numberFormat](#) Field**Syntax: Visual Basic**

```
Protected Shared numberFormat As NumberFormatInfo
```

**Syntax: C#**

```
protected static NumberFormatInfo numberFormat;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.2. [MySqlValue](#) Constructor

Initializes a new instance of the [MySqlValue](#) class.

**Syntax: Visual Basic**

```
Public Sub New()
```

**Syntax: C#**

```
public MySqlValue();
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.3. [ValueAsObject](#) Property

Returns the value of this field as an object

**Syntax: Visual Basic**

```
Public ReadOnly Property ValueAsObject As Object
```

**Syntax: C#**

```
public object ValueAsObject {get;}
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.4. [MySqlValue.ToString](#) Method

Returns a string representation of this value

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.5. [MySqlValue.classType](#) Field

The system type represented by this value

**Syntax: Visual Basic**

```
Protected classType As Type
```

**Syntax: C#**

```
protected Type classType;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.6. [MySqlValue.dbType](#) Field

The generic dbtype of this value

**Syntax: Visual Basic**

```
Protected dbType As DbType
```

**Syntax: C#**

```
protected DbType dbType;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.7. [MySqlValue.mysqlDbType](#) Field

The specific MySQL db type

**Syntax: Visual Basic**

```
Protected mysqlDbType As MySqlDbType
```

**Syntax: C#**

```
protected MySqlDbType mysqlDbType;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.8. [MySqlValue.mySqlTypeName](#) Field

The MySQL-specific typename of this value

**Syntax: Visual Basic**

```
Protected mysqlTypeName As String
```

**Syntax: C#**

```
protected string mysqlTypeName;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.4.1.1.9. [MySqlValue.objectValue](#) Field**Syntax: Visual Basic**

```
Protected objectValue As Object
```

**Syntax: C#**

```
protected object objectValue;
```

**See Also**

[MySqlValue Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.5. [IsValidDateTime](#) Property

Indicates if this object contains a value that can be represented as a DateTime

**Syntax: Visual Basic**

```
Public ReadOnly Property IsValidDateTime As Boolean
```

**Syntax: C#**

```
public bool IsValidDateTime {get;}
```

**See Also**

[MySqlDateTime Class](#), [MySql.Data.Types Namespace](#)

21.2.7.2.3.1.6. [Millisecond](#) Property

Returns the millisecond portion of this datetime

**Syntax: Visual Basic**

```
Public Property Millisecond As Integer
```

**Syntax: C#**

```
public int Millisecond {get; set;}
```

**See Also**

[MySqlDateTime Class](#), [MySql.Data.Types Namespace](#)

#### 21.2.7.2.3.1.7. Minute Property

Returns the minute portion of this datetime

**Syntax: Visual Basic**

```
Public Property Minute As Integer
```

**Syntax: C#**

```
public int Minute {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

#### 21.2.7.2.3.1.8. Month Property

Returns the month portion of this datetime

**Syntax: Visual Basic**

```
Public Property Month As Integer
```

**Syntax: C#**

```
public int Month {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

#### 21.2.7.2.3.1.9. Second Property

Returns the second portion of this datetime

**Syntax: Visual Basic**

```
Public Property Second As Integer
```

**Syntax: C#**

```
public int Second {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

#### 21.2.7.2.3.1.10. Year Property

Returns the year portion of this datetime

**Syntax: Visual Basic**

```
Public Property Year As Integer
```

**Syntax: C#**

```
public int Year {get; set;}
```

**See Also**

[MySQLDateTime Class](#), [MySQL.Data.Types Namespace](#)

#### 21.2.7.2.3.1.11. [MySQLDateTime.GetDateTime](#) Method

Returns this value as a DateTime



**Syntax: Visual Basic**

```
Public Function GetDateTime() As Date
```

**Syntax: C#**

```
public DateTime GetDateTime();
```

**See Also**

[MySqlDateTime Class](#), [MySql.Data.Types Namespace](#)

**21.2.7.2.3.1.12. `MySqlDateTime.ToString` Method**

Returns a MySQL-specific string representation of this value

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**See Also**

[MySqlDateTime Class](#), [MySql.Data.Types Namespace](#)

## 21.2.8. Connector/Net Support

The developers of Connector/Net greatly value the input of our users in the software development process. If you find Connector/Net lacking some feature important to you, or if you discover a bug and need to file a bug report, please use the instructions in [Section 1.7, "How to Report Bugs or Problems"](#).

### 21.2.8.1. Connector/Net Community Support

- Community support for Connector/Net can be found through the forums at <http://forums.mysql.com>.
- Community support for Connector/Net can also be found through the mailing lists at <http://lists.mysql.com>.
- Paid support is available from Oracle. Additional information is available at <http://dev.mysql.com/support/>.

### 21.2.8.2. How to Report Connector/Net Problems or Bugs

If you encounter difficulties or problems with Connector/Net, contact the Connector/Net community, as explained in [Section 21.2.8.1, "Connector/Net Community Support"](#).

First try to execute the same SQL statements and commands from the `mysql` client program or from `admindemo`. This helps you determine whether the error is in Connector/Net or MySQL.

If reporting a problem, ideally include the following information with the email:

- Operating system and version.
- Connector/Net version.
- MySQL server version.
- Copies of error messages or other unexpected output.
- Simple reproducible sample.

Remember that the more information you can supply to us, the more likely it is that we can fix the problem.

If you believe the problem to be a bug, then you must report the bug through <http://bugs.mysql.com/>.

## 21.2.9. Connector/Net FAQ

### Questions

- [22.2.9.1: \[2634\]](#) How do I obtain the value of an auto-incremented column?

### Questions and Answers

#### 22.2.9.1: How do I obtain the value of an auto-incremented column?

When using `CommandBuilder`, setting `ReturnGeneratedIdentifiers` property to `true` no longer works, as `CommandBuilder` does not add `last_insert_id()` by default.

`CommandBuilder` hooks up to the `DataAdapter.RowUpdating` event handler, which means it will get called for every row. It examines the command object and, if it is the same referenced object, it essentially rebuilds the object, thereby destroying your command text changes.

One approach to solving this problem is to clone the command object so you have a different actual reference:

```
dataAdapter.InsertCommand = cb.GetInsertCommand().Clone()
```

This will work, but since the `CommandBuilder` is still connected to the `DataAdapter`, the `RowUpdating` event will still fire and performance will be hit. To stop that, once all your commands have been added you need to disconnect the `CommandBuilder` from the `DataAdapter`:

```
cb.DataAdapter = null;
```

The last requirement is to make sure the `id` that is returned by `last_insert_id()` has the correct name. For example:

```
SELECT last_insert_id() AS id
```

A complete working example is shown here:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

using System.Data;
using MySql.Data;
using MySql.Data.MySqlClient;

namespace GetAutoIncId
{
    class Program
    {
        static void Main(string[] args)
        {
            string connStr = "server=localhost;user=root;database=TestDB;port=3306;password=*****";
            MySqlConnection conn = new MySqlConnection(connStr);

            try
            {
                Console.WriteLine("Connecting to MySQL...");
                conn.Open();

                string sql = "SELECT * FROM TestTable";

                MySqlDataAdapter da = new MySqlDataAdapter(sql, conn);
```

```

MySQLCommandBuilder cb = new MySQLCommandBuilder(da);

MySQLCommand cmd = new MySQLCommand();
cmd.Connection = conn;
cmd.CommandText = sql;
// use Cloned object to avoid .NET rebuilding the object, and
// thereby throwing away our command text additions.
MySQLCommand insertCmd = cb.GetInsertCommand().Clone();
insertCmd.CommandText = insertCmd.CommandText + ";SELECT last_insert_id() AS id";
insertCmd.UpdatedRowSource = UpdateRowSource.FirstReturnedRecord;
da.InsertCommand = insertCmd;
cb.DataAdapter = null; // Unhook RowUpdating event handler

DataTable dt = new DataTable();
da.Fill(dt);

DataRow row = dt.NewRow();
row["name"] = "Joe Smith";

dt.Rows.Add(row);
da.Update(dt);

System.Console.WriteLine("ID after update: " + row["id"]);
}
catch (Exception ex)
{
    Console.WriteLine(ex.ToString());
}

conn.Close();
Console.WriteLine("Done.");
}
}
}

```

## 21.3. MySQL Connector/J

This section explains how to configure and develop Java applications with MySQL Connector/J, the JDBC driver that is integrated with MySQL.

For release notes detailing the changes in each release of Connector/J, see [MySQL Connector/J Release Notes](#).

### 21.3.1. Overview of MySQL Connector/J

MySQL provides connectivity for client applications developed in the Java programming language through a JDBC driver, which is called MySQL Connector/J.

MySQL Connector/J is a JDBC Type 4 driver. Different versions are available that are compatible with the JDBC 3.0 and JDBC 4.0 specifications. The Type 4 designation means that the driver is a pure Java implementation of the MySQL protocol and does not rely on the MySQL client libraries.

For large-scale programs that use common design patterns of data access, consider using one of the popular persistence frameworks such as [Hibernate](#), [Spring's JDBC templates](#) or [Ibatis SQL Maps](#) to reduce the amount of JDBC code for you to debug, tune, secure, and maintain.

#### Key Topics

- For help with connection strings, connection options, and setting up your connection through JDBC, see [Section 21.3.5.1, "Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J"](#).

### 21.3.2. Connector/J Versions

There are currently four versions of MySQL Connector/J available:

- Connector/J 5.1 is the Type 4 pure Java JDBC driver, which conforms to the JDBC 3.0 and JDBC 4.0 specifications. It provides compatibility with all the functionality of MySQL, including 4.1, 5.0, 5.1, 5.4 and 5.5. Connector/J 5.1 provides ease of development features, including auto-registration with the Driver Manager, standardized validity checks, categorized SQLExceptions, support for the JDBC-4.0 XML processing, per connection client information, [NCHAR](#), [NVARCHAR](#) and [NCLOB](#) types. This release also includes all bug fixes up to and including Connector/J 5.0.6.
- Connector/J 5.0 provides support for all the functionality offered by Connector/J 3.1 and includes distributed transaction (XA) support.
- Connector/J 3.1 was designed for connectivity to MySQL 4.1 and MySQL 5.0 servers and provides support for all the functionality in MySQL 5.0 except distributed transaction (XA) support.
- Connector/J 3.0 provides core functionality and was designed for connectivity to MySQL 3.x or MySQL 4.1 servers, although it provides basic compatibility with later versions of MySQL. Connector/J 3.0 does not support server-side prepared statements, and does not support any of the features in versions of MySQL later than 4.1.

The following table summarizes the Connector/J versions available, along with the details of JDBC driver type, what version of the JDBC API it supports, what versions of MySQL Server it works with, and whether it is currently supported or not:

**Table 21.22. Summary of Connector/J Versions**

Connector/J version	Driver Type	JDBC version	MySQL Server version	Status
5.1	4	3.0, 4.0	4.1, 5.0, 5.1, 5.4, 5.5	Recommended version
5.0	4	3.0	4.1, 5.0	Released version
3.1	4	3.0	4.1, 5.0	Obsolete
3.0	4	3.0	3.x, 4.1	Obsolete

The current recommended version for Connector/J is 5.1. This guide covers all four connector versions, with specific notes given where a setting applies to a specific option.

### 21.3.2.1. Connector/J Release Notes and Change History

For details of new features and bug fixes in each Connector/J release, see the [MySQL Connector/J Release Notes](#).

### 21.3.2.2. Java Versions Supported

The following table summarizes what version of Java RTE is required to use Connector/J with Java applications, and what version of JDK is required to build Connector/J source code:

**Table 21.23. Summary of Java Versions Required by Connector/J**

Connector/J version	Java RTE required	JDK required (to build source code)
5.1	1.5.x, 1.6.x, 1.7.x	1.6.x and 1.5.x
5.0	1.3.x, 1.4.x, 1.5.x, 1.6.x	1.4.2, 1.5.x, 1.6.x
3.1	1.2.x, 1.3.x, 1.4.x, 1.5.x, 1.6.x	1.4.2, 1.5.x, 1.6.x
3.0	1.2.x, 1.3.x, 1.4.x, 1.5.x, 1.6.x	1.4.2, 1.5.x, 1.6.x

If you are building Connector/J from source code using the source distribution (see [Section 21.3.3.4, "Installing from the Development Source Tree"](#)), you must use JDK 1.4.2 or newer to compile the Connector package. For Connector/J 5.1, you must have both JDK-1.6.x and JDK-1.5.x installed to be able to build the source code.

Java 1.7 support requires Connector/J 5.1.21 and higher. Several JDBC 4.1 methods were implemented for the first time in Connector/J 5.1.21.

Because of the implementation of `java.sql.Savepoint`, Connector/J 3.1.0 and newer will not run on a Java runtime older than 1.4 unless the class verifier is turned off (by setting the `-Xverify:none` option to the Java runtime). This is because the class verifier will try to load the class definition for `java.sql.Savepoint` even though it is not accessed by the driver unless you actually use savepoint functionality.

Caching functionality provided by Connector/J 3.1.0 or newer is also not available on JVMs older than 1.4.x, as it relies on `java.util.LinkedHashMap` which was first available in JDK-1.4.0.

MySQL Connector/J does not support JDK-1.1.x or JDK-1.0.x.

### 21.3.3. Connector/J Installation

You can install the Connector/J package using either the binary or source distribution. The binary distribution provides the easiest method for installation; the source distribution lets you customize your installation further. With either solution, you manually add the Connector/J location to your Java `CLASSPATH`.

If you are upgrading from a previous version, read the upgrade information in [Section 21.3.3.3, “Upgrading from an Older Version”](#) before continuing.

Connector/J is also available as part of the Maven project. For more information, and to download the Connector/J JAR files, see the [Maven repository](#).

#### 21.3.3.1. Installing Connector/J from a Binary Distribution

For the easiest method of installation, use the binary distribution of the Connector/J package. The binary distribution is available either as a tar/gzip or zip file. Extract it to a suitable location, then optionally make the information about the package available by changing your `CLASSPATH` (see [Section 21.3.3.2, “Installing the Driver and Configuring the CLASSPATH”](#)).

MySQL Connector/J is distributed as a `.zip` or `.tar.gz` archive containing the sources, the class files, and the JAR archive named `mysql-connector-java-version-bin.jar`.

Starting with Connector/J 3.1.9, the `.class` files that constitute the JAR files are only included as part of the driver JAR file.

Starting with Connector/J 3.1.8, the archive also includes a debug build of the driver in a file named `mysql-connector-java-version-bin-g.jar`. Do not use the debug build of the driver unless instructed to do so when reporting a problem or a bug, as it is not designed to be run in production environments, and will have adverse performance impact when used. The debug binary also depends on the Aspect/J runtime library, which is located in the `src/lib/aspectjrt.jar` file that comes with the Connector/J distribution.

Use the appropriate graphical or command-line utility to extract the distribution (for example, WinZip for the `.zip` archive, and `tar` for the `.tar.gz` archive). Because there are potentially long file names in the distribution, we use the GNU tar archive format. Use GNU tar (or an application that understands the GNU tar archive format) to unpack the `.tar.gz` variant of the distribution.

#### 21.3.3.2. Installing the Driver and Configuring the CLASSPATH

Once you have extracted the distribution archive, you can install the driver by placing `mysql-connector-java-version-bin.jar` in your classpath, either by adding the full path to it to your `CLASSPATH` environment variable, or by directly specifying it with the command line switch `-cp` when starting the JVM.

To use the driver with the JDBC `DriverManager`, use `com.mysql.jdbc.Driver` as the class that implements `java.sql.Driver`.

You can set the `CLASSPATH` environment variable under UNIX, Linux or Mac OS X either locally for a user within their `.profile`, `.login` or other login file. You can also set it globally by editing the global `/etc/profile` file.

For example, add the Connector/J driver to your `CLASSPATH` using one of the following forms, depending on your command shell:

```
# Bourne-compatible shell (sh, ksh, bash, zsh):
shell> export CLASSPATH=/path/mysql-connector-java-ver-bin.jar:$CLASSPATH

# C shell (csh, tcsh):
shell> setenv CLASSPATH /path/mysql-connector-java-ver-bin.jar:$CLASSPATH
```

Within Windows 2000, Windows XP, Windows Server 2003 and Windows Vista, you set the environment variable through the System Control Panel.

To use MySQL Connector/J with an application server such as GlassFish, Tomcat or JBoss, read your vendor's documentation for more information on how to configure third-party class libraries, as most application servers ignore the `CLASSPATH` environment variable. For configuration examples for some J2EE application servers, see [Section 21.3.7, "Connection Pooling with Connector/J"](#) [Section 21.3.8, "Load Balancing with Connector/J"](#), and [Section 21.3.9, "Failover with Connector/J"](#). However, the authoritative source for JDBC connection pool configuration information for your particular application server is the documentation for that application server.

If you are developing servlets or JSPs, and your application server is J2EE-compliant, you can put the driver's `.jar` file in the `WEB-INF/lib` subdirectory of your webapp, as this is a standard location for third party class libraries in J2EE web applications.

You can also use the `MysqlDataSource` or `MysqlConnectionPoolDataSource` classes in the `com.mysql.jdbc.jdbc2.optional` package, if your J2EE application server supports or requires them. Starting with Connector/J 5.0.0, the `javax.sql.XADataSource` interface is implemented using the `com.mysql.jdbc.jdbc2.optional.MysqlXADataSource` class, which supports XA distributed transactions when used in combination with MySQL server version 5.0.

The various `MysqlDataSource` classes support the following parameters (through standard set mutators):

- `user`
- `password`
- `serverName` (see the previous section about fail-over hosts)
- `databaseName`
- `port`

### 21.3.3.3. Upgrading from an Older Version

This section has information for users who are upgrading from one version of Connector/J to another, or to a new version of the MySQL server that supports a more recent level of JDBC. A newer version of Connector/J might include changes to support new features, improve existing functionality, or comply with new standards.

#### 21.3.3.3.1. Upgrading to MySQL Connector/J 5.1.x

- In Connector/J 5.0.x and earlier, the alias for a table in a `SELECT` statement is returned when accessing the result set metadata using `ResultSetMetaData.getColumnName()`. This behavior however is not JDBC compliant, and in Connector/J 5.1 this behavior was changed so that the original table name, rather than the alias, is returned.

The JDBC-compliant behavior is designed to let API users reconstruct the DML statement based on the metadata within `ResultSet` and `ResultSetMetaData`.

You can get the alias for a column in a result set by calling `ResultSetMetaData.getColumnLabel()`. To use the old noncompliant behavior with `ResultSetMetaData.getColumnName()`, use the `useOldAliasMetadataBehavior` option and set the value to `true`.

In Connector/J 5.0.x, the default value of `useOldAliasMetadataBehavior` was `true`, but in Connector/J 5.1 this was changed to a default value of `false`.

### 21.3.3.3.2. JDBC-Specific Issues When Upgrading to MySQL Server 4.1 or Newer

- *Using the UTF-8 Character Encoding* - Prior to MySQL server version 4.1, the UTF-8 character encoding was not supported by the server, however the JDBC driver could use it, allowing storage of multiple character sets in `latin1` tables on the server.

Starting with MySQL-4.1, this functionality is deprecated. If you have applications that rely on this functionality, and can not upgrade them to use the official Unicode character support in MySQL server version 4.1 or newer, add the following property to your connection URL:

```
useOldUTF8Behavior=true
```

- *Server-side Prepared Statements* - Connector/J 3.1 will automatically detect and use server-side prepared statements when they are available (MySQL server version 4.1.0 and newer). If your application encounters issues with server-side prepared statements, you can revert to the older client-side emulated prepared statement code that is still presently used for MySQL servers older than 4.1.0 with the following connection property:

```
useServerPrepStmts=false
```

### 21.3.3.3.3. Upgrading from MySQL Connector/J 3.0 to 3.1

Connector/J 3.1 is designed to be backward-compatible with Connector/J 3.0 as much as possible. Major changes are isolated to new functionality exposed in MySQL-4.1 and newer, which includes Unicode character sets, server-side prepared statements, `SQLState` codes returned in error messages by the server and various performance enhancements that can be enabled or disabled using configuration properties.

- **Unicode Character Sets:** See the next section, as well as [Section 10.1, "Character Set Support"](#), for information on this MySQL feature. If you have something misconfigured, it will usually show up as an error with a message similar to `Illegal mix of collations`.
- **Server-side Prepared Statements:** Connector/J 3.1 will automatically detect and use server-side prepared statements when they are available (MySQL server version 4.1.0 and newer).

Starting with version 3.1.7, the driver scans SQL you are preparing using all variants of `Connection.prepareStatement()` to determine if it is a supported type of statement to prepare on the server side, and if it is not supported by the server, it instead prepares it as a client-side emulated prepared statement. You can disable this feature by passing `emulateUnsupportedPstmts=false` in your JDBC URL.

If your application encounters issues with server-side prepared statements, you can revert to the older client-side emulated prepared statement code that is still presently used for MySQL servers older than 4.1.0 with the connection property `useServerPrepStmts=false`.

- **Datetimes** with all-zero components (`0000-00-00 ...`): These values cannot be represented reliably in Java. Connector/J 3.0.x always converted them to `NULL` when being read from a `ResultSet`.

Connector/J 3.1 throws an exception by default when these values are encountered, as this is the most correct behavior according to the JDBC and SQL standards. This behavior can be modified using the `zeroDateTimeBehavior` configuration property. The permissible values are:



- `exception` (the default), which throws an `SQLException` with an `SQLState` of `S1009`.
- `convertToNull`, which returns `NULL` instead of the date.
- `round`, which rounds the date to the nearest closest value which is `0001-01-01`.

Starting with Connector/J 3.1.7, `ResultSet.getString()` can be decoupled from this behavior using `noDatetimeStringSync=true` (the default value is `false`) so that you can retrieve the unaltered all-zero value as a `String`. Note that this also precludes using any time zone conversions, therefore the driver will not allow you to enable `noDatetimeStringSync` and `useTimezone` at the same time.

- **New SQLState Codes:** Connector/J 3.1 uses SQL:1999 `SQLState` codes returned by the MySQL server (if supported), which are different from the legacy X/Open state codes that Connector/J 3.0 uses. If connected to a MySQL server older than MySQL-4.1.0 (the oldest version to return `SQLStates` as part of the error code), the driver will use a built-in mapping. You can revert to the old mapping by using the configuration property `useSqlStateCodes=false`.
- **`ResultSet.getString()`:** Calling `ResultSet.getString()` on a `BLOB` column will now return the address of the `byte[]` array that represents it, instead of a `String` representation of the `BLOB`. `BLOB` values have no character set, so they cannot be converted to `java.lang.Strings` without data loss or corruption.

To store strings in MySQL with LOB behavior, use one of the `TEXT` types, which the driver will treat as a `java.sql.Clob`.

- **Debug builds:** Starting with Connector/J 3.1.8 a debug build of the driver in a file named `mysql-connector-java-version-bin-g.jar` is shipped alongside the normal binary jar file that is named `mysql-connector-java-version-bin.jar`.

Starting with Connector/J 3.1.9, we do not ship the `.class` files unbundled, they are only available in the JAR archives that ship with the driver.

Do not use the debug build of the driver unless instructed to do so when reporting a problem or bug, as it is not designed to be run in production environments, and will have adverse performance impact when used. The debug binary also depends on the Aspect/J runtime library, which is located in the `src/lib/aspectjrt.jar` file that comes with the Connector/J distribution.

#### 21.3.3.4. Installing from the Development Source Tree



##### Caution

Read this section only if you are interested in helping us test our new code. To just get MySQL Connector/J up and running on your system, use a standard binary release distribution.

To install MySQL Connector/J from the development source tree, make sure that you have the following prerequisites:

- A Bazaar client, to check out the sources from our Launchpad repository (available from <http://bazaar-vcs.org/>).
- Apache Ant version 1.7 or newer (available from <http://ant.apache.org/>).
- JDK 1.4.2 or later. Although MySQL Connector/J can be used with older JDKs, compiling it from source requires at least JDK 1.4.2. To build Connector/J 5.1 requires JDK 1.6.x *and* an older JDK such as JDK 1.5.x; point your `JAVA_HOME` environment variable at the older installation.

To check out and compile a specific branch of MySQL Connector/J, follow these steps:

1. Check out the latest code from the branch that you want with one of the following commands.



The source code repository for MySQL Connector/J is located on Launchpad at <https://code.launchpad.net/connectorj>. To check out the latest development branch, use:

```
shell> bazaar branch lp:connectorj
```

This creates a `connectorj` subdirectory in the current directory that contains the latest sources for the requested branch.

To check out the latest 5.1 code, use:

```
shell> bazaar branch lp:connectorj/5.1
```

This creates a `5.1` subdirectory in the current directory containing the latest 5.1 code.

2. To build Connector/J 5.1, make sure that you have both JDK 1.6.x installed and an older JDK such as JDK 1.5.x. This is because Connector/J supports both JDBC 3.0 (which was prior to JDK 1.6.x) and JDBC 4.0. Set your `JAVA_HOME` environment variable to the path of the older JDK installation.
3. Change your current working directory to either the `connectorj` or `5.1` directory, depending on which branch you intend to build.
4. To build Connector/J 5.1, edit the `build.xml` to reflect the location of your JDK 1.6.x installation. The lines to change are:

```
<property name="com.mysql.jdbc.java6.javac" value="C:\jvms\jdk1.6.0\bin\javac.exe" />
<property name="com.mysql.jdbc.java6.rtar" value="C:\jvms\jdk1.6.0\jre\lib\rt.jar" />
```

Alternatively, you can set the value of these property names through the Ant `-D` option.

5. Issue the following command to compile the driver and create a `.jar` file suitable for installation:

```
shell> ant dist
```

This creates a `build` directory in the current directory, where all build output will go. A directory is created in the `build` directory that includes the version number of the sources you are building from. This directory contains the sources, compiled `.class` files, and a `.jar` file suitable for deployment. For other possible targets, including ones that will create a fully packaged distribution, issue the following command:

```
shell> ant -projecthelp
```

6. A newly created `.jar` file containing the JDBC driver will be placed in the directory `build/mysql-connector-java-version`.

Install the newly created JDBC driver as you would a binary `.jar` file that you download from MySQL, by following the instructions in [Section 21.3.3.2, "Installing the Driver and Configuring the CLASSPATH"](#).

A package containing both the binary and source code for Connector/J 5.1 can also be found at the following location: [Connector/J 5.1 Download](#)

## 21.3.4. Connector/J Examples

Examples of using Connector/J are located throughout this document. This section provides a summary and links to these examples.

- [Example 21.1, "Connector/J: Obtaining a connection from the DriverManager"](#)
- [Example 21.2, "Connector/J: Using java.sql.Statement to execute a SELECT query"](#)
- [Example 21.3, "Connector/J: Calling Stored Procedures"](#)
- [Example 21.4, "Connector/J: Using Connection.prepareCall\(\)"](#)

- [Example 21.5, “Connector/J: Registering output parameters”](#)
- [Example 21.6, “Connector/J: Setting `CallableStatement` input parameters”](#)
- [Example 21.7, “Connector/J: Retrieving results and output parameter values”](#)
- [Example 21.8, “Connector/J: Retrieving `AUTO\_INCREMENT` column values using `Statement.getGeneratedKeys\(\)`”](#)
- [Example 21.9, “Connector/J: Retrieving `AUTO\_INCREMENT` column values using `SELECT LAST\_INSERT\_ID\(\)`”](#)
- [Example 21.10, “Connector/J: Retrieving `AUTO\_INCREMENT` column values in `Updatable ResultSets`”](#)
- [Example 21.11, “Connector/J: Using a connection pool with a J2EE application server”](#)
- [Example 21.12, “Connector/J: Example of transaction with retry logic”](#)

## 21.3.5. Connector/J (JDBC) Reference

This section of the manual contains reference material for MySQL Connector/J.

### 21.3.5.1. Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J

The name of the class that implements `java.sql.Driver` in MySQL Connector/J is `com.mysql.jdbc.Driver`. The `org.gjt.mm.mysql.Driver` class name is also usable for backward compatibility with MM.MySQL, the predecessor of Connector/J. Use this class name when registering the driver, or when otherwise configuring software to use MySQL Connector/J.

#### JDBC URL Format

The JDBC URL format for MySQL Connector/J is as follows, with items in square brackets ([, ]) being optional:

```
jdbc:mysql://[host][,failoverhost...][:port]/[database] »
[?propertyName1][=propertyValue1][&propertyName2][=propertyValue2]...
```

If the host name is not specified, it defaults to `127.0.0.1`. If the port is not specified, it defaults to `3306`, the default port number for MySQL servers.

```
jdbc:mysql://[host:port],[host:port].../[database] »
[?propertyName1][=propertyValue1][&propertyName2][=propertyValue2]...
```

Here is a sample connection URL:

```
jdbc:mysql://localhost:3306/sakila?profileSQL=true
```

#### IPv6 Connections

For IPv6 connections, use this alternative syntax to specify hosts in the URL, `address=(key=value)`. Supported keys are:

- `(protocol=tcp)`, or `(protocol=pipe)` for named pipes on Windows.
- `(path=path_to_pipe)` for named pipes.
- `(host=hostname)` for TCP connections.
- `(port=port_number)` for TCP connections.

For example:

```
jdbc:mysql://address=(protocol=tcp)(host=localhost)(port=3306)(user=test)/db
```

Any other parameters are treated as host-specific properties that follow the conventions of the JDBC URL properties. This now allows per-host overrides of any configuration property for multi-host connections (that is, when using failover, load balancing, or replication). Limit the overrides to user, password, network timeouts and statement and metadata cache sizes; the results of other per-host overrides are not defined.

## Initial Database for Connection

If the database is not specified, the connection is made with no default database. In this case, either call the `setCatalog()` method on the Connection instance, or fully specify table names using the database name (that is, `SELECT dbname.tablename.colname FROM dbname.tablename...`) in your SQL. Opening a connection without specifying the database to use is generally only useful when building tools that work with multiple databases, such as GUI database managers.



### Note

Always use the `Connection.setCatalog()` method to specify the desired database in JDBC applications, rather than the `USE database` statement.

## Failover Support

MySQL Connector/J has failover support. This enables the driver to fail over to any number of slave hosts and still perform read-only queries. Failover only happens when the connection is in an `autoCommit(true)` state, because failover cannot happen reliably when a `transaction` is in progress. Most application servers and connection pools set `autoCommit` to `true` at the end of every transaction/connection use.

The failover functionality has the following behavior:

- If the URL property `autoReconnect` is `false`: Failover only happens at connection initialization, and fallback occurs when the driver determines that the first host has become available again.
- If the URL property `autoReconnect` is `true`: Failover happens when the driver determines that the connection has failed (checked before every query), and falls back to the first host when it determines that the host has become available again (after `queriesBeforeRetryMaster` queries have been issued).

In either case, whenever you are connected to a “failed-over” server, the connection is set to read-only state, so queries that attempt to modify data will throw exceptions (the query will **never** be processed by the MySQL server).

## Setting Configuration Properties

Configuration properties define how Connector/J will make a connection to a MySQL server. Unless otherwise noted, properties can be set for a `DataSource` object or for a `Connection` object.

Configuration properties can be set in one of the following ways:

- Using the `set*()` methods on MySQL implementations of `java.sql.DataSource` (which is the preferred method when using implementations of `java.sql.DataSource`):
  - `com.mysql.jdbc.jdbc2.optional.MysqlDataSource`
  - `com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource`
- As a key/value pair in the `java.util.Properties` instance passed to `DriverManager.getConnection()` or `Driver.connect()`
- As a JDBC URL parameter in the URL given to `java.sql.DriverManager.getConnection()`, `java.sql.Driver.connect()` or the MySQL implementations of the `javax.sql.DataSource.setURL()` method.

**Note**

If the mechanism you use to configure a JDBC URL is XML-based, use the XML character literal `&amp;` to separate configuration parameters, as the ampersand is a reserved character for XML.

The properties are listed in the following tables.

**Connection/Authentication.**

Property Name	Definition	Default Value	Since Version
user	The user to connect as		all versions
password	The password to use when connecting		all versions
socketFactory	The name of the class that the driver should use for creating socket connections to the server. This class must implement the interface 'com.mysql.jdbc.SocketFactory' and have public no-args constructor.	com.mysql.jdbc.StandardSocketFactory	
connectTimeout	Timeout for socket connect (in milliseconds), with 0 being no timeout. Only works on JDK-1.4 or newer. Defaults to '0'.	0	3.0.1
socketTimeout	Timeout on network socket operations (0, the default means no timeout).	0	3.0.1
connectionLifecycleInterceptors	A comma-delimited list of classes that implement "com.mysql.jdbc.ConnectionLifecycleInterceptor" that should notified of connection lifecycle events (creation, destruction, commit, rollback, setCatalog and setAutoCommit) and potentially alter the execution of these commands. ConnectionLifecycleInterceptors are "stackable", more than one interceptor may be specified via the configuration property as a comma-delimited list, with the interceptors executed in order from left to right.		5.1.4
useConfigs	Load the comma-delimited list of configuration properties before parsing the URL or applying user-specified properties. These configurations are explained in the 'Configurations' of the documentation.		3.1.5
authenticationPlugins	Comma-delimited list of classes that implement com.mysql.jdbc.AuthenticationPlugin and which will be used for authentication unless disabled by "disabledAuthenticationPlugins" property.		5.1.19
defaultAuthenticationPlugin	Name of a class implementing com.mysql.jdbc.AuthenticationPlugin which will be used as the default authentication plugin (see below). It is an error to use a class which is not listed in "authenticationPlugins" nor it is one of the built-in plugins. It is an error to set as default a plugin which was disabled with "disabledAuthenticationPlugins" property. It is an error to set this value to null or the empty	com.mysql.jdbc.AuthenticationPlugin	5.1.19

	string (i.e. there must be at least a valid default authentication plugin specified for the connection, meeting all constraints listed above).		
disabledAuthenticationPlugins	Comma-delimited list of classes implementing com.mysql.jdbc.AuthenticationPlugin or mechanisms, i.e. "mysql_native_password". The authentication plugins or mechanisms listed will not be used for authentication which will fail if it requires one of them. It is an error to disable the default authentication plugin (either the one named by "defaultAuthenticationPlugin" property or the hard-coded one if "defaultAuthenticationPlugin" property is not set).		5.1.19
disconnectOnExpiredPasswords	If "disconnectOnExpiredPasswords" is set to "false" and password is expired then server enters "sandbox" mode and sends ERR(08001, ER_MUST_CHANGE_PASSWORD) for all commands that are not needed to set a new password until a new password is set.	true	5.1.23
interactiveClient	Set the CLIENT_INTERACTIVE flag, which tells MySQL to timeout connections based on INTERACTIVE_TIMEOUT instead of WAIT_TIMEOUT	false	3.1.0
localSocketAddress	Hostname or IP address given to explicitly configure the interface that the driver will bind the client side of the TCP/IP connection to when connecting.		5.0.5
propertiesTransform	An implementation of com.mysql.jdbc.ConnectionPropertiesTransform that the driver will use to modify URL properties passed to the driver before attempting a connection		3.1.4
useCompression	Use zlib compression when communicating with the server (true/false)? Defaults to 'false'.	false	3.0.17

**Networking.**

Property Name	Definition	Default Value	Since Version
maxAllowedPacket	Maximum allowed packet size to send to server. If not set, the value of system variable 'max_allowed_packet' will be used to initialize this upon connecting. This value will not take effect if set larger than the value of 'max_allowed_packet'.	-1	5.1.8
tcpKeepAlive	If connecting using TCP/IP, should the driver set SO_KEEPALIVE?	true	5.0.7
tcpNoDelay	If connecting using TCP/IP, should the driver set SO_TCP_NODELAY (disabling the Nagle Algorithm)?	true	5.0.7
tcpRcvBuf	If connecting using TCP/IP, should the driver set SO_RCV_BUF to the given value? The default value of '0', means use the platform default value for this property)	0	5.0.7

tcpSndBuf	If connecting using TCP/IP, should the driver set SO_SND_BUF to the given value? The default value of '0', means use the platform default value for this property)	0	5.0.7
tcpTrafficClass	If connecting using TCP/IP, should the driver set traffic class or type-of-service fields ?See the documentation for java.net.Socket.setTrafficClass() for more information.	0	5.0.7

**High Availability and Clustering.**

Property Name	Definition	Default Value	Since Version
autoReconnect	Should the driver try to re-establish stale and/or dead connections? If enabled the driver will throw an exception for a queries issued on a stale or dead connection, which belong to the current transaction, but will attempt reconnect before the next query issued on the connection in a new transaction. The use of this feature is not recommended, because it has side effects related to session state and data consistency when applications don't handle SQLExceptions properly, and is only designed to be used when you are unable to configure your application to handle SQLExceptions resulting from dead and stale connections properly. Alternatively, as a last option, investigate setting the MySQL server variable "wait_timeout" to a high value, rather than the default of 8 hours.	false	1.1
autoReconnectForPools	Use a reconnection strategy appropriate for connection pools (defaults to 'false')	false	3.1.3
failOverReadOnly	When failing over in autoReconnect mode, should the connection be set to 'read-only'?	true	3.0.12
maxReconnects	Maximum number of reconnects to attempt if autoReconnect is true, default is '3'.	3	1.1
reconnectAtTxEnd	If autoReconnect is set to true, should the driver attempt reconnections at the end of every transaction?	false	3.0.10
retriesAllDown	When using loadbalancing, the number of times the driver should cycle through available hosts, attempting to connect. Between cycles, the driver will pause for 250ms if no servers are available.	120	5.1.6
initialTimeout	If autoReconnect is enabled, the initial time to wait between re-connect attempts (in seconds, defaults to '2').	2	1.1
roundRobinLoadBalance	When autoReconnect is enabled, and failoverReadOnly is false, should we pick hosts to connect to on a round-robin basis?	false	3.1.2
queriesBeforeRetryMaster	Number of queries to issue before falling back to master when failed over (when using multi-host failover). Whichever condition is met first, 'queriesBeforeRetryMaster' or	50	3.0.2

	'secondsBeforeRetryMaster' will cause an attempt to be made to reconnect to the master. Defaults to 50.		
secondsBeforeRetryMaster	How long should the driver wait, when failed over, before attempting	30	3.0.2
selfDestructOnPingMaxOperations	If set to a non-zero value, the driver will report close the connection and report failure when Connection.ping() or Connection.isValid(int) is called if the connection's count of commands sent to the server exceeds this value.	0	5.1.6
selfDestructOnPingSecondsLifetime	If set to a non-zero value, the driver will report close the connection and report failure when Connection.ping() or Connection.isValid(int) is called if the connection's lifetime exceeds this value.	0	5.1.6
resourceId	A globally unique name that identifies the resource that this datasource or connection is connected to, used for XAResource.isSameRM() when the driver can't determine this value based on hostnames used in the URL		5.0.1

**Security.**

Property Name	Definition	Default Value	Since Version
allowMultiQueries	Allow the use of ';' to delimit multiple queries during one statement (true/false), defaults to 'false', and does not affect the addBatch() and executeBatch() methods, which instead rely on rewriteBatchStatements.	false	3.1.1
useSSL	Use SSL when communicating with the server (true/false), defaults to 'false'	false	3.0.2
requireSSL	Require SSL connection if useSSL=true? (defaults to 'false').	false	3.1.0
verifyServerCertificate	If "useSSL" is set to "true", should the driver verify the server's certificate? When using this feature, the keystore parameters should be specified by the "clientCertificateKeyStore*" properties, rather than system properties.	true	5.1.6
clientCertificateKeyStoreUrl	URL to the client certificate KeyStore (if not specified, use defaults)		5.1.0
clientCertificateKeyStoreType	KeyStore type for client certificates (NULL or empty means use the default, which is "JKS". Standard keystore types supported by the JVM are "JKS" and "PKCS12", your environment may have more available depending on what security products are installed and available to the JVM.	JKS	5.1.0
clientCertificateKeyStorePassword	Password for the client certificates KeyStore		5.1.0
trustCertificateKeyStoreUrl	URL to the trusted root certificate KeyStore (if not specified, use defaults)		5.1.0
trustCertificateKeyStoreType	KeyStore type for trusted root certificates (NULL or empty means use the default, which is "JKS". Standard keystore types supported by the JVM	JKS	5.1.0

	are "JKS" and "PKCS12", your environment may have more available depending on what security products are installed and available to the JVM.		
trustCertificateKeyStorePassword	Password for the trusted root certificates KeyStore		5.1.0
allowLoadLocalInfile	Should the driver allow use of 'LOAD DATA LOCAL INFILE...' (defaults to 'true').	true	3.0.3
allowUrlInLocalInfile	Should the driver allow URLs in 'LOAD DATA LOCAL INFILE' statements?	false	3.1.4
paranoid	Take measures to prevent exposure sensitive information in error messages and clear data structures holding sensitive data when possible? (defaults to 'false')	false	3.0.1
passwordCharacterEncoding	What character encoding is used for passwords? Leaving this set to the default value (null), uses the platform character set, which works for ISO8859_1 (i.e. "latin1") passwords. For passwords in other character encodings, the encoding will have to be specified with this property, as it's not possible for the driver to auto-detect this.		5.1.7

**Performance Extensions.**

Property Name	Definition	Default Value	Since Version
callableStmtCacheSize	If 'cacheCallableStmts' is enabled, how many callable statements should be cached?	100	3.1.2
metadataCacheSize	The number of queries to cache ResultSetMetadata for if cacheResultSetMetaData is set to 'true' (default 50)	50	3.1.1
useLocalSessionState	Should the driver refer to the internal values of autocommit and transaction isolation that are set by Connection.setAutoCommit() and Connection.setTransactionIsolation() and transaction state as maintained by the protocol, rather than querying the database or blindly sending commands to the database for commit() or rollback() method calls?	false	3.1.7
useLocalTransactionState	Should the driver use the in-transaction state provided by the MySQL protocol to determine if a commit() or rollback() should actually be sent to the database?	false	5.1.7
prepStmtCacheSize	If prepared statement caching is enabled, how many prepared statements should be cached?	25	3.0.10
prepStmtCacheSqlLimit	If prepared statement caching is enabled, what's the largest SQL the driver will cache the parsing for?	256	3.0.10
parseInfoCacheFactory	Name of a class implementing com.mysql.jdbc.CacheAdapterFactory, which will be used to create caches for the parsed representation of client-side prepared statements.	com.mysql.jdbc.PerConnectionL	
serverConfigCacheFactory	Name of a class implementing com.mysql.jdbc.CacheAdapterFactory<String,	com.mysql.jdbc.PerVmServerCo	



	Map<String, String>>, which will be used to create caches for MySQL server configuration values		
alwaysSendSetIsolation	Should the driver always communicate with the database when Connection.setTransactionIsolation() is called? If set to false, the driver will only communicate with the database when the requested transaction isolation is different than the whichever is newer, the last value that was set via Connection.setTransactionIsolation(), or the value that was read from the server when the connection was established. Note that useLocalSessionState=true will force the same behavior as alwaysSendSetIsolation=false, regardless of how alwaysSendSetIsolation is set.	true	3.1.7
maintainTimeStats	Should the driver maintain various internal timers to enable idle time calculations as well as more verbose error messages when the connection to the server fails? Setting this property to false removes at least two calls to System.getCurrentTimeMillis() per query.	true	3.1.9
useCursorFetch	If connected to MySQL > 5.0.2, and setFetchSize() > 0 on a statement, should that statement use cursor-based fetching to retrieve rows?	false	5.0.0
blobSendChunkSize	Chunk to use when sending BLOB/CLOBs via ServerPreparedStatements	1048576	3.1.9
cacheCallableStmts	Should the driver cache the parsing stage of CallableStatements	false	3.1.2
cachePrepStmts	Should the driver cache the parsing stage of PreparedStatements of client-side prepared statements, the "check" for suitability of server-side prepared and server-side prepared statements themselves?	false	3.0.10
cacheResultSetMetadata	Should the driver cache ResultSetMetaData for Statements and PreparedStatements? (Req. JDK-1.4+, true/false, default 'false')	false	3.1.1
cacheServerConfiguration	Should the driver cache the results of 'SHOW VARIABLES' and 'SHOW COLLATION' on a per-URL basis?	false	3.1.5
defaultFetchSize	The driver will call setFetchSize(n) with this value on all newly-created Statements	0	3.1.9
dontTrackOpenResources	The JDBC specification requires the driver to automatically track and close resources, however if your application doesn't do a good job of explicitly calling close() on statements or result sets, this can cause memory leakage. Setting this property to true relaxes this constraint, and can be more memory efficient for some applications.	false	3.1.7
dynamicCalendars	Should the driver retrieve the default calendar when required, or cache it per connection/session?	false	3.1.5

elideSetAutoCommits	If using MySQL-4.1 or newer, should the driver only issue 'set autocommit=n' queries when the server's state doesn't match the requested state by Connection.setAutoCommit(boolean)?	false	3.1.3
enableQueryTimeouts	When enabled, query timeouts set via Statement.setQueryTimeout() use a shared java.util.Timer instance for scheduling. Even if the timeout doesn't expire before the query is processed, there will be memory used by the TimerTask for the given timeout which won't be reclaimed until the time the timeout would have expired if it hadn't been cancelled by the driver. High-load environments might want to consider disabling this functionality.	true	5.0.6
holdResultsOpenOverStatementClose	Should the driver close result sets on Statement.close() as required by the JDBC specification?	false	3.1.7
largeRowSizeThreshold	What size result set row should the JDBC driver consider "large", and thus use a more memory-efficient way of representing the row internally?	2048	5.1.1
loadBalanceStrategy	If using a load-balanced connection to connect to SQL nodes in a MySQL Cluster/ NDB configuration (by using the URL prefix "jdbc:mysql:loadbalance://"), which load balancing algorithm should the driver use: (1) "random" - the driver will pick a random host for each request. This tends to work better than round-robin, as the randomness will somewhat account for spreading loads where requests vary in response time, while round-robin can sometimes lead to overloaded nodes if there are variations in response times across the workload. (2) "bestResponseTime" - the driver will route the request to the host that had the best response time for the previous transaction.	random	5.0.6
locatorFetchBufferSize	If 'emulateLocators' is configured to 'true', what size buffer should be used when fetching BLOB data for getBinaryInputStream?	1048576	3.2.1
rewriteBatchedStatements	Should the driver use multiqueries (irregardless of the setting of "allowMultiQueries") as well as rewriting of prepared statements for INSERT into multi-value inserts when executeBatch() is called? Notice that this has the potential for SQL injection if using plain java.sql.Statements and your code doesn't sanitize input correctly. Notice that for prepared statements, server-side prepared statements can not currently take advantage of this rewrite option, and that if you don't specify stream lengths when using PreparedStatement.set*Stream(), the driver won't be able to determine the optimum number of parameters per batch and you might receive an error from the driver that the resultant packet is too large. Statement.getGeneratedKeys() for	false	3.1.13

	these rewritten statements only works when the entire batch includes INSERT statements.		
useDirectRowUnpack	Use newer result set row unpacking code that skips a copy from network buffers to a MySQL packet instance and instead reads directly into the result set row data buffers.	true	5.1.1
useDynamicCharsetInfo	Should the driver use a per-connection cache of character set information queried from the server when necessary, or use a built-in static mapping that is more efficient, but isn't aware of custom character sets or character sets implemented after the release of the JDBC driver?	true	5.0.6
useFastDateParsing	Use internal String->Date/Time/Timestamp conversion routines to avoid excessive object creation?	true	5.0.5
useFastIntParsing	Use internal String->Integer conversion routines to avoid excessive object creation?	true	3.1.4
useJvmCharsetConverters	Always use the character encoding routines built into the JVM, rather than using lookup tables for single-byte character sets?	false	5.0.1
useReadAheadInput	Use newer, optimized non-blocking, buffered input stream when reading from the server?	true	3.1.5

**Debugging/Profiling.**

Property Name	Definition	Default Value	Since Version
logger	The name of a class that implements "com.mysql.jdbc.log.Log" that will be used to log messages to. (default is "com.mysql.jdbc.log.StandardLogger", which logs to STDERR)	com.mysql.jdbc.log.Standard	
gatherPerfMetrics	Should the driver gather performance metrics, and report them via the configured logger every 'reportMetricsIntervalMillis' milliseconds?	false	3.1.2
profileSQL	Trace queries and their execution/fetch times to the configured logger (true/false) defaults to 'false'	false	3.1.0
profileSql	Deprecated, use 'profileSQL' instead. Trace queries and their execution/fetch times on STDERR (true/false) defaults to 'false'		2.0.14
reportMetricsIntervalMillis	If 'gatherPerfMetrics' is enabled, how often should they be logged (in ms)?	30000	3.1.2
maxQuerySizeToLog	Controls the maximum length/size of a query that will get logged when profiling or tracing	2048	3.1.3
packetDebugBufferSize	The maximum number of packets to retain when 'enablePacketDebug' is true	20	3.1.3
slowQueryThresholdMillis	If 'logSlowQueries' is enabled, how long should a query (in ms) before it is logged as 'slow'?	2000	3.1.2
slowQueryThresholdNanos	If 'useNanosForElapsedTime' is set to true, and this property is set to a non-zero value, the driver will use this threshold (in nanosecond units) to determine if a query was slow.	0	5.0.7

useUsageAdvisor	Should the driver issue 'usage' warnings advising proper and efficient usage of JDBC and MySQL Connector/J to the log (true/false, defaults to 'false')?	false	3.1.1	
autoGenerateTestcaseScript	Should the driver dump the SQL it is executing, including server-side prepared statements to STDERR?	false	3.1.9	
autoSlowLog	Instead of using slowQueryThreshold* to determine if a query is slow enough to be logged, maintain statistics that allow the driver to determine queries that are outside the 99th percentile?	true	5.1.4	
clientInfoProvider	The name of a class that implements the com.mysql.jdbc.JDBC4ClientInfoProvider interface in order to support JDBC-4.0's Connection.get/setClientInfo() methods	com.mysql.jdbc.JDBC4ClientInfoProvider		JDBC4Comments
dumpMetadataOnColumnNotFoundException	Should the driver dump the field-level metadata of a result set into the exception message when ResultSet.findColumn() fails?	false	3.1.13	
dumpQueriesOnException	Should the driver dump the contents of the query sent to the server in the message for SQLExceptions?	false	3.1.3	
enablePacketDebug	When enabled, a ring-buffer of 'packetDebugBufferSize' packets will be kept, and dumped when exceptions are thrown in key areas in the driver's code	false	3.1.3	
explainSlowQueries	If 'logSlowQueries' is enabled, should the driver automatically issue an 'EXPLAIN' on the server and send the results to the configured log at a WARN level?	false	3.1.2	
includeInnodbStatusInDeadlockExceptions	Should the output of "SHOW ENGINE INNODB STATUS" in exception messages when deadlock exceptions are detected?	false	5.0.7	
includeThreadDumpInDeadlockExceptions	Should the current Java thread dump in exception messages when deadlock exceptions are detected?	false	5.1.15	
includeThreadNamesAsStatementComments	Should the name of the current thread as a comment visible in "SHOW PROCESSLIST", or in Innodb deadlock dumps, useful in correlation with "includeInnodbStatusInDeadlockExceptions=true" and "includeThreadDumpInDeadlockExceptions=true".	false	5.1.15	
logSlowQueries	Should queries that take longer than 'slowQueryThresholdMillis' be logged?	false	3.1.2	
logXaCommands	Should the driver log XA commands sent by MySQLXaConnection to the server, at the DEBUG level of logging?	false	5.0.5	
profilerEventHandler	Name of a class that implements the interface com.mysql.jdbc.profiler.ProfilerEventHandler that will be used to handle profiling/tracing events.	com.mysql.jdbc.profiler.LoggingProfilerEventHandler		Logging

resultSetSizeThreshold	If the usage advisor is enabled, how many rows should a result set contain before the driver warns that it is suspiciously large?	100	5.0.5
traceProtocol	Should trace-level network protocol be logged?	false	3.1.2
useNanosForElapsedTime	For profiling/debugging functionality that measures elapsed time, should the driver try to use nanoseconds resolution if available (JDK >= 1.5)?	false	5.0.7

**Miscellaneous.**

Property Name	Definition	Default Value	Since Version
useUnicode	Should the driver use Unicode character encodings when handling strings? Should only be used when the driver can't determine the character set mapping, or you are trying to 'force' the driver to use a character set that MySQL either doesn't natively support (such as UTF-8), true/false, defaults to 'true'	true	1.1g
characterEncoding	If 'useUnicode' is set to true, what character encoding should the driver use when dealing with strings? (defaults is to 'autodetect')		1.1g
characterSetResults	Character set to tell the server to return results as.		3.0.13
connectionAttributes	A comma-delimited list of user-defined key:value pairs (in addition to standard MySQL-defined key:value pairs) to be passed to MySQL Server for display as connection attributes in the PERFORMANCE_SCHEMA.SESSION_CONNECT_ATTRS table. Example usage: connectionAttributes=key1:value1,key2:value2 This functionality is available for use with MySQL Server version 5.6 or later only. Earlier versions of MySQL Server do not support connection attributes, causing this configuration option will be ignored. Setting connectionAttributes=none will cause connection attribute processing to be bypassed, for situations where Connection creation/initialization speed is critical.		5.1.25
connectionCollation	If set, tells the server to use this collation via 'set collation_connection'		3.0.13
useBlobToStoreUTF8OutsideBMP	Tells the driver to treat [MEDIUM/LONG]BLOB columns as [LONG]VARCHAR columns holding text encoded in UTF-8 that has characters outside the BMP (4-byte encodings), which MySQL server can't handle natively.	false	5.1.3
utf8OutsideBmpExcludedColumnNamePattern	When useBlobToStoreUTF8OutsideBMP is set to "true", column names matching the given regex will still be treated as BLOBs unless they match the regex specified for "utf8OutsideBmpIncludedColumnNamePattern". The regex must follow the patterns used for the java.util.regex package.		5.1.3
utf8OutsideBmpIncludedColumnNamePattern	Used to specify exclusion rules to "utf8OutsideBmpExcludedColumnNamePattern".		5.1.3

	The regex must follow the patterns used for the java.util.regex package.		
loadBalanceEnableJMX	Enables JMX-based management of load-balanced connection groups, including live addition/removal of hosts from load-balancing pool.	false	5.1.13
sessionVariables	A comma-separated list of name/value pairs to be sent as SET SESSION ... to the server when the driver connects.		3.1.8
useColumnNamesInFindColumn	Prior to JDBC-4.0, the JDBC specification had a bug related to what could be given as a "column name" to ResultSet methods like findColumn(), or getters that took a String property. JDBC-4.0 clarified "column name" to mean the label, as given in an "AS" clause and returned by ResultSetMetaData.getColumnLabel(), and if no AS clause, the column name. Setting this property to "true" will give behavior that is congruent to JDBC-3.0 and earlier versions of the JDBC specification, but which because of the specification bug could give unexpected results. This property is preferred over "useOldAliasMetadataBehavior" unless you need the specific behavior that it provides with respect to ResultSetMetadata.	false	5.1.7
allowNanAndInf	Should the driver allow NaN or +/- INF values in PreparedStatement.setDouble()?	false	3.1.5
autoClosePstmtStreams	Should the driver automatically call .close() on streams/readers passed as arguments via set*() methods?	false	3.1.12
autoDeserialize	Should the driver automatically detect and de-serialize objects stored in BLOB fields?	false	3.1.5
blobsAreStrings	Should the driver always treat BLOBs as Strings - specifically to work around dubious metadata returned by the server for GROUP BY clauses?	false	5.0.8
capitalizeTypeNames	Capitalize type names in DatabaseMetaData? (usually only useful when using WebObjects, true/false, defaults to 'false')	true	2.0.7
clobCharacterEncoding	The character encoding to use for sending and retrieving TEXT, MEDIUMTEXT and LONGTEXT values instead of the configured connection characterEncoding		5.0.0
clobberStreamingResults	This will cause a 'streaming' ResultSet to be automatically closed, and any outstanding data still streaming from the server to be discarded if another query is executed before all the data has been read from the server.	false	3.0.9
compensateOnDuplicateKeyUpdateCounts	Should the driver compensate for the update counts of "ON DUPLICATE KEY" INSERT statements (2 = 1, 0 = 1) when using prepared statements?	false	5.1.7

continueBatchOnError	Should the driver continue processing batch commands if one statement fails. The JDBC spec allows either way (defaults to 'true').	true	3.0.3
createDatabaseIfNotExist	Creates the database given in the URL if it doesn't yet exist. Assumes the configured user has permissions to create databases.	false	3.1.9
emptyStringsConvertToZero	Should the driver allow conversions from empty string fields to numeric values of '0'?	true	3.1.8
emulateLocators	Should the driver emulate java.sql.Blobs with locators? With this feature enabled, the driver will delay loading the actual Blob data until the one of the retrieval methods (getInputStream(), getBytes(), and so forth) on the blob data stream has been accessed. For this to work, you must use a column alias with the value of the column to the actual name of the Blob. The feature also has the following restrictions: The SELECT that created the result set must reference only one table, the table must have a primary key; the SELECT must alias the original blob column name, specified as a string, to an alternate name; the SELECT must cover all columns that make up the primary key.	false	3.1.0
emulateUnsupportedPstmts	Should the driver detect prepared statements that are not supported by the server, and replace them with client-side emulated versions?	true	3.1.7
exceptionInterceptors	Comma-delimited list of classes that implement com.mysql.jdbc.ExceptionInterceptor. These classes will be instantiated one per Connection instance, and all SQLExceptions thrown by the driver will be allowed to be intercepted by these interceptors, in a chained fashion, with the first class listed as the head of the chain.		5.1.8
functionsNeverReturnBlobs	Should the driver always treat data from functions returning BLOBs as Strings - specifically to work around dubious metadata returned by the server for GROUP BY clauses?	false	5.0.8
generateSimpleParameterMetadata	Should the driver generate simplified parameter metadata for PreparedStatement when no metadata is available either because the server couldn't support preparing the statement, or server-side prepared statements are disabled?	false	5.0.5
ignoreNonTxTables	Ignore non-transactional table warning for rollback? (defaults to 'false').	false	3.0.9
jdbcCompliantTruncation	Should the driver throw java.sql.DataTruncation exceptions when data is truncated as is required by the JDBC specification when connected to a server that supports warnings (MySQL 4.1.0 and newer)? This property has no effect if the server sql-mode includes STRICT_TRANS_TABLES.	true	3.1.2
loadBalanceAutoCommitStatements	When load-balancing is enabled for auto-commit statements (via loadBalanceAutoCommitStatementThreshold), the statement counter will only increment when the		5.1.15

Connector/J (JDBC) Reference

	SQL matches the regular expression. By default, every statement issued matches.		
loadBalanceAutoCommitStatementThreshold	When the auto-commit is enabled, the number of statements which should be executed before triggering load-balancing to rebalance. Default value of 0 causes load-balanced connections to only rebalance when exceptions are encountered, or auto-commit is disabled and transactions are explicitly committed or rolled back.	0	5.1.15
loadBalanceBlacklistTimeout	Time in milliseconds between checks of servers which are unavailable, by controlling how long a server lives in the global blacklist.	0	5.1.0
loadBalanceConnectionGroup	Logical group of load-balanced connections within a classloader, used to manage different groups independently. If not specified, live management of load-balanced connections is disabled.		5.1.13
loadBalanceExceptionChecker	Fully-qualified class name of custom exception checker. The class must implement com.mysql.jdbc.LoadBalanceExceptionChecker interface, and is used to inspect SQLExceptions and determine whether they should trigger fail-over to another host in a load-balanced deployment.	com.mysql.jdbc.StandardLoadBalanceExceptionChecker	
loadBalancePingTimeout	Time in milliseconds to wait for ping response from each of load-balanced physical connections when using load-balanced Connection.	0	5.1.13
loadBalanceSQLExceptionSubClassFailover	Comma-delimited list of classes/interfaces used by default load-balanced exception checker to determine whether a given SQLException should trigger failover. The comparison is done using Class.isInstance(SQLException) using the thrown SQLException.		5.1.13
loadBalanceSQLStateFailover	Comma-delimited list of SQLState codes used by default load-balanced exception checker to determine whether a given SQLException should trigger failover. The SQLState of a given SQLException is evaluated to determine whether it begins with any value in the comma-delimited list.		5.1.13
loadBalanceValidateConnectionOnSwapStatement	Should a load-balanced Connection explicitly check whether the connection is live when swapping to a new physical connection at commit/rollback?	false	5.1.13
maxRows	The maximum number of rows to return (0, the default means return all rows).	-1	all versions
netTimeoutForStreamingResults	What value should the driver automatically set the server setting 'net_write_timeout' to when the streaming result sets feature is in use? (value has unit of seconds, the value '0' means the driver will not try and adjust this value)	600	5.1.0
noAccessToProcedureBodies	When determining procedure parameter types for CallableStatements, and the connected user can't access procedure bodies through "SHOW CREATE PROCEDURE" or select on mysql.proc	false	5.0.3



	should the driver instead create basic metadata (all parameters reported as IN VARCHARs, but allowing registerOutParameter() to be called on them anyway) instead of throwing an exception?		
noDatetimeStringSync	Don't ensure that ResultSet.getDatetimeType().toString().equals(ResultSet.getString())	false	3.1.7
noTimezoneConversionForTimeType	Do not convert TIME values using the server timezone if 'useTimezone'='true'	false	5.0.0
nullCatalogMeansCurrent	When DatabaseMetadataMethods ask for a 'catalog' parameter, does the value null mean use the current catalog? (this is not JDBC-compliant, but follows legacy behavior from earlier versions of the driver)	true	3.1.8
nullNamePatternMatchesAll	Should DatabaseMetaData methods that accept *pattern parameters treat null the same as '%' (this is not JDBC-compliant, however older versions of the driver accepted this departure from the specification)	true	3.1.8
overrideSupportsIntegrityEnhancementFacility	Should the driver return "true" for DatabaseMeta-data.supportsIntegrityEnhancementFacility() even if the database doesn't support it to workaround applications that require this method to return "true" to signal support of foreign keys, even though the SQL specification states that this facility contains much more than just foreign key support (one such application being OpenOffice)?	false	3.1.12
padCharsWithSpace	If a result set column has the CHAR type and the value does not fill the amount of characters specified in the DDL for the column, should the driver pad the remaining characters with space (for ANSI compliance)?	false	5.0.6
pedantic	Follow the JDBC spec to the letter.	false	3.0.0
pinGlobalTxToPhysicalConnection	When using XAConnections, should the driver ensure that operations on a given XID are always routed to the same physical connection? This allows the XAConnection to support "XA START ... JOIN" after "XA END" has been called	false	5.0.1
populateInsertRowWithDefaultValues	When using ResultSets that are CONCUR_UPDATABLE, should the driver pre-populate the "insert" row with default values from the DDL for the table used in the query so those values are immediately available for ResultSet accessors? This functionality requires a call to the database for metadata each time a result set of this type is created. If disabled (the default), the default values will be populated by the an internal call to refreshRow() which pulls back default values and/or values changed by triggers.	false	5.0.5
processEscapeCodesForPreparedStatements	Should the driver process escape codes in queries that are prepared?	true	3.1.12
queryTimeoutKillsConnection	If the timeout given in Statement.setQueryTimeout() expires, should	false	5.1.9

## Connector/J (JDBC) Reference

	the driver forcibly abort the Connection instead of attempting to abort the query?		
relaxAutoCommit	If the version of MySQL the driver connects to does not support transactions, still allow calls to commit(), rollback() and setAutoCommit() (true/false, defaults to 'false')?	false	2.0.13
retainStatementAfterResultSetClose	Should the driver retain the Statement reference in a ResultSet after ResultSet.close() has been called. This is not JDBC-compliant after JDBC-4.0.	false	3.1.11
rollbackOnPooledClose	Should the driver issue a rollback() when the logical connection in a pool is closed?	true	3.0.15
runningCTS13	Enables workarounds for bugs in Sun's JDBC compliance testsuite version 1.3	false	3.1.7
serverTimezone	Override detection/mapping of timezone. Used when timezone from server doesn't map to Java timezone		3.0.2
statementInterceptors	A comma-delimited list of classes that implement "com.mysql.jdbc.StatementInterceptor" that should be placed "in between" query execution to influence the results. StatementInterceptors are "chainable", the results returned by the "current" interceptor will be passed on to the next in the chain, from left-to-right order, as specified in this property.		5.1.1
strictFloatingPoint	Used only in older versions of compliance test	false	3.0.0
strictUpdates	Should the driver do strict checking (all primary keys selected) of updatable result sets (true, false, defaults to 'true')?	true	3.0.4
tinyInt1isBit	Should the driver treat the datatype TINYINT(1) as the BIT type (because the server silently converts BIT -> TINYINT(1) when creating tables)?	true	3.0.16
transformedBitsIsBoolean	If the driver converts TINYINT(1) to a different type, should it use BOOLEAN instead of BIT for future compatibility with MySQL-5.0, as MySQL-5.0 has a BIT type?	false	3.1.9
treatUtilDateAsTimestamp	Should the driver treat java.util.Date as a TIMESTAMP for the purposes of PreparedStatement.setObject()?	true	5.0.5
ultraDevHack	Create PreparedStatements for prepareCall() when required, because UltraDev is broken and issues a prepareCall() for _all_ statements? (true/false, defaults to 'false')	false	2.0.3
useAffectedRows	Don't set the CLIENT_FOUND_ROWS flag when connecting to the server (not JDBC-compliant, will break most applications that rely on "found" rows vs. "affected rows" for DML statements), but does cause "correct" update counts from "INSERT ... ON DUPLICATE KEY UPDATE" statements to be returned by the server.	false	5.1.7
useGmtMillisForDatetimes	Convert between session timezone and GMT before creating Date and Timestamp instances	false	3.1.12

	(value of "false" is legacy behavior, "true" leads to more JDBC-compliant behavior.		
useHostsInPrivileges	Add '@hostname' to users in DatabaseMetaData.getColumn/TablePrivileges() (true/false), defaults to 'true'.	true	3.0.2
useInformationSchema	When connected to MySQL-5.0.7 or newer, should the driver use the INFORMATION_SCHEMA to derive information used by DatabaseMetaData?	false	5.0.0
useJDBCCompliantTimezoneShift	Should the driver use JDBC-compliant rules when converting TIME/TIMESTAMP/DATETIME values' timezone information for those JDBC arguments which take a java.util.Calendar argument? (Notice that this option is exclusive of the "useTimezone=true" configuration option.)	false	5.0.0
useLegacyDatetimeCode	Use code for DATE/TIME/DATETIME/TIMESTAMP handling in result sets and statements that consistently handles timezone conversions from client to server and back again, or use the legacy code for these datatypes that has been in the driver for backwards-compatibility?	true	5.1.6
useOldAliasMetadataBehavior	Should the driver use the legacy behavior for "AS" clauses on columns and tables, and only return aliases (if any) for ResultSetMetaData.getColumn() or ResultSetMetaData.getTable() rather than the original column/table name? In 5.0.x, the default value was true.	false	5.0.4
useOldUTF8Behavior	Use the UTF-8 behavior the driver did when communicating with 4.0 and older servers	false	3.1.6
useOnlyServerErrorMessages	Don't prepend 'standard' SQLState error messages to error messages returned by the server.	true	3.0.15
useSSPSCompatibleTimezoneShift	When migrating from an environment that was using server-side prepared statements, and the configuration property "useJDBCCompliantTimezoneShift" set to "true", use compatible behavior when not using server-side prepared statements when sending TIMESTAMP values to the MySQL server.	false	5.0.5
useServerPrepStmts	Use server-side prepared statements if the server supports them?	false	3.1.0
useSqlStateCodes	Use SQL Standard state codes instead of 'legacy' X/Open/SQL state codes (true/false), default is 'true'	true	3.1.3
useStreamLengthsInPrepStmts	Honor stream length parameter in PreparedStatement/ResultSet.setXXXStream() method calls (true/false, defaults to 'true')?	true	3.0.2
useTimezone	Convert time/date types between client and server timezones (true/false, defaults to 'false')?	false	3.0.2
useUnbufferedInput	Don't use BufferedInputStream for reading data from the server	true	3.0.11

yearIsDateType	Should the JDBC driver treat the MySQL type "YEAR" as a java.sql.Date, or as a SHORT?	true	3.1.9
zeroDateTimeBehavior	What should happen when the driver encounters DATETIME values that are composed entirely of zeros (used by MySQL to represent invalid dates)? Valid values are "exception", "round" and "convertToNull".	exception	3.1.4

Connector/J also supports access to MySQL using named pipes on Windows NT, Windows 2000, or Windows XP using the `NamedPipeSocketFactory` as a plugin-socket factory using the `socketFactory` property. If you do not use a `namedPipePath` property, the default of `'\\.\pipe\MySQL'` is used. If you use the `NamedPipeSocketFactory`, the host name and port number values in the JDBC url are ignored. To enable this feature, use:

```
socketFactory=com.mysql.jdbc.NamedPipeSocketFactory
```

Named pipes only work when connecting to a MySQL server on the same physical machine where the JDBC driver is running. In simple performance tests, named pipe access is between 30%-50% faster than the standard TCP/IP access. However, this varies per system, and named pipes are slower than TCP/IP in many Windows configurations.

To create your own socket factories, follow the example code in `com.mysql.jdbc.NamedPipeSocketFactory`, or `com.mysql.jdbc.StandardSocketFactory`.

### 21.3.5.1.1. Properties Files for the `useConfigs` Option

The `useConfigs` connection option is a convenient shorthand for specifying combinations of options for particular scenarios. The argument values you can use with this option correspond to the names of `.properties` files within the Connector/J `mysql-connector-java-version-bin.jar` JAR file. For example, the Connector/J 5.1.9 driver includes the following configuration properties files:

```
$ unzip mysql-connector-java-5.1.19-bin.jar '*/configs/*'
Archive:  mysql-connector-java-5.1.19-bin.jar
  creating:  com/mysql/jdbc/configs/
  inflating:  com/mysql/jdbc/configs/3-0-Compat.properties
  inflating:  com/mysql/jdbc/configs/5-0-Compat.properties
  inflating:  com/mysql/jdbc/configs/clusterBase.properties
  inflating:  com/mysql/jdbc/configs/coldFusion.properties
  inflating:  com/mysql/jdbc/configs/fullDebug.properties
  inflating:  com/mysql/jdbc/configs/maxPerformance.properties
  inflating:  com/mysql/jdbc/configs/solarisMaxPerformance.properties
```

To specify one of these combinations of options, specify `useConfigs=3-0-Compat`, `useConfigs=maxPerformance`, and so on. The following sections show the options that are part of each `useConfigs` setting. For the details of why each one is included, see the comments in the `.properties` files.

#### 3-0-Compat

```
emptyStringsConvertToZero=true
jdbcCompliantTruncation=false
noDatetimeStringSync=true
nullCatalogMeansCurrent=true
nullNamePatternMatchesAll=true
transformedBitIsBoolean=false
dontTrackOpenResources=true
zeroDateTimeBehavior=convertToNull
useServerPrepStmts=false
autoClosePstmtStreams=true
processEscapeCodesForPrepStmts=false
useFastDateParsing=false
populateInsertRowWithDefaultValues=false
useDirectRowUnpack=false
```

## 5-0-Compat

```
useDirectRowUnpack=false
```

## clusterBase

```
autoReconnect=true
failOverReadOnly=false
roundRobinLoadBalance=true
```

## coldFusion

```
useDynamicCharsetInfo=false
alwaysSendSetIsolation=false
useLocalSessionState=true
autoReconnect=true
```

## fullDebug

```
profileSQL=true
gatherPerMetrics=true
useUsageAdvisor=true
logSlowQueries=true
explainSlowQueries=true
```

## maxPerformance

```
cachePrepStmts=true
cacheCallableStmts=true
cacheServerConfiguration=true
useLocalSessionState=true
elideSetAutoCommits=true
alwaysSendSetIsolation=false
enableQueryTimeouts=false
```

## solarisMaxPerformance

```
useUnbufferedInput=false
useReadAheadInput=false
maintainTimeStats=false
```

## 21.3.5.2. JDBC API Implementation Notes

MySQL Connector/J passes all of the tests in the publicly available version of Sun's JDBC compliance test suite. This section gives details on a interface-by-interface level about implementation decisions that might affect how you code applications with MySQL Connector/J. The JDBC specification is vague about how certain functionality should be implemented, or the specification enables leeway in implementation.

### • BLOB

Starting with Connector/J version 3.1.0, you can emulate BLOBs with locators by adding the property `emulateLocators=true` to your JDBC URL. Using this method, the driver will delay loading the actual BLOB data until you retrieve the other data and then use retrieval methods (`getInputStream()`, `getBytes()`, and so forth) on the BLOB data stream.

You must use a column alias with the value of the column to the actual name of the BLOB, for example:

```
SELECT id, 'data' as blob_data from blobtable
```

You must also follow these rules:

- The `SELECT` must reference only one table. The table must have a [primary key](#).
- The `SELECT` must alias the original BLOB column name, specified as a string, to an alternate name.

- The `SELECT` must cover all columns that make up the primary key.

The BLOB implementation does not allow in-place modification (they are copies, as reported by the `DatabaseMetaData.locatorsUpdateCopies()` method). Because of this, use the corresponding `PreparedStatement.setBlob()` or `ResultSet.updateBlob()` (in the case of updatable result sets) methods to save changes back to the database.

- **CallableStatement**

Starting with Connector/J 3.1.1, stored procedures are supported when connecting to MySQL version 5.0 or newer using the `CallableStatement` interface. Currently, the `getParameterMetaData()` method of `CallableStatement` is not supported.

- **CLOB**

The CLOB implementation does not allow in-place modification (they are copies, as reported by the `DatabaseMetaData.locatorsUpdateCopies()` method). Because of this, use the `PreparedStatement.setClob()` method to save changes back to the database. The JDBC API does not have a `ResultSet.updateClob()` method.

- **Connection**

Unlike the pre-Connector/J JDBC driver (`MM.MySQL`), the `isClosed()` method does not ping the server to determine if it is available. In accordance with the JDBC specification, it only returns true if `closed()` has been called on the connection. If you need to determine if the connection is still valid, issue a simple query, such as `SELECT 1`. The driver will throw an exception if the connection is no longer valid.

- **DatabaseMetaData**

**Foreign key** information (`getImportedKeys()`/`getExportedKeys()` and `getCrossReference()`) is only available from `InnoDB` tables. The driver uses `SHOW CREATE TABLE` to retrieve this information, so if any other storage engines add support for foreign keys, the driver would transparently support them as well.

- **PreparedStatement**

`PreparedStatement`s are implemented by the driver, as MySQL does not have a prepared statement feature. Because of this, the driver does not implement `getParameterMetaData()` or `getMetaData()` as it would require the driver to have a complete SQL parser in the client.

Starting with version 3.1.0 MySQL Connector/J, server-side prepared statements and binary-encoded result sets are used when the server supports them.

Take care when using a server-side prepared statement with **large** parameters that are set using `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setBlob()`, or `setClob()`. To re-execute the statement with any large parameter changed to a nonlarge parameter, call `clearParameters()` and set all parameters again. The reason for this is as follows:

- During both server-side prepared statements and client-side emulation, large data is exchanged only when `PreparedStatement.execute()` is called.
- Once that has been done, the stream used to read the data on the client side is closed (as per the JDBC spec), and cannot be read from again.
- If a parameter changes from large to nonlarge, the driver must reset the server-side state of the prepared statement to allow the parameter that is being changed to take the place of the prior large value. This removes all of the large data that has already been sent to the server,

thus requiring the data to be re-sent, using the `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setBlob()` or `setClob()` method.

Consequently, to change the type of a parameter to a nonlarge one, you must call `clearParameters()` and set all parameters of the prepared statement again before it can be re-executed.

- **ResultSet**

By default, ResultSets are completely retrieved and stored in memory. In most cases this is the most efficient way to operate, and due to the design of the MySQL network protocol is easier to implement. If you are working with ResultSets that have a large number of rows or large values, and cannot allocate heap space in your JVM for the memory required, you can tell the driver to stream the results back one row at a time.

To enable this functionality, create a `Statement` instance in the following manner:

```
stmt = conn.createStatement( java.sql.ResultSet.TYPE_FORWARD_ONLY,
                             java.sql.ResultSet.CONCUR_READ_ONLY );
stmt.setFetchSize( Integer.MIN_VALUE );
```

The combination of a forward-only, read-only result set, with a fetch size of `Integer.MIN_VALUE` serves as a signal to the driver to stream result sets row-by-row. After this, any result sets created with the statement will be retrieved row-by-row.

There are some caveats with this approach. You must read all of the rows in the result set (or close it) before you can issue any other queries on the connection, or an exception will be thrown.

The earliest the locks these statements hold can be released (whether they be `MyISAM` table-level locks or row-level locks in some other storage engine such as `InnoDB`) is when the statement completes.

If the statement is within scope of a transaction, then locks are released when the transaction completes (which implies that the statement needs to complete first). As with most other databases, statements are not complete until all the results pending on the statement are read or the active result set for the statement is closed.

Therefore, if using streaming results, process them as quickly as possible if you want to maintain concurrent access to the tables referenced by the statement producing the result set.

- **ResultSetMetaData**

The `isAutoIncrement()` method only works when using MySQL servers 4.0 and newer.

- **Statement**

When using versions of the JDBC driver earlier than 3.2.1, and connected to server versions earlier than 5.0.3, the `setFetchSize()` method has no effect, other than to toggle result set streaming as described above.

Connector/J 5.0.0 and later include support for both `Statement.cancel()` and `Statement.setQueryTimeout()`. Both require MySQL 5.0.0 or newer server, and require a separate connection to issue the `KILL QUERY` statement. In the case of `setQueryTimeout()`, the implementation creates an additional thread to handle the timeout functionality.

**Note**

Failures to cancel the statement for `setQueryTimeout()` may manifest themselves as `RuntimeException` rather than failing silently, as there is currently no way to unblock the thread that is executing the query being cancelled due to timeout expiration and have it throw the exception instead.

**Note**

The MySQL statement `KILL QUERY` (which is what the driver uses to implement `Statement.cancel()`) is non-deterministic; thus, avoid the use of `Statement.cancel()` if possible. If no query is in process, the next query issued will be killed by the server. This race condition is guarded against as of Connector/J 5.1.18.

MySQL does not support SQL cursors, and the JDBC driver doesn't emulate them, so `setCursorName()` has no effect.

Connector/J 5.1.3 and later include two additional methods:

- `setLocalInfileInputStream()` sets an `InputStream` instance that will be used to send data to the MySQL server for a `LOAD DATA LOCAL INFILE` statement rather than a `FileInputStream` or `URLInputStream` that represents the path given as an argument to the statement.

This stream will be read to completion upon execution of a `LOAD DATA LOCAL INFILE` statement, and will automatically be closed by the driver, so it needs to be reset before each call to `execute*()` that would cause the MySQL server to request data to fulfill the request for `LOAD DATA LOCAL INFILE`.

If this value is set to `NULL`, the driver will revert to using a `FileInputStream` or `URLInputStream` as required.

- `getLocalInfileInputStream()` returns the `InputStream` instance that will be used to send data in response to a `LOAD DATA LOCAL INFILE` statement.

This method returns `NULL` if no such stream has been set using `setLocalInfileInputStream()`.

### 21.3.5.3. Java, JDBC and MySQL Types

MySQL Connector/J is flexible in the way it handles conversions between MySQL data types and Java data types.

In general, any MySQL data type can be converted to a `java.lang.String`, and any numeric type can be converted to any of the Java numeric types, although round-off, overflow, or loss of precision may occur.

**Note**

All `TEXT` types return `Types.LONGVARCHAR` with different `getPrecision()` values (65535, 255, 16777215, and 2147483647 respectively) with `getColumnType()` returning `-1`. This behavior is intentional even though `TINYTEXT` does not fall, regarding to its size, within the `LONGVARCHAR` category. This is to avoid different handling inside the same base type. And `getColumnType()` returns `-1` because the internal server handling is of type `TEXT`, which is similar to `BLOB`.

Also note that `getColumnTypeName()` will return `VARCHAR` even though `getColumnType()` returns `Types.LONGVARCHAR`, because `VARCHAR` is the designated column database-specific name for this type.

Starting with Connector/J 3.1.0, the JDBC driver issues warnings or throws `DataTruncation` exceptions as is required by the JDBC specification unless the connection was configured not to do so by using the property `jdbcCompliantTruncation` and setting it to `false`.



The conversions that are always guaranteed to work are listed in the following table. The first column lists one or more MySQL data types, and the second column lists one or more Java types to which the MySQL types can be converted.

**Table 21.24. Connection Properties - Miscellaneous**

These MySQL Data Types	Can always be converted to these Java types
CHAR, VARCHAR, BLOB, TEXT, ENUM, and SET	java.lang.String, java.io.InputStream, java.io.Reader, java.sql.Blob, java.sql.Clob
FLOAT, REAL, DOUBLE PRECISION, NUMERIC, DECIMAL, TINYINT, SMALLINT, MEDIUMINT, INTEGER, BIGINT	java.lang.String, java.lang.Short, java.lang.Integer, java.lang.Long, java.lang.Double, java.math.BigDecimal
DATE, TIME, DATETIME, TIMESTAMP	java.lang.String, java.sql.Date, java.sql.Timestamp

**Note**

Round-off, overflow or loss of precision may occur if you choose a Java numeric data type that has less precision or capacity than the MySQL data type you are converting to/from.

The `ResultSet.getObject()` method uses the type conversions between MySQL and Java types, following the JDBC specification where appropriate. The value returned by `ResultSetMetaData.GetColumnName()` is also shown below. For more information on the `java.sql.Types` classes see [Java 2 Platform Types](#).

**Table 21.25. MySQL Types to Java Types for ResultSet.getObject()**

MySQL Type Name	Return value of <code>GetColumnName</code>	Returned as Java Class
BIT(1) (new in MySQL-5.0)	BIT	java.lang.Boolean
BIT( > 1) (new in MySQL-5.0)	BIT	byte[]
TINYINT	TINYINT	java.lang.Boolean if the configuration property <code>tinyInt1isBit</code> is set to <code>true</code> (the default) and the storage size is 1, or java.lang.Integer if not.
BOOL, BOOLEAN	TINYINT	See TINYINT, above as these are aliases for TINYINT(1), currently.
SMALLINT[ (M) ] [UNSIGNED]	SMALLINT [UNSIGNED]	java.lang.Integer (regardless if UNSIGNED or not)
MEDIUMINT[ (M) ] [UNSIGNED]	MEDIUMINT [UNSIGNED]	java.lang.Integer, if UNSIGNED java.lang.Long (C/J 3.1 and earlier), or java.lang.Integer for C/J 5.0 and later
INT, INTEGER[ (M) ] [UNSIGNED]	INTEGER [UNSIGNED]	java.lang.Integer, if UNSIGNED java.lang.Long
BIGINT[ (M) ] [UNSIGNED]	BIGINT [UNSIGNED]	java.lang.Long, if UNSIGNED java.math.BigInteger
FLOAT[ (M,D) ]	FLOAT	java.lang.Float
DOUBLE[ (M,B) ]	DOUBLE	java.lang.Double
DECIMAL[ (M[,D) ] ]	DECIMAL	java.math.BigDecimal

MySQL Type Name	Return value of <code>GetColumnName</code>	Returned as Java Class
DATE	DATE	<code>java.sql.Date</code>
DATETIME	DATETIME	<code>java.sql.Timestamp</code>
TIMESTAMP[ (M) ]	TIMESTAMP	<code>java.sql.Timestamp</code>
TIME	TIME	<code>java.sql.Time</code>
YEAR[ ( 2   4 ) ]	YEAR	If <code>yearIsDateType</code> configuration property is set to <code>false</code> , then the returned object type is <code>java.sql.Short</code> . If set to <code>true</code> (the default), then the returned object is of type <code>java.sql.Date</code> with the date set to January 1st, at midnight.
CHAR(M)	CHAR	<code>java.lang.String</code> (unless the character set for the column is <code>BINARY</code> , then <code>byte[]</code> is returned.
VARCHAR(M) [BINARY]	VARCHAR	<code>java.lang.String</code> (unless the character set for the column is <code>BINARY</code> , then <code>byte[]</code> is returned.
BINARY(M)	BINARY	<code>byte[]</code>
VARBINARY(M)	VARBINARY	<code>byte[]</code>
TINYBLOB	TINYBLOB	<code>byte[]</code>
TINYTEXT	VARCHAR	<code>java.lang.String</code>
BLOB	BLOB	<code>byte[]</code>
TEXT	VARCHAR	<code>java.lang.String</code>
MEDIUMBLOB	MEDIUMBLOB	<code>byte[]</code>
MEDIUMTEXT	VARCHAR	<code>java.lang.String</code>
LOBLOB	LOBLOB	<code>byte[]</code>
LONGTEXT	VARCHAR	<code>java.lang.String</code>
ENUM('value1', 'value2', ...)	CHAR...	<code>java.lang.String</code>
SET('value1', 'value2', ...)	CHAR..	<code>java.lang.String</code>

#### 21.3.5.4. Using Character Sets and Unicode

All strings sent from the JDBC driver to the server are converted automatically from native Java Unicode form to the client character encoding, including all queries sent using `Statement.execute()`, `Statement.executeUpdate()`, `Statement.executeQuery()` as well as all `PreparedStatement` and `CallableStatement` parameters with the exclusion of parameters set using `setBytes()`, `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()` and `setBlob()`.

#### Number of Encodings Per Connection

In MySQL Server 4.1 and higher, Connector/J supports a single character encoding between client and server, and any number of character encodings for data returned by the server to the client in `ResultSets`.

Prior to MySQL Server 4.1, Connector/J supported a single character encoding per connection, which could either be automatically detected from the server configuration, or could be configured by the user through the `useUnicode` and `characterEncoding` properties.

#### Setting the Character Encoding

The character encoding between client and server is automatically detected upon connection. You specify the encoding on the server using the `character_set_server` [468] for server versions 4.1.0

and newer, and `character_set` system variable for server versions older than 4.1.0. The driver automatically uses the encoding specified by the server. For more information, see [Section 10.1.3.1, “Server Character Set and Collation”](#).

For example, to use 4-byte UTF-8 character sets with Connector/J, configure the MySQL server with `character_set_server=utf8mb4` [468], and leave `characterEncoding` out of the Connector/J connection string. Connector/J will then autodetect the UTF-8 setting.

To override the automatically detected encoding on the client side, use the `characterEncoding` property in the URL used to connect to the server.

To allow multiple character sets to be sent from the client, use the UTF-8 encoding, either by configuring `utf8` as the default server character set, or by configuring the JDBC driver to use UTF-8 through the `characterEncoding` property.

When specifying character encodings on the client side, use Java-style names. The following table lists MySQL character set names and the corresponding Java-style names:

**Table 21.26. MySQL to Java Encoding Name Translations**

MySQL Character Set Name	Java-Style Character Encoding Name
<code>ascii</code>	<code>US-ASCII</code>
<code>big5</code>	<code>Big5</code>
<code>gbk</code>	<code>GBK</code>
<code>sjis</code>	<code>SJIS</code> (or <code>Cp932</code> or <code>MS932</code> for MySQL Server < 4.1.11)
<code>cp932</code>	<code>Cp932</code> or <code>MS932</code> (MySQL Server > 4.1.11)
<code>gb2312</code>	<code>EUC_CN</code>
<code>ujis</code>	<code>EUC_JP</code>
<code>euckr</code>	<code>EUC_KR</code>
<code>latin1</code>	<code>Cp1252</code>
<code>latin2</code>	<code>ISO8859_2</code>
<code>greek</code>	<code>ISO8859_7</code>
<code>hebrew</code>	<code>ISO8859_8</code>
<code>cp866</code>	<code>Cp866</code>
<code>tis620</code>	<code>TIS620</code>
<code>cp1250</code>	<code>Cp1250</code>
<code>cp1251</code>	<code>Cp1251</code>
<code>cp1257</code>	<code>Cp1257</code>
<code>macroman</code>	<code>MacRoman</code>
<code>macce</code>	<code>MacCentralEurope</code>
<code>utf8</code>	<code>UTF-8</code>
<code>ucs2</code>	<code>UnicodeBig</code>



#### Warning

Do not issue the query `set names` with Connector/J, as the driver will not detect that the character set has changed, and will continue to use the character set detected during the initial connection setup.

### 21.3.5.5. Connecting Securely Using SSL

SSL in MySQL Connector/J encrypts all data (other than the initial handshake) between the JDBC driver and the server. The performance penalty for enabling SSL is an increase in query processing time between 35% and 50%, depending on the size of the query, and the amount of data it returns.

For SSL support to work, you must have the following:

- A JDK that includes JSSE (Java Secure Sockets Extension), like JDK-1.4.1 or newer. SSL does not currently work with a JDK that you can add JSSE to, like JDK-1.2.x or JDK-1.3.x due to the following JSSE bug: <http://developer.java.sun.com/developer/bugParade/bugs/4273544.html>
- A MySQL server that supports SSL and has been compiled and configured to do so, which is MySQL 4.0.4 or later. For more information, see [Section 6.3.9.2, “Configuring MySQL for SSL”](#).
- A client certificate (covered later in this section)

The system works through two Java truststore files, one file contains the certificate information for the server (`truststore` in the examples below). The other file contains the certificate for the client (`keystore` in the examples below). All Java truststore files are password protected by supplying a suitable password to the `keytool` when you create the files. You need the file names and associated passwords to create an SSL connection.

You will first need to import the MySQL server CA Certificate into a Java truststore. A sample MySQL server CA Certificate is located in the [SSL](#) subdirectory of the MySQL source distribution. This is what SSL will use to determine if you are communicating with a secure MySQL server. Alternatively, use the CA Certificate that you have generated or been provided with by your SSL provider.

To use Java's `keytool` to create a truststore in the current directory, and import the server's CA certificate (`cacert.pem`), you can do the following (assuming that `keytool` is in your path. The `keytool` is typically located in the `bin` subdirectory of your JDK or JRE):

```
shell> keytool -import -alias mysqlServerCACert \
  -file cacert.pem -keystore truststore
```

Enter the password when prompted for the keystore file. Interaction with `keytool` looks like this:

```
Enter keystore password: *****
Owner: EMAILADDRESS=walrus@example.com, CN=Walrus,
  O=MySQL AB, L=Orenburg, ST=Some-State, C=RU
Issuer: EMAILADDRESS=walrus@example.com, CN=Walrus,
  O=MySQL AB, L=Orenburg, ST=Some-State, C=RU
Serial number: 0
Valid from:
  Fri Aug 02 16:55:53 CDT 2002 until: Sat Aug 02 16:55:53 CDT 2003
Certificate fingerprints:
  MD5: 61:91:A0:F2:03:07:61:7A:81:38:66:DA:19:C4:8D:AB
  SHA1: 25:77:41:05:D5:AD:99:8C:14:8C:CA:68:9C:2F:B8:89:C3:34:4D:6C
Trust this certificate? [no]: yes
Certificate was added to keystore
```

You then have two options: either import the client certificate that matches the CA certificate you just imported, or create a new client certificate.

Importing an existing certificate requires the certificate to be in DER format. You can use `openssl` to convert an existing certificate into the new format. For example:

```
shell> openssl x509 -outform DER -in client-cert.pem -out client.cert
```

Now import the converted certificate into your keystore using `keytool`:

```
shell> keytool -import -file client.cert -keystore keystore -alias mysqlClientCertificate
```

To generate your own client certificate, use `keytool` to create a suitable certificate and add it to the `keystore` file:

```
shell> keytool -genkey -keyalg rsa \
  -alias mysqlClientCertificate -keystore keystore
```

Keytool will prompt you for the following information, and create a keystore named `keystore` in the current directory.

Respond with information that is appropriate for your situation:

```
Enter keystore password: *****
What is your first and last name?
  [Unknown]: Matthews
What is the name of your organizational unit?
  [Unknown]: Software Development
What is the name of your organization?
  [Unknown]: MySQL AB
What is the name of your City or Locality?
  [Unknown]: Flossmoor
What is the name of your State or Province?
  [Unknown]: IL
What is the two-letter country code for this unit?
  [Unknown]: US
Is <CN=Matthews, OU=Software Development, O=MySQL AB,
L=Flossmoor, ST=IL, C=US> correct?
  [no]: y

Enter key password for <mysqlClientCertificate>
  (RETURN if same as keystore password):
```

Finally, to get JSSE to use the keystore and truststore that you have generated, you need to set the following system properties when you start your JVM, replacing `path_to_keystore_file` with the full path to the keystore file you created, `path_to_truststore_file` with the path to the truststore file you created, and using the appropriate password values for each property. You can do this either on the command line:

```
-Djavax.net.ssl.keyStore=path_to_keystore_file
-Djavax.net.ssl.keyStorePassword=password
-Djavax.net.ssl.trustStore=path_to_truststore_file
-Djavax.net.ssl.trustStorePassword=password
```

Or you can set the values directly within the application:

```
System.setProperty("javax.net.ssl.keyStore", "path_to_keystore_file");
System.setProperty("javax.net.ssl.keyStorePassword", "password");
System.setProperty("javax.net.ssl.trustStore", "path_to_truststore_file");
System.setProperty("javax.net.ssl.trustStorePassword", "password");
```

You will also need to set `useSSL` to `true` in your connection parameters for MySQL Connector/J, either by adding `useSSL=true` to your URL, or by setting the property `useSSL` to `true` in the `java.util.Properties` instance you pass to `DriverManager.getConnection()`.

You can test that SSL is working by turning on JSSE debugging (as detailed below), and look for the following key events:

```
...
*** ClientHello, v3.1
RandomCookie: GMT: 1018531834 bytes = { 199, 148, 180, 215, 74, 12, »
  54, 244, 0, 168, 55, 103, 215, 64, 16, 138, 225, 190, 132, 153, 2, »
  217, 219, 239, 202, 19, 121, 78 }
Session ID: {}
Cipher Suites: { 0, 5, 0, 4, 0, 9, 0, 10, 0, 18, 0, 19, 0, 3, 0, 17 }
Compression Methods: { 0 }
***
[write] MD5 and SHA1 hashes: len = 59
0000: 01 00 00 37 03 01 3D B6 90 FA C7 94 B4 D7 4A 0C ...7..=.....J.
0010: 36 F4 00 A8 37 67 D7 40 10 8A E1 BE 84 99 02 D9 6...7g.@.....
0020: DB EF CA 13 79 4E 00 00 10 00 05 00 04 00 09 00 ...yN.....
0030: 0A 00 12 00 13 00 03 00 11 01 00 .....
main, WRITE: SSL v3.1 Handshake, length = 59
main, READ: SSL v3.1 Handshake, length = 74
*** ServerHello, v3.1
RandomCookie: GMT: 1018577560 bytes = { 116, 50, 4, 103, 25, 100, 58, »
  202, 79, 185, 178, 100, 215, 66, 254, 21, 83, 187, 190, 42, 170, 3, »
```

```

    132, 110, 82, 148, 160, 92 }
Session ID: {163, 227, 84, 53, 81, 127, 252, 254, 178, 179, 68, 63, »
    182, 158, 30, 11, 150, 79, 170, 76, 255, 92, 15, 226, 24, 17, 177, »
    219, 158, 177, 187, 143}
Cipher Suite: { 0, 5 }
Compression Method: 0
***
%% Created: [Session-1, SSL_RSA_WITH_RC4_128_SHA]
** SSL_RSA_WITH_RC4_128_SHA
[read] MD5 and SHA1 hashes: len = 74
0000: 02 00 00 46 03 01 3D B6 43 98 74 32 04 67 19 64 ...F..=.C.t2.g.d
0010: 3A CA 4F B9 B2 64 D7 42 FE 15 53 BB BE 2A AA 03 :O..d.B..S..*..
0020: 84 6E 52 94 A0 5C 20 A3 E3 54 35 51 7F FC FE B2 .nR..\ ..T5Q....
0030: B3 44 3F B6 9E 1E 0B 96 4F AA 4C FF 5C 0F E2 18 .D?.....O.L.\...
0040: 11 B1 DB 9E B1 BB 8F 00 05 00 .....
main, READ: SSL v3.1 Handshake, length = 1712
...

```

JSSSE provides debugging (to `stdout`) when you set the following system property: - `Djavax.net.debug=all` This will tell you what keystores and truststores are being used, as well as what is going on during the SSL handshake and certificate exchange. It will be helpful when trying to determine what is not working when trying to get an SSL connection to happen.

### 21.3.5.6. Connecting Using PAM Authentication

Java applications using Connector/J 5.1.21 and higher can connect to MySQL servers that use the pluggable authentication module (PAM) authentication scheme.

For PAM authentication to work, you must have the following:

- A MySQL server that supports PAM authentication: a commercial distribution of MySQL 5.5.16 or higher. See [Section 6.3.7.3, “The PAM Authentication Plugin”](#) for more information. Connector/J implements the same cleartext authentication method as in [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).
- SSL capability, as explained in [Section 21.3.5.5, “Connecting Securely Using SSL”](#). Because the PAM authentication scheme sends the original password to the server, the connection to the server must be encrypted.

PAM authentication support is enabled by default in Connector/J 5.1.21 and up, so no extra configuration is needed.

To disable the PAM authentication feature, specify `mysql_clear_password` (the method) or `com.mysql.jdbc.authentication.MySqlClearPasswordPlugin` (the class name) in the comma-separated list of arguments for the `disabledAuthenticationPlugins` connection option. See [Section 21.3.5.1, “Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J”](#) for details about that connection option.

### 21.3.5.7. Using Master/Slave Replication with ReplicationConnection

Connector/J 3.1.7 and higher includes a variant of the driver that will automatically send queries to a read/write master, or a failover or round-robin loadbalanced set of slaves based on the state of `Connection.getReadOnly()`.

An application signals that it wants a transaction to be read-only by calling `Connection.setReadOnly(true)`, this replication-aware connection will use one of the slave connections, which are load-balanced per-vm using a round-robin scheme (a given connection is sticky to a slave unless that slave is removed from service). If you have a write transaction, or if you have a read that is time-sensitive (remember, replication in MySQL is asynchronous), set the connection to be not read-only, by calling `Connection.setReadOnly(false)` and the driver will ensure that further calls are sent to the master MySQL server. The driver takes care of propagating the current state of autocommit, isolation level, and catalog between all of the connections that it uses to accomplish this load balancing functionality.

To enable this functionality, use the `com.mysql.jdbc.ReplicationDriver` class when configuring your application server's connection pool or when creating an instance of a JDBC driver for your standalone application. Because it accepts the same URL format as the standard MySQL JDBC driver, `ReplicationDriver` does not currently work with `java.sql.DriverManager`-based connection creation unless it is the only MySQL JDBC driver registered with the `DriverManager`.

Here is a short example of how `ReplicationDriver` might be used in a standalone application:

```
import java.sql.Connection;
import java.sql.ResultSet;
import java.util.Properties;

import com.mysql.jdbc.ReplicationDriver;

public class ReplicationDriverDemo {

    public static void main(String[] args) throws Exception {
        ReplicationDriver driver = new ReplicationDriver();

        Properties props = new Properties();

        // We want this for failover on the slaves
        props.put("autoReconnect", "true");

        // We want to load balance between the slaves
        props.put("roundRobinLoadBalance", "true");

        props.put("user", "foo");
        props.put("password", "bar");

        //
        // Looks like a normal MySQL JDBC url, with a
        // comma-separated list of hosts, the first
        // being the 'master', the rest being any number
        // of slaves that the driver will load balance against
        //

        Connection conn =
            driver.connect("jdbc:mysql:replication://master,slave1,slave2,slave3/test",
                props);

        //
        // Perform read/write work on the master
        // by setting the read-only flag to "false"
        //

        conn.setReadOnly(false);
        conn.setAutoCommit(false);
        conn.createStatement().executeUpdate("UPDATE some_table ...");
        conn.commit();

        //
        // Now, do a query from a slave, the driver automatically picks one
        // from the list
        //

        conn.setReadOnly(true);

        ResultSet rs =
            conn.createStatement().executeQuery("SELECT a,b FROM alt_table");

        .....
    }
}
```

Consider investigating the Load Balancing JDBC Pool (`lbpool`) tool, which provides a wrapper around the standard JDBC driver and enables you to use DB connection pools that includes checks for system failures and uneven load distribution. For more information, see [Load Balancing JDBC Pool \(lbpool\)](#).

### 21.3.5.8. Mapping MySQL Error Numbers to JDBC SQLState Codes

The table below provides a mapping of the MySQL error numbers to JDBC `SQLState` values.

**Table 21.27. Mapping of MySQL Error Numbers to SQLStates**

MySQL Error Number	MySQL Error Name	Legal (X/Open) SQLState	Standard SQLState
1022	ER_DUPSCHEMA	23000	23000
1037	ER_OUT_OF_MEMORY	SF001	SF001
1038	ER_OUT_OF_SORTMEMORY	SF001	SF001
1040	ER_COND_COUNT_ERROR	23004	23004
1042	ER_BAD_DB_ERROR	42000	42000
1043	ER_HANDSHAKE_ERROR	08S01	08S01
1044	ER_DBACCESS_DENIED_ERROR	42000	42000
1045	ER_ACCESS_DENIED_ERROR	28S01	28S01
1047	ER_UNKNOWN_COOM_ERROR	08S01	08S01
1050	ER_TABLE_EXISTS_ERROR	42000	42000
1051	ER_BAD_TABLE_ERROR	42000	42000
1052	ER_NON_UNIQUE_ERROR	23000	23000
1053	ER_SERVER_SHUTDOWN	08S01	08S01
1054	ER_BAD_FIELD_ERROR	42000	42000
1055	ER_WRONG_FIELD_WITH_GROUP	42000	42000
1056	ER_WRONG_GROUP_FIELD	42000	42000
1057	ER_WRONG_SELECT	42000	42000
1058	ER_WRONG_COUNT	42000	42000
1059	ER_TOO_LONG_IDENT	42000	42000
1060	ER_DUPSCHEMA	23000	23000
1061	ER_DUPSCHEMA	23000	23000
1062	ER_DUPSCHEMA	23000	23000
1063	ER_WRONG_FIELD_SPEC	42000	42000
1064	ER_PARSE_ERROR	42000	42000
1065	ER_EMPTY_QUERY	42000	42000
1066	ER_NON_TEMP_TABLE	42000	42000
1067	ER_INVALID_DEFAULT	42000	42000
1068	ER_MULTIPLE_PRI_KEY	42000	42000
1069	ER_TOO_MANY_KEYS	42000	42000
1070	ER_TOO_MANY_KEY_PARTS	42000	42000
1071	ER_TOO_MANY_KEY	42000	42000
1072	ER_KEY_DOES_NOT_EXISTS	42000	42000
1073	ER_BLOB_AS_KEY	42000	42000
1074	ER_TOO_LONG_FIELD	42000	42000
1075	ER_WRONG_KEY	42000	42000
1080	ER_FORCE_CLOSE	42000	42000
1081	ER_IPS_ERROR	08S01	08S01



MySQL Error Number	MySQL Error Name	Legacy (X/Open) SQL State	SQL Standard SQL State	Description
1082	ER_NO_SUCH_INDEX	22002	22002	INDEX
1083	ER_WRONG_FIELD_TERMINATORS	22002	22002	FIELD_TERMINATORS
1084	ER_BLOCK_SIZE_NO_TERMINATED	22002	22002	NO_TERMINATED
1090	ER_CANNOT_REMOVE_ALL_FIELDS	22002	22002	REMOVE_ALL_FIELDS
1091	ER_CANNOT_DROP_FIELD_OR_KEY	22002	22002	FIELD_OR_KEY
1101	ER_BLOCK_SIZE_HAVE_DEFAULT	22002	22002	HAVE_DEFAULT
1102	ER_WRONG_DB_NAME	22002	22002	DB_NAME
1103	ER_WRONG_TABLE_NAME	22002	22002	TABLE_NAME
1104	ER_TOO_BIG_OBJECT	22002	22002	OBJECT
1106	ER_UNKNOWN_PROCEDURE	22002	22002	PROCEDURE
1107	ER_WRONG_PARAMCOUNT_TO_PROCEDURE	22002	22002	PARAMCOUNT_TO_PROCEDURE
1109	ER_UNKNOWN_TABLE	22002	22002	TABLE
1110	ER_FIELD_SPECIFIED_TWICE	22002	22002	SPECIFIED_TWICE
1112	ER_UNSUPPORTED_EXTENSION	22002	22002	EXTENSION
1113	ER_TABLE_MUST_HAVE_COLUMNS	22002	22002	MUST_HAVE_COLUMNS
1115	ER_UNKNOWN_CHARACTER_SET	22002	22002	CHARACTER_SET
1118	ER_TOO_BIG_ROWSIZE	22002	22002	ROWSIZE
1120	ER_WRONG_ORDER_JOIN	22002	22002	ORDER_JOIN
1121	ER_NULL_IN_INDEX	22002	22002	NULL_IN_INDEX
1129	ER_HOST_IS_BLOCKED	22002	22002	HOST_IS_BLOCKED
1130	ER_HOST_NOT_PRIVILEGED	22002	22002	HOST_NOT_PRIVILEGED
1131	ER_PASSWORD_ANONYMOUS_USER	22002	22002	PASSWORD_ANONYMOUS_USER
1132	ER_PASSWORD_NOT_ALLOWED	22002	22002	PASSWORD_NOT_ALLOWED
1133	ER_PASSWORD_NO_MATCH	22002	22002	PASSWORD_NO_MATCH
1136	ER_WRONG_VALUE_COUNT_ON_ROW	22002	22002	VALUE_COUNT_ON_ROW
1138	ER_INVALID_USE_OF_NULL	22002	22002	INVALID_USE_OF_NULL
1139	ER_REGEXP_ERROR	22002	22002	REGEXP_ERROR
1140	ER_MIXING_GROUP_FUNC_AND_FIELDS	22002	22002	MIXING_GROUP_FUNC_AND_FIELDS
1141	ER_NONEXISTING_GRANT	22002	22002	NONEXISTING_GRANT
1142	ER_TABLEACCESS_DENIED_ERROR	22002	22002	TABLEACCESS_DENIED_ERROR
1143	ER_COLUMNACCESS_DENIED_ERROR	22002	22002	COLUMNACCESS_DENIED_ERROR
1144	ER_ILLEGAL_TABLE_FOR_TABLE	22002	22002	ILLEGAL_TABLE_FOR_TABLE
1145	ER_GRANTING_HOST_OR_USER	22002	22002	GRANTING_HOST_OR_USER
1146	ER_NO_SUCH_TABLE	22002	22002	SUCH_TABLE
1147	ER_NONEXISTING_TABLE_GRANT	22002	22002	NONEXISTING_TABLE_GRANT
1148	ER_NOT_ALLOWED_COMMAND	22002	22002	NOT_ALLOWED_COMMAND
1149	ER_SYNTAX_ERROR	22002	22002	SYNTAX_ERROR
1152	ER_ABORTING_CONNECTION	22002	22002	ABORTING_CONNECTION
1153	ER_NET_PACKET_TOO_LARGE	22002	22002	NET_PACKET_TOO_LARGE

MySQL Error Number	MySQL Error Name	Legacy (X/Open) SQL State	SQL Standard SQL State	SQL State
1154	ER_NET_READ_ERROR_FROM_PIPE	SR002	SR002	SR002
1155	ER_NET_READ_ERROR	SR002	SR002	SR002
1156	ER_NET_PACKETS_OUT_OF_ORDER	SR002	SR002	SR002
1157	ER_NET_SINCE_SOCKETS_ERROR	SR002	SR002	SR002
1158	ER_NET_READ_ERROR	SR002	SR002	SR002
1159	ER_NET_READ_INTERRUPTED	SR002	SR002	SR002
1160	ER_NET_SESSION_WRITE	SR002	SR002	SR002
1161	ER_NET_SOCKET_INTERRUPTED	SR002	SR002	SR002
1162	ER_TOO_LONG_STRING	22003	22003	22003
1163	ER_TABLE_CANT_HANDLE_BLOB	40000	40000	40000
1164	ER_TABLE_CANT_HANDLE_AUTO_INCREMENT	40000	40000	40000
1166	ER_WRONG_COLUMN_NAME	40000	40000	40000
1167	ER_WRONG_COLUMN	40000	40000	40000
1169	ER_DUPES	40000	40000	40000
1170	ER_BLOB_KEY_WITHOUT_LENGTH	40000	40000	40000
1171	ER_PRIMARY_KEY_HAVE_NULL	40000	40000	40000
1172	ER_TOO_MANY_ROWS	40000	40000	40000
1173	ER_REQUEST_PRIMARY_KEY	40000	40000	40000
1177	ER_CHECKSUCH_TABLE	40000	40000	40000
1178	ER_CHECKSUCH_IMPLEMENTED	40000	40000	40000
1179	ER_CANT_DO_THIS_DURING_AN_TRANSACTION	40000	40000	40000
1184	ER_NEW_ABORTING_CONNECTION	40000	40000	40000
1189	ER_MASTER_READ	40000	40000	40000
1190	ER_MASTER_WRITE	40000	40000	40000
1203	ER_TOO_MANY_USER_CONNECTIONS	40000	40000	40000
1205	ER_LOCK_TIMEOUT	40000	40000	40000
1207	ER_READ_ONLY_TRANSACTION	40000	40000	40000
1211	ER_NO_PERMISSION_TO_CREATE_USER	40000	40000	40000
1213	ER_LOCK_DEADLOCK	40000	40000	40000
1216	ER_NO_REFERENCED_ROW	40000	40000	40000
1217	ER_ROW_REFERENCED	40000	40000	40000
1218	ER_CONSISTENT_MASTER	40000	40000	40000
1222	ER_WRONG_NUMBER_OF_COLUMNS_IN_SELECT	40000	40000	40000
1226	ER_USER_REACHED	40000	40000	40000
1230	ER_NO_DEFAULT	40000	40000	40000
1231	ER_WRONG_TYPE_FOR_VAR	40000	40000	40000
1232	ER_WRONG_TYPE_FOR_VAR	40000	40000	40000
1234	ER_CANT_OPTION_HERE	40000	40000	40000
1235	ER_NOT_SUPPORTED_YET	40000	40000	40000

MySQL Error Number	MySQL Error Name	Legacy (X/Open) SQL State	SQL Standard SQL State
1239	ER_WRONG_OBJECT	22000	DEF
1241	ER_OPEN_TABLE	24000	COLUMNS
1242	ER_SUBQUERY_NO_1_ROW	22000	
1247	ER_ILLEGAL_REFERENCE	22000	
1248	ER_DERIVED_HAS_ALIAS	22000	
1249	ER_SELECT_REDUCED	22000	
1250	ER_TABLESPACE_NOT_ALLOWED_HERE	22000	
1251	ER_NOT_SUPPORTED_AUTH_MODE	22000	
1252	ER_SPALIST_NOT_HAVE_NULL	22000	
1253	ER_COLLATE_CHARSET_MISMATCH	22000	
1261	ER_WARN_TOO_FEW_RECORDS	22000	
1262	ER_WARN_TOO_MANY_RECORDS	22000	
1263	ER_WARN_NULL_TO_NOTNULL	22000	
1264	ER_WARN_DATA_OUT_OF_RANGE	22000	
1265	ER_WARN_DATA_TRUNCATED	22000	
1280	ER_WRONG_NAME_FOR_INDEX	22000	
1281	ER_WRONG_NAME_FOR_CATALOG	22000	
1286	ER_UNKNOWN_STORAGE_ENGINE	22000	

### 21.3.6. JDBC Concepts

This section provides some general JDBC background.

#### 21.3.6.1. Connecting to MySQL Using the JDBC `DriverManager` Interface

When you are using JDBC outside of an application server, the `DriverManager` class manages the establishment of Connections.

Specify to the `DriverManager` which JDBC drivers to try to make Connections with. The easiest way to do this is to use `Class.forName()` on the class that implements the `java.sql.Driver` interface. With MySQL Connector/J, the name of this class is `com.mysql.jdbc.Driver`. With this method, you could use an external configuration file to supply the driver class name and driver parameters to use when connecting to a database.

The following section of Java code shows how you might register MySQL Connector/J from the `main()` method of your application. If testing this code, first read the installation section at [Section 21.3.3, “Connector/J Installation”](#), to make sure you have connector installed correctly and the `CLASSPATH` set up. Also, ensure that MySQL is configured to accept external TCP/IP connections.

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

// Notice, do not import com.mysql.jdbc.*
// or you will have problems!

public class LoadDriver {
    public static void main(String[] args) {
        try {
            // The newInstance() call is a work around for some
```

```

        // broken Java implementations

        Class.forName("com.mysql.jdbc.Driver").newInstance();
    } catch (Exception ex) {
        // handle the error
    }
}

```

After the driver has been registered with the `DriverManager`, you can obtain a `Connection` instance that is connected to a particular database by calling `DriverManager.getConnection()`:

#### Example 21.1. Connector/J: Obtaining a connection from the `DriverManager`

If you have not already done so, please review the section [Section 21.3.6.1, "Connecting to MySQL Using the JDBC `DriverManager` Interface"](#) before working with these examples.

This example shows how you can obtain a `Connection` instance from the `DriverManager`. There are a few different signatures for the `getConnection()` method. Consult the API documentation that comes with your JDK for more specific information on how to use them.

```

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

Connection conn = null;
...
try {
    conn =
        DriverManager.getConnection("jdbc:mysql://localhost/test?" +
                                   "user=monty&password=greatsqldb");

    // Do something with the Connection

    ...
} catch (SQLException ex) {
    // handle any errors
    System.out.println("SQLException: " + ex.getMessage());
    System.out.println("SQLState: " + ex.getSQLState());
    System.out.println("VendorError: " + ex.getErrorCode());
}

```

Once a `Connection` is established, it can be used to create `Statement` and `PreparedStatement` objects, as well as retrieve metadata about the database. This is explained in the following sections.

### 21.3.6.2. Using JDBC `Statement` Objects to Execute SQL

`Statement` objects allow you to execute basic SQL queries and retrieve the results through the `ResultSet` class, which is described later.

To create a `Statement` instance, you call the `createStatement()` method on the `Connection` object you have retrieved using one of the `DriverManager.getConnection()` or `DataSource.getConnection()` methods described earlier.

Once you have a `Statement` instance, you can execute a `SELECT` query by calling the `executeQuery(String)` method with the SQL you want to use.

To update data in the database, use the `executeUpdate(String SQL)` method. This method returns the number of rows matched by the update statement, not the number of rows that were modified.

If you do not know ahead of time whether the SQL statement will be a `SELECT` or an `UPDATE/INSERT`, then you can use the `execute(String SQL)` method. This method will return true if the SQL query was a `SELECT`, or false if it was an `UPDATE`, `INSERT`, or `DELETE` statement. If the statement was a `SELECT` query, you can retrieve the results by calling the `getResultSet()` method. If the statement was an `UPDATE`, `INSERT`, or `DELETE` statement, you can retrieve the affected rows count by calling `getUpdateCount()` on the `Statement` instance.

**Example 21.2. Connector/J: Using `java.sql.Statement` to execute a `SELECT` query**

```

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;
import java.sql.Statement;
import java.sql.ResultSet;

// assume that conn is an already created JDBC connection (see previous examples)

Statement stmt = null;
ResultSet rs = null;

try {
    stmt = conn.createStatement();
    rs = stmt.executeQuery("SELECT foo FROM bar");

    // or alternatively, if you don't know ahead of time that
    // the query will be a SELECT...

    if (stmt.execute("SELECT foo FROM bar")) {
        rs = stmt.getResultSet();
    }

    // Now do something with the ResultSet ....
}
catch (SQLException ex){
    // handle any errors
    System.out.println("SQLException: " + ex.getMessage());
    System.out.println("SQLState: " + ex.getSQLState());
    System.out.println("VendorError: " + ex.getErrorCode());
}
finally {
    // it is a good idea to release
    // resources in a finally{} block
    // in reverse-order of their creation
    // if they are no-longer needed

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException sqlEx) { } // ignore

        rs = null;
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) { } // ignore

        stmt = null;
    }
}

```

**21.3.6.3. Using JDBC `CallableStatements` to Execute Stored Procedures**

Starting with MySQL server version 5.0 when used with Connector/J 3.1.1 or newer, the `java.sql.CallableStatement` interface is fully implemented with the exception of the `getParameterMetaData()` method.

For more information on MySQL stored procedures, please refer to <http://dev.mysql.com/doc/mysql/en/stored-routines.html>.

Connector/J exposes stored procedure functionality through JDBC's `CallableStatement` interface.

**Note**

Current versions of MySQL server do not return enough information for the JDBC driver to provide result set metadata for callable statements. This means

that when using `CallableStatement`, `ResultSetMetaData` may return `NULL`.

The following example shows a stored procedure that returns the value of `inOutParam` incremented by 1, and the string passed in using `inputParam` as a `ResultSet`:

### Example 21.3. Connector/J: Calling Stored Procedures

```
CREATE PROCEDURE demoSp(IN inputParam VARCHAR(255), \
                        INOUT inOutParam INT)
BEGIN
  DECLARE z INT;
  SET z = inOutParam + 1;
  SET inOutParam = z;

  SELECT inputParam;

  SELECT CONCAT('zyxw', inputParam);
END
```

To use the `demoSp` procedure with Connector/J, follow these steps:

1. Prepare the callable statement by using `Connection.prepareCall()`.

Notice that you have to use JDBC escape syntax, and that the parentheses surrounding the parameter placeholders are not optional:

### Example 21.4. Connector/J: Using `Connection.prepareCall()`

```
import java.sql.CallableStatement;

...

//
// Prepare a call to the stored procedure 'demoSp'
// with two parameters
//
// Notice the use of JDBC-escape syntax ({call ...})
//

CallableStatement cStmt = conn.prepareCall("{call demoSp(?, ?)}");

cStmt.setString(1, "abcdefg");
```



#### Note

`Connection.prepareCall()` is an expensive method, due to the metadata retrieval that the driver performs to support output parameters. For performance reasons, minimize unnecessary calls to `Connection.prepareCall()` by reusing `CallableStatement` instances in your code.

2. Register the output parameters (if any exist)

To retrieve the values of output parameters (parameters specified as `OUT` or `INOUT` when you created the stored procedure), JDBC requires that they be specified before statement execution using the various `registerOutputParameter()` methods in the `CallableStatement` interface:

### Example 21.5. Connector/J: Registering output parameters

```
import java.sql.Types;

...
//
// Connector/J supports both named and indexed
// output parameters. You can register output
```

```

// parameters using either method, as well
// as retrieve output parameters using either
// method, regardless of what method was
// used to register them.
//
// The following examples show how to use
// the various methods of registering
// output parameters (you should of course
// use only one registration per parameter).
//
//
// Registers the second parameter as output, and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter(2, Types.INTEGER);
//
// Registers the named parameter 'inOutParam', and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter("inOutParam", Types.INTEGER);
...

```

### 3. Set the input parameters (if any exist)

Input and in/out parameters are set as for [PreparedStatement](#) objects. However, [CallableStatement](#) also supports setting parameters by name:

#### Example 21.6. Connector/J: Setting [CallableStatement](#) input parameters

```

...
//
// Set a parameter by index
//
cStmt.setString(1, "abcdefg");
//
// Alternatively, set a parameter using
// the parameter name
//
cStmt.setString("inputParameter", "abcdefg");
//
// Set the 'in/out' parameter using an index
//
cStmt.setInt(2, 1);
//
// Alternatively, set the 'in/out' parameter
// by name
//
cStmt.setInt("inOutParam", 1);
...

```

### 4. Execute the [CallableStatement](#), and retrieve any result sets or output parameters.

Although [CallableStatement](#) supports calling any of the [Statement](#) execute methods ([executeUpdate\(\)](#), [executeQuery\(\)](#) or [execute\(\)](#)), the most flexible method to call is [execute\(\)](#), as you do not need to know ahead of time if the stored procedure returns result sets:

**Example 21.7. Connector/J: Retrieving results and output parameter values**

```

...

boolean hadResults = cStmt.execute();

//
// Process all returned result sets
//

while (hadResults) {
    ResultSet rs = cStmt.getResultSet();

    // process result set
    ...

    hadResults = cStmt.getMoreResults();
}

//
// Retrieve output parameters
//
// Connector/J supports both index-based and
// name-based retrieval
//

int outputValue = cStmt.getInt(2); // index-based

outputValue = cStmt.getInt("inOutParam"); // name-based

...

```

**21.3.6.4. Retrieving `AUTO_INCREMENT` Column Values through JDBC**

Before version 3.0 of the JDBC API, there was no standard way of retrieving key values from databases that supported auto increment or identity columns. With older JDBC drivers for MySQL, you could always use a MySQL-specific method on the `Statement` interface, or issue the query `SELECT LAST_INSERT_ID()` after issuing an `INSERT` to a table that had an `AUTO_INCREMENT` key. Using the MySQL-specific method call isn't portable, and issuing a `SELECT` to get the `AUTO_INCREMENT` key's value requires another round-trip to the database, which isn't as efficient as possible. The following code snippets demonstrate the three different ways to retrieve `AUTO_INCREMENT` values. First, we demonstrate the use of the new JDBC 3.0 method `getGeneratedKeys()` which is now the preferred method to use if you need to retrieve `AUTO_INCREMENT` keys and have access to JDBC 3.0. The second example shows how you can retrieve the same value using a standard `SELECT LAST_INSERT_ID()` query. The final example shows how updatable result sets can retrieve the `AUTO_INCREMENT` value when using the `insertRow()` method.

**Example 21.8. Connector/J: Retrieving `AUTO_INCREMENT` column values using `Statement.getGeneratedKeys()`**

```

Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets assuming you have a
    // Connection 'conn' to a MySQL database already
    // available

    stmt = conn.createStatement(java.sql.ResultSet.TYPE_FORWARD_ONLY,
                                java.sql.ResultSet.CONCUR_UPDATABLE);

    //
    // Issue the DDL queries for the table for this example

```



```

//

stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
stmt.executeUpdate(
    "CREATE TABLE autoIncTutorial ("
    + "priKey INT NOT NULL AUTO_INCREMENT, "
    + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

//
// Insert one row that will generate an AUTO INCREMENT
// key in the 'priKey' field
//
stmt.executeUpdate(
    "INSERT INTO autoIncTutorial (dataField) "
    + "values ('Can I Get the Auto Increment Field?')",
    Statement.RETURN_GENERATED_KEYS);

//
// Example of using Statement.getGeneratedKeys()
// to retrieve the value of an auto-increment
// value
//

int autoIncKeyFromApi = -1;

rs = stmt.getGeneratedKeys();

if (rs.next()) {
    autoIncKeyFromApi = rs.getInt(1);
} else {

    // throw an exception from here
}

rs.close();

rs = null;

System.out.println("Key returned from getGeneratedKeys():"
    + autoIncKeyFromApi);
} finally {

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}
}

```

### Example 21.9. Connector/J: Retrieving `AUTO_INCREMENT` column values using `SELECT LAST_INSERT_ID()`

```

Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets.

```

```

stmt = conn.createStatement();

//
// Issue the DDL queries for the table for this example
//

stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
stmt.executeUpdate(
    "CREATE TABLE autoIncTutorial ("
    + "priKey INT NOT NULL AUTO_INCREMENT, "
    + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

//
// Insert one row that will generate an AUTO INCREMENT
// key in the 'priKey' field
//

stmt.executeUpdate(
    "INSERT INTO autoIncTutorial (dataField) "
    + "values ('Can I Get the Auto Increment Field?')");

//
// Use the MySQL LAST_INSERT_ID()
// function to do the same thing as getGeneratedKeys()
//

int autoIncKeyFromFunc = -1;
rs = stmt.executeQuery("SELECT LAST_INSERT_ID()");

if (rs.next()) {
    autoIncKeyFromFunc = rs.getInt(1);
} else {
    // throw an exception from here
}

rs.close();

System.out.println("Key returned from " +
    "'SELECT LAST_INSERT_ID()': " +
    autoIncKeyFromFunc);
} finally {

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}
}

```

### Example 21.10. Connector/J: Retrieving `AUTO_INCREMENT` column values in `Updatable ResultSets`

```

Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets as well as an 'updatable'

```

```

// one, assuming you have a Connection 'conn' to
// a MySQL database already available
//

stmt = conn.createStatement(java.sql.ResultSet.TYPE_FORWARD_ONLY,
                             java.sql.ResultSet.CONCUR_UPDATABLE);

//
// Issue the DDL queries for the table for this example
//

stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
stmt.executeUpdate(
    "CREATE TABLE autoIncTutorial ( "
    + "priKey INT NOT NULL AUTO_INCREMENT, "
    + "dataField VARCHAR(64), PRIMARY KEY (priKey)");

//
// Example of retrieving an AUTO INCREMENT key
// from an updatable result set
//

rs = stmt.executeQuery("SELECT priKey, dataField "
    + "FROM autoIncTutorial");

rs.moveToInsertRow();

rs.updateString("dataField", "AUTO INCREMENT here?");
rs.insertRow();

//
// the driver adds rows at the end
//

rs.last();

//
// We should now be on the row we just inserted
//

int autoIncKeyFromRS = rs.getInt("priKey");

rs.close();

rs = null;

System.out.println("Key returned for inserted row: "
    + autoIncKeyFromRS);
} finally {

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}
}

```

Running the preceding example code should produce the following output:

```

Key returned from getGeneratedKeys(): 1
Key returned from SELECT LAST_INSERT_ID(): 1

```

```
Key returned for inserted row: 2
```

At times, it can be tricky to use the `SELECT LAST_INSERT_ID()` query, as that function's value is scoped to a connection. So, if some other query happens on the same connection, the value is overwritten. On the other hand, the `getGeneratedKeys()` method is scoped by the `Statement` instance, so it can be used even if other queries happen on the same connection, but not on the same `Statement` instance.

### 21.3.7. Connection Pooling with Connector/J

Connection pooling is a technique of creating and managing a pool of connections that are ready for use by any `thread` that needs them. Connection pooling can greatly increase the performance of your Java application, while reducing overall resource usage.

#### How Connection Pooling Works

Most applications only need a thread to have access to a JDBC connection when they are actively processing a `transaction`, which often takes only milliseconds to complete. When not processing a transaction, the connection sits idle. Connection pooling enables the idle connection to be used by some other thread to do useful work.

In practice, when a thread needs to do work against a MySQL or other database with JDBC, it requests a connection from the pool. When the thread is finished using the connection, it returns it to the pool, so that it can be used by any other threads.

When the connection is loaned out from the pool, it is used exclusively by the thread that requested it. From a programming point of view, it is the same as if your thread called `DriverManager.getConnection()` every time it needed a JDBC connection. With connection pooling, your thread may end up using either a new connection or an already-existing connection.

#### Benefits of Connection Pooling

The main benefits to connection pooling are:

- Reduced connection creation time.

Although this is not usually an issue with the quick connection setup that MySQL offers compared to other databases, creating new JDBC connections still incurs networking and JDBC driver overhead that will be avoided if connections are recycled.

- Simplified programming model.

When using connection pooling, each individual thread can act as though it has created its own JDBC connection, allowing you to use straightforward JDBC programming techniques.

- Controlled resource usage.

If you create a new connection every time a thread needs one, rather than using connection pooling, your application's resource usage can be wasteful and lead to unpredictable behavior under load.

#### Using Connection Pooling with Connector/J

Sun has standardized the concept of connection pooling in JDBC through the JDBC 2.0 Optional interfaces, and all major application servers have implementations of these APIs that work with MySQL Connector/J.

Generally, you configure a connection pool in your application server configuration files, and access it through the Java Naming and Directory Interface (JNDI). The following code shows how you might use a connection pool from an application deployed in a J2EE application server:

**Example 21.11. Connector/J: Using a connection pool with a J2EE application server**

```
import java.sql.Connection;
import java.sql.SQLException;
import java.sql.Statement;

import javax.naming.InitialContext;
import javax.sql.DataSource;

public class MyServletJspOrEjb {

    public void doSomething() throws Exception {
        /*
         * Create a JNDI Initial context to be able to
         * lookup the DataSource
         *
         * In production-level code, this should be cached as
         * an instance or static variable, as it can
         * be quite expensive to create a JNDI context.
         *
         * Note: This code only works when you are using servlets
         * or EJBs in a J2EE application server. If you are
         * using connection pooling in standalone Java code, you
         * will have to create/configure datasources using whatever
         * mechanisms your particular connection pooling library
         * provides.
         */

        InitialContext ctx = new InitialContext();

        /*
         * Lookup the DataSource, which will be backed by a pool
         * that the application server provides. DataSource instances
         * are also a good candidate for caching as an instance
         * variable, as JNDI lookups can be expensive as well.
         */

        DataSource ds =
            (DataSource)ctx.lookup("java:comp/env/jdbc/MySQLDB");

        /*
         * The following code is what would actually be in your
         * Servlet, JSP or EJB 'service' method...where you need
         * to work with a JDBC connection.
         */

        Connection conn = null;
        Statement stmt = null;

        try {
            conn = ds.getConnection();

            /*
             * Now, use normal JDBC programming to work with
             * MySQL, making sure to close each resource when you're
             * finished with it, which permits the connection pool
             * resources to be recovered as quickly as possible
             */

            stmt = conn.createStatement();
            stmt.execute("SOME SQL QUERY");

            stmt.close();
            stmt = null;

            conn.close();
            conn = null;
        } finally {
            /*
             * close any jdbc instances here that weren't
             * explicitly closed during normal code path, so

```

```
    * that we don't 'leak' resources...
    */

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) {
            // ignore, as we can't do anything about it here
        }

        stmt = null;
    }

    if (conn != null) {
        try {
            conn.close();
        } catch (SQLException sqlEx) {
            // ignore, as we can't do anything about it here
        }

        conn = null;
    }
}
}
```

As shown in the example above, after obtaining the JNDI `InitialContext`, and looking up the `DataSource`, the rest of the code follows familiar JDBC conventions.

When using connection pooling, always make sure that connections, and anything created by them (such as statements or result sets) are closed. This rule applies no matter what happens in your code (exceptions, flow-of-control, and so forth). When these objects are closed, they can be re-used; otherwise, they will be stranded, which means that the MySQL server resources they represent (such as buffers, locks, or sockets) are tied up for some time, or in the worst case can be tied up forever.

## Sizing the Connection Pool

Each connection to MySQL has overhead (memory, CPU, context switches, and so forth) on both the client and server side. Every connection limits how many resources there are available to your application as well as the MySQL server. Many of these resources will be used whether or not the connection is actually doing any useful work! Connection pools can be tuned to maximize performance, while keeping resource utilization below the point where your application will start to fail rather than just run slower.

The optimal size for the connection pool depends on anticipated load and average database transaction time. In practice, the optimal connection pool size can be smaller than you might expect. If you take Sun's Java Petstore blueprint application for example, a connection pool of 15-20 connections can serve a relatively moderate load (600 concurrent users) using MySQL and Tomcat with acceptable response times.

To correctly size a connection pool for your application, create load test scripts with tools such as Apache JMeter or The Grinder, and load test your application.

An easy way to determine a starting point is to configure your connection pool's maximum number of connections to be unbounded, run a load test, and measure the largest amount of concurrently used connections. You can then work backward from there to determine what values of minimum and maximum pooled connections give the best performance for your particular application.

## Validating Connections

MySQL Connector/J can validate the connection by executing a lightweight ping against a server. In the case of load-balanced connections, this is performed against all active pooled internal connections that are retained. This is beneficial to Java applications using connection pools, as the pool can

use this feature to validate connections. Depending on your connection pool and configuration, this validation can be carried out at different times:

1. Before the pool returns a connection to the application.
2. When the application returns a connection to the pool.
3. During periodic checks of idle connections.

To use this feature, specify a validation query in your connection pool that starts with `/* ping */`. Note that the syntax must be exactly as specified. This will cause the driver send a ping to the server and return a dummy lightweight result set. When using a [ReplicationConnection](#) or [LoadBalancedConnection](#), the ping will be sent across all active connections.

It is critical that the syntax be specified correctly. The syntax needs to be exact for reasons of efficiency, as this test is done for every statement that is executed:

```
protected static final String PING_MARKER = "/* ping */";
...
if (sql.charAt(0) == '/') {
if (sql.startsWith(PING_MARKER)) {
doPingInstead();
...
}
```

None of the following snippets will work, because the ping syntax is sensitive to whitespace, capitalization, and placement:

```
sql = "/* PING */ SELECT 1";
sql = "SELECT 1 /* ping*/";
sql = "/*ping*/ SELECT 1";
sql = " /* ping */ SELECT 1";
sql = "/*to ping or not to ping*/ SELECT 1";
```

All of the previous statements will issue a normal [SELECT](#) statement and will **not** be transformed into the lightweight ping. Further, for load-balanced connections, the statement will be executed against one connection in the internal pool, rather than validating each underlying physical connection. This results in the non-active physical connections assuming a stale state, and they may die. If Connector/J then re-balances, it might select a dead connection, resulting in an exception being passed to the application. To help prevent this, you can use [loadBalanceValidateConnectionOnSwapServer](#) to validate the connection before use.

If your Connector/J deployment uses a connection pool that allows you to specify a validation query, take advantage of it, but ensure that the query starts *exactly* with `/* ping */`. This is particularly important if you are using the load-balancing or replication-aware features of Connector/J, as it will help keep alive connections which otherwise will go stale and die, causing problems later.

### 21.3.8. Load Balancing with Connector/J

Connector/J has long provided an effective means to distribute read/write load across multiple MySQL server instances for Cluster or master-master replication deployments. Starting with Connector/J 5.1.3, you can now dynamically configure load-balanced connections, with no service outage. In-process transactions are not lost, and no application exceptions are generated if any application is trying to use that particular server instance.

There are two connection string options associated with this functionality:

- [loadBalanceConnectionGroup](#) – This provides the ability to group connections from different sources. This allows you to manage these JDBC sources within a single class loader in any combination you choose. If they use the same configuration, and you want to manage them as a logical single group, give them the same name. This is the key property for management: if you do not define a name (string) for [loadBalanceConnectionGroup](#), you cannot manage the

connections. All load-balanced connections sharing the same `loadBalanceConnectionGroup` value, regardless of how the application creates them, will be managed together.

- `loadBalanceEnableJMX` – The ability to manage the connections is exposed when you define a `loadBalanceConnectionGroup`, but if you want to manage this externally, enable JMX by setting this property to `true`. This enables a JMX implementation, which exposes the management and monitoring operations of a connection group. Further, start your application with the `-Dcom.sun.management.jmxremote` JVM flag. You can then perform connect and perform operations using a JMX client such as `jconsole`.

Once a connection has been made using the correct connection string options, a number of monitoring properties are available:

- Current active host count.
- Current active physical connection count.
- Current active logical connection count.
- Total logical connections created.
- Total transaction count.

The following management operations can also be performed:

- Add host.
- Remove host.

The JMX interface, `com.mysql.jdbc.jmx.LoadBalanceConnectionGroupManagerMBean`, has the following methods:

- `int getActiveHostCount(String group);`
- `int getTotalHostCount(String group);`
- `long getTotalLogicalConnectionCount(String group);`
- `long getActiveLogicalConnectionCount(String group);`
- `long getActivePhysicalConnectionCount(String group);`
- `long getTotalPhysicalConnectionCount(String group);`
- `long getTotalTransactionCount(String group);`
- `void removeHost(String group, String host) throws SQLException;`
- `void stopNewConnectionsToHost(String group, String host) throws SQLException;`
- `void addHost(String group, String host, boolean forExisting);`
- `String getActiveHostsList(String group);`
- `String getRegisteredConnectionGroups();`

The `getRegisteredConnectionGroups()` method returns the names of all connection groups defined in that class loader.

You can test this setup with the following code:

---



```

public class Test {

    private static String URL = "jdbc:mysql:loadbalance://" +
        "localhost:3306,localhost:3310/test?" +
        "loadBalanceConnectionGroup=first&loadBalanceEnableJMX=true";

    public static void main(String[] args) throws Exception {
        new Thread(new Repeater()).start();
        new Thread(new Repeater()).start();
        new Thread(new Repeater()).start();
    }

    static Connection getNewConnection() throws SQLException, ClassNotFoundException {
        Class.forName("com.mysql.jdbc.Driver");
        return DriverManager.getConnection(URL, "root", "");
    }

    static void executeSimpleTransaction(Connection c, int conn, int trans){
        try {
            c.setAutoCommit(false);
            Statement s = c.createStatement();
            s.executeQuery("SELECT SLEEP(1) /* Connection: " + conn + ", transaction: " + trans + " */");
            c.commit();
        } catch (SQLException e) {
            e.printStackTrace();
        }
    }

    public static class Repeater implements Runnable {
        public void run() {
            for(int i=0; i < 100; i++){
                try {
                    Connection c = getNewConnection();
                    for(int j=0; j < 10; j++){
                        executeSimpleTransaction(c, i, j);
                        Thread.sleep(Math.round(100 * Math.random()));
                    }
                    c.close();
                    Thread.sleep(100);
                } catch (Exception e) {
                    e.printStackTrace();
                }
            }
        }
    }
}

```

After compiling, the application can be started with the `-Dcom.sun.management.jmxremote` flag, to enable remote management. `jconsole` can then be started. The `Test` main class will be listed by `jconsole`. Select this and click **Connect**. You can then navigate to the `com.mysql.jdbc.jmx.LoadBalanceConnectionGroupManager` bean. At this point, you can click on various operations and examine the returned result.

If you now had an additional instance of MySQL running on port 3309, you could ensure that Connector/J starts using it by using the `addHost()`, which is exposed in `jconsole`. Note that these operations can be performed dynamically without having to stop the application running.

For further information on the combination of load balancing and failover, see [Section 21.3.9, “Failover with Connector/J”](#).

### 21.3.9. Failover with Connector/J

Connector/J provides a useful load-balancing implementation for Cluster or multi-master deployments, as explained in [Section 21.3.8, “Load Balancing with Connector/J”](#). As of Connector/J 5.1.12, this same implementation is used for balancing load between read-only slaves with `ReplicationDriver`. When trying to balance workload between multiple servers, the driver has to determine when it is safe to swap servers, doing so in the middle of a transaction, for example, could cause problems. It is

important not to lose state information. For this reason, Connector/J will only try to pick a new server when one of the following happens:

1. At transaction boundaries (transactions are explicitly committed or rolled back).
2. A communication exception (SQL State starting with "08") is encountered.
3. When a `SQLException` matches conditions defined by user, using the extension points defined by the `loadBalanceSQLStateFailover`, `loadBalanceSQLExceptionSubclassFailover` or `loadBalanceExceptionChecker` properties.

The third condition revolves around three new properties introduced with Connector/J 5.1.13. It allows you to control which `SQLExceptions` trigger failover.

- `loadBalanceExceptionChecker` - The `loadBalanceExceptionChecker` property is really the key. This takes a fully-qualified class name which implements the new `com.mysql.jdbc.LoadBalanceExceptionChecker` interface. This interface is very simple, and you only need to implement the following method:

```
public boolean shouldExceptionTriggerFailover(SQLException ex)
```

A `SQLException` is passed in, and a boolean returned. A value of `true` triggers a failover, `false` does not.

You can use this to implement your own custom logic. An example where this might be useful is when dealing with transient errors when using MySQL Cluster, where certain buffers may become overloaded. The following code snippet illustrates this:

```
public class NdbLoadBalanceExceptionChecker
    extends StandardLoadBalanceExceptionChecker {

    public boolean shouldExceptionTriggerFailover(SQLException ex) {
        return super.shouldExceptionTriggerFailover(ex)
            || checkNdbException(ex);
    }

    private boolean checkNdbException(SQLException ex){
        // Have to parse the message since most NDB errors
        // are mapped to the same DEMC.
        return (ex.getMessage().startsWith("Lock wait timeout exceeded") ||
            (ex.getMessage().startsWith("Got temporary error")
            && ex.getMessage().endsWith("from NDB")));
    }
}
```

The code above extends `com.mysql.jdbc.StandardLoadBalanceExceptionChecker`, which is the default implementation. There are a few convenient shortcuts built into this, for those who want to have some level of control using properties, without writing Java code. This default implementation uses the two remaining properties: `loadBalanceSQLStateFailover` and `loadBalanceSQLExceptionSubclassFailover`.

- `loadBalanceSQLStateFailover` - allows you to define a comma-delimited list of `SQLState` code prefixes, against which a `SQLException` is compared. If the prefix matches, failover is triggered. So, for example, the following would trigger a failover if a given `SQLException` starts with "00", or is "12345":

```
loadBalanceSQLStateFailover=00,12345
```

- `loadBalanceSQLExceptionSubclassFailover` - can be used in conjunction with `loadBalanceSQLStateFailover` or on its own. If you want certain subclasses of `SQLException` to trigger failover, simply provide a comma-delimited list of fully-qualified class or interface names to check against. For example, if you want all `SQLTransientConnectionExceptions` to trigger failover, you would specify:

```
loadBalanceSQLExceptionSubclassFailover=java.sql.SQLException
```

While the three fail-over conditions enumerated earlier suit most situations, if `autoCommit` is enabled, Connector/J never re-balances, and continues using the same physical connection. This can be problematic, particularly when load-balancing is being used to distribute read-only load across multiple slaves. However, Connector/J can be configured to re-balance after a certain number of statements are executed, when `autoCommit` is enabled. This functionality is dependent upon the following properties:

- `loadBalanceAutoCommitStatementThreshold` – defines the number of matching statements which will trigger the driver to potentially swap physical server connections. The default value, 0, retains the behavior that connections with `autoCommit` enabled are never balanced.
- `loadBalanceAutoCommitStatementRegex` – the regular expression against which statements must match. The default value, blank, matches all statements. So, for example, using the following properties will cause Connector/J to re-balance after every third statement that contains the string “test”:

```
loadBalanceAutoCommitStatementThreshold=3
loadBalanceAutoCommitStatementRegex=.*test.*
```

`loadBalanceAutoCommitStatementRegex` can prove useful in a number of situations. Your application may use temporary tables, server-side session state variables, or connection state, where letting the driver arbitrarily swap physical connections before processing is complete could cause data loss or other problems. This allows you to identify a trigger statement that is only executed when it is safe to swap physical connections.

### 21.3.10. Using the Connector/J Interceptor Classes

An interceptor is a software design pattern that provides a transparent way to extend or modify some aspect of a program, similar to a user exit. No recompiling is required. With Connector/J, the interceptors are enabled and disabled by updating the connection string to refer to different sets of interceptor classes that you instantiate.

The connection properties that control the interceptors are explained in [Section 21.3.5.1, “Driver/ Datasource Class Names, URL Syntax and Configuration Properties for Connector/J”](#):

- `connectionLifecycleInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.jdbc.ConnectionLifecycleInterceptor` interface. In these kinds of interceptor classes, you might log events such as rollbacks, measure the time between transaction start and end, or count events such as calls to `setAutoCommit()`.
- `exceptionInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.jdbc.ExceptionInterceptor` interface. In these kinds of interceptor classes, you might add extra diagnostic information to exceptions that can have multiple causes or indicate a problem with server settings. Because `exceptionInterceptors` classes are only called when handling a `SQLException` thrown from Connector/J code, they can be used even in production deployments without substantial performance overhead.
- `statementInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.jdbc.StatementInterceptorV2` interface. In these kinds of interceptor classes, you might change or augment the processing done by certain kinds of statements, such as automatically checking for queried data in a `memcached` server, rewriting slow queries, logging information about statement execution, or route requests to remote servers.

### 21.3.11. Using Connector/J with Tomcat

The following instructions are based on the instructions for Tomcat-5.x, available at <http://tomcat.apache.org/tomcat-5.5-doc/jndi-datasource-examples-howto.html> which is current at the time this document was written.

First, install the `.jar` file that comes with Connector/J in `$(CATALINA_HOME)/common/lib` so that it is available to all applications installed in the container.

Next, configure the JNDI DataSource by adding a declaration resource to `$(CATALINA_HOME)/conf/server.xml` in the context that defines your web application:

```
<Context ....>

...

<Resource name="jdbc/MySQLDB"
          auth="Container"
          type="javax.sql.DataSource"/>

<ResourceParams name="jdbc/MySQLDB">
  <parameter>
    <name>factory</name>
    <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
  </parameter>

  <parameter>
    <name>maxActive</name>
    <value>10</value>
  </parameter>

  <parameter>
    <name>maxIdle</name>
    <value>5</value>
  </parameter>

  <parameter>
    <name>validationQuery</name>
    <value>SELECT 1</value>
  </parameter>

  <parameter>
    <name>testOnBorrow</name>
    <value>true</value>
  </parameter>

  <parameter>
    <name>testWhileIdle</name>
    <value>true</value>
  </parameter>

  <parameter>
    <name>timeBetweenEvictionRunsMillis</name>
    <value>10000</value>
  </parameter>

  <parameter>
    <name>minEvictableIdleTimeMillis</name>
    <value>60000</value>
  </parameter>

  <parameter>
    <name>username</name>
    <value>someuser</value>
  </parameter>

  <parameter>
    <name>password</name>
    <value>somepass</value>
  </parameter>

  <parameter>
    <name>driverClassName</name>
    <value>com.mysql.jdbc.Driver</value>
  </parameter>

  <parameter>
    <name>url</name>
```

```

    <value>jdbc:mysql://localhost:3306/test</value>
  </parameter>

</ResourceParams>
</Context>

```

Note that Connector/J 5.1.3 introduced a facility whereby, rather than use a `validationQuery` value of `SELECT 1`, it is possible to use `validationQuery` with a value set to `/* ping */`. This sends a ping to the server which then returns a fake result set. This is a lighter weight solution. It also has the advantage that if using `ReplicationConnection` or `LoadBalancedConnection` type connections, the ping will be sent across all active connections. The following XML snippet illustrates how to select this option:

```

<parameter>
  <name>validationQuery</name>
  <value>/* ping */</value>
</parameter>

```

Note that `/* ping */` has to be specified exactly.

In general, follow the installation instructions that come with your version of Tomcat, as the way you configure datasources in Tomcat changes from time to time, and if you use the wrong syntax in your XML file, you will most likely end up with an exception similar to the following:

```

Error: java.sql.SQLException: Cannot load JDBC driver class 'null ' SQL
state: null

```

Note that the auto-loading of drivers having the `META-INF/service/java.sql.Driver` class in JDBC 4.0 causes an improper undeployment of the Connector/J driver in Tomcat on Windows. Namely, the Connector/J jar remains locked. This is an initialization problem that is not related to the driver. The possible workarounds, if viable, are as follows: use `antiResourceLocking=true` as a Tomcat Context attribute, or remove the `META-INF/` directory.

## 21.3.12. Using Connector/J with JBoss

These instructions cover JBoss-4.x. To make the JDBC driver classes available to the application server, copy the `.jar` file that comes with Connector/J to the `lib` directory for your server configuration (which is usually called `default`). Then, in the same configuration directory, in the subdirectory named `deploy`, create a datasource configuration file that ends with `-ds.xml`, which tells JBoss to deploy this file as a JDBC Datasource. The file should have the following contents:

```

<datasources>
  <local-tx-datasource>

    <jndi-name>MySQLDB</jndi-name>
    <connection-url>jdbc:mysql://localhost:3306/dbname</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <user-name>user</user-name>
    <password>pass</password>

    <min-pool-size>5</min-pool-size>

    <max-pool-size>20</max-pool-size>

    <idle-timeout-minutes>5</idle-timeout-minutes>

    <exception-sorter-class-name>
com.mysql.jdbc.integration.jboss.ExtendedMysqlExceptionSorter
    </exception-sorter-class-name>
    <valid-connection-checker-class-name>
com.mysql.jdbc.integration.jboss.MysqlValidConnectionChecker
    </valid-connection-checker-class-name>

  </local-tx-datasource>
</datasources>

```

### 21.3.13. Using Connector/J with Spring

The Spring Framework is a Java-based application framework designed for assisting in application design by providing a way to configure components. The technique used by Spring is a well known design pattern called Dependency Injection (see [Inversion of Control Containers and the Dependency Injection pattern](#)). This article will focus on Java-oriented access to MySQL databases with Spring 2.0. For those wondering, there is a .NET port of Spring appropriately named Spring.NET.

Spring is not only a system for configuring components, but also includes support for aspect oriented programming (AOP). This is one of the main benefits and the foundation for Spring's resource and transaction management. Spring also provides utilities for integrating resource management with JDBC and Hibernate.

For the examples in this section the MySQL world sample database will be used. The first task is to set up a MySQL data source through Spring. Components within Spring use the "bean" terminology. For example, to configure a connection to a MySQL server supporting the world sample database, you might use:

```
<util:map id="dbProps">
  <entry key="db.driver" value="com.mysql.jdbc.Driver"/>
  <entry key="db.jdbcurl" value="jdbc:mysql://localhost/world"/>
  <entry key="db.username" value="myuser"/>
  <entry key="db.password" value="mypass"/>
</util:map>
```

In the above example, we are assigning values to properties that will be used in the configuration. For the datasource configuration:

```
<bean id="dataSource"
  class="org.springframework.jdbc.datasource.DriverManagerDataSource">
  <property name="driverClassName" value="${db.driver}"/>
  <property name="url" value="${db.jdbcurl}"/>
  <property name="username" value="${db.username}"/>
  <property name="password" value="${db.password}"/>
</bean>
```

The placeholders are used to provide values for properties of this bean. This means that you can specify all the properties of the configuration in one place instead of entering the values for each property on each bean. We do, however, need one more bean to pull this all together. The last bean is responsible for actually replacing the placeholders with the property values.

```
<bean
  class="org.springframework.beans.factory.config.PropertyPlaceholderConfigurer">
  <property name="properties" ref="dbProps"/>
</bean>
```

Now that we have our MySQL data source configured and ready to go, we write some Java code to access it. The example below will retrieve three random cities and their corresponding country using the data source we configured with Spring.

```
// Create a new application context. this processes the Spring config
ApplicationContext ctx =
  new ClassPathXmlApplicationContext("exlappContext.xml");
// Retrieve the data source from the application context
DataSource ds = (DataSource) ctx.getBean("dataSource");
// Open a database connection using Spring's DataSourceUtils
Connection c = DataSourceUtils.getConnection(ds);
try {
  // retrieve a list of three random cities
  PreparedStatement ps = c.prepareStatement(
    "select City.Name as 'City', Country.Name as 'Country' " +
    "from City inner join Country on City.CountryCode = Country.Code " +
```

```

        "order by rand() limit 3");
    ResultSet rs = ps.executeQuery();
    while(rs.next()) {
        String city = rs.getString("City");
        String country = rs.getString("Country");
        System.out.printf("The city %s is in %s\n", city, country);
    }
} catch (SQLException ex) {
    // something has failed and we print a stack trace to analyse the error
    ex.printStackTrace();
    // ignore failure closing connection
    try { c.close(); } catch (SQLException e) { }
} finally {
    // properly release our connection
    DataSourceUtils.releaseConnection(c, ds);
}

```

This is very similar to normal JDBC access to MySQL with the main difference being that we are using `DataSourceUtils` instead of the `DriverManager` to create the connection.

While it may seem like a small difference, the implications are somewhat far reaching. Spring manages this resource in a way similar to a container managed data source in a J2EE application server. When a connection is opened, it can be subsequently accessed in other parts of the code if it is synchronized with a transaction. This makes it possible to treat different parts of your application as transactional instead of passing around a database connection.

### 21.3.13.1. Using `JdbcTemplate`

Spring makes extensive use of the Template method design pattern (see [Template Method Pattern](#)). Our immediate focus will be on the `JdbcTemplate` and related classes, specifically `NamedParameterJdbcTemplate`. The template classes handle obtaining and releasing a connection for data access when one is needed.

The next example shows how to use `NamedParameterJdbcTemplate` inside of a DAO (Data Access Object) class to retrieve a random city given a country code.

```

public class Ex2JdbcDao {
    /**
     * Data source reference which will be provided by Spring.
     */
    private DataSource dataSource;

    /**
     * Our query to find a random city given a country code. Notice
     * the ":country" parameter toward the end. This is called a
     * named parameter.
     */
    private String queryString = "select Name from City " +
        "where CountryCode = :country order by rand() limit 1";

    /**
     * Retrieve a random city using Spring JDBC access classes.
     */
    public String getRandomCityByCountryCode(String centryCode) {
        // A template that permits using queries with named parameters
        NamedParameterJdbcTemplate template =
            new NamedParameterJdbcTemplate(dataSource);
        // A java.util.Map is used to provide values for the parameters
        Map params = new HashMap();
        params.put("country", centryCode);
        // We query for an Object and specify what class we are expecting
        return (String)template.queryForObject(queryString, params, String.class);
    }

    /**
     * A JavaBean setter-style method to allow Spring to inject the data source.
     * @param dataSource
     */
}

```



```

public void setDataSource(DataSource dataSource) {
    this.dataSource = dataSource;
}
}

```

The focus in the above code is on the `getRandomCityByCountryCode()` method. We pass a country code and use the `NamedParameterJdbcTemplate` to query for a city. The country code is placed in a Map with the key "country", which is the parameter is named in the SQL query.

To access this code, you need to configure it with Spring by providing a reference to the data source.

```

<bean id="dao" class="code.Ex2JdbcDao">
    <property name="dataSource" ref="dataSource"/>
</bean>

```

At this point, we can just grab a reference to the DAO from Spring and call `getRandomCityByCountryCode()`.

```

// Create the application context
ApplicationContext ctx =
new ClassPathXmlApplicationContext("ex2appContext.xml");
// Obtain a reference to our DAO
Ex2JdbcDao dao = (Ex2JdbcDao) ctx.getBean("dao");

String countryCode = "USA";

// Find a few random cities in the US
for(int i = 0; i < 4; ++i)
    System.out.printf("A random city in %s is %s%n", countryCode,
        dao.getRandomCityByCountryCode(countryCode));

```

This example shows how to use Spring's JDBC classes to completely abstract away the use of traditional JDBC classes including `Connection` and `PreparedStatement`.

### 21.3.13.2. Transactional JDBC Access

You might be wondering how we can add transactions into our code if we do not deal directly with the JDBC classes. Spring provides a transaction management package that not only replaces JDBC transaction management, but also enables declarative transaction management (configuration instead of code).

To use transactional database access, we will need to change the storage engine of the tables in the world database. The downloaded script explicitly creates MyISAM tables which do not support transactional semantics. The InnoDB storage engine does support transactions and this is what we will be using. We can change the storage engine with the following statements.

```

ALTER TABLE City ENGINE=InnoDB;
ALTER TABLE Country ENGINE=InnoDB;
ALTER TABLE CountryLanguage ENGINE=InnoDB;

```

A good programming practice emphasized by Spring is separating interfaces and implementations. What this means is that we can create a Java interface and only use the operations on this interface without any internal knowledge of what the actual implementation is. We will let Spring manage the implementation and with this it will manage the transactions for our implementation.

First you create a simple interface:

```

public interface Ex3Dao {
    Integer createCity(String name, String countryCode,
        String district, Integer population);
}

```

This interface contains one method that will create a new city record in the database and return the id of the new record. Next you need to create an implementation of this interface.



```

public class Ex3DaoImpl implements Ex3Dao {
    protected DataSource dataSource;
    protected SqlUpdate updateQuery;
    protected SqlFunction idQuery;

    public Integer createCity(String name, String countryCode,
        String district, Integer population) {
        updateQuery.update(new Object[] { name, countryCode,
            district, population });
        return getLastId();
    }

    protected Integer getLastId() {
        return idQuery.run();
    }
}

```

You can see that we only operate on abstract query objects here and do not deal directly with the JDBC API. Also, this is the complete implementation. All of our transaction management will be dealt with in the configuration. To get the configuration started, we need to create the DAO.

```

<bean id="dao" class="code.Ex3DaoImpl">
    <property name="dataSource" ref="dataSource"/>
    <property name="updateQuery">...</property>
    <property name="idQuery">...</property>
</bean>

```

Now you need to set up the transaction configuration. The first thing you must do is create transaction manager to manage the data source and a specification of what transaction properties are required for the `dao` methods.

```

<bean id="transactionManager"
    class="org.springframework.jdbc.datasource.DataSourceTransactionManager">
    <property name="dataSource" ref="dataSource"/>
</bean>

<tx:advice id="txAdvice" transaction-manager="transactionManager">
    <tx:attributes>
        <tx:method name="*" />
    </tx:attributes>
</tx:advice>

```

The preceding code creates a transaction manager that handles transactions for the data source provided to it. The `txAdvice` uses this transaction manager and the attributes specify to create a transaction for all methods. Finally you need to apply this advice with an AOP pointcut.

```

<aop:config>
    <aop:pointcut id="daoMethods"
        expression="execution(* code.Ex3Dao.*(..))"/>
    <aop:advisor advice-ref="txAdvice" pointcut-ref="daoMethods"/>
</aop:config>

```

This basically says that all methods called on the `Ex3Dao` interface will be wrapped in a transaction. To make use of this, you only have to retrieve the `dao` from the application context and call a method on the `dao` instance.

```

Ex3Dao dao = (Ex3Dao) ctx.getBean("dao");
Integer id = dao.createCity(name, countryCode, district, pop);

```

We can verify from this that there is no transaction management happening in our Java code and it is all configured with Spring. This is a very powerful notion and regarded as one of the most beneficial features of Spring.

### 21.3.13.3. Connection Pooling with Spring

In many situations, such as web applications, there will be a large number of small database transactions. When this is the case, it usually makes sense to create a pool of database connections available for web requests as needed. Although MySQL does not spawn an extra process when a connection is made, there is still a small amount of overhead to create and set up the connection. Pooling of connections also alleviates problems such as collecting large amounts of sockets in the `TIME_WAIT` state.

Setting up pooling of MySQL connections with Spring is as simple as changing the data source configuration in the application context. There are a number of configurations that we can use. The first example is based on the [Jakarta Commons DBCP library](#). The example below replaces the source configuration that was based on `DriverManagerDataSource` with DBCP's `BasicDataSource`.

```
<bean id="dataSource" destroy-method="close"
  class="org.apache.commons.dbcp.BasicDataSource">
  <property name="driverClassName" value="${db.driver}"/>
  <property name="url" value="${db.jdbcurl}"/>
  <property name="username" value="${db.username}"/>
  <property name="password" value="${db.password}"/>
  <property name="initialSize" value="3"/>
</bean>
```

The configuration of the two solutions is very similar. The difference is that DBCP will pool connections to the database instead of creating a new connection every time one is requested. We have also set a parameter here called `initialSize`. This tells DBCP that we want three connections in the pool when it is created.

Another way to configure connection pooling is to configure a data source in our J2EE application server. Using JBoss as an example, you can set up the MySQL connection pool by creating a file called `mysql-local-ds.xml` and placing it in the `server/default/deploy` directory in JBoss. Once we have this setup, we can use JNDI to look it up. With Spring, this lookup is very simple. The data source configuration looks like this.

```
<jee:jndi-lookup id="dataSource" jndi-name="java:MySQL_DS"/>
```

## 21.3.14. Using Connector/J with GlassFish

This section explains how to use MySQL Connector/J with Glassfish™ Server Open Source Edition 3.0.1. Glassfish can be downloaded from the [Glassfish website](#).

Once Glassfish is installed you will need to make sure it can access MySQL Connector/J. To do this copy the MySQL Connector/J JAR file to the directory `GLASSFISH_INSTALL/glassfish/lib`. For example, copy `mysql-connector-java-5.1.12-bin.jar` to `C:\glassfishv3\glassfish\lib`. Restart the Glassfish Application Server.

You are now ready to create JDBC Connection Pools and JDBC Resources.

### Creating a Connection Pool

1. In the Glassfish Administration Console, using the navigation tree navigate to **Resources, JDBC, Connection Pools**.
2. In the **JDBC Connection Pools** frame click **New**. You will enter a two step wizard.
3. In the **Name** field under **General Settings** enter the name for the connection pool, for example enter `MySQLConnPool`.
4. In the **Resource Type** field, select `javax.sql.DataSource` from the drop-down listbox.
5. In the **Database Vendor** field, select `MySQL` from the drop-down listbox. Click **Next** to go to the next page of the wizard.

6. You can accept the default settings for General Settings, Pool Settings and Transactions for this example. Scroll down to Additional Properties.
7. In Additional Properties you will need to ensure the following properties are set:
  - **ServerName** - The server to connect to. For local testing this will be `localhost`.
  - **User** - The user name with which to connect to MySQL.
  - **Password** - The corresponding password for the user.
  - **DatabaseName** - The database to connect to, for example the sample MySQL database `world`.
8. Click **Finish** to exit the wizard. You will be taken to the **JDBC Connection Pools** page where all current connection pools, including the one you just created, will be displayed.
9. In the **JDBC Connection Pools** frame click on the connection pool you just created. Here you can review and edit information about the connection pool.
10. To test your connection pool click the **Ping** button at the top of the frame. A message will be displayed confirming correct operation or otherwise. If an error message is received recheck the previous steps, and ensure that MySQL Connector/J has been correctly copied into the previously specified location.

Now that you have created a connection pool you will also need to create a JDBC Resource (data source) for use by your application.

### Creating a JDBC Resource

Your Java application will usually reference a data source object to establish a connection with the database. This needs to be created first using the following procedure.

- Using the navigation tree in the Glassfish Administration Console, navigate to **Resources**, **JDBC**, **JDBC Resources**. A list of resources will be displayed in the **JDBC Resources** frame.
- Click **New**. The **New JDBC Resource** frame will be displayed.
- In the **JNDI Name** field, enter the JNDI name that will be used to access this resource, for example enter `jdbc/MySQLDataSource`.
- In the **Pool Name** field, select a connection pool you want this resource to use from the drop-down listbox.
- Optionally, you can enter a description into the **Description** field.
- Additional properties can be added if required.
- Click **OK** to create the new JDBC resource. The **JDBC Resources** frame will list all available JDBC Resources.

#### 21.3.14.1. A Simple JSP Application with Glassfish, Connector/J and MySQL

This section shows how to deploy a simple JSP application on Glassfish, that connects to a MySQL database.

This example assumes you have already set up a suitable Connection Pool and JDBC Resource, as explained in the preceding sections. It is also assumed you have a sample database installed, such as `world`.

The main application code, `index.jsp` is presented here:

```

<%@ page import="java.sql.*, javax.sql.*, java.io.*, javax.naming.*" %>
<html>
<head><title>Hello world from JSP</title></head>
<body>
<%
    InitialContext ctx;
    DataSource ds;
    Connection conn;
    Statement stmt;
    ResultSet rs;

    try {
        ctx = new InitialContext();
        ds = (DataSource) ctx.lookup("java:comp/env/jdbc/MySQLDataSource");
        //ds = (DataSource) ctx.lookup("jdbc/MySQLDataSource");
        conn = ds.getConnection();
        stmt = conn.createStatement();
        rs = stmt.executeQuery("SELECT * FROM Country");

        while(rs.next()) {
%>
            <h3>Name: <%= rs.getString("Name") %></h3>
            <h3>Population: <%= rs.getString("Population") %></h3>
<%
        }
    }
    catch (SQLException se) {
%>
        <%= se.getMessage() %>
<%
    }
    catch (NamingException ne) {
%>
        <%= ne.getMessage() %>
<%
    }
%>
</body>
</html>

```

In addition two XML files are required: [web.xml](#), and [sun-web.xml](#). There may be other files present, such as classes and images. These files are organized into the directory structure as follows:

```

index.jsp
WEB-INF
|
- web.xml
- sun-web.xml

```

The code for [web.xml](#) is:

```

<?xml version="1.0" encoding="UTF-8"?>
<web-app version="2.4" xmlns="http://java.sun.com/xml/ns/j2ee" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
<display-name>HelloWebApp</display-name>
<distributable/>
<resource-ref>
    <res-ref-name>jdbc/MySQLDataSource</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
    <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
</web-app>

```

The code for [sun-web.xml](#) is:

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE sun-web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Application Server 8.1 Servlet 2.4//EN" "http:
<sun-web-app>
    <context-root>HelloWebApp</context-root>

```

```
<resource-ref>
  <res-ref-name>jdbc/MySQLDataSource</res-ref-name>
  <jndi-name>jdbc/MySQLDataSource</jndi-name>
</resource-ref>
</sun-web-app>
```

These XML files illustrate a very important aspect of running JDBC applications on Glassfish. On Glassfish it is important to map the string specified for a JDBC resource to its JNDI name, as set up in the Glassfish administration console. In this example, the JNDI name for the JDBC resource, as specified in the Glassfish Administration console when creating the JDBC Resource, was `jdbc/MySQLDataSource`. This must be mapped to the name given in the application. In this example the name specified in the application, `jdbc/MySQLDataSource`, and the JNDI name, happen to be the same, but this does not necessarily have to be the case. Note that the XML element `<res-ref-name>` is used to specify the name as used in the application source code, and this is mapped to the JNDI name specified using the `<jndi-name>` element, in the file `sun-web.xml`. The resource also has to be created in the `web.xml` file, although the mapping of the resource to a JNDI name takes place in the `sun-web.xml` file.

If you do not have this mapping set up correctly in the XML files you will not be able to lookup the data source using a JNDI lookup string such as:

```
ds = (DataSource) ctx.lookup("java:comp/env/jdbc/MySQLDataSource");
```

You will still be able to access the data source directly using:

```
ds = (DataSource) ctx.lookup("jdbc/MySQLDataSource");
```

With the source files in place, in the correct directory structure, you are ready to deploy the application:

1. In the navigation tree, navigate to **Applications** - the **Applications** frame will be displayed. Click **Deploy**.
2. You can now deploy an application packaged into a single WAR file from a remote client, or you can choose a packaged file or directory that is locally accessible to the server. If you are simply testing an application locally you can simply point Glassfish at the directory that contains your application, without needing to package the application into a WAR file.
3. Now select the application type from the **Type** drop-down listbox, which in this example is `Web application`.
4. Click OK.

Now, when you navigate to the **Applications** frame, you will have the option to **Launch**, **Redeploy**, or **Restart** your application. You can test your application by clicking **Launch**. The application will connection to the MySQL database and display the Name and Population of countries in the `Country` table.

### 21.3.14.2. A Simple Servlet with Glassfish, Connector/J and MySQL

This section describes a simple servlet that can be used in the Glassfish environment to access a MySQL database. As with the previous section, this example assumes the sample database `world` is installed.

The project is set up with the following directory structure:

```
index.html
WEB-INF
|
- web.xml
- sun-web.xml
- classes
|
- HelloWorldServlet.java
```

```
- HelloWebServlet.class
```

The code for the servlet, located in `HelloWebServlet.java`, is as follows:

```
import javax.servlet.http.*;
import javax.servlet.*;
import java.io.*;
import java.sql.*;
import javax.sql.*;
import javax.naming.*;

public class HelloWebServlet extends HttpServlet {

    InitialContext ctx = null;
    DataSource ds = null;
    Connection conn = null;
    PreparedStatement ps = null;
    ResultSet rs = null;

    String sql = "SELECT Name, Population FROM Country WHERE Name=?";

    public void init () throws ServletException {
        try {
            ctx = new InitialContext();
            ds = (DataSource) ctx.lookup("java:comp/env/jdbc/MySQLDataSource");
            conn = ds.getConnection();
            ps = conn.prepareStatement(sql);
        }
        catch (SQLException se) {
            System.out.println("SQLException: "+se.getMessage());
        }
        catch (NamingException ne) {
            System.out.println("NamingException: "+ne.getMessage());
        }
    }

    public void destroy () {
        try {
            if (rs != null)
                rs.close();
            if (ps != null)
                ps.close();
            if (conn != null)
                conn.close();
            if (ctx != null)
                ctx.close();
        }
        catch (SQLException se) {
            System.out.println("SQLException: "+se.getMessage());
        }
        catch (NamingException ne) {
            System.out.println("NamingException: "+ne.getMessage());
        }
    }

    public void doPost(HttpServletRequest req, HttpServletResponse resp){
        try {
            String country_name = req.getParameter("country_name");
            resp.setContentType("text/html");
            PrintWriter writer = resp.getWriter();
            writer.println("<html><body>");
            writer.println("<p>Country: "+country_name+"</p>");
            ps.setString(1, country_name);
            rs = ps.executeQuery();
            if (!rs.next()){
                writer.println("<p>Country does not exist!</p>");
            }
            else {
                rs.beforeFirst();
                while(rs.next()) {
                    writer.println("<p>Name: "+rs.getString("Name")+"</p>");
                }
            }
        }
    }
}
```

```

        writer.println("<p>Population: "+rs.getString("Population")+ "</p>");
    }
}
writer.println("</body></html>");
writer.close();
}
catch (Exception e) {
    e.printStackTrace();
}
}

public void doGet(HttpServletRequest req, HttpServletResponse resp){
    try {
        resp.setContentType("text/html");
        PrintWriter writer = resp.getWriter();
        writer.println("<html><body>");
        writer.println("<p>Hello from servlet doGet()</p>");
        writer.println("</body></html>");
        writer.close();
    }
    catch (Exception e) {
        e.printStackTrace();
    }
}
}
}

```

In the preceding code a basic `doGet()` method is implemented, but is not used in the example. The code to establish the connection with the database is as shown in the previous example, [Section 21.3.14.1, “A Simple JSP Application with Glassfish, Connector/J and MySQL”](#), and is most conveniently located in the servlet `init()` method. The corresponding freeing of resources is located in the destroy method. The main functionality of the servlet is located in the `doPost()` method. If the user enters into the input form a country name that can be located in the database, the population of the country is returned. The code is invoked using a POST action associated with the input form. The form is defined in the file `index.html`:

```

<html>
  <head><title>HelloWebServlet</title></head>

  <body>
    <h1>HelloWebServlet</h1>

    <p>Please enter country name:</p>

    <form action="HelloWebServlet" method="POST">
      <input type="text" name="country_name" length="50" />
      <input type="submit" value="Submit" />
    </form>

  </body>
</html>

```

The XML files `web.xml` and `sun-web.xml` are as for the example in the preceding section, [Section 21.3.14.1, “A Simple JSP Application with Glassfish, Connector/J and MySQL”](#), no additional changes are required.

When compiling the Java source code, you will need to specify the path to the file `javaee.jar`. On Windows, this can be done as follows:

```
shell> javac -classpath c:\glassfishv3\glassfish\lib\javaee.jar HelloWebServlet.java
```

Once the code is correctly located within its directory structure, and compiled, the application can be deployed in Glassfish. This is done in exactly the same way as described in the preceding section, [Section 21.3.14.1, “A Simple JSP Application with Glassfish, Connector/J and MySQL”](#).

Once deployed the application can be launched from within the Glassfish Administration Console. Enter a country name such as “England”, and the application will return “Country does not exist!”. Enter “France”, and the application will return a population of 59225700.

## 21.3.15. Troubleshooting Connector/J Applications

This section explains the symptoms and resolutions for the most commonly encountered issues with applications using MySQL Connector/J.

### Questions

- [22.3.15.1: \[2705\]](#) When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

What is going on? I can connect just fine with the MySQL command-line client.

- [22.3.15.2: \[2705\]](#) My application throws an SQLException 'No Suitable Driver'. Why is this happening?
- [22.3.15.3: \[2705\]](#) I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
Is there a MySQL server running on the machine/port you
are trying to connect to?

(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

- [22.3.15.4: \[2706\]](#) I have a servlet/application that works fine for a day, and then stops working overnight
- [22.3.15.5: \[2708\]](#) I'm trying to use JDBC 2.0 updatable result sets, and I get an exception saying my result set is not updatable.
- [22.3.15.6: \[2708\]](#) I cannot connect to the MySQL server using Connector/J, and I'm sure the connection parameters are correct.
- [22.3.15.7: \[2708\]](#) I am trying to connect to my MySQL server within my application, but I get the following error and stack trace:

```
java.net.SocketException
MESSAGE: Software caused connection abort: recv failed

STACKTRACE:

java.net.SocketException: Software caused connection abort: recv failed
at java.net.SocketInputStream.socketRead0(Native Method)
at java.net.SocketInputStream.read(Unknown Source)
at com.mysql.jdbc.MysqlIO.readFully(MysqlIO.java:1392)
at com.mysql.jdbc.MysqlIO.readPacket(MysqlIO.java:1414)
at com.mysql.jdbc.MysqlIO.doHandshake(MysqlIO.java:625)
at com.mysql.jdbc.Connection.createNewIO(Connection.java:1926)
at com.mysql.jdbc.Connection.<init>(Connection.java:452)
at com.mysql.jdbc.NonRegisteringDriver.connect(NonRegisteringDriver.java:411)
```

- [22.3.15.8: \[2709\]](#) My application is deployed through JBoss and I am using transactions to handle the statements on the MySQL database. Under heavy loads, I am getting an error and stack trace, but these only occur after a fixed period of heavy activity.
- [22.3.15.9: \[2709\]](#) When using `gcj`, a `java.io.CharConversionException` exception is raised when working with certain character sequences.
- [22.3.15.10: \[2709\]](#) Updating a table that contains a `primary key` that is either `FLOAT` or compound primary key that uses `FLOAT` fails to update the table and raises an exception.



- [22.3.15.11: \[2709\]](#) You get an `ER_NET_PACKET_TOO_LARGE [3670]` exception, even though the binary blob size you want to insert using JDBC is safely below the `max_allowed_packet [502]` size.
- [22.3.15.12: \[2709\]](#) What should you do if you receive error messages similar to the following: “Communications link failure – Last packet sent to the server was X ms ago”?
- [22.3.15.13: \[2710\]](#) Why does Connector/J not reconnect to MySQL and re-issue the statement after a communication failure, instead of throwing an Exception, even though I use the `autoReconnect` connection string option?
- [22.3.15.14: \[2711\]](#) How can I use 3-byte UTF8 with Connector/J?
- [22.3.15.15: \[2711\]](#) How can I use 4-byte UTF8, `utf8mb4` with Connector/J?
- [22.3.15.16: \[2711\]](#) Using `useServerPrepStmts=false` and certain character encodings can lead to corruption when inserting BLOBs. How can this be avoided?

## Questions and Answers

### 22.3.15.1: When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

#### What is going on? I can connect just fine with the MySQL command-line client.

MySQL Connector/J must use TCP/IP sockets to connect to MySQL, as Java does not support Unix Domain Sockets. Therefore, when MySQL Connector/J connects to MySQL, the security manager in MySQL server will use its grant tables to determine whether the connection is permitted.

You must add the necessary security credentials to the MySQL server for this to happen, using the `GRANT` statement to your MySQL Server. See [Section 13.7.1.4, “GRANT Syntax”](#), for more information.



#### Note

Testing your connectivity with the `mysql` command-line client will not work unless you add the "host" flag, and use something other than `localhost` for the host. The `mysql` command-line client will use Unix domain sockets if you use the special host name `localhost`. If you are testing connectivity to `localhost`, use `127.0.0.1` as the host name instead.



#### Warning

Changing privileges and permissions improperly in MySQL can potentially cause your server installation to not have optimal security properties.

### 22.3.15.2: My application throws an SQLException 'No Suitable Driver'. Why is this happening?

There are three possible causes for this error:

- The Connector/J driver is not in your `CLASSPATH`, see [Section 21.3.3, “Connector/J Installation”](#).
- The format of your connection URL is incorrect, or you are referencing the wrong JDBC driver.
- When using DriverManager, the `jdbc.drivers` system property has not been populated with the location of the Connector/J driver.

### 22.3.15.3: I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
Is there a MySQL server running on the machine/port you
are trying to connect to?
```

```
(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

Either you're running an Applet, your MySQL server has been installed with the "skip-networking" option set, or your MySQL server has a firewall sitting in front of it.

Applets can only make network connections back to the machine that runs the web server that served the .class files for the applet. This means that MySQL must run on the same machine (or you must have some sort of port re-direction) for this to work. This also means that you will not be able to test applets from your local file system, you must always deploy them to a web server.

MySQL Connector/J can only communicate with MySQL using TCP/IP, as Java does not support Unix domain sockets. TCP/IP communication with MySQL might be affected if MySQL was started with the "skip-networking" flag, or if it is firewalled.

If MySQL has been started with the "skip-networking" option set (the Debian Linux package of MySQL server does this for example), you need to comment it out in the file `/etc/mysql/my.cnf` or `/etc/my.cnf`. Of course your `my.cnf` file might also exist in the `data` directory of your MySQL server, or anywhere else (depending on how MySQL was compiled for your system). Binaries created by us always look in `/etc/my.cnf` and `datadir/my.cnf`. If your MySQL server has been firewalled, you will need to have the firewall configured to allow TCP/IP connections from the host where your Java code is running to the MySQL server on the port that MySQL is listening to (by default, 3306).

#### 22.3.15.4: I have a servlet/application that works fine for a day, and then stops working overnight

MySQL closes connections after 8 hours of inactivity. You either need to use a connection pool that handles stale connections or use the `autoReconnect` parameter (see [Section 21.3.5.1, "Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J"](#)).

Also, catch `SQLExceptions` in your application and deal with them, rather than propagating them all the way until your application exits. This is just good programming practice. MySQL Connector/J will set the `SQLState` (see `java.sql.SQLException.getSQLState()` in your API docs) to `08S01` when it encounters network-connectivity issues during the processing of a query. Attempt to reconnect to MySQL at this point.

The following (simplistic) example shows what code that can handle these exceptions might look like:

#### Example 21.12. Connector/J: Example of transaction with retry logic

```
public void doBusinessOp() throws SQLException {
    Connection conn = null;
    Statement stmt = null;
    ResultSet rs = null;

    //
    // How many times do you want to retry the transaction
    // (or at least _getting_ a connection)?
    //
    int retryCount = 5;

    boolean transactionCompleted = false;

    do {
        try {
            conn = getConnection(); // assume getting this from a
                                   // javax.sql.DataSource, or the
                                   // java.sql.DriverManager

            conn.setAutoCommit(false);

            //
            // Okay, at this point, the 'retry-ability' of the
            // transaction really depends on your application logic,
            // whether or not you're using autocommit (in this case
```

```

// not), and whether you're using transactional storage
// engines
//
// For this example, we'll assume that it's _not_ safe
// to retry the entire transaction, so we set retry
// count to 0 at this point
//
// If you were using exclusively transaction-safe tables,
// or your application could recover from a connection going
// bad in the middle of an operation, then you would not
// touch 'retryCount' here, and just let the loop repeat
// until retryCount == 0.
//
retryCount = 0;

stmt = conn.createStatement();

String query = "SELECT foo FROM bar ORDER BY baz";

rs = stmt.executeQuery(query);

while (rs.next()) {
}

rs.close();
rs = null;

stmt.close();
stmt = null;

conn.commit();
conn.close();
conn = null;

transactionCompleted = true;
} catch (SQLException sqlEx) {

//
// The two SQL states that are 'retry-able' are 08S01
// for a communications error, and 40001 for deadlock.
//
// Only retry if the error was due to a stale connection,
// communications problem or deadlock
//

String sqlState = sqlEx.getSQLState();

if ("08S01".equals(sqlState) || "40001".equals(sqlState)) {
    retryCount -= 1;
} else {
    retryCount = 0;
}
} finally {
    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this...
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this as well...
        }
    }

    if (conn != null) {
        try {
            //

```

```

        // If we got here, and conn is not null, the
        // transaction should be rolled back, as not
        // all work has been done

        try {
            conn.rollback();
        } finally {
            conn.close();
        }
    } catch (SQLException sqlEx) {
        //
        // If we got an exception here, something
        // pretty serious is going on, so we better
        // pass it up the stack, rather than just
        // logging it...

        throw sqlEx;
    }
}
} while (!transactionCompleted && (retryCount > 0));
}
}

```



### Note

Use of the `autoReconnect` option is not recommended because there is no safe method of reconnecting to the MySQL server without risking some corruption of the connection state or database state information. Instead, use a connection pool, which will enable your application to connect to the MySQL server using an available connection from the pool. The `autoReconnect` facility is deprecated, and may be removed in a future release.

#### 22.3.15.5: I'm trying to use JDBC 2.0 updatable result sets, and I get an exception saying my result set is not updatable.

Because MySQL does not have row identifiers, MySQL Connector/J can only update result sets that have come from queries on tables that have at least one [primary key](#), the query must select every primary key column, and the query can only span one table (that is, no joins). This is outlined in the JDBC specification.

Note that this issue only occurs when using updatable result sets, and is caused because Connector/J is unable to guarantee that it can identify the correct rows within the result set to be updated without having a unique reference to each row. There is no requirement to have a unique field on a table if you are using `UPDATE` or `DELETE` statements on a table where you can individually specify the criteria to be matched using a `WHERE` clause.

#### 22.3.15.6: I cannot connect to the MySQL server using Connector/J, and I'm sure the connection parameters are correct.

Make sure that the `skip-networking` [440] option has not been enabled on your server. Connector/J must be able to communicate with your server over TCP/IP; named sockets are not supported. Also ensure that you are not filtering connections through a firewall or other network security system. For more information, see [Section C.5.2.2, "Can't connect to \[local\] MySQL server"](#).

#### 22.3.15.7: I am trying to connect to my MySQL server within my application, but I get the following error and stack trace:

```

java.net.SocketException
MESSAGE: Software caused connection abort: recv failed

STACKTRACE:

java.net.SocketException: Software caused connection abort: recv failed
at java.net.SocketInputStream.socketRead0(Native Method)
at java.net.SocketInputStream.read(Unknown Source)

```

```
at com.mysql.jdbc.MySqlIO.readFully(MySqlIO.java:1392)
at com.mysql.jdbc.MySqlIO.readPacket(MySqlIO.java:1414)
at com.mysql.jdbc.MySqlIO.doHandshake(MySqlIO.java:625)
at com.mysql.jdbc.Connection.createNewIO(Connection.java:1926)
at com.mysql.jdbc.Connection.<init>(Connection.java:452)
at com.mysql.jdbc.NonRegisteringDriver.connect(NonRegisteringDriver.java:411)
```

The error probably indicates that you are using a older version of the Connector/J JDBC driver (2.0.14 or 3.0.x) and you are trying to connect to a MySQL server with version 4.1x or newer. The older drivers are not compatible with 4.1 or newer of MySQL as they do not support the newer authentication mechanisms.

It is likely that the older version of the Connector/J driver exists within your application directory or your `CLASSPATH` includes the older Connector/J package.

**22.3.15.8: My application is deployed through JBoss and I am using transactions to handle the statements on the MySQL database. Under heavy loads, I am getting an error and stack trace, but these only occur after a fixed period of heavy activity.**

This is a JBoss, not Connector/J, issue and is connected to the use of transactions. Under heavy loads the time taken for transactions to complete can increase, and the error is caused because you have exceeded the predefined timeout.

You can increase the timeout value by setting the `TransactionTimeout` attribute to the `TransactionManagerService` within the `/conf/jboss-service.xml` file (pre-4.0.3) or `/deploy/jta-service.xml` for JBoss 4.0.3 or later. See `TransactionTimeout` within the JBoss wiki for more information.

**22.3.15.9: When using gcj, a java.io.CharConversionException exception is raised when working with certain character sequences.**

This is a known issue with `gcj` which raises an exception when it reaches an unknown character or one it cannot convert. Add `useJvmCharsetConverters=true` to your connection string to force character conversion outside of the `gcj` libraries, or try a different JDK.

**22.3.15.10: Updating a table that contains a primary key that is either FLOAT or compound primary key that uses FLOAT fails to update the table and raises an exception.**

Connector/J adds conditions to the `WHERE` clause during an `UPDATE` to check the old values of the primary key. If there is no match, then Connector/J considers this a failure condition and raises an exception.

The problem is that rounding differences between supplied values and the values stored in the database may mean that the values never match, and hence the update fails. The issue will affect all queries, not just those from Connector/J.

To prevent this issue, use a primary key that does not use `FLOAT`. If you have to use a floating point column in your primary key, use `DOUBLE` or `DECIMAL` types in place of `FLOAT`.

**22.3.15.11: You get an ER\_NET\_PACKET\_TOO\_LARGE [3670] exception, even though the binary blob size you want to insert using JDBC is safely below the max\_allowed\_packet [502] size.**

This is because the `hexEscapeBlock()` method in `com.mysql.jdbc.PreparedStatement.streamToBytes()` may almost double the size of your data.

**22.3.15.12: What should you do if you receive error messages similar to the following: “Communications link failure – Last packet sent to the server was X ms ago”?**

Generally speaking, this error suggests that the network connection has been closed. There can be several root causes:

- Firewalls or routers may clamp down on idle connections (the MySQL client/server protocol does not ping).
- The MySQL Server may be closing idle connections that exceed the `wait_timeout` or `interactive_timeout` threshold.

To help troubleshoot these issues, the following tips can be used. If a recent (5.1.13+) version of Connector/J is used, you will see an improved level of information compared to earlier versions. Older versions simply display the last time a packet was sent to the server, which is frequently 0 ms ago. This is of limited use, as it may be that a packet was just sent, while a packet from the server has not been received for several hours. Knowing the period of time since Connector/J last received a packet from the server is useful information, so if this is not displayed in your exception message, it is recommended that you update Connector/J.

Further, if the time a packet was last sent/received exceeds the `wait_timeout` or `interactive_timeout` threshold, this is noted in the exception message.

Although network connections can be volatile, the following can be helpful in avoiding problems:

- Ensure connections are valid when used from the connection pool. Use a query that starts with `/* ping */` to execute a lightweight ping instead of full query. Note, the syntax of the ping needs to be exactly as specified here.
- Minimize the duration a connection object is left idle while other application logic is executed.
- Explicitly validate the connection before using it if the connection has been left idle for an extended period of time.
- Ensure that `wait_timeout` and `interactive_timeout` are set sufficiently high.
- Ensure that `tcpKeepalive` is enabled.
- Ensure that any configurable firewall or router timeout settings allow for the maximum expected connection idle time.



#### Note

Do not expect to be able to reuse a connection without problems, if it has been lying idle for a period. If a connection is to be reused after being idle for any length of time, ensure that you explicitly test it before reusing it.

#### 22.3.15.13: Why does Connector/J not reconnect to MySQL and re-issue the statement after a communication failure, instead of throwing an Exception, even though I use the `autoReconnect` connection string option?

There are several reasons for this. The first is transactional integrity. The MySQL Reference Manual states that “there is no safe method of reconnecting to the MySQL server without risking some corruption of the connection state or database state information”. Consider the following series of statements for example:

```
conn.createStatement().execute(
    "UPDATE checking_account SET balance = balance - 1000.00 WHERE customer='Smith'");
conn.createStatement().execute(
    "UPDATE savings_account SET balance = balance + 1000.00 WHERE customer='Smith'");
conn.commit();
```

Consider the case where the connection to the server fails after the `UPDATE` to `checking_account`. If no exception is thrown, and the application never learns about the problem, it will continue executing. However, the server did not commit the first transaction in this case, so that will get rolled back. But execution continues with the next transaction, and increases the `savings_account` balance by 1000. The application did not receive an exception, so it continued regardless, eventually committing the

second transaction, as the commit only applies to the changes made in the new connection. Rather than a transfer taking place, a deposit was made in this example.

Note that running with `autocommit` enabled does not solve this problem. When Connector/J encounters a communication problem, there is no means to determine whether the server processed the currently executing statement or not. The following theoretical states are equally possible:

- The server never received the statement, and therefore no related processing occurred on the server.
- The server received the statement, executed it in full, but the response was not received by the client.

If you are running with `autocommit` enabled, it is not possible to guarantee the state of data on the server when a communication exception is encountered. The statement may have reached the server, or it may not. All you know is that communication failed at some point, before the client received confirmation (or data) from the server. This does not only affect `autocommit` statements though. If the communication problem occurred during `Connection.commit()`, the question arises of whether the transaction was committed on the server before the communication failed, or whether the server received the commit request at all.

The second reason for the generation of exceptions is that transaction-scoped contextual data may be vulnerable, for example:

- Temporary tables.
- User-defined variables.
- Server-side prepared statements.

These items are lost when a connection fails, and if the connection silently reconnects without generating an exception, this could be detrimental to the correct execution of your application.

In summary, communication errors generate conditions that may well be unsafe for Connector/J to simply ignore by silently reconnecting. It is necessary for the application to be notified. It is then for the application developer to decide how to proceed in the event of connection errors and failures.

#### **22.3.15.14: How can I use 3-byte UTF8 with Connector/J?**

To use 3-byte UTF8 with Connector/J set `characterEncoding=utf8` and set `useUnicode=true` in the connection string.

#### **22.3.15.15: How can I use 4-byte UTF8, `utf8mb4` with Connector/J?**

To use 4-byte UTF8 with Connector/J configure the MySQL server with `character_set_server=utf8mb4`. Connector/J will then use that setting as long as `characterEncoding` has not been set in the connection string. This is equivalent to autodetection of the character set.

#### **22.3.15.16: Using `useServerPrepStmts=false` and certain character encodings can lead to corruption when inserting BLOBs. How can this be avoided?**

When using certain character encodings, such as SJIS, CP932, and BIG5, it is possible that BLOB data contains characters that can be interpreted as control characters, for example, backslash, '\'. This can lead to corrupted data when inserting BLOBs into the database. There are two things that need to be done to avoid this:

1. Set the connection string option `useServerPrepStmts` to `true`.
2. Set `SQL_MODE` to `NO_BACKSLASH_ESCAPES`.



---

## 21.3.16. Connector/J Support

### 21.3.16.1. Connector/J Community Support

Oracle provides assistance to the user community by means of its mailing lists. For Connector/J related issues, you can get help from experienced users by using the MySQL and Java mailing list. Archives and subscription information is available online at <http://lists.mysql.com/java>.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See [Section 1.6.1, "MySQL Mailing Lists"](#).

Community support from experienced users is also available through the [JDBC Forum](#). You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See [Section 1.6.2, "MySQL Community Support at the MySQL Forums"](#).

### 21.3.16.2. How to Report Connector/J Bugs or Problems

The normal place to report bugs is <http://bugs.mysql.com/>, which is the address for our bugs database. This database is public, and can be browsed and searched by anyone. If you log in to the system, you will also be able to enter new reports.

If you find a sensitive security bug in MySQL Server, please let us know immediately by sending an email message to [<secalert\\_us@oracle.com>](mailto:secalert_us@oracle.com). Exception: Support customers should report all problems, including security bugs, to Oracle Support at <http://support.oracle.com/>.

Writing a good bug report takes patience, but doing it right the first time saves time both for us and for yourself. A good bug report, containing a full test case for the bug, makes it very likely that we will fix the bug in the next release.

This section will help you write your report correctly so that you do not waste your time doing things that may not help us much or at all.

If you have a repeatable bug report, please report it to the bugs database at <http://bugs.mysql.com/>. Any bug that we are able to repeat has a high chance of being fixed in the next MySQL release.

To report other problems, you can use one of the MySQL mailing lists.

Remember that it is possible for us to respond to a message containing too much information, but not to one containing too little. People often omit facts because they think they know the cause of a problem and assume that some details do not matter.

A good principle is this: If you are in doubt about stating something, state it. It is faster and less troublesome to write a couple more lines in your report than to wait longer for the answer if we must ask you to provide information that was missing from the initial report.

The most common errors made in bug reports are (a) not including the version number of Connector/J or MySQL used, and (b) not fully describing the platform on which Connector/J is installed (including the JVM version, and the platform type and version number that MySQL itself is installed on).

This is highly relevant information, and in 99 cases out of 100, the bug report is useless without it. Very often we get questions like, "Why doesn't this work for me?" Then we find that the feature requested wasn't implemented in that MySQL version, or that a bug described in a report has already been fixed in newer MySQL versions.

Sometimes the error is platform-dependent; in such cases, it is next to impossible for us to fix anything without knowing the operating system and the version number of the platform.

If at all possible, create a repeatable, standalone testcase that doesn't involve any third-party classes.

To streamline this process, we ship a base class for testcases with Connector/J, named `'com.mysql.jdbc.util.BaseBugReport'`. To create a testcase for Connector/J using this class,



create your own class that inherits from `com.mysql.jdbc.util.BaseBugReport` and override the methods `setUp()`, `tearDown()` and `runTest()`.

In the `setUp()` method, create code that creates your tables, and populates them with any data needed to demonstrate the bug.

In the `runTest()` method, create code that demonstrates the bug using the tables and data you created in the `setUp` method.

In the `tearDown()` method, drop any tables you created in the `setUp()` method.

In any of the above three methods, use one of the variants of the `getConnection()` method to create a JDBC connection to MySQL:

- `getConnection()` - Provides a connection to the JDBC URL specified in `getUrl()`. If a connection already exists, that connection is returned, otherwise a new connection is created.
- `getNewConnection()` - Use this if you need to get a new connection for your bug report (that is, there is more than one connection involved).
- `getConnection(String url)` - Returns a connection using the given URL.
- `getConnection(String url, Properties props)` - Returns a connection using the given URL and properties.

If you need to use a JDBC URL that is different from 'jdbc:mysql:///test', override the method `getUrl()` as well.

Use the `assertTrue(boolean expression)` and `assertTrue(String failureMessage, boolean expression)` methods to create conditions that must be met in your testcase demonstrating the behavior you are expecting (vs. the behavior you are observing, which is why you are most likely filing a bug report).

Finally, create a `main()` method that creates a new instance of your testcase, and calls the `run` method:

```
public static void main(String[] args) throws Exception {
    new MyBugReport().run();
}
```

Once you have finished your testcase, and have verified that it demonstrates the bug you are reporting, upload it with your bug report to <http://bugs.mysql.com/>.

## 21.4. MySQL Connector/C++

MySQL Connector/C++ is a MySQL database connector for C++. It lets you develop C++ applications that connect to the MySQL Server.

For release notes detailing the changes in each release of MySQL Connector/C++, see [MySQL Connector/C++ Release Notes](#).

### MySQL Connector/C++ Benefits

MySQL Connector/C++ offers the following benefits for C++ users compared to the MySQL C API (MySQL client library):

- Convenience of pure C++; no C function calls required
- Supports JDBC 4.0, an industry standard API
- Supports the object-oriented programming paradigm
- Reduces development time

- Licensed under the GPL with the FLOSS License Exception
- Available under a commercial license upon request

## JDBC Compatibility

MySQL Connector/C++ is compatible with the JDBC 4.0 API. MySQL Connector/C++ does not implement the entire JDBC 4.0 API, but does feature the following classes:

- `Connection`
- `DatabaseMetaData`
- `Driver`
- `PreparedStatement`
- `ResultSet`
- `ResultSetMetaData`
- `Savepoint`
- `Statement`

The JDBC 4.0 API defines approximately 450 methods for the classes just mentioned. MySQL Connector/C++ implements approximately 80% of these.

## Platform Support and Prerequisites

MySQL Connector/C++ requires:

- MySQL 5.1 or later
- Microsoft Visual Studio 2003 or later on Windows

The release has been successfully compiled and tested on the following platforms:

### AIX

- 5.2 (PPC32, PPC64)
- 5.3 (PPC32, PPC64)

### FreeBSD

- 6.0 (x86, x86\_64)

### HPUX

- 11.11 (PA-RISC 32bit, PA-RISC 64bit)

### Linux

- Debian 3.1 (PPC32, x86)
- FC4 (x86)
- RHEL 3 (ia64, x86, x86\_64)
- RHEL 4 (ia64, x86, x86\_64)
- RHEL 5 (ia64, x86, x86\_64)

- SLES 9 (ia64, x86, x86\_64)
- SLES 10 (ia64, x86\_64)
- SuSE 10.3, (x86\_64)
- Ubuntu 8.04 (x86)
- Ubuntu 8.10 (x86\_64)

#### **Mac OS X**

- Mac OS X 10.3 (PPC32, PPC64)
- Mac OS X 10.4 (PPC32, PPC64, x86)
- Mac OS X 10.5 (PPC32, PPC64, x86, x86\_64)

#### **SunOS**

- Solaris 8 (SPARC32, SPARC64, x86)
- Solaris 9 (SPARC32, SPARC64, x86)
- Solaris 10 (SPARC32, SPARC64, x86, x86\_64)

#### **Windows**

- XP Professional (32bit)
- 2003 (64bit)

## **Release Status**

MySQL Connector/C++ is available as a Generally Available (GA) version. We ask users and developers to try it out and provide feedback.

Note that MySQL Workbench is successfully using MySQL Connector/C++.

If you have any queries, please [contact us](#).

## **Licensing**

MySQL Connector/C++ is licensed under the terms of the GPL, like most MySQL Connectors. There are special exceptions to the terms and conditions of the GPL as applied to this software; see FLOSS License Exception. If you need a non-GPL license for commercial distribution, please contact us.

### **21.4.1. How to Get MySQL Connector/C++**

Binary and source packages can be obtained from [MySQL Connector/C++ downloads](#).

#### **MySQL Connector/C++ Binary Distributions**

Binary distributions are available for these platforms:

Microsoft Windows:

- MSI installer package
- Without installer (a Zip file)

Other platforms:

- Compressed GNU TAR archive ([tar.gz](#))

## MySQL Connector/C++ Source Distributions

Source packages use compressed GNU TAR file (`tar.gz`) format and can be used on any supported platform.

## MySQL Connector/C++ Source Repository

The latest development sources are available through [Launchpad](#).

The MySQL Connector/C++ code repository uses Bazaar. To check out the latest source code, use the `bzr` command-line tool:

```
shell> bzr branch lp:~mysql/mysql-connector-cpp/trunk .
```

### 21.4.2. Installing MySQL Connector/C++ from a Binary Distribution



#### Caution

One problem that can occur is when the tools you use to build your application are not compatible with the tools used to build the binary versions of MySQL Connector/C++. Ideally, build your application with the same tools that were used to build the MySQL Connector/C++ binaries. To help with this, the following resources are provided.

All distributions include a [README](#) file that contains platform-specific notes. At the end of the [README](#) file contained in the binary distribution, you will find the settings used to build the binaries. If you experience build-related issues on a platform, it may help to check the settings used on the platform to build the binary.

Developers using Microsoft Windows must satisfy the following requirements:

1. Use a supported version of Visual Studio, either Visual Studio 2005 or Visual Studio 2008.
2. Ensure that your application uses the same runtime library as that used to build MySQL Connector/C++. Visual Studio 2005 builds use `Microsoft.VC80.CRT (8.0.50727.762)`, and Visual Studio 2008 builds use `Microsoft.VC90.CRT (9.0.21022.8)`.
3. Your application should use the same linker configuration as MySQL Connector/C++. For example, use one of `/MD`, `/MDd`, `/MT`, or `/MTd`.

To use a variation of the requirements previously listed, such as a different compiler version, release configuration, or runtime library, compile MySQL Connector/C++ from source using your desired settings and ensure that your application is built using these same settings. To avoid issues, ensure that the three variables of compiler version, runtime library, and runtime linker configuration settings are the same for both application and MySQL Connector/C++ itself.

A better solution that ensures compatibility is to build your MySQL Connector/C++ libraries from the source code using the same tools that you use to build your application.

#### Downloading MySQL Connector/C++

Binary packages can be obtained from [MySQL Connector/C++ downloads](#).

#### Archive Package

Unpack the distribution archive into an appropriate directory. If you plan to use a dynamically linked version of MySQL Connector/C++, make sure that your system can reference the MySQL client library (MySQL Connector/C++ is linked against and thus requires the MySQL client library). Consult your operating system documentation on how to modify and expand the search path for libraries. If you cannot modify the library search path, it may help to copy your application, the MySQL Connector/C++ library and the MySQL client library into the same directory. Most systems search for libraries in the current directory.

Windows users can choose between two binary packaging formats:

- Windows MSI Installer (.msi file): To use the MSI Installer, launch it and follow the prompts in the screens it presents to install MySQL Connector/C in the location of your choosing.
- Zip archive without installer (.zip file): To use a Zip archive, unpack it in the directory where you intend to install it using WinZip or another tool that can read .zip files.

### Windows MSI Installer

Using the MSI Installer may be the easiest solution. The MSI Installer does not require any administrative permissions as it simply copies files.

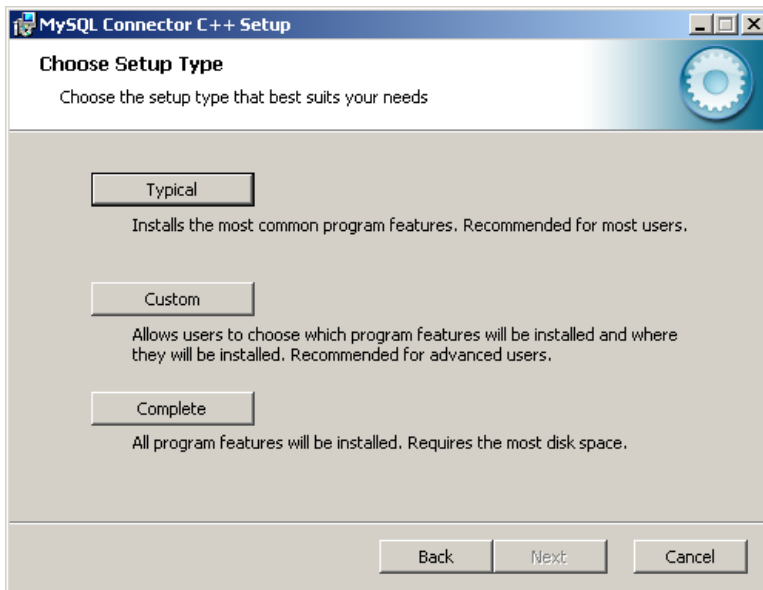
The MSI Installer begins by presenting a welcome screen.

**Figure 21.70. MSI Installer Welcome Screen**



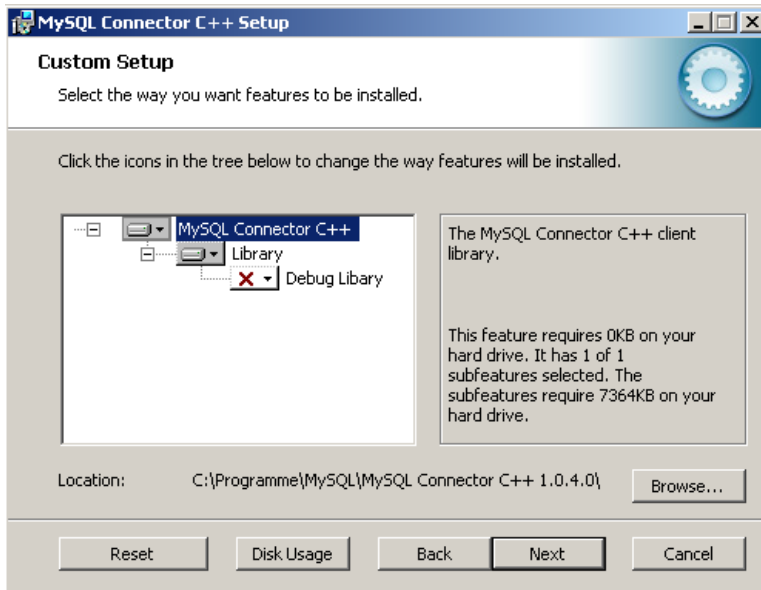
The MSI Installer overview screen enables you to select the type of installation you want to perform. The "Typical" installation consists of all required header files and the Release libraries. The "Custom" installation enables you to install additional Debug versions of the connector libraries.

**Figure 21.71. MSI Installer Overview Screen**



If you select a “Custom” installation, the MSI Installer presents a Custom Setup screen that enables you to select which features to install and where to install them.

**Figure 21.72. MSI Installer Custom Setup Screen**



### 21.4.3. Installing MySQL Connector/C++ from Source

MySQL Connector/C++ is based on the MySQL client library (MySQL C API) and is linked against it. Thus, to compile MySQL Connector/C++, the MySQL Client Library must be installed.

You also need the cross-platform build tool [CMake](#) 2.4, or newer, and [GLib](#) 2.2.3 or newer. Check the [README](#) file included with the distribution for platform-specific notes on building for Windows and SunOS.

Typically, the MySQL client library is installed when the MySQL Server is installed. However, check your operating system documentation for other installation options.

As of MySQL Connector/C++ 1.1.0, the Boost C++ libraries 1.34.0 or newer must be installed. Boost is required only to build the connector, required to use the connector. You can obtain Boost and its installation instructions from [the official site](#). Once Boost is installed, tell the build system where the Boost files are by defining the `BOOST_ROOT:STRING` option. This can be done when you invoke [CMake](#). For example:

```
shell> cmake . -DBOOST_ROOT:STRING=/usr/local/boost_1_40_0
```

Change `/usr/local/boost_1_40_0/` as necessary to match your installation. For further details, see [Section 21.4.3.1, “Building MySQL Connector/C++ from Source on Unix, Solaris, and Mac OS X”](#) and [Section 21.4.3.2, “Building MySQL Connector/C++ from Source on Windows”](#).

#### 21.4.3.1. Building MySQL Connector/C++ from Source on Unix, Solaris, and Mac OS X

1. Change location to the top-level directory of the source distribution:

```
shell> cd /path/to/mysql-connector-cpp
```

2. Run [CMake](#) to build a [Makefile](#):

```
shell> cmake .
-- Check for working C compiler: /usr/local/bin/gcc
-- Check for working C compiler: /usr/local/bin/gcc -- works
[...]
-- Generating done
-- Build files have been written to: /path/to/mysql-connector-cpp/
```

On non-Windows systems, `CMake` first checks to see if the `CMake` variable `MYSQL_CONFIG_EXECUTABLE` is set. If it is not found, `CMake` tries to locate `mysql_config` in the default locations.

If you have any problems with the configuration process, check the troubleshooting instructions given later.

- Use `make` to build the libraries. First make sure you have a clean build:

```
shell> make clean
```

Then build the connector:

```
shell> make
[ 1%] Building CXX object »
driver/CMakeFiles/mysqlcppconn.dir/mysql_connection.o
[ 3%] Building CXX object »
driver/CMakeFiles/mysqlcppconn.dir/mysql_constructed_resultset.o
[...]
[100%] Building CXX object examples/CMakeFiles/statement.dir/statement.o
Linking CXX executable statement
```

If all goes well, you will find the MySQL Connector/C++ library in the `driver` directory.

- Install the header and library files:

```
shell> make install
```

Unless you have changed the location in the configuration step, `make install` copies the header files to the directory `/usr/local/include`. The header files copied are `mysql_connection.h` and `mysql_driver.h`.

Again, unless you have specified otherwise, `make install` copies the library files to `/usr/local/lib`. The files copied are the dynamic library `libmysqlcppconn.so`, and the static library `libmysqlcppconn-static.a`. The extension of the dynamic library might be different on your system (for example, `.dylib` on Mac OS X).

If you encounter any errors, please first carry out these checks:

- `CMake` options: MySQL installation path, debug version and more

In case of configuration or compilation problems, check the list of `CMake` options:

```
shell> cmake -L
[...]
CMAKE_BACKWARDS_COMPATIBILITY:STRING=2.4
CMAKE_BUILD_TYPE:STRING=
CMAKE_INSTALL_PREFIX:PATH=/usr/local
EXECUTABLE_OUTPUT_PATH:PATH=
LIBRARY_OUTPUT_PATH:PATH=
MYSQLCPPCONN_GCOV_ENABLE:BOOL=0
MYSQLCPPCONN_TRACE_ENABLE:BOOL=0
MYSQL_CONFIG_EXECUTABLE:FILEPATH=/usr/bin/mysql_config
```

For example, if your MySQL Server installation path is not `/usr/local/mysql` and you want to build a debug version of the MySQL Connector/C++, use this command:

```
shell> cmake \
  -D CMAKE_BUILD_TYPE:STRING=Debug \
  -D MYSQL_CONFIG_EXECUTABLE=/path/to/my/mysql/server/bin/mysql_config .
```

- Verify your settings with `cmake -L`:

```
shell> cmake -L
[...]
CMAKE_BACKWARDS_COMPATIBILITY:STRING=2.4
```

```

CMAKE_BUILD_TYPE:STRING=
CMAKE_INSTALL_PREFIX:PATH=/usr/local
EXECUTABLE_OUTPUT_PATH:PATH=
LIBRARY_OUTPUT_PATH:PATH=
MYSQLCPPCONN_GCOV_ENABLE:BOOL=0
MYSQLCPPCONN_TRACE_ENABLE:BOOL=0
MYSQL_CONFIG_EXECUTABLE=/path/to/my/mysql/server/bin/mysql_config

```

Proceed by executing a `make clean` command followed by a `make` command, as described previously.

Once you have installed MySQL Connector/C++, you can carry out a quick test to check the installation. To do this, compile and run one of the example programs, such as `examples/standalone_example.cpp`. This example is discussed in more detail later, but for now, you can use it to test whether the connector has been correctly installed. This procedure assumes that you have a working MySQL Server to which you can connect. It also assumes header and library locations of `/usr/local/include` and `/usr/local/lib`, respectively; adjust these as necessary for your system.

1. Compile the example program. To do this, change location to the `examples` directory and enter this command:

```

shell> g++ -o test_install \
-I/usr/local/include -I/usr/local/include/cppconn \
-Wl,-Bdynamic -lmysqlcppconn standalone_example.cpp

```

2. Make sure the dynamic library which is used in this case can be found at runtime:

```

shell> export LD_LIBRARY_PATH=/usr/local/lib

```

3. Now run the program to test your installation, substituting the appropriate host, user, password, and database names for your system:

```

shell> ./test_install localhost root password database

```

You should see output similar to the following:

```

Connector/C++ standalone program example...

... running 'SELECT 'Welcome to Connector/C++' AS _message'
... MySQL replies: Welcome to Connector/C++
... say it again, MySQL
...MySQL replies: Welcome to Connector/C++

... find more at http://www.mysql.com

```

If you see any errors, take note of them and go through the troubleshooting procedures discussed earlier.

### 21.4.3.2. Building MySQL Connector/C++ from Source on Windows



#### Note

The only compiler formally supported for Windows is Microsoft Visual Studio 2003 and above.

The basic steps for building the connector on Windows are the same as for Unix. It is important to use `CMake` 2.6.2 or newer to generate build files for your compiler and to invoke the compiler.



#### Note

On Windows, `mysql_config` is not present, so `CMake` attempts to retrieve the location of MySQL from the environment variable `$ENV{MYSQL_DIR}`. If `MYSQL_DIR` is not set, `CMake` then proceeds to check for MySQL in the following locations: `$ENV{ProgramFiles}/MySQL/*/include`, and `$ENV{SystemDrive}/MySQL/*/include`.



**CMake** makes it easy for you to try other compilers. However, you may experience compile warnings, compile errors or linking issues not detected by Visual Studio. Patches are gratefully accepted to fix issues with other compilers.

Consult the **CMake** manual or check `cmake --help` to find out which build systems are supported by your **CMake** version:

```
C:\>cmake --help
cmake version 2.6-patch 2
Usage
[...]
Generators

The following generators are available on this platform:
  Borland Makefiles           = Generates Borland makefiles.
  MSYS Makefiles              = Generates MSYS makefiles.
  MinGW Makefiles             = Generates a make file for use with
                               mingw32-make.
  NMake Makefiles             = Generates NMake makefiles.
  Unix Makefiles              = Generates standard UNIX makefiles.
  Visual Studio 6              = Generates Visual Studio 6 project files.
  Visual Studio 7              = Generates Visual Studio .NET 2002 project
                               files.
  Visual Studio 7 .NET 2003    = Generates Visual Studio .NET 2003 project
                               files.
  Visual Studio 8 2005         = Generates Visual Studio .NET 2005 project
                               files.
  Visual Studio 8 2005 Win64   = Generates Visual Studio .NET 2005 Win64
                               project files.
  Visual Studio 9 2008         = Generates Visual Studio 9 2008 project fil
  Visual Studio 9 2008 Win64   = Generates Visual Studio 9 2008 Win64 proje
                               files.
[...]
```

It is likely that your **CMake** binary supports more compilers, known by **CMake** as *generators*, than can actually be used to build MySQL Connector/C++. We have built the connector using the following generators:

- Microsoft Visual Studio 8 (Visual Studio 2005)
- Microsoft Visual Studio 9 (Visual Studio 2008, Visual Studio 2008 Express)
- NMake

Please see the building instructions for Unix, Solaris and Mac OS X for troubleshooting and configuration hints.

Use these steps to build the connector:

1. Change location to the top-level directory of the source distribution:

```
shell> cd C:\path_to_mysql_cpp
```

2. Run **CMake** to generate build files for your generator:

### Visual Studio

```
C:\>cmake -G "Visual Studio 9 2008"
-- Check for working C compiler: cl
-- Check for working C compiler: cl -- works
-- Detecting C compiler ABI info
-- Detecting C compiler ABI info - done
-- Check for working CXX compiler: cl
-- Check for working CXX compiler: cl -- works
-- Detecting CXX compiler ABI info
-- Detecting CXX compiler ABI info - done
-- ENV{MYSQL_DIR} =
-- MySQL Include dir: C:/Programme/MySQL/MySQL Server 5.5/include
```

```

-- MySQL Library      : C:/Programs/MySQL/MySQL Server 5.5/lib/opt/mysqlclient.lib
-- MySQL Library dir: C:/Programs/MySQL/MySQL Server 5.5/lib/opt
-- MySQL CFLAGS:
-- MySQL Link flags:
-- MySQL Include dir: C:/Programs/MySQL/MySQL Server 5.5/include
-- MySQL Library dir: C:/Programs/MySQL/MySQL Server 5.5/lib/opt
-- MySQL CFLAGS:
-- MySQL Link flags:
-- Configuring cppconn
-- Configuring test cases
-- Looking for isinf
-- Looking for isinf - not found
-- Looking for isinf
-- Looking for isinf - not found.
-- Looking for finite
-- Looking for finite - not found.
-- Configuring C/J junit tests port
-- Configuring examples
-- Configuring done
-- Generating done
-- Build files have been written to: C:\path_to_mysql_cpp
C:\>dir *.sln *.vcproj
[... ]
19.11.2008  12:16                23.332  MYSQLCPPCONN.sln
[... ]
19.11.2008  12:16                27.564  ALL_BUILD.vcproj
19.11.2008  12:16                27.869  INSTALL.vcproj
19.11.2008  12:16                28.073  PACKAGE.vcproj
19.11.2008  12:16                27.495  ZERO_CHECK.vcproj

```

### NMake

```

C:\>cmake -G "NMake Makefiles"
-- The C compiler identification is MSVC
-- The CXX compiler identification is MSVC
[... ]
-- Build files have been written to: C:\path_to_mysql_cpp

```

3. Use your compiler to build MySQL Connector/C++.

### Visual Studio - GUI

Open the newly generated project files in the Visual Studio GUI or use a Visual Studio command line to build the driver. The project files contain a variety of different configurations, debug and nondebug versions among them.

### Visual Studio - NMake

```

C:\>nmake

Microsoft (R) Program Maintenance Utility Version 9.00.30729.01
Copyright (C) Microsoft Corporation. All rights reserved.

Scanning dependencies of target mysqlcppconn
[ 2%] Building CXX object driver/CMakeFiles/mysqlcppconn.dir/mysql_connection.obj
mysql_connection.cpp
[... ]
Linking CXX executable statement.exe
[100%] Built target statement

```

## 21.4.3.3. Dynamically Linking MySQL Connector/C++ Against the MySQL Client Library



### Note

This section refers to dynamic linking of MySQL Connector/C++ with the client library, not dynamic linking of the application to MySQL Connector/C++.

Precompiled binaries of MySQL Connector/C++ use static binding with the client library by default.

An application that uses MySQL Connector/C++ can be either statically or dynamically linked to the MySQL Connector/C++ libraries. MySQL Connector/C++ is usually statically linked to the underlying MySQL client library (or Connector/C). Note that, unless otherwise stated, reference to the MySQL client library is also taken to include Connector/C, which is a separately packaged, standalone version of the MySQL client library. From MySQL Connector/C++ 1.1.0 on, it is possible to also dynamically link to the underlying MySQL client library. The ability of MySQL Connector/C++ to dynamically link to MySQL client library is not enabled by default. Enabling this feature is done through a compile-time option when compiling the MySQL Connector/C++ source code.

To use the ability to dynamically link the client library to MySQL Connector/C++, define the `MYSQLCLIENT_STATIC_BINDING:BOOL` when building the MySQL Connector/C++ source code:

```
shell> rm CMakeCache.txt
shell> cmake -DMYSQLCLIENT_STATIC_BINDING:BOOL=1 .
shell> make clean
shell> make
shell> make install
```

Now, in your application, when creating a connection, MySQL Connector/C++ will select and load a client library at runtime. It will choose the client library by searching defined locations and environment variables depending on the host operating system. It is also possible when creating a connection in an application to define an absolute path to the client library to be loaded at runtime. This can be convenient if you have defined a standard location from which you want the client library to be loaded. This is sometimes done to circumvent possible conflicts with other versions of the client library that may be located on the system.

## 21.4.4. Building MySQL Connector/C++ Windows Applications with Microsoft Visual Studio

MySQL Connector/C++ is available as a static or dynamic library to use with your application. This section describes how to link the library to your application.



### Note

To avoid potential crashes, the build configuration of MySQL Connector/C++ should match the build configuration of the application using it. For example, do not use the release build of MySQL Connector/C++ with a debug build of the client application.

### Static Library

The MySQL Connector/C++ static library file is `mysqlcppconn-static.lib`. You link this library statically with your application. Also link against the files `libmysql.dll` and `libmysql.lib`. At runtime, the application will require access to `libmysql.dll`.

### Dynamic Library

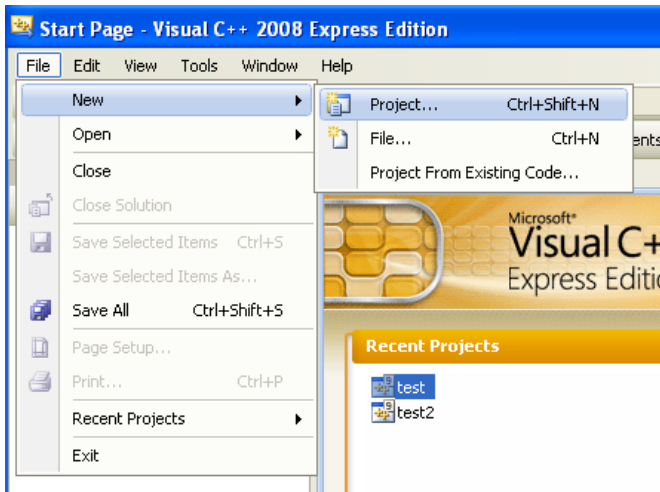
The MySQL Connector/C++ dynamic library file is `mysqlcppconn.dll`. To build your client application, link it with the file `mysqlcppconn.lib`. At runtime, the application will require access to the files `mysqlcppconn.dll` and `libmysql.dll`.

### Building a MySQL Connector/C++ Application with Microsoft Visual Studio

The initial steps for building an application to use either the static or dynamic library are the same. Some additional steps vary, depend on whether you are building your application to use the [static \[2729\]](#) or [dynamic \[2730\]](#) library.

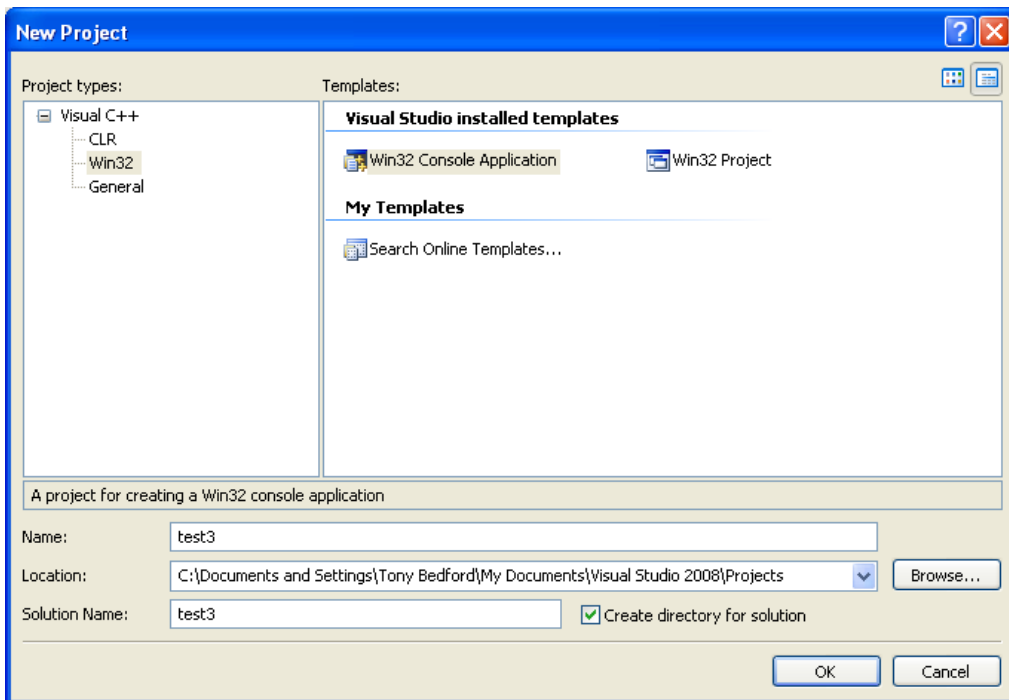
1. Select **File**, **New**, **Project** from the main menu.

**Figure 21.73. Creating a New Project**



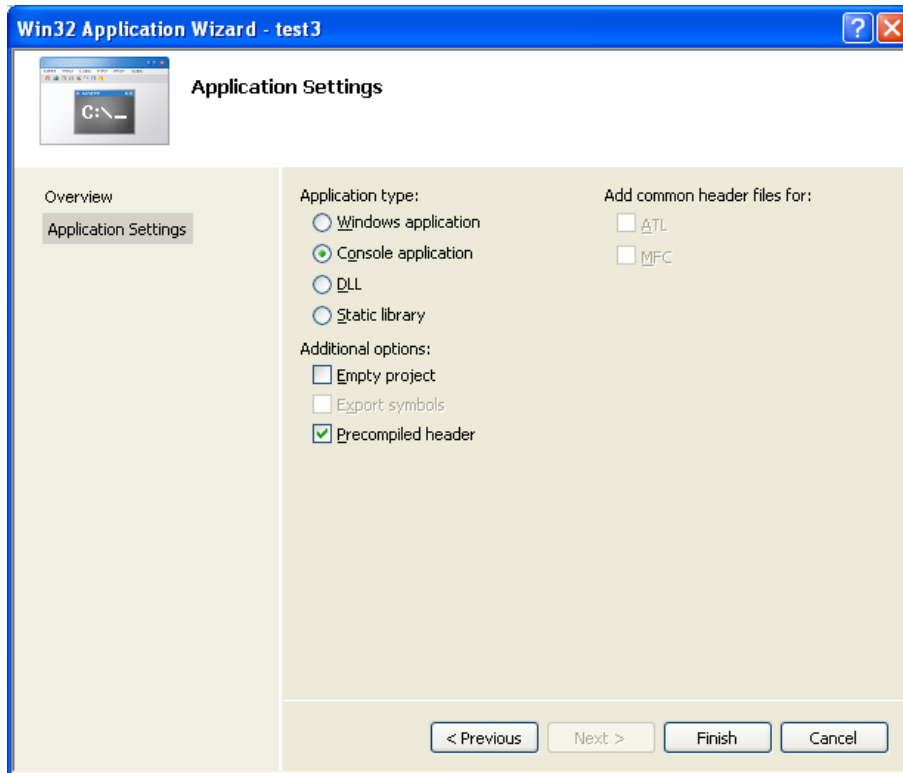
2. In the wizard, select **Visual C++**, **Win32**. From **Visual Studio Installed Templates**, select the application type **Win32 Console Application**. Enter a name for the application, then click **OK**, to move to the Win32 Application Wizard.

**Figure 21.74. The New Project Dialog Box**




3. In the Win32 Application Wizard, click **Application Settings** and ensure the defaults are selected. The radio button **Console application** and the check box **Precompiled headers** are selected. Click **Finish** to close the wizard.

**Figure 21.75. The Win32 Application Wizard**



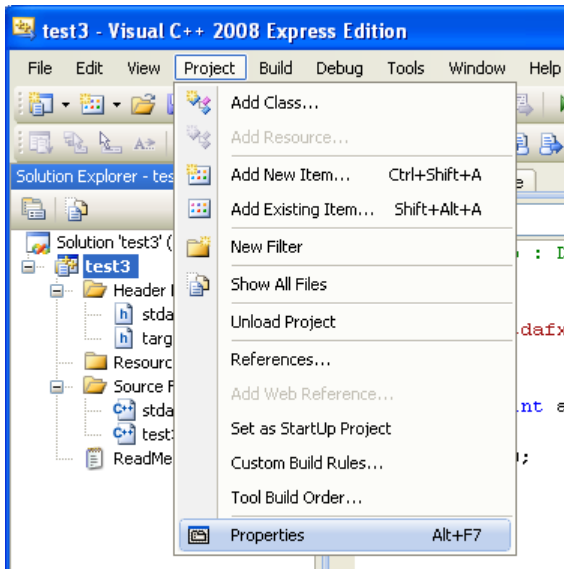
4. From the drop down list box on the toolbar, change from the default **Debug** build to the **Release** build.

**Figure 21.76. Selecting the Release Build**



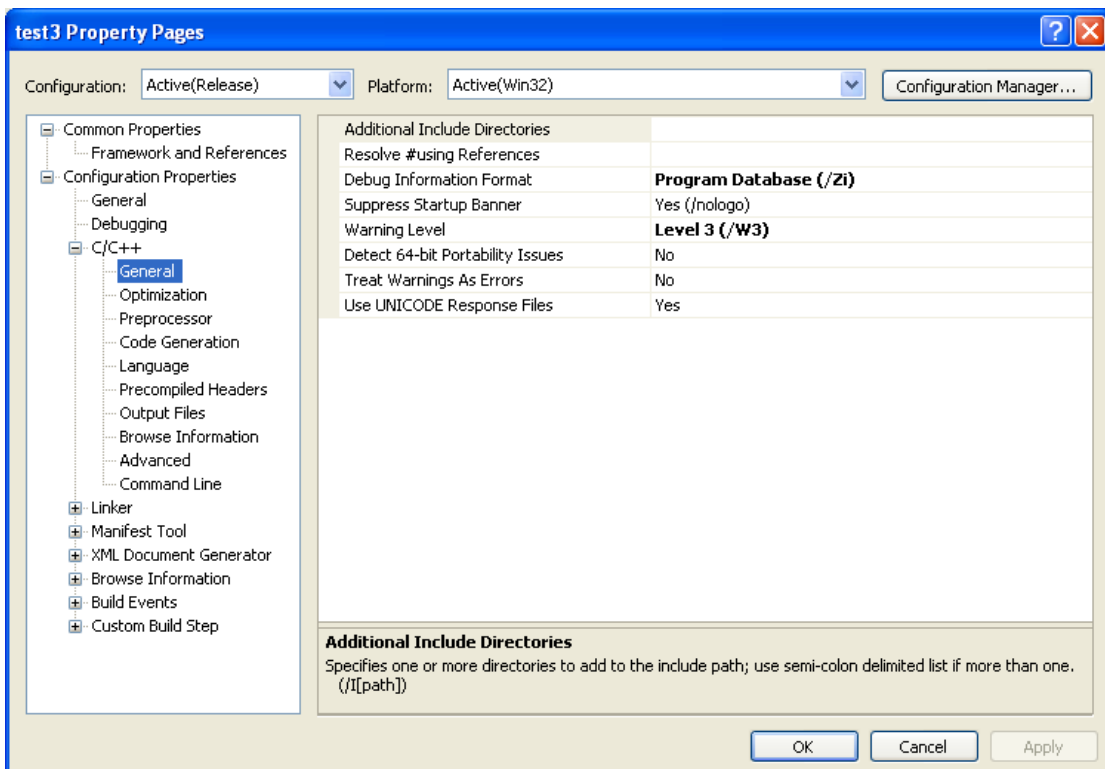
5. From the main menu select **Project**, **Properties**. This can also be accessed using the hot key **ALT + F7**.

**Figure 21.77. Selecting Project Properties from the Main Menu**



6. Under **Configuration Properties**, open the tree view.
7. Select **C++**, **General** in the tree view.

**Figure 21.78. Setting Properties**



8. Ensure that Visual Studio can find the MySQL include directory. This directory includes header files that can optionally be installed when installing MySQL Server.

Figure 21.79. MySQL Include Directory

```

C:\mysql\include>dir
Volume in drive C has no label.
Volume Serial Number is 9484-B7C3

Directory of C:\mysql\include

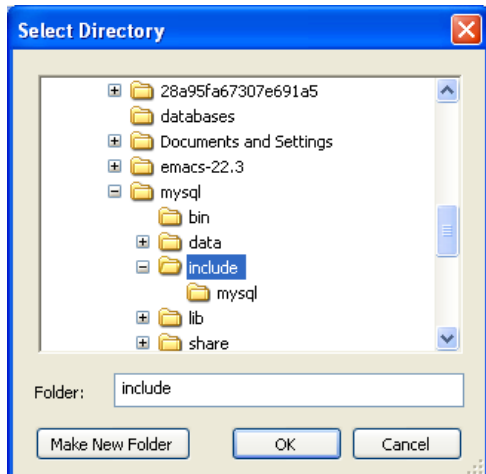
10/12/2008  13:56    <DIR>          .
10/12/2008  13:56    <DIR>          ..
15/11/2008  06:53                10,935  config-win.h
15/11/2008  06:53                4,058  decimal.h
15/11/2008  06:53                3,854  errmsg.h
15/11/2008  06:53                6,954  keycache.h
15/11/2008  06:53                2,819  libmysql.def
10/12/2008  13:56    <DIR>          mysql
15/11/2008  06:53            33,654  mysql.h
15/11/2008  06:53            23,155  mysqld_errname.h
15/11/2008  06:53            23,212  mysqld_error.h
15/11/2008  06:53            20,174  mysql_com.h
15/11/2008  06:53             1,130  mysql_embed.h
15/11/2008  06:53             2,097  mysql_time.h
15/11/2008  06:53             860  mysql_version.h
15/11/2008  06:53             1,804  my_alloc.h
15/11/2008  06:53             1,941  my_attribute.h
15/11/2008  06:53             4,576  my_debug.h
15/11/2008  06:53             3,460  my_dir.h
15/11/2008  06:53             3,278  my_getopt.h
15/11/2008  06:53            50,497  my_global.h
15/11/2008  06:53             1,488  my_list.h
15/11/2008  06:53             3,656  my_net.h
15/11/2008  06:53             1,542  my_no_pthread.h
15/11/2008  06:53            26,468  my_pthread.h
15/11/2008  06:53            40,097  my_sys.h
15/11/2008  06:53             2,677  my_xml.h
15/11/2008  06:53            22,584  mctype.h
15/11/2008  06:53             8,256  mstring.h
15/11/2008  06:53            1,957  sql_common.h
15/11/2008  06:53            12,484  sql_state.h
15/11/2008  06:53             1,006  sslopt-case.h
15/11/2008  06:53             2,205  sslopt-longopts.h
15/11/2008  06:53             1,107  sslopt-vars.h
15/11/2008  06:53             1,501  typelib.h
                32 File(s)          325,486 bytes
                 3 Dir(s)  4,091,248,640 bytes free

C:\mysql\include>

```


- In the **Additional Include Directories** text field, add the MySQL `include/` directory.

Figure 21.80. Select Directory Dialog



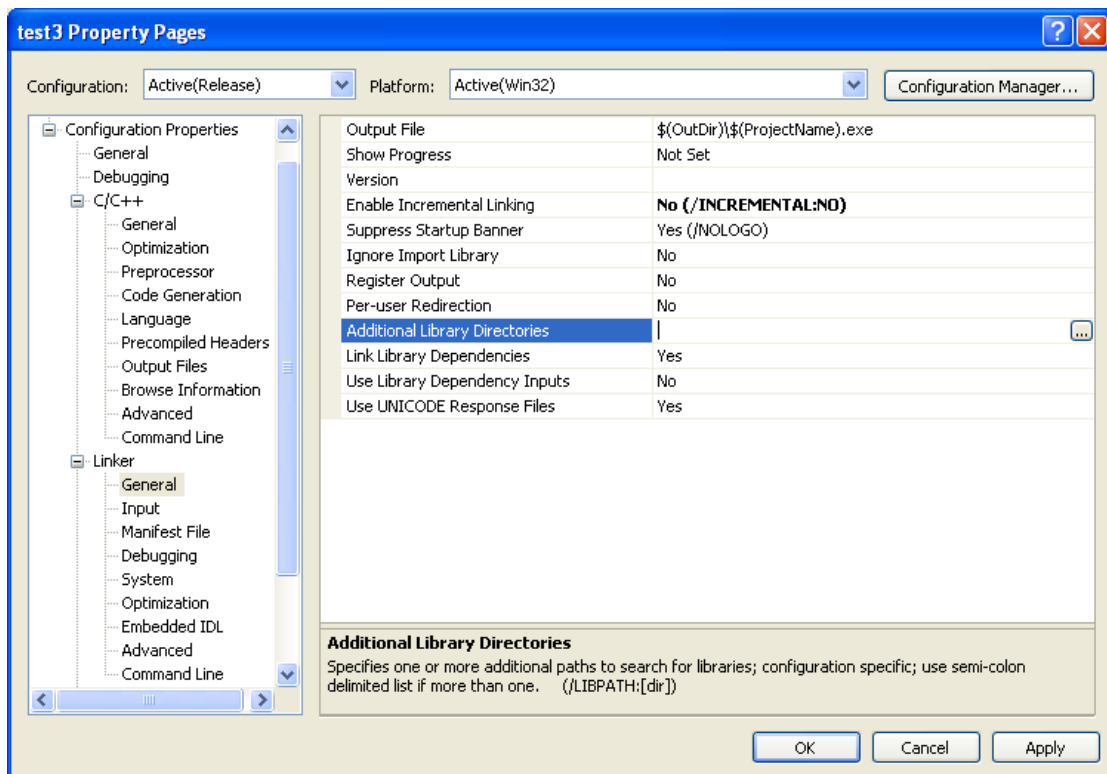
- Also set the location of additional libraries that Visual Studio needs to build the application. These are located in the MySQL `lib/opt` directory, a subdirectory of the MySQL Server installation directory.

Figure 21.81. Typical Contents of MySQL lib/opt Directory



11. In the tree view, open **Linker, General, Additional Library Directories**.

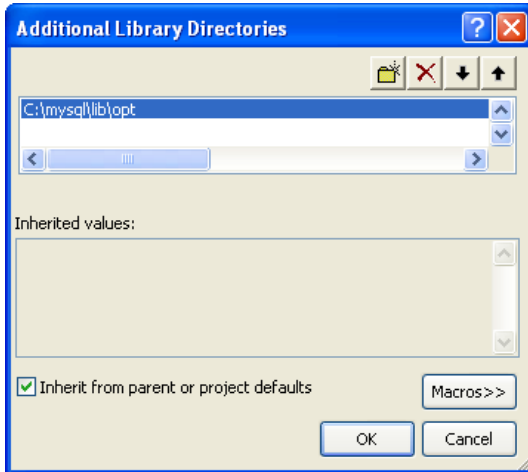
Figure 21.82. Additional Library Directories



12. Add the `lib/opt` directory into the **Additional Library Directories** text field. This enables the library file `libmysql.lib` to be found.



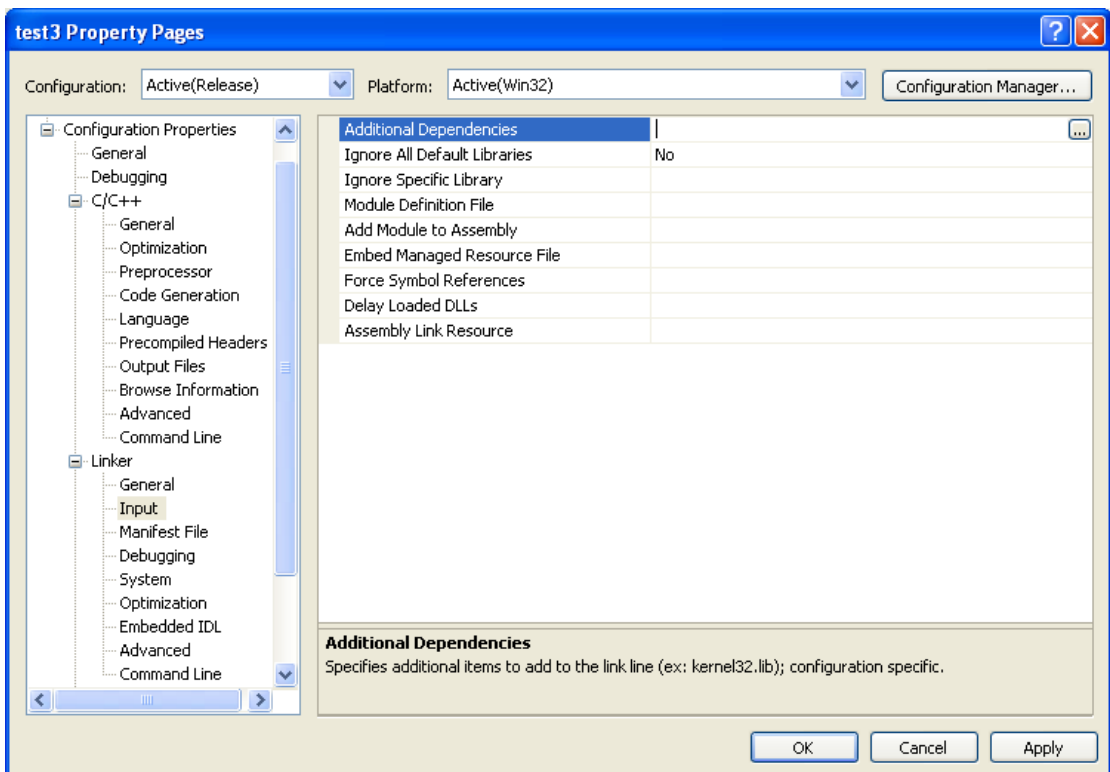
**Figure 21.83. Additional Library Directories Dialog**



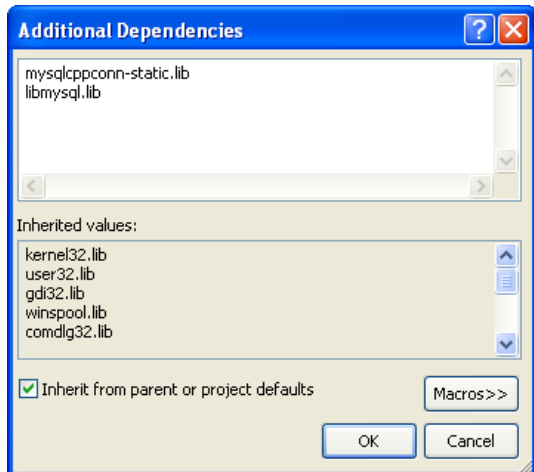
The remaining steps depend on whether you are building an application to use the MySQL Connector/C++ static or dynamic library. If you are building your application to use the dynamic library, see [Section 21.4.4, “Building MySQL Connector/C++ Windows Applications with Microsoft Visual Studio” \[2730\]](#). If you are building your application to use the static library, carry out the following steps:

1. Open **Linker, Input, Additional Dependencies**.

**Figure 21.84. Additional Dependencies**

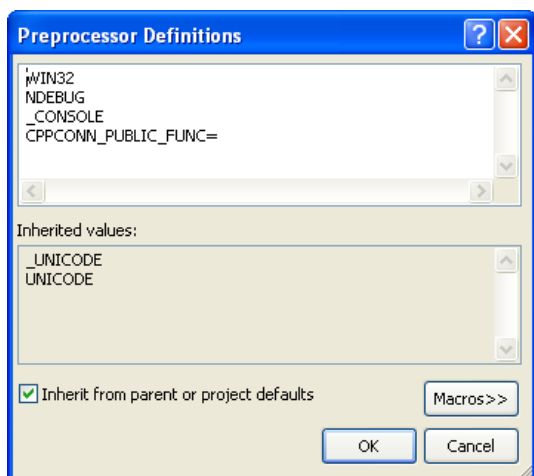


2. Enter `mysqlcppconn-static.lib` and `libmysql.lib`.

**Figure 21.85. Adding Additional Dependencies**

- By default `CPPCONN_PUBLIC_FUNC` is defined to declare functions to be compatible with an application that calls a DLL. If you are building an application to call the static library, ensure that function prototypes are compatible with this. In this case, define `CPPCONN_PUBLIC_FUNC` to be an empty string, so that functions are declared with the correct prototype.

In the **Project, Properties** tree view, under **C++**, **Preprocessor**, enter `CPPCONN_PUBLIC_FUNC=` into the **Preprocessor Definitions** text field.

**Figure 21.86. Setting the CPPCONN\_PUBLIC\_FUNC Define****Note**

Make sure you enter `CPPCONN_PUBLIC_FUNC=` and not `CPPCONN_PUBLIC_FUNC`, so that it is defined as an empty string.

If you are building an application to use the MySQL Connector/C++ dynamically linked library, carry out these steps:

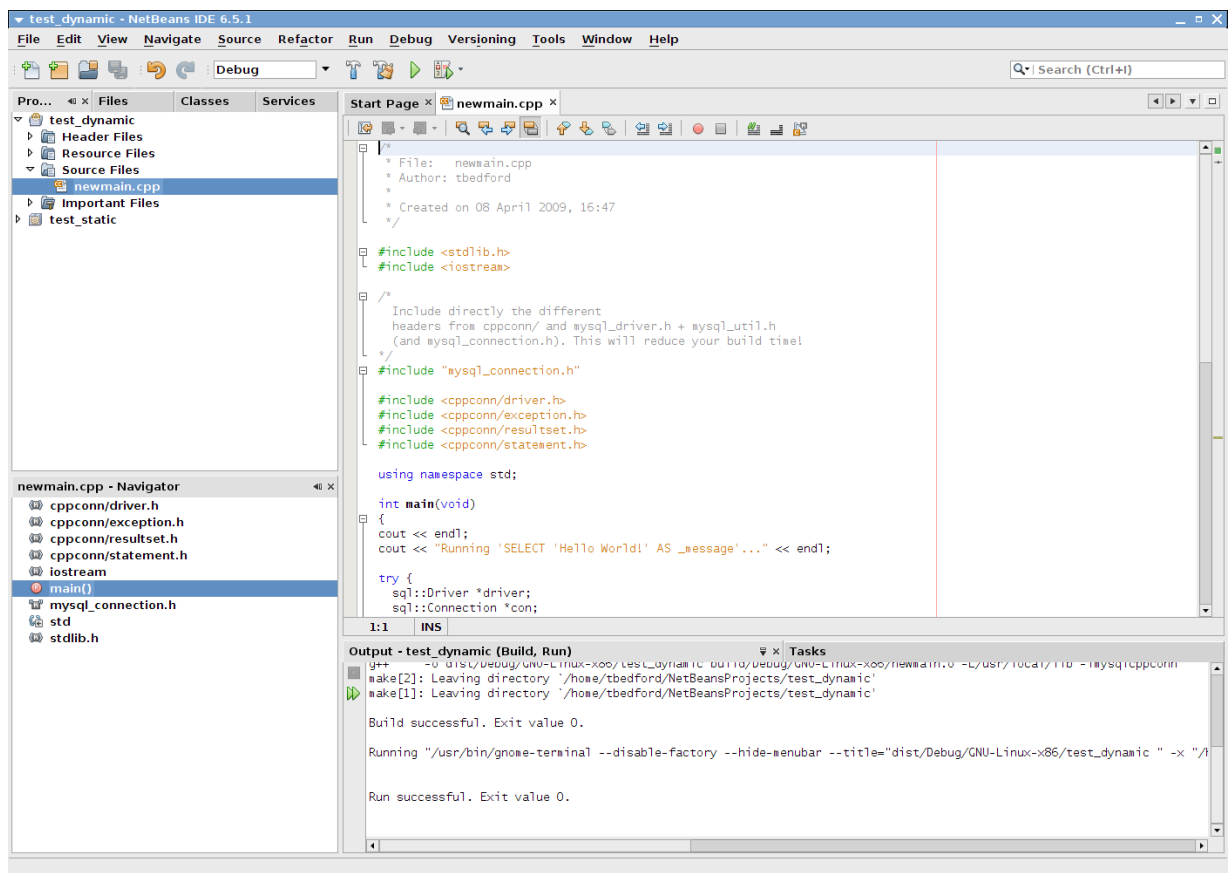
- Under **Linker, Input**, add `mysqlcppconn.lib` into the **Additional Dependencies** text field.
- `mysqlcppconn.dll` must be in the same directory as the application executable, or somewhere on the system's path, so that the application can access the MySQL Connector/C++ Dynamic Linked Library at runtime.

Copy `mysqlcppconn.dll` to the same directory as the application. Alternatively, extend the `PATH` environment variable using `SET PATH=%PATH%;C:\path\to\cpp`. Alternatively, you can copy `mysqlcppconn.dll` to the Windows installation Directory, typically `c:\windows`.

## 21.4.5. Building MySQL Connector/C++ Linux Applications with NetBeans

This section describes how to build MySQL Connector/C++ applications for Linux using the NetBeans IDE.

Figure 21.87. The NetBeans IDE



### Note

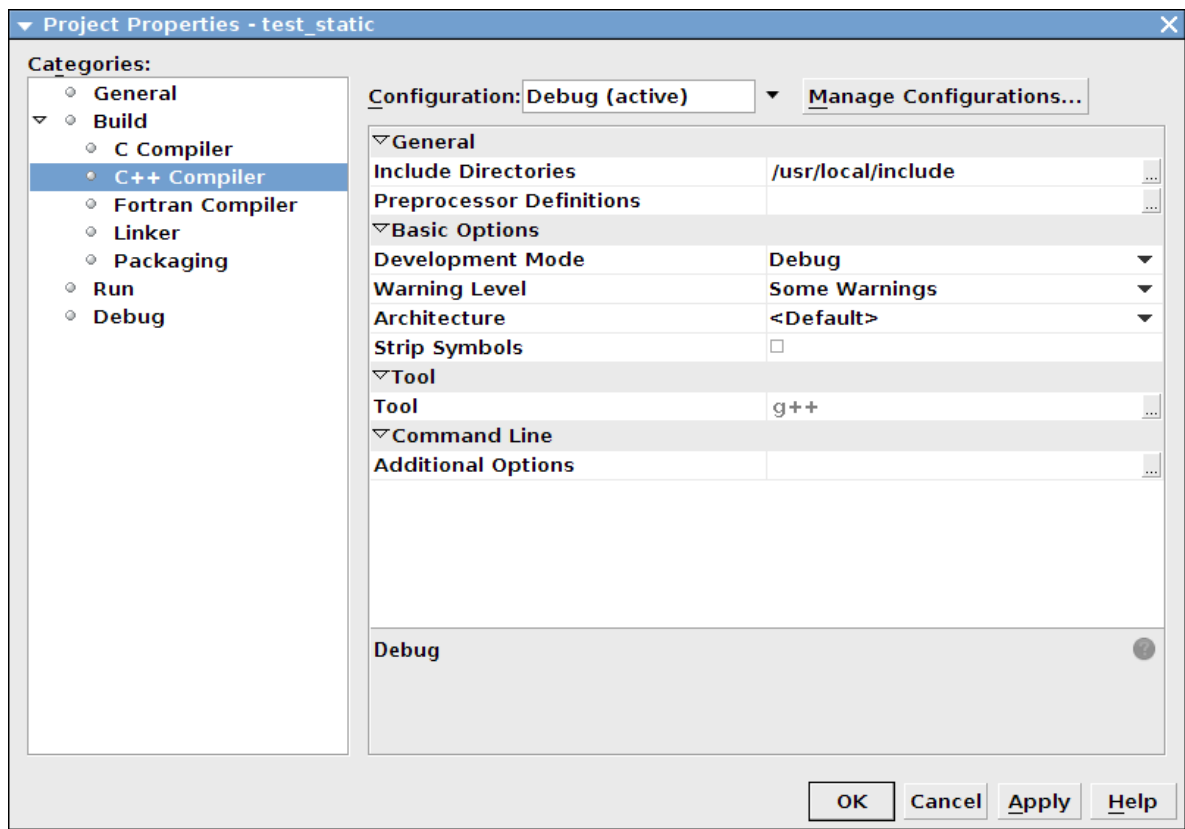
To avoid potential crashes, the build configuration of MySQL Connector/C++ should match the build configuration of the application using it. For example, do not use the release build of MySQL Connector/C++ with a debug build of the client application.

1. Create a new project. Select **File**, **New Project**. Choose a **C/C++ Application** and click **Next**.
2. Give the project a name and click **Finish**. A new project is created.
3. In the **Projects** tab, right-click **Source Files** and select **New**, then **Main C++ File...**
4. Change the filename, or simply select the defaults and click **Finish** to add the new file to the project.
5. Now add some working code to your main source file. Explore your MySQL Connector/C++ installation and navigate to the `examples` directory.
6. Select a suitable example, such as `standalone_example_docs1.cpp`. Copy all the code in this file, and use it to replace the code in your existing main source file. Amend the code to reflect the

connection properties required for your test database. You now have a working example that will access a MySQL database using MySQL Connector/C++.

7. At this point, NetBeans shows some errors in the source code. Direct NetBeans to the necessary header files to include. Select **File, Project Properties** from the main menu.
8. In the **Categories:** tree view panel, navigate to **Build, C++ Compiler**.
9. In the **General** panel, select **Include Directories**.
10. Click the **...** button.
11. Click **Add**, then navigate to the directory where the MySQL Connector/C++ header files are located. This is `/usr/local/include` unless you have installed the files to a different location. Click **Select**. Click **OK**.

**Figure 21.88. Setting the Header Include Directory**



12. Click **OK** again to close the **Project Properties** dialog.

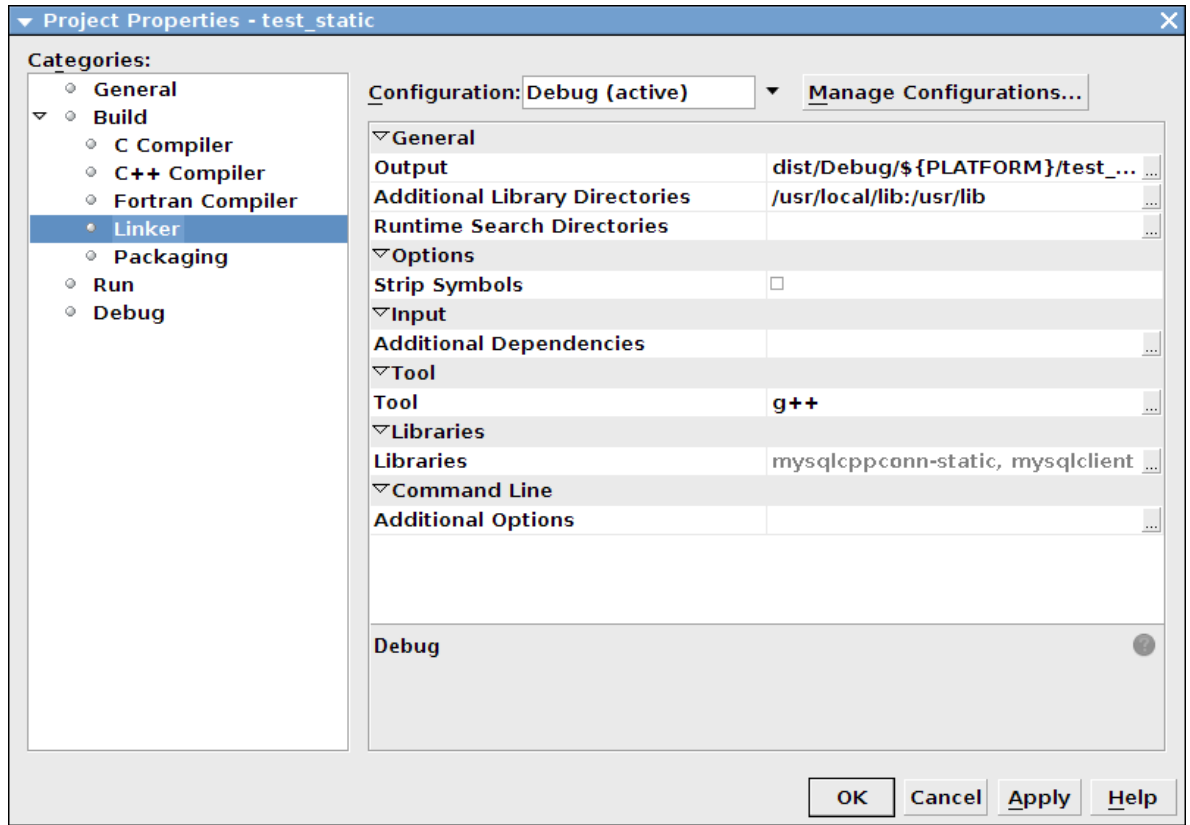
At this point, you have created a NetBeans project containing a single C++ source file. You have also ensured that the necessary include files are accessible. Before continuing, decide whether your project is to use the MySQL Connector/C++ static or dynamic library. The project settings are slightly different in each case, because you link against a different library.

### Using the Static Library

To use the static MySQL Connector/C++ library, link against two library files, `libmysqlcppconn-static.a` and `libmysqlclient.a`. The locations of the files depend on your setup, but typically the former are in `/usr/local/lib` and the latter in `/usr/lib`. The file `libmysqlclient.a` is not part of MySQL Connector/C++, but is the MySQL client library file distributed with MySQL Server. (Remember, the MySQL client library is an optional component as part of the MySQL Server installation process.) The MySQL Client Library is also available as part of the MySQL Connector/C distribution.

1. Set the project to link the necessary library files. Select **File**, **Project Properties** from the main menu.
2. In the **Categories:** tree view, navigate to **Linker**.
3. In the **General** panel, select **Additional Library Directories**. Click the **...** button.
4. Select and add the `/usr/lib` and `/usr/local/lib` directories.
5. In the same panel, add the two library files required for static linking as discussed earlier. The properties panel should then look similar to the following screenshot.

**Figure 21.89. Setting the Static Library Directories and File Names**

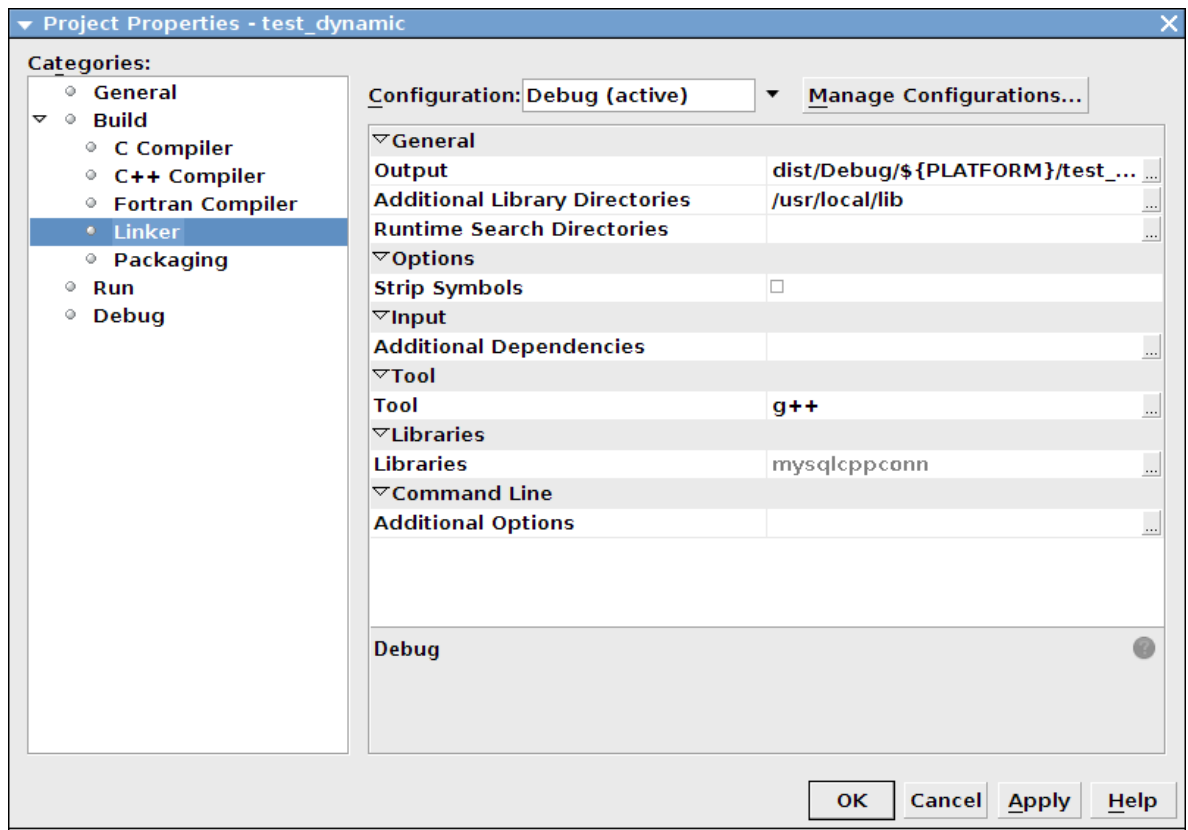


6. Click OK to close the **Project Properties** dialog.

### Using the Dynamic Library

To use the MySQL Connector/C++ dynamic library, link your project with a single library file, `libmysqlcppconn.so`. The location of this file depends on how you configured your installation of MySQL Connector/C++, but typically is `/usr/local/lib`.


1. Set the project to link the necessary library file. Select **File**, **Project Properties** from the main menu.
2. In the **Categories:** tree view, navigate to **Linker**.
3. In the **General** panel, select **Additional Library Directories**. Click the **...** button.
4. Select and add the `/usr/local/lib` directories.
5. In the same panel, add the library file required for static linking as discussed earlier. The properties panel should look similar to the following screenshot.

**Figure 21.90. Setting the Dynamic Library Directory and File Name**

6. Click OK to close the Project Properties dialog.

After configuring your project, build it by selecting **Run, Build Main Project** from the main menu. You then run the project using **Run, Run Main Project**.

On running the application, you should see a screen similar to the following (this is actually the static version of the application shown):

**Figure 21.91. The Example Application Running**A screenshot of a terminal window titled "dist/Debug/GNU-Linux-x86/test\_static". The window contains the following text:

```
Running 'SELECT 'Hello World!' AS _message'...  
... MySQL replies: Hello World!  
... MySQL says it again: Hello World!  
[Press Enter to close window]
```

**Note**

The preceding settings and procedures were carried out for the default [Debug](#) configuration. To create a [Release](#) configuration, select that configuration before setting the Project Properties.

## 21.4.6. MySQL Connector/C++ Getting Started: Usage Examples

The download package contains usage examples in the directory [examples/](#). These examples explain how to use the following classes:

- [Connection](#)
- [Driver](#)
- [PreparedStatement](#)
- [ResultSet](#)
- [ResultSetMetaData](#)
- [Statement](#)

The examples cover:

- Using the [Driver](#) class to connect to MySQL
- Creating tables, inserting rows, fetching rows using (simple) statements
- Creating tables, inserting rows, fetching rows using prepared statements
- Hints for working around prepared statement limitations
- Accessing result set metadata

The examples in this document are only code snippets, not complete programs. The code snippets provide a brief overview on the API. For complete programs, check the [examples/](#) directory of your MySQL Connector/C++ installation. Please also read the [README](#) file in that directory. To test

the example code, edit the `examples.h` file in the `examples/` directory to add your connection information, then rebuild the code by issuing a `make` command.

The examples in the `examples/` directory include:

- `examples/connect.cpp`:

How to create a connection, insert data into MySQL and handle exceptions.

- `examples/connection_meta_schemaobj.cpp`:

How to obtain metadata associated with a connection object, for example, a list of tables, databases, MySQL version, connector version.

- `examples/debug_output.cpp`:

How to activate and deactivate the MySQL Connector/C++ debug protocol.

- `examples/exceptions.cpp`:

A closer look at the exceptions thrown by the connector and how to fetch error information.

- `examples/prepared_statements.cpp`:

How to run Prepared Statements including an example how to handle SQL statements that cannot be prepared by the MySQL Server.

- `examples/resultset.cpp`:

How to use a cursor to fetch data and iterate over a result set.

- `examples/resultset_meta.cpp`:

How to obtain metadata associated with a result set, for example, number of columns and column types.

- `examples/resultset_types.cpp`:

Result sets returned from metadata methods. (This is more a test than an example.)

- `examples/standalone_example.cpp`:

Simple standalone program not integrated into regular `CMake` builds.

- `examples/statements.cpp`:

How to execute SQL statements without using Prepared Statements.

- `examples/cpp_trace_analyzer.cpp`:

This example shows how to filter the output of the `debug trace`. Please see the inline comments for further documentation. This script is unsupported.

### 21.4.6.1. MySQL Connector/C++ Connecting to MySQL

To establish a connection to MySQL Server, retrieve an instance of `sql::Connection` from a `sql::mysql::MySQL_Driver` object. A `sql::mysql::MySQL_Driver` object is returned by `sql::mysql::MySQL_Driver::get_mysql_driver_instance()`.

```
sql::mysql::MySQL_Driver *driver;
sql::Connection *con;

driver = sql::mysql::MySQL_Driver::get_mysql_driver_instance();
con = driver->connect("tcp://127.0.0.1:3306", "user", "password");
```



```
delete con;
```

Make sure that you free, `con`, the `sql::Connection` object as soon as you do not need it any more. But do not explicitly free `driver`, the connector object! The connector will take care of freeing that.

### 21.4.6.2. MySQL Connector/C++ Running a Simple Query

For running simple queries, you can use the methods `sql::Statement::execute()`, `sql::Statement::executeQuery()` and `sql::Statement::executeUpdate()`. Use the method `sql::Statement::execute()` if your query does not return a result set or if your query returns more than one result set. See the `examples/` directory for more information.

```
sql::mysql::MySQL_Driver *driver;
sql::Connection *con;
sql::Statement *stmt;

driver = sql::mysql::get_mysql_driver_instance();
con = driver->connect("tcp://127.0.0.1:3306", "user", "password");

stmt = con->createStatement();
stmt->execute("USE " EXAMPLE_DB);
stmt->execute("DROP TABLE IF EXISTS test");
stmt->execute("CREATE TABLE test(id INT, label CHAR(1))");
stmt->execute("INSERT INTO test(id, label) VALUES (1, 'a')");

delete stmt;
delete con;
```

Note that you must free `sql::Statement` and `sql::Connection` objects explicitly using `delete`.

### 21.4.6.3. MySQL Connector/C++ Fetching Results

The API for fetching result sets is identical for (simple) statements and prepared statements. If your query returns one result set, use `sql::Statement::executeQuery()` or `sql::PreparedStatement::executeQuery()` to run your query. Both methods return `sql::ResultSet` objects. The preview version does buffer all result sets on the client to support cursors.

```
// ...
sql::Connection *con;
sql::Statement *stmt;
sql::ResultSet *res;
// ...
stmt = con->createStatement();
// ...

res = stmt->executeQuery("SELECT id, label FROM test ORDER BY id ASC");
while (res->next()) {
    // You can use either numeric offsets...
    cout << "id = " << res->getInt(1); // getInt(1) returns the first column
    // ... or column names for accessing results.
    // The latter is recommended.
    cout << ", label = '" << res->getString("label") << "' << endl;
}

delete res;
delete stmt;
delete con;
```



#### Note

In the preceding code snippet, column indexing starts from 1.

Note that you must free `sql::Statement`, `sql::Connection` and `sql::ResultSet` objects explicitly using `delete`.

Cursor usage is demonstrated in the examples contained in the download package.

### 21.4.6.4. MySQL Connector/C++ Using Prepared Statements

If you are not familiar with Prepared Statements on MySQL, take a look at the source code comments and explanations in the file [examples/prepared\\_statement.cpp](#).

`sql::PreparedStatement` is created by passing an SQL query to `sql::Connection::prepareStatement()`. As `sql::PreparedStatement` is derived from `sql::Statement`, you will feel familiar with the API once you have learned how to use (simple) statements (`sql::Statement`). For example, the syntax for fetching results is identical.

```
// ...
sql::Connection *con;
sql::PreparedStatement *prep_stmt
// ...

prep_stmt = con->prepareStatement("INSERT INTO test(id, label) VALUES (?, ?)");

prep_stmt->setInt(1, 1);
prep_stmt->setString(2, "a");
prep_stmt->execute();

prep_stmt->setInt(1, 2);
prep_stmt->setString(2, "b");
prep_stmt->execute();

delete prep_stmt;
delete con;
```

As usual, you must free `sql::PreparedStatement` and `sql::Connection` objects explicitly.

### 21.4.6.5. MySQL Connector/C++ Complete Example 1

The following code shows a complete example of how to use MySQL Connector/C++:

```
/* Copyright 2008, 2010, Oracle and/or its affiliates. All rights reserved.

This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; version 2 of the License.

There are special exceptions to the terms and conditions of the GPL
as it is applied to this software. View the full text of the
exception in file EXCEPTIONS-CONNECTOR-C++ in the directory of this
software distribution.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
*/

/* Standard C++ includes */
#include <stdlib.h>
#include <iostream>

/*
   Include directly the different
   headers from cppconn/ and mysql_driver.h + mysql_util.h
   (and mysql_connection.h). This will reduce your build time!
*/
#include "mysql_connection.h"

#include <cppconn/driver.h>
#include <cppconn/exception.h>
#include <cppconn/resultset.h>
#include <cppconn/statement.h>
```

```

using namespace std;

int main(void)
{
cout << endl;
cout << "Running 'SELECT 'Hello World!' »
    AS _message'..." << endl;

try {
    sql::Driver *driver;
    sql::Connection *con;
    sql::Statement *stmt;
    sql::ResultSet *res;

    /* Create a connection */
    driver = get_driver_instance();
    con = driver->connect("tcp://127.0.0.1:3306", "root", "root");
    /* Connect to the MySQL test database */
    con->setSchema("test");

    stmt = con->createStatement();
    res = stmt->executeQuery("SELECT 'Hello World!' AS _message");
    while (res->next()) {
        cout << "\t... MySQL replies: ";
        /* Access column data by alias or column name */
        cout << res->getString("_message") << endl;
        cout << "\t... MySQL says it again: ";
        /* Access column data by numeric offset, 1 is the first column */
        cout << res->getString(1) << endl;
    }
    delete res;
    delete stmt;
    delete con;

} catch (sql::SQLException &e) {
    cout << "# ERR: SQLException in " << __FILE__;
    cout << "(" << __FUNCTION__ << " on line " <<
        << __LINE__ << endl;
    cout << "# ERR: " << e.what();
    cout << " (MySQL error code: " << e.getErrorCode();
    cout << ", SQLState: " << e.getSQLState() << " )" << endl;
}

cout << endl;

return EXIT_SUCCESS;
}

```

### 21.4.6.6. MySQL Connector/C++ Complete Example 2

The following code shows a complete example of how to use MySQL Connector/C++:

```

/* Copyright 2008, 2010, Oracle and/or its affiliates. All rights reserved.

This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; version 2 of the License.

There are special exceptions to the terms and conditions of the GPL
as it is applied to this software. View the full text of the
exception in file EXCEPTIONS-CONNECTOR-C++ in the directory of this
software distribution.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

```

```

*/

/* Standard C++ includes */
#include <stdlib.h>
#include <iostream>

/*
   Include directly the different
   headers from cppconn/ and mysql_driver.h + mysql_util.h
   (and mysql_connection.h). This will reduce your build time!
*/
#include "mysql_connection.h"

#include <cppconn/driver.h>
#include <cppconn/exception.h>
#include <cppconn/resultset.h>
#include <cppconn/statement.h>
#include <cppconn/prepared_statement.h>

using namespace std;

int main(void)
{
    cout << endl;
    cout << "Let's have MySQL count from 10 to 1..." << endl;

    try {
        sql::Driver *driver;
        sql::Connection *con;
        sql::Statement *stmt;
        sql::ResultSet *res;
        sql::PreparedStatement *pstmt;

        /* Create a connection */
        driver = get_driver_instance();
        con = driver->connect("tcp://127.0.0.1:3306", "root", "root");
        /* Connect to the MySQL test database */
        con->setSchema("test");

        stmt = con->createStatement();
        stmt->execute("DROP TABLE IF EXISTS test");
        stmt->execute("CREATE TABLE test(id INT)");
        delete stmt;

        /* '?' is the supported placeholder syntax */
        pstmt = con->prepareStatement("INSERT INTO test(id) VALUES (?)");
        for (int i = 1; i <= 10; i++) {
            pstmt->setInt(1, i);
            pstmt->executeUpdate();
        }
        delete pstmt;

        /* Select in ascending order */
        pstmt = con->prepareStatement("SELECT id FROM test ORDER BY id ASC");
        res = pstmt->executeQuery();

        /* Fetch in reverse = descending order! */
        res->afterLast();
        while (res->previous())
            cout << "\t... MySQL counts: " << res->getInt("id") << endl;
        delete res;

        delete pstmt;
        delete con;
    } catch (sql::SQLException &e) {
        cout << "# ERR: SQLException in " << __FILE__;
        cout << "(" << __FUNCTION__ << " on line " <<
            << __LINE__ << endl;
        cout << "# ERR: " << e.what();
        cout << " (MySQL error code: " << e.getErrorCode();
        cout << ", SQLState: " << e.getSQLState() << >

```

```

    " )" << endl;
}

cout << endl;

return EXIT_SUCCESS;
}

```

## 21.4.7. MySQL Connector/C++ Tutorials

The following tutorials illustrate various aspects of using MySQL Connector/C++. Also consult the examples in [Section 21.4.6, “MySQL Connector/C++ Getting Started: Usage Examples”](#).

### 21.4.7.1. Prerequisites and Background Information

This section describes the prerequisites that must be satisfied before you work through the remaining tutorial sections, and shows how to set up the framework code that serves as the basis for the tutorial applications.

These tutorials refer to tables and sample data from the `world` database, which you can download from the “Example Databases” section of the [MySQL Documentation](#) page.

Each tutorial application uses a framework consisting of the following code. The examples vary at the line that says `/* INSERT TUTORIAL CODE HERE! */` within the `try` block, which is replaced for each application with the application-specific code.

```

#include <stdlib.h>
#include <iostream>
#include <sstream>
#include <stdexcept>
/* uncomment for applications that use vectors */
/*#include <vector>*/

#include "mysql_connection.h"

#include <cppconn/driver.h>
#include <cppconn/exception.h>
#include <cppconn/resultset.h>
#include <cppconn/statement.h>
#include <cppconn/prepared_statement.h>

#define EXAMPLE_HOST "localhost"
#define EXAMPLE_USER "worlduser"
#define EXAMPLE_PASS "worldpass"
#define EXAMPLE_DB "world"

using namespace std;

int main(int argc, const char **argv)
{
    string url(argc >= 2 ? argv[1] : EXAMPLE_HOST);
    const string user(argc >= 3 ? argv[2] : EXAMPLE_USER);
    const string pass(argc >= 4 ? argv[3] : EXAMPLE_PASS);
    const string database(argc >= 5 ? argv[4] : EXAMPLE_DB);

    cout << "Connector/C++ tutorial framework..." << endl;
    cout << endl;

    try {

        /* INSERT TUTORIAL CODE HERE! */

    } catch (sql::SQLException &e) {
        /*
         * MySQL Connector/C++ throws three different exceptions:
         *
         * - sql::MethodNotImplementedException (derived from sql::SQLException)
         * - sql::InvalidArgumentException (derived from sql::SQLException)
         */
    }
}

```

```

- sql::SQLException (derived from std::runtime_error)
*/
cout << "# ERR: SQLException in " << __FILE__;
cout << "(" << __FUNCTION__ << ") on line " << __LINE__ << endl;
/* what() (derived from std::runtime_error) fetches error message */
cout << "# ERR: " << e.what();
cout << " (MySQL error code: " << e.getErrorCode();
cout << ", SQLState: " << e.getSQLState() << " )" << endl;

return EXIT_FAILURE;
}

cout << "Done." << endl;
return EXIT_SUCCESS;
}

```

Try the framework code as a standalone program using this procedure:

1. Copy and paste the framework code to a file such as `framework.cpp`. Edit the `#define` statements to reflect your connection parameters (server, user, password, database). Also, because the file contains those parameters, set its access mode to be readable only to yourself.
2. Compile the framework. For example, on Mac OS X, the command might look like this (enter the command on one line):

```

shell> g++ -o framework
-I/usr/local/include -I/usr/local/include/cppconn
-lmysqlcppconn framework.cpp

```

Adapt the command as necessary for your system. A similar command is needed for the tutorial applications that follow.

3. To run the framework, enter the following:

```

shell> ./framework

```

You will see a simple message:

```

Connector/C++ tutorial framework...
Done.

```

You are now ready to continue to the tutorials.

### 21.4.7.2. Calling Stored Procedures with `Statement` Objects

A stored procedure can be called using a `Statement` or `PreparedStatement` object. This section shows how to call stored procedures using `Statement` objects. To see how to use `PreparedStatement` objects, see [Section 21.4.7.3, "Calling Stored Procedures with PreparedStatement Objects"](#).

You can construct and call different types of stored procedures:

1. A stored procedure that returns no result. For example, such a stored procedure can log non-critical information, or change database data in a straightforward way.
2. A stored procedure that returns one or more values using output parameters. For example, such a procedure can indicate success or failure, or retrieve and return data items.
3. A stored procedure that returns one or more result sets. The procedure can execute one or more queries, each of which returns an arbitrary number of rows. Your application loops through each result set to display, transform, or otherwise process each row in it.

The following stored procedures illustrate each of these scenarios.

The following procedure adds a country to the `world` database, but does not return a result. This corresponds to Scenario 1 described earlier.

```
CREATE PROCEDURE add_country (IN country_code CHAR(3),
                             IN country_name CHAR(52),
                             IN continent_name CHAR(30))
BEGIN
  INSERT INTO Country(Code, Name, Continent)
    VALUES (country_code, country_name, continent_name);
END;
```

The next procedures use an output parameter to return the population of a specified country or continent, or the entire world. These correspond to Scenario 2 described earlier.

```
CREATE PROCEDURE get_pop (IN country_name CHAR(52),
                          OUT country_pop BIGINT)
BEGIN
  SELECT Population INTO country_pop FROM Country
    WHERE Name = country_name;
END;
```

```
CREATE PROCEDURE get_pop_continent (IN continent_name CHAR(30),
                                     OUT continent_pop BIGINT)
BEGIN
  SELECT SUM(Population) INTO continent_pop FROM Country
    WHERE Continent = continent_name;
END;
```

```
CREATE PROCEDURE get_pop_world (OUT world_pop BIGINT)
BEGIN
  SELECT SUM(Population) INTO world_pop FROM Country;
END;
```

The next procedure returns several result sets. This corresponds to Scenario 3 described earlier.

```
CREATE PROCEDURE get_data ()
BEGIN
  SELECT Code, Name, Population, Continent FROM Country
    WHERE Continent = 'Oceania' AND Population < 10000;
  SELECT Code, Name, Population, Continent FROM Country
    WHERE Continent = 'Europe' AND Population < 10000;
  SELECT Code, Name, Population, Continent FROM Country
    WHERE Continent = 'North America' AND Population < 10000;
END;
```

Enter and test the stored procedures manually to ensure that they will be available to your C++ applications. (Select `world` as the default database before you create them.) You are now ready to start writing applications using Connector/C++ that call stored procedures.

### Scenario 1: Using a [Statement](#) for a Stored Procedure That Returns No Result

This example shows how to call a stored procedure that returns no result set.

1. Make a copy of the tutorial framework code:

```
shell> cp framework.cpp sp_scenario1.cpp
```

2. Add the following code to the `try` block of the tutorial framework:

```
sql::Driver* driver = get_driver_instance();
std::auto_ptr<sql::Connection> con(driver->connect(url, user, pass));
con->setSchema(database);
std::auto_ptr<sql::Statement> stmt(con->createStatement());

// We need not check the return value explicitly. If it indicates
// an error, Connector/C++ generates an exception.
stmt->execute("CALL add_country('ATL', 'Atlantis', 'North America')");
```

3. Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
4. Run the program:

```
shell> ./sp_scenario1
```

- Using the `mysql` command-line client or other suitable program, check the `world` database to determine that it has been updated correctly. You can use this query:

```
mysql> SELECT Code, Name, Continent FROM Country WHERE Code='ATL';
+-----+-----+-----+
| Code | Name      | Continent |
+-----+-----+-----+
| ATL  | Atlantis | North America |
+-----+-----+-----+
```

The code in this application simply invokes the `execute` method, passing to it a statement that calls the stored procedure. The procedure itself returns no value, although it is important to note there is always a return value from the `CALL` statement; this is the `execute` status. MySQL Connector/C++ handles this status for you, so you need not handle it explicitly. If the `execute` call fails for some reason, it raises an exception that the `catch` block handles.

## Scenario 2: Using a `Statement` for a Stored Procedure That Returns an Output Parameter

This example shows how to handle a stored procedure that returns an output parameter.

- Make a copy of the tutorial framework code:

```
shell> cp framework.cpp sp_scenario2.cpp
```

- Add the following code to the `try` block of the tutorial framework:

```
sql::Driver* driver = get_driver_instance();
std::auto_ptr<sql::Connection> con(driver->connect(url, user, pass));
con->setSchema(database);
std::auto_ptr<sql::Statement> stmt(con->createStatement());

stmt->execute("CALL get_pop('Uganda', @pop)");

std::auto_ptr<sql::ResultSet> res(stmt->executeQuery("SELECT @pop AS _reply"));
while (res->next())
    cout << "Population of Uganda: " << res->getString("_reply") << endl;

stmt->execute("CALL get_pop_continent('Asia', @pop)");

res.reset(stmt->executeQuery("SELECT @pop AS _reply"));
while (res->next())
    cout << "Population of Asia: " << res->getString("_reply") << endl;

stmt->execute("CALL get_pop_world(@pop)");

res.reset(stmt->executeQuery("SELECT @pop AS _reply"));
while (res->next())
    cout << "Population of World: " << res->getString("_reply") << endl;
```

- Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
- Run the program:

```
shell> ./sp_scenario2
Connector/C++ tutorial framework...

Population of Uganda: 21778000
Population of Asia: 3705025700
Population of World: 6078749450
Done.
```

In this scenario, each stored procedure sets the value of an output parameter. This is not returned directly to the `execute` method, but needs to be obtained using a subsequent query. If you were executing the SQL statements directly, you might use statements similar to these:



```
CALL get_pop('Uganda', @pop);
SELECT @pop;
CALL get_pop_continent('Asia', @pop);
SELECT @pop;
CALL get_pop_world(@pop);
SELECT @pop;
```

In the C++ code, a similar sequence is carried out for each procedure call:

- Execute the `CALL` statement.
- Obtain the output parameter by executing an additional query. The query produces a `ResultSet` object.
- Retrieve the data using a `while` loop. The simplest way to do this is to use a `getString` method on the `ResultSet`, passing the name of the variable to access. In this example `_reply` is used as a placeholder for the variable and therefore is used as the key to access the correct element of the result dictionary.

Although the query used to obtain the output parameter returns only a single row, it is important to use the `while` loop to catch more than one row, to avoid the possibility of the connection becoming unstable.

### Scenario 3: Using a `Statement` for a Stored Procedure That Returns a Result Set

This example shows how to handle result sets produced by a stored procedure.



#### Note

This scenario requires MySQL 5.5.3 or higher. The client/server protocol does not support fetching multiple result sets from stored procedures prior to 5.5.3.

1. Make a copy of the tutorial framework code:

```
shell> cp framework.cpp sp_scenario3.cpp
```

2. Add the following code to the `try` block of the tutorial framework:

```
sql::Driver* driver = get_driver_instance();
std::auto_ptr<sql::Connection> con(driver->connect(url, user, pass));
con->setSchema(database);
std::auto_ptr<sql::Statement> stmt(con->createStatement());

stmt->execute("CALL get_data()");
std::auto_ptr< sql::ResultSet > res;
do {
    res.reset(stmt->getResultSet());
    while (res->next()) {
        cout << "Name: " << res->getString("Name")
            << " Population: " << res->getInt("Population")
            << endl;
    }
} while (stmt->getMoreResults());
```

3. Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
4. Run the program:

```
shell> ./sp_scenario3
Connector/C++ tutorial framework...

Name: Cocos (Keeling) Islands Population: 600
Name: Christmas Island Population: 2500
Name: Norfolk Island Population: 2000
Name: Niue Population: 2000
Name: Pitcairn Population: 50
```

```
Name: Tokelau Population: 2000
Name: United States Minor Outlying Islands Population: 0
Name: Svalbard and Jan Mayen Population: 3200
Name: Holy See (Vatican City State) Population: 1000
Name: Anguilla Population: 8000
Name: Atlantis Population: 0
Name: Saint Pierre and Miquelon Population: 7000
Done.
```

The code is similar to the examples shown previously. The code of particular interest here is:

```
do {
    res.reset(stmt->getResultSet());
    while (res->next()) {
        cout << "Name: " << res->getString("Name")
            << " Population: " << res->getInt("Population")
            << endl;
    }
} while (stmt->getMoreResults());
```

The `CALL` is executed as before, but this time the results are returned into multiple `ResultSet` objects because the stored procedure executes multiple `SELECT` statements. In this example, the output shows that three result sets are processed, because there are three `SELECT` statements in the stored procedure. Each result set returns more than one row.

The results are processed using this code pattern:

```
do {
    Get Result Set
    while (Get Result) {
        Process Result
    }
} while (Get More Result Sets);
```



#### Note

Use this pattern even if the stored procedure executes only a single `SELECT` and produces only one result set. This is a requirement of the underlying protocol.

### 21.4.7.3. Calling Stored Procedures with `PreparedStatement` Objects

This section shows how to call stored procedures using prepared statements. It is recommended that, before working through it, you first work through the previous tutorial [Section 21.4.7.2, "Calling Stored Procedures with `Statement` Objects"](#). That section shows the stored procedures required by the applications in this section.

#### Scenario 1: Using a `PreparedStatement` for a Stored Procedure That Returns No Result

This example shows how to call a stored procedure that returns no result set.

1. Make a copy of the tutorial framework code:

```
shell> cp framework.cpp ps_scenario1.cpp
```

2. Add the following code to the `try` block of the tutorial framework:

```
vector<string> code_vector;
code_vector.push_back("SLD");
code_vector.push_back("DSN");
code_vector.push_back("ATL");

vector<string> name_vector;
name_vector.push_back("Sealand");
```

```

name_vector.push_back("Disneyland");
name_vector.push_back("Atlantis");

vector<string> cont_vector;
cont_vector.push_back("Europe");
cont_vector.push_back("North America");
cont_vector.push_back("Oceania");

sql::Driver * driver = get_driver_instance();

std::auto_ptr< sql::Connection > con(driver->connect(url, user, pass));
con->setSchema(database);

std::auto_ptr< sql::PreparedStatement > pstmt;

pstmt.reset(con->prepareStatement("CALL add_country(?,?,?)"));
for (int i=0; i<3; i++)
{
    pstmt->setString(1,code_vector[i]);
    pstmt->setString(2,name_vector[i]);
    pstmt->setString(3,cont_vector[i]);

    pstmt->execute();
}

```

Also, uncomment `#include <vector>` near the top of the code, because vectors are used to store sample data.

3. Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
4. Run the program:

```
shell> ./ps_scenario1
```

5. You can check whether the database has been updated correctly by using this query:

```

mysql> SELECT Code, Name, Continent FROM Country
       -> WHERE Code IN('DSN','ATL','SLD');
+-----+-----+-----+
| Code | Name       | Continent |
+-----+-----+-----+
| ATL  | Atlantis  | Oceania   |
| DSN  | Disneyland| North America |
| SLD  | Sealand   | Europe    |
+-----+-----+-----+

```

The code is relatively simple, as no processing is required to handle result sets. The procedure call, `CALL add_country(?,?,?)`, is made using placeholders for input parameters denoted by '?'. These placeholders are replaced by the appropriate data values using the `PreparedStatement` object's `setString` method. The `for` loop is set up to iterate 3 times, as there are three data sets in this example. The same `PreparedStatement` is executed three times, each time with different input parameters.

## Scenario 2: Using a `PreparedStatement` for a Stored Procedure That Returns an Output Parameter

This example shows how to handle a stored procedure that returns an output parameter.

1. Make a copy of the tutorial framework code:

```
shell> cp framework.cpp ps_scenario2.cpp
```

2. Add the following code to the `try` block of the tutorial framework:

```

vector<string> cont_vector;
cont_vector.push_back("Europe");
cont_vector.push_back("North America");

```

```

cont_vector.push_back("Oceania");

sql::Driver * driver = get_driver_instance();

std::auto_ptr< sql::Connection > con(driver->connect(url, user, pass));
con->setSchema(database);

std::auto_ptr< sql::Statement > stmt(con->createStatement());
std::auto_ptr< sql::PreparedStatement > pstmt;
std::auto_ptr< sql::ResultSet > res;

pstmt.reset(con->prepareStatement("CALL get_pop_continent(?,@pop)"));

for (int i=0; i<3; i++)
{
    pstmt->setString(1,cont_vector[i]);
    pstmt->execute();
    res.reset(stmt->executeQuery("SELECT @pop AS _population"));
    while (res->next())
        cout << "Population of "
            << cont_vector[i]
            << " is "
            << res->getString("_population") << endl;
}

```

Also, uncomment `#include <vector>` near the top of the code, because vectors are used to store sample data.

3. Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
4. Run the program:

```

shell> ./ps_scenario2
Connector/C++ tutorial framework...

Population of Europe is 730074600
Population of North America is 482993000
Population of Oceania is 30401150
Done.

```

In this scenario a `PreparedStatement` object is created that calls the `get_pop_continent` stored procedure. This procedure takes an input parameter, and also returns an output parameter. The approach used is to create another statement that can be used to fetch the output parameter using a `SELECT` query. Note that when the `PreparedStatement` is created, the input parameter to the stored procedure is denoted by '?'. Prior to execution of the prepared statement, it is necessary to replace this placeholder by an actual value. This is done using the `setString` method:

```
pstmt->setString(1,cont_vector[i]);
```

Although the query used to obtain the output parameter returns only a single row, it is important to use the `while` loop to catch more than one row, to avoid the possibility of the connection becoming unstable.

### Scenario 3: Using a `PreparedStatement` for a Stored Procedure That Returns a Result Set

This example shows how to handle result sets produced by a stored procedure.



#### Note

This scenario requires MySQL 5.5.3 or higher. The client/server protocol does not support fetching multiple result sets from stored procedures prior to 5.5.3.

1. Make a copy of the tutorial framework code:

```
shell> cp framework.cpp ps_scenario3.cpp
```

2. Add the following code to the `try` block of the tutorial framework:

```

sql::Driver * driver = get_driver_instance();

std::auto_ptr< sql::Connection > con(driver->connect(url, user, pass));
con->setSchema(database);

std::auto_ptr< sql::PreparedStatement > pstmt;
std::auto_ptr< sql::ResultSet > res;

pstmt.reset(con->prepareStatement("CALL get_data()"));
res.reset(pstmt->executeQuery());

do {
    res.reset(pstmt->getResultSet());
    while (res->next()) {
        cout << "Name: " << res->getString("Name")
              << " Population: " << res->getInt("Population")
              << endl;
    }
} while (pstmt->getMoreResults());

```

3. Compile the program as described in [Section 21.4.7.1, "Prerequisites and Background Information"](#).
4. Run the program:

```
shell> ./ps_scenario3
```

5. Make a note of the output generated.

The code executes the stored procedure using a [PreparedStatement](#) object. The standard [do/while](#) construct is used to ensure that all result sets are fetched. The returned values are fetched from the result sets using the [getInt](#) and [getString](#) methods.

## 21.4.8. MySQL Connector/C++ Debug Tracing

Although a debugger can be used to debug your application, you may find it beneficial to turn on the debug traces of the connector. Some problems happen randomly which makes them difficult to debug using a debugger. In such cases, debug traces and protocol files are more useful because they allow you to trace the activities of all instances of your program.

DTrace is a very powerful technology to trace any application without having to develop an extra trace module for your application. Unfortunately, DTrace is currently only available on Solaris, Mac OS X 10.5, and FreeBSD.

MySQL Connector/C++ can write two trace files:

1. Trace file generated by the MySQL client library
2. Trace file generated internally by MySQL Connector/C++

The first trace file can be generated by the underlying MySQL client library ([libmysqlclient](#)). To enable this trace, the connector will call the [mysql\\_debug\(\)](#) C API function internally. Because only debug versions of the MySQL client library are capable of writing a trace file, compile MySQL Connector/C++ against a debug version of the library to use this trace. The trace shows the internal function calls and the addresses of internal objects as shown here:

```

>mysql_stmt_init
| >_mymalloc
| | enter: Size: 816
| | exit: ptr: 0x68e7b8
| <_mymalloc | >init_alloc_root
| | enter: root: 0x68e7b8
| | >_mymalloc
| | | enter: Size: 2064
| | | exit: ptr: 0x68eb28
[...]
```

The second trace is the MySQL Connector/C++ internal trace. It is available with debug and nondebug builds of the connector as long as you have enabled the tracing module at compile time using `cmake -DMYSQLCPPCONN_TRACE_ENABLE:BOOL=1`. By default, tracing functionality is not available and calls to trace functions are removed by the preprocessor.

Compiling the connector with tracing functionality enabled causes two additional tracing function calls per each connector function call. For example:

```
| INF: Tracing enabled
| <MySQL_Connection::setClientOption
| >MySQL_Prepared_Statement::setInt
| INF: this=0x69a2e0
| >MySQL_Prepared_Statement::checkClosed
| <MySQL_Prepared_Statement::checkClosed
| <MySQL_Prepared_Statement::setInt
| [...]

```

Run your own benchmark to find out how much this will impact the performance of your application.

A simple test using a loop running 30,000 `INSERT` SQL statements showed no significant real-time impact. The two variants of this application using a trace enabled and trace disabled version of the connector performed equally well. The runtime measured in real time was not significantly impacted as long as writing a debug trace was not enabled. However, there will be a difference in the time spent in the application. When writing a debug trace, the I/O subsystem may become a bottleneck.

In summary, use connector builds with tracing enabled carefully. Trace-enabled versions may cause higher CPU usage even if the overall runtime of your application is not impacted significantly.

The example from `examples/debug_output.cpp` demonstrates how to activate the debug traces in your program. Currently they can only be activated through API calls. The traces are controlled on a per-connection basis. You can use the `setClientOptions()` method of a connection object to activate and deactivate trace generation. The MySQL client library trace is always written to a file, whereas the connector's protocol messages are printed to the standard output.

```
sql::Driver *driver;
int on_off = 1;

/* Using the Driver to create a connection */
driver = get_driver_instance();
std::auto_ptr< sql::Connection > con(driver->connect(host, user, pass));

/*
Activate debug trace of the MySQL client library (C API)
Only available with a debug build of the MySQL client library!
*/
con->setClientOption("libmysql_debug", "d:t:O,client.trace");

/*
Connector/C++ tracing is available if you have compiled the
driver using cmake -DMYSQLCPPCONN_TRACE_ENABLE:BOOL=1
*/
con->setClientOption("client_trace", &on_off);

```

## 21.4.9. MySQL Connector/C++ Usage Notes

MySQL Connector/C++ is compatible with the JDBC 4.0 API. See the [JDBC overview](#) for information on JDBC 4.0. Please also check the `examples/` directory of the download package.

- The MySQL Connector/C++ `sql::DataType` class defines the following JDBC standard data types: `UNKNOWN`, `BIT`, `TINYINT`, `SMALLINT`, `MEDIUMINT`, `INTEGER`, `BIGINT`, `REAL`, `DOUBLE`, `DECIMAL`, `NUMERIC`, `CHAR`, `BINARY`, `VARCHAR`, `VARBINARY`, `LONGVARCHAR`, `LONGVARBINARY`, `TIMESTAMP`, `DATE`, `TIME`, `GEOMETRY`, `ENUM`, `SET`, `SQLNULL`.

MySQL Connector/C++ does not support the following JDBC standard data types: `ARRAY`, `BLOB`, `CLOB`, `DISTINCT`, `FLOAT`, `OTHER`, `REF`, `STRUCT`.

- `DatabaseMetaData::supportsBatchUpdates()` returns `true` because MySQL supports batch updates in general. However, the MySQL Connector/C++ API provides no API calls for batch updates.
- Two non-JDBC methods let you fetch and set unsigned integers: `getUInt64()` and `getUInt()`. These are available for `ResultSet` and `PreparedStatement`:
  - `ResultSet::getUInt64()`
  - `ResultSet::getUInt()`
  - `PreparedStatement::setUInt64()`
  - `PreparedStatement::setUInt()`

The corresponding `getLong()` and `setLong()` methods have been removed.

- The `DatabaseMetaData::getColumns()` method has 23 columns in its result set, rather than the 22 columns defined by JDBC. The first 22 columns are as described in the JDBC documentation, but column 23 is new:

23. `IS_AUTOINCREMENT`: A string which is “YES” if the column is an auto-increment column, “NO” otherwise.

- MySQL Connector/C++ may return different metadata for the same column, depending on the method you call.

Suppose that you have a column that accepts a character set and a collation in its specification and you specify a binary collation, such as:

```
VARCHAR(20) CHARACTER SET utf8 COLLATE utf8_bin
```

The server sets the `BINARY` flag in the result set metadata of this column. The `ResultSetMetadata::getColumnTypeName()` method uses the metadata and will report, due to the `BINARY` flag, that the column type name is `BINARY`. This is illustrated here:

```
mysql> CREATE TABLE varbin (a VARCHAR(20) CHARACTER SET utf8 COLLATE utf8_bin);
Query OK, 0 rows affected (0.00 sec)

mysql> select * from varbin;
Field  1: `a`
Catalog: `def`
Database: `test`
Table: `varbin`
Org_table: `varbin`
Type: VAR_STRING
Collation: latin1_swedish_ci (8)
Length: 20
Max_length: 0
Decimals: 0
Flags: BINARY

0 rows in set (0.00 sec)

mysql> SELECT * FROM INFORMATION_SCHEMA.COLUMNS WHERE TABLE_NAME='varbin'\G
***** 1. row *****
      TABLE_CATALOG: NULL
      TABLE_SCHEMA: test
      TABLE_NAME: varbin
      COLUMN_NAME: a
      ORDINAL_POSITION: 1
      COLUMN_DEFAULT: NULL
      IS_NULLABLE: YES
      DATA_TYPE: varchar
CHARACTER_MAXIMUM_LENGTH: 20
CHARACTER_OCTET_LENGTH: 60
```

```

NUMERIC_PRECISION: NULL
NUMERIC_SCALE: NULL
CHARACTER_SET_NAME: utf8
COLLATION_NAME: utf8_bin
COLUMN_TYPE: varchar(20)
COLUMN_KEY:
EXTRA:
PRIVILEGES: select,insert,update,references
COLUMN_COMMENT:
1 row in set (0.01 sec)

```

However, `INFORMATION_SCHEMA` gives no hint in its `COLUMNS` table that metadata will contain the `BINARY` flag. `DatabaseMetaData::getColumns()` uses `INFORMATION_SCHEMA` and will report the type name `VARCHAR` for the same column. It also returns a different type code.

- When inserting or updating `BLOB` or `TEXT` columns, MySQL Connector/C++ developers are advised not to use `setString()`. Instead, use the dedicated `setBlob()` API function.

The use of `setString()` can cause a `Packet too large` error message. The error occurs if the length of the string passed to the connector using `setString()` exceeds `max_allowed_packet [502]` (minus a few bytes reserved in the protocol for control purposes). This situation is not handled in MySQL Connector/C++, because it could lead to security issues, such as extremely large memory allocation requests due to malevolently long strings.

If `setBlob()` is used, this problem does not arise because `setBlob()` takes a streaming approach based on `std::istream`. When sending the data from the stream to MySQL Server, MySQL Connector/C++ splits the stream into chunks appropriate for MySQL Server using the current `max_allowed_packet [502]` setting.



### Caution

When using `setString()`, it is not possible to set `max_allowed_packet [502]` to a value large enough for the string prior to passing it to MySQL Connector/C++. That configuration option cannot be changed within a session.

This difference from the JDBC specification ensures that MySQL Connector/C++ is not vulnerable to memory flooding attacks.

- In general, MySQL Connector/C++ works with MySQL 5.0, but it is not completely supported. Some methods may not be available when connecting to MySQL 5.0. This is because the Information Schema is used to obtain the requested information. There are no plans to improve the support for 5.0 because the current GA version of MySQL Server is 5.6. MySQL Connector/C++ is primarily targeted at the MySQL Server GA version that is available on its release.

The following methods throw a `sql::MethodNotImplemented` exception when you connect to a MySQL server earlier than 5.1:

- `DatabaseMetaData::getCrossReference()`
- `DatabaseMetaData::getExportedKeys()`
- MySQL Connector/C++ includes a `Connection::getClientOption()` method that is not included in the JDBC API specification. The prototype is:

```
void getClientOption(const std::string & optionName, void * optionValue)
```

The method can be used to check the value of connection properties set when establishing a database connection. The values are returned through the `optionValue` argument passed to the method with the type `void *`.

Currently, `getClientOption()` supports fetching the `optionValue` of the following options:



- `metadataUseInfoSchema`
- `defaultStatementResultType`
- `defaultPreparedStatementResultType`

The `metadataUseInfoSchema` connection option controls whether to use the `Information_Schemata` for returning the metadata of `SHOW` statements:

- For `metadataUseInfoSchema`, interpret the `optionValue` argument as a boolean upon return.
- For `defaultStatementResultType` and `defaultPreparedStatementResultType`, interpret the `optionValue` argument as an integer upon return.

The connection property can be set either when establishing the connection through the connection property map, or using `void Connection::setClientOption(const std::string & optionName, const void * optionValue)` where `optionName` is assigned the value `metadataUseInfoSchema`.

Some examples:

```
bool isInfoSchemaUsed;
conn->getClientOption("metadataUseInfoSchema", (void *) &isInfoSchemaUsed);

int defaultStmtResType;
int defaultPStmtResType;
conn->getClientOption("defaultStatementResultType", (void *) &defaultStmtResType);
conn->getClientOption("defaultPreparedStatementResultType", (void *) &defaultPStmtResType);
```

- MySQL Connector/C++ supports the following methods not found in the JDBC API standard:

```
std::string MySQL_Connection::getSessionVariable(const std::string & varname)
```

```
void MySQL_Connection::setSessionVariable(const std::string & varname, const std::string & value)
```

These methods get and set MySQL session variables. Both are members of the `MySQL_Connection` class.

`getSessionVariable()` is equivalent to executing the following and fetching the first return value:

```
SHOW SESSION VARIABLES LIKE "<varname>"
```

You can use the “%” and “\_” SQL pattern characters in `<varname>`.

`setSessionVariable()` is equivalent to executing:

```
SET SESSION <varname> = <value>
```

- Fetching the value of a column can sometimes return different values depending on whether the call is made from a Statement or Prepared Statement. This is because the protocol used to communicate with the server differs depending on whether a Statement or Prepared Statement is used.

To illustrate this, consider the case where a column has been defined as type `BIGINT`. The most negative `BIGINT` value is then inserted into the column. If a Statement and Prepared Statement are created that perform a `GetUInt64()` call, then the results will be different in each case. The Statement returns the maximum positive value for `BIGINT`. The Prepared Statement returns 0.

The difference results from the fact that Statements use a text protocol, and Prepared Statements use a binary protocol. With the binary protocol in this case, a binary value is returned from the server that can be interpreted as an `int64`. In the preceding scenario, a very large negative value is fetched with `GetUInt64()`, which fetches unsigned integers. Because the large negative value cannot be sensibly converted to an unsigned value, 0 is returned.

In the case of the Statement, which uses the text protocol, values are returned from the server as strings, and then converted as required. When a string value is returned from the server in the preceding scenario, the large negative value must be converted by the runtime library function `strtoul()`, which `GetUInt64()` calls. The behavior of `strtoul()` is dependent upon the specific runtime and host operating system, so the results can be platform dependent. In the case, given a large positive value was actually returned.

Although it is very rare, there are some cases where Statements and Prepared Statements can return different values unexpectedly, but this usually only happens in extreme cases such as the one mentioned.

- The JDBC documentation [lists many fields](#) for the `DatabaseMetaData` class. JDBC also appears to [define certain values](#) for those fields. However, MySQL Connector/C++ does not define certain values for those fields. Internally enumerations are used and the compiler determines the values to assign to a field.

To compare a value with the field, use code such as the following, rather than making assumptions about specific values for the attribute:

```
// dbmeta is an instance of DatabaseMetaData
if (myvalue == dbmeta->attributeNoNulls) {
    ...
}
```

Usually `myvalue` will be a column from a result set holding metadata information. MySQL Connector/C++ does not guarantee that `attributeNoNulls` is 0. It can be any value.

- When programming stored procedures, JDBC has available an extra class, an extra abstraction layer for callable statements, the `CallableStatement` class. As this class is not present in MySQL Connector/C++, use the methods from the `Statement` and `PreparedStatement` classes to execute a stored procedure using `CALL`.

## 21.4.10. MySQL Connector/C++ Known Bugs and Issues

Please report bugs through the MySQL Bug System. See [Section 1.7, “How to Report Bugs or Problems”](#).

For release notes detailing the changes in each release of Connector/C++, see [MySQL Connector/C++ Release Notes](#).

- When linking against a static library for 1.0.3 on Windows, define `CPPDBC_PUBLIC_FUNC` either in the compiler options (preferable) or with `/D "CPPCONN_PUBLIC_FUNC="`. You can also explicitly define it in your code by placing `#define CPPCONN_PUBLIC_FUNC` before the header inclusions.
- Generally speaking, C++ library binaries are less portable than C library binaries. Issues can be caused by name mangling, different Standard Template Library (STL) versions, and using different compilers and linkers for linking against the libraries than were used for building the library itself.

Even a small change in the compiler version can cause problems. If you obtain error messages that you suspect are related to binary incompatibilities, build MySQL Connector/C++ from source, using the same compiler and linker that you will use to build and link your application.

Due to the variations between Linux distributions, compiler and linker versions and STL versions, it is not possible to provide binaries for each and every possible configuration. However, the MySQL Connector/C++ binary distributions contain a [README](#) file that describes the environment and settings used to build the binary versions of the libraries.

- To avoid potential crashes, the build configuration of MySQL Connector/C++ should match the build configuration of the application using it. For example, do not use the release build of MySQL Connector/C++ with a debug build of the client application.

## 21.4.11. MySQL Connector/C++ Support

For general discussion of the MySQL Connector/C++ please use the [C/C++ community forum](#) or join the [MySQL Connector/C++ mailing list](#).

Bugs can be reported at the [MySQL bug Web site](#).

For release notes detailing the changes in each release of MySQL Connector/C++, see [MySQL Connector/C++ Release Notes](#).

For Licensing questions, and to purchase MySQL Products and Services, please see <http://www.mysql.com/buy-mysql/>

## 21.5. MySQL Connector/C

MySQL Connector/C is a client library that implements the C API for client/server communication. It is a standalone replacement for the MySQL client library shipped with MySQL Server distributions. See [Section 21.8.1, “MySQL C API Implementations”](#).

Reasons to use MySQL Connector/C:

- If you need only the client library, MySQL Connector/C provides everything required. There is no need to compile or install the MySQL Server package, which is much larger.
- MySQL Connector/C does not rely on the MySQL Server release cycle, so bug fixes and new features can be distributed independently of MySQL Server releases.

For documentation of the C API implemented by MySQL Connector/C, see [Section 21.8, “MySQL C API”](#).

For release notes detailing the changes in each release of MySQL Connector/C, see [MySQL Connector/C Release Notes](#).

The following discussion covers these topics:

- MySQL Connector/C supported platforms
- MySQL Connector/C distribution contents
- Obtaining and installing MySQL Connector/C
- Building client programs that use MySQL Connector/C

### 21.5.1. MySQL Connector/C Supported Platforms

The following table indicates the platforms on which MySQL Connector/C has been tested.

**Table 21.28. Supported Platforms for MySQL Connector/C 6.0.2**

Platform	Versions
Windows	Windows 2000, XP, Vista, Server 2003, Server 2008
Linux	RHEL 4, 5; SuSE LES 9, 10; generic <code>glibc</code> 2-based Linux systems
Solaris	8, 9, 10
Mac OS X	10.4, 10.5
FreeBSD	6, 7
HP-UX	11.11, 11.23, 11.31
IBM AIX	5.3

## 21.5.2. MySQL Connector/C Distribution Contents

A MySQL Connector/C distribution contains the header, library, and utility files necessary to build MySQL client applications that communicate with MySQL Server using the C API.

Distributions are available in binary and source formats. A binary distribution contains the header, library, and utility components discussed following, compiled and ready for use in writing client programs. A source distribution contains the source files required to produce the same headers, libraries, and utilities included in a binary distribution, but you compile them yourself.

MySQL Connector/C distributions include these components:

- A set of `.h` header files that C applications include at compile time. These files are located in the `include` directory.
- Static and dynamic libraries that C applications link to at link time. These libraries are located in the `lib` directory. The library names depend on the library type and platform for which a distribution is built:
  - On Unix (and Unix-like) systems, the static library is `libmysqlclient.a`. The dynamic library is `libmysqlclient.so` on most Unix systems and `libmysqlclient.dylib` on Mac OS X.
  - On Windows, the static library is `mysqlclient.lib` and the dynamic library is `libmysql.dll`. Windows distributions also include `libmysql.lib`, a static import library needed for using the dynamic library.
- Utilities. MySQL Connector/C includes the `mysql_config` utility, located in the `bin` directory.

## 21.5.3. Obtaining MySQL Connector/C

To obtain a distribution, visit [MySQL Connector/C downloads](#). For a list of supported platforms, see [Section 21.5.1, “MySQL Connector/C Supported Platforms”](#).

## 21.5.4. Installing MySQL Connector/C from a Binary Distribution

On Unix and Unix-like systems, a MySQL Connector/C binary distribution is packaged as a compressed `tar` file, denoted here as `PACKAGE.tar.gz`. Download the distribution file and unpack it in the intended installation directory using this command:

```
shell> tar zxvf PACKAGE.tar.gz
```

Windows users can choose between two binary packaging formats:

- Windows MSI Installer (`.msi` file): To use the MSI Installer, launch it and follow the prompts in the screens it presents to install MySQL Connector/C in the location of your choosing.
- Zip archive without installer (`.zip` file): To use a Zip archive, unpack it in the intended installation directory using `WinZip` or another tool that can read `.zip` files.

`.tar.gz` and `.zip` packages unpack into a directory named `PACKAGE`, which is a value such as `mysql-connector-c-6.0.2-linux-rhel5-x86-64bit`. If you want a simpler name to work with, rename the directory. On Unix, an alternative is to create a symbolic link with a simpler name:

```
shell> ln -s mysql-connector-c-6.0.2-linux-rhel5-x86-64bit connector-c
```

## 21.5.5. Installing MySQL Connector/C from Source

A MySQL Connector/C source distribution is packaged as a compressed `tar` file or Zip archive, denoted here as `PACKAGE.tar.gz` or `PACKAGE.zip`. A source distribution can be used on any supported platform listed in [Section 21.5.1, “MySQL Connector/C Supported Platforms”](#).

Download the distribution file and unpack it in the intended installation directory. To unpack a compressed `tar` file, use this command:

```
shell> tar zxvf PACKAGE.tar.gz
```

To unpack a Zip archive, use [WinZip](#) or another tool that can read `.zip` files.

After unpacking the distribution, build it using the following instructions appropriate for your platform.

### 21.5.5.1. Building MySQL Connector/C from Source on Unix

If the native compiler toolset for the target platform is available (for example, SunStudio for Solaris), you can use that for compilation. Alternatively, the GNU toolset can be used on all platforms.

You also need [CMake](#) 2.6 or newer, which is available from [cmake.org](#).

To build and install the source distribution, use the following procedure:

1. Change location to the top-level directory of the source distribution.
2. Generate the [Makefile](#):

```
shell> cmake -G "Unix Makefiles"
```

Or, for a Debug build:

```
shell> cmake -G "Unix Makefiles" -DCMAKE_BUILD_TYPE=Debug
```

By default, the installation location for MySQL Connector/C is `/usr/local/mysql`. To change this location, use the `CMAKE_INSTALL_PREFIX` option to specify a different directory when generating the [Makefile](#). For example:

```
shell> cmake -G "Unix Makefiles" -DCMAKE_INSTALL_PREFIX=/opt/local/mysql
```

For other [CMake](#) options that you might find useful, see [Section 21.5.5.3, "Other Build Options"](#).

3. Build the project:

```
shell> make
```

4. As `root`, install the MySQL Connector/C headers, libraries, and utilities:

```
root-shell> make install
```

### 21.5.5.2. Building MySQL Connector/C from Source on Microsoft Windows

MySQL Connector/C can be built on several versions of Microsoft Windows, listed in [Section 21.5.1, "MySQL Connector/C Supported Platforms"](#).

Microsoft Visual Studio 8 or 9 is recommended. The Express Edition of Visual Studio and other compilers might work, but are untested.

You also need [CMake](#) 2.6 or newer, which is available from [cmake.org](#).

To build and install the source distribution, use the following procedure:

1. Set the environment variables for the Visual Studio toolchain. Visual Studio includes a batch file to set these for you, and installs a shortcut in the **Start** menu to open a command prompt with these variables set.
2. Change location to the top-level directory of the source distribution.
3. Generate the [Makefile](#) by entering the following command in a command-prompt window:

```
shell> cmake -G "Visual Studio 9 2008"
```

For other [CMake](#) options that you might find useful, see [Section 21.5.5.3, "Other Build Options"](#).

The result of the `cmake` command is a project (solution) file, `libmysql.sln`, that you can open with Visual Studio. Alternatively, build from the command line with either of these commands:

```
shell> devenv.com libmysql.sln /build Release
```

```
shell> devenv.com libmysql.sln /build RelWithDebInfo
```

For other versions of Visual Studio or for an `nmake`-based build, use the following command to check which generators can be specified with the `-G` option:

```
shell> cmake --help
```

To compile a Debug build, you must set the `CMake` build type so the correct external library versions are used, then compile using the `Debug` solution configuration:

```
shell> cmake -G "Visual Studio 9 2008" -DCMAKE_BUILD_TYPE=Debug
shell> devenv.com libmysql.sln /build Debug
```

A normal build builds the C API libraries for the `lib` directory. A Debug build additionally builds debug libraries for the `lib/debug` directory. You must use the debug libraries when compiling clients built using the debug C runtime.

4. Use the install operation provided by your development environment to install the MySQL Connector/C headers, libraries, and utilities. You can also use this `CMake` command:

```
shell> cmake --build . --target INSTALL --config RelWithDebInfo
```

### 21.5.5.3. Other Build Options

The following table shows other options that can be used when building MySQL Connector/C from source.

**Table 21.29. Build Options for MySQL Connector/C**

Build Option	Description
<code>-DWITH_OPENSSL=1</code>	Enable dynamic linking to the system OpenSSL library.
<code>-DWITH_EXTERNAL_ZLIB=1</code>	Enable dynamic linking to the system Zlib library.

### 21.5.6. Testing MySQL Connector/C

To test MySQL Connector/C, a running MySQL server instance must be available. Before you run the test suite, specify the following environment variables:

- `MYSQL_TEST_HOST`: The host where the MySQL server is running (default `localhost`)
- `MYSQL_TEST_USER`: The user name of the MySQL account to use
- `MYSQL_TEST_PASSWD`: The password of the MySQL account to use
- `MYSQL_TEST_PORT`: The TCP/IP port to connect to
- `MYSQL_TEST_SOCKET`: The socket file to connect to
- `MYSQL_TEST_DB`: The default database to use (default `test`)

To run the test suite, execute `ctest` from the command line:

```
shell> ctest
```

### 21.5.7. Building MySQL Connector/C Applications

For instructions on building C applications to use MySQL Connector/C, see [Section 21.8.3.1, "Building C API Client Programs"](#). To enable your compiler to find the header and library files under the directory

where you installed MySQL Connector/C, specify the appropriate compile-time options, as indicated in that section.

## 21.6. MySQL Connector/Python

MySQL Connector/Python allows Python programs to access MySQL databases, using an API that is compliant with the [Python DB API version 2.0](#). It is written in pure Python and does not have any dependencies except for the [Python Standard Library](#).

For release notes detailing the changes in each release of Connector/Python, see [MySQL Connector/Python Release Notes](#).

MySQL Connector/Python includes support for:

- Almost all features provided by MySQL Server up to and including MySQL Server version 5.5.
- Converting parameter values back and forth between Python and MySQL data types, for example Python `datetime` and MySQL `DATETIME`. You can turn automatic conversion on for convenience, or off for optimal performance.
- All MySQL extensions to standard SQL syntax.
- Protocol compression, which enables compressing the data stream between the client and server.
- Connections using TCP/IP sockets and on Unix using Unix sockets.
- Secure TCP/IP connections using SSL.
- Self-contained driver. Connector/Python does not require the MySQL client library or any Python modules outside the standard library.

MySQL Connector/Python supports from Python version 2.4 through 2.7, and Python 3.1 and later. Note that Connector/Python does not support the old MySQL Server authentication methods, which means that MySQL versions prior to 4.1 will not work.

### 21.6.1. Guidelines for Python Developers

The following guidelines cover aspects of developing MySQL applications that might not be immediately obvious to developers coming from a Python background:

- For security, do not hardcode the values needed to connect and log into the database in your main script. Python has the convention of a `config.py` module, where you can keep such values separate from the rest of your code.
- Python scripts often build up and tear down large data structures in memory, up to the limits of available RAM. Because MySQL often deals with data sets that are many times larger than available memory, techniques that optimize storage space and disk I/O are especially important. For example, in MySQL tables, you typically use numeric IDs rather than string-based dictionary keys, so that the key values are compact and have a predictable length. This is especially important for columns that make up the [primary key](#) for an [InnoDB](#) table, because those column values are duplicated within each [secondary index](#).
- Any application that accepts input must expect to handle bad data.

The bad data might be accidental, such as out-of-range values or misformatted strings. The application can use server-side checks such as [unique constraints](#) and [NOT NULL constraints](#), to keep the bad data from ever reaching the database. On the client side, use techniques such as exception handlers to report any problems and take corrective action.

The bad data might also be deliberate, representing a “SQL injection” attack. For example, input values might contain quotation marks, semicolons, `%` and `_` wildcard characters and other characters



significant in SQL statements. Validate input values to make sure they have only the expected characters. Escape any special characters that could change the intended behavior when substituted into a SQL statement. Never concatenate a user input value into a SQL statement without doing validation and escaping first. Even when accepting input generated by some other program, expect that the other program could also be hacked and be sending you incorrect or malicious data.

- Because the result sets from SQL queries can be very large, use the appropriate method to retrieve items from the result set as you loop through them. `fetchone()` retrieves a single item, when you know the result set contains a single row. `fetchall()` retrieves all the items, when you know the result set contains a limited number of rows that can fit comfortably into memory. `fetchmany()` is the general-purpose method when you cannot predict the size of the result set: you keep calling it and looping through the returned items, until there are no more results to process.
- Since Python already has convenient modules such as `pickle` and `cPickle` to read and write data structures on disk, the data that you choose store in MySQL instead is likely to have special characteristics:
  - **Too large to all fit in memory at one time.** You use `SELECT` statements to query only the precise items you need, and `aggregate functions` to perform calculations across multiple items. You configure the `innodb_buffer_pool_size [1708]` option within the MySQL server to dedicate a certain amount of RAM for caching query results.
  - **Too complex to be represented by a single data structure.** You divide the data between different SQL tables. You can recombine data from multiple tables by using a `join` query. You make sure that related data is kept in sync between different tables by setting up `foreign key` relationships.
  - **Updated frequently, perhaps by multiple users simultaneously.** The updates might only affect a small portion of the data, making it wasteful to write the whole structure each time. You use the SQL `INSERT`, `UPDATE`, and `DELETE` statements to update different items concurrently, writing only the changed values to disk. You use `InnoDB` tables and `transactions` to keep write operations from conflicting with each other, and to return consistent query results even as the underlying data is being updated.
- Building in MySQL best practices for performance can help your application to scale without requiring major rewrites and architectural changes. See [Chapter 8, Optimization](#) for best practices for MySQL performance. It offers guidelines and tips for SQL tuning, database design, and server configuration.
- You can avoid reinventing the wheel by learning the MySQL SQL statements for common operations: operators to use in queries, techniques for bulk loading data, and so on. Some statements and clauses are extensions to the basic ones defined by the SQL standard. See [Section 13.2, “Data Manipulation Statements”](#), [Section 13.1, “Data Definition Statements”](#), and [Section 13.2.9, “SELECT Syntax”](#) for the main classes of statements.
- Issuing SQL statements from Python typically involves declaring very long, possibly multi-line string literals. Because string literals within the SQL statements could be enclosed by single quotation, double quotation marks, or contain either of those characters, for simplicity you can use Python's triple-quoting mechanism to enclose the entire statement. For example:

```
'''It doesn't matter if this string contains 'single'
or "double" quotes, as long as there aren't 3 in a
row.'''
```

You can use either of the `'` or `"` characters for triple-quoting multi-line string literals.

- Many of the secrets to a fast, scalable MySQL application involve using the right syntax at the very start of your setup procedure, in the `CREATE TABLE` statements. For example, Oracle recommends the `ENGINE=INNODB` clause for most tables, and makes it the default in MySQL 5.5 and up. Using `InnoDB` tables enables transactional behavior that helps scalability of read-write workloads and offers automatic `crash recovery`. A follow-on recommendation is to declare a numeric `primary key` for each table, which offers the fastest way to look up values and can act as a pointer to associated



values in other tables (a [foreign key](#)). Also within the `CREATE TABLE` statement, using the most compact column data types that meet your application requirements helps performance and scalability, as the database server moves large amounts of data back and forth between memory and disk.

## 21.6.2. Connector/Python Versions

MySQL Connector/Python v1.0.x series went through a series of beta releases, leading to the first generally available (GA) version 1.0.7. Any development releases prior to general availability are not supported now that the GA version is released.

The following table summarizes the available Connector/Python versions:

**Table 21.30. Connector/Python Version Reference**

Connector/Python Version	MySQL Server Versions	Python Versions	Support Status for Connector
1.0	5.6, 5.5 (5.1, 5.0, 4.1)	2.7, 2.6 (2.5, 2.4); 3.1 and later	Recommended version



### Note

MySQL server and Python versions within brackets are known to work with Connector/Python, but are not officially supported. Bugs might not get fixed for those versions.

## 21.6.3. Connector/Python Installation

Connector/Python runs on any platform where Python is installed. Python comes pre-installed on almost any Linux distribution or UNIX-like system such as Apple Mac OS X and FreeBSD. On Microsoft Windows systems, you can install Python using the installer found on the [Python Download website](#).

Connector/Python is a pure Python implementation of the MySQL Client/Server protocol, meaning it does not require any other MySQL client libraries or other components. It also has no third-party dependencies. If you need SSL support, verify that your Python installation has been compiled using the [OpenSSL](#) libraries.

The installation of Connector/Python is similar on every platform and follows the standard [Python Distribution Utilities](#) or [Distutils](#). Some platforms have specific packaging, for example RPM, and, when made available, the installation of these will be covered in this manual.

Python terminology regarding distributions:

- **Source Distribution** is a distribution that contains only source files and is generally platform independent.
- **Built Distribution** can be regarded as a binary package. It contains both sources and platform-independent bytecode.

### 21.6.3.1. Installing Connector/Python Source Distribution on Linux, UNIX, or Mac OS X

On UNIX-like systems such as Linux distributions, Solaris, Apple Mac OS X, and FreeBSD, you can download Connector/Python as a `tar` archive from <http://dev.mysql.com/downloads/connector/python/>.

To install Connector/Python from the `.tar.gz` file, download the latest version and follow these steps:

```
shell> gunzip mysql-connector-python-1.0.6b1.tar.gz
shell> tar xf mysql-connector-python-1.0.6b1.tar
shell> cd mysql-connector-python-1.0.6b1
shell> sudo python setup.py install
```

On UNIX-like systems, Connector/Python gets installed in the default location `/prefix/lib/pythonX.Y/site-packages/`, where `prefix` is the location where Python was installed and `X.Y` is the version of Python. See [How installation works](#) in the Python manual.

If you are not sure where Connector/Python was installed, do the following to retrieve the location:

```
shell> python
>>> from distutils.sysconfig import get_python_lib

>>> print get_python_lib()          # Python v2.x
/Library/Python/2.7/site-packages

>>> print(get_python_lib())        # Python v3.x
/Library/Frameworks/Python.framework/Versions/3.1/lib/python3.1/site-packages
```



#### Note

The above example shows the default installation location on Mac OS X 10.7.

### 21.6.3.2. Installing Connector/Python Source Distribution on Microsoft Windows

On Microsoft Windows systems, you can download Connector/Python as a [zip](http://dev.mysql.com/downloads/connector/python/) archive from <http://dev.mysql.com/downloads/connector/python/>.

Make sure that the Python executable is available in the Windows `%PATH%` setting. For more information about installation and configuration of Python on Windows, see the section [Using Python on Windows](#) in the Python documentation.

To install Connector/Python from the `.zip` file, download the latest version and follow these steps:

1. Unpack the downloaded `zip` archive into a directory of your choice. For example, into the folder `C:\mysql-connector\`. Use the appropriate unzip command for your system, for example, `unzip`, `pkunzip`, and so on.
2. Start a console window (or a DOS window) and change to the folder where you unpacked the Connector/Python `zip` archive.

```
shell> cd C:\mysql-connector\
```

3. Once inside the Connector/Python folder, do the following:

```
shell> python setup.py install
```

On Windows, Connector/Python gets installed in the default location `C:\PythonX.Y\Lib\site-packages\` where `X.Y` is the Python version you used to install the connector.

If you are not sure where Connector/Python ended up, do the following to retrieve the location where packages get installed:

```
shell> python
>>> from distutils.sysconfig import get_python_lib
>>> print get_python_lib()          # Python v2.x
>>> print(get_python_lib())        # Python v3.x
```

### 21.6.3.3. Verifying Your Connector/Python Installation

To test that your Connector/Python installation is working and is able to connect to a MySQL database server, you can run a very simple program where you substitute the login credentials and host information of the MySQL server. See [Section 21.6.4.1, “Connecting to MySQL Using Connector/Python”](#) for an example.

## 21.6.4. Connector/Python Coding Examples

These coding examples illustrate how to develop Python applications and scripts which connect to a MySQL Server using MySQL Connector/Python.

### 21.6.4.1. Connecting to MySQL Using Connector/Python

The `connect()` constructor is used for creating a connection to the MySQL server and returns a `MySQLConnection` object.

The following example shows how to connect to the MySQL server:

```
import mysql.connector

cnx = mysql.connector.connect(user='scott', password='tiger',
                             host='127.0.0.1',
                             database='employees')

cnx.close()
```

See [Section 21.6.6, “Connector/Python Connection Arguments”](#) for all possible connection arguments.

It is also possible to create connection objects using the `connection.MySQLConnection()` class. Both methods, using the `connect()` constructor, or the class directly, are valid and functionally equal, but using `connector()` is preferred and will be used in most examples in this manual.

To handle connection errors, use the `try` statement and catch all errors using the `errors.Error` exception:

```
import mysql.connector
from mysql.connector import errorcode

try:
    cnx = mysql.connector.connect(user='scott',
                                 database='testt')
except mysql.connector.Error as err:
    if err.errno == errorcode.ER_ACCESS_DENIED_ERROR:
        print("Something is wrong your username or password")
    elif err.errno == errorcode.ER_BAD_DB_ERROR:
        print("Database does not exists")
    else:
        print(err)
else:
    cnx.close()
```

If you have lots of connection arguments, it's best to keep them in a dictionary and use the `**`-operator. Here is an example:

```
import mysql.connector

config = {
    'user': 'scott',
    'password': 'tiger',
    'host': '127.0.0.1',
    'database': 'employees',
    'raise_on_warnings': True,
}

cnx = mysql.connector.connect(**config)

cnx.close()
```

### 21.6.4.2. Creating Tables Using Connector/Python

All **DDL** (Data Definition Language) statements are executed using a handle structure known as a cursor. The following examples show how to create the tables of the `employees` database. You will need them for the other examples.

In a MySQL server, tables are very long-lived objects, and are often accessed by multiple applications written in different languages. You might typically work with tables that are already set up, rather than creating them within your own application. Avoid setting up and dropping tables over and over

again, as that is an expensive operation. The exception is [temporary tables](#), which can be created and dropped quickly within an application.

```

from __future__ import print_function

import mysql.connector
from mysql.connector import errorcode

DB_NAME = 'employees'

TABLES = {}
TABLES['employees'] = (
    "CREATE TABLE `employees` ("
    "  `emp_no` int(11) NOT NULL AUTO_INCREMENT,"
    "  `birth_date` date NOT NULL,"
    "  `first_name` varchar(14) NOT NULL,"
    "  `last_name` varchar(16) NOT NULL,"
    "  `gender` enum('M','F') NOT NULL,"
    "  `hire_date` date NOT NULL,"
    "  PRIMARY KEY (`emp_no`)"
    ") ENGINE=InnoDB")

TABLES['departments'] = (
    "CREATE TABLE `departments` ("
    "  `dept_no` char(4) NOT NULL,"
    "  `dept_name` varchar(40) NOT NULL,"
    "  PRIMARY KEY (`dept_no`), UNIQUE KEY `dept_name` (`dept_name`)"
    ") ENGINE=InnoDB")

TABLES['salaries'] = (
    "CREATE TABLE `salaries` ("
    "  `emp_no` int(11) NOT NULL,"
    "  `salary` int(11) NOT NULL,"
    "  `from_date` date NOT NULL,"
    "  `to_date` date NOT NULL,"
    "  PRIMARY KEY (`emp_no`,`from_date`), KEY `emp_no` (`emp_no`),"
    "  CONSTRAINT `salaries_ibfk_1` FOREIGN KEY (`emp_no`) "
    "    REFERENCES `employees` (`emp_no`) ON DELETE CASCADE"
    ") ENGINE=InnoDB")

TABLES['dept_emp'] = (
    "CREATE TABLE `dept_emp` ("
    "  `emp_no` int(11) NOT NULL,"
    "  `dept_no` char(4) NOT NULL,"
    "  `from_date` date NOT NULL,"
    "  `to_date` date NOT NULL,"
    "  PRIMARY KEY (`emp_no`,`dept_no`), KEY `emp_no` (`emp_no`),"
    "  KEY `dept_no` (`dept_no`),"
    "  CONSTRAINT `dept_emp_ibfk_1` FOREIGN KEY (`emp_no`) "
    "    REFERENCES `employees` (`emp_no`) ON DELETE CASCADE,"
    "  CONSTRAINT `dept_emp_ibfk_2` FOREIGN KEY (`dept_no`) "
    "    REFERENCES `departments` (`dept_no`) ON DELETE CASCADE"
    ") ENGINE=InnoDB")

TABLES['dept_manager'] = (
    "CREATE TABLE `dept_manager` ("
    "  `dept_no` char(4) NOT NULL,"
    "  `emp_no` int(11) NOT NULL,"
    "  `from_date` date NOT NULL,"
    "  `to_date` date NOT NULL,"
    "  PRIMARY KEY (`emp_no`,`dept_no`),"
    "  KEY `emp_no` (`emp_no`),"
    "  KEY `dept_no` (`dept_no`),"
    "  CONSTRAINT `dept_manager_ibfk_1` FOREIGN KEY (`emp_no`) "
    "    REFERENCES `employees` (`emp_no`) ON DELETE CASCADE,"
    "  CONSTRAINT `dept_manager_ibfk_2` FOREIGN KEY (`dept_no`) "
    "    REFERENCES `departments` (`dept_no`) ON DELETE CASCADE"
    ") ENGINE=InnoDB")

TABLES['titles'] = (
    "CREATE TABLE `titles` ("
    "  `emp_no` int(11) NOT NULL,"

```

```
" `title` varchar(50) NOT NULL,"
" `from_date` date NOT NULL,"
" `to_date` date DEFAULT NULL,"
" PRIMARY KEY (`emp_no`,`title`,`from_date`), KEY `emp_no` (`emp_no`),"
" CONSTRAINT `titles_ibfk_1` FOREIGN KEY (`emp_no`)
" REFERENCES `employees` (`emp_no`) ON DELETE CASCADE"
") ENGINE=InnoDB")
```

The above code shows how we are storing the `CREATE` statements in a Python dictionary called `TABLES`. We also define the database in a global variable called `DB_NAME`, which allows you to easily use a different schema.

```
cnx = mysql.connector.connect(user='scott')
cursor = cnx.cursor()
```

A single MySQL server can contain multiple `databases`. Typically, you specify the database to switch to when connecting to the MySQL server. This example does not connect to the database upon connection, so that it can make sure the database exists, and create it if not.

```
def create_database(cursor):
    try:
        cursor.execute(
            "CREATE DATABASE {} DEFAULT CHARACTER SET 'utf8'".format(DB_NAME))
    except mysql.connector.Error as err:
        print("Failed creating database: {}".format(err))
        exit(1)

try:
    cnx.database = DB_NAME
except mysql.connector.Error as err:
    if err.errno == errorcode.ER_BAD_DB_ERROR:
        create_database(cursor)
        cnx.database = DB_NAME
    else:
        print(err)
        exit(1)
```

We first try to change to a particular database using the `database` property of the connection object `cnx`. If there is an error, we examine the error number to check if the database does not exist. If so, we call the `create_database` function to create it for us.

On any other error, the application exits and displays the error message.

```
for name, ddl in TABLES.iteritems():
    try:
        print("Creating table {}: ".format(name), end='')
        cursor.execute(ddl)
    except mysql.connector.Error as err:
        if err.errno == errorcode.ER_TABLE_EXISTS_ERROR:
            print("already exists.")
        else:
            print(err.errmsg)
    else:
        print("OK")

cursor.close()
cnx.close()
```

After we successfully created or changed to the target database, we create the tables by iterating over the items of the `TABLES` dictionary.

We handle the error when the table already exists by simply notifying the user that it was already there. Other errors are printed, but we simply continue creating tables. (We show how to handle the “table already exists” condition for illustration purposes. In a real application, we would typically avoid the error condition entirely by using the `IF NOT EXISTS` clause of the `CREATE TABLE` statement.)

The output would be something like this:

```

Creating table employees: already exists.
Creating table salaries: already exists.
Creating table titles: OK
Creating table departments: already exists.
Creating table dept_manager: already exists.
Creating table dept_emp: already exists.

```

To populate the employees tables, use the dump files of the [Employee Sample Database](#). Note that you only need the data dump files that you will find in an archive named like [employees\\_db-dump-files-1.0.5.tar.bz2](#). After downloading the dump files, do the following from the command line, adding connection options to the `mysql` commands if necessary:

```

shell> tar xzf employees_db-dump-files-1.0.5.tar.bz2
shell> cd employees_db
shell> mysql employees < load_employees.dump
shell> mysql employees < load_titles.dump
shell> mysql employees < load_departments.dump
shell> mysql employees < load_salaries.dump
shell> mysql employees < load_dept_emp.dump
shell> mysql employees < load_dept_manager.dump

```

### 21.6.4.3. Inserting Data Using Connector/Python

Inserting or updating data is also done using the handler structure known as a cursor. When you use a transactional storage engine such as [InnoDB](#) (which is the default in MySQL 5.5 and later), you must [commit](#) the data after a sequence of [INSERT](#), [DELETE](#), and [UPDATE](#) statements.

In this example we show how to insert new data. The second [INSERT](#) depends on the value of the newly created [primary key](#) of the first. We are also demonstrating how to use extended formats. The task is to add a new employee starting to work tomorrow with a salary set to 50000.



#### Note

The following example uses tables created in the example [Section 21.6.4.2, “Creating Tables Using Connector/Python”](#). The `AUTO_INCREMENT` column option for the primary key of the `employees` table is important to ensure reliable, easily searchable data.

```

from __future__ import print_function
from datetime import date, datetime, timedelta
import mysql.connector

cnx = mysql.connector.connect(user='scott', database='employees')
cursor = cnx.cursor()

tomorrow = datetime.now().date() + timedelta(days=1)

add_employee = ("INSERT INTO employees "
               "(first_name, last_name, hire_date, gender, birth_date) "
               "VALUES (%s, %s, %s, %s, %s)")
add_salary = ("INSERT INTO salaries "
             "(emp_no, salary, from_date, to_date) "
             "VALUES (%(emp_no)s, %(salary)s, %(from_date)s, %(to_date)s)")

data_employee = ('Geert', 'Vanderkelen', tomorrow, 'M', date(1977, 6, 14))

# Insert new employee
cursor.execute(add_employee, data_employee)
emp_no = cursor.lastrowid

# Insert salary information
data_salary = {
    'emp_no': emp_no,
    'salary': 50000,
    'from_date': tomorrow,
    'to_date': date(9999, 1, 1),
}

```

```

cursor.execute(add_salary, data_salary)

# Make sure data is committed to the database
cnx.commit()

cursor.close()
cnx.close()

```

We first open a connection to the MySQL server and store the [connection object](#) in the variable `cnx`. We then create a new cursor, by default a [MySQLCursor](#) object, using the connection's `cursor()` method.

We could calculate tomorrow by calling a database function, but for clarity we do it in Python using the [datetime](#) module.

Both `INSERT` statements are stored in the variables called `add_employee` and `add_salary`. Note that the second `INSERT` statement uses extended Python format codes.

The information of the new employee is stored in the tuple `data_employee`. The query to insert the new employee is executed and we retrieve the newly inserted value for the column `emp_no` using the `lastrowid` property of the cursor object.

Next, we insert the new salary for the new employee. We are using the `emp_no` variable in the directory holding the data. This directory is passed to the `execute()` method of the cursor object.

Since by default Connector/Python turns [autocommit](#) off, and MySQL 5.5 and later uses transactional [InnoDB](#) tables by default, it is necessary to commit your changes using the connection's `commit()` method. You could also [roll back](#) using the `rollback()` method.

#### 21.6.4.4. Querying Data Using Connector/Python

The following example shows how to [query](#) data using a cursor created using the connection's `cursor()` method. The data returned is formatted and printed on the console.

The task is to select all employees hired in the year 1999 and print their names with their hire date to the console.

```

import datetime
import mysql.connector

cnx = mysql.connector.connect(user='scott', database='employees')
cursor = cnx.cursor()

query = ("SELECT first_name, last_name, hire_date FROM employees "
        "WHERE hire_date BETWEEN %s AND %s")

hire_start = datetime.date(1999, 1, 1)
hire_end = datetime.date(1999, 12, 31)

cursor.execute(query, (hire_start, hire_end))

for (first_name, last_name, hire_date) in cursor:
    print("{} , {} was hired on {:%d %b %Y}".format(
        last_name, first_name, hire_date))

cursor.close()
cnx.close()

```

We first open a connection to the MySQL server and store the [connection object](#) in the variable `cnx`. We then create a new cursor, by default a [MySQLCursor](#) object, using the connection's `cursor()` method.

In the preceding example, we store the `SELECT` statement in the variable `query`. Note that we are using unquoted `%s`-markers where dates should have been. Connector/Python converts `hire_start`

and `hire_end` from Python types to a data type that MySQL understands and adds the required quotes. In this case, it replaces the first `%s` with `'1999-01-01'`, and the second with `'1999-12-31'`.

We then execute the operation stored in the `query` variable using the `execute()` method. The data used to replace the `%s`-markers in the query is passed as a tuple: `(hire_start, hire_end)`.

After executing the query, the MySQL server is ready to send the data. The result set could be zero rows, one row, or 100 million rows. Depending on the expected volume, you can use different techniques to process this result set. In this example, we use the `cursor` object as an iterator. The first column in the row will be stored in the variable `first_name`, the second in `last_name`, and the third in `hire_date`.

We print the result, formatting the output using Python's built-in `format()` function. Note that `hire_date` was converted automatically by Connector/Python to a Python `datetime.date` object. This means that we can easily format the date in a more human-readable form.

The output should be something like this:

```
..
Wilharm, LiMin was hired on 16 Dec 1999
Wielonsky, Lalit was hired on 16 Dec 1999
Kamble, Dannz was hired on 18 Dec 1999
DuBourdieu, Zhongwei was hired on 19 Dec 1999
Fujisawa, Rosita was hired on 20 Dec 1999
..
```

## 21.6.5. Connector/Python Tutorials

These tutorials illustrate how to develop Python applications and scripts that connect to a MySQL database server using MySQL Connector/Python.

### 21.6.5.1. Tutorial: Raise employee's salary using a buffering cursor

The following example script will give a long-overdue raise effective tomorrow to all employees who joined in the year 2000 and are still with the company.

We are using buffered cursors to iterate through the selected employees. This way we do not have to fetch the rows in a new variables, but can instead use the cursor as an iterator.

Note that the script is an example; there are other ways of doing this simple task.

```
from __future__ import print_function

from decimal import Decimal
from datetime import datetime, date, timedelta

import mysql.connector

# Connect with the MySQL Server
cnx = mysql.connector.connect(user='scott', database='employees')

# Get two buffered cursors
curA = cnx.cursor(buffered=True)
curB = cnx.cursor(buffered=True)

# Query to get employees who joined in a period defined by two dates
query = (
    "SELECT s.emp_no, salary, from_date, to_date FROM employees AS e "
    "LEFT JOIN salaries AS s USING (emp_no) "
    "WHERE to_date = DATE('9999-01-01') "
    "AND e.hire_date BETWEEN DATE(%s) AND DATE(%s)"
)

# UPDATE and INSERT statements for the old and new salary
update_old_salary = (
    "UPDATE salaries SET to_date = %s "
    "WHERE emp_no = %s AND from_date = %s"
)
insert_new_salary = (
```



```

"INSERT INTO salaries (emp_no, from_date, to_date, salary) "
"VALUES (%s, %s, %s, %s)"

# Select the employees getting a raise
curA.execute(query, (date(2000, 1, 1), date(2001, 1, 1)))

# Iterate through the result of curA
for (emp_no, salary, from_date, to_date) in curA:

    # Update the old and insert the new salary
    new_salary = int(round(salary * Decimal('1.15')))
    curB.execute(update_old_salary, (tomorrow, emp_no, from_date))
    curB.execute(insert_new_salary,
                  (emp_no, tomorrow, date(9999, 1, 1), new_salary))

# Commit the changes
cnx.commit()

cnx.close()

```

## 21.6.6. Connector/Python Connection Arguments

The following lists the arguments that can be used to initiate a connection with the MySQL server using either:

- Function `mysql.connector.connect()`
- Class `mysql.connector.MySQLConnection()`

**Table 21.31. Connection Arguments for Connector/Python**

Argument Name	Default	Description
<code>user</code> ( <code>username*</code> )		The username used to authenticate with the MySQL Server.
<code>password</code> ( <code>passwd*</code> )		The password to authenticate the user with the MySQL Server.
<code>database</code> ( <code>db*</code> )		Database name to use when connecting with the MySQL Server.
<code>host</code>	127.0.0.1	Hostname or IP address of the MySQL Server.
<code>port</code>	3306	TCP/IP port of the MySQL Server. Must be an integer.
<code>unix_socket</code>		The location of the Unix socket file.
<code>use_unicode</code>	True	Whether to use Unicode or not.
<code>charset</code>	utf8	Which MySQL character set to use.
<code>collation</code>	utf8_general_ci	Which MySQL collation to use.
<code>autocommit</code>	False	Whether to <a href="#">autocommit</a> transactions.
<code>time_zone</code>		Set the <code>time_zone</code> session variable at connection.
<code>sql_mode</code>		Set the <code>sql_mode</code> session variable at connection.
<code>get_warnings</code>	False	Whether to fetch warnings.
<code>raise_on_warnings</code>	False	Whether to raise an exception on warnings.
<code>connection_timeout</code> ( <code>connect_timeout*</code> )		Timeout for the TCP and Unix socket connections.
<code>client_flags</code>		MySQL client flags.
<code>buffered</code>	False	Whether cursor object fetches the result immediately after executing query.
<code>raw</code>	False	Whether MySQL results are returned as-is, rather than converted to Python types.

Argument Name	Default	Description
<code>ssl_ca</code>		File containing the SSL certificate authority.
<code>ssl_cert</code>		File containing the SSL certificate file.
<code>ssl_key</code>		File containing the SSL key.
<code>ssl_verify_cert</code>	False	When set to <code>True</code> , checks the server certificate against the certificate file specified by the <code>ssl_ca</code> option. Any mismatch causes a <code>ValueError</code> exception.
<code>force_ipv6</code>	False	When set to <code>True</code> , uses IPv6 when an address resolves to both IPv4 and IPv6. By default, IPv4 is used in such cases.
<code>dsn</code>		Not supported (raises <code>NotSupportedError</code> when used).

\* Synonymous argument name, available only for compatibility with other Python MySQL drivers. Oracle recommends not to use these alternative names.

## MySQL Authentication

Authentication with MySQL will use `username` and `password`. Note that MySQL Connector/Python does not support the old, insecure password protocols of MySQL versions prior to 4.1.

When the `database` parameter is given, the current database is set to the given value. To later change the database, execute the MySQL `USE` command or set the `database` property of the `MySQLConnection` instance.

By default, Connector/Python tries to connect to a MySQL server running on `localhost` using TCP/IP. The `host` argument defaults to IP address 127.0.0.1 and `port` to 3306. Unix sockets are supported by setting `unix_socket`. Named pipes on the Windows platform are not supported.

## Character Encoding

Strings coming from MySQL are by default returned as Python Unicode literals. To change this behavior, set `use_unicode` to `False`. You can change the character setting for the client connection through the `charset` argument. To change the character set after connecting to MySQL, set the `charset` property of the `MySQLConnection` instance. This technique is preferred over using the MySQL `SET NAMES` statement directly. Similar to the `charset` property, you can set the `collation` for the current MySQL session.

## Transactions

Transactions are not automatically committed; call the `commit()` method of the `MySQLConnection` instance within your application after doing a set of related insert, update, and delete operations. For data consistency and high throughput for write operations, it is best to leave the `autocommit` configuration option turned off when using `InnoDB` or other transactional tables.

## Time Zones

The time zone can be set per connection using the `time_zone` argument. This is useful if the MySQL server is set, for example, to UTC and `TIMESTAMP` values should be returned by MySQL converted to the `PST` time zone.

## SQL Modes

MySQL supports so called SQL Modes, which will change the behavior of the server globally or per connection. For example, to have warnings raised as errors, set `sql_mode` to `TRADITIONAL`. For more information, see [Section 5.1.7, "Server SQL Modes"](#).

## Troubleshooting and Error Handling

Warnings generated by queries are fetched automatically when `get_warnings` is set to `True`. You can also immediately raise an exception by setting `raise_on_warnings` to `True`. Consider using the MySQL `sql_mode` setting for turning warnings into errors.

To set a timeout value for connections, use `connection_timeout`.

## Enabling and Disabling Features through Client Flags

MySQL uses `client flags` to enable or disable features. Using the `client_flags` argument, you have control of what is set. To find out what flags are available, use the following:

```
from mysql.connector.constants import ClientFlag
print '\n'.join(ClientFlag.get_full_info())
```

If `client_flags` is not specified (that is, it is zero), defaults are used for MySQL v4.1 and later. If you specify an integer greater than 0, make sure all flags are set. A better way to set and unset flags is to use a list. For example, to set `FOUND_ROWS`, but disable the default `LONG_FLAG`:

```
flags = [ClientFlag.FOUND_ROWS, -ClientFlag.LONG_FLAG]
mysql.connector.connect(client_flags=flags)
```

## Buffering for Result Sets

By default, MySQL Connector/Python does not buffer or pre-fetch results. This means that after a query is executed, your program is responsible of fetching the data. This avoids using excessive memory when queries return large result sets. If you know that the result set is small enough to handle all at once, fetching the results immediately by setting `buffered` to `True`. It is also possible to set this per cursor (see cursor manual).

## Type Conversions

MySQL types will be converted automatically to Python types. For example, a `DATETIME` column becomes a `datetime.datetime` object. When conversion should be done differently, for example to get better performance, set `raw` to `True`.

## Connecting through SSL

Using SSL connections is possible when your [Python installation supports SSL](#), that is, when it is compiled against the OpenSSL libraries. When you provide the arguments `ssl_ca`, `ssl_key` and `ssl_cert`, the connection switches to SSL, and the `client_flags` option includes the `ClientFlag.SSL` value automatically. You can use this in combination with the `compressed` argument set to `True`.

```
# Note (Example is valid for Python v2 and v3)
from __future__ import print_function

import sys

#sys.path.insert(0, 'python{0}/'.format(sys.version_info[0]))

import mysql.connector
from mysql.connector.constants import ClientFlag

config = {
    'user': 'ssluser',
    'password': 'asecret',
    'host': '127.0.0.1',
    'client_flags': [ClientFlag.SSL],
    'ssl_ca': '/opt/mysql/ssl/ca-cert.pem',
    'ssl_cert': '/opt/mysql/ssl/client-cert.pem',
    'ssl_key': '/opt/mysql/ssl/client-key.pem',
```

```

}

cnx = mysql.connector.connect(**config)
cur = cnx.cursor(buffered=True)
cur.execute("SHOW STATUS LIKE 'Ssl_cipher'")
print(cur.fetchone())
cur.close()
cnx.close()

```

## Compatibility with Other Connection Interfaces

`passwd`, `db` and `connect_timeout` are valid for compatibility with other MySQL interfaces and are respectively the same as `password`, `database` and `connection_timeout`. The latter take precedence. Data source name syntax or `dsn` is not used; if specified, it raises a `NotSupportedError` exception.

## 21.6.7. Connector/Python API Reference

This section contains the public API reference of Connector/Python. Although valid for both Python 2 and Python 3, examples should be considered working for Python 2.7, and Python 3.1 and greater.

The following overview shows the `mysql.connector` package with its modules. Currently, only the most useful modules, classes and functions for end users are documented.

```

mysql.connector
  errorcode
  errors
  connection
  constants
  conversion
  cursor
  dbapi
  locales
    eng
      client_error
  protocol
  utils

```

### 21.6.7.1. Errors and Exceptions

The `mysql.connector.errors` module defines exception classes for errors and warnings raised by MySQL Connector/Python. Most classes defined in this module are available when you import `mysql.connector`.

The exception classes defined in this module follow mostly the Python Database Specification v2.0 (PEP-249). For some MySQL client or server errors it is not always clear which exception to raise. It is good to discuss whether an error should be reclassified by opening a bug report.

MySQL Server errors are mapped with Python exception based on their SQLState (see [Section C.3, “Server Error Codes and Messages”](#)). The following list shows the `SQLState` classes and the exception Connector/Python will raise. It is, however, possible to redefine which exception is raised for each server error. Note that the default exception is `DatabaseError`.

- `02`: `DataError`
- `07`: `DatabaseError`
- `08`: `OperationalError`
- `0A`: `NotSupportedError`
- `21`: `DataError`
- `22`: `DataError`

- [23](#): IntegrityError
- [24](#): ProgrammingError
- [25](#): ProgrammingError
- [26](#): ProgrammingError
- [27](#): ProgrammingError
- [28](#): ProgrammingError
- [2A](#): ProgrammingError
- [2B](#): DatabaseError
- [2C](#): ProgrammingError
- [2D](#): DatabaseError
- [2E](#): DatabaseError
- [33](#): DatabaseError
- [34](#): ProgrammingError
- [35](#): ProgrammingError
- [37](#): ProgrammingError
- [3C](#): ProgrammingError
- [3D](#): ProgrammingError
- [3F](#): ProgrammingError
- [40](#): InternalError
- [42](#): ProgrammingError
- [44](#): InternalError
- [HZ](#): OperationalError
- [XA](#): IntegrityError
- [0K](#): OperationalError
- [HY](#): DatabaseError

#### 21.6.7.1.1. Module `errorcode`

This module contains both MySQL server and client error codes defined as module attributes with the error number as value. Using error codes instead of error numbers could make reading the source code a bit easier.

```
>>> from mysql.connector import errorcode
>>> errorcode.ER_BAD_TABLE_ERROR
1051
```

See [Section C.3](#), “Server Error Codes and Messages” and [Section C.4](#), “Client Error Codes and Messages”.

### 21.6.7.1.2. Exception `errors.Error`

This exception is the base class for all other exceptions in the `errors` module. It can be used to catch all errors in a single `except` statement.

The following example shows how we could catch syntax errors:

```
import mysql.connector
try:
    cnx = mysql.connector.connect(user='scott', database='employees')
    cursor = cnx.cursor()
    cursor.execute("SELECT * FORM employees") # Syntax error in query
    cnx.close()
except mysql.connector.Error as err:
    print("Something went wrong: {}".format(err))
```

Initializing the exception supports a few optional arguments, namely `msg`, `errno`, `values` and `sqlstate`. All of them are optional and default to `None`. `errors.Error` is internally used by Connector/Python to raise MySQL client and server errors and should not be used by your application to raise exceptions.

The following examples show the result when using no or a combination of the arguments:

```
>>> from mysql.connector.errors import Error
>>> str(Error())
'Unknown error'

>>> str(Error("Oops! There was an error. "))
'Oops! There was an error.'

>>> str(Error(errno=2006))
'2006: MySQL server has gone away'

>>> str(Error(errno=2002, values=('/tmp/mysql.sock', 2)))
'2002: Can't connect to local MySQL server through socket '/tmp/mysql.sock' (2)'

>>> str(Error(errno=1146, sqlstate='42S02', msg="Table 'test.spam' doesn't exist"))
'1146 (42S02): Table 'test.spam' doesn't exist'
```

The example which uses error number 1146 is used when Connector/Python receives an error packet from the MySQL Server. The information is parsed and passed to the `Error` exception as shown.

Each exception subclassing from `Error` can be initialized using the above mentioned arguments. Additionally, each instance has the attributes `errno`, `msg` and `sqlstate` which can be used in your code.

The following example shows how to handle errors when dropping a table which does not exist (when you do not want to use the `IF EXISTS` clause):

```
import mysql.connector
from mysql.connector import errorcode

cnx = mysql.connector.connect(user='scott', database='test')
try:
    cur.execute("DROP TABLE spam")
except mysql.connector.Error as err:
    if err.errno == errorcode.ER_BAD_TABLE_ERROR:
        print("Creating table spam")
    else:
        raise
```

`errors.Error` is a subclass of the Python `StandardError`.

### 21.6.7.1.3. Exception `errors.Warning`

This exception is used for reporting important warnings, however, Connector/Python does not use it. It is included to be compliant with the Python Database Specification v2.0 (PEP-249).

Consider using either more strict [Server SQL Modes](#) or the [raise\\_on\\_warnings](#) connection argument to make Connector/Python raise errors when your queries produce warnings.

`errors.Warning` is a subclass of the Python `StandardError`.

#### 21.6.7.1.4. Exception `errors.InterfaceError`

This exception is raised for errors originating from Connector/Python itself, not related to the MySQL server.

`errors.InterfaceError` is a subclass of `errors.Error`.

#### 21.6.7.1.5. Exception `errors.DatabaseError`

This exception is the default for any MySQL error which does not fit the other exceptions.

`errors.DatabaseError` is a subclass of `errors.Error`.

#### 21.6.7.1.6. Exception `errors.InternalError`

This exception is raised when the MySQL server encounters an internal error, for example, when a deadlock occurred.

`errors.InternalError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.7. Exception `errors.OperationalError`

This exception is raised for errors which are related to MySQL's operations. For example, to many connections, a hostname could not be resolved, bad handshake, server is shutting down, communication errors, and so on.

`errors.OperationalError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.8. Exception `errors.ProgrammingError`

This exception is raised on programming errors, for example when you have a syntax error in your SQL or a table was not found.

The following example shows how to handle syntax errors:

```
try:
    cursor.execute("CREATE DESK t1 (id int, PRIMARY KEY (id))")
except mysql.connector.ProgrammingError as err:
    if err.errno == errorcode.ER_SYNTAX_ERROR:
        print("Check your syntax!")
    else:
        print("Error: {}".format(err))
```

`errors.ProgrammingError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.9. Exception `errors.IntegrityError`

This exception is raised when the relational integrity of the data is affected. For example, a duplicate key was inserted or a foreign key constraint would fail.

The following example shows a duplicate key error raised as `IntegrityError`:

```
cursor.execute("CREATE TABLE t1 (id int, PRIMARY KEY (id))")
try:
    cursor.execute("INSERT INTO t1 (id) VALUES (1)")
    cursor.execute("INSERT INTO t1 (id) VALUES (1)")
except mysql.connector.IntegrityError as err:
```

```
print("Error: {}".format(err))
```

`errors.IntegrityError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.10. Exception `errors.DataError`

This exception is raised when there were problems with the data. Examples are a column set to NULL when it can not, out of range values for a column, division by zero, column count does not match value count, and so on.

`errors.DataError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.11. Exception `errors.NotSupportedError`

This exception is raised in case some feature was used but not supported by the version of MySQL which returned the error. It is also raised when using functions or statements which are not supported by stored routines.

`errors.NotSupportedError` is a subclass of `errors.DatabaseError`.

#### 21.6.7.1.12. Function `errors.custom_error_exception(error=None, exception=None)`

This function defines custom exceptions for MySQL server errors and returns current customizations.

If `error` is a MySQL Server error number, then you have to pass also the `exception` class. The `error` argument can also be a dictionary in which case the key is the server error number, and value the class of the exception to be raised.

To reset the customizations, simply supply an empty dictionary.

```
import mysql.connector
from mysql.connector import errorcode

# Server error 1028 should raise a DatabaseError
mysql.connector.custom_error_exception(1028, mysql.connector.DatabaseError)

# Or using a dictionary:
mysql.connector.custom_error_exception({
    1028: mysql.connector.DatabaseError,
    1029: mysql.connector.OperationalError,
})

# To reset, pass an empty dictionary:
mysql.connector.custom_error_exception({})
```

### 21.6.7.2. Class `connection.MySQLConnection`

The `MySQLConnection` class is used to open and manage a connection to a MySQL server. It also used to send commands and SQL queries and read result.

#### 21.6.7.2.1. Constructor `connection.MySQLConnection(**kwargs)`

The `MySQLConnection` constructor initializes the attributes and when at least one argument is passed, it tries to connect with the MySQL server.

For a complete list of arguments, see [Section 21.6.6, “Connector/Python Connection Arguments”](#).

#### 21.6.7.2.2. Method `MySQLConnection.close()`

See [disconnect\(\)](#).

Returns a tuple.



### 21.6.7.2.3. Method `MySQLConnection.config(**kwargs)`

Allows to configure a `MySQLConnection` instance after it was instantiated. See [Section 21.6.6, “Connector/Python Connection Arguments”](#) for a complete list of possible arguments.

You could use the `config()` method to change, for example, the username and call `reconnect()`.

```
cnx = MySQLConnection(user='joe', database='test')
# Connected as 'joe'
cnx.config(user='jane')
cnx.reconnect()
# Now connected as 'jane'
```

### 21.6.7.2.4. Method `MySQLConnection.connect(**kwargs)`

This method sets up the connection to the MySQL server. If no arguments are given, it uses the already configured or default values. See [Section 21.6.6, “Connector/Python Connection Arguments”](#) for a complete list of possible arguments.

### 21.6.7.2.5. Method `MySQLConnection.commit()`

This method sends the `COMMIT` command to the MySQL server, committing the current transaction. Since by default, Connector/Python does not auto commit, it is important to call this method after every transaction which updates data for tables using transactional storage engines.

See the `rollback()` method for rolling back transactions.

```
>>> cursor.execute("INSERT INTO employees (first_name) VALUES (%s)", ('Jane'))
>>> cnx.commit()
```

### 21.6.7.2.6. Method `MySQLConnection.cursor(buffered=None, raw=None, cursor_class=None)`

This method returns a `MySQLCursor()` object, or a subclass of it depending the passed arguments.

When `buffered` is `True`, the cursor will fetch all rows after the operation is executed. This is useful when queries return small result sets. Setting `raw` will skip the conversion from MySQL data types to Python types when fetching rows. Raw is usually used when you want to have more performance and/or you want to do the conversion yourself.

The `cursor_class` argument can be used to pass a class to use for instantiating a new cursor. It has to be a subclass of `cursor.CursorBase`.

The returned object depends on the combination of the `buffered` and `raw` arguments.

- If not buffered and not raw: `cursor.MySQLCursor`
- If buffered and not raw: `cursor.MySQLCursorBuffered`
- If buffered and raw: `cursor.MySQLCursorBufferedRaw`
- If not buffered and raw: `cursor.MySQLCursorRaw`

Returns a `CursorBase` instance.

### 21.6.7.2.7. Method `MySQLConnection.cmd_change_user(username='', password='', database='', charset=33)`

Changes the user using `username` and `password`. It also causes the specified `database` to become the default (current) database. It is also possible to change the character set using the `charset` argument.

Returns a dictionary containing the OK packet information.

#### 21.6.7.2.8. Method `MySQLConnection.cmd_debug()`

Instructs the server to write some debug information to the log. For this to work, the connected user must have the `SUPER` privilege.

Returns a dictionary containing the OK packet information.

#### 21.6.7.2.9. Method `MySQLConnection.cmd_init_db(database)`

This method makes specified database the default (current) database. In subsequent queries, this database is the default for table references that do not include an explicit database specifier.

Returns a dictionary containing the OK packet information.

#### 21.6.7.2.10. Method `MySQLConnection.cmd_ping()`

Checks whether the connection to the server is working.

This method is not to be used directly. Use `ping()` or `is_connected()` instead.

Returns a dictionary containing the OK packet information.

#### 21.6.7.2.11. Method `MySQLConnection.cmd_process_info()`

This method raises the `NotSupportedError` exception. Instead, use the `SHOW PROCESSLIST` statement or query the tables found in the database `INFORMATION_SCHEMA`.

#### 21.6.7.2.12. Method `MySQLConnection.cmd_process_kill(mysql_pid)`

Asks the server to kill the thread specified by `mysql_pid`. Although still available, it's better to use the SQL `KILL` command.

Returns a dictionary containing the OK packet information.

The following two lines do the same:

```
>>> cnx.cmd_process_kill(123)
>>> cnx.cmd_query('KILL 123')
```

#### 21.6.7.2.13. Method `MySQLConnection.cmd_quit()`

This method sends the `QUIT` command to the MySQL server, closing the current connection. Since there is no response from the MySQL, the packet that was sent is returned.

#### 21.6.7.2.14. Method `MySQLConnection.cmd_query(statement)`

This method sends the given `statement` to the MySQL server and returns a result. If you need to send multiple statements, you have to use the `cmd_query_iter()` method.

The returned dictionary contains information depending on what kind of query was executed. If the query is a `SELECT` statement, the result contains information about columns. Other statements return a dictionary containing OK or EOF packet information.

Errors received from the MySQL server are raised as exceptions. An `InterfaceError` is raised when multiple results are found.

Returns a dictionary.

### 21.6.7.2.15. Method `MySQLConnection.cmd_query_iter(statement)`

Similar to the `cmd_query()` method, but returns a generator object to iterate through results. Use `cmd_query_iter()` when sending multiple statements, and separate the statements with semicolons.

The following example shows how to iterate through the results after sending multiple statements:

```
statement = 'SELECT 1; INSERT INTO t1 VALUES (); SELECT 2'
for result in cnx.cmd_query(statement, iterate=True):
    if 'columns' in result:
        columns = result['columns']
        rows = cnx.get_rows()
    else:
        # do something useful with INSERT result
```

Returns a generator object.

### 21.6.7.2.16. Method `MySQLConnection.cmd_refresh(options)`

This method flushes tables or caches, or resets replication server information. The connected user must have the `RELOAD` privilege.

The `options` argument should be a bitwise value using constants from the class `constants.RefreshOption`.

See [Section 21.6.7.9, “Class `constants.RefreshOption`”](#) for a list of options.

Example:

```
>>> from mysql.connector import RefreshOption
>>> refresh = RefreshOption.LOG | RefreshOption.THREADS
>>> cnx.cmd_refresh(refresh)
```

### 21.6.7.2.17. Method `MySQLConnection.cmd_shutdown()`

Asks the database server to shut down. The connected user must have the `SHUTDOWN` privilege.

Returns a dictionary containing the OK packet information.

### 21.6.7.2.18. Method `MySQLConnection.cmd_statistics()`

Returns a dictionary containing information about the MySQL server including uptime in seconds and the number of running threads, questions, reloads, and open tables.

### 21.6.7.2.19. Method `MySQLConnection.disconnect()`

This method tries to send the `QUIT` command and close the socket. It does not raise any exceptions.

`MySQLConnection.close()` is a synonymous method name and more commonly used.

### 21.6.7.2.20. Method `MySQLConnection.get_rows(count=None)`

This method retrieves all or remaining rows of a query result set, returning a tuple containing the rows as sequence and the EOF packet information. The `count` argument can be used to get a given amount of rows. If `count` is not specified or is `None`, all rows are retrieved.

The tuple returned by `get_rows()` consists of:

- A list of tuples containing the row data as byte objects, or an empty list when no rows are available.

- EOF packet information as a dictionary containing `status_flag` and `warning_count`.

An `InterfaceError` is raised when all rows have been retrieved.

The `get_rows()` method is used by `MySQLCursor` to fetch rows.

Returns a tuple.

#### 21.6.7.2.21. Method `MySQLConnection.get_row()`

This method retrieves the next row of a query result set, returning a tuple.

The tuple returned by `get_row()` consists of:

- The row as a tuple containing byte objects, or `None` when no more rows are available.
- EOF packet information as a dictionary containing `status_flag` and `warning_count`, or `None` when the row returned is not the last row.

The `get_row()` method is used by `MySQLCursor` to fetch rows.

#### 21.6.7.2.22. Method `MySQLConnection.get_server_info()`

This method returns the MySQL server information verbatim, for example `'5.5.24-log'`, or `None` when not connected.

Returns a string or `None`.

#### 21.6.7.2.23. Method `MySQLConnection.get_server_version()`

This method returns the MySQL server version as a tuple, or `None` when not connected.

Returns a tuple or `None`.

#### 21.6.7.2.24. Method `MySQLConnection.is_connected()`

Reports whether the connection to MySQL Server is available.

This method checks whether the connection to MySQL is available using the `ping()` method, but unlike `ping()`, `is_connected()` returns `True` when the connection is available, `False` otherwise.

Returns `True` or `False`.

#### 21.6.7.2.25. Method `MySQLConnection.isset_client_flag(flag)`

This method returns `True` if the client flag was set, `False` otherwise.

Returns `True` or `False`.

#### 21.6.7.2.26. Method `MySQLConnection.ping(attempts=1, delay=0)`

Check whether the connection to the MySQL server is still available.

When `reconnect` is set to `True`, one or more attempts are made to try to reconnect to the MySQL server using the `reconnect()` method. Use the `delay` argument (seconds) if you want to wait between each retry.

When the connection is not available, an `InterfaceError` is raised. Use the `is_connected()` method to check the connection without raising an error.

Raises `InterfaceError` on errors.

### 21.6.7.2.27. Method `MySQLConnection.reconnect(attempts=1, delay=0)`

Attempt to reconnect with the MySQL server.

The argument `attempts` specifies the number of times a reconnect is tried. The `delay` argument is the number of seconds to wait between each retry.

You might set the number of attempts higher and use a longer delay when you expect the MySQL server to be down for maintenance, or when you expect the network to be temporarily unavailable.

### 21.6.7.2.28. Method `MySQLConnection.rollback()`

This method sends the `ROLLBACK` command to the MySQL server, undoing all data changes from the current `transaction`. Since by default, Connector/Python does not auto commit, it is possible to cancel transactions when using transactional storage engines such as `InnoDB`.

See the `commit()` method for `committing` transactions.

```
>>> cursor.execute("INSERT INTO employees (first_name) VALUES (%s)", ('Jane'))
>>> cnx.rollback()
```

### 21.6.7.2.29. Method `MySQLConnection.set_charset_collation(charset=None, collation=None)`

This method sets the character set and collation to be used for the current connection. The `charset` argument can be either the name of a character set, or the numerical equivalent as defined in `constants.CharacterSet`.

When `collation` is `None`, the default will be looked up and used.

The `charset` argument then be either:

In the following example, we set the character set to `latin1` and the collation will be set to the default `latin1_swedish_ci`:

```
>>> cnx = mysql.connector.connect(user='scott')
>>> cnx.set_charset('latin1')
```

Specify a specific collation as follows:

```
>>> cnx = mysql.connector.connect(user='scott')
>>> cnx.set_charset('latin1', 'latin1_general_ci')
```

### 21.6.7.2.30. Method `MySQLConnection.set_client_flags(flags)`

This method sets the client flags which are used when connecting with the MySQL server and returns the new value. The `flags` argument can be either an integer or a sequence of valid client flag values (see [Section 21.6.7.5, "Class constants.ClientFlag"](#)).

If `flags` is a sequence, each item in the sequence will set the flag when the value is positive or unset it when negative. For example, to unset `LONG_FLAG` and set the `FOUND_ROWS` flags:

```
>>> from mysql.connector.constants import ClientFlag
>>> cnx.set_client_flags([ClientFlag.FOUND_ROWS, -ClientFlag.LONG_FLAG])
>>> cnx.reconnect()
```

Note that client flags are only set or used when connecting with the MySQL server. It is therefore necessary to reconnect after making changes.

Returns an integer.

### 21.6.7.2.31. Property `MySQLConnection.autocommit`

This property is used to toggle the auto commit feature of MySQL and retrieve the current state. When the value evaluates to True, auto commit will be turned, otherwise it is turned off.

Note that auto commit is disabled by default when connecting through Connector/Python. This can be toggled using the [connection parameter](#) `autocommit`.

When the auto commit is turned off, you have to [commit](#) transactions when using transactional storage engines such as InnoDB or NDBCluster.

```
>>> cnx.autocommit
False
>>> cnx.autocommit = True
>>> cnx.autocommit
True
```

Returns True or False.

### 21.6.7.2.32. Property `MySQLConnection.charset_name`

This property returns which character set is used for the connection whether it is connected or not.

Returns a string.

### 21.6.7.2.33. Property `MySQLConnection.collation_name`

This property returns which collation is used for the connection whether it is connected or not.

Returns a string.

### 21.6.7.2.34. Property `MySQLConnection.connection_id`

This property returns the connection ID (thread ID or session ID) for the current connection or None when not connected.

Returns a integer or None.

### 21.6.7.2.35. Property `MySQLConnection.database`

This property is used to set current (active) database executing the [USE](#) command. The property can also be used to retrieve the current database name.

```
>>> cnx.database = 'test'
>>> cnx.database = 'mysql'
>>> cnx.database
u'mysql'
```

Returns a string.

### 21.6.7.2.36. Property `MySQLConnection.get_warnings`

This property is used to toggle whether warnings should be fetched automatically or not. It accepts True or False (default).

Fetching warnings automatically could be useful when debugging queries. Cursors will make warnings available through the method [MySQLCursor.fetchwarnings\(\)](#).

```
>>> cnx.get_warnings = True
>>> cursor.execute('SELECT "a"+1')
```

```
>>> cursor.fetchall()
[(1.0,)]
>>> cursor.fetchwarnings()
[(u'Warning', 1292, u"Truncated incorrect DOUBLE value: 'a'")]
```

Returns True or False.

#### 21.6.7.2.37. Property `MySQLConnection.raise_on_warnings`

This property is used to toggle whether warnings should raise exceptions or not. It accepts True or False (default).

toggling `raise_on_warnings` will also toggle `get_warnings` since warnings need to be fetched so they can be raised as exceptions.

Note that you might always want to check setting SQL Mode if you would like to have the MySQL server directly report warnings as errors. It is also good to use transactional engines so transactions can be rolled back when catching the exception.

Result sets needs to be fetched completely before any exception can be raised. The following example shows the execution of a query which produces a warning

```
>>> cnx.raise_on_warnings = True
>>> cursor.execute('SELECT "a"+1')
>>> cursor.fetchall()
..
mysql.connector.errors.DataError: 1292: Truncated incorrect DOUBLE value: 'a'
```

Returns True or False.

#### 21.6.7.2.38. Property `MySQLConnection.server_host`

This read-only property returns the hostname or IP address used for connecting with the MySQL server.

Returns a string.

#### 21.6.7.2.39. Property `MySQLConnection.server_port`

This read-only property returns the TCP/IP port used for connecting with the MySQL server.

Returns an integer.

#### 21.6.7.2.40. Property `MySQLConnection.sql_mode`

This property is used to retrieve and set the SQL Modes for the current. The value should be list of different modes separated by comma (","), or a sequence of modes, preferably using the constants.SQLMode class.

To unset all modes, pass an empty string or an empty sequence.

```
>>> cnx.sql_mode = 'TRADITIONAL,NO_ENGINE_SUBSTITUTION'
>>> cnx.sql_mode.split(',')
[u'STRICT_TRANS_TABLES', u'STRICT_ALL_TABLES', u'NO_ZERO_IN_DATE',
u'NO_ZERO_DATE', u'ERROR_FOR_DIVISION_BY_ZERO', u'TRADITIONAL',
u'NO_AUTO_CREATE_USER', u'NO_ENGINE_SUBSTITUTION']

>>> from mysql.connector.constants import SQLMode
>>> cnx.sql_mode = [ SQLMode.NO_ZERO_DATE, SQLMode.REAL_AS_FLOAT]
>>> cnx.sql_mode
u'REAL_AS_FLOAT,NO_ZERO_DATE'
```

Returns a string.

#### 21.6.7.2.41. Property `MySQLConnection.time_zone`

This property is used to set the time zone session variable for the current connection and retrieve it.

```
>>> cnx.time_zone = '+00:00'
>>> cur.execute('SELECT NOW()') ; cur.fetchone()
(datetime.datetime(2012, 6, 15, 11, 24, 36),)
>>> cnx.time_zone = '-09:00'
>>> cur.execute('SELECT NOW()') ; cur.fetchone()
(datetime.datetime(2012, 6, 15, 2, 24, 44),)
>>> cnx.time_zone
u'-09:00'
```

Returns a string.

#### 21.6.7.2.42. Property `MySQLConnection.unix_socket`

This read-only property returns the UNIX socket user for connecting with the MySQL server.

Returns a string.

#### 21.6.7.2.43. Property `MySQLConnection.user`

This read-only property returns the username used for connecting with the MySQL server.

Returns a string.

### 21.6.7.3. Class `cursor.MySQLCursor`

The `MySQLCursor` class is used to instantiate object which can execute operation such as SQL queries. They interact with the MySQL server using a `MySQLConnection` object.

#### 21.6.7.3.1. Constructor `cursor.MySQLCursor`

The constructor initializes the instance with the optional `connection`, which should be an instance of `MySQLConnection`.

In most cases, the `MySQLConnection` method `cursor()` is used to instantiate a `MySQLCursor` object.

#### 21.6.7.3.2. Method `MySQLCursor.callproc(procname, args=())`

This method calls a stored procedure with the given name. The `args` sequence of parameters must contain one entry for each argument that the routine expects. The result is returned as modified copy of the input sequence. Input parameters are left untouched, output and input/output parameters replaced with possibly new values.

Result set provided by the stored procedure are automatically fetched and stored as `MySQLBufferedCursor` instances. See `stored_results()` for more information.

The following example shows how to execute a stored procedure which takes two parameters, multiplies the values and returns the product:

```
# Definition of the multiply stored procedure:
# CREATE PROCEDURE multiply(IN pFac1 INT, IN pFac2 INT, OUT pProd INT)
# BEGIN
#   SET pProd := pFac1 * pFac2;
# END

>>> args = (5, 5, 0) # 0 is to hold value of the OUT parameter pProd
>>> cursor.callproc('multiply', args)
('5', '5', 25L)
```



### 21.6.7.3.3. Method `MySQLCursor.close()`

This method will close the MySQL cursor, resetting all results and removing the connection.

Use `close()` every time you are done using the cursor.

### 21.6.7.3.4. Method `MySQLCursor.execute(operation, params=None, multi=False)`

This method prepare the given database `operation` (query or command). The parameters found in the tuple or dictionary `params` will be bound to the variables in the operation. Variables are specified using `%s` markers or named markers `%(name)s`.

For example, insert information about a new employee and selecting again the data of this person:

```

insert = (
    "INSERT INTO employees (emp_no, first_name, last_name, hire_date) "
    "VALUES (%s, %s, %s, %s)"
    data = (2, 'Jane', 'Doe', datetime.date(2012, 3, 23))
    cursor.execute(insert, data)

select = "SELECT * FROM employees WHERE emp_no = %(emp_no)s"
cursor.execute(select, { 'emp_no': 2 })

```

Note that the data is converted from Python object to something MySQL understand. In the above example, the `datetime.date()` instance is converted to '2012-03-23' in the above example.

When `multi` is set to `True`, `execute()` will be able to execute multiple statements. It will return an iterator which makes it possible to go through all results for each statement. Note that using parameters is not working well in this case, and it's usually a good idea to execute each statement on its own.

In the following example we select and insert data in one operation and display the result:

```

operation = 'SELECT 1; INSERT INTO t1 VALUES (); SELECT 2'
for result in cursor.execute(operation):
    if result.with_rows:
        print("Statement '{}' has following rows:".format(
            result.statement))
        print(result.fetchall())
    else:
        print("Affected row(s) by query '{}' was {}".format(
            result.statement, result.rowcount))

```

If the connection was configured to fetch warnings, warnings generated by the operation will be available through the method `MySQLCursor.fetchwarnings()`.

Returns an iterator when `multi` is `True`.

### 21.6.7.3.5. Method `MySQLCursor.executemany(operation, seq_params)`

This method prepares a database operation (query or command) and then execute it against all parameter sequences or mappings found in the sequence `seq_of_params`.

The `executemany()` is simply iterating through the sequence of parameters calling the `execute()` method. Inserting data, however, is optimized by batching them using the multiple rows syntax.

In the following example we are inserting 3 records:

```

data = [
    ('Jane', date(2005, 2, 12)),
    ('Joe', date(2006, 5, 23)),
    ('John', date(2010, 10, 3)),
]
stmt = "INSERT INTO employees (first_name, hire_date) VALUES (%s, %s)"
cursor.executemany(stmt, data)

```

In the above example, the INSERT statement sent to MySQL would be as follows: `INSERT INTO employees (first_name, hire_date) VALUES ('Jane', '2005-02-12'), ('Joe', '2006-05-23'), ('John', '2010-10-03')`.

Note that it is not possible to execute multiple statements using the `executemany()` method. Doing so will raise an `InternalError` exception.

#### 21.6.7.3.6. Method `MySQLCursor.fetchall()`

The method fetches all or remaining rows of a query result set, returning a list of tuples. An empty list is returned when no rows are (anymore) available.

The following examples shows how to retrieve the first 2 rows of a result set, and then retrieve the remaining rows:

```
>>> cursor.execute("SELECT * FROM employees ORDER BY emp_no")
>>> head_rows = cursor.fetchmany(size=2)
>>> remaining_rows = cursor.fetchall()
```

Note that you have to fetch all rows before being able to execute new queries using the same connection.

Returns a list of tuples or empty list when no rows available.

#### 21.6.7.3.7. Method `MySQLCursor.fetchmany(size=1)`

This method fetches the next set of rows of a query results, returning a list of tuples. An empty list is returned when no more rows are available.

The number of rows returned can be specified using the `size` argument, which defaults to one. Fewer rows might be returned, when there are not more rows available than specified by the argument.

Note that you have to fetch all rows before being able to execute new queries using the same connection.

Returns a list of tuples or empty list when no rows available.

#### 21.6.7.3.8. Method `MySQLCursor.fetchone()`

This method retrieves the next row of a query result set, returning a single sequence, or `None` when no more data is available. The returned tuple consists of data returned by the MySQL server converted to Python objects.

The `fetchone()` method is used by `fetchmany()` and `fetchall()`. It is also used when using the `MySQLCursor` instance as an iterator.

The following examples show how to iterate through the result of a query using `fetchone()`:

```
# Using a while-loop
cursor.execute("SELECT * FROM employees")
row = cursor.fetchone()
while row is not None:
    print(row)
    row = cursor.fetchone()

# Using the cursor as iterator
cursor.execute("SELECT * FROM employees")
for row in cursor:
    print(row)
```

Note that you have to fetch all rows before being able to execute new queries using the same connection.

Returns a tuple or `None`.

### 21.6.7.3.9. Method `MySQLCursor.fetchwarnings()`

This method returns a list of tuples containing warnings generated by previously executed statement. Use the connection's `get_warnings` property to toggle whether warnings has to be fetched.

The following example shows a SELECT statement which generated a warning:

```
>>> cnx.get_warnings = True
>>> cursor.execute('SELECT "a"+1')
>>> cursor.fetchall()
[(1.0,)]
>>> cursor.fetchwarnings()
[(u'Warning', 1292, u"Truncated incorrect DOUBLE value: 'a'")]
```

It is also possible to raise errors when warnings are found. See the `MySQLConnection` property `raise_on_warnings`.

Returns a list of tuples.

### 21.6.7.3.10. Method `MySQLCursor.stored_results()`

This method returns an list iterator object which can be used to go through result sets provided by stored procedures after calling them using the `callproc()` method.

In the following example we execute a stored procedure which will provide two result sets. We use `stored_results()` to retrieve them:

```
>>> cursor.callproc('spl')
()
>>> for result in cursor.stored_results():
...     print result.fetchall()
...
[(1,)]
[(2,)]
```

Note that the result sets stay available until you executed another operation or call another stored procedure.

Returns a listiterator.

### 21.6.7.3.11. Property `MySQLCursor.column_names`

This read-only property returns the column names of a result set as sequence of (unicode) strings.

The following example shows how you can create a dictionary out of a tuple containing data with keys using `column_names`:

```
cursor.execute("SELECT last_name, first_name, hire_date "
              "FROM employees WHERE emp_no = %s", (123,))
row = dict(zip(cursor.column_names, cursor.fetchone())
print("{last_name}, {first_name}: {hire_date}".format(row))
```

Returns a tuple.

### 21.6.7.3.12. Property `MySQLCursor.statement`

This read-only property returns the last executed statement. In case multiple statements where executed, it will show the actual statement.

The `statement` property might be useful for debugging and showing what was send to the MySQL server.

Returns a string.

### 21.6.7.3.13. Property `MySQLCursor.with_rows`

This read-only property will return True when the result of the executed operation provides rows.

The `with_rows` property is useful when executing multiple statements and you need to fetch rows. In the following example we only report the affected rows by the `UPDATE` statement:

```
import mysql.connector
cnx = mysql.connector.connect(user='scott', database='test')
cursor = cnx.cursor()
operation = 'SELECT 1; UPDATE t1 SET c1 = 2; SELECT 2'
for result in cursor.execute(operation, multi=True):
    if result.with_rows:
        result.fetchall()
    else:
        print("Updated row(s): {}".format(result.rowcount))
```

### 21.6.7.4. Class `cursor.MySQLCursorBuffered`

This class is inheriting from `cursor.MySQLCursor` and if needed automatically retrieves rows after an operation has been executed.

`MySQLCursorBuffered` can be useful in situations where two queries, with small result sets, need to be combined or computed with each other.

You can either use the `buffered` argument when using the connection's `cursor()` method, or you can use the `buffered connection option` to make all created cursors by default buffering.

```
import mysql.connector

cnx = mysql.connector.connect()

# Only this particular cursor will be buffering results
cursor.cursor(buffered=True)

# All cursors by default buffering
cnx = mysql.connector.connect(buffered=True)
```

See [Section 21.6.5.1, "Tutorial: Raise employee's salary using a buffering cursor"](#) for a practical use case.

### 21.6.7.5. Class `constants.ClientFlag`

This class provides constants defining MySQL client flags which can be used upon connection to configure the session. The `ClientFlag` class is available when importing `mysql.connector`.

```
>>> import mysql.connector
>>> mysql.connector.constants.ClientFlag.FOUND_ROWS
2
```

See [Section 21.6.7.2.30, "Method `MySQLConnection.set\_client\_flags\(flags\)`"](#) and the [connection argument `client\_flag`](#).

Note that the `ClientFlag` class can not be instantiated.

### 21.6.7.6. Class `constants.FieldType`

This class provides all supported MySQL field or data types. They can be useful when dealing with raw data or defining your own converters. The field type is stored with every cursor in the description for each column.

The following example shows how you can print the name of the data types for each of the columns in the result set.

```

from __future__ import print_function
import mysql.connector
from mysql.connector import FieldType

cnx = mysql.connector.connect(user='scott', database='test')
cursor = cnx.cursor()

cursor.execute(
    "SELECT DATE(NOW()) AS `c1`, TIME(NOW()) AS `c2`, "
    "NOW() AS `c3`, 'a string' AS `c4`, 42 AS `c5`")
rows = cursor.fetchall()

for desc in cursor.description:
    colname = desc[0]
    coltype = desc[1]
    print("Column {} has type {}".format(
        colname, FieldType.get_info(coltype)))

cursor.close()
cnx.close()

```

Note that the `FieldType` class can not be instantiated.

#### 21.6.7.7. Class constants `.SQLMode`

This class provides all known MySQL [Server SQL Modes](#). It is mostly used when setting the SQL modes at connection time using the connection's property `sql_mode`. See [Section 21.6.7.2.40](#), "[Property `MySQLConnection.sql\_mode`](#)".

Note that the `SQLMode` class can not be instantiated.

#### 21.6.7.8. Class constants `.CharacterSet`

This class provides all known MySQL characters sets and their default collations. See [Section 21.6.7.2.29](#), "[Method `MySQLConnection.set\_charset\_collation\(charset=None, collation=None\)`](#)" for examples.

Note that the `CharacterSet` class can not be instantiated.

#### 21.6.7.9. Class constants `.RefreshOption`

- `RefreshOption.GRANT`  
Refresh the grant tables, like `FLUSH PRIVILEGES`.
- `RefreshOption.LOG`  
Flush the logs, like `FLUSH LOGS`.
- `RefreshOption.TABLES`  
Flush the table cache, like `FLUSH TABLES`.
- `RefreshOption.HOSTS`  
Flush the host cache, like `FLUSH HOSTS`.
- `RefreshOption.STATUS`  
Reset status variables, like `FLUSH STATUS`.
- `RefreshOption.THREADS`  
Flush the thread cache.

- `RefreshOption.SLAVE`

On a slave replication server, reset the master server information and restart the slave, like `RESET SLAVE`.

- `RefreshOption.MASTER`

On a master replication server, remove the binary log files listed in the binary log index and truncate the index file, like `RESET MASTER`.

## 21.6.8. Connector/Python FAQ

Frequently asked questions.

## 21.7. libmysqld, the Embedded MySQL Server Library

The embedded MySQL server library makes it possible to run a full-featured MySQL server inside a client application. The main benefits are increased speed and more simple management for embedded applications.

The embedded server library is based on the client/server version of MySQL, which is written in C/C++. Consequently, the embedded server also is written in C/C++. There is no embedded server available in other languages.

The API is identical for the embedded MySQL version and the client/server version. To change a threaded application to use the embedded library, you normally only have to add calls to the following functions.

**Table 21.32. MySQL Embedded Server Library Functions**

Function	When to Call
<code>mysql_library_init()</code>	Call it before any other MySQL function is called, preferably early in the <code>main()</code> function.
<code>mysql_library_end()</code>	Call it before your program exits.
<code>mysql_thread_init()</code>	Call it in each thread you create that accesses MySQL.
<code>mysql_thread_end()</code>	Call it before calling <code>pthread_exit()</code> .

Then, link your code with `libmysqld.a` instead of `libmysqlclient.a`. To ensure binary compatibility between your application and the server library, always compile your application against headers for the same series of MySQL that was used to compile the server library. For example, if `libmysqld` was compiled against MySQL 5.1 headers, do not compile your application against MySQL 5.5 headers, or vice versa.

Because the `mysql_library_xxx()` functions are also included in `libmysqlclient.a`, you can change between the embedded and the client/server version by just linking your application with the right library. See [Section 21.8.6.40](#), “`mysql_library_init()`”.

One difference between the embedded server and the standalone server is that for the embedded server, authentication for connections is disabled by default. To use authentication for the embedded server, define the `HAVE_EMBEDDED_PRIVILEGE_CONTROL` [137] compiler flag when you invoke `CMake` to configure your MySQL distribution. See [Section 2.9.4](#), “MySQL Source-Configuration Options”.

### 21.7.1. Compiling Programs with `libmysqld`

In precompiled binary MySQL distributions that include `libmysqld`, the embedded server library, MySQL builds the library using the appropriate vendor compiler if there is one.

To get a `libmysqld` library if you build MySQL from source yourself, you should configure MySQL with the `-DWITH_EMBEDDED_SERVER=1` [136] option. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

When you link your program with `libmysqld`, you must also include the system-specific `pthread` libraries and some libraries that the MySQL server uses. You can get the full list of libraries by executing `mysql_config --libmysqld-libs`.

The correct flags for compiling and linking a threaded program must be used, even if you do not directly call any thread functions in your code.

To compile a C program to include the necessary files to embed the MySQL server library into an executable version of a program, the compiler will need to know where to find various files and need instructions on how to compile the program. The following example shows how a program could be compiled from the command line, assuming that you are using `gcc`, use the GNU C compiler:

```
gcc mysql_test.c -o mysql_test \  
`/usr/local/mysql/bin/mysql_config --include --libmysqld-libs`
```

Immediately following the `gcc` command is the name of the C program source file. After it, the `-o` option is given to indicate that the file name that follows is the name that the compiler is to give to the output file, the compiled program. The next line of code tells the compiler to obtain the location of the include files and libraries and other settings for the system on which it is compiled. The `mysql_config` command is contained in backticks, not single quotation marks.

On some non-`gcc` platforms, the embedded library depends on C++ runtime libraries and linking against the embedded library might result in missing-symbol errors. To solve this, link using a C++ compiler or explicitly list the required libraries on the link command line.

## 21.7.2. Restrictions When Using the Embedded MySQL Server

The embedded server has the following limitations:

- No user-defined functions (UDFs).
- No stack trace on core dump.
- You cannot set this up as a master or a slave (no replication).
- Very large result sets may be unusable on low memory systems.
- You cannot connect to an embedded server from an outside process with sockets or TCP/IP. However, you can connect to an intermediate application, which in turn can connect to an embedded server on the behalf of a remote client or outside process.
- [InnoDB](#) is not reentrant in the embedded server and cannot be used for multiple connections, either successively or simultaneously.
- The Event Scheduler is not available. Because of this, the `event_scheduler` [480] system variable is disabled.

Some of these limitations can be changed by editing the `mysql_embed.h` include file and recompiling MySQL.

## 21.7.3. Options with the Embedded Server

Any options that may be given with the `mysqld` server daemon, may be used with an embedded server library. Server options may be given in an array as an argument to the `mysql_library_init()`, which initializes the server. They also may be given in an option file like `my.cnf`. To specify an option file for a C program, use the `--defaults-file` [219] option as one of the elements of the second argument of the `mysql_library_init()`

function. See [Section 21.8.6.40](#), “`mysql_library_init()`”, for more information on the `mysql_library_init()` function.

Using option files can make it easier to switch between a client/server application and one where MySQL is embedded. Put common options under the `[server]` group. These are read by both MySQL versions. Client/server-specific options should go under the `[mysqld]` section. Put options specific to the embedded MySQL server library in the `[embedded]` section. Options specific to applications go under section labeled `[ApplicationName_SERVER]`. See [Section 4.2.3.3](#), “Using Option Files”.

## 21.7.4. Embedded Server Examples

These two example programs should work without any changes on a Linux or FreeBSD system. For other operating systems, minor changes are needed, mostly with file paths. These examples are designed to give enough details for you to understand the problem, without the clutter that is a necessary part of a real application. The first example is very straightforward. The second example is a little more advanced with some error checking. The first is followed by a command-line entry for compiling the program. The second is followed by a GNUmake file that may be used for compiling instead.

### Example 1

`test1_libmysqld.c`

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>
#include "mysql.h"

MYSQL *mysql;
MYSQL_RES *results;
MYSQL_ROW record;

static char *server_options[] = \
    { "mysql_test", "--defaults-file=my.cnf", NULL };
int num_elements = (sizeof(server_options) / sizeof(char *)) - 1;

static char *server_groups[] = { "libmysqld_server",
    "libmysqld_client", NULL };

int main(void)
{
    mysql_library_init(num_elements, server_options, server_groups);
    mysql = mysql_init(NULL);
    mysql_options(mysql, MYSQL_READ_DEFAULT_GROUP, "libmysqld_client");
    mysql_options(mysql, MYSQL_OPT_USE_EMBEDDED_CONNECTION, NULL);

    mysql_real_connect(mysql, NULL,NULL,NULL, "database1", 0,NULL,0);

    mysql_query(mysql, "SELECT column1, column2 FROM table1");

    results = mysql_store_result(mysql);

    while((record = mysql_fetch_row(results))) {
        printf("%s - %s \n", record[0], record[1]);
    }

    mysql_free_result(results);
    mysql_close(mysql);
    mysql_library_end();

    return 0;
}
```

Here is the command line for compiling the above program:

```
gcc test1_libmysqld.c -o test1_libmysqld \
    `usr/local/mysql/bin/mysql_config --include --libmysqld-libs`
```



**Example 2**

To try the example, create an `test2_libmysqld` directory at the same level as the MySQL source directory. Save the `test2_libmysqld.c` source and the `GNUmakefile` in the directory, and run GNU `make` from inside the `test2_libmysqld` directory.

`test2_libmysqld.c`

```

/*
 * A simple example client, using the embedded MySQL server library
 */

#include <mysql.h>
#include <stdarg.h>
#include <stdio.h>
#include <stdlib.h>

MYSQL *db_connect(const char *dbname);
void db_disconnect(MYSQL *db);
void db_do_query(MYSQL *db, const char *query);

const char *server_groups[] = {
    "test2_libmysqld_SERVER", "embedded", "server", NULL
};

int
main(int argc, char **argv)
{
    MYSQL *one, *two;

    /* mysql_library_init() must be called before any other mysql
     * functions.
     *
     * You can use mysql_library_init(0, NULL, NULL), and it
     * initializes the server using groups = {
     *     "server", "embedded", NULL
     * }.
     *
     * In your $HOME/.my.cnf file, you probably want to put:
     *
     * [test2_libmysqld_SERVER]
     * language = /path/to/source/of/mysql/sql/share/english
     *
     * You could, of course, modify argc and argv before passing
     * them to this function. Or you could create new ones in any
     * way you like. But all of the arguments in argv (except for
     * argv[0], which is the program name) should be valid options
     * for the MySQL server.
     *
     * If you link this client against the normal mysqlclient
     * library, this function is just a stub that does nothing.
     */
    mysql_library_init(argc, argv, (char **)server_groups);

    one = db_connect("test");
    two = db_connect(NULL);

    db_do_query(one, "SHOW TABLE STATUS");
    db_do_query(two, "SHOW DATABASES");

    mysql_close(two);
    mysql_close(one);

    /* This must be called after all other mysql functions */
    mysql_library_end();

    exit(EXIT_SUCCESS);
}

static void
die(MYSQL *db, char *fmt, ...)

```

```

{
    va_list ap;
    va_start(ap, fmt);
    vfprintf(stderr, fmt, ap);
    va_end(ap);
    (void)putc('\n', stderr);
    if (db)
        db_disconnect(db);
    exit(EXIT_FAILURE);
}

MYSQL *
db_connect(const char *dbname)
{
    MYSQL *db = mysql_init(NULL);
    if (!db)
        die(db, "mysql_init failed: no memory");
    /*
     * Notice that the client and server use separate group names.
     * This is critical, because the server does not accept the
     * client's options, and vice versa.
     */
    mysql_options(db, MYSQL_READ_DEFAULT_GROUP, "test2_libmysqld_CLIENT");
    if (!mysql_real_connect(db, NULL, NULL, NULL, dbname, 0, NULL, 0))
        die(db, "mysql_real_connect failed: %s", mysql_error(db));

    return db;
}

void
db_disconnect(MYSQL *db)
{
    mysql_close(db);
}

void
db_do_query(MYSQL *db, const char *query)
{
    if (mysql_query(db, query) != 0)
        goto err;

    if (mysql_field_count(db) > 0)
    {
        MYSQL_RES *res;
        MYSQL_ROW row, end_row;
        int num_fields;

        if (!(res = mysql_store_result(db)))
            goto err;
        num_fields = mysql_num_fields(res);
        while ((row = mysql_fetch_row(res)))
        {
            (void)fputs(">> ", stdout);
            for (end_row = row + num_fields; row < end_row; ++row)
                (void)printf("%s\t", row ? (char*)*row : "NULL");
            (void)fputc('\n', stdout);
        }
        (void)fputc('\n', stdout);
        mysql_free_result(res);
    }
    else
        (void)printf("Affected rows: %lld\n", mysql_affected_rows(db));

    return;

err:
    die(db, "db_do_query failed: %s [%s]", mysql_error(db), query);
}

```

### GNUMakefile

```
# This assumes the MySQL software is installed in /usr/local/mysql
```

```

inc      := /usr/local/mysql/include/mysql
lib      := /usr/local/mysql/lib

# If you have not installed the MySQL software yet, try this instead
#inc     := $(HOME)/mysql-5.6/include
#lib     := $(HOME)/mysql-5.6/libmysqld

CC       := gcc
CPPFLAGS := -I$(inc) -D_THREAD_SAFE -D_REENTRANT
CFLAGS   := -g -W -Wall
LDFLAGS  := -static
# You can change -lmysqld to -mysqlclient to use the
# client/server library
LDLIBS   = -L$(lib) -lmysqld -lm -ldl -lcrypt

ifneq (,$(shell grep FreeBSD /COPYRIGHT 2>/dev/null))
# FreeBSD
LDFLAGS += -pthread
else
# Assume Linux
LDLIBS += -lpthread
endif

# This works for simple one-file test programs
sources := $(wildcard *.c)
objects := $(patsubst %c,%o,$(sources))
targets := $(basename $(sources))

all: $(targets)

clean:
    rm -f $(targets) $(objects) *.core

```

## 21.8. MySQL C API

The C API provides low-level access to the MySQL client/server protocol and enables C programs to access database contents. The C API code is distributed with MySQL and implemented in the `libmysqlclient` library. See [Section 21.8.1, “MySQL C API Implementations”](#).

Most other client APIs use the `libmysqlclient` library to communicate with the MySQL server. (Exceptions are except Connector/J and Connector/Net.) This means that, for example, you can take advantage of many of the same environment variables that are used by other client programs because they are referenced from the library. For a list of these variables, see [Section 4.1, “Overview of MySQL Programs”](#).

For instructions on building client programs using the C API, see [Section 21.8.3.1, “Building C API Client Programs”](#). For programming with threads, see [Section 21.8.3.2, “Writing C API Threaded Client Programs”](#). To create a standalone application which includes the “server” and “client” in the same program (and does not communicate with an external MySQL server), see [Section 21.7, “libmysqld, the Embedded MySQL Server Library”](#).



### Note

If, after an upgrade, you experience problems with compiled client programs, such as [Commands out of sync](#) or unexpected core dumps, the programs were probably compiled using old header or library files. In this case, check the date of the `mysql.h` file and `libmysqlclient.a` library used for compilation to verify that they are from the new MySQL distribution. If not, recompile the programs with the new headers and libraries. Recompilation might also be necessary for programs compiled against the shared client library if the library major version number has changed (for example, from `libmysqlclient.so.15` to `libmysqlclient.so.16`).

Clients have a maximum communication buffer size. The size of the buffer that is allocated initially (16KB) is automatically increased up to the maximum size (16MB by default). Because buffer sizes

are increased only as demand warrants, simply increasing the maximum limit does not in itself cause more resources to be used. This size check is mostly a precaution against erroneous statements and communication packets.

The communication buffer must be large enough to contain a single SQL statement (for client-to-server traffic) and one row of returned data (for server-to-client traffic). Each session's communication buffer is dynamically enlarged to handle any query or row up to the maximum limit. For example, if you have `BLOB` values that contain up to 16MB of data, you must have a communication buffer limit of at least 16MB (in both server and client). The client's default maximum is 16MB, but the default maximum in the server is 1MB. You can increase this by changing the value of the `max_allowed_packet` [502] parameter at server startup. See [Section 8.11.2, "Tuning Server Parameters"](#).

The MySQL server shrinks each communication buffer to `net_buffer_length` [515] bytes after each query. For clients, the size of the buffer associated with a connection is not decreased until the connection is closed, at which time client memory is reclaimed.

## 21.8.1. MySQL C API Implementations

The MySQL C API is a C-based API that client applications written in C can use to communicate with MySQL Server. Client programs refer to C API header files at compile time and link to a C API library file at link time. The library comes in two versions, depending on how the application is intended to communicate with the server:

- `libmysqlclient`: The client version of the library, used for applications that communicate over a network connection as a client of a standalone server process.
- `libmysqld`: The embedded server version of the library, used for applications intended to include an embedded MySQL server within the application itself. The application communicates with its own private server instance.

Both libraries have the same interface. In terms of C API calls, an application communicates with a standalone server the same way it communicates with an embedded server. A given client can be built to communicate with a standalone or embedded server, depending on whether it is linked against `libmysqlclient` or `libmysqld` at build time.

There are two ways to obtain the C API header and library files required to build C API client programs:

- Install a MySQL Server distribution. Server distributions include both `libmysqlclient` and `libmysqld`.
- Install a MySQL Connector/C distribution. Connector/C distributions include only `libmysqlclient`. They do not include `libmysqld`.

For both MySQL Server and MySQL Connector/C, you can install a binary distribution that contains the C API files pre-built, or you can use a source distribution and build the C API files yourself.

Normally, you install either a MySQL Server distribution or a MySQL Connector/C distribution, but not both. It is possible that installing both can cause problems.

The names of the library files to use when linking C API client applications depend on the library type and platform for which a distribution is built:

- On Unix (and Unix-like) systems, the static library is `libmysqlclient.a`. The dynamic library is `libmysqlclient.so` on most Unix systems and `libmysqlclient.dylib` on Mac OS X.

For distributions that include embedded server libraries, the corresponding library names begin with `libmysqld` rather than `libmysqlclient`.

- On Windows, the static library is `mysqlclient.lib` and the dynamic library is `libmysql.dll`. Windows distributions also include `libmysql.lib`, a static import library needed for using the dynamic library.

For distributions that include embedded server libraries, the corresponding library names are `mysqlserver.lib`, `libmysqld.dll`, and `libmysqld.lib`.

Windows distributions also include a set of debug libraries. These have the same names as the nondebug libraries, but are located in the `lib/debug` library. You must use the debug libraries when compiling clients built using the debug C runtime.

On Unix, you may also see libraries that include `_r` in the names. Before MySQL 5.5, these were built as thread-safe (re-entrant) libraries separately from the non-`_r` libraries. As of 5.5, both libraries are the same and the `_r` names are symbolic links to the corresponding non-`_r` names. There is no need to use the `_r` libraries. For example, if you use `mysql_config` to obtain linker flags, you can use `mysql_config --libs` in all cases, even for threaded clients. There is no need to use `mysql_config --libs_r`.

## 21.8.2. Example C API Client Programs

Many of the clients in MySQL source distributions are written in C, such as `mysql`, `mysqladmin`, and `mysqlshow`. If you are looking for examples that demonstrate how to use the C API, take a look at these clients: Obtain a source distribution and look in its `client` directory. See [Section 2.1.3, “How to Get MySQL”](#).

## 21.8.3. Building and Running C API Client Programs

The following sections provide information on building client programs that use the C API. Topics include compiling and linking clients, writing threaded clients, and troubleshooting runtime problems.

### 21.8.3.1. Building C API Client Programs

This section provides guidelines for compiling C programs that use the MySQL C API.

#### Compiling MySQL Clients on Unix

You may need to specify an `-I` option when you compile client programs that use MySQL header files, so that the compiler can find them. For example, if the header files are installed in `/usr/local/mysql/include`, use this option in the compile command:

```
-I/usr/local/mysql/include
```

MySQL clients must be linked using the `-lmysqlclient` option in the link command. You may also need to specify a `-L` option to tell the linker where to find the library. For example, if the library is installed in `/usr/local/mysql/lib`, use these options in the link command:

```
-L/usr/local/mysql/lib -lmysqlclient
```

The path names may differ on your system. Adjust the `-I` and `-L` options as necessary.

To make it simpler to compile MySQL programs on Unix, use the `mysql_config` script. See [Section 4.7.2, “mysql\\_config — Display Options for Compiling Clients”](#).

`mysql_config` displays the options needed for compiling or linking:

```
shell> mysql_config --cflags
shell> mysql_config --libs
```

You can run those commands to get the proper options and add them manually to compilation or link commands. Alternatively, include the output from `mysql_config` directly within command lines using backticks:

```
shell> gcc -c `mysql_config --cflags` progname.c
```

```
shell> gcc -o progname progname.o `mysql_config --libs`
```

## Compiling MySQL Clients on Microsoft Windows

To specify the header and library file locations, use the facilities provided by your development environment.

On Windows, you can link your code with either the dynamic or static client library. The static library is named `mysqlclient.lib` and the dynamic library is named `libmysql.dll`. In addition, the `libmysql.lib` static import library is needed for using the dynamic library.

If you link with the static library, failure can occur unless these conditions are satisfied:

- The client application must be compiled with the same version of Visual Studio used to compile the library.
- The client application should link the C runtime statically by using the `/MT` compiler option.

If the client application is built in in debug mode and uses the static debug C runtime (`/MTd` compiler option), it can link to the `mysqlclient.lib` static library if that library was built using the same option. If the client application uses the dynamic C runtime (`/MD` option, or `/MDd` option in debug mode), it must be linked to the `libmysql.dll` dynamic library. It cannot link to the static client library.

The MSDN page describing the link options can be found here: <http://msdn.microsoft.com/en-us/library/2kzt1wy3.aspx>

## Troubleshooting Problems Linking to the MySQL Client Library

If the linker cannot find the MySQL client library, you might get undefined-reference errors for symbols that start with `mysql_`, such as those shown here:

```
/tmp/ccFKsdPa.o: In function `main':
/tmp/ccFKsdPa.o(.text+0xb): undefined reference to `mysql_init'
/tmp/ccFKsdPa.o(.text+0x31): undefined reference to `mysql_real_connect'
/tmp/ccFKsdPa.o(.text+0x69): undefined reference to `mysql_error'
/tmp/ccFKsdPa.o(.text+0x9a): undefined reference to `mysql_close'
```

You should be able to solve this problem by adding `-Ldir_path -lmysqlclient` at the end of your link command, where `dir_path` represents the path name of the directory where the client library is located. To determine the correct directory, try this command:

```
shell> mysql_config --libs
```

The output from `mysql_config` might indicate other libraries that should be specified on the link command as well. You can include `mysql_config` output directly in your compile or link command using backticks. For example:

```
shell> gcc -o progname progname.o `mysql_config --libs`
```

If an error occurs at link time that the `floor` symbol is undefined, link to the math library by adding `-lm` to the end of the compile/link line. Similarly, if you get undefined-reference errors for other functions that should exist on your system, such as `connect()`, check the manual page for the function in question to determine which libraries you should add to the link command.

If you get undefined-reference errors such as the following for functions that don't exist on your system, it usually means that your MySQL client library was compiled on a system that is not 100% compatible with yours:

```
mf_format.o(.text+0x201): undefined reference to `__lxstat'
```

In this case, you should download the latest MySQL or MySQL Connector/C source distribution and compile the MySQL client library yourself. See [Section 2.9, "Installing MySQL from Source"](#), and [Section 21.5, "MySQL Connector/C"](#).

### 21.8.3.2. Writing C API Threaded Client Programs

The client library is almost thread-safe. The biggest problem is that the subroutines in `net.c` that read from sockets are not interrupt-safe. This was done with the thought that you might want to have your own alarm that can break a long read to a server. If you install interrupt handlers for the `SIGPIPE` interrupt, socket handling should be thread-safe.

To avoid aborting the program when a connection terminates, MySQL blocks `SIGPIPE` on the first call to `mysql_library_init()`, `mysql_init()`, or `mysql_connect()`. To use your own `SIGPIPE` handler, first call `mysql_library_init()`, then install your handler.

If “undefined symbol” errors occur when linking against the `libmysqlclient` client library, in most cases this is because you have not included the thread libraries on the link/compile command.

The client library is thread-safe per connection. You can let two threads share the same connection with the following caveats:

- Multiple threads cannot send a query to the MySQL server at the same time on the same connection. In particular, you must ensure that between calls to `mysql_query()` and `mysql_store_result()` in one thread, no other thread uses the same connection. You must have a mutex lock around your pair of `mysql_query()` and `mysql_store_result()` calls. After `mysql_store_result()` returns, the lock can be released and other threads may query the same connection.

If you use POSIX threads, you can use `pthread_mutex_lock()` and `pthread_mutex_unlock()` to establish and release a mutex lock.

- Many threads can access different result sets that are retrieved with `mysql_store_result()`.
- To use `mysql_use_result()`, you must ensure that no other thread is using the same connection until the result set is closed. However, it really is best for threaded clients that share the same connection to use `mysql_store_result()`.

You need to know the following if you have a thread that did not create the connection to the MySQL database but is calling MySQL functions:

When you call `mysql_init()`, MySQL creates a thread-specific variable for the thread that is used by the debug library (among other things). If you call a MySQL function before the thread has called `mysql_init()`, the thread does not have the necessary thread-specific variables in place and you are likely to end up with a core dump sooner or later. To avoid problems, you must do the following:

1. Call `mysql_library_init()` before any other MySQL functions. It is not thread-safe, so call it before threads are created, or protect the call with a mutex.
2. Arrange for `mysql_thread_init()` to be called early in the thread handler before calling any MySQL function. If you call `mysql_init()`, it will call `mysql_thread_init()` for you.
3. In the thread, call `mysql_thread_end()` before calling `pthread_exit()`. This frees the memory used by MySQL thread-specific variables.

The preceding notes regarding `mysql_init()` also apply to `mysql_connect()`, which calls `mysql_init()`.

### 21.8.3.3. Running C API Clients

Undefined-reference errors might occur at runtime when you try to execute a MySQL program. If these errors specify symbols that start with `mysql_` or indicate that the `libmysqlclient` library cannot be found, it means that your system cannot find the shared `libmysqlclient.so` library. The solution to this problem is to tell your system to search for shared libraries in the directory where that library is located. Use whichever of the following methods is appropriate for your system:



- Add the path of the directory where `libmysqlclient.so` is located to the `LD_LIBRARY_PATH` or `LD_LIBRARY` environment variable.
- On Mac OS X, add the path of the directory where `libmysqlclient.dylib` is located to the `DYLD_LIBRARY_PATH` environment variable.
- Copy the shared-library files (such as `libmysqlclient.so`) to some directory that is searched by your system, such as `/lib`, and update the shared library information by executing `ldconfig`. Be sure to copy all related files. A shared library might exist under several names, using symlinks to provide the alternate names.

If the application is linked to the embedded server library, runtime error messages will indicate the `libmysqld` rather than `libmysqlclient` library, but the solution to the problem is the same as just described.

## 21.8.4. C API Data Structures

This section describes C API data structures other than those used for prepared statements. For information about the latter, see [Section 21.8.8, "C API Prepared Statement Data Structures"](#).

- `MYSQL`

This structure represents a handle to one database connection. It is used for almost all MySQL functions. Do not try to make a copy of a `MYSQL` structure. There is no guarantee that such a copy will be usable.

- `MYSQL_RES`

This structure represents the result of a query that returns rows (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`). The information returned from a query is called the *result set* in the remainder of this section.

- `MYSQL_ROW`

This is a type-safe representation of one row of data. It is currently implemented as an array of counted byte strings. (You cannot treat these as null-terminated strings if field values may contain binary data, because such values may contain null bytes internally.) Rows are obtained by calling `mysql_fetch_row()`.

- `MYSQL_FIELD`

This structure contains metadata: information about a field, such as the field's name, type, and size. Its members are described in more detail later in this section. You may obtain the `MYSQL_FIELD` structures for each field by calling `mysql_fetch_field()` repeatedly. Field values are not part of this structure; they are contained in a `MYSQL_ROW` structure.

- `MYSQL_FIELD_OFFSET`

This is a type-safe representation of an offset into a MySQL field list. (Used by `mysql_field_seek()`.) Offsets are field numbers within a row, beginning at zero.

- `my_ulonglong`

The type used for the number of rows and for `mysql_affected_rows()`, `mysql_num_rows()`, and `mysql_insert_id()`. This type provides a range of 0 to `1.84e19`.

On some systems, attempting to print a value of type `my_ulonglong` does not work. To print such a value, convert it to `unsigned long` and use a `%lu` print format. Example:

```
printf ("Number of rows: %lu\n",
       (unsigned long) mysql_num_rows(result));
```



- `my_bool`

A boolean type, for values that are true (nonzero) or false (zero).

The `MYSQL_FIELD` structure contains the members described in the following list. The definitions apply primarily for columns of result sets such as those produced by `SELECT` statements. In MySQL 5.6, `MYSQL_FIELD` structures are also used to provide metadata for `OUT` and `INOUT` parameters returned from stored procedures executed using prepared `CALL` statements. For such parameters, some of the structure members have a meaning different from the meaning for column values.

- `char * name`

The name of the field, as a null-terminated string. If the field was given an alias with an `AS` clause, the value of `name` is the alias. For a procedure parameter, the parameter name.

- `char * org_name`

The name of the field, as a null-terminated string. Aliases are ignored. For expressions, the value is an empty string. For a procedure parameter, the parameter name.

- `char * table`

The name of the table containing this field, if it isn't a calculated field. For calculated fields, the `table` value is an empty string. If the column is selected from a view, `table` names the view. If the table or view was given an alias with an `AS` clause, the value of `table` is the alias. For a `UNION`, the value is the empty string. For a procedure parameter, the procedure name.

- `char * org_table`

The name of the table, as a null-terminated string. Aliases are ignored. If the column is selected from a view, `org_table` names the underlying table. For a `UNION`, the value is the empty string. For a procedure parameter, the procedure name.

- `char * db`

The name of the database that the field comes from, as a null-terminated string. If the field is a calculated field, `db` is an empty string. For a `UNION`, the value is the empty string. For a procedure parameter, the name of the database containing the procedure.

- `char * catalog`

The catalog name. This value is always `"def"`.

- `char * def`

The default value of this field, as a null-terminated string. This is set only if you use `mysql_list_fields()`.

- `unsigned long length`

The width of the field. This corresponds to the display length, in bytes.

The server determines the `length` value before it generates the result set, so this is the minimum length required for a data type capable of holding the largest possible value from the result column, without knowing in advance the actual values that will be produced by the query for the result set.

- `unsigned long max_length`

The maximum width of the field for the result set (the length in bytes of the longest field value for the rows actually in the result set). If you use `mysql_store_result()` or `mysql_list_fields()`, this contains the maximum length for the field. If you use `mysql_use_result()`, the value of this variable is zero.

The value of `max_length` is the length of the string representation of the values in the result set. For example, if you retrieve a `FLOAT` column and the “widest” value is `-12.345`, `max_length` is 7 (the length of `'-12.345'`).

If you are using prepared statements, `max_length` is not set by default because for the binary protocol the lengths of the values depend on the types of the values in the result set. (See [Section 21.8.8, “C API Prepared Statement Data Structures”](#).) If you want the `max_length` values anyway, enable the `STMT_ATTR_UPDATE_MAX_LENGTH` option with `mysql_stmt_attr_set()` and the lengths will be set when you call `mysql_stmt_store_result()`. (See [Section 21.8.10.3, “mysql\\_stmt\\_attr\\_set\(\)”](#), and [Section 21.8.10.28, “mysql\\_stmt\\_store\\_result\(\)”](#).)

- `unsigned int name_length`

The length of `name`.

- `unsigned int org_name_length`

The length of `org_name`.

- `unsigned int table_length`

The length of `table`.

- `unsigned int org_table_length`

The length of `org_table`.

- `unsigned int db_length`

The length of `db`.

- `unsigned int catalog_length`

The length of `catalog`.

- `unsigned int def_length`

The length of `def`.

- `unsigned int flags`

Bit-flags that describe the field. The `flags` value may have zero or more of the bits set that are shown in the following table.

Flag Value	Flag Description
<code>NOT_NULL_FLAG</code>	Field can't be <code>NULL</code>
<code>PRI_KEY_FLAG</code>	Field is part of a primary key
<code>UNIQUE_KEY_FLAG</code>	Field is part of a unique key
<code>MULTIPLE_KEY_FLAG</code>	Field is part of a nonunique key
<code>UNSIGNED_FLAG</code>	Field has the <code>UNSIGNED</code> attribute
<code>ZEROFILL_FLAG</code>	Field has the <code>ZEROFILL</code> attribute
<code>BINARY_FLAG</code>	Field has the <code>BINARY</code> attribute
<code>AUTO_INCREMENT_FLAG</code>	Field has the <code>AUTO_INCREMENT</code> attribute
<code>ENUM_FLAG</code>	Field is an <code>ENUM</code> (deprecated)
<code>SET_FLAG</code>	Field is a <code>SET</code> (deprecated)
<code>BLOB_FLAG</code>	Field is a <code>BLOB</code> or <code>TEXT</code> (deprecated)

Flag Value	Flag Description
<code>TIMESTAMP_FLAG</code>	Field is a <code>TIMESTAMP</code> (deprecated)
<code>NUM_FLAG</code>	Field is numeric; see additional notes following table
<code>NO_DEFAULT_VALUE_FLAG</code>	Field has no default value; see additional notes following table

Some of these flags indicate data type information and are superseded by or used in conjunction with the `MYSQL_TYPE_xxx` value in the `field->type` member described later:

- To check for `BLOB` or `TIMESTAMP` values, check whether `type` is `MYSQL_TYPE_BLOB` or `MYSQL_TYPE_TIMESTAMP`. (The `BLOB_FLAG` and `TIMESTAMP_FLAG` flags are unneeded.)
- `ENUM` and `SET` values are returned as strings. For these, check that the `type` value is `MYSQL_TYPE_STRING` and that the `ENUM_FLAG` or `SET_FLAG` flag is set in the `flags` value.

`NUM_FLAG` indicates that a column is numeric. This includes columns with a type of `MYSQL_TYPE_DECIMAL`, `MYSQL_TYPE_NEWDECIMAL`, `MYSQL_TYPE_TINY`, `MYSQL_TYPE_SHORT`, `MYSQL_TYPE_LONG`, `MYSQL_TYPE_FLOAT`, `MYSQL_TYPE_DOUBLE`, `MYSQL_TYPE_NULL`, `MYSQL_TYPE_LONGLONG`, `MYSQL_TYPE_INT24`, and `MYSQL_TYPE_YEAR`.

`NO_DEFAULT_VALUE_FLAG` indicates that a column has no `DEFAULT` clause in its definition. This does not apply to `NULL` columns (because such columns have a default of `NULL`), or to `AUTO_INCREMENT` columns (which have an implied default value).

The following example illustrates a typical use of the `flags` value:

```
if (field->flags & NOT_NULL_FLAG)
    printf("Field can't be null\n");
```

You may use the convenience macros shown in the following table to determine the boolean status of the `flags` value.

Flag Status	Description
<code>IS_NOT_NULL(flags)</code>	True if this field is defined as <code>NOT NULL</code>
<code>IS_PRI_KEY(flags)</code>	True if this field is a primary key
<code>IS_BLOB(flags)</code>	True if this field is a <code>BLOB</code> or <code>TEXT</code> (deprecated; test <code>field-&gt;type</code> instead)

- `unsigned int decimals`

The number of decimals for numeric fields, and (as of MySQL 5.6.4) the fractional seconds precision for temporal fields.

- `unsigned int charsetnr`

An ID number that indicates the character set/collation pair for the field.

To distinguish between binary and nonbinary data for string data types, check whether the `charsetnr` value is 63. If so, the character set is `binary`, which indicates binary rather than nonbinary data. This enables you to distinguish `BINARY` from `CHAR`, `VARBINARY` from `VARCHAR`, and the `BLOB` types from the `TEXT` types.

`charsetnr` values are the same as those displayed in the `Id` column of the `SHOW COLLATION` statement or the `ID` column of the `INFORMATION_SCHEMA.COLLATIONS` table. You can use those information sources to see which character set and collation specific `charsetnr` values indicate:

```
mysql> SHOW COLLATION WHERE Id = 63;
+-----+-----+-----+-----+-----+-----+
| Collation | Charset | Id | Default | Compiled | Sortlen |
```

```

+-----+-----+-----+-----+-----+
| binary | binary | 63 | Yes | Yes | 1 |
+-----+-----+-----+-----+-----+

mysql> SELECT COLLATION_NAME, CHARACTER_SET_NAME
-> FROM INFORMATION_SCHEMA.COLLATIONS WHERE ID = 33;
+-----+-----+
| COLLATION_NAME | CHARACTER_SET_NAME |
+-----+-----+
| utf8_general_ci | utf8 |
+-----+-----+

```

- enum `enum_field_types type`

The type of the field. The `type` value may be one of the `MYSQL_TYPE_` symbols shown in the following table.

Type Value	Type Description
<code>MYSQL_TYPE_TINY</code>	<code>TINYINT</code> field
<code>MYSQL_TYPE_SHORT</code>	<code>SMALLINT</code> field
<code>MYSQL_TYPE_LONG</code>	<code>INTEGER</code> field
<code>MYSQL_TYPE_INT24</code>	<code>MEDIUMINT</code> field
<code>MYSQL_TYPE_LONGLONG</code>	<code>BIGINT</code> field
<code>MYSQL_TYPE_DECIMAL</code>	<code>DECIMAL</code> or <code>NUMERIC</code> field
<code>MYSQL_TYPE_NEWDECIMAL</code>	Precision math <code>DECIMAL</code> or <code>NUMERIC</code>
<code>MYSQL_TYPE_FLOAT</code>	<code>FLOAT</code> field
<code>MYSQL_TYPE_DOUBLE</code>	<code>DOUBLE</code> or <code>REAL</code> field
<code>MYSQL_TYPE_BIT</code>	<code>BIT</code> field
<code>MYSQL_TYPE_TIMESTAMP</code>	<code>TIMESTAMP</code> field
<code>MYSQL_TYPE_DATE</code>	<code>DATE</code> field
<code>MYSQL_TYPE_TIME</code>	<code>TIME</code> field
<code>MYSQL_TYPE_DATETIME</code>	<code>DATETIME</code> field
<code>MYSQL_TYPE_YEAR</code>	<code>YEAR</code> field
<code>MYSQL_TYPE_STRING</code>	<code>CHAR</code> or <code>BINARY</code> field
<code>MYSQL_TYPE_VAR_STRING</code>	<code>VARCHAR</code> or <code>VARBINARY</code> field
<code>MYSQL_TYPE_BLOB</code>	<code>BLOB</code> or <code>TEXT</code> field (use <code>max_length</code> to determine the maximum length)
<code>MYSQL_TYPE_SET</code>	<code>SET</code> field
<code>MYSQL_TYPE_ENUM</code>	<code>ENUM</code> field
<code>MYSQL_TYPE_GEOMETRY</code>	Spatial field
<code>MYSQL_TYPE_NULL</code>	<code>NULL</code> -type field

The `MYSQL_TYPE_TIME2`, `MYSQL_TYPE_DATETIME2`, and `MYSQL_TYPE_TIMESTAMP2` type codes are used only on the server side. Clients see the `MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATETIME`, and `MYSQL_TYPE_TIMESTAMP` codes.

You can use the `IS_NUM( )` macro to test whether a field has a numeric type. Pass the `type` value to `IS_NUM( )` and it evaluates to `TRUE` if the field is numeric:

```

if (IS_NUM(field->type))
    printf("Field is numeric\n");

```

`ENUM` and `SET` values are returned as strings. For these, check that the `type` value is `MYSQL_TYPE_STRING` and that the `ENUM_FLAG` or `SET_FLAG` flag is set in the `flags` value.

## 21.8.5. C API Function Overview

The functions available in the C API are summarized here and described in greater detail in a later section. See [Section 21.8.6, “C API Function Descriptions”](#).

**Table 21.33. C API Function Names and Descriptions**

Function	Description
<code>my_init()</code>	Initialize global variables, and thread handler in thread-safe programs
<code>mysql_affected_rows()</code>	Returns the number of rows changed/deleted/inserted by the last <code>UPDATE</code> , <code>DELETE</code> , or <code>INSERT</code> query
<code>mysql_autocommit()</code>	Toggles autocommit mode on/off
<code>mysql_change_user()</code>	Changes user and database on an open connection
<code>mysql_character_set_name()</code>	Return default character set name for current connection
<code>mysql_client_find_plugin()</code>	Return pointer to plugin
<code>mysql_client_register_plugin()</code>	Register a plugin
<code>mysql_close()</code>	Closes a server connection
<code>mysql_commit()</code>	Commits the transaction
<code>mysql_connect()</code>	Connects to a MySQL server (this function is deprecated; use <code>mysql_real_connect()</code> instead)
<code>mysql_create_db()</code>	Creates a database (this function is deprecated; use the SQL statement <code>CREATE DATABASE</code> instead)
<code>mysql_data_seek()</code>	Seeks to an arbitrary row number in a query result set
<code>mysql_debug()</code>	Does a <code>DEBUG_PUSH</code> with the given string
<code>mysql_drop_db()</code>	Drops a database (this function is deprecated; use the SQL statement <code>DROP DATABASE</code> instead)
<code>mysql_dump_debug_info()</code>	Makes the server write debug information to the log
<code>mysql_eof()</code>	Determines whether the last row of a result set has been read (this function is deprecated; <code>mysql_errno()</code> or <code>mysql_error()</code> may be used instead)
<code>mysql_errno()</code>	Returns the error number for the most recently invoked MySQL function
<code>mysql_error()</code>	Returns the error message for the most recently invoked MySQL function
<code>mysql_escape_string()</code>	Escapes special characters in a string for use in an SQL statement
<code>mysql_fetch_field()</code>	Returns the type of the next table field
<code>mysql_fetch_field_direct()</code>	Returns the type of a table field, given a field number
<code>mysql_fetch_fields()</code>	Returns an array of all field structures
<code>mysql_fetch_lengths()</code>	Returns the lengths of all columns in the current row
<code>mysql_fetch_row()</code>	Fetches the next row from the result set
<code>mysql_field_count()</code>	Returns the number of result columns for the most recent statement
<code>mysql_field_seek()</code>	Puts the column cursor on a specified column
<code>mysql_field_tell()</code>	Returns the position of the field cursor used for the last <code>mysql_fetch_field()</code>

Function	Description
<code>mysql_free_result()</code>	Frees memory used by a result set
<code>mysql_get_character_set_name()</code>	Returns information about default character set
<code>mysql_get_client_info()</code>	Returns client version information as a string
<code>mysql_get_client_version()</code>	Returns client version information as an integer
<code>mysql_get_host_info()</code>	Returns a string describing the connection
<code>mysql_get_proto_info()</code>	Returns the protocol version used by the connection
<code>mysql_get_server_info()</code>	Returns the server version number
<code>mysql_get_server_version()</code>	Returns version number of server as an integer
<code>mysql_get_ssl_cipher()</code>	Return current SSL cipher
<code>mysql_hex_string()</code>	Encode string in hexadecimal format
<code>mysql_info()</code>	Returns information about the most recently executed query
<code>mysql_init()</code>	Gets or initializes a <code>MYSQL</code> structure
<code>mysql_insert_id()</code>	Returns the ID generated for an <code>AUTO_INCREMENT</code> column by the previous query
<code>mysql_kill()</code>	Kills a given thread
<code>mysql_library_end()</code>	Finalize the MySQL C API library
<code>mysql_library_init()</code>	Initialize the MySQL C API library
<code>mysql_list_dbs()</code>	Returns database names matching a simple regular expression
<code>mysql_list_fields()</code>	Returns field names matching a simple regular expression
<code>mysql_list_processes()</code>	Returns a list of the current server threads
<code>mysql_list_tables()</code>	Returns table names matching a simple regular expression
<code>mysql_load_plugin()</code>	Load a plugin
<code>mysql_load_plugin_v()</code>	Load a plugin
<code>mysql_more_results()</code>	Checks whether any more results exist
<code>mysql_next_result()</code>	Returns/initiates the next result in multiple-result executions
<code>mysql_num_fields()</code>	Returns the number of columns in a result set
<code>mysql_num_rows()</code>	Returns the number of rows in a result set
<code>mysql_options()</code>	Sets connect options for <code>mysql_real_connect()</code>
<code>mysql_options4()</code>	Sets connect options for <code>mysql_real_connect()</code>
<code>mysql_ping()</code>	Checks whether the connection to the server is working, reconnecting as necessary
<code>mysql_plugin_options()</code>	Set a plugin option
<code>mysql_query()</code>	Executes an SQL query specified as a null-terminated string
<code>mysql_real_connect()</code>	Connects to a MySQL server
<code>mysql_real_escape_string()</code>	Escapes special characters in a string for use in an SQL statement, taking into account the current character set of the connection
<code>mysql_real_query()</code>	Executes an SQL query specified as a counted string
<code>mysql_refresh()</code>	Flush or reset tables and caches
<code>mysql_reload()</code>	Tells the server to reload the grant tables
<code>mysql_rollback()</code>	Rolls back the transaction
<code>mysql_row_seek()</code>	Seeks to a row offset in a result set, using value returned from <code>mysql_row_tell()</code>

Function	Description
<code>mysql_row_tell()</code>	Returns the row cursor position
<code>mysql_select_db()</code>	Selects a database
<code>mysql_server_end()</code>	Finalize the MySQL C API library
<code>mysql_server_init()</code>	Initialize the MySQL C API library
<code>mysql_set_character_set()</code>	Set default character set for current connection
<code>mysql_set_local_infile_defaults()</code>	Set the <code>LOAD DATA LOCAL INFILE</code> handler callbacks to their default values
<code>mysql_set_local_infile_handler()</code>	Install application-specific <code>LOAD DATA LOCAL INFILE</code> handler callbacks
<code>mysql_set_server_option()</code>	Sets an option for the connection (like <code>multi-statements</code> )
<code>mysql_sqlstate()</code>	Returns the SQLSTATE error code for the last error
<code>mysql_shutdown()</code>	Shuts down the database server
<code>mysql_ssl_set()</code>	Prepare to establish SSL connection to server
<code>mysql_stat()</code>	Returns the server status as a string
<code>mysql_store_result()</code>	Retrieves a complete result set to the client
<code>mysql_thread_end()</code>	Finalize thread handler
<code>mysql_thread_id()</code>	Returns the current thread ID
<code>mysql_thread_init()</code>	Initialize thread handler
<code>mysql_thread_safe()</code>	Returns 1 if the clients are compiled as thread-safe
<code>mysql_use_result()</code>	Initiates a row-by-row result set retrieval
<code>mysql_warning_count()</code>	Returns the warning count for the previous SQL statement

Application programs should use this general outline for interacting with MySQL:

1. Initialize the MySQL library by calling `mysql_library_init()`. This function exists in both the `libmysqlclient` C client library and the `libmysqld` embedded server library, so it is used whether you build a regular client program by linking with the `-libmysqlclient` flag, or an embedded server application by linking with the `-libmysqld` flag.
2. Initialize a connection handler by calling `mysql_init()` and connect to the server by calling `mysql_real_connect()`.
3. Issue SQL statements and process their results. (The following discussion provides more information about how to do this.)
4. Close the connection to the MySQL server by calling `mysql_close()`.
5. End use of the MySQL library by calling `mysql_library_end()`.

The purpose of calling `mysql_library_init()` and `mysql_library_end()` is to provide proper initialization and finalization of the MySQL library. For applications that are linked with the client library, they provide improved memory management. If you don't call `mysql_library_end()`, a block of memory remains allocated. (This does not increase the amount of memory used by the application, but some memory leak detectors will complain about it.) For applications that are linked with the embedded server, these calls start and stop the server.

In a nonmulti-threaded environment, the call to `mysql_library_init()` may be omitted, because `mysql_init()` will invoke it automatically as necessary. However, `mysql_library_init()` is not thread-safe in a multi-threaded environment, and thus neither is `mysql_init()`, which calls `mysql_library_init()`. You must either call `mysql_library_init()` prior to spawning any threads, or else use a mutex to protect the call, whether you invoke `mysql_library_init()` or indirectly through `mysql_init()`. This should be done prior to any other client library call.



To connect to the server, call `mysql_init()` to initialize a connection handler, then call `mysql_real_connect()` with that handler (along with other information such as the host name, user name, and password). Upon connection, `mysql_real_connect()` sets the `reconnect` flag (part of the `MYSQL` structure) to a value of `1` in versions of the API older than 5.0.3, or `0` in newer versions. A value of `1` for this flag indicates that if a statement cannot be performed because of a lost connection, to try reconnecting to the server before giving up. You can use the `MYSQL_OPT_RECONNECT` option to `mysql_options()` to control reconnection behavior. When you are done with the connection, call `mysql_close()` to terminate it.

While a connection is active, the client may send SQL statements to the server using `mysql_query()` or `mysql_real_query()`. The difference between the two is that `mysql_query()` expects the query to be specified as a null-terminated string whereas `mysql_real_query()` expects a counted string. If the string contains binary data (which may include null bytes), you must use `mysql_real_query()`.

For each non-`SELECT` query (for example, `INSERT`, `UPDATE`, `DELETE`), you can find out how many rows were changed (affected) by calling `mysql_affected_rows()`.

For `SELECT` queries, you retrieve the selected rows as a result set. (Note that some statements are `SELECT`-like in that they return rows. These include `SHOW`, `DESCRIBE`, and `EXPLAIN`. Treat these statements the same way as `SELECT` statements.)

There are two ways for a client to process result sets. One way is to retrieve the entire result set all at once by calling `mysql_store_result()`. This function acquires from the server all the rows returned by the query and stores them in the client. The second way is for the client to initiate a row-by-row result set retrieval by calling `mysql_use_result()`. This function initializes the retrieval, but does not actually get any rows from the server.

In both cases, you access rows by calling `mysql_fetch_row()`. With `mysql_store_result()`, `mysql_fetch_row()` accesses rows that have previously been fetched from the server. With `mysql_use_result()`, `mysql_fetch_row()` actually retrieves the row from the server. Information about the size of the data in each row is available by calling `mysql_fetch_lengths()`.

After you are done with a result set, call `mysql_free_result()` to free the memory used for it.

The two retrieval mechanisms are complementary. Choose the approach that is most appropriate for each client application. In practice, clients tend to use `mysql_store_result()` more commonly.

An advantage of `mysql_store_result()` is that because the rows have all been fetched to the client, you not only can access rows sequentially, you can move back and forth in the result set using `mysql_data_seek()` or `mysql_row_seek()` to change the current row position within the result set. You can also find out how many rows there are by calling `mysql_num_rows()`. On the other hand, the memory requirements for `mysql_store_result()` may be very high for large result sets and you are more likely to encounter out-of-memory conditions.

An advantage of `mysql_use_result()` is that the client requires less memory for the result set because it maintains only one row at a time (and because there is less allocation overhead, `mysql_use_result()` can be faster). Disadvantages are that you must process each row quickly to avoid tying up the server, you don't have random access to rows within the result set (you can only access rows sequentially), and you don't know how many rows are in the result set until you have retrieved them all. Furthermore, you *must* retrieve all the rows even if you determine in mid-retrieval that you've found the information you were looking for.

The API makes it possible for clients to respond appropriately to statements (retrieving rows only as necessary) without knowing whether the statement is a `SELECT`. You can do this by calling `mysql_store_result()` after each `mysql_query()` (or `mysql_real_query()`). If the result set call succeeds, the statement was a `SELECT` and you can read the rows. If the result set call fails, call `mysql_field_count()` to determine whether a result was actually to be expected. If `mysql_field_count()` returns zero, the statement returned no data (indicating that it was an `INSERT`, `UPDATE`, `DELETE`, and so forth), and was not expected to return rows. If `mysql_field_count()` is nonzero, the statement should have returned rows,



but didn't. This indicates that the statement was a `SELECT` that failed. See the description for `mysql_field_count()` for an example of how this can be done.

Both `mysql_store_result()` and `mysql_use_result()` enable you to obtain information about the fields that make up the result set (the number of fields, their names and types, and so forth). You can access field information sequentially within the row by calling `mysql_fetch_field()` repeatedly, or by field number within the row by calling `mysql_fetch_field_direct()`. The current field cursor position may be changed by calling `mysql_field_seek()`. Setting the field cursor affects subsequent calls to `mysql_fetch_field()`. You can also get information for fields all at once by calling `mysql_fetch_fields()`.

For detecting and reporting errors, MySQL provides access to error information by means of the `mysql_errno()` and `mysql_error()` functions. These return the error code or error message for the most recently invoked function that can succeed or fail, enabling you to determine when an error occurred and what it was.

## 21.8.6. C API Function Descriptions

In the descriptions here, a parameter or return value of `NULL` means `NULL` in the sense of the C programming language, not a MySQL `NULL` value.

Functions that return a value generally return a pointer or an integer. Unless specified otherwise, functions returning a pointer return a non-`NULL` value to indicate success or a `NULL` value to indicate an error, and functions returning an integer return zero to indicate success or nonzero to indicate an error. Note that “nonzero” means just that. Unless the function description says otherwise, do not test against a value other than zero:

```
if (result)                /* correct */
    ... error ...

if (result < 0)           /* incorrect */
    ... error ...

if (result == -1)        /* incorrect */
    ... error ...
```

When a function returns an error, the **Errors** subsection of the function description lists the possible types of errors. You can find out which of these occurred by calling `mysql_errno()`. A string representation of the error may be obtained by calling `mysql_error()`.

### 21.8.6.1. `mysql_affected_rows()`

```
my_ulonglong mysql_affected_rows(MYSQL *mysql)
```

#### Description

`mysql_affected_rows()` may be called immediately after executing a statement with `mysql_query()` or `mysql_real_query()`. It returns the number of rows changed, deleted, or inserted by the last statement if it was an `UPDATE`, `DELETE`, or `INSERT`. For `SELECT` statements, `mysql_affected_rows()` works like `mysql_num_rows()`.

For `UPDATE` statements, the affected-rows value by default is the number of rows actually changed. If you specify the `CLIENT_FOUND_ROWS` flag to `mysql_real_connect()` when connecting to `mysqld`, the affected-rows value is the number of rows “found”; that is, matched by the `WHERE` clause.

For `REPLACE` statements, the affected-rows value is 2 if the new row replaced an old row, because in this case, one row was inserted after the duplicate was deleted.

For `INSERT ... ON DUPLICATE KEY UPDATE` statements, the affected-rows value is 1 if the row is inserted as a new row and 2 if an existing row is updated.

Following a `CALL` statement for a stored procedure, `mysql_affected_rows()` returns the value that it would return for the last statement executed within the procedure, or 0 if that statement would return

-1. Within the procedure, you can use `ROW_COUNT()` [1272] at the SQL level to obtain the affected-rows value for individual statements.

In MySQL 5.6, `mysql_affected_rows()` returns a meaningful value for a wider range of statements. For details, see the description for `ROW_COUNT()` [1272] in [Section 12.14, “Information Functions”](#).

## Return Values

An integer greater than zero indicates the number of rows affected or retrieved. Zero indicates that no records were updated for an `UPDATE` statement, no rows matched the `WHERE` clause in the query or that no query has yet been executed. -1 indicates that the query returned an error or that, for a `SELECT` query, `mysql_affected_rows()` was called prior to calling `mysql_store_result()`.

Because `mysql_affected_rows()` returns an unsigned value, you can check for -1 by comparing the return value to `(my_ulonglong)-1` (or to `(my_ulonglong)~0`, which is equivalent).

## Errors

None.

## Example

```
char *stmt = "UPDATE products SET cost=cost*1.25
             WHERE group=10";
mysql_query(&mysql, stmt);
printf("%ld products updated",
       (long) mysql_affected_rows(&mysql));
```

### 21.8.6.2. `mysql_autocommit()`

```
my_bool mysql_autocommit(MYSQL *mysql, my_bool mode)
```

## Description

Sets autocommit mode on if `mode` is 1, off if `mode` is 0.

## Return Values

Zero if successful. Nonzero if an error occurred.

## Errors

None.

### 21.8.6.3. `mysql_change_user()`

```
my_bool mysql_change_user(MYSQL *mysql, const char *user, const char
                          *password, const char *db)
```

## Description

Changes the user and causes the database specified by `db` to become the default (current) database on the connection specified by `mysql`. In subsequent queries, this database is the default for table references that do not include an explicit database specifier.

`mysql_change_user()` fails if the connected user cannot be authenticated or doesn't have permission to use the database. In this case, the user and database are not changed.

The `db` parameter may be set to `NULL` if you don't want to have a default database.

This command resets the state as if one had done a new connect. (See [Section 21.8.15, “Controlling Automatic Reconnection Behavior”](#).) It always performs a `ROLLBACK` of any active transactions, closes

and drops all temporary tables, and unlocks all locked tables. Session system variables are reset to the values of the corresponding global system variables. Prepared statements are released and [HANDLER](#) variables are closed. Locks acquired with [GET\\_LOCK\(\)](#) [1275] are released. These effects occur even if the user didn't change.

## Return Values

Zero for success. Nonzero if an error occurred.

## Errors

The same that you can get from [mysql\\_real\\_connect\(\)](#).

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC](#) [3719]

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR](#) [3718]

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST](#) [3719]

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR](#) [3718]

An unknown error occurred.

- [ER\\_UNKNOWN\\_COM\\_ERROR](#) [3663]

The MySQL server doesn't implement this command (probably an old server).

- [ER\\_ACCESS\\_DENIED\\_ERROR](#) [3663]

The user or password was wrong.

- [ER\\_BAD\\_DB\\_ERROR](#) [3664]

The database didn't exist.

- [ER\\_DBACCESS\\_DENIED\\_ERROR](#) [3663]

The user did not have access rights to the database.

- [ER\\_WRONG\\_DB\\_NAME](#) [3667]

The database name was too long.

## Example

```
if (mysql_change_user(&mysql, "user", "password", "new_database"))
{
    fprintf(stderr, "Failed to change user. Error: %s\n",
            mysql_error(&mysql));
}
```

### 21.8.6.4. [mysql\\_character\\_set\\_name\(\)](#)

```
const char *mysql_character_set_name(MYSQL *mysql)
```

## Description

Returns the default character set name for the current connection.

**Return Values**

The default character set name

**Errors**

None.

**21.8.6.5. `mysql_close()`**

```
void mysql_close(MYSQL *mysql)
```

**Description**

Closes a previously opened connection. `mysql_close()` also deallocates the connection handle pointed to by `mysql` if the handle was allocated automatically by `mysql_init()` or `mysql_connect()`.

**Return Values**

None.

**Errors**

None.

**21.8.6.6. `mysql_commit()`**

```
my_bool mysql_commit(MYSQL *mysql)
```

**Description**

Commits the current transaction.

The action of this function is subject to the value of the `completion_type` [470] system variable. In particular, if the value of `completion_type` [470] is `RELEASE` (or 2), the server performs a release after terminating a transaction and closes the client connection. Call `mysql_close()` from the client program to close the connection from the client side.

**Return Values**

Zero if successful. Nonzero if an error occurred.

**Errors**

None.

**21.8.6.7. `mysql_connect()`**

```
MYSQL *mysql_connect(MYSQL *mysql, const char *host, const char *user, const char *passwd)
```

**Description**

This function is deprecated. Use `mysql_real_connect()` instead.

`mysql_connect()` attempts to establish a connection to a MySQL database engine running on `host`. `mysql_connect()` must complete successfully before you can execute any of the other API functions, with the exception of `mysql_get_client_info()`.

The meanings of the parameters are the same as for the corresponding parameters for `mysql_real_connect()` with the difference that the connection parameter may be `NULL`. In this case, the C API allocates memory for the connection structure automatically and frees it when you call

`mysql_close()`. The disadvantage of this approach is that you can't retrieve an error message if the connection fails. (To get error information from `mysql_errno()` or `mysql_error()`, you must provide a valid `MYSQL` pointer.)

### Return Values

Same as for `mysql_real_connect()`.

### Errors

Same as for `mysql_real_connect()`.

### 21.8.6.8. `mysql_create_db()`

```
int mysql_create_db(MYSQL *mysql, const char *db)
```

### Description

Creates the database named by the `db` parameter.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `CREATE DATABASE` statement instead.

### Return Values

Zero if the database was created successfully. Nonzero if an error occurred.

### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

### Example

```
if(mysql_create_db(&mysql, "my_database"))
{
    fprintf(stderr, "Failed to create new database. Error: %s\n",
            mysql_error(&mysql));
}
```

### 21.8.6.9. `mysql_data_seek()`

```
void mysql_data_seek(MYSQL_RES *result, my_ulonglong offset)
```

### Description

Seeks to an arbitrary row in a query result set. The `offset` value is a row number. Specify a value in the range from 0 to `mysql_num_rows(result)-1`.

This function requires that the result set structure contains the entire result of the query, so `mysql_data_seek()` may be used only in conjunction with `mysql_store_result()`, not with `mysql_use_result()`.

**Return Values**

None.

**Errors**

None.

**21.8.6.10. `mysql_debug()`**

```
void mysql_debug(const char *debug)
```

**Description**

Does a `DEBUG_PUSH` with the given string. `mysql_debug()` uses the Fred Fish debug library. To use this function, you must compile the client library to support debugging. See [MySQL Internals: Porting to Other Systems](#).

**Return Values**

None.

**Errors**

None.

**Example**

The call shown here causes the client library to generate a trace file in `/tmp/client.trace` on the client machine:

```
mysql_debug("d:t:0,/tmp/client.trace");
```

**21.8.6.11. `mysql_drop_db()`**

```
int mysql_drop_db(MYSQL *mysql, const char *db)
```

**Description**

Drops the database named by the `db` parameter.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `DROP DATABASE` statement instead.

**Return Values**

Zero if the database was dropped successfully. Nonzero if an error occurred.

**Errors**

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### Example

```
if(mysql_drop_db(&mysql, "my_database"))
    fprintf(stderr, "Failed to drop the database: Error: %s\n",
            mysql_error(&mysql));
```

### 21.8.6.12. [mysql\\_dump\\_debug\\_info\(\)](#)

```
int mysql_dump_debug_info(MYSQL *mysql)
```

#### Description

Instructs the server to write some debug information to the log. For this to work, the connected user must have the [SUPER \[756\]](#) privilege.

#### Return Values

Zero if the command was successful. Nonzero if an error occurred.

#### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.13. [mysql\\_eof\(\)](#)

```
my_bool mysql_eof(MYSQL_RES *result)
```

#### Description

This function is deprecated. [mysql\\_errno\(\)](#) or [mysql\\_error\(\)](#) may be used instead.

[mysql\\_eof\(\)](#) determines whether the last row of a result set has been read.

If you acquire a result set from a successful call to [mysql\\_store\\_result\(\)](#), the client receives the entire set in one operation. In this case, a `NULL` return from [mysql\\_fetch\\_row\(\)](#) always means the end of the result set has been reached and it is unnecessary to call [mysql\\_eof\(\)](#). When used with [mysql\\_store\\_result\(\)](#), [mysql\\_eof\(\)](#) always returns true.

On the other hand, if you use [mysql\\_use\\_result\(\)](#) to initiate a result set retrieval, the rows of the set are obtained from the server one by one as you call [mysql\\_fetch\\_row\(\)](#) repeatedly. Because an error may occur on the connection during this process, a `NULL` return value from [mysql\\_fetch\\_row\(\)](#) does not necessarily mean the end of the result set was reached normally. In this case, you can use [mysql\\_eof\(\)](#) to determine what happened. [mysql\\_eof\(\)](#) returns a nonzero value if the end of the result set was reached and zero if an error occurred.

Historically, `mysql_eof()` predates the standard MySQL error functions `mysql_errno()` and `mysql_error()`. Because those error functions provide the same information, their use is preferred over `mysql_eof()`, which is deprecated. (In fact, they provide more information, because `mysql_eof()` returns only a boolean value whereas the error functions indicate a reason for the error when one occurs.)

## Return Values

Zero if no error occurred. Nonzero if the end of the result set has been reached.

## Errors

None.

## Example

The following example shows how you might use `mysql_eof()`:

```
mysql_query(&mysql, "SELECT * FROM some_table");
result = mysql_use_result(&mysql);
while((row = mysql_fetch_row(result)))
{
    // do something with data
}
if(!mysql_eof(result)) // mysql_fetch_row() failed due to an error
{
    fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
}
```

However, you can achieve the same effect with the standard MySQL error functions:

```
mysql_query(&mysql, "SELECT * FROM some_table");
result = mysql_use_result(&mysql);
while((row = mysql_fetch_row(result)))
{
    // do something with data
}
if(mysql_errno(&mysql)) // mysql_fetch_row() failed due to an error
{
    fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
}
```

### 21.8.6.14. `mysql_errno()`

```
unsigned int mysql_errno(MYSQL *mysql)
```

## Description

For the connection specified by `mysql`, `mysql_errno()` returns the error code for the most recently invoked API function that can succeed or fail. A return value of zero means that no error occurred. Client error message numbers are listed in the MySQL `errmsg.h` header file. Server error message numbers are listed in `mysqld_error.h`. Errors also are listed at [Appendix C, Errors, Error Codes, and Common Problems](#).

Note that some functions like `mysql_fetch_row()` don't set `mysql_errno()` if they succeed.

A rule of thumb is that all functions that have to ask the server for information reset `mysql_errno()` if they succeed.

MySQL-specific error numbers returned by `mysql_errno()` differ from SQLSTATE values returned by `mysql_sqlstate()`. For example, the `mysql` client program displays errors using the following format, where 1146 is the `mysql_errno()` value and '42S02' is the corresponding `mysql_sqlstate()` value:

```
shell> SELECT * FROM no_such_table;
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```



## Return Values

An error code value for the last `mysql_xxx()` call, if it failed. zero means no error occurred.

## Errors

None.

### 21.8.6.15. `mysql_error()`

```
const char *mysql_error(MYSQL *mysql)
```

## Description

For the connection specified by `mysql`, `mysql_error()` returns a null-terminated string containing the error message for the most recently invoked API function that failed. If a function didn't fail, the return value of `mysql_error()` may be the previous error or an empty string to indicate no error.

A rule of thumb is that all functions that have to ask the server for information reset `mysql_error()` if they succeed.

For functions that reset `mysql_error()`, the following two tests are equivalent:

```
if(*mysql_error(&mysql))
{
    // an error occurred
}

if(mysql_error(&mysql)[0])
{
    // an error occurred
}
```

The language of the client error messages may be changed by recompiling the MySQL client library. Currently, you can choose error messages in several different languages. See [Section 10.2, "Setting the Error Message Language"](#).

## Return Values

A null-terminated character string that describes the error. An empty string if no error occurred.

## Errors

None.

### 21.8.6.16. `mysql_escape_string()`

Use `mysql_real_escape_string()` instead!

This function is identical to `mysql_real_escape_string()` except that `mysql_real_escape_string()` takes a connection handler as its first argument and escapes the string according to the current character set. `mysql_escape_string()` does not take a connection argument and does not respect the current character set.

### 21.8.6.17. `mysql_fetch_field()`

```
MYSQL_FIELD *mysql_fetch_field(MYSQL_RES *result)
```

## Description

Returns the definition of one column of a result set as a `MYSQL_FIELD` structure. Call this function repeatedly to retrieve information about all columns in the result set. `mysql_fetch_field()` returns `NULL` when no more fields are left.

`mysql_fetch_field()` is reset to return information about the first field each time you execute a new `SELECT` query. The field returned by `mysql_fetch_field()` is also affected by calls to `mysql_field_seek()`.

If you've called `mysql_query()` to perform a `SELECT` on a table but have not called `mysql_store_result()`, MySQL returns the default blob length (8KB) if you call `mysql_fetch_field()` to ask for the length of a `BLOB` field. (The 8KB size is chosen because MySQL doesn't know the maximum length for the `BLOB`. This should be made configurable sometime.) Once you've retrieved the result set, `field->max_length` contains the length of the largest value for this column in the specific query.

## Return Values

The `MYSQL_FIELD` structure for the current column. `NULL` if no columns are left.

## Errors

None.

## Example

```
MYSQL_FIELD *field;

while((field = mysql_fetch_field(result)))
{
    printf("field name %s\n", field->name);
}
```

### 21.8.6.18. `mysql_fetch_field_direct()`

```
MYSQL_FIELD *mysql_fetch_field_direct(MYSQL_RES *result, unsigned int
fieldnr)
```

## Description

Given a field number `fieldnr` for a column within a result set, returns that column's field definition as a `MYSQL_FIELD` structure. Use this function to retrieve the definition for an arbitrary column. Specify a value for `fieldnr` in the range from 0 to `mysql_num_fields(result)-1`.

## Return Values

The `MYSQL_FIELD` structure for the specified column.

## Errors

None.

## Example

```
unsigned int num_fields;
unsigned int i;
MYSQL_FIELD *field;

num_fields = mysql_num_fields(result);
for(i = 0; i < num_fields; i++)
{
    field = mysql_fetch_field_direct(result, i);
    printf("Field %u is %s\n", i, field->name);
}
```

### 21.8.6.19. `mysql_fetch_fields()`

```
MYSQL_FIELD *mysql_fetch_fields(MYSQL_RES *result)
```

## Description

Returns an array of all `MYSQL_FIELD` structures for a result set. Each structure provides the field definition for one column of the result set.

## Return Values

An array of `MYSQL_FIELD` structures for all columns of a result set.

## Errors

None.

## Example

```
unsigned int num_fields;
unsigned int i;
MYSQL_FIELD *fields;

num_fields = mysql_num_fields(result);
fields = mysql_fetch_fields(result);
for(i = 0; i < num_fields; i++)
{
    printf("Field %u is %s\n", i, fields[i].name);
}
```

### 21.8.6.20. `mysql_fetch_lengths()`

```
unsigned long *mysql_fetch_lengths(MYSQL_RES *result)
```

## Description

Returns the lengths of the columns of the current row within a result set. If you plan to copy field values, this length information is also useful for optimization, because you can avoid calling `strlen()`. In addition, if the result set contains binary data, you **must** use this function to determine the size of the data, because `strlen()` returns incorrect results for any field containing null characters.

The length for empty columns and for columns containing `NULL` values is zero. To see how to distinguish these two cases, see the description for `mysql_fetch_row()`.

## Return Values

An array of unsigned long integers representing the size of each column (not including any terminating null characters). `NULL` if an error occurred.

## Errors

`mysql_fetch_lengths()` is valid only for the current row of the result set. It returns `NULL` if you call it before calling `mysql_fetch_row()` or after retrieving all rows in the result.

## Example

```
MYSQL_ROW row;
unsigned long *lengths;
unsigned int num_fields;
unsigned int i;

row = mysql_fetch_row(result);
if (row)
{
    num_fields = mysql_num_fields(result);
    lengths = mysql_fetch_lengths(result);
    for(i = 0; i < num_fields; i++)
    {
        printf("Column %u is %lu bytes in length.\n",
            i, lengths[i]);
    }
}
```

```
}
}
```

### 21.8.6.21. `mysql_fetch_row()`

```
MYSQL_ROW mysql_fetch_row(MYSQL_RES *result)
```

#### Description

Retrieves the next row of a result set. When used after `mysql_store_result()`, `mysql_fetch_row()` returns `NULL` when there are no more rows to retrieve. When used after `mysql_use_result()`, `mysql_fetch_row()` returns `NULL` when there are no more rows to retrieve or if an error occurred.

The number of values in the row is given by `mysql_num_fields(result)`. If `row` holds the return value from a call to `mysql_fetch_row()`, pointers to the values are accessed as `row[0]` to `row[mysql_num_fields(result)-1]`. `NULL` values in the row are indicated by `NULL` pointers.

The lengths of the field values in the row may be obtained by calling `mysql_fetch_lengths()`. Empty fields and fields containing `NULL` both have length 0; you can distinguish these by checking the pointer for the field value. If the pointer is `NULL`, the field is `NULL`; otherwise, the field is empty.

#### Return Values

A `MYSQL_ROW` structure for the next row. `NULL` if there are no more rows to retrieve or if an error occurred.

#### Errors

Note that error is not reset between calls to `mysql_fetch_row()`

- `CR_SERVER_LOST` [3719]

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR` [3718]

An unknown error occurred.

#### Example

```
MYSQL_ROW row;
unsigned int num_fields;
unsigned int i;

num_fields = mysql_num_fields(result);
while ((row = mysql_fetch_row(result)))
{
    unsigned long *lengths;
    lengths = mysql_fetch_lengths(result);
    for(i = 0; i < num_fields; i++)
    {
        printf("[%.*s] ", (int) lengths[i],
            row[i] ? row[i] : "NULL");
    }
    printf("\n");
}
```

### 21.8.6.22. `mysql_field_count()`

```
unsigned int mysql_field_count(MYSQL *mysql)
```

#### Description

Returns the number of columns for the most recent query on the connection.

The normal use of this function is when `mysql_store_result()` returned `NULL` (and thus you have no result set pointer). In this case, you can call `mysql_field_count()` to determine whether `mysql_store_result()` should have produced a nonempty result. This enables the client program to take proper action without knowing whether the query was a `SELECT` (or `SELECT`-like) statement. The example shown here illustrates how this may be done.

See [Section 21.8.14.1, “Why `mysql\_store\_result\(\)` Sometimes Returns `NULL` After `mysql\_query\(\)` Returns Success”](#).

## Return Values

An unsigned integer representing the number of columns in a result set.

## Errors

None.

## Example

```
MYSQL_RES *result;
unsigned int num_fields;
unsigned int num_rows;

if (mysql_query(&mysql, query_string))
{
    // error
}
else // query succeeded, process any data returned by it
{
    result = mysql_store_result(&mysql);
    if (result) // there are rows
    {
        num_fields = mysql_num_fields(result);
        // retrieve rows, then call mysql_free_result(result)
    }
    else // mysql_store_result() returned nothing; should it have?
    {
        if(mysql_field_count(&mysql) == 0)
        {
            // query does not return data
            // (it was not a SELECT)
            num_rows = mysql_affected_rows(&mysql);
        }
        else // mysql_store_result() should have returned data
        {
            fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
        }
    }
}
}
```

An alternative is to replace the `mysql_field_count(&mysql)` call with `mysql_errno(&mysql)`. In this case, you are checking directly for an error from `mysql_store_result()` rather than inferring from the value of `mysql_field_count()` whether the statement was a `SELECT`.

### 21.8.6.23. `mysql_field_seek()`

```
MYSQL_FIELD_OFFSET mysql_field_seek(MYSQL_RES *result, MYSQL_FIELD_OFFSET
offset)
```

## Description

Sets the field cursor to the given offset. The next call to `mysql_fetch_field()` retrieves the field definition of the column associated with that offset.

To seek to the beginning of a row, pass an `offset` value of zero.

**Return Values**

The previous value of the field cursor.

**Errors**

None.

**21.8.6.24. `mysql_field_tell()`**

```
MYSQL_FIELD_OFFSET mysql_field_tell(MYSQL_RES *result)
```

**Description**

Returns the position of the field cursor used for the last `mysql_fetch_field()`. This value can be used as an argument to `mysql_field_seek()`.

**Return Values**

The current offset of the field cursor.

**Errors**

None.

**21.8.6.25. `mysql_free_result()`**

```
void mysql_free_result(MYSQL_RES *result)
```

**Description**

Frees the memory allocated for a result set by `mysql_store_result()`, `mysql_use_result()`, `mysql_list_dbs()`, and so forth. When you are done with a result set, you must free the memory it uses by calling `mysql_free_result()`.

Do not attempt to access a result set after freeing it.

**Return Values**

None.

**Errors**

None.

**21.8.6.26. `mysql_get_character_set_info()`**

```
void mysql_get_character_set_info(MYSQL *mysql, MY_CHARSET_INFO *cs)
```

**Description**

This function provides information about the default client character set. The default character set may be changed with the `mysql_set_character_set()` function.

**Example**

This example shows the fields that are available in the `MY_CHARSET_INFO` structure:

```
if (!mysql_set_character_set(&mysql, "utf8"))
{
    MY_CHARSET_INFO cs;
    mysql_get_character_set_info(&mysql, &cs);
    printf("character set information:\n");
    printf("character set+collation number: %d\n", cs.number);
    printf("character set name: %s\n", cs.name);
}
```

```
printf("collation name: %s\n", cs.csname);
printf("comment: %s\n", cs.comment);
printf("directory: %s\n", cs.dir);
printf("multi byte character min. length: %d\n", cs.mbminlen);
printf("multi byte character max. length: %d\n", cs.mbmaxlen);
}
```

### 21.8.6.27. `mysql_get_client_info()`

```
const char *mysql_get_client_info(void)
```

#### Description

Returns a string that represents the client library version.

#### Return Values

A character string that represents the MySQL client library version.

#### Errors

None.

### 21.8.6.28. `mysql_get_client_version()`

```
unsigned long mysql_get_client_version(void)
```

#### Description

Returns an integer that represents the client library version. The value has the format `XYZZ` where `X` is the major version, `YY` is the release level, and `ZZ` is the version number within the release level. For example, a value of `40102` represents a client library version of `4.1.2`.

#### Return Values

An integer that represents the MySQL client library version.

#### Errors

None.

### 21.8.6.29. `mysql_get_host_info()`

```
const char *mysql_get_host_info(MYSQL *mysql)
```

#### Description

Returns a string describing the type of connection in use, including the server host name.

#### Return Values

A character string representing the server host name and the connection type.

#### Errors

None.

### 21.8.6.30. `mysql_get_proto_info()`

```
unsigned int mysql_get_proto_info(MYSQL *mysql)
```

#### Description

Returns the protocol version used by current connection.

**Return Values**

An unsigned integer representing the protocol version used by the current connection.

**Errors**

None.

**21.8.6.31. `mysql_get_server_info()`**

```
const char *mysql_get_server_info(MYSQL *mysql)
```

**Description**

Returns a string that represents the server version number.

**Return Values**

A character string that represents the server version number.

**Errors**

None.

**21.8.6.32. `mysql_get_server_version()`**

```
unsigned long mysql_get_server_version(MYSQL *mysql)
```

**Description**

Returns the version number of the server as an integer.

**Return Values**

A number that represents the MySQL server version in this format:

```
major_version*10000 + minor_version *100 + sub_version
```

For example, 5.1.5 is returned as 50105.

This function is useful in client programs for quickly determining whether some version-specific server capability exists.

**Errors**

None.

**21.8.6.33. `mysql_get_ssl_cipher()`**

```
const char *mysql_get_ssl_cipher(MYSQL *mysql)
```

**Description**

`mysql_get_ssl_cipher()` returns the SSL cipher used for the given connection to the server. `mysql` is the connection handler returned from `mysql_init()`.

**Return Values**

A string naming the SSL cipher used for the connection, or `NULL` if no cipher is being used.

**21.8.6.34. `mysql_hex_string()`**

```
unsigned long mysql_hex_string(char *to, const char *from, unsigned long length)
```



## Description

This function is used to create a legal SQL string that you can use in an SQL statement. See [Section 9.1.1, "String Literals"](#).

The string in `from` is encoded to hexadecimal format, with each character encoded as two hexadecimal digits. The result is placed in `to` and a terminating null byte is appended.

The string pointed to by `from` must be `length` bytes long. You must allocate the `to` buffer to be at least `length*2+1` bytes long. When `mysql_hex_string()` returns, the contents of `to` is a null-terminated string. The return value is the length of the encoded string, not including the terminating null character.

The return value can be placed into an SQL statement using either `0xvalue` or `X'value'` format. However, the return value does not include the `0x` or `X'...`. The caller must supply whichever of those is desired.

## Example

```
char query[1000],*end;

end = strmov(query,"INSERT INTO test_table values(");
end = strmov(end,"0x");
end += mysql_hex_string(end,"What is this",12);
end = strmov(end,"0x");
end += mysql_hex_string(end,"binary data: \0\r\n",16);
*end++ = '\0';

if (mysql_real_query(&mysql,query,(unsigned int) (end - query)))
{
    fprintf(stderr, "Failed to insert row, Error: %s\n",
            mysql_error(&mysql));
}
```

The `strmov()` function used in the example is included in the `libmysqlclient` library and works like `strcpy()` but returns a pointer to the terminating null of the first parameter.

## Return Values

The length of the value placed into `to`, not including the terminating null character.

## Errors

None.

### 21.8.6.35. `mysql_info()`

```
const char *mysql_info(MYSQL *mysql)
```

## Description

Retrieves a string providing information about the most recently executed statement, but only for the statements listed here. For other statements, `mysql_info()` returns `NULL`. The format of the string varies depending on the type of statement, as described here. The numbers are illustrative only; the string contains values appropriate for the statement.

- `INSERT INTO ... SELECT ...`

String format: `Records: 100 Duplicates: 0 Warnings: 0`

- `INSERT INTO ... VALUES (...),(...),(...)...`

String format: `Records: 3 Duplicates: 0 Warnings: 0`

- `LOAD DATA INFILE ...`

String format: `Records: 1 Deleted: 0 Skipped: 0 Warnings: 0`

- `ALTER TABLE`

String format: `Records: 3 Duplicates: 0 Warnings: 0`

- `UPDATE`

String format: `Rows matched: 40 Changed: 40 Warnings: 0`

Note that `mysql_info()` returns a non-NULL value for `INSERT ... VALUES` only for the multiple-row form of the statement (that is, only if multiple value lists are specified).

## Return Values

A character string representing additional information about the most recently executed statement. `NULL` if no information is available for the statement.

## Errors

None.

### 21.8.6.36. `mysql_init()`

```
MYSQL *mysql_init(MYSQL *mysql)
```

## Description

Allocates or initializes a `MYSQL` object suitable for `mysql_real_connect()`. If `mysql` is a `NULL` pointer, the function allocates, initializes, and returns a new object. Otherwise, the object is initialized and the address of the object is returned. If `mysql_init()` allocates a new object, it is freed when `mysql_close()` is called to close the connection.

## Return Values

An initialized `MYSQL*` handle. `NULL` if there was insufficient memory to allocate a new object.

## Errors

In case of insufficient memory, `NULL` is returned.

### 21.8.6.37. `mysql_insert_id()`

```
my_ulonglong mysql_insert_id(MYSQL *mysql)
```

## Description

Returns the value generated for an `AUTO_INCREMENT` column by the previous `INSERT` or `UPDATE` statement. Use this function after you have performed an `INSERT` statement into a table that contains an `AUTO_INCREMENT` field, or have used `INSERT` or `UPDATE` to set a column value with `LAST_INSERT_ID(expr)` [1269].

The return value of `mysql_insert_id()` is always zero unless explicitly updated under one of the following conditions:

- `INSERT` statements that store a value into an `AUTO_INCREMENT` column. This is true whether the value is automatically generated by storing the special values `NULL` or `0` into the column, or is an explicit nonspecial value.
- In the case of a multiple-row `INSERT` statement, `mysql_insert_id()` returns the first automatically generated `AUTO_INCREMENT` value that was successfully inserted.

If no rows are successfully inserted, `mysql_insert_id()` returns 0.

- If an `INSERT ... SELECT` statement is executed, and no automatically generated value is successfully inserted, `mysql_insert_id()` returns the ID of the last inserted row.
- If an `INSERT ... SELECT` statement uses `LAST_INSERT_ID(expr)` [1269], `mysql_insert_id()` returns `expr`.
- `INSERT` statements that generate an `AUTO_INCREMENT` value by inserting `LAST_INSERT_ID(expr)` [1269] into any column or by updating any column to `LAST_INSERT_ID(expr)` [1269].
- If the previous statement returned an error, the value of `mysql_insert_id()` is undefined.

The return value of `mysql_insert_id()` can be simplified to the following sequence:

1. If there is an `AUTO_INCREMENT` column, and an automatically generated value was successfully inserted, return the first such value.
2. If `LAST_INSERT_ID(expr)` [1269] occurred in the statement, return `expr`, even if there was an `AUTO_INCREMENT` column in the affected table.
3. The return value varies depending on the statement used. When called after an `INSERT` statement:
  - If there is an `AUTO_INCREMENT` column in the table, and there were some explicit values for this column that were successfully inserted into the table, return the last of the explicit values.

When called after an `INSERT ... ON DUPLICATE KEY UPDATE` statement:

- If there is an `AUTO_INCREMENT` column in the table and there were some explicit successfully inserted values or some updated values, return the last of the inserted or updated values.

`mysql_insert_id()` returns 0 if the previous statement does not use an `AUTO_INCREMENT` value. If you need to save the value for later, be sure to call `mysql_insert_id()` immediately after the statement that generates the value.

The value of `mysql_insert_id()` is affected only by statements issued within the current client connection. It is not affected by statements issued by other clients.

The `LAST_INSERT_ID()` [1269] SQL function will contain the value of the first automatically generated value that was successfully inserted. `LAST_INSERT_ID()` [1269] is not reset between statements because the value of that function is maintained in the server. Another difference from `mysql_insert_id()` is that `LAST_INSERT_ID()` [1269] is not updated if you set an `AUTO_INCREMENT` column to a specific nonspecial value. See [Section 12.14, "Information Functions"](#).

`mysql_insert_id()` returns 0 following a `CALL` statement for a stored procedure that generates an `AUTO_INCREMENT` value because in this case `mysql_insert_id()` applies to `CALL` and not the statement within the procedure. Within the procedure, you can use `LAST_INSERT_ID()` [1269] at the SQL level to obtain the `AUTO_INCREMENT` value.

The reason for the differences between `LAST_INSERT_ID()` [1269] and `mysql_insert_id()` is that `LAST_INSERT_ID()` [1269] is made easy to use in scripts while `mysql_insert_id()` tries to provide more exact information about what happens to the `AUTO_INCREMENT` column.

## Return Values

Described in the preceding discussion.

## Errors

None.

### 21.8.6.38. `mysql_kill()`

```
int mysql_kill(MYSQL *mysql, unsigned long pid)
```

#### Description

Asks the server to kill the thread specified by `pid`.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `KILL` statement instead.

`mysql_kill()` cannot handle values larger than 32 bits, but as of MySQL 5.6.9 to guard against killing the wrong thread returns an error in these cases:

- If given an ID larger than 32 bits, `mysql_kill()` returns a `CR_INVALID_CONN_HANDLE` [3721] error.
- After the server's internal thread ID counter reaches a value larger than 32 bits, it returns an `ER_DATA_OUT_OF_RANGE` [3705] error for any `mysql_kill()` invocation and `mysql_kill()` fails.

#### Return Values

Zero for success. Nonzero if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC` [3719]  
Commands were executed in an improper order.
- `CR_INVALID_CONN_HANDLE` [3721]  
The `pid` was larger than 32 bits.
- `CR_SERVER_GONE_ERROR` [3718]  
The MySQL server has gone away.
- `CR_SERVER_LOST` [3719]  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR` [3718]  
An unknown error occurred.
- `ER_DATA_OUT_OF_RANGE` [3705]  
The server's internal thread ID counter has reached a value larger than 32 bits, at which point it rejects all `mysql_kill()` invocations.

### 21.8.6.39. `mysql_library_end()`

```
void mysql_library_end(void)
```

#### Description

This function finalizes the MySQL library. Call it when you are done using the library (for example, after disconnecting from the server). The action taken by the call depends on whether your application is linked to the MySQL client library or the MySQL embedded server library. For a client program linked against the `libmysqlclient` library by using the `-lmysqlclient` flag, `mysql_library_end()` performs some memory management to clean up. For an embedded server application linked against

the `libmysqld` library by using the `-lmysqld` flag, `mysql_library_end()` shuts down the embedded server and then cleans up.

For usage information, see [Section 21.8.5, “C API Function Overview”](#), and [Section 21.8.6.40, “mysql\\_library\\_init\(\)”](#).

#### 21.8.6.40. `mysql_library_init()`

```
int mysql_library_init(int argc, char **argv, char **groups)
```

##### Description

Call this function to initialize the MySQL library before you call any other MySQL function, whether your application is a regular client program or uses the embedded server. If the application uses the embedded server, this call starts the server and initializes any subsystems (`mysys`, `InnoDB`, and so forth) that the server uses.

After your application is done using the MySQL library, call `mysql_library_end()` to clean up. See [Section 21.8.6.39, “mysql\\_library\\_end\(\)”](#).

The choice of whether the application operates as a regular client or uses the embedded server depends on whether you use the `libmysqlclient` or `libmysqld` library at link time to produce the final executable. For additional information, see [Section 21.8.5, “C API Function Overview”](#).

In a nonmulti-threaded environment, the call to `mysql_library_init()` may be omitted, because `mysql_init()` will invoke it automatically as necessary. However, `mysql_library_init()` is not thread-safe in a multi-threaded environment, and thus neither is `mysql_init()`, which calls `mysql_library_init()`. You must either call `mysql_library_init()` prior to spawning any threads, or else use a mutex to protect the call, whether you invoke `mysql_library_init()` or indirectly through `mysql_init()`. Do this prior to any other client library call.

The `argc` and `argv` arguments are analogous to the arguments to `main()`, and enable passing of options to the embedded server. For convenience, `argc` may be 0 (zero) if there are no command-line arguments for the server. This is the usual case for applications intended for use only as regular (nonembedded) clients, and the call typically is written as `mysql_library_init(0, NULL, NULL)`.

```
#include <mysql.h>
#include <stdlib.h>

int main(void) {
    if (mysql_library_init(0, NULL, NULL)) {
        fprintf(stderr, "could not initialize MySQL library\n");
        exit(1);
    }

    /* Use any MySQL API functions here */

    mysql_library_end();

    return EXIT_SUCCESS;
}
```

When arguments are to be passed (`argc` is greater than 0), the first element of `argv` is ignored (it typically contains the program name). `mysql_library_init()` makes a copy of the arguments so it is safe to destroy `argv` or `groups` after the call.

For embedded applications, if you want to connect to an external server without starting the embedded server, you have to specify a negative value for `argc`.

The `groups` argument is an array of strings that indicate the groups in option files from which to read options. See [Section 4.2.3.3, “Using Option Files”](#). Make the final entry in the array `NULL`. For convenience, if the `groups` argument itself is `NULL`, the `[server]` and `[embedded]` groups are used by default.

```

#include <mysql.h>
#include <stdlib.h>

static char *server_args[] = {
    "this_program",          /* this string is not used */
    "--datadir=",
    "--key_buffer_size=32M"
};
static char *server_groups[] = {
    "embedded",
    "server",
    "this_program_SERVER",
    (char *)NULL
};

int main(void) {
    if (mysql_library_init(sizeof(server_args) / sizeof(char *),
        server_args, server_groups)) {
        fprintf(stderr, "could not initialize MySQL library\n");
        exit(1);
    }

    /* Use any MySQL API functions here */

    mysql_library_end();

    return EXIT_SUCCESS;
}

```

## Return Values

Zero if successful. Nonzero if an error occurred.

### 21.8.6.41. `mysql_list_dbs()`

```
MYSQL_RES *mysql_list_dbs(MYSQL *mysql, const char *wild)
```

## Description

Returns a result set consisting of database names on the server that match the simple regular expression specified by the `wild` parameter. `wild` may contain the wildcard characters “%” or “\_”, or may be a `NULL` pointer to match all databases. Calling `mysql_list_dbs()` is similar to executing the query `SHOW DATABASES [LIKE wild]`.

You must free the result set with `mysql_free_result()`.

## Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)  
Out of memory.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

#### 21.8.6.42. `mysql_list_fields()`

```
MYSQL_RES *mysql_list_fields(MYSQL *mysql, const char *table, const char *wild)
```

##### Description

Returns a result set consisting of field names in the given table that match the simple regular expression specified by the `wild` parameter. `wild` may contain the wildcard characters “%” or “\_”, or may be a `NULL` pointer to match all fields. Calling `mysql_list_fields()` is similar to executing the query `SHOW COLUMNS FROM tbl_name [LIKE wild]`.

It is preferable to use `SHOW COLUMNS FROM tbl_name` instead of `mysql_list_fields()`.

You must free the result set with `mysql_free_result()`.

##### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

##### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

#### 21.8.6.43. `mysql_list_processes()`

```
MYSQL_RES *mysql_list_processes(MYSQL *mysql)
```

##### Description

Returns a result set describing the current server threads. This is the same kind of information as that reported by `mysqladmin processlist` or a `SHOW PROCESSLIST` query.

You must free the result set with `mysql_free_result()`.

##### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

##### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

#### 21.8.6.44. `mysql_list_tables()`

```
MYSQL_RES *mysql_list_tables(MYSQL *mysql, const char *wild)
```

##### Description

Returns a result set consisting of table names in the current database that match the simple regular expression specified by the `wild` parameter. `wild` may contain the wildcard characters “%” or “\_”, or may be a `NULL` pointer to match all tables. Calling `mysql_list_tables()` is similar to executing the query `SHOW TABLES [LIKE wild]`.

You must free the result set with `mysql_free_result()`.

##### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

##### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

#### 21.8.6.45. `mysql_more_results()`

```
my_bool mysql_more_results(MYSQL *mysql)
```

##### Description

This function is used when you execute multiple statements specified as a single statement string, or when you execute `CALL` statements, which can return multiple result sets.

`mysql_more_results()` true if more results exist from the currently executed statement, in which case the application must call `mysql_next_result()` to fetch the results.

##### Return Values

`TRUE` (1) if more results exist. `FALSE` (0) if no more results exist.

In most cases, you can call `mysql_next_result()` instead to test whether more results exist and initiate retrieval if so.



See [Section 21.8.16, “C API Support for Multiple Statement Execution”](#), and [Section 21.8.6.46, “mysql\\_next\\_result\(\)”](#).

## Errors

None.

### 21.8.6.46. `mysql_next_result()`

```
int mysql_next_result(MYSQL *mysql)
```

## Description

This function is used when you execute multiple statements specified as a single statement string, or when you use `CALL` statements to execute stored procedures, which can return multiple result sets.

`mysql_next_result()` reads the next statement result and returns a status to indicate whether more results exist. If `mysql_next_result()` returns an error, there are no more results.

Before each call to `mysql_next_result()`, you must call `mysql_free_result()` for the current statement if it is a statement that returned a result set (rather than just a result status).

After calling `mysql_next_result()` the state of the connection is as if you had called `mysql_real_query()` or `mysql_query()` for the next statement. This means that you can call `mysql_store_result()`, `mysql_warning_count()`, `mysql_affected_rows()`, and so forth.

If your program uses `CALL` statements to execute stored procedures, the `CLIENT_MULTI_RESULTS` flag must be enabled. This is because each `CALL` returns a result to indicate the call status, in addition to any result sets that might be returned by statements executed within the procedure. Because `CALL` can return multiple results, process them using a loop that calls `mysql_next_result()` to determine whether there are more results.

`CLIENT_MULTI_RESULTS` can be enabled when you call `mysql_real_connect()`, either explicitly by passing the `CLIENT_MULTI_RESULTS` flag itself, or implicitly by passing `CLIENT_MULTI_STATEMENTS` (which also enables `CLIENT_MULTI_RESULTS`). In MySQL 5.6, `CLIENT_MULTI_RESULTS` is enabled by default.

It is also possible to test whether there are more results by calling `mysql_more_results()`. However, this function does not change the connection state, so if it returns true, you must still call `mysql_next_result()` to advance to the next result.

For an example that shows how to use `mysql_next_result()`, see [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

## Return Values

Return Value	Description
0	Successful and there are more results
-1	Successful and there are no more results
>0	An error occurred

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order. For example, if you didn't call `mysql_use_result()` for a previous result set.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.47. `mysql_num_fields()`

```
unsigned int mysql_num_fields(MYSQL_RES *result)
```

To pass a `MYSQL*` argument instead, use `unsigned int mysql_field_count(MYSQL *mysql)`.

#### Description

Returns the number of columns in a result set.

Note that you can get the number of columns either from a pointer to a result set or to a connection handle. You would use the connection handle if `mysql_store_result()` or `mysql_use_result()` returned `NULL` (and thus you have no result set pointer). In this case, you can call `mysql_field_count()` to determine whether `mysql_store_result()` should have produced a nonempty result. This enables the client program to take proper action without knowing whether the query was a `SELECT` (or `SELECT`-like) statement. The example shown here illustrates how this may be done.

See [Section 21.8.14.1, "Why `mysql\_store\_result\(\)` Sometimes Returns `NULL` After `mysql\_query\(\)` Returns Success"](#).

#### Return Values

An unsigned integer representing the number of columns in a result set.

#### Errors

None.

#### Example

```
MYSQL_RES *result;
unsigned int num_fields;
unsigned int num_rows;

if (mysql_query(&mysql, query_string))
{
    // error
}
else // query succeeded, process any data returned by it
{
    result = mysql_store_result(&mysql);
    if (result) // there are rows
    {
        num_fields = mysql_num_fields(result);
        // retrieve rows, then call mysql_free_result(result)
    }
    else // mysql_store_result() returned nothing; should it have?
    {
        if (mysql_errno(&mysql))
        {
            fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
        }
        else if (mysql_field_count(&mysql) == 0)
        {
            // query does not return data
            // (it was not a SELECT)
            num_rows = mysql_affected_rows(&mysql);
        }
    }
}
```

```
}  
}
```

An alternative (if you know that your query should have returned a result set) is to replace the `mysql_errno(&mysql)` call with a check whether `mysql_field_count(&mysql)` returns 0. This happens only if something went wrong.

#### 21.8.6.48. `mysql_num_rows()`

```
my_ulonglong mysql_num_rows(MYSQL_RES *result)
```

##### Description

Returns the number of rows in the result set.

The use of `mysql_num_rows()` depends on whether you use `mysql_store_result()` or `mysql_use_result()` to return the result set. If you use `mysql_store_result()`, `mysql_num_rows()` may be called immediately. If you use `mysql_use_result()`, `mysql_num_rows()` does not return the correct value until all the rows in the result set have been retrieved.

`mysql_num_rows()` is intended for use with statements that return a result set, such as `SELECT`. For statements such as `INSERT`, `UPDATE`, or `DELETE`, the number of affected rows can be obtained with `mysql_affected_rows()`.

##### Return Values

The number of rows in the result set.

##### Errors

None.

#### 21.8.6.49. `mysql_options()`

```
int mysql_options(MYSQL *mysql, enum mysql_option option, const void *arg)
```

##### Description

Can be used to set extra connect options and affect behavior for a connection. This function may be called multiple times to set several options.

Call `mysql_options()` after `mysql_init()` and before `mysql_connect()` or `mysql_real_connect()`.

The `option` argument is the option that you want to set; the `arg` argument is the value for the option. If the option is an integer, specify a pointer to the value of the integer as the `arg` argument.

The following list describes the possible options, their effect, and how `arg` is used for each option. Several of the options apply only when the application is linked against the `libmysqld` embedded server library and are unused for applications linked against the `libmysqlclient` client library. For option descriptions that indicate `arg` is unused, its value is irrelevant; it is conventional to pass 0.

- `MYSQL_DEFAULT_AUTH` (argument type: `char *`)

The name of the authentication plugin to use.

- `MYSQL_ENABLE_CLEARTEXT_PLUGIN` (argument type: `my_bool *`)

Enable the `mysql_clear_password` cleartext authentication plugin. (See [Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”](#).) This option was added in MySQL 5.6.7.

- `MYSQL_INIT_COMMAND` (argument type: `char *`)

SQL statement to execute when connecting to the MySQL server. Automatically re-executed if reconnection occurs.

- `MYSQL_OPT_BIND` (argument: `char *`)

The network interface from which to connect to the server. This is used when the client host has multiple network interfaces. The argument is a host name or IP address (specified as a string). This option was added in MySQL 5.6.1.

- `MYSQL_OPT_CAN_HANDLE_EXPIRED_PASSWORDS` (argument type: `my_bool *`)

Indicate whether the client can handle expired passwords. For more information, see [Section 6.3.6, “Password Expiration and Sandbox Mode”](#). This option was added in MySQL 5.6.10.

- `MYSQL_OPT_COMPRESS` (argument: not used)

Use the compressed client/server protocol.

- `MYSQL_OPT_CONNECT_ATTR_DELETE` (argument types: `char *`)

Given a key name, this option deletes a key/value pair from the current set of connection attributes to be passed to the server at connect time. The argument is a pointer to a null-terminated string naming the key. Comparison of the key name with existing keys is case sensitive.

See also the description for the `MYSQL_OPT_CONNECT_ATTR_RESET` option, as well as the description for the `MYSQL_OPT_CONNECT_ATTR_ADD` option in the description of the `mysql_options4()` function. That function description also includes a usage example. These options were added in MySQL 5.6.6.

Connection attributes are exposed through the `session_connect_attrs` and `session_account_connect_attrs` Performance Schema tables. See [Section 20.9.7, “Performance Schema Connection Attribute Tables”](#).

- `MYSQL_OPT_CONNECT_ATTR_RESET` (argument not used)

This option resets (clears) the current set of connection attributes to be passed to the server at connect time.

See also the description for the `MYSQL_OPT_CONNECT_ATTR_DELETE` option, as well as the description for the `MYSQL_OPT_CONNECT_ATTR_ADD` option in the description of the `mysql_options4()` function. That function description also includes a usage example. These options were added in MySQL 5.6.6.

Connection attributes are exposed through the `session_connect_attrs` and `session_account_connect_attrs` Performance Schema tables. See [Section 20.9.7, “Performance Schema Connection Attribute Tables”](#).

- `MYSQL_OPT_CONNECT_TIMEOUT` (argument type: `unsigned int *`)

Connect timeout in seconds.

- `MYSQL_OPT_GUESS_CONNECTION` (argument: not used)

For an application linked against the `libmysqld` embedded server library, this enables the library to guess whether to use the embedded server or a remote server. “Guess” means that if the host name is set and is not `localhost`, it uses a remote server. This behavior is the default. `MYSQL_OPT_USE_EMBEDDED_CONNECTION` and `MYSQL_OPT_USE_REMOTE_CONNECTION` can be used to override it. This option is ignored for applications linked against the `libmysqlclient` client library.

- `MYSQL_OPT_LOCAL_INFILE` (argument type: optional pointer to `unsigned int`)

If no pointer is given or if pointer points to an `unsigned int` that has a nonzero value, the `LOAD LOCAL INFILE` statement is enabled.

- `MYSQL_OPT_NAMED_PIPE` (argument: not used)

Use named pipes to connect to a MySQL server on Windows, if the server permits named-pipe connections.

- `MYSQL_OPT_PROTOCOL` (argument type: `unsigned int *`)

Type of protocol to use. Specify one of the enum values of `mysql_protocol_type` defined in `mysql.h`.

- `MYSQL_OPT_READ_TIMEOUT` (argument type: `unsigned int *`)

The timeout in seconds for attempts to read from the server. Each attempt uses this timeout value and there are retries if necessary, so the total effective timeout value is three times the option value. You can set the value so that a lost connection can be detected earlier than the TCP/IP `Close_Wait_Timeout` value of 10 minutes.

Implementation of this timeout uses mechanisms that may not be available on all platforms. On such a platform, a client that issues a read call might under certain circumstances wait without timing out. For example, a client might not time out if the server is not responding because it is waiting for a “disk full” condition to clear.

- `MYSQL_OPT_RECONNECT` (argument type: `my_bool *`)

Enable or disable automatic reconnection to the server if the connection is found to have been lost. Reconnect is off by default; this option provides a way to set reconnection behavior explicitly.

- `MYSQL_OPT_SSL_CA` (argument type: `char *`)

The path to a file that contains a list of trusted SSL CAs. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_CAPATH` (argument type: `char *`)

The path to a directory that contains trusted SSL CA certificates in PEM format. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_CERT` (argument type: `char *`)

The name of the SSL certificate file to use for establishing a secure connection. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_CIPHER` (argument type: `char *`)

A list of permissible ciphers to use for SSL encryption. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_CRL` (argument type: `char *`)

The path to a file containing certificate revocation lists in PEM format. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_CRLPATH` (argument type: `char *`)

The path to a directory that contains files containing certificate revocation lists in PEM format. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_KEY` (argument type: `char *`)

The name of the SSL key file to use for establishing a secure connection. This option was added in MySQL 5.6.3.

- `MYSQL_OPT_SSL_VERIFY_SERVER_CERT` (argument type: `my_bool *`)

Enable or disable verification of the server's Common Name value in its certificate against the host name used when connecting to the server. The connection is rejected if there is a mismatch. This feature can be used to prevent man-in-the-middle attacks. Verification is disabled by default.

- `MYSQL_OPT_USE_EMBEDDED_CONNECTION` (argument: not used)

For an application linked against the `libmysqld` embedded server library, this forces the use of the embedded server for the connection. This option is ignored for applications linked against the `libmysqlclient` client library.

- `MYSQL_OPT_USE_REMOTE_CONNECTION` (argument: not used)

For an application linked against the `libmysqld` embedded server library, this forces the use of a remote server for the connection. This option is ignored for applications linked against the `libmysqlclient` client library.

- `MYSQL_OPT_USE_RESULT` (argument: not used)

This option is unused.

- `MYSQL_OPT_WRITE_TIMEOUT` (argument type: `unsigned int *`)

The timeout in seconds for attempts to write to the server. Each attempt uses this timeout value and there are `net_retry_count` [515] retries if necessary, so the total effective timeout value is `net_retry_count` [515] times the option value.

- `MYSQL_PLUGIN_DIR` (argument type: `char *`)

The directory in which to look for client plugins.

- `MYSQL_READ_DEFAULT_FILE` (argument type: `char *`)

Read options from the named option file instead of from `my.cnf`.

- `MYSQL_READ_DEFAULT_GROUP` (argument type: `char *`)

Read options from the named group from `my.cnf` or the file specified with `MYSQL_READ_DEFAULT_FILE`.

- `MYSQL_REPORT_DATA_TRUNCATION` (argument type: `my_bool *`)

Enable or disable reporting of data truncation errors for prepared statements using the `error` member of `MYSQL_BIND` structures. (Default: enabled.)

- `MYSQL_SECURE_AUTH` (argument type: `my_bool *`)

Whether to connect to a server that does not support the password hashing used in MySQL 4.1.1 and later. As of MySQL 5.6.7, this option is enabled by default.

- `MYSQL_SERVER_PUBLIC_KEY` (argument type: `char *`)

The path name to a file containing the server RSA public key. The file must be in PEM format. The public key is used for RSA encryption of the client password for connections to the server made using accounts that authenticate with the `sha256_password` plugin. This option is ignored for client accounts that do not authenticate with that plugin. It is also ignored if password encryption is not needed, as is the case when the client connects to the server using an SSL connection.

The server sends the public key to the client as needed, so it is not necessary to use this option for RSA password encryption to occur. It is more efficient to do so because then the server need not send the key.

For additional discussion regarding use of the `sha256_password` plugin, including how to get the RSA public key, see [Section 6.3.7.2, “The SHA-256 Authentication Plugin”](#).

This option was added in MySQL 5.6.6.

- `MYSQL_SET_CHARSET_DIR` (argument type: `char *`)

The path name to the directory that contains character set definition files.

- `MYSQL_SET_CHARSET_NAME` (argument type: `char *`)

The name of the character set to use as the default character set. The argument can be `MYSQL_AUTODETECT_CHARSET_NAME` to cause the character set to be autodetected based on the operating system setting (see [Section 10.1.4, “Connection Character Sets and Collations”](#)).

- `MYSQL_SET_CLIENT_IP` (argument type: `char *`)

For an application linked against the `libmysqld` embedded server library (when `libmysqld` is compiled with authentication support), this means that the user is considered to have connected from the specified IP address (specified as a string) for authentication purposes. This option is ignored for applications linked against the `libmysqlclient` client library.

- `MYSQL_SHARED_MEMORY_BASE_NAME` (argument type: `char *`)

The name of the shared-memory object for communication to the server on Windows, if the server supports shared-memory connections. Specify the same value as the `--shared-memory-base-name` [\[439\]](#) option used for the `mysqld` server you want to connect to.

The `client` group is always read if you use `MYSQL_READ_DEFAULT_FILE` or `MYSQL_READ_DEFAULT_GROUP`.

The specified group in the option file may contain the following options.

Option	Description
<code>character-sets-dir=path</code>	The directory where character sets are installed.
<code>compress</code>	Use the compressed client/server protocol.
<code>connect-timeout=seconds</code>	Connect timeout in seconds. On Linux this timeout is also used for waiting for the first answer from the server.
<code>database=db_name</code>	Connect to this database if no database was specified in the connect command.
<code>debug</code>	Debug options.
<code>default-character-set=charset_name</code>	The default character set to use.
<code>disable-local-infile</code>	Disable use of <code>LOAD DATA LOCAL</code> .
<code>enable-cleartext-plugin</code>	Enable the <code>mysql_clear_password</code> cleartext authentication plugin. Added in MySQL 5.6.7.
<code>host=host_name</code>	Default host name.
<code>init-command=stmt</code>	Statement to execute when connecting to MySQL server. Automatically re-executed if reconnection occurs.
<code>interactive-timeout=seconds</code>	Same as specifying <code>CLIENT_INTERACTIVE</code> to <code>mysql_real_connect()</code> . See <a href="#">Section 21.8.6.53, “mysql_real_connect()”</a> .
<code>local-infile[={0 1}]</code>	If no argument or nonzero argument, enable use of <code>LOAD DATA LOCAL</code> ; otherwise disable.

Option	Description
<code>max_allowed_packet=bytes</code>	Maximum size of packet that client can read from server.
<code>multi-queries, multi-results</code>	Enable multiple result sets from multiple-statement executions or stored procedures.
<code>multi-statements</code>	Enable the client to send multiple statements in a single string (separated by ";").
<code>password=password</code>	Default password.
<code>pipe</code>	Use named pipes to connect to a MySQL server on Windows.
<code>port=port_num</code>	Default port number.
<code>protocol={TCP   SOCKET   PIPE   MEMORY}</code>	The protocol to use when connecting to the server.
<code>return-found-rows</code>	Tell <code>mysql_info()</code> to return found rows instead of updated rows when using <code>UPDATE</code> .
<code>shared-memory-base-name=name</code>	Shared-memory name to use to connect to server.
<code>socket=path</code>	Default socket file.
<code>ssl-ca=file_name</code>	Certificate Authority file.
<code>ssl-capath=path</code>	Certificate Authority directory.
<code>ssl-cert=file_name</code>	Certificate file.
<code>ssl-cipher=cipher_list</code>	Permissible SSL ciphers.
<code>ssl-key=file_name</code>	Key file.
<code>timeout=seconds</code>	Like <code>connect-timeout</code> .
<code>user</code>	Default user.

`timeout` has been replaced by `connect-timeout`, but `timeout` is still supported in MySQL 5.6 for backward compatibility.

For more information about option files, see [Section 4.2.3.3, "Using Option Files"](#).

## Return Values

Zero for success. Nonzero if you specify an unknown option.

## Example

The following `mysql_options()` calls request the use of compression in the client/server protocol, cause options to be read from the `[odbc]` group of option files, and disable transaction autocommit mode:

```
MYSQL mysql;

mysql_init(&mysql);
mysql_options(&mysql, MYSQL_OPT_COMPRESS, 0);
mysql_options(&mysql, MYSQL_READ_DEFAULT_GROUP, "odbc");
mysql_options(&mysql, MYSQL_INIT_COMMAND, "SET autocommit=0");
if (!mysql_real_connect(&mysql, "host", "user", "passwd", "database", 0, NULL, 0))
{
    fprintf(stderr, "Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
```

This code requests that the client use the compressed client/server protocol and read the additional options from the `odbc` section in the `my.cnf` file.

### 21.8.6.50. `mysql_options4()`



```
int mysql_options4(MYSQL *mysql, enum mysql_option option, const void *arg1,
const void *arg2)
```

## Description

`mysql_options4()` is similar to `mysql_options()` but has an extra fourth argument so that two values can be passed for the option specified in the second argument. This function was added in MySQL 5.6.6.

The following list describes the permitted options, their effect, and how `arg1` and `arg2` are used.

- `MYSQL_OPT_CONNECT_ATTR_ADD` (argument types: `char *`, `char *`)

This option adds a key/value pair to the current set of connection attributes to be passed to the server at connect time. Both arguments are pointers to null-terminated strings. The first and second strings indicate the key and value, respectively. If the key already exists in the current set of connection attributes, the new value replaces the existing one. Comparison of the key name with existing keys is case sensitive.

Key names that begin with an underscore (`_`) are reserved for internal use and should not be used by application programs.

See also the descriptions for the `MYSQL_OPT_CONNECT_ATTR_RESET`, `MYSQL_OPT_CONNECT_ATTR_DELETE` options in the description of the `mysql_options()` function.

Connection attributes are exposed through the `session_connect_attrs` and `session_account_connect_attrs` Performance Schema tables. See [Section 20.9.7, "Performance Schema Connection Attribute Tables"](#).

## Return Values

Zero for success. Nonzero if you specify an unknown option.

## Example

This example demonstrates the calls that specify connection attributes:

```
MYSQL mysql;

mysql_init(&mysql);
mysql_options4(&mysql,MYSQL_OPT_CONNECT_ATTR_RESET, 0);
mysql_options4(&mysql,MYSQL_OPT_CONNECT_ATTR_ADD, "key1", "value1");
mysql_options4(&mysql,MYSQL_OPT_CONNECT_ATTR_ADD, "key2", "value2");
mysql_options4(&mysql,MYSQL_OPT_CONNECT_ATTR_ADD, "key3", "value3");
mysql_options4(&mysql,MYSQL_OPT_CONNECT_ATTR_DELETE, "key1");
if (!mysql_real_connect(&mysql,"host","user","passwd","database",0,NULL,0))
{
    fprintf(stderr, "Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
mysql_options(&mysql,MYSQL_OPT_CONNECT_ATTR_RESET, 0);
```

### 21.8.6.51. `mysql_ping()`

```
int mysql_ping(MYSQL *mysql)
```

## Description

Checks whether the connection to the server is working. If the connection has gone down and auto-reconnect is enabled an attempt to reconnect is made. If the connection is down and auto-reconnect is disabled, `mysql_ping()` returns an error.

Auto-reconnect is disabled by default. To enable it, call `mysql_options()` with the `MYSQL_OPT_RECONNECT` option. For details, see [Section 21.8.6.49](#), “`mysql_options()`”.

`mysql_ping()` can be used by clients that remain idle for a long while, to check whether the server has closed the connection and reconnect if necessary.

If `mysql_ping()` does cause a reconnect, there is no explicit indication of it. To determine whether a reconnect occurs, call `mysql_thread_id()` to get the original connection identifier before calling `mysql_ping()`, then call `mysql_thread_id()` again to see whether the identifier has changed.

If reconnect occurs, some characteristics of the connection will have been reset. For details about these characteristics, see [Section 21.8.15](#), “[Controlling Automatic Reconnection Behavior](#)”.

## Return Values

Zero if the connection to the server is active. Nonzero if an error occurred. A nonzero return does not indicate whether the MySQL server itself is down; the connection might be broken for other reasons such as network problems.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

### 21.8.6.52. `mysql_query()`

```
int mysql_query(MYSQL *mysql, const char *stmt_str)
```

## Description

Executes the SQL statement pointed to by the null-terminated string `stmt_str`. Normally, the string must consist of a single SQL statement without a terminating semicolon (“;”) or `\g`. If multiple-statement execution has been enabled, the string can contain several statements separated by semicolons. See [Section 21.8.16](#), “[C API Support for Multiple Statement Execution](#)”.

`mysql_query()` cannot be used for statements that contain binary data; you must use `mysql_real_query()` instead. (Binary data may contain the “\0” character, which `mysql_query()` interprets as the end of the statement string.)

If you want to know whether the statement returns a result set, you can use `mysql_field_count()` to check for this. See [Section 21.8.6.22](#), “[mysql\\_field\\_count\(\)](#)”.

## Return Values

Zero if the statement was successful. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- `CR_SERVER_LOST` [3719]

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR` [3718]

An unknown error occurred.

### 21.8.6.53. `mysql_real_connect()`

```
MYSQL *mysql_real_connect(MYSQL *mysql, const char *host, const char
*user, const char *passwd, const char *db, unsigned int port, const char
*unix_socket, unsigned long client_flag)
```

#### Description

`mysql_real_connect()` attempts to establish a connection to a MySQL database engine running on `host`. `mysql_real_connect()` must complete successfully before you can execute any other API functions that require a valid `MYSQL` connection handle structure.

The parameters are specified as follows:

- For the first parameter, specify the address of an existing `MYSQL` structure. Before calling `mysql_real_connect()`, call `mysql_init()` to initialize the `MYSQL` structure. You can change a lot of connect options with the `mysql_options()` call. See [Section 21.8.6.49](#), “`mysql_options()`”.
- The value of `host` may be either a host name or an IP address. If `host` is `NULL` or the string “localhost”, a connection to the local host is assumed. For Windows, the client connects using a shared-memory connection, if the server has shared-memory connections enabled. Otherwise, TCP/IP is used. For Unix, the client connects using a Unix socket file. For local connections, you can also influence the type of connection to use with the `MYSQL_OPT_PROTOCOL` or `MYSQL_OPT_NAMED_PIPE` options to `mysql_options()`. The type of connection must be supported by the server. For a `host` value of “.” on Windows, the client connects using a named pipe, if the server has named-pipe connections enabled. If named-pipe connections are not enabled, an error occurs.
- The `user` parameter contains the user's MySQL login ID. If `user` is `NULL` or the empty string “”, the current user is assumed. Under Unix, this is the current login name. Under Windows ODBC, the current user name must be specified explicitly. See the Connector/ODBC section of [Chapter 21](#), *Connectors and APIs*.
- The `passwd` parameter contains the password for `user`. If `passwd` is `NULL`, only entries in the `user` table for the user that have a blank (empty) password field are checked for a match. This enables the database administrator to set up the MySQL privilege system in such a way that users get different privileges depending on whether they have specified a password.



#### Note

Do not attempt to encrypt the password before calling `mysql_real_connect()`; password encryption is handled automatically by the client API.

- The `user` and `passwd` parameters use whatever character set has been configured for the `MYSQL` object. By default, this is `latin1`, but can be changed by calling `mysql_options(mysql, MYSQL_SET_CHARSET_NAME, "charset_name")` prior to connecting.
- `db` is the database name. If `db` is not `NULL`, the connection sets the default database to this value.

- If `port` is not 0, the value is used as the port number for the TCP/IP connection. Note that the `host` parameter determines the type of the connection.
- If `unix_socket` is not `NULL`, the string specifies the socket or named pipe to use. Note that the `host` parameter determines the type of the connection.
- The value of `client_flag` is usually 0, but can be set to a combination of the following flags to enable certain features.

Flag Name	Flag Description
<code>CAN_HANDLE_EXPIRED_PASSWORDS</code>	The client can handle expired passwords. For more information, see <a href="#">Section 6.3.6, “Password Expiration and Sandbox Mode”</a> . This option was added in MySQL 5.6.10.
<code>CLIENT_COMPRESS</code>	Use compression protocol.
<code>CLIENT_FOUND_ROWS</code>	Return the number of found (matched) rows, not the number of changed rows.
<code>CLIENT_IGNORE_SIGPIPE</code>	Prevents the client library from installing a <code>SIGPIPE</code> signal handler. This can be used to avoid conflicts with a handler that the application has already installed.
<code>CLIENT_IGNORE_SPACE</code>	Permit spaces after function names. Makes all functions names reserved words.
<code>CLIENT_INTERACTIVE</code>	Permit <code>interactive_timeout</code> [490] seconds (instead of <code>wait_timeout</code> [563] seconds) of inactivity before closing the connection. The client's session <code>wait_timeout</code> [563] variable is set to the value of the session <code>interactive_timeout</code> [490] variable.
<code>CLIENT_LOCAL_FILES</code>	Enable <code>LOAD DATA LOCAL</code> handling.
<code>CLIENT_MULTI_RESULTS</code>	Tell the server that the client can handle multiple result sets from multiple-statement executions or stored procedures. This flag is automatically enabled if <code>CLIENT_MULTI_STATEMENTS</code> is enabled. See the note following this table for more information about this flag.
<code>CLIENT_MULTI_STATEMENTS</code>	Tell the server that the client may send multiple statements in a single string (separated by “;”). If this flag is not set, multiple-statement execution is disabled. See the note following this table for more information about this flag.
<code>CLIENT_NO_SCHEMA</code>	Do not permit the <code>db_name.tbl_name.col_name</code> syntax. This is for ODBC. It causes the parser to generate an error if you use that syntax, which is useful for trapping bugs in some ODBC programs.
<code>CLIENT_ODBC</code>	Unused.
<code>CLIENT_SSL</code>	Use SSL (encrypted protocol). Do not set this option within an application program; it is set internally in the client library. Instead, use <code>mysql_ssl_set()</code> before calling <code>mysql_real_connect()</code> .
<code>CLIENT_REMEMBER_OPTIONS</code>	Remember options specified by calls to <code>mysql_options()</code> . Without this option, if <code>mysql_real_connect()</code> fails, you must repeat the <code>mysql_options()</code> calls before trying to connect again. With this option, the <code>mysql_options()</code> calls need not be repeated.

If your program uses `CALL` statements to execute stored procedures, the `CLIENT_MULTI_RESULTS` flag must be enabled. This is because each `CALL` returns a result to indicate the call status, in addition to any result sets that might be returned by statements executed within the procedure. Because `CALL`

can return multiple results, process them using a loop that calls `mysql_next_result()` to determine whether there are more results.

`CLIENT_MULTI_RESULTS` can be enabled when you call `mysql_real_connect()`, either explicitly by passing the `CLIENT_MULTI_RESULTS` flag itself, or implicitly by passing `CLIENT_MULTI_STATEMENTS` (which also enables `CLIENT_MULTI_RESULTS`). In MySQL 5.6, `CLIENT_MULTI_RESULTS` is enabled by default.

If you enable `CLIENT_MULTI_STATEMENTS` or `CLIENT_MULTI_RESULTS`, process the result for every call to `mysql_query()` or `mysql_real_query()` by using a loop that calls `mysql_next_result()` to determine whether there are more results. For an example, see [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

For some parameters, it is possible to have the value taken from an option file rather than from an explicit value in the `mysql_real_connect()` call. To do this, call `mysql_options()` with the `MYSQL_READ_DEFAULT_FILE` or `MYSQL_READ_DEFAULT_GROUP` option before calling `mysql_real_connect()`. Then, in the `mysql_real_connect()` call, specify the “no-value” value for each parameter to be read from an option file:

- For `host`, specify a value of `NULL` or the empty string (`" "`).
- For `user`, specify a value of `NULL` or the empty string.
- For `passwd`, specify a value of `NULL`. (For the password, a value of the empty string in the `mysql_real_connect()` call cannot be overridden in an option file, because the empty string indicates explicitly that the MySQL account must have an empty password.)
- For `db`, specify a value of `NULL` or the empty string.
- For `port`, specify a value of 0.
- For `unix_socket`, specify a value of `NULL`.

If no value is found in an option file for a parameter, its default value is used as indicated in the descriptions given earlier in this section.

## Return Values

A `MYSQL*` connection handle if the connection was successful, `NULL` if the connection was unsuccessful. For a successful connection, the return value is the same as the value of the first parameter.

## Errors

- `CR_CONN_HOST_ERROR` [3718]  
Failed to connect to the MySQL server.
- `CR_CONNECTION_ERROR` [3718]  
Failed to connect to the local MySQL server.
- `CR_IPSOCK_ERROR` [3718]  
Failed to create an IP socket.
- `CR_OUT_OF_MEMORY` [3718]  
Out of memory.
- `CR_SOCKET_CREATE_ERROR` [3718]  
Failed to create a Unix socket.

- [CR\\_UNKNOWN\\_HOST \[3718\]](#)

Failed to find the IP address for the host name.

- [CR\\_VERSION\\_ERROR \[3718\]](#)

A protocol mismatch resulted from attempting to connect to a server with a client library that uses a different protocol version.

- [CR\\_NAMEDPIPEOPEN\\_ERROR \[3719\]](#)

Failed to create a named pipe on Windows.

- [CR\\_NAMEDPIPEWAIT\\_ERROR \[3719\]](#)

Failed to wait for a named pipe on Windows.

- [CR\\_NAMEDPIPESETSTATE\\_ERROR \[3719\]](#)

Failed to get a pipe handler on Windows.

- [CR\\_SERVER\\_LOST \[3719\]](#)

If [connect\\_timeout \[472\]](#) > 0 and it took longer than [connect\\_timeout \[472\]](#) seconds to connect to the server or if the server died while executing the `init-command`.

- [CR\\_ALREADY\\_CONNECTED \[3721\]](#)

The `MYSQL` connection handle is already connected.

## Example

```
MYSQL mysql;

mysql_init(&mysql);
mysql_options(&mysql, MYSQL_READ_DEFAULT_GROUP, "your_prog_name");
if (!mysql_real_connect(&mysql, "host", "user", "passwd", "database", 0, NULL, 0))
{
    fprintf(stderr, "Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
```

By using `mysql_options()` the MySQL library reads the `[client]` and `[your_prog_name]` sections in the `my.cnf` file which ensures that your program works, even if someone has set up MySQL in some nonstandard way.

Note that upon connection, `mysql_real_connect()` sets the `reconnect` flag (part of the `MYSQL` structure) to a value of `1` in versions of the API older than 5.0.3, or `0` in newer versions. A value of `1` for this flag indicates that if a statement cannot be performed because of a lost connection, to try reconnecting to the server before giving up. You can use the `MYSQL_OPT_RECONNECT` option to `mysql_options()` to control reconnection behavior.

### 21.8.6.54. `mysql_real_escape_string()`

```
unsigned long mysql_real_escape_string(MYSQL *mysql, char *to, const char
*from, unsigned long length)
```

Note that `mysql` must be a valid, open connection. This is needed because the escaping depends on the character set in use by the server.

## Description

This function is used to create a legal SQL string that you can use in an SQL statement. See [Section 9.1.1, "String Literals"](#).

The string in `from` is encoded to an escaped SQL string, taking into account the current character set of the connection. The result is placed in `to` and a terminating null byte is appended. Characters encoded are “\”, “,”, “'”, `NUL` (ASCII 0), “\n”, “\r”, and Control+Z. Strictly speaking, MySQL requires only that backslash and the quote character used to quote the string in the query be escaped. `mysql_real_escape_string()` quotes the other characters to make them easier to read in log files. For comparison, see the quoting rules for literal strings and the `QUOTE()` [1183] SQL function in [Section 9.1.1, “String Literals”](#), and [Section 12.5, “String Functions”](#).

The string pointed to by `from` must be `length` bytes long. You must allocate the `to` buffer to be at least `length*2+1` bytes long. (In the worst case, each character may need to be encoded as using two bytes, and you need room for the terminating null byte.) When `mysql_real_escape_string()` returns, the contents of `to` is a null-terminated string. The return value is the length of the encoded string, not including the terminating null character.

If you need to change the character set of the connection, use the `mysql_set_character_set()` function rather than executing a `SET NAMES` (or `SET CHARACTER SET`) statement. `mysql_set_character_set()` works like `SET NAMES` but also affects the character set used by `mysql_real_escape_string()`, which `SET NAMES` does not.

## Example

```
char query[1000],*end;

end = strmov(query,"INSERT INTO test_table values(");
*end++ = '\';
end += mysql_real_escape_string(&mysql, end,"What is this",12);
*end++ = '\';
*end++ = ',';
*end++ = '\';
end += mysql_real_escape_string(&mysql, end,"binary data: \0\r\n",16);
*end++ = '\';
*end++ = ')';

if (mysql_real_query(&mysql,query,(unsigned int) (end - query)))
{
    fprintf(stderr, "Failed to insert row, Error: %s\n",
            mysql_error(&mysql));
}
```

The `strmov()` function used in the example is included in the `libmysqlclient` library and works like `strcpy()` but returns a pointer to the terminating null of the first parameter.

## Return Values

The length of the value placed into `to`, not including the terminating null character.

## Errors

None.

### 21.8.6.55. `mysql_real_query()`

```
int mysql_real_query(MYSQL *mysql, const char *stmt_str, unsigned long
length)
```

## Description

Executes the SQL statement pointed to by `stmt_str`, a string `length` bytes long. Normally, the string must consist of a single SQL statement without a terminating semicolon (“;”) or `\g`. If multiple-statement execution has been enabled, the string can contain several statements separated by semicolons. See [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

`mysql_query()` cannot be used for statements that contain binary data; you must use `mysql_real_query()` instead. (Binary data may contain the “\0” character, which `mysql_query()`

interprets as the end of the statement string.) In addition, `mysql_real_query()` is faster than `mysql_query()` because it does not call `strlen()` on the statement string.

If you want to know whether the statement returns a result set, you can use `mysql_field_count()` to check for this. See [Section 21.8.6.22](#), “`mysql_field_count()`”.

## Return Values

Zero if the statement was successful. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

### 21.8.6.56. `mysql_refresh()`

```
int mysql_refresh(MYSQL *mysql, unsigned int options)
```

## Description

This function flushes tables or caches, or resets replication server information. The connected user must have the [RELOAD \[755\]](#) privilege.

The `options` argument is a bit mask composed from any combination of the following values. Multiple values can be OR'ed together to perform multiple operations with a single call.

- [REFRESH\\_GRANT](#)  
Refresh the grant tables, like `FLUSH PRIVILEGES`.
- [REFRESH\\_LOG](#)  
Flush the logs, like `FLUSH LOGS`.
- [REFRESH\\_TABLES](#)  
Flush the table cache, like `FLUSH TABLES`.
- [REFRESH\\_HOSTS](#)  
Flush the host cache, like `FLUSH HOSTS`.
- [REFRESH\\_STATUS](#)  
Reset status variables, like `FLUSH STATUS`.
- [REFRESH\\_THREADS](#)  
Flush the thread cache.



- [REFRESH\\_SLAVE](#)

On a slave replication server, reset the master server information and restart the slave, like [RESET SLAVE](#).

- [REFRESH\\_MASTER](#)

On a master replication server, remove the binary log files listed in the binary log index and truncate the index file, like [RESET MASTER](#).

## Return Values

Zero for success. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.57. [mysql\\_reload\(\)](#)

```
int mysql_reload(MYSQL *mysql)
```

## Description

Asks the MySQL server to reload the grant tables. The connected user must have the [RELOAD \[755\]](#) privilege.

This function is deprecated. It is preferable to use [mysql\\_query\(\)](#) to issue an SQL [FLUSH PRIVILEGES](#) statement instead.

## Return Values

Zero for success. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.58. `mysql_rollback()`

```
my_bool mysql_rollback(MYSQL *mysql)
```

#### Description

Rolls back the current transaction.

The action of this function is subject to the value of the `completion_type` [470] system variable. In particular, if the value of `completion_type` [470] is `RELEASE` (or 2), the server performs a release after terminating a transaction and closes the client connection. Call `mysql_close()` from the client program to close the connection from the client side.

#### Return Values

Zero if successful. Nonzero if an error occurred.

#### Errors

None.

### 21.8.6.59. `mysql_row_seek()`

```
MYSQL_ROW_OFFSET mysql_row_seek(MYSQL_RES *result, MYSQL_ROW_OFFSET offset)
```

#### Description

Sets the row cursor to an arbitrary row in a query result set. The `offset` value is a row offset, typically a value returned from `mysql_row_tell()` or from `mysql_row_seek()`. This value is not a row number; to seek to a row within a result set by number, use `mysql_data_seek()` instead.

This function requires that the result set structure contains the entire result of the query, so `mysql_row_seek()` may be used only in conjunction with `mysql_store_result()`, not with `mysql_use_result()`.

#### Return Values

The previous value of the row cursor. This value may be passed to a subsequent call to `mysql_row_seek()`.

#### Errors

None.

### 21.8.6.60. `mysql_row_tell()`

```
MYSQL_ROW_OFFSET mysql_row_tell(MYSQL_RES *result)
```

#### Description

Returns the current position of the row cursor for the last `mysql_fetch_row()`. This value can be used as an argument to `mysql_row_seek()`.

Use `mysql_row_tell()` only after `mysql_store_result()`, not after `mysql_use_result()`.

#### Return Values

The current offset of the row cursor.

## Errors

None.

### 21.8.6.61. `mysql_select_db()`

```
int mysql_select_db(MYSQL *mysql, const char *db)
```

#### Description

Causes the database specified by `db` to become the default (current) database on the connection specified by `mysql`. In subsequent queries, this database is the default for table references that do not include an explicit database specifier.

`mysql_select_db()` fails unless the connected user can be authenticated as having permission to use the database.

#### Return Values

Zero for success. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

### 21.8.6.62. `mysql_set_character_set()`

```
int mysql_set_character_set(MYSQL *mysql, const char *csname)
```

#### Description

This function is used to set the default character set for the current connection. The string `csname` specifies a valid character set name. The connection collation becomes the default collation of the character set. This function works like the `SET NAMES` statement, but also sets the value of `mysql->charset`, and thus affects the character set used by `mysql_real_escape_string()`

#### Return Values

Zero for success. Nonzero if an error occurred.

## Example

```
MYSQL mysql;

mysql_init(&mysql);
if (!mysql_real_connect(&mysql, "host", "user", "passwd", "database", 0, NULL, 0))
{
    fprintf(stderr, "Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
```

```

}

if (!mysql_set_character_set(&mysql, "utf8"))
{
    printf("New client character set: %s\n",
        mysql_character_set_name(&mysql));
}

```

### 21.8.6.63. `mysql_set_local_infile_default()`

```
void mysql_set_local_infile_default(MYSQL *mysql);
```

#### Description

Sets the `LOAD LOCAL DATA INFILE` handler callback functions to the defaults used internally by the C client library. The library calls this function automatically if `mysql_set_local_infile_handler()` has not been called or does not supply valid functions for each of its callbacks.

#### Return Values

None.

#### Errors

None.

### 21.8.6.64. `mysql_set_local_infile_handler()`

```
void mysql_set_local_infile_handler(MYSQL *mysql, int (*local_infile_init)
(void **, const char *, void *), int (*local_infile_read)(void *, char *,
unsigned int), void (*local_infile_end)(void *), int (*local_infile_error)
(void *, char*, unsigned int), void *userdata);
```

#### Description

This function installs callbacks to be used during the execution of `LOAD DATA LOCAL INFILE` statements. It enables application programs to exert control over local (client-side) data file reading. The arguments are the connection handler, a set of pointers to callback functions, and a pointer to a data area that the callbacks can use to share information.

To use `mysql_set_local_infile_handler()`, you must write the following callback functions:

```
int
local_infile_init(void **ptr, const char *filename, void *userdata);
```

The initialization function. This is called once to do any setup necessary, open the data file, allocate data structures, and so forth. The first `void**` argument is a pointer to a pointer. You can set the pointer (that is, `*ptr`) to a value that will be passed to each of the other callbacks (as a `void*`). The callbacks can use this pointed-to value to maintain state information. The `userdata` argument is the same value that is passed to `mysql_set_local_infile_handler()`.

Make the initialization function return zero for success, nonzero for an error.

```
int
local_infile_read(void *ptr, char *buf, unsigned int buf_len);
```

The data-reading function. This is called repeatedly to read the data file. `buf` points to the buffer where the read data is stored, and `buf_len` is the maximum number of bytes that the callback can read and store in the buffer. (It can read fewer bytes, but should not read more.)

The return value is the number of bytes read, or zero when no more data could be read (this indicates EOF). Return a value less than zero if an error occurs.

```
void
local_infile_end(void *ptr)
```

The termination function. This is called once after `local_infile_read()` has returned zero (EOF) or an error. Within this function, deallocate any memory allocated by `local_infile_init()` and perform any other cleanup necessary. It is invoked even if the initialization function returns an error.

```
int
local_infile_error(void *ptr,
                  char *error_msg,
                  unsigned int error_msg_len);
```

The error-handling function. This is called to get a textual error message to return to the user in case any of your other functions returns an error. `error_msg` points to the buffer into which the message is written, and `error_msg_len` is the length of the buffer. Write the message as a null-terminated string, at most `error_msg_len-1` bytes long.

The return value is the error number.

Typically, the other callbacks store the error message in the data structure pointed to by `ptr`, so that `local_infile_error()` can copy the message from there into `error_msg`.

After calling `mysql_set_local_infile_handler()` in your C code and passing pointers to your callback functions, you can then issue a `LOAD DATA LOCAL INFILE` statement (for example, by using `mysql_query()`). The client library automatically invokes your callbacks. The file name specified in `LOAD DATA LOCAL INFILE` will be passed as the second parameter to the `local_infile_init()` callback.

## Return Values

None.

## Errors

None.

### 21.8.6.65. `mysql_set_server_option()`

```
int mysql_set_server_option(MYSQL *mysql, enum enum_mysql_set_option option)
```

## Description

Enables or disables an option for the connection. `option` can have one of the following values.

Option	Description
<code>MYSQL_OPTION_MULTI_STATEMENTS_ON</code>	Enable multiple-statement support
<code>MYSQL_OPTION_MULTI_STATEMENTS_OFF</code>	Disable multiple-statement support

If you enable multiple-statement support, you should retrieve results from calls to `mysql_query()` or `mysql_real_query()` by using a loop that calls `mysql_next_result()` to determine whether there are more results. For an example, see [Section 21.8.16, “C API Support for Multiple Statement Execution”](#).

Enabling multiple-statement support with `MYSQL_OPTION_MULTI_STATEMENTS_ON` does not have quite the same effect as enabling it by passing the `CLIENT_MULTI_STATEMENTS` flag to `mysql_real_connect()`: `CLIENT_MULTI_STATEMENTS` also enables `CLIENT_MULTI_RESULTS`. If you are using the `CALL SQL` statement in your programs, multiple-result support must be enabled; this means that `MYSQL_OPTION_MULTI_STATEMENTS_ON` by itself is insufficient to permit the use of `CALL`.

## Return Values

Zero for success. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [ER\\_UNKNOWN\\_COM\\_ERROR \[3663\]](#)

The server didn't support `mysql_set_server_option()` (which is the case that the server is older than 4.1.1) or the server didn't support the option one tried to set.

### 21.8.6.66. `mysql_shutdown()`

```
int mysql_shutdown(MYSQL *mysql, enum mysql_enum_shutdown_level
shutdown_level)
```

## Description

Asks the database server to shut down. The connected user must have the [SHUTDOWN \[756\]](#) privilege. MySQL 5.6 servers support only one type of shutdown; `shutdown_level` must be equal to `SHUTDOWN_DEFAULT`. Additional shutdown levels are planned to make it possible to choose the desired level. Dynamically linked executables which have been compiled with older versions of the `libmysqlclient` headers and call `mysql_shutdown()` need to be used with the old `libmysqlclient` dynamic library.

The shutdown process is described in [Section 5.1.12, "The Shutdown Process"](#).

## Return Values

Zero for success. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.67. `mysql_sqlstate()`

```
const char *mysql_sqlstate(MYSQL *mysql)
```

## Description

Returns a null-terminated string containing the SQLSTATE error code for the most recently executed SQL statement. The error code consists of five characters. '00000' means "no error." The values are specified by ANSI SQL and ODBC. For a list of possible values, see [Appendix C, Errors, Error Codes, and Common Problems](#).

SQLSTATE values returned by `mysql_sqlstate()` differ from MySQL-specific error numbers returned by `mysql_errno()`. For example, the `mysql` client program displays errors using the following format, where 1146 is the `mysql_errno()` value and '42S02' is the corresponding `mysql_sqlstate()` value:

```
shell> SELECT * FROM no_such_table;
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```

Not all MySQL error numbers are mapped to SQLSTATE error codes. The value 'HY000' (general error) is used for unmapped error numbers.

If you call `mysql_sqlstate()` after `mysql_real_connect()` fails, `mysql_sqlstate()` might not return a useful value. For example, this happens if a host is blocked by the server and the connection is closed without any SQLSTATE value being sent to the client.

## Return Values

A null-terminated character string containing the SQLSTATE error code.

## See Also

See [Section 21.8.6.14, "mysql\\_errno\(\)"](#), [Section 21.8.6.15, "mysql\\_error\(\)"](#), and [Section 21.8.10.27, "mysql\\_stmt\\_sqlstate\(\)"](#).

### 21.8.6.68. `mysql_ssl_set()`

```
my_bool mysql_ssl_set(MYSQL *mysql, const char *key, const char *cert, const
char *ca, const char *capath, const char *cipher)
```

## Description

`mysql_ssl_set()` is used for establishing secure connections using SSL. It must be called before `mysql_real_connect()`.

`mysql_ssl_set()` does nothing unless SSL support is enabled in the client library.

`mysql` is the connection handler returned from `mysql_init()`. The other parameters are specified as follows:

- `key` is the path name to the key file.
- `cert` is the path name to the certificate file.
- `ca` is the path name to the certificate authority file.
- `capath` is the path name to a directory that contains trusted SSL CA certificates in PEM format.
- `cipher` is a list of permissible ciphers to use for SSL encryption.

Any unused SSL parameters may be given as `NULL`.

## Return Values

This function always returns 0. If SSL setup is incorrect, `mysql_real_connect()` returns an error when you attempt to connect.

### 21.8.6.69. `mysql_stat()`

```
const char *mysql_stat(MYSQL *mysql)
```

#### Description

Returns a character string containing information similar to that provided by the `mysqladmin status` command. This includes uptime in seconds and the number of running threads, questions, reloads, and open tables.

#### Return Values

A character string describing the server status. `NULL` if an error occurred.

#### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

### 21.8.6.70. `mysql_store_result()`

```
MYSQL_RES *mysql_store_result(MYSQL *mysql)
```

#### Description

After invoking `mysql_query()` or `mysql_real_query()`, you must call `mysql_store_result()` or `mysql_use_result()` for every statement that successfully produces a result set (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`, `CHECK TABLE`, and so forth). You must also call `mysql_free_result()` after you are done with the result set.

You don't have to call `mysql_store_result()` or `mysql_use_result()` for other statements, but it does not do any harm or cause any notable performance degradation if you call `mysql_store_result()` in all cases. You can detect whether the statement has a result set by checking whether `mysql_store_result()` returns a nonzero value (more about this later).

If you enable multiple-statement support, you should retrieve results from calls to `mysql_query()` or `mysql_real_query()` by using a loop that calls `mysql_next_result()` to determine whether there are more results. For an example, see [Section 21.8.16, "C API Support for Multiple Statement Execution"](#).

If you want to know whether a statement should return a result set, you can use `mysql_field_count()` to check for this. See [Section 21.8.6.22, "mysql\\_field\\_count\(\)"](#).

`mysql_store_result()` reads the entire result of a query to the client, allocates a `MYSQL_RES` structure, and places the result into this structure.

`mysql_store_result()` returns a null pointer if the statement didn't return a result set (for example, if it was an `INSERT` statement).



`mysql_store_result()` also returns a null pointer if reading of the result set failed. You can check whether an error occurred by checking whether `mysql_error()` returns a nonempty string, `mysql_errno()` returns nonzero, or `mysql_field_count()` returns zero.

An empty result set is returned if there are no rows returned. (An empty result set differs from a null pointer as a return value.)

After you have called `mysql_store_result()` and gotten back a result that isn't a null pointer, you can call `mysql_num_rows()` to find out how many rows are in the result set.

You can call `mysql_fetch_row()` to fetch rows from the result set, or `mysql_row_seek()` and `mysql_row_tell()` to obtain or set the current row position within the result set.

See [Section 21.8.14.1, “Why `mysql\_store\_result\(\)` Sometimes Returns `NULL` After `mysql\_query\(\)` Returns Success”](#).

## Return Values

A `MYSQL_RES` result structure with the results. `NULL` (0) if an error occurred.

## Errors

`mysql_store_result()` resets `mysql_error()` and `mysql_errno()` if it succeeds.

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)

Out of memory.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.71. `mysql_thread_id()`

```
unsigned long mysql_thread_id(MYSQL *mysql)
```

## Description

Returns the thread ID of the current connection. This value can be used as an argument to `mysql_kill()` to kill the thread.

If the connection is lost and you reconnect with `mysql_ping()`, the thread ID changes. This means you should not get the thread ID and store it for later. You should get it when you need it.



### Note

This function does not work correctly if thread IDs become larger than 32 bits, which can occur on some systems. To avoid problems with `mysql_thread_id()`, do not use it. To get the connection ID, execute a `SELECT CONNECTION_ID()` query and retrieve the result.

## Return Values

The thread ID of the current connection.

## Errors

None.

### 21.8.6.72. `mysql_use_result()`

```
MYSQL_RES *mysql_use_result(MYSQL *mysql)
```

## Description

After invoking `mysql_query()` or `mysql_real_query()`, you must call `mysql_store_result()` or `mysql_use_result()` for every statement that successfully produces a result set (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`, `CHECK TABLE`, and so forth). You must also call `mysql_free_result()` after you are done with the result set.

`mysql_use_result()` initiates a result set retrieval but does not actually read the result set into the client like `mysql_store_result()` does. Instead, each row must be retrieved individually by making calls to `mysql_fetch_row()`. This reads the result of a query directly from the server without storing it in a temporary table or local buffer, which is somewhat faster and uses much less memory than `mysql_store_result()`. The client allocates memory only for the current row and a communication buffer that may grow up to `max_allowed_packet` [502] bytes.

On the other hand, you shouldn't use `mysql_use_result()` if you are doing a lot of processing for each row on the client side, or if the output is sent to a screen on which the user may type a `^S` (stop scroll). This ties up the server and prevent other threads from updating any tables from which the data is being fetched.

When using `mysql_use_result()`, you must execute `mysql_fetch_row()` until a `NULL` value is returned, otherwise, the unfetched rows are returned as part of the result set for your next query. The C API gives the error `Commands out of sync; you can't run this command now` if you forget to do this!

You may not use `mysql_data_seek()`, `mysql_row_seek()`, `mysql_row_tell()`, `mysql_num_rows()`, or `mysql_affected_rows()` with a result returned from `mysql_use_result()`, nor may you issue other queries until `mysql_use_result()` has finished. (However, after you have fetched all the rows, `mysql_num_rows()` accurately returns the number of rows fetched.)

You must call `mysql_free_result()` once you are done with the result set.

When using the `libmysqld` embedded server, the memory benefits are essentially lost because memory usage incrementally increases with each row retrieved until `mysql_free_result()` is called.

## Return Values

A `MYSQL_RES` result structure. `NULL` if an error occurred.

## Errors

`mysql_use_result()` resets `mysql_error()` and `mysql_errno()` if it succeeds.

- `CR_COMMANDS_OUT_OF_SYNC` [3719]  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY` [3718]  
Out of memory.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.6.73. `mysql_warning_count()`

```
unsigned int mysql_warning_count(MYSQL *mysql)
```

#### Description

Returns the number of errors, warnings, and notes generated during execution of the previous SQL statement.

#### Return Values

The warning count.

#### Errors

None.

## 21.8.7. C API Prepared Statements

The MySQL client/server protocol provides for the use of prepared statements. This capability uses the `MYSQL_STMT` statement handler data structure returned by the `mysql_stmt_init()` initialization function. Prepared execution is an efficient way to execute a statement more than once. The statement is first parsed to prepare it for execution. Then it is executed one or more times at a later time, using the statement handle returned by the initialization function.

Prepared execution is faster than direct execution for statements executed more than once, primarily because the query is parsed only once. In the case of direct execution, the query is parsed every time it is executed. Prepared execution also can provide a reduction of network traffic because for each execution of the prepared statement, it is necessary only to send the data for the parameters.

Prepared statements might not provide a performance increase in some situations. For best results, test your application both with prepared and nonprepared statements and choose whichever yields best performance.

Another advantage of prepared statements is that it uses a binary protocol that makes data transfer between client and server more efficient.

For a list of SQL statements that can be used as prepared statements, see [Section 13.5, “SQL Syntax for Prepared Statements”](#).

Metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

## 21.8.8. C API Prepared Statement Data Structures

Prepared statements use several data structures:

- To obtain a statement handle, pass a `MYSQL` connection handler to `mysql_stmt_init()`, which returns a pointer to a `MYSQL_STMT` data structure. This structure is used for further operations with

the statement. To specify the statement to prepare, pass the `MYSQL_STMT` pointer and the statement string to `mysql_stmt_prepare()`.

- To provide input parameters for a prepared statement, set up `MYSQL_BIND` structures and pass them to `mysql_stmt_bind_param()`. To receive output column values, set up `MYSQL_BIND` structures and pass them to `mysql_stmt_bind_result()`.
- The `MYSQL_TIME` structure is used to transfer temporal data in both directions.

The following discussion describes the prepared statement data types in detail. For examples that show how to use them, see [Section 21.8.10.10](#), “`mysql_stmt_execute()`”, and [Section 21.8.10.11](#), “`mysql_stmt_fetch()`”.

- `MYSQL_STMT`

This structure is a handle for a prepared statement. A handle is created by calling `mysql_stmt_init()`, which returns a pointer to a `MYSQL_STMT`. The handle is used for all subsequent operations with the statement until you close it with `mysql_stmt_close()`, at which point the handle becomes invalid.

The `MYSQL_STMT` structure has no members intended for application use. Applications should not try to copy a `MYSQL_STMT` structure. There is no guarantee that such a copy will be usable.

Multiple statement handles can be associated with a single connection. The limit on the number of handles depends on the available system resources.

- `MYSQL_BIND`

This structure is used both for statement input (data values sent to the server) and output (result values returned from the server):

- For input, use `MYSQL_BIND` structures with `mysql_stmt_bind_param()` to bind parameter data values to buffers for use by `mysql_stmt_execute()`.
- For output, use `MYSQL_BIND` structures with `mysql_stmt_bind_result()` to bind buffers to result set columns, for use in fetching rows with `mysql_stmt_fetch()`.

To use a `MYSQL_BIND` structure, zero its contents to initialize it, then set its members appropriately. For example, to declare and initialize an array of three `MYSQL_BIND` structures, use this code:

```
MYSQL_BIND bind[3];
memset(bind, 0, sizeof(bind));
```

The `MYSQL_BIND` structure contains the following members for use by application programs. For several of the members, the manner of use depends on whether the structure is used for input or output.

- `enum enum_field_types buffer_type`

The type of the buffer. This member indicates the data type of the C language variable bound to a statement parameter or result set column. For input, `buffer_type` indicates the type of the variable containing the value to be sent to the server. For output, it indicates the type of the variable into which a value received from the server should be stored. For permissible `buffer_type` values, see [Section 21.8.8.1](#), “C API Prepared Statement Type Codes”.

- `void *buffer`

A pointer to the buffer to be used for data transfer. This is the address of a C language variable.

For input, `buffer` is a pointer to the variable in which you store the data value for a statement parameter. When you call `mysql_stmt_execute()`, MySQL use the value stored in the variable

in place of the corresponding parameter marker in the statement (specified with `?` in the statement string).

For output, `buffer` is a pointer to the variable in which to return a result set column value. When you call `mysql_stmt_fetch()`, MySQL stores a column value from the current row of the result set in this variable. You can access the value when the call returns.

To minimize the need for MySQL to perform type conversions between C language values on the client side and SQL values on the server side, use C variables that have types similar to those of the corresponding SQL values:

- For numeric data types, `buffer` should point to a variable of the proper numeric C type. For integer variables (which can be `char` for single-byte values or an integer type for larger values), you should also indicate whether the variable has the `unsigned` attribute by setting the `is_unsigned` member, described later.
- For character (nonbinary) and binary string data types, `buffer` should point to a character buffer.
- For date and time data types, `buffer` should point to a `MYSQL_TIME` structure.

For guidelines about mapping between C types and SQL types and notes about type conversions, see [Section 21.8.8.1, “C API Prepared Statement Type Codes”](#), and [Section 21.8.8.2, “C API Prepared Statement Type Conversions”](#).

- `unsigned long buffer_length`

The actual size of `*buffer` in bytes. This indicates the maximum amount of data that can be stored in the buffer. For character and binary C data, the `buffer_length` value specifies the length of `*buffer` when used with `mysql_stmt_bind_param()` to specify input values, or the maximum number of output data bytes that can be fetched into the buffer when used with `mysql_stmt_bind_result()`.

- `unsigned long *length`

A pointer to an `unsigned long` variable that indicates the actual number of bytes of data stored in `*buffer`. `length` is used for character or binary C data.

For input parameter data binding, set `*length` to indicate the actual length of the parameter value stored in `*buffer`. This is used by `mysql_stmt_execute()`.

For output value binding, MySQL sets `*length` when you call `mysql_stmt_fetch()`. The `mysql_stmt_fetch()` return value determines how to interpret the length:

- If the return value is 0, `*length` indicates the actual length of the parameter value.
- If the return value is `MYSQL_DATA_TRUNCATED`, `*length` indicates the nontruncated length of the parameter value. In this case, the minimum of `*length` and `buffer_length` indicates the actual length of the value.

`length` is ignored for numeric and temporal data types because the `buffer_type` value determines the length of the data value.

If you must determine the length of a returned value before fetching it, see [Section 21.8.10.11, “mysql\\_stmt\\_fetch\(\)”](#), for some strategies.

- `my_bool *is_null`

This member points to a `my_bool` variable that is true if a value is `NULL`, false if it is not `NULL`. For input, set `*is_null` to true to indicate that you are passing a `NULL` value as a statement parameter.

`is_null` is a *pointer* to a boolean scalar, not a boolean scalar, to provide flexibility in how you specify `NULL` values:

- If your data values are always `NULL`, use `MYSQL_TYPE_NULL` as the `buffer_type` value when you bind the column. The other `MYSQL_BIND` members, including `is_null`, do not matter.
- If your data values are always `NOT NULL`, set `is_null = (my_bool*) 0`, and set the other members appropriately for the variable you are binding.
- In all other cases, set the other members appropriately and set `is_null` to the address of a `my_bool` variable. Set that variable's value to true or false appropriately between executions to indicate whether the corresponding data value is `NULL` or `NOT NULL`, respectively.

For output, when you fetch a row, MySQL sets the value pointed to by `is_null` to true or false according to whether the result set column value returned from the statement is or is not `NULL`.

- `my_bool is_unsigned`

This member applies for C variables with data types that can be `unsigned` (`char`, `short int`, `int`, `long long int`). Set `is_unsigned` to true if the variable pointed to by `buffer` is `unsigned` and false otherwise. For example, if you bind a `signed char` variable to `buffer`, specify a type code of `MYSQL_TYPE_TINY` and set `is_unsigned` to false. If you bind an `unsigned char` instead, the type code is the same but `is_unsigned` should be true. (For `char`, it is not defined whether it is signed or unsigned, so it is best to be explicit about signedness by using `signed char` or `unsigned char`.)

`is_unsigned` applies only to the C language variable on the client side. It indicates nothing about the signedness of the corresponding SQL value on the server side. For example, if you use an `int` variable to supply a value for a `BIGINT UNSIGNED` column, `is_unsigned` should be false because `int` is a signed type. If you use an `unsigned int` variable to supply a value for a `BIGINT` column, `is_unsigned` should be true because `unsigned int` is an unsigned type. MySQL performs the proper conversion between signed and unsigned values in both directions, although a warning occurs if truncation results.

- `my_bool *error`

For output, set this member to point to a `my_bool` variable to have truncation information for the parameter stored there after a row fetching operation. When truncation reporting is enabled, `mysql_stmt_fetch()` returns `MYSQL_DATA_TRUNCATED` and `*error` is true in the `MYSQL_BIND` structures for parameters in which truncation occurred. Truncation indicates loss of sign or significant digits, or that a string was too long to fit in a column. Truncation reporting is enabled by default, but can be controlled by calling `mysql_options()` with the `MYSQL_REPORT_DATA_TRUNCATION` option.

- `MYSQL_TIME`

This structure is used to send and receive `DATE`, `TIME`, `DATETIME`, and `TIMESTAMP` data directly to and from the server. Set the `buffer` member to point to a `MYSQL_TIME` structure, and set the `buffer_type` member of a `MYSQL_BIND` structure to one of the temporal types (`MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATE`, `MYSQL_TYPE_DATETIME`, `MYSQL_TYPE_TIMESTAMP`).

The `MYSQL_TIME` structure contains the members listed in the following table.

Member	Description
<code>unsigned int year</code>	The year
<code>unsigned int month</code>	The month of the year
<code>unsigned int day</code>	The day of the month

Member	Description
<code>unsigned int hour</code>	The hour of the day
<code>unsigned int minute</code>	The minute of the hour
<code>unsigned int second</code>	The second of the minute
<code>my_bool neg</code>	A boolean flag indicating whether the time is negative
<code>unsigned long second_part</code>	The fractional part of the second in microseconds (unused before MySQL 5.6.4)

Only those parts of a `MYSQL_TIME` structure that apply to a given type of temporal value are used. The `year`, `month`, and `day` elements are used for `DATE`, `DATETIME`, and `TIMESTAMP` values. The `hour`, `minute`, and `second` elements are used for `TIME`, `DATETIME`, and `TIMESTAMP` values. See [Section 21.8.18, “C API Prepared Statement Handling of Date and Time Values”](#).

### 21.8.8.1. C API Prepared Statement Type Codes

The `buffer_type` member of `MYSQL_BIND` structures indicates the data type of the C language variable bound to a statement parameter or result set column. For input, `buffer_type` indicates the type of the variable containing the value to be sent to the server. For output, it indicates the type of the variable into which a value received from the server should be stored.

The following table shows the permissible values for the `buffer_type` member of `MYSQL_BIND` structures for input values sent to the server. The table shows the C variable types that you can use, the corresponding type codes, and the SQL data types for which the supplied value can be used without conversion. Choose the `buffer_type` value according to the data type of the C language variable that you are binding. For the integer types, you should also set the `is_unsigned` member to indicate whether the variable is signed or unsigned.

Input Variable C Type	<code>buffer_type</code> Value	SQL Type of Destination Value
<code>signed char</code>	<code>MYSQL_TYPE_TINY</code>	<code>TINYINT</code>
<code>short int</code>	<code>MYSQL_TYPE_SHORT</code>	<code>SMALLINT</code>
<code>int</code>	<code>MYSQL_TYPE_LONG</code>	<code>INT</code>
<code>long long int</code>	<code>MYSQL_TYPE_LONGLONG</code>	<code>BIGINT</code>
<code>float</code>	<code>MYSQL_TYPE_FLOAT</code>	<code>FLOAT</code>
<code>double</code>	<code>MYSQL_TYPE_DOUBLE</code>	<code>DOUBLE</code>
<code>MYSQL_TIME</code>	<code>MYSQL_TYPE_TIME</code>	<code>TIME</code>
<code>MYSQL_TIME</code>	<code>MYSQL_TYPE_DATE</code>	<code>DATE</code>
<code>MYSQL_TIME</code>	<code>MYSQL_TYPE_DATETIME</code>	<code>DATETIME</code>
<code>MYSQL_TIME</code>	<code>MYSQL_TYPE_TIMESTAMP</code>	<code>TIMESTAMP</code>
<code>char[ ]</code>	<code>MYSQL_TYPE_STRING</code>	<code>TEXT</code> , <code>CHAR</code> , <code>VARCHAR</code>
<code>char[ ]</code>	<code>MYSQL_TYPE_BLOB</code>	<code>BLOB</code> , <code>BINARY</code> , <code>VARBINARY</code>
	<code>MYSQL_TYPE_NULL</code>	<code>NULL</code>

Use `MYSQL_TYPE_NULL` as indicated in the description for the `is_null` member in [Section 21.8.8, “C API Prepared Statement Data Structures”](#).

For input string data, use `MYSQL_TYPE_STRING` or `MYSQL_TYPE_BLOB` depending on whether the value is a character (nonbinary) or binary string:

- `MYSQL_TYPE_STRING` indicates character input string data. The value is assumed to be in the character set indicated by the `character_set_client` [467] system variable. If the server stores the value into a column with a different character set, it converts the value to that character set.



- `MYSQL_TYPE_BLOB` indicates binary input string data. The value is treated as having the `binary` character set. That is, it is treated as a byte string and no conversion occurs.

The following table shows the permissible values for the `buffer_type` member of `MYSQL_BIND` structures for output values received from the server. The table shows the SQL types of received values, the corresponding type codes that such values have in result set metadata, and the recommended C language data types to bind to the `MYSQL_BIND` structure to receive the SQL values without conversion. Choose the `buffer_type` value according to the data type of the C language variable that you are binding. For the integer types, you should also set the `is_unsigned` member to indicate whether the variable is signed or unsigned.

SQL Type of Received Value	<code>buffer_type</code> Value	Output Variable C Type
TINYINT	<code>MYSQL_TYPE_TINY</code>	signed char
SMALLINT	<code>MYSQL_TYPE_SHORT</code>	short int
MEDIUMINT	<code>MYSQL_TYPE_INT24</code>	int
INT	<code>MYSQL_TYPE_LONG</code>	int
BIGINT	<code>MYSQL_TYPE_LONGLONG</code>	long long int
FLOAT	<code>MYSQL_TYPE_FLOAT</code>	float
DOUBLE	<code>MYSQL_TYPE_DOUBLE</code>	double
DECIMAL	<code>MYSQL_TYPE_NEWDECIMAL</code>	char[ ]
YEAR	<code>MYSQL_TYPE_SHORT</code>	short int
TIME	<code>MYSQL_TYPE_TIME</code>	<code>MYSQL_TIME</code>
DATE	<code>MYSQL_TYPE_DATE</code>	<code>MYSQL_TIME</code>
DATETIME	<code>MYSQL_TYPE_DATETIME</code>	<code>MYSQL_TIME</code>
TIMESTAMP	<code>MYSQL_TYPE_TIMESTAMP</code>	<code>MYSQL_TIME</code>
CHAR, BINARY	<code>MYSQL_TYPE_STRING</code>	char[ ]
VARCHAR, VARBINARY	<code>MYSQL_TYPE_VAR_STRING</code>	char[ ]
TINYBLOB, TINYTEXT	<code>MYSQL_TYPE_TINY_BLOB</code>	char[ ]
BLOB, TEXT	<code>MYSQL_TYPE_BLOB</code>	char[ ]
MEDIUMBLOB, MEDIUMTEXT	<code>MYSQL_TYPE_MEDIUM_BLOB</code>	char[ ]
LOBLOB, LONGTEXT	<code>MYSQL_TYPE_LONG_BLOB</code>	char[ ]
BIT	<code>MYSQL_TYPE_BIT</code>	char[ ]

### 21.8.8.2. C API Prepared Statement Type Conversions

Prepared statements transmit data between the client and server using C language variables on the client side that correspond to SQL values on the server side. If there is a mismatch between the C variable type on the client side and the corresponding SQL value type on the server side, MySQL performs implicit type conversions in both directions.

MySQL knows the type code for the SQL value on the server side. The `buffer_type` value in the `MYSQL_BIND` structure indicates the type code of the C variable that holds the value on the client side. The two codes together tell MySQL what conversion must be performed, if any. Here are some examples:

- If you use `MYSQL_TYPE_LONG` with an `int` variable to pass an integer value to the server that is to be stored into a `FLOAT` column, MySQL converts the value to floating-point format before storing it.
- If you fetch an SQL `MEDIUMINT` column value, but specify a `buffer_type` value of `MYSQL_TYPE_LONGLONG` and use a C variable of type `long long int` as the destination buffer,



MySQL converts the `MEDIUMINT` value (which requires less than 8 bytes) for storage into the `long long int` (an 8-byte variable).

- If you fetch a numeric column with a value of 255 into a `char[4]` character array and specify a `buffer_type` value of `MYSQL_TYPE_STRING`, the resulting value in the array is a 4-byte string `'255\0'`.
- MySQL returns `DECIMAL` values as the string representation of the original server-side value, which is why the corresponding C type is `char[]`. For example, `12.345` is returned to the client as `'12.345'`. If you specify `MYSQL_TYPE_NEWDECIMAL` and bind a string buffer to the `MYSQL_BIND` structure, `mysql_stmt_fetch()` stores the value in the buffer as a string without conversion. If instead you specify a numeric variable and type code, `mysql_stmt_fetch()` converts the string-format `DECIMAL` value to numeric form.
- For the `MYSQL_TYPE_BIT` type code, `BIT` values are returned into a string buffer, which is why the corresponding C type is `char[]`. The value represents a bit string that requires interpretation on the client side. To return the value as a type that is easier to deal with, you can cause the value to be cast to integer using either of the following types of expressions:

```
SELECT bit_col + 0 FROM t
SELECT CAST(bit_col AS UNSIGNED) FROM t
```

To retrieve the value, bind an integer variable large enough to hold the value and specify the appropriate corresponding integer type code.

Before binding variables to the `MYSQL_BIND` structures that are to be used for fetching column values, you can check the type codes for each column of the result set. This might be desirable if you want to determine which variable types would be best to use to avoid type conversions. To get the type codes, call `mysql_stmt_result_metadata()` after executing the prepared statement with `mysql_stmt_execute()`. The metadata provides access to the type codes for the result set as described in [Section 21.8.10.23](#), “`mysql_stmt_result_metadata()`”, and [Section 21.8.4](#), “C API Data Structures”.

To determine whether output string values in a result set returned from the server contain binary or nonbinary data, check whether the `charsetnr` value of the result set metadata is 63 (see [Section 21.8.4](#), “C API Data Structures”). If so, the character set is `binary`, which indicates binary rather than nonbinary data. This enables you to distinguish `BINARY` from `CHAR`, `VARBINARY` from `VARCHAR`, and the `BLOB` types from the `TEXT` types.

If you cause the `max_length` member of the `MYSQL_FIELD` column metadata structures to be set (by calling `mysql_stmt_attr_set()`), be aware that the `max_length` values for the result set indicate the lengths of the longest string representation of the result values, not the lengths of the binary representation. That is, `max_length` does not necessarily correspond to the size of the buffers needed to fetch the values with the binary protocol used for prepared statements. Choose the size of the buffers according to the types of the variables into which you fetch the values. For example, a `TINYINT` column containing the value -128 might have a `max_length` value of 4. But the binary representation of any `TINYINT` value requires only 1 byte for storage, so you can supply a `signed char` variable in which to store the value and set `is_unsigned` to indicate that values are signed.

Metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed. For more information, see [Section 8.9.4](#), “Caching of Prepared Statements and Stored Programs”.

## 21.8.9. C API Prepared Statement Function Overview

The functions available for prepared statement processing are summarized here and described in greater detail in a later section. See [Section 21.8.10](#), “C API Prepared Statement Function Descriptions”.

Function	Description
<code>mysql_stmt_affected_rows()</code>	Returns the number of rows changed, deleted, or inserted by prepared <code>UPDATE</code> , <code>DELETE</code> , or <code>INSERT</code> statement
<code>mysql_stmt_attr_get()</code>	Gets value of an attribute for a prepared statement
<code>mysql_stmt_attr_set()</code>	Sets an attribute for a prepared statement
<code>mysql_stmt_bind_param()</code>	Associates application data buffers with the parameter markers in a prepared SQL statement
<code>mysql_stmt_bind_result()</code>	Associates application data buffers with columns in a result set
<code>mysql_stmt_close()</code>	Frees memory used by a prepared statement
<code>mysql_stmt_data_seek()</code>	Seeks to an arbitrary row number in a statement result set
<code>mysql_stmt_errno()</code>	Returns the error number for the last statement execution
<code>mysql_stmt_error()</code>	Returns the error message for the last statement execution
<code>mysql_stmt_execute()</code>	Executes a prepared statement
<code>mysql_stmt_fetch()</code>	Fetches the next row of data from a result set and returns data for all bound columns
<code>mysql_stmt_fetch_column()</code>	Fetch data for one column of the current row of a result set
<code>mysql_stmt_field_count()</code>	Returns the number of result columns for the most recent statement
<code>mysql_stmt_free_result()</code>	Free the resources allocated to a statement handle
<code>mysql_stmt_init()</code>	Allocates memory for a <code>MYSQL_STMT</code> structure and initializes it
<code>mysql_stmt_insert_id()</code>	Returns the ID generated for an <code>AUTO_INCREMENT</code> column by a prepared statement
<code>mysql_stmt_next_result()</code>	Returns/initiates the next result in a multiple-result execution
<code>mysql_stmt_num_rows()</code>	Returns the row count from a buffered statement result set
<code>mysql_stmt_param_count()</code>	Returns the number of parameters in a prepared statement
<code>mysql_stmt_param_metadata()</code>	(Return parameter metadata in the form of a result set) Currently, this function does nothing
<code>mysql_stmt_prepare()</code>	Prepares an SQL statement string for execution
<code>mysql_stmt_reset()</code>	Resets the statement buffers in the server
<code>mysql_stmt_result_metadata()</code>	Returns prepared statement metadata in the form of a result set
<code>mysql_stmt_row_seek()</code>	Seeks to a row offset in a statement result set, using value returned from <code>mysql_stmt_row_tell()</code>
<code>mysql_stmt_row_tell()</code>	Returns the statement row cursor position
<code>mysql_stmt_send_long_data()</code>	Sends long data in chunks to server
<code>mysql_stmt_sqlstate()</code>	Returns the <code>SQLSTATE</code> error code for the last statement execution
<code>mysql_stmt_store_result()</code>	Retrieves a complete result set to the client

Call `mysql_stmt_init()` to create a statement handle, then `mysql_stmt_prepare()` to prepare the statement string, `mysql_stmt_bind_param()` to supply the parameter data, and `mysql_stmt_execute()` to execute the statement. You can repeat the `mysql_stmt_execute()` by changing parameter values in the respective buffers supplied through `mysql_stmt_bind_param()`.

You can send text or binary data in chunks to server using `mysql_stmt_send_long_data()`. See [Section 21.8.10.26](#), “`mysql_stmt_send_long_data()`”.

If the statement is a `SELECT` or any other statement that produces a result set, `mysql_stmt_prepare()` also returns the result set metadata information in the form of a `MYSQL_RES` result set through `mysql_stmt_result_metadata()`.

You can supply the result buffers using `mysql_stmt_bind_result()`, so that the `mysql_stmt_fetch()` automatically returns data to these buffers. This is row-by-row fetching.

When statement execution has been completed, close the statement handle using `mysql_stmt_close()` so that all resources associated with it can be freed.

If you obtained a `SELECT` statement's result set metadata by calling `mysql_stmt_result_metadata()`, you should also free the metadata using `mysql_free_result()`.

## Execution Steps

To prepare and execute a statement, an application follows these steps:

1. Create a prepared statement handle with `mysql_stmt_init()`. To prepare the statement on the server, call `mysql_stmt_prepare()` and pass it a string containing the SQL statement.
2. If the statement will produce a result set, call `mysql_stmt_result_metadata()` to obtain the result set metadata. This metadata is itself in the form of result set, albeit a separate one from the one that contains the rows returned by the query. The metadata result set indicates how many columns are in the result and contains information about each column.
3. Set the values of any parameters using `mysql_stmt_bind_param()`. All parameters must be set. Otherwise, statement execution returns an error or produces unexpected results.
4. Call `mysql_stmt_execute()` to execute the statement.
5. If the statement produces a result set, bind the data buffers to use for retrieving the row values by calling `mysql_stmt_bind_result()`.
6. Fetch the data into the buffers row by row by calling `mysql_stmt_fetch()` repeatedly until no more rows are found.
7. Repeat steps 3 through 6 as necessary, by changing the parameter values and re-executing the statement.

When `mysql_stmt_prepare()` is called, the MySQL client/server protocol performs these actions:

- The server parses the statement and sends the okay status back to the client by assigning a statement ID. It also sends total number of parameters, a column count, and its metadata if it is a result set oriented statement. All syntax and semantics of the statement are checked by the server during this call.
- The client uses this statement ID for the further operations, so that the server can identify the statement from among its pool of statements.

When `mysql_stmt_execute()` is called, the MySQL client/server protocol performs these actions:

- The client uses the statement handle and sends the parameter data to the server.
- The server identifies the statement using the ID provided by the client, replaces the parameter markers with the newly supplied data, and executes the statement. If the statement produces a result set, the server sends the data back to the client. Otherwise, it sends an okay status and the number of rows changed, deleted, or inserted.

When `mysql_stmt_fetch()` is called, the MySQL client/server protocol performs these actions:

- The client reads the data from the current row of the result set and places it into the application data buffers by doing the necessary conversions. If the application buffer type is same as that of the field type returned from the server, the conversions are straightforward.

If an error occurs, you can get the statement error number, error message, and SQLSTATE code using `mysql_stmt_errno()`, `mysql_stmt_error()`, and `mysql_stmt_sqlstate()`, respectively.

## Prepared Statement Logging

For prepared statements that are executed with the `mysql_stmt_prepare()` and `mysql_stmt_execute()` C API functions, the server writes `Prepare` and `Execute` lines to the general query log so that you can tell when statements are prepared and executed.

Suppose that you prepare and execute a statement as follows:

1. Call `mysql_stmt_prepare()` to prepare the statement string `"SELECT ?"`.
2. Call `mysql_stmt_bind_param()` to bind the value `3` to the parameter in the prepared statement.
3. Call `mysql_stmt_execute()` to execute the prepared statement.

As a result of the preceding calls, the server writes the following lines to the general query log:

```
Prepare [1] SELECT ?
Execute [1] SELECT 3
```

Each `Prepare` and `Execute` line in the log is tagged with a `[N]` statement identifier so that you can keep track of which prepared statement is being logged. `N` is a positive integer. If there are multiple prepared statements active simultaneously for the client, `N` may be greater than 1. Each `Execute` lines shows a prepared statement after substitution of data values for `?` parameters.

## 21.8.10. C API Prepared Statement Function Descriptions

To prepare and execute queries, use the functions described in detail in the following sections.

All functions that operate with a `MYSQL_STMT` structure begin with the prefix `mysql_stmt_`.

To create a `MYSQL_STMT` handle, use the `mysql_stmt_init()` function.

### 21.8.10.1. `mysql_stmt_affected_rows()`

```
my_ulonglong mysql_stmt_affected_rows(MYSQL_STMT *stmt)
```

#### Description

`mysql_stmt_affected_rows()` may be called immediately after executing a statement with `mysql_stmt_execute()`. It is like `mysql_affected_rows()` but for prepared statements. For a description of what the affected-rows value returned by this function means, See [Section 21.8.6.1](#), “`mysql_affected_rows()`”.

#### Errors

None.

#### Example

See the Example in [Section 21.8.10.10](#), “`mysql_stmt_execute()`”.

### 21.8.10.2. `mysql_stmt_attr_get()`

```
my_bool mysql_stmt_attr_get(MYSQL_STMT *stmt, enum enum_stmt_attr_type
option, void *arg)
```

#### Description

Can be used to get the current value for a statement attribute.

The `option` argument is the option that you want to get; the `arg` should point to a variable that should contain the option value. If the option is an integer, `arg` should point to the value of the integer.

See [Section 21.8.10.3](#), “`mysql_stmt_attr_set()`”, for a list of options and option types.

## Return Values

Zero if successful. Nonzero if `option` is unknown.

## Errors

None.

### 21.8.10.3. `mysql_stmt_attr_set()`

```
my_bool mysql_stmt_attr_set(MYSQL_STMT *stmt, enum enum_stmt_attr_type
option, const void *arg)
```

## Description

Can be used to affect behavior for a prepared statement. This function may be called multiple times to set several options.

The `option` argument is the option that you want to set. The `arg` argument is the value for the option. `arg` should point to a variable that is set to the desired attribute value. The variable type is as indicated in the following table.

The following table shows the possible `option` values.

Option	Argument Type	Function
<code>STMT_ATTR_UPDATE_MAX_LENGTH</code>	<code>my_bool *</code>	If set to 1, causes <code>mysql_stmt_store_result()</code> to update the metadata <code>MYSQL_FIELD-&gt;max_length</code> value.
<code>STMT_ATTR_CURSOR_TYPE</code>	<code>unsigned long *</code>	Type of cursor to open for statement when <code>mysql_stmt_execute()</code> is invoked. <code>*arg</code> can be <code>CURSOR_TYPE_NO_CURSOR</code> (the default) or <code>CURSOR_TYPE_READ_ONLY</code> .
<code>STMT_ATTR_PREFETCH_ROWS</code>	<code>unsigned long *</code>	Number of rows to fetch from server at a time when using a cursor. <code>*arg</code> can be in the range from 1 to the maximum value of <code>unsigned long</code> . The default is 1.

If you use the `STMT_ATTR_CURSOR_TYPE` option with `CURSOR_TYPE_READ_ONLY`, a cursor is opened for the statement when you invoke `mysql_stmt_execute()`. If there is already an open cursor from a previous `mysql_stmt_execute()` call, it closes the cursor before opening a new one. `mysql_stmt_reset()` also closes any open cursor before preparing the statement for re-execution. `mysql_stmt_free_result()` closes any open cursor.

If you open a cursor for a prepared statement, `mysql_stmt_store_result()` is unnecessary, because that function causes the result set to be buffered on the client side.

## Return Values

Zero if successful. Nonzero if `option` is unknown.

## Errors

None.

## Example

The following example opens a cursor for a prepared statement and sets the number of rows to fetch at a time to 5:

```
MYSQL_STMT *stmt;
int rc;
unsigned long type;
unsigned long prefetch_rows = 5;

stmt = mysql_stmt_init(mysql);
type = (unsigned long) CURSOR_TYPE_READ_ONLY;
rc = mysql_stmt_attr_set(stmt, STMT_ATTR_CURSOR_TYPE, (void*) &type);
/* ... check return value ... */
rc = mysql_stmt_attr_set(stmt, STMT_ATTR_PREFETCH_ROWS,
                        (void*) &prefetch_rows);
/* ... check return value ... */
```

### 21.8.10.4. `mysql_stmt_bind_param()`

```
my_bool mysql_stmt_bind_param(MYSQL_STMT *stmt, MYSQL_BIND *bind)
```

## Description

`mysql_stmt_bind_param()` is used to bind input data for the parameter markers in the SQL statement that was passed to `mysql_stmt_prepare()`. It uses `MYSQL_BIND` structures to supply the data. `bind` is the address of an array of `MYSQL_BIND` structures. The client library expects the array to contain one element for each `?` parameter marker that is present in the query.

Suppose that you prepare the following statement:

```
INSERT INTO mytbl VALUES(?,?,?)
```

When you bind the parameters, the array of `MYSQL_BIND` structures must contain three elements, and can be declared like this:

```
MYSQL_BIND bind[3];
```

Section 21.8.8, “C API Prepared Statement Data Structures”, describes the members of each `MYSQL_BIND` element and how they should be set to provide input values.

## Return Values

Zero if the bind operation was successful. Nonzero if an error occurred.

## Errors

- `CR_UNSUPPORTED_PARAM_TYPE` [3720]

The conversion is not supported. Possibly the `buffer_type` value is invalid or is not one of the supported types.

- `CR_OUT_OF_MEMORY` [3718]

Out of memory.

- `CR_UNKNOWN_ERROR` [3718]

An unknown error occurred.

## Example

See the Example in Section 21.8.10.10, “`mysql_stmt_execute()`”.

### 21.8.10.5. `mysql_stmt_bind_result()`

```
my_bool mysql_stmt_bind_result(MYSQL_STMT *stmt, MYSQL_BIND *bind)
```

## Description

`mysql_stmt_bind_result()` is used to associate (that is, bind) output columns in the result set to data buffers and length buffers. When `mysql_stmt_fetch()` is called to fetch data, the MySQL client/server protocol places the data for the bound columns into the specified buffers.

All columns must be bound to buffers prior to calling `mysql_stmt_fetch()`. `bind` is the address of an array of `MYSQL_BIND` structures. The client library expects the array to contain one element for each column of the result set. If you do not bind columns to `MYSQL_BIND` structures, `mysql_stmt_fetch()` simply ignores the data fetch. The buffers should be large enough to hold the data values, because the protocol doesn't return data values in chunks.

A column can be bound or rebound at any time, even after a result set has been partially retrieved. The new binding takes effect the next time `mysql_stmt_fetch()` is called. Suppose that an application binds the columns in a result set and calls `mysql_stmt_fetch()`. The client/server protocol returns data in the bound buffers. Then suppose that the application binds the columns to a different set of buffers. The protocol places data into the newly bound buffers when the next call to `mysql_stmt_fetch()` occurs.

To bind a column, an application calls `mysql_stmt_bind_result()` and passes the type, address, and length of the output buffer into which the value should be stored. [Section 21.8.8, "C API Prepared Statement Data Structures"](#), describes the members of each `MYSQL_BIND` element and how they should be set to receive output values.

## Return Values

Zero if the bind operation was successful. Nonzero if an error occurred.

## Errors

- [CR\\_UNSUPPORTED\\_PARAM\\_TYPE \[3720\]](#)  
The conversion is not supported. Possibly the `buffer_type` value is invalid or is not one of the supported types.
- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)  
Out of memory.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

## Example

See the Example in [Section 21.8.10.11, "mysql\\_stmt\\_fetch\(\)"](#).

### 21.8.10.6. `mysql_stmt_close()`

```
my_bool mysql_stmt_close(MYSQL_STMT *)
```

## Description

Closes the prepared statement. `mysql_stmt_close()` also deallocates the statement handle pointed to by `stmt`.

If the current statement has pending or unread results, this function cancels them so that the next query can be executed.

## Return Values

Zero if the statement was freed successfully. Nonzero if an error occurred.



## Errors

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

## Example

See the Example in [Section 21.8.10.10](#), “`mysql_stmt_execute()`”.

### 21.8.10.7. `mysql_stmt_data_seek()`

```
void mysql_stmt_data_seek(MYSQL_STMT *stmt, my_ulonglong offset)
```

## Description

Seeks to an arbitrary row in a statement result set. The `offset` value is a row number and should be in the range from 0 to `mysql_stmt_num_rows(stmt)-1`.

This function requires that the statement result set structure contains the entire result of the last executed query, so `mysql_stmt_data_seek()` may be used only in conjunction with `mysql_stmt_store_result()`.

## Return Values

None.

## Errors

None.

### 21.8.10.8. `mysql_stmt_errno()`

```
unsigned int mysql_stmt_errno(MYSQL_STMT *stmt)
```

## Description

For the statement specified by `stmt`, `mysql_stmt_errno()` returns the error code for the most recently invoked statement API function that can succeed or fail. A return value of zero means that no error occurred. Client error message numbers are listed in the MySQL `errmsg.h` header file. Server error message numbers are listed in `mysql_error.h`. Errors also are listed at [Appendix C, Errors, Error Codes, and Common Problems](#).

## Return Values

An error code value. Zero if no error occurred.

## Errors

None.

### 21.8.10.9. `mysql_stmt_error()`

```
const char *mysql_stmt_error(MYSQL_STMT *stmt)
```

## Description

For the statement specified by `stmt`, `mysql_stmt_error()` returns a null-terminated string containing the error message for the most recently invoked statement API function that can succeed or fail. An empty string (“”) is returned if no error occurred. This means the following two tests are equivalent:



```
if (*mysql_stmt_errno(stmt))
{
    // an error occurred
}

if (mysql_stmt_error(stmt)[0])
{
    // an error occurred
}
```

The language of the client error messages may be changed by recompiling the MySQL client library. Currently, you can choose error messages in several different languages.

## Return Values

A character string that describes the error. An empty string if no error occurred.

## Errors

None.

### 21.8.10.10. `mysql_stmt_execute()`

```
int mysql_stmt_execute(MYSQL_STMT *stmt)
```

## Description

`mysql_stmt_execute()` executes the prepared query associated with the statement handle. The currently bound parameter marker values are sent to server during this call, and the server replaces the markers with this newly supplied data.

Statement processing following `mysql_stmt_execute()` depends on the type of statement:

- For an `UPDATE`, `DELETE`, or `INSERT`, the number of changed, deleted, or inserted rows can be found by calling `mysql_stmt_affected_rows()`.
- For a statement such as `SELECT` that generates a result set, you must call `mysql_stmt_fetch()` to fetch the data prior to calling any other functions that result in query processing. For more information on how to fetch the results, refer to [Section 21.8.10.11, “mysql\\_stmt\\_fetch\(\)”](#).

Do not following invocation of `mysql_stmt_execute()` with a call to `mysql_store_result()` or `mysql_use_result()`. Those functions are not intended for processing results from prepared statements.

For statements that generate a result set, you can request that `mysql_stmt_execute()` open a cursor for the statement by calling `mysql_stmt_attr_set()` before executing the statement. If you execute a statement multiple times, `mysql_stmt_execute()` closes any open cursor before opening a new one.

Metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

## Return Values

Zero if execution was successful. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)

Out of memory.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

## Example

The following example demonstrates how to create and populate a table using `mysql_stmt_init()`, `mysql_stmt_prepare()`, `mysql_stmt_param_count()`, `mysql_stmt_bind_param()`, `mysql_stmt_execute()`, and `mysql_stmt_affected_rows()`. The `mysql` variable is assumed to be a valid connection handle. For an example that shows how to retrieve data, see [Section 21.8.10.11](#), “`mysql_stmt_fetch()`”.

```
#define STRING_SIZE 50

#define DROP_SAMPLE_TABLE "DROP TABLE IF EXISTS test_table"
#define CREATE_SAMPLE_TABLE "CREATE TABLE test_table(coll INT,\
                                col2 VARCHAR(40),\
                                col3 SMALLINT,\
                                col4 TIMESTAMP)"

#define INSERT_SAMPLE "INSERT INTO \
    test_table(coll,col2,col3) \
    VALUES(?,?,?)"

MYSQL_STMT *stmt;
MYSQL_BIND bind[3];
my_ulonglong affected_rows;
int param_count;
short small_data;
int int_data;
char str_data[STRING_SIZE];
unsigned long str_length;
my_bool is_null;

if (mysql_query(mysql, DROP_SAMPLE_TABLE))
{
    fprintf(stderr, " DROP TABLE failed\n");
    fprintf(stderr, " %s\n", mysql_error(mysql));
    exit(0);
}

if (mysql_query(mysql, CREATE_SAMPLE_TABLE))
{
    fprintf(stderr, " CREATE TABLE failed\n");
    fprintf(stderr, " %s\n", mysql_error(mysql));
    exit(0);
}

/* Prepare an INSERT query with 3 parameters */
/* (the TIMESTAMP column is not named; the server */
/* sets it to the current date and time) */
stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, " mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(stmt, INSERT_SAMPLE, strlen(INSERT_SAMPLE)))
{
    fprintf(stderr, " mysql_stmt_prepare(), INSERT failed\n");
```

```

fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
exit(0);
}
fprintf(stdout, " prepare, INSERT successful\n");

/* Get the parameter count from the statement */
param_count= mysql_stmt_param_count(stmt);
fprintf(stdout, " total parameters in INSERT: %d\n", param_count);

if (param_count != 3) /* validate parameter count */
{
    fprintf(stderr, " invalid parameter count returned by MySQL\n");
    exit(0);
}

/* Bind the data for all 3 parameters */

memset(bind, 0, sizeof(bind));

/* INTEGER PARAM */
/* This is a number type, so there is no need
   to specify buffer_length */
bind[0].buffer_type= MYSQL_TYPE_LONG;
bind[0].buffer= (char *)&int_data;
bind[0].is_null= 0;
bind[0].length= 0;

/* STRING PARAM */
bind[1].buffer_type= MYSQL_TYPE_STRING;
bind[1].buffer= (char *)str_data;
bind[1].buffer_length= STRING_SIZE;
bind[1].is_null= 0;
bind[1].length= &str_length;

/* SMALLINT PARAM */
bind[2].buffer_type= MYSQL_TYPE_SHORT;
bind[2].buffer= (char *)&small_data;
bind[2].is_null= &is_null;
bind[2].length= 0;

/* Bind the buffers */
if (mysql_stmt_bind_param(stmt, bind))
{
    fprintf(stderr, " mysql_stmt_bind_param() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Specify the data values for the first row */
int_data= 10; /* integer */
strncpy(str_data, "MySQL", STRING_SIZE); /* string */
str_length= strlen(str_data);

/* INSERT SMALLINT data as NULL */
is_null= 1;

/* Execute the INSERT statement - 1*/
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute(), 1 failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get the number of affected rows */
affected_rows= mysql_stmt_affected_rows(stmt);
fprintf(stdout, " total affected rows(insert 1): %lu\n",
        (unsigned long) affected_rows);

if (affected_rows != 1) /* validate affected rows */
{
    fprintf(stderr, " invalid affected rows by MySQL\n");

```

```

    exit(0);
}

/* Specify data values for second row,
   then re-execute the statement */
int_data= 1000;
strncpy(str_data, "
        The most popular Open Source database",
        STRING_SIZE);
str_length= strlen(str_data);
small_data= 1000;          /* smallint */
is_null= 0;              /* reset */

/* Execute the INSERT statement - 2*/
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute, 2 failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get the total rows affected */
affected_rows= mysql_stmt_affected_rows(stmt);
fprintf(stdout, " total affected rows(insert 2): %lu\n",
        (unsigned long) affected_rows);

if (affected_rows != 1) /* validate affected rows */
{
    fprintf(stderr, " invalid affected rows by MySQL\n");
    exit(0);
}

/* Close the statement */
if (mysql_stmt_close(stmt))
{
    fprintf(stderr, " failed while closing the statement\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

```



### Note

For complete examples on the use of prepared statement functions, refer to the file `tests/mysql_client_test.c`. This file can be obtained from a MySQL source distribution or from the Bazaar source repository.

#### 21.8.10.11. `mysql_stmt_fetch()`

```
int mysql_stmt_fetch(MYSQL_STMT *stmt)
```

#### Description

`mysql_stmt_fetch()` returns the next row in the result set. It can be called only while the result set exists; that is, after a call to `mysql_stmt_execute()` for a statement such as `SELECT` that produces a result set.

`mysql_stmt_fetch()` returns row data using the buffers bound by `mysql_stmt_bind_result()`. It returns the data in those buffers for all the columns in the current row set and the lengths are returned to the `length` pointer. All columns must be bound by the application before it calls `mysql_stmt_fetch()`.

By default, result sets are fetched unbuffered a row at a time from the server. To buffer the entire result set on the client, call `mysql_stmt_store_result()` after binding the data buffers and before calling `mysql_stmt_fetch()`.

If a fetched data value is a `NULL` value, the `*is_null` value of the corresponding `MYSQL_BIND` structure contains `TRUE` (1). Otherwise, the data and its length are returned in the `*buffer` and `*length` elements based on the buffer type specified by the application. Each numeric and temporal

type has a fixed length, as listed in the following table. The length of the string types depends on the length of the actual data value, as indicated by `data_length`.

Type	Length
<code>MYSQL_TYPE_TINY</code>	1
<code>MYSQL_TYPE_SHORT</code>	2
<code>MYSQL_TYPE_LONG</code>	4
<code>MYSQL_TYPE_LONGLONG</code>	8
<code>MYSQL_TYPE_FLOAT</code>	4
<code>MYSQL_TYPE_DOUBLE</code>	8
<code>MYSQL_TYPE_TIME</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_DATE</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_DATETIME</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_STRING</code>	<code>data_length</code>
<code>MYSQL_TYPE_BLOB</code>	<code>data_length</code>

In some cases you might want to determine the length of a column value before fetching it with `mysql_stmt_fetch()`. For example, the value might be a long string or `BLOB` value for which you want to know how much space must be allocated. To accomplish this, you can use these strategies:

- Before invoking `mysql_stmt_fetch()` to retrieve individual rows, pass `STMT_ATTR_UPDATE_MAX_LENGTH` to `mysql_stmt_attr_set()`, then invoke `mysql_stmt_store_result()` to buffer the entire result on the client side. Setting the `STMT_ATTR_UPDATE_MAX_LENGTH` attribute causes the maximal length of column values to be indicated by the `max_length` member of the result set metadata returned by `mysql_stmt_result_metadata()`.
- Invoke `mysql_stmt_fetch()` with a zero-length buffer for the column in question and a pointer in which the real length can be stored. Then use the real length with `mysql_stmt_fetch_column()`.

```
real_length= 0;

bind[0].buffer= 0;
bind[0].buffer_length= 0;
bind[0].length= &real_length
mysql_stmt_bind_result(stmt, bind);

mysql_stmt_fetch(stmt);
if (real_length > 0)
{
    data= malloc(real_length);
    bind[0].buffer= data;
    bind[0].buffer_length= real_length;
    mysql_stmt_fetch_column(stmt, bind, 0, 0);
}
```

## Return Values

Return Value	Description
0	Successful, the data has been fetched to application data buffers.
1	Error occurred. Error code and message can be obtained by calling <code>mysql_stmt_errno()</code> and <code>mysql_stmt_error()</code> .
<code>MYSQL_NO_DATA</code>	No more rows/data exists
<code>MYSQL_DATA_TRUNCATED</code>	Data truncation occurred

`MYSQL_DATA_TRUNCATED` is returned when truncation reporting is enabled. To determine which column values were truncated when this value is returned, check the `error` members of the

`MYSQL_BIND` structures used for fetching values. Truncation reporting is enabled by default, but can be controlled by calling `mysql_options()` with the `MYSQL_REPORT_DATA_TRUNCATION` option.

## Errors

- `CR_COMMANDS_OUT_OF_SYNC` [3719]  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY` [3718]  
Out of memory.
- `CR_SERVER_GONE_ERROR` [3718]  
The MySQL server has gone away.
- `CR_SERVER_LOST` [3719]  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR` [3718]  
An unknown error occurred.
- `CR_UNSUPPORTED_PARAM_TYPE` [3720]  
The buffer type is `MYSQL_TYPE_DATE`, `MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATETIME`, or `MYSQL_TYPE_TIMESTAMP`, but the data type is not `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP`.
- All other unsupported conversion errors are returned from `mysql_stmt_bind_result()`.

## Example

The following example demonstrates how to fetch data from a table using `mysql_stmt_result_metadata()`, `mysql_stmt_bind_result()`, and `mysql_stmt_fetch()`. (This example expects to retrieve the two rows inserted by the example shown in [Section 21.8.10.10](#), “`mysql_stmt_execute()`”.) The `mysql` variable is assumed to be a valid connection handle.

```
#define STRING_SIZE 50

#define SELECT_SAMPLE "SELECT col1, col2, col3, col4 \
                      FROM test_table"

MYSQL_STMT      *stmt;
MYSQL_BIND      bind[4];
MYSQL_RES       *prepare_meta_result;
MYSQL_TIME      ts;
unsigned long   length[4];
int             param_count, column_count, row_count;
short          small_data;
int            int_data;
char           str_data[STRING_SIZE];
my_bool       is_null[4];
my_bool       error[4];

/* Prepare a SELECT query to fetch data from test_table */
stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, " mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(stmt, SELECT_SAMPLE, strlen(SELECT_SAMPLE)))
{
    fprintf(stderr, " mysql_stmt_prepare(), SELECT failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}
```

```
}
fprintf(stdout, " prepare, SELECT successful\n");

/* Get the parameter count from the statement */
param_count= mysql_stmt_param_count(stmt);
fprintf(stdout, " total parameters in SELECT: %d\n", param_count);

if (param_count != 0) /* validate parameter count */
{
    fprintf(stderr, " invalid parameter count returned by MySQL\n");
    exit(0);
}

/* Fetch result set meta information */
prepare_meta_result = mysql_stmt_result_metadata(stmt);
if (!prepare_meta_result)
{
    fprintf(stderr,
            " mysql_stmt_result_metadata(), \
            returned no meta information\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get total columns in the query */
column_count= mysql_num_fields(prepare_meta_result);
fprintf(stdout,
        " total columns in SELECT statement: %d\n",
        column_count);

if (column_count != 4) /* validate column count */
{
    fprintf(stderr, " invalid column count returned by MySQL\n");
    exit(0);
}

/* Execute the SELECT query */
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute(), failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Bind the result buffers for all 4 columns before fetching them */

memset(bind, 0, sizeof(bind));

/* INTEGER COLUMN */
bind[0].buffer_type= MYSQL_TYPE_LONG;
bind[0].buffer= (char *)&int_data;
bind[0].is_null= &is_null[0];
bind[0].length= &length[0];
bind[0].error= &error[0];

/* STRING COLUMN */
bind[1].buffer_type= MYSQL_TYPE_STRING;
bind[1].buffer= (char *)str_data;
bind[1].buffer_length= STRING_SIZE;
bind[1].is_null= &is_null[1];
bind[1].length= &length[1];
bind[1].error= &error[1];

/* SMALLINT COLUMN */
bind[2].buffer_type= MYSQL_TYPE_SHORT;
bind[2].buffer= (char *)&small_data;
bind[2].is_null= &is_null[2];
bind[2].length= &length[2];
bind[2].error= &error[2];

/* TIMESTAMP COLUMN */
bind[3].buffer_type= MYSQL_TYPE_TIMESTAMP;
```

```
bind[3].buffer= (char *)&ts;
bind[3].is_null= &is_null[3];
bind[3].length= &length[3];
bind[3].error= &error[3];

/* Bind the result buffers */
if (mysql_stmt_bind_result(stmt, bind))
{
    fprintf(stderr, " mysql_stmt_bind_result() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Now buffer all results to client (optional step) */
if (mysql_stmt_store_result(stmt))
{
    fprintf(stderr, " mysql_stmt_store_result() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Fetch all rows */
row_count= 0;
fprintf(stdout, "Fetching results ...\n");
while (!mysql_stmt_fetch(stmt))
{
    row_count++;
    fprintf(stdout, " row %d\n", row_count);

    /* column 1 */
    fprintf(stdout, " column1 (integer) : ");
    if (is_null[0])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %d(%ld)\n", int_data, length[0]);

    /* column 2 */
    fprintf(stdout, " column2 (string) : ");
    if (is_null[1])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %s(%ld)\n", str_data, length[1]);

    /* column 3 */
    fprintf(stdout, " column3 (smallint) : ");
    if (is_null[2])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %d(%ld)\n", small_data, length[2]);

    /* column 4 */
    fprintf(stdout, " column4 (timestamp): ");
    if (is_null[3])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %04d-%02d-%02d %02d:%02d:%02d (%ld)\n",
            ts.year, ts.month, ts.day,
            ts.hour, ts.minute, ts.second,
            length[3]);
    fprintf(stdout, "\n");
}

/* Validate rows fetched */
fprintf(stdout, " total rows fetched: %d\n", row_count);
if (row_count != 2)
{
    fprintf(stderr, " MySQL failed to return all rows\n");
    exit(0);
}

/* Free the prepared result metadata */
mysql_free_result(prepare_meta_result);
```



```
/* Close the statement */
if (mysql_stmt_close(stmt))
{
    fprintf(stderr, " failed while closing the statement\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}
```

### 21.8.10.12. `mysql_stmt_fetch_column()`

```
int mysql_stmt_fetch_column(MYSQL_STMT *stmt, MYSQL_BIND *bind, unsigned int
column, unsigned long offset)
```

#### Description

Fetch one column from the current result set row. `bind` provides the buffer where data should be placed. It should be set up the same way as for `mysql_stmt_bind_result()`. `column` indicates which column to fetch. The first column is numbered 0. `offset` is the offset within the data value at which to begin retrieving data. This can be used for fetching the data value in pieces. The beginning of the value is offset 0.

#### Return Values

Zero if the value was fetched successfully. Nonzero if an error occurred.

#### Errors

- [CR\\_INVALID\\_PARAMETER\\_NO \[3720\]](#)  
Invalid column number.
- [CR\\_NO\\_DATA \[3721\]](#)  
The end of the result set has already been reached.

### 21.8.10.13. `mysql_stmt_field_count()`

```
unsigned int mysql_stmt_field_count(MYSQL_STMT *stmt)
```

#### Description

Returns the number of columns for the most recent statement for the statement handler. This value is zero for statements such as `INSERT` or `DELETE` that do not produce result sets.

`mysql_stmt_field_count()` can be called after you have prepared a statement by invoking `mysql_stmt_prepare()`.

#### Return Values

An unsigned integer representing the number of columns in a result set.

#### Errors

None.

### 21.8.10.14. `mysql_stmt_free_result()`

```
my_bool mysql_stmt_free_result(MYSQL_STMT *stmt)
```

#### Description

Releases memory associated with the result set produced by execution of the prepared statement. If there is a cursor open for the statement, `mysql_stmt_free_result()` closes it.

## Return Values

Zero if the result set was freed successfully. Nonzero if an error occurred.

## Errors

### 21.8.10.15. `mysql_stmt_init()`

```
MYSQL_STMT *mysql_stmt_init(MYSQL *mysql)
```

## Description

Create a `MYSQL_STMT` handle. The handle should be freed with `mysql_stmt_close(MYSQL_STMT *)`.

See also [Section 21.8.8, “C API Prepared Statement Data Structures”](#), for more information.

## Return Values

A pointer to a `MYSQL_STMT` structure in case of success. `NULL` if out of memory.

## Errors

- `CR_OUT_OF_MEMORY` [3718]

Out of memory.

### 21.8.10.16. `mysql_stmt_insert_id()`

```
my_ulonglong mysql_stmt_insert_id(MYSQL_STMT *stmt)
```

## Description

Returns the value generated for an `AUTO_INCREMENT` column by the prepared `INSERT` or `UPDATE` statement. Use this function after you have executed a prepared `INSERT` statement on a table which contains an `AUTO_INCREMENT` field.

See [Section 21.8.6.37, “`mysql\_insert\_id\(\)`”](#), for more information.

## Return Values

Value for `AUTO_INCREMENT` column which was automatically generated or explicitly set during execution of prepared statement, or value generated by `LAST_INSERT_ID(expr)` [1269] function. Return value is undefined if statement does not set `AUTO_INCREMENT` value.

## Errors

None.

### 21.8.10.17. `mysql_stmt_next_result()`

```
int mysql_stmt_next_result(MYSQL_STMT *mysql)
```

## Description

This function is used when you use prepared `CALL` statements to execute stored procedures, which can return multiple result sets. Use a loop that calls `mysql_stmt_next_result()` to determine whether there are more results. If a procedure has `OUT` or `INOUT` parameters, their values will be returned as a single-row result set following any other result sets. The values will appear in the order in which they are declared in the procedure parameter list.

`mysql_stmt_next_result()` returns a status to indicate whether more results exist. If `mysql_stmt_next_result()` returns an error, there are no more results.

Before each call to `mysql_stmt_next_result()`, you must call `mysql_stmt_free_result()` for the current result if it produced a result set (rather than just a result status).

After calling `mysql_stmt_next_result()` the state of the connection is as if you had called `mysql_stmt_execute()`. This means that you can call `mysql_stmt_bind_result()`, `mysql_stmt_affected_rows()`, and so forth.

It is also possible to test whether there are more results by calling `mysql_more_results()`. However, this function does not change the connection state, so if it returns true, you must still call `mysql_stmt_next_result()` to advance to the next result.

For an example that shows how to use `mysql_stmt_next_result()`, see [Section 21.8.19, “C API Support for Prepared CALL Statements”](#).

## Return Values

Return Value	Description
0	Successful and there are more results
-1	Successful and there are no more results
>0	An error occurred

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.10.18. `mysql_stmt_num_rows()`

```
my_ulonglong mysql_stmt_num_rows(MYSQL_STMT *stmt)
```

## Description

Returns the number of rows in the result set.

The use of `mysql_stmt_num_rows()` depends on whether you used `mysql_stmt_store_result()` to buffer the entire result set in the statement handle. If you use `mysql_stmt_store_result()`, `mysql_stmt_num_rows()` may be called immediately. Otherwise, the row count is unavailable unless you count the rows as you fetch them.

`mysql_stmt_num_rows()` is intended for use with statements that return a result set, such as `SELECT`. For statements such as `INSERT`, `UPDATE`, or `DELETE`, the number of affected rows can be obtained with `mysql_stmt_affected_rows()`.

**Return Values**

The number of rows in the result set.

**Errors**

None.

**21.8.10.19. `mysql_stmt_param_count()`**

```
unsigned long mysql_stmt_param_count(MYSQL_STMT *stmt)
```

**Description**

Returns the number of parameter markers present in the prepared statement.

**Return Values**

An unsigned long integer representing the number of parameters in a statement.

**Errors**

None.

**Example**

See the Example in [Section 21.8.10.10](#), "`mysql_stmt_execute()`".

**21.8.10.20. `mysql_stmt_param_metadata()`**

```
MYSQL_RES *mysql_stmt_param_metadata(MYSQL_STMT *stmt)
```

This function currently does nothing.

**Description****Return Values****Errors****21.8.10.21. `mysql_stmt_prepare()`**

```
int mysql_stmt_prepare(MYSQL_STMT *stmt, const char *stmt_str, unsigned long length)
```

**Description**

Given the statement handle returned by `mysql_stmt_init()`, prepares the SQL statement pointed to by the string `stmt_str` and returns a status value. The string length should be given by the `length` argument. The string must consist of a single SQL statement. You should not add a terminating semicolon ("`;`") or `\g` to the statement.

The application can include one or more parameter markers in the SQL statement by embedding question mark (`?`) characters into the SQL string at the appropriate positions.

The markers are legal only in certain places in SQL statements. For example, they are permitted in the `VALUES()` list of an `INSERT` statement (to specify column values for a row), or in a comparison with a column in a `WHERE` clause to specify a comparison value. However, they are not permitted for identifiers (such as table or column names), or to specify both operands of a binary operator such as the `=` equal sign. The latter restriction is necessary because it would be impossible to determine

the parameter type. In general, parameters are legal only in Data Manipulation Language (DML) statements, and not in Data Definition Language (DDL) statements.

The parameter markers must be bound to application variables using `mysql_stmt_bind_param()` before executing the statement.

Metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed. For more information, see [Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#).

## Return Values

Zero if the statement was prepared successfully. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)

Out of memory.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

If the prepare operation was unsuccessful (that is, `mysql_stmt_prepare()` returns nonzero), the error message can be obtained by calling `mysql_stmt_error()`.

## Example

See the Example in [Section 21.8.10.10, “mysql\\_stmt\\_execute\(\)”](#).

### 21.8.10.22. `mysql_stmt_reset()`

```
my_bool mysql_stmt_reset(MYSQL_STMT *stmt)
```

## Description

Resets a prepared statement on client and server to state after prepare. It resets the statement on the server, data sent using `mysql_stmt_send_long_data()`, unbuffered result sets and current errors. It does not clear bindings or stored result sets. Stored result sets will be cleared when executing the prepared statement (or closing it).

To re-prepare the statement with another query, use `mysql_stmt_prepare()`.

## Return Values

Zero if the statement was reset successfully. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)

Commands were executed in an improper order.

- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)

The MySQL server has gone away.

- [CR\\_SERVER\\_LOST \[3719\]](#)

The connection to the server was lost during the query

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### 21.8.10.23. `mysql_stmt_result_metadata()`

```
MYSQL_RES *mysql_stmt_result_metadata(MYSQL_STMT *stmt)
```

#### Description

If a statement passed to `mysql_stmt_prepare()` is one that produces a result set, `mysql_stmt_result_metadata()` returns the result set metadata in the form of a pointer to a `MYSQL_RES` structure that can be used to process the meta information such as number of fields and individual field information. This result set pointer can be passed as an argument to any of the field-based API functions that process result set metadata, such as:

- `mysql_num_fields()`
- `mysql_fetch_field()`
- `mysql_fetch_field_direct()`
- `mysql_fetch_fields()`
- `mysql_field_count()`
- `mysql_field_seek()`
- `mysql_field_tell()`
- `mysql_free_result()`

The result set structure should be freed when you are done with it, which you can do by passing it to `mysql_free_result()`. This is similar to the way you free a result set obtained from a call to `mysql_store_result()`.

The result set returned by `mysql_stmt_result_metadata()` contains only metadata. It does not contain any row results. The rows are obtained by using the statement handle with `mysql_stmt_fetch()`.

#### Return Values

A `MYSQL_RES` result structure. `NULL` if no meta information exists for the prepared query.

#### Errors

- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)

Out of memory.

- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)

An unknown error occurred.

### Example

See the Example in [Section 21.8.10.11](#), “`mysql_stmt_fetch()`”.

#### 21.8.10.24. `mysql_stmt_row_seek()`

```
MYSQL_ROW_OFFSET mysql_stmt_row_seek(MYSQL_STMT *stmt, MYSQL_ROW_OFFSET
offset)
```

### Description

Sets the row cursor to an arbitrary row in a statement result set. The `offset` value is a row offset that should be a value returned from `mysql_stmt_row_tell()` or from `mysql_stmt_row_seek()`. This value is not a row number; if you want to seek to a row within a result set by number, use `mysql_stmt_data_seek()` instead.

This function requires that the result set structure contains the entire result of the query, so `mysql_stmt_row_seek()` may be used only in conjunction with `mysql_stmt_store_result()`.

### Return Values

The previous value of the row cursor. This value may be passed to a subsequent call to `mysql_stmt_row_seek()`.

### Errors

None.

#### 21.8.10.25. `mysql_stmt_row_tell()`

```
MYSQL_ROW_OFFSET mysql_stmt_row_tell(MYSQL_STMT *stmt)
```

### Description

Returns the current position of the row cursor for the last `mysql_stmt_fetch()`. This value can be used as an argument to `mysql_stmt_row_seek()`.

You should use `mysql_stmt_row_tell()` only after `mysql_stmt_store_result()`.

### Return Values

The current offset of the row cursor.

### Errors

None.

#### 21.8.10.26. `mysql_stmt_send_long_data()`

```
my_bool mysql_stmt_send_long_data(MYSQL_STMT *stmt, unsigned int
parameter_number, const char *data, unsigned long length)
```

### Description

Enables an application to send parameter data to the server in pieces (or “chunks”). Call this function after `mysql_stmt_bind_param()` and before `mysql_stmt_execute()`. It can be called multiple times to send the parts of a character or binary data value for a column, which must be one of the `TEXT` or `BLOB` data types.

`parameter_number` indicates which parameter to associate the data with. Parameters are numbered beginning with 0. `data` is a pointer to a buffer containing data to be sent, and `length` indicates the number of bytes in the buffer.



#### Note

The next `mysql_stmt_execute()` call ignores the bind buffer for all parameters that have been used with `mysql_stmt_send_long_data()` since last `mysql_stmt_execute()` or `mysql_stmt_reset()`.

If you want to reset/forget the sent data, you can do it with `mysql_stmt_reset()`. See [Section 21.8.10.22](#), “`mysql_stmt_reset()`”.

As of MySQL 5.6.3, the `max_allowed_packet` [502] system variable controls the maximum size of parameter values that can be sent with `mysql_stmt_send_long_data()`.

## Return Values

Zero if the data is sent successfully to server. Nonzero if an error occurred.

## Errors

- [CR\\_INVALID\\_BUFFER\\_USE \[3720\]](#)  
The parameter does not have a string or binary type.
- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)  
Out of memory.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

## Example

The following example demonstrates how to send the data for a `TEXT` column in chunks. It inserts the data value `'MySQL - The most popular Open Source database'` into the `text_column` column. The `mysql` variable is assumed to be a valid connection handle.

```
#define INSERT_QUERY "INSERT INTO \  
                    test_long_data(text_column) VALUES(?)"  
  
MYSQL_BIND bind[1];  
long        length;  
  
stmt = mysql_stmt_init(mysql);  
if (!stmt)  
{  
    fprintf(stderr, "mysql_stmt_init(), out of memory\n");  
    exit(0);  
}  
if (mysql_stmt_prepare(stmt, INSERT_QUERY, strlen(INSERT_QUERY))  
{  
    fprintf(stderr, "\nmysql_stmt_prepare(), INSERT failed");  
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
```



```

    exit(0);
}
memset(bind, 0, sizeof(bind));
bind[0].buffer_type= MYSQL_TYPE_STRING;
bind[0].length= &length;
bind[0].is_null= 0;

/* Bind the buffers */
if (mysql_stmt_bind_param(stmt, bind))
{
    fprintf(stderr, "\n param bind failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Supply data in chunks to server */
if (mysql_stmt_send_long_data(stmt,0,"MySQL",5))
{
    fprintf(stderr, "\n send_long_data failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Supply the next piece of data */
if (mysql_stmt_send_long_data(stmt,0,
    " - The most popular Open Source database",40))
{
    fprintf(stderr, "\n send_long_data failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Now, execute the query */
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, "\n mysql_stmt_execute failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

```

### 21.8.10.27. `mysql_stmt_sqlstate()`

```
const char *mysql_stmt_sqlstate(MYSQL_STMT *stmt)
```

#### Description

For the statement specified by `stmt`, `mysql_stmt_sqlstate()` returns a null-terminated string containing the SQLSTATE error code for the most recently invoked prepared statement API function that can succeed or fail. The error code consists of five characters. "00000" means "no error." The values are specified by ANSI SQL and ODBC. For a list of possible values, see [Appendix C, Errors, Error Codes, and Common Problems](#).

Note that not all MySQL errors are yet mapped to SQLSTATE codes. The value "HY000" (general error) is used for unmapped errors.

#### Return Values

A null-terminated character string containing the SQLSTATE error code.

### 21.8.10.28. `mysql_stmt_store_result()`

```
int mysql_stmt_store_result(MYSQL_STMT *stmt)
```

#### Description

Result sets are produced by calling `mysql_stmt_execute()` to executed prepared statements for SQL statements such as `SELECT`, `SHOW`, `DESCRIBE`, and `EXPLAIN`. By default,

result sets for successfully executed prepared statements are not buffered on the client and `mysql_stmt_fetch()` fetches them one at a time from the server. To cause the complete result set to be buffered on the client, call `mysql_stmt_store_result()` after binding data buffers with `mysql_stmt_bind_result()` and before calling `mysql_stmt_fetch()` to fetch rows. (For an example, see [Section 21.8.10.11](#), “`mysql_stmt_fetch()`”.)

`mysql_stmt_store_result()` is optional for result set processing, unless you will call `mysql_stmt_data_seek()`, `mysql_stmt_row_seek()`, or `mysql_stmt_row_tell()`. Those functions require a seekable result set.

It is unnecessary to call `mysql_stmt_store_result()` after executing an SQL statement that does not produce a result set, but if you do, it does not harm or cause any notable performance problem. You can detect whether the statement produced a result set by checking if `mysql_stmt_result_metadata()` returns `NULL`. For more information, refer to [Section 21.8.10.23](#), “`mysql_stmt_result_metadata()`”.



#### Note

MySQL doesn't by default calculate `MYSQL_FIELD->max_length` for all columns in `mysql_stmt_store_result()` because calculating this would slow down `mysql_stmt_store_result()` considerably and most applications don't need `max_length`. If you want `max_length` to be updated, you can call `mysql_stmt_attr_set(MYSQL_STMT, STMT_ATTR_UPDATE_MAX_LENGTH, &flag)` to enable this. See [Section 21.8.10.3](#), “`mysql_stmt_attr_set()`”.

## Return Values

Zero if the results are buffered successfully. Nonzero if an error occurred.

## Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]](#)  
Commands were executed in an improper order.
- [CR\\_OUT\\_OF\\_MEMORY \[3718\]](#)  
Out of memory.
- [CR\\_SERVER\\_GONE\\_ERROR \[3718\]](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST \[3719\]](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR \[3718\]](#)  
An unknown error occurred.

## 21.8.11. C API Threaded Function Descriptions

To create a threaded client, use the functions described in the following sections. See also [Section 21.8.3.2](#), “[Writing C API Threaded Client Programs](#)”.

### 21.8.11.1. `my_init()`

```
void my_init(void)
```

## Description

`my_init()` initializes some global variables that MySQL needs. It also calls `mysql_thread_init()` for this thread.

It is necessary for `my_init()` to be called early in the initialization phase of a program's use of the MySQL library. However, `my_init()` is automatically called by `mysql_init()`, `mysql_library_init()`, `mysql_server_init()`, and `mysql_connect()`. If you ensure that your program invokes one of those functions before any other MySQL calls, there is no need to invoke `my_init()` explicitly.

To access the prototype for `my_init()`, your program should include these header files:

```
#include <my_global.h>
#include <my_sys.h>
```

## Return Values

None.

### 21.8.11.2. `mysql_thread_end()`

```
void mysql_thread_end(void)
```

## Description

This function needs to be called before calling `pthread_exit()` to free memory allocated by `mysql_thread_init()`.

`mysql_thread_end()` is not invoked automatically by the client library. It must be called explicitly to avoid a memory leak.

## Return Values

None.

### 21.8.11.3. `mysql_thread_init()`

```
my_bool mysql_thread_init(void)
```

## Description

This function must be called early within each created thread to initialize thread-specific variables. However, you may not necessarily need to invoke it explicitly: `mysql_thread_init()` is automatically called by `my_init()`, which itself is automatically called by `mysql_init()`, `mysql_library_init()`, `mysql_server_init()`, and `mysql_connect()`. If you invoke any of those functions, `mysql_thread_init()` will be called for you.

## Return Values

Zero if successful. Nonzero if an error occurred.

### 21.8.11.4. `mysql_thread_safe()`

```
unsigned int mysql_thread_safe(void)
```

## Description

This function indicates whether the client library is compiled as thread-safe.

## Return Values

1 if the client library is thread-safe, 0 otherwise.

## 21.8.12. C API Embedded Server Function Descriptions

MySQL applications can be written to use an embedded server. See [Section 21.7, “libmysqld, the Embedded MySQL Server Library”](#). To write such an application, you must link it against the `libmysqld` library by using the `-lmysqld` flag rather than linking it against the `libmysqlclient` client library by using the `-libmysqlclient` flag. However, the calls to initialize and finalize the library are the same whether you write a client application or one that uses the embedded server: Call `mysql_library_init()` to initialize the library and `mysql_library_end()` when you are done with it. See [Section 21.8.5, “C API Function Overview”](#).

### 21.8.12.1. `mysql_server_init()`

```
int mysql_server_init(int argc, char **argv, char **groups)
```

#### Description

This function initializes the MySQL library, which must be done before you call any other MySQL function. However, `mysql_server_init()` is deprecated and you should call `mysql_library_init()` instead. See [Section 21.8.6.40, “mysql\\_library\\_init\(\)”](#).

#### Return Values

Zero if successful. Nonzero if an error occurred.

### 21.8.12.2. `mysql_server_end()`

```
void mysql_server_end(void)
```

#### Description

This function finalizes the MySQL library, which should be done when you are done using the library. However, `mysql_server_end()` is deprecated and `mysql_library_end()` should be used instead. See [Section 21.8.6.39, “mysql\\_library\\_end\(\)”](#).

#### Return Values

None.

## 21.8.13. C API Client Plugin Functions

This section describes functions used for the client-side plugin API. They enable management of client plugins. For a description of the `st_mysql_client_plugin` structure used by these functions, see [Section 22.2.4.2.3, “Client Plugin Descriptors”](#).

It is unlikely that a client program needs to call the functions in this section. For example, a client that supports the use of authentication plugins normally causes a plugin to be loaded by calling `mysql_options()` to set the `MYSQL_DEFAULT_AUTH` and `MYSQL_PLUGIN_DIR` options:

```
char *plugin_dir = "path_to_plugin_dir";
char *default_auth = "plugin_name";

/* ... process command-line options ... */

mysql_options(&mysql, MYSQL_PLUGIN_DIR, plugin_dir);
mysql_options(&mysql, MYSQL_DEFAULT_AUTH, default_auth);
```

Typically, the program will also accept `--plugin-dir` and `--default-auth` options that enable users to override the default values.

### 21.8.13.1. `mysql_client_find_plugin()`

```
struct st_mysql_client_plugin *mysql_client_find_plugin(MYSQL *mysql, const
char *name, int type)
```

#### Description

Returns a pointer to a loaded plugin, loading the plugin first if necessary. An error occurs if the type is invalid or the plugin cannot be found or loaded.

Specify the parameters as follows:

- `mysql`: A pointer to a `MYSQL` structure. The plugin API does not require a connection to a MySQL server, but this structure must be properly initialized. The structure is used to obtain connection-related information.
- `name`: The plugin name.
- `type`: The plugin type.

#### Return Values

A pointer to the plugin for success. `NULL` if an error occurred.

#### Errors

To check for errors, call the `mysql_error()` or `mysql_errno()` function. See [Section 21.8.6.15](#), “`mysql_error()`”, and [Section 21.8.6.14](#), “`mysql_errno()`”.

#### Example

```
MYSQL mysql;
struct st_mysql_client_plugin *p;

if ((p = mysql_client_find_plugin(&mysql, "myplugin",
                                MYSQL_CLIENT_AUTHENTICATION_PLUGIN, 0)))
{
    printf("Plugin version: %d.%d.%d\n", p->version[0], p->version[1], p->version[2]);
}
```

### 21.8.13.2. `mysql_client_register_plugin()`

```
struct st_mysql_client_plugin *mysql_client_register_plugin(MYSQL *mysql,
struct st_mysql_client_plugin *plugin)
```

#### Description

Adds a plugin structure to the list of loaded plugins. An error occurs if the plugin is already loaded.

Specify the parameters as follows:

- `mysql`: A pointer to a `MYSQL` structure. The plugin API does not require a connection to a MySQL server, but this structure must be properly initialized. The structure is used to obtain connection-related information.
- `plugin`: A pointer to the plugin structure.

#### Return Values

A pointer to the plugin for success. `NULL` if an error occurred.

## Errors

To check for errors, call the `mysql_error()` or `mysql_errno()` function. See [Section 21.8.6.15](#), “`mysql_error()`”, and [Section 21.8.6.14](#), “`mysql_errno()`”.

### 21.8.13.3. `mysql_load_plugin()`

```
struct st_mysql_client_plugin *mysql_load_plugin(MYSQL *mysql, const char
*name, int type, int argc, ...)
```

## Description

Loads a MySQL client plugin, specified by name and type. An error occurs if the type is invalid or the plugin cannot be loaded.

It is not possible to load multiple plugins of the same type. An error occurs if you try to load a plugin of a type already loaded.

Specify the parameters as follows:

- `mysql`: A pointer to a `MYSQL` structure. The plugin API does not require a connection to a MySQL server, but this structure must be properly initialized. The structure is used to obtain connection-related information.
- `name`: The name of the plugin to load.
- `type`: The type of plugin to load, or `-1` to disable type checking. If type is not `-1`, only plugins matching the type are considered for loading.
- `argc`: The number of following arguments (0 if there are none). Interpretation of any following arguments depends on the plugin type.

Another way to cause plugins to be loaded is to set the `LIBMYSQL_PLUGINS` environment variable to a semicolon-separated list of plugin names. For example:

```
shell> export LIBMYSQL_PLUGINS="myplugin1;myplugin2"
```

Plugins named by `LIBMYSQL_PLUGINS` are loaded when the client program calls `mysql_library_init()`. No error is reported if problems occur loading these plugins.

As of MySQL 5.6.10, the `LIBMYSQL_PLUGIN_DIR` environment variable can be set to the path name of the directory in which to look for client plugins. This variable is used in two ways:

- During client plugin preloading, the value of the `--plugin-dir` option is not available, so client plugin loading fails unless the plugins are located in the hardwired default directory. If the plugins are located elsewhere, `LIBMYSQL_PLUGIN_DIR` environment variable can be set to the proper directory to enable plugin preloading to succeed.
- For explicit client plugin loading, the `mysql_load_plugin()` and `mysql_load_plugin_v()` C API functions use the `LIBMYSQL_PLUGIN_DIR` value if it exists and the `--plugin-dir` option was not given. If `--plugin-dir` is given, `mysql_load_plugin()` and `mysql_load_plugin_v()` ignore `LIBMYSQL_PLUGIN_DIR`.

## Return Values

A pointer to the plugin if it was loaded successfully. `NULL` if an error occurred.

## Errors

To check for errors, call the `mysql_error()` or `mysql_errno()` function. See [Section 21.8.6.15](#), “`mysql_error()`”, and [Section 21.8.6.14](#), “`mysql_errno()`”.

## Example

```
MYSQL mysql;

if(!mysql_load_plugin(&mysql, "myplugin",
                     MYSQL_CLIENT_AUTHENTICATION_PLUGIN, 0))
{
    fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
    exit(-1);
}
```

## See Also

See also [Section 21.8.13.3](#), “`mysql_load_plugin()`”, [Section 21.8.6.15](#), “`mysql_error()`”, [Section 21.8.6.14](#), “`mysql_errno()`”.

### 21.8.13.4. `mysql_load_plugin_v()`

```
struct st_mysql_client_plugin *mysql_load_plugin_v(MYSQL *mysql, const char
*name, int type, int argc, va_list args)
```

## Description

This function is equivalent to `mysql_load_plugin()`, but it accepts a `va_list` instead of a variable list of parameters.

## See Also

See also [Section 21.8.13.3](#), “`mysql_load_plugin()`”.

### 21.8.13.5. `mysql_plugin_options()`

```
int mysql_plugin_options(struct st_mysql_client_plugin *plugin, const char
*option, const void *value)
```

## Description

Passes an option type and value to a plugin. This function can be called multiple times to set several options. If the plugin does not have an option handler, an error occurs.

Specify the parameters as follows:

- `plugin`: A pointer to the plugin structure.
- `option`: The option to be set.
- `value`: A pointer to the option value.

## Return Values

Zero for success, 1 if an error occurred. If the plugin has an option handler, that handler should also return zero for success and 1 if an error occurred.

## 21.8.14. Common Questions and Problems When Using the C API

### 21.8.14.1. Why `mysql_store_result()` Sometimes Returns `NULL` After `mysql_query()` Returns Success

It is possible for `mysql_store_result()` to return `NULL` following a successful call to `mysql_query()`. When this happens, it means one of the following conditions occurred:

- There was a `malloc()` failure (for example, if the result set was too large).

- The data could not be read (an error occurred on the connection).
- The query returned no data (for example, it was an `INSERT`, `UPDATE`, or `DELETE`).

You can always check whether the statement should have produced a nonempty result by calling `mysql_field_count()`. If `mysql_field_count()` returns zero, the result is empty and the last query was a statement that does not return values (for example, an `INSERT` or a `DELETE`). If `mysql_field_count()` returns a nonzero value, the statement should have produced a nonempty result. See the description of the `mysql_field_count()` function for an example.

You can test for an error by calling `mysql_error()` or `mysql_errno()`.

### 21.8.14.2. What Results You Can Get from a Query

In addition to the result set returned by a query, you can also get the following information:

- `mysql_affected_rows()` returns the number of rows affected by the last query when doing an `INSERT`, `UPDATE`, or `DELETE`.

For a fast re-create, use `TRUNCATE TABLE`.

- `mysql_num_rows()` returns the number of rows in a result set. With `mysql_store_result()`, `mysql_num_rows()` may be called as soon as `mysql_store_result()` returns. With `mysql_use_result()`, `mysql_num_rows()` may be called only after you have fetched all the rows with `mysql_fetch_row()`.
- `mysql_insert_id()` returns the ID generated by the last query that inserted a row into a table with an `AUTO_INCREMENT` index. See [Section 21.8.6.37](#), “`mysql_insert_id()`”.
- Some queries (`LOAD DATA INFILE ...`, `INSERT INTO ... SELECT ...`, `UPDATE`) return additional information. The result is returned by `mysql_info()`. See the description for `mysql_info()` for the format of the string that it returns. `mysql_info()` returns a `NULL` pointer if there is no additional information.

### 21.8.14.3. How to Get the Unique ID for the Last Inserted Row

If you insert a record into a table that contains an `AUTO_INCREMENT` column, you can obtain the value stored into that column by calling the `mysql_insert_id()` function.

You can check from your C applications whether a value was stored in an `AUTO_INCREMENT` column by executing the following code (which assumes that you've checked that the statement succeeded). It determines whether the query was an `INSERT` with an `AUTO_INCREMENT` index:

```
if ((result = mysql_store_result(&mysql)) == 0 &&
    mysql_field_count(&mysql) == 0 &&
    mysql_insert_id(&mysql) != 0)
{
    used_id = mysql_insert_id(&mysql);
}
```

When a new `AUTO_INCREMENT` value has been generated, you can also obtain it by executing a `SELECT LAST_INSERT_ID()` statement with `mysql_query()` and retrieving the value from the result set returned by the statement.

When inserting multiple values, the last automatically incremented value is returned.

For `LAST_INSERT_ID()` [1269], the most recently generated ID is maintained in the server on a per-connection basis. It is not changed by another client. It is not even changed if you update another `AUTO_INCREMENT` column with a nonmagic value (that is, a value that is not `NULL` and not `0`). Using `LAST_INSERT_ID()` [1269] and `AUTO_INCREMENT` columns simultaneously from multiple clients is perfectly valid. Each client will receive the last inserted ID for the last statement *that* client executed.



If you want to use the ID that was generated for one table and insert it into a second table, you can use SQL statements like this:

```
INSERT INTO foo (auto,text)
  VALUES(NULL,'text');           # generate ID by inserting NULL
INSERT INTO foo2 (id,text)
  VALUES(LAST_INSERT_ID(),'text'); # use ID in second table
```

Note that `mysql_insert_id()` returns the value stored into an `AUTO_INCREMENT` column, whether that value is automatically generated by storing `NULL` or `0` or was specified as an explicit value. `LAST_INSERT_ID()` [1269] returns only automatically generated `AUTO_INCREMENT` values. If you store an explicit value other than `NULL` or `0`, it does not affect the value returned by `LAST_INSERT_ID()` [1269].

For more information on obtaining the last ID in an `AUTO_INCREMENT` column:

- For information on `LAST_INSERT_ID()` [1269], which can be used within an SQL statement, see [Section 12.14, “Information Functions”](#).
- For information on `mysql_insert_id()`, the function you use from within the C API, see [Section 21.8.6.37, “mysql\\_insert\\_id\(\)”](#).
- For information on obtaining the auto-incremented value when using Connector/J, see [Section 21.3.6.4, “Retrieving AUTO\\_INCREMENT Column Values through JDBC”](#).
- For information on obtaining the auto-incremented value when using Connector/ODBC, see [Section 21.1.7.1.1, “Obtaining Auto-Increment Values”](#).

## 21.8.15. Controlling Automatic Reconnection Behavior

The MySQL client library can perform an automatic reconnection to the server if it finds that the connection is down when you attempt to send a statement to the server to be executed. In this case, the library tries once to reconnect to the server and send the statement again.

In MySQL 5.6, auto-reconnect is disabled by default.

If it is important for your application to know that the connection has been dropped (so that it can exit or take action to adjust for the loss of state information), be sure that auto-reconnect is disabled. To ensure this, call `mysql_options()` with the `MYSQL_OPT_RECONNECT` option:

```
my_bool reconnect = 0;
mysql_options(&mysql, MYSQL_OPT_RECONNECT, &reconnect);
```

If the connection has gone down, the effect of `mysql_ping()` depends on the auto-reconnect state. If auto-reconnect is enabled, `mysql_ping()` performs a reconnect. Otherwise, it returns an error.

Some client programs might provide the capability of controlling automatic reconnection. For example, `mysql` reconnects by default, but the `--skip-reconnect` [252] option can be used to suppress this behavior.

If an automatic reconnection does occur (for example, as a result of calling `mysql_ping()`), there is no explicit indication of it. To check for reconnection, call `mysql_thread_id()` to get the original connection identifier before calling `mysql_ping()`, then call `mysql_thread_id()` again to see whether the identifier has changed.

Automatic reconnection can be convenient because you need not implement your own reconnect code, but if a reconnection does occur, several aspects of the connection state are reset on the server side and your application will not know about it. The connection-related state is affected as follows:

- Any active transactions are rolled back and autocommit mode is reset.
- All table locks are released.

- All `TEMPORARY` tables are closed (and dropped).
- Session variables are reinitialized to the values of the corresponding variables. This also affects variables that are set implicitly by statements such as `SET NAMES`.
- User variable settings are lost.
- Prepared statements are released.
- `HANDLER` variables are closed.
- The value of `LAST_INSERT_ID()` [1269] is reset to 0.
- Locks acquired with `GET_LOCK()` [1275] are released.
- The association of the client with the Performance Schema `threads` table row that determines connection thread instrumentation is lost. If the client reconnects after a disconnect, the session will be associated with a new row in the `threads` table and the thread monitoring state may be different. See [Section 20.9.9.3, “The threads Table”](#).

If the connection drops, it is possible that the session associated with the connection on the server side will still be running if the server has not yet detected that the client is no longer connected. In this case, any locks held by the original connection still belong to that session, so you may want to kill it by calling `mysql_kill()`.

## 21.8.16. C API Support for Multiple Statement Execution

By default, `mysql_query()` and `mysql_real_query()` interpret their statement string argument as a single statement to be executed, and you process the result according to whether the statement produces a result set (a set of rows, as for `SELECT`) or an affected-rows count (as for `INSERT`, `UPDATE`, and so forth).

MySQL 5.6 also supports the execution of a string containing multiple statements separated by semicolon (“;”) characters. This capability is enabled by special options that are specified either when you connect to the server with `mysql_real_connect()` or after connecting by calling `mysql_set_server_option()`.

Executing a multiple-statement string can produce multiple result sets or row-count indicators. Processing these results involves a different approach than for the single-statement case: After handling the result from the first statement, it is necessary to check whether more results exist and process them in turn if so. To support multiple-result processing, the C API includes the `mysql_more_results()` and `mysql_next_result()` functions. These functions are used at the end of a loop that iterates as long as more results are available. *Failure to process the result this way may result in a dropped connection to the server.*

Multiple-result processing also is required if you execute `CALL` statements for stored procedures. Results from a stored procedure have these characteristics:

- Statements within the procedure may produce result sets (for example, if it executes `SELECT` statements). These result sets are returned in the order that they are produced as the procedure executes.

In general, the caller cannot know how many result sets a procedure will return. Procedure execution may depend on loops or conditional statements that cause the execution path to differ from one call to the next. Therefore, you must be prepared to retrieve multiple results.

- The final result from the procedure is a status result that includes no result set. The status indicates whether the procedure succeeded or an error occurred.

The multiple statement and result capabilities can be used only with `mysql_query()` or `mysql_real_query()`. They cannot be used with the prepared statement interface. Prepared

statement handles are defined to work only with strings that contain a single statement. See [Section 21.8.7, “C API Prepared Statements”](#).

To enable multiple-statement execution and result processing, the following options may be used:

- The `mysql_real_connect()` function has a `flags` argument for which two option values are relevant:
  - `CLIENT_MULTI_RESULTS` enables the client program to process multiple results. This option *must* be enabled if you execute `CALL` statements for stored procedures that produce result sets. Otherwise, such procedures result in an error `Error 1312 (0A000): PROCEDURE proc_name can't return a result set in the given context`. In MySQL 5.6, `CLIENT_MULTI_RESULTS` is enabled by default.
  - `CLIENT_MULTI_STATEMENTS` enables `mysql_query()` and `mysql_real_query()` to execute statement strings containing multiple statements separated by semicolons. This option also enables `CLIENT_MULTI_RESULTS` implicitly, so a `flags` argument of `CLIENT_MULTI_STATEMENTS` to `mysql_real_connect()` is equivalent to an argument of `CLIENT_MULTI_STATEMENTS | CLIENT_MULTI_RESULTS`. That is, `CLIENT_MULTI_STATEMENTS` is sufficient to enable multiple-statement execution and all multiple-result processing.
- After the connection to the server has been established, you can use the `mysql_set_server_option()` function to enable or disable multiple-statement execution by passing it an argument of `MYSQL_OPTION_MULTI_STATEMENTS_ON` or `MYSQL_OPTION_MULTI_STATEMENTS_OFF`. Enabling multiple-statement execution with this function also enables processing of “simple” results for a multiple-statement string where each statement produces a single result, but is *not* sufficient to permit processing of stored procedures that produce result sets.

The following procedure outlines a suggested strategy for handling multiple statements:

1. Pass `CLIENT_MULTI_STATEMENTS` to `mysql_real_connect()`, to fully enable multiple-statement execution and multiple-result processing.
2. After calling `mysql_query()` or `mysql_real_query()` and verifying that it succeeds, enter a loop within which you process statement results.
3. For each iteration of the loop, handle the current statement result, retrieving either a result set or an affected-rows count. If an error occurs, exit the loop.
4. At the end of the loop, call `mysql_next_result()` to check whether another result exists and initiate retrieval for it if so. If no more results are available, exit the loop.

One possible implementation of the preceding strategy is shown following. The final part of the loop can be reduced to a simple test of whether `mysql_next_result()` returns nonzero. The code as written distinguishes between no more results and an error, which enables a message to be printed for the latter occurrence.

```

/* connect to server with the CLIENT_MULTI_STATEMENTS option */
if (mysql_real_connect (mysql, host_name, user_name, password,
    db_name, port_num, socket_name, CLIENT_MULTI_STATEMENTS) == NULL)
{
    printf("mysql_real_connect() failed\n");
    mysql_close(mysql);
    exit(1);
}

/* execute multiple statements */
status = mysql_query(mysql,
    "DROP TABLE IF EXISTS test_table;\
    CREATE TABLE test_table(id INT);\
    INSERT INTO test_table VALUES(10);\
    UPDATE test_table SET id=20 WHERE id=10;\

```

```

        SELECT * FROM test_table;\
        DROP TABLE test_table");
if (status)
{
    printf("Could not execute statement(s)");
    mysql_close(mysql);
    exit(0);
}

/* process each statement result */
do {
    /* did current statement return data? */
    result = mysql_store_result(mysql);
    if (result)
    {
        /* yes; process rows and free the result set */
        process_result_set(mysql, result);
        mysql_free_result(result);
    }
    else /* no result set or error */
    {
        if (mysql_field_count(mysql) == 0)
        {
            printf("%lld rows affected\n",
                mysql_affected_rows(mysql));
        }
        else /* some error occurred */
        {
            printf("Could not retrieve result set\n");
            break;
        }
    }
}
/* more results? -1 = no, >0 = error, 0 = yes (keep looping) */
if ((status = mysql_next_result(mysql)) > 0)
    printf("Could not execute statement\n");
} while (status == 0);

mysql_close(mysql);

```

## 21.8.17. C API Prepared Statement Problems

Here follows a list of the currently known problems with prepared statements:

- [TIME](#), [TIMESTAMP](#), and [DATETIME](#) do not support parts of seconds (for example, from [DATE\\_FORMAT\(\)](#) [1215]).
- When converting an integer to string, [ZEROFILL](#) is honored with prepared statements in some cases where the MySQL server doesn't print the leading zeros. (For example, with [MIN\(number-with-zerofill\)](#) [1284]).
- When converting a floating-point number to a string in the client, the rightmost digits of the converted value may differ slightly from those of the original value.
- Prepared statements use the query cache under the conditions described in [Section 8.9.3.1, "How the Query Cache Operates"](#).
- Prepared statements do not support multi-statements (that is, multiple statements within a single string separated by ";" characters).
- The capabilities of prepared [CALL](#) statements are described in [Section 21.8.19, "C API Support for Prepared CALL Statements"](#).

## 21.8.18. C API Prepared Statement Handling of Date and Time Values

The binary (prepared statement) protocol enables you to send and receive date and time values ([DATE](#), [TIME](#), [DATETIME](#), and [TIMESTAMP](#)), using the [MYSQL\\_TIME](#) structure. The members of this structure are described in [Section 21.8.8, "C API Prepared Statement Data Structures"](#).

To send temporal data values, create a prepared statement using `mysql_stmt_prepare()`. Then, before calling `mysql_stmt_execute()` to execute the statement, use the following procedure to set up each temporal parameter:

1. In the `MYSQL_BIND` structure associated with the data value, set the `buffer_type` member to the type that indicates what kind of temporal value you're sending. For `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` values, set `buffer_type` to `MYSQL_TYPE_DATE`, `MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATETIME`, or `MYSQL_TYPE_TIMESTAMP`, respectively.
2. Set the `buffer` member of the `MYSQL_BIND` structure to the address of the `MYSQL_TIME` structure in which you pass the temporal value.
3. Fill in the members of the `MYSQL_TIME` structure that are appropriate for the type of temporal value to be passed.

Use `mysql_stmt_bind_param()` to bind the parameter data to the statement. Then you can call `mysql_stmt_execute()`.

To retrieve temporal values, the procedure is similar, except that you set the `buffer_type` member to the type of value you expect to receive, and the `buffer` member to the address of a `MYSQL_TIME` structure into which the returned value should be placed. Use `mysql_stmt_bind_result()` to bind the buffers to the statement after calling `mysql_stmt_execute()` and before fetching the results.

Here is a simple example that inserts `DATE`, `TIME`, and `TIMESTAMP` data. The `mysql` variable is assumed to be a valid connection handle.

```

MYSQL_TIME  ts;
MYSQL_BIND  bind[3];
MYSQL_STMT  *stmt;

strmov(query, "INSERT INTO test_table(date_field, time_field, \
              timestamp_field) VALUES(?,?,?);");

stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, "mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(mysql, query, strlen(query)))
{
    fprintf(stderr, "\n mysql_stmt_prepare(), INSERT failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* set up input buffers for all 3 parameters */
bind[0].buffer_type= MYSQL_TYPE_DATE;
bind[0].buffer= (char *)&ts;
bind[0].is_null= 0;
bind[0].length= 0;
...
bind[1]= bind[2]= bind[0];
...

mysql_stmt_bind_param(stmt, bind);

/* supply the data to be sent in the ts structure */
ts.year= 2002;
ts.month= 02;
ts.day= 03;

ts.hour= 10;
ts.minute= 45;
ts.second= 20;

mysql_stmt_execute(stmt);

```

## 21.8.19. C API Support for Prepared `CALL` Statements

This section describes prepared-statement support in the C API for stored procedures executed using `CALL` statements:

In MySQL 5.6, stored procedures executed using prepared `CALL` statements can be used in the following ways:

- A stored procedure can produce any number of result sets. The number of columns and the data types of the columns need not be the same for all result sets.
- The final values of `OUT` and `INOUT` parameters are available to the calling application after the procedure returns. These parameters are returned as an extra single-row result set following any result sets produced by the procedure itself. The row contains the values of the `OUT` and `INOUT` parameters in the order in which they are declared in the procedure parameter list.

The following discussion shows how to use these capabilities through the C API for prepared statements. To use prepared `CALL` statements through the `PREPARE` and `EXECUTE` statements, see [Section 13.2.1, “CALL Syntax”](#).

If an application might be compiled or executed in a context where a version of MySQL older than 5.5.3 is used, prepared `CALL` capabilities for multiple result sets and `OUT` or `INOUT` parameters might not be available:

- For the client side, the application will not compile unless the libraries are from MySQL 5.5.3 or higher (the API function and symbols introduced in that version will not be present).
- To verify at runtime that the server is recent enough, a client can use this test:

```
if (mysql_get_server_version(mysql) < 50503)
{
    fprintf(stderr,
           "Server does not support required CALL capabilities\n");
    mysql_close(mysql);
    exit (1);
}
```

An application that executes a prepared `CALL` statement should use a loop that fetches a result and then invokes `mysql_stmt_next_result()` to determine whether there are more results. The results consist of any result sets produced by the stored procedure followed by a final status value that indicates whether the procedure terminated successfully.

If the procedure has `OUT` or `INOUT` parameters, the result set preceding the final status value contains their values. To determine whether a result set contains parameter values, test whether the `SERVER_PS_OUT_PARAMS` bit is set in the `server_status` member of the `MYSQL` connection handler:

```
mysql->server_status & SERVER_PS_OUT_PARAMS
```

The following example uses a prepared `CALL` statement to execute a stored procedure that produces multiple result sets and that provides parameter values back to the caller by means of `OUT` and `INOUT` parameters. The procedure takes parameters of all three types (`IN`, `OUT`, `INOUT`), displays their initial values, assigns new values, displays the updated values, and returns. The expected return information from the procedure therefore consists of multiple result sets and a final status:

- One result set from a `SELECT` that displays the initial parameter values: `10`, `NULL`, `30`. (The `OUT` parameter is assigned a value by the caller, but this assignment is expected to be ineffective: `OUT` parameters are seen as `NULL` within a procedure until assigned a value within the procedure.)
- One result set from a `SELECT` that displays the modified parameter values: `100`, `200`, `300`.
- One result set containing the final `OUT` and `INOUT` parameter values: `200`, `300`.

- A final status packet.

The code to execute the procedure:

```

MYSQL_STMT *stmt;
MYSQL_BIND ps_params[3]; /* input parameter buffers */
int         int_data[3]; /* input/output values */
my_bool     is_null[3]; /* output value nullability */
int         status;

/* set up stored procedure */
status = mysql_query(mysql, "DROP PROCEDURE IF EXISTS p1");
test_error(mysql, status);

status = mysql_query(mysql,
    "CREATE PROCEDURE p1("
    "  IN p_in INT, "
    "  OUT p_out INT, "
    "  INOUT p_inout INT) "
    "BEGIN "
    "  SELECT p_in, p_out, p_inout; "
    "  SET p_in = 100, p_out = 200, p_inout = 300; "
    "  SELECT p_in, p_out, p_inout; "
    "END");
test_error(mysql, status);

/* initialize and prepare CALL statement with parameter placeholders */
stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    printf("Could not initialize statement\n");
    exit(1);
}
status = mysql_stmt_prepare(stmt, "CALL p1(?, ?, ?)", 16);
test_stmt_error(stmt, status);

/* initialize parameters: p_in, p_out, p_inout (all INT) */
memset(ps_params, 0, sizeof (ps_params));

ps_params[0].buffer_type = MYSQL_TYPE_LONG;
ps_params[0].buffer = (char *) &int_data[0];
ps_params[0].length = 0;
ps_params[0].is_null = 0;

ps_params[1].buffer_type = MYSQL_TYPE_LONG;
ps_params[1].buffer = (char *) &int_data[1];
ps_params[1].length = 0;
ps_params[1].is_null = 0;

ps_params[2].buffer_type = MYSQL_TYPE_LONG;
ps_params[2].buffer = (char *) &int_data[2];
ps_params[2].length = 0;
ps_params[2].is_null = 0;

/* bind parameters */
status = mysql_stmt_bind_param(stmt, ps_params);
test_stmt_error(stmt, status);

/* assign values to parameters and execute statement */
int_data[0] = 10; /* p_in */
int_data[1] = 20; /* p_out */
int_data[2] = 30; /* p_inout */

status = mysql_stmt_execute(stmt);
test_stmt_error(stmt, status);

/* process results until there are no more */
do {
    int i;
    int num_fields; /* number of columns in result */
    MYSQL_FIELD *fields; /* for result set metadata */
    MYSQL_BIND *rs_bind; /* for output buffers */

```

```
/* the column count is > 0 if there is a result set */
/* 0 if the result is only the final status packet */
num_fields = mysql_stmt_field_count(stmt);

if (num_fields > 0)
{
    /* there is a result set to fetch */
    printf("Number of columns in result: %d\n", (int) num_fields);

    /* what kind of result set is this? */
    printf("Data: ");
    if(mysql->server_status & SERVER_PS_OUT_PARAMS)
        printf("this result set contains OUT/INOUT parameters\n");
    else
        printf("this result set is produced by the procedure\n");

    MYSQL_RES *rs_metadata = mysql_stmt_result_metadata(stmt);
    test_stmt_error(stmt, rs_metadata == NULL);

    fields = mysql_fetch_fields(rs_metadata);

    rs_bind = (MYSQL_BIND *) malloc(sizeof (MYSQL_BIND) * num_fields);
    if (!rs_bind)
    {
        printf("Cannot allocate output buffers\n");
        exit(1);
    }
    memset(rs_bind, 0, sizeof (MYSQL_BIND) * num_fields);

    /* set up and bind result set output buffers */
    for (i = 0; i < num_fields; ++i)
    {
        rs_bind[i].buffer_type = fields[i].type;
        rs_bind[i].is_null = &is_null[i];

        switch (fields[i].type)
        {
            case MYSQL_TYPE_LONG:
                rs_bind[i].buffer = (char *) &(int_data[i]);
                rs_bind[i].buffer_length = sizeof (int_data);
                break;

            default:
                fprintf(stderr, "ERROR: unexpected type: %d.\n", fields[i].type);
                exit(1);
        }
    }
}

status = mysql_stmt_bind_result(stmt, rs_bind);
test_stmt_error(stmt, status);

/* fetch and display result set rows */
while (1)
{
    status = mysql_stmt_fetch(stmt);

    if (status == 1 || status == MYSQL_NO_DATA)
        break;

    for (i = 0; i < num_fields; ++i)
    {
        switch (rs_bind[i].buffer_type)
        {
            case MYSQL_TYPE_LONG:
                if (*rs_bind[i].is_null)
                    printf(" val[%d] = NULL;", i);
                else
                    printf(" val[%d] = %ld;",
                        i, (long) *((int *) rs_bind[i].buffer));
                break;
        }
    }
}
```



```

        default:
            printf(" unexpected type (%d)\n",
                rs_bind[i].buffer_type);
        }
    }
    printf("\n");
}

mysql_free_result(rs_metadata); /* free metadata */
free(rs_bind);                /* free output buffers */
}
else
{
    /* no columns = final status packet */
    printf("End of procedure output\n");
}

/* more results? -1 = no, >0 = error, 0 = yes (keep looking) */
status = mysql_stmt_next_result(stmt);
if (status > 0)
    test_stmt_error(stmt, status);
} while (status == 0);

mysql_stmt_close(stmt);

```

Execution of the procedure should produce the following output:

```

Number of columns in result: 3
Data: this result set is produced by the procedure
val[0] = 10; val[1] = NULL; val[2] = 30;
Number of columns in result: 3
Data: this result set is produced by the procedure
val[0] = 100; val[1] = 200; val[2] = 300;
Number of columns in result: 2
Data: this result set contains OUT/INOUT parameters
val[0] = 200; val[1] = 300;
End of procedure output

```

The code uses two utility routines, `test_error()` and `test_stmt_error()`, to check for errors and terminate after printing diagnostic information if an error occurred:

```

static void test_error(MYSQL *mysql, int status)
{
    if (status)
    {
        fprintf(stderr, "Error: %s (errno: %d)\n",
            mysql_error(mysql), mysql_errno(mysql));
        exit(1);
    }
}

static void test_stmt_error(MYSQL_STMT *stmt, int status)
{
    if (status)
    {
        fprintf(stderr, "Error: %s (errno: %d)\n",
            mysql_stmt_error(stmt), mysql_stmt_errno(stmt));
        exit(1);
    }
}

```

## 21.9. MySQL PHP API

PHP is a server-side, HTML-embedded scripting language that may be used to create dynamic Web pages. It is available for most operating systems and Web servers, and can access most common databases, including MySQL. PHP may be run as a separate program or compiled as a module for use with a Web server.

PHP provides three different MySQL API extensions:

- [Section 21.9.3, “MySQL Improved Extension \(`Mysqli`\)”](#): Stands for “MySQL, Improved”; this extension is available as of PHP 5.0.0. It is intended for use with MySQL 4.1.1 and later. This extension fully supports the authentication protocol used in MySQL 5.0, as well as the Prepared Statements and Multiple Statements APIs. In addition, this extension provides an advanced, object-oriented programming interface.
- [Section 21.9.4, “MySQL Functions \(`PDO\_MYSQL`\) \(`MySQL \(PDO\)`\)”](#): Not its own API, but instead it's a MySQL driver for the PHP database abstraction layer PDO (PHP Data Objects). The PDO MySQL driver sits in the layer below PDO itself, and provides MySQL-specific functionality. This extension is available as of PHP 5.1.0.
- [Section 21.9.2, “Original MySQL API \(`Mysql`\)”](#): Available for PHP versions 4 and 5, this extension is intended for use with MySQL versions prior to MySQL 4.1. This extension does not support the improved authentication protocol used in MySQL 4.1, nor does it support prepared statements or multiple statements. To use this extension with MySQL 4.1, you will likely configure the MySQL server to set the `old_passwords` [517] system variable to 1 (see [Section C.5.2.4, “Client does not support authentication protocol”](#)).

This extension is deprecated as of PHP 5.5.0, and its use is highly discouraged. All users are encouraged to migrate to either `mysqli` or `PDO_MYSQL`. For further information, see [Section 21.9.1.3, “Choosing an API”](#).

The PHP distribution and documentation are available from the [PHP Web site](#).

*Portions of this section are Copyright (c) 1997-2012 the PHP Documentation Group* This material may be distributed only subject to the terms and conditions set forth in the Creative Commons Attribution 3.0 License or later. A copy of the Creative Commons Attribution 3.0 license is distributed with this manual. The latest version is presently available at <http://creativecommons.org/licenses/by/3.0/>.

## 21.9.1. Overview of the MySQL PHP drivers

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

### 21.9.1.1. Introduction

There are three PHP APIs for accessing the MySQL database. This guide explains the [terminology](#) used to describe each API, information about [choosing which API](#) to use, and also information to help choose which MySQL [library to use](#) with the API.

### 21.9.1.2. Terminology overview

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

This section provides an introduction to the options available to you when developing a PHP application that needs to interact with a MySQL database.

*What is an API?*

An Application Programming Interface, or API, defines the classes, methods, functions and variables that your application will need to call in order to carry out its desired task. In the case of PHP applications that need to communicate with databases the necessary APIs are usually exposed via PHP extensions.

APIs can be procedural or object-oriented. With a procedural API you call functions to carry out tasks, with the object-oriented API you instantiate classes and then call methods on the resulting objects. Of the two the latter is usually the preferred interface, as it is more modern and leads to better organized code.

When writing PHP applications that need to connect to the MySQL server there are several API options available. This document discusses what is available and how to select the best solution for your application.

### *What is a Connector?*

In the MySQL documentation, the term *connector* refers to a piece of software that allows your application to connect to the MySQL database server. MySQL provides connectors for a variety of languages, including PHP.

If your PHP application needs to communicate with a database server you will need to write PHP code to perform such activities as connecting to the database server, querying the database and other database-related functions. Software is required to provide the API that your PHP application will use, and also handle the communication between your application and the database server, possibly using other intermediate libraries where necessary. This software is known generically as a connector, as it allows your application to *connect* to a database server.

### *What is a Driver?*

A driver is a piece of software designed to communicate with a specific type of database server. The driver may also call a library, such as the MySQL Client Library or the MySQL Native Driver. These libraries implement the low-level protocol used to communicate with the MySQL database server.

By way of an example, the [PHP Data Objects \(PDO\) \[2985\]](#) database abstraction layer may use one of several database-specific drivers. One of the drivers it has available is the PDO MySQL driver, which allows it to interface with the MySQL server.

Sometimes people use the terms connector and driver interchangeably, this can be confusing. In the MySQL-related documentation the term “driver” is reserved for software that provides the database-specific part of a connector package.

### *What is an Extension?*

In the PHP documentation you will come across another term - *extension*. The PHP code consists of a core, with optional extensions to the core functionality. PHP's MySQL-related extensions, such as the [mysqli](#) extension, and the [mysql](#) extension, are implemented using the PHP extension framework.

An extension typically exposes an API to the PHP programmer, to allow its facilities to be used programmatically. However, some extensions which use the PHP extension framework do not expose an API to the PHP programmer.

The PDO MySQL driver extension, for example, does not expose an API to the PHP programmer, but provides an interface to the PDO layer above it.

The terms API and extension should not be taken to mean the same thing, as an extension may not necessarily expose an API to the programmer.

## 21.9.1.3. Choosing an API

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

PHP offers three different APIs to connect to MySQL. Below we show the APIs provided by the `mysql`, `mysqli`, and `PDO` extensions. Each code snippet creates a connection to a MySQL server running on "example.com" using the username "user" and the password "password". And a query is run to greet the user.

### **Example 21.13. Comparing the three MySQL APIs**

```
<?php
// mysqli
$mysqli = new mysqli("example.com", "user", "password", "database");
$result = $mysqli->query("SELECT 'Hello, dear MySQL user!' AS _message FROM DUAL");
$row = $result->fetch_assoc();
echo htmlentities($row['_message']);
```

```
// PDO
$pdo = new PDO('mysql:host=example.com;dbname=database', 'user', 'password');
$stmt = $pdo->query("SELECT 'Hello, dear MySQL user!' AS _message FROM DUAL");
$row = $stmt->fetch(PDO::FETCH_ASSOC);
echo htmlentities($row['_message']);

// mysql
$c = mysql_connect("example.com", "user", "password");
mysql_select_db("database");
$result = mysql_query("SELECT 'Hello, dear MySQL user!' AS _message FROM DUAL");
$row = mysql_fetch_assoc($result);
echo htmlentities($row['_message']);
?>
```

### Recommended API

It is recommended to use either the [mysqli](#) or [PDO\\_MySQL](#) extensions. It is not recommended to use the old [mysql](#) extension for new development, as it has been deprecated as of PHP 5.5.0 and will be removed in the future. A detailed feature comparison matrix is provided below. The overall performance of all three extensions is considered to be about the same. Although the performance of the extension contributes only a fraction of the total run time of a PHP web request. Often, the impact is as low as 0.1%.

### Feature comparison

	<b>ext/mysqli</b>	<b>PDO_MySQL</b>	<b>ext/mysql</b>
PHP version introduced	5.0	5.1	2.0
Included with PHP 5.x	Yes	Yes	Yes
Development status	Active	Active	Maintenance only
Lifecycle	Active	Active	Deprecated
Recommended for new projects	Yes	Yes	No
OOP Interface	Yes	Yes	No
Procedural Interface	Yes	No	Yes
API supports non-blocking, asynchronous queries with mysqlnd	Yes	No	No
Persistent Connections	Yes	Yes	Yes
API supports Charsets	Yes	Yes	Yes
API supports server-side Prepared Statements	Yes	Yes	No
API supports client-side Prepared Statements	No	Yes	No
API supports Stored Procedures	Yes	Yes	No
API supports Multiple Statements	Yes	Most	No
API supports Transactions	Yes	Yes	No
Transactions can be controlled with SQL	Yes	Yes	Yes
Supports all MySQL 5.1+ functionality	Yes	Most	No

### 21.9.1.4. Choosing a library

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqli`, `PDO_MySQL` and `mysql` PHP extensions are lightweight wrappers on top of a C client library. The extensions can either use the `mysqlnd` library or the `libmysqlclient` library. Choosing a library is a compile time decision.

The `mysqlnd` library is part of the PHP distribution since 5.3.0. It offers features like lazy connections and query caching, features that are not available with `libmysqlclient`, so using the built-in `mysqlnd` library is highly recommended. See the [mysqlnd documentation](#) for additional details, and a listing of features and functionality that it offers.

#### Example 21.14. Configure commands for using `mysqlnd` or `libmysqlclient`

```
// Recommended, compiles with mysqlnd
$ ./configure --with-mysqli=mysqlnd --with-pdo-mysql=mysqlnd --with-mysql=mysqlnd

// Not recommended, compiles with libmysqlclient
$ ./configure --with-mysqli=/path/to/mysql_config --with-pdo-mysql=/path/to/mysql_config --with-mysql=
```

#### Library feature comparison

It is recommended to use the `mysqlnd` library instead of the MySQL Client Server library (`libmysqlclient`). Both libraries are supported and constantly being improved.

	MySQL native driver ( <code>mysqlnd</code> )	MySQL client server library ( <code>libmysqlclient</code> )
Part of the PHP distribution	Yes	No
PHP version introduced	5.3.0	N/A
License	PHP License 3.01	Dual-License
Development status	Active	Active
Lifecycle	No end announced	No end announced
PHP 5.4 compile default (for all MySQL extensions)	Yes	No
PHP 5.3 compile default (for all MySQL extensions)	No	Yes
Compression protocol support	Yes (5.3.1+)	Yes
SSL support	Yes (5.3.3+)	Yes
Named pipe support	Yes (5.3.4+)	Yes
Non-blocking, asynchronous queries	Yes	No
Performance statistics	Yes	No
LOAD LOCAL INFILE respects the <a href="#">open_basedir directive</a>	Yes	No
Uses PHP's native memory management system (e.g., follows PHP memory limits)	Yes	No
Return numeric column as double (COM_QUERY)	Yes	No
Return numeric column as string (COM_QUERY)	Yes	Yes

	MySQL native driver ( <a href="#">mysqlnd</a> )	MySQL client server library ( <a href="#">libmysqlclient</a> )
Plugin API	Yes	Limited
Read/Write splitting for MySQL Replication	Yes, with plugin	No
Load Balancing	Yes, with plugin	No
Fail over	Yes, with plugin	No
Lazy connections	Yes, with plugin	No
Query caching	Yes, with plugin	No
Transparent query manipulations (E.g., auto-EXPLAIN or monitoring)	Yes, with plugin	No

### 21.9.1.5. Concepts

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

These concepts are specific to the MySQL drivers for PHP.

#### 21.9.1.5.1. Buffered and Unbuffered queries

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Queries are using the buffered mode by default. This means that query results are immediately transferred from the MySQL Server to PHP in is then kept in the memory of the PHP process. This allows additional operations like counting the number of rows, and moving (seeking) the current result pointer. It also allows issuing further queries on the same connection while working on the result set. The downside of the buffered mode is that larger result sets might require quite a lot memory. The memory will be kept occupied till all references to the result set are unset or the result set was explicitly freed, which will automatically happen during request end the latest. The terminology "store result" is also used for buffered mode, as the whole result set is stored at once.



#### Note

When using `libmysqlclient` as library PHP's memory limit won't count the memory used for result sets unless the data is fetched into PHP variables. With `mysqlnd` the memory accounted for will include the full result set.

Unbuffered MySQL queries execute the query and then return a resource while the data is still waiting on the MySQL server for being fetched. This uses less memory on the PHP-side, but can increase the load on the server. Unless the full result set was fetched from the server no further queries can be sent over the same connection. Unbuffered queries can also be referred to as "use result".

Following these characteristics buffered queries should be used in cases where you expect only a limited result set or need to know the amount of returned rows before reading all rows. Unbuffered mode should be used when you expect larger results.

Because buffered queries are the default, the examples below will demonstrate how to execute unbuffered queries with each API.

#### Example 21.15. Unbuffered query example: `mysqli`

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");
$result = $mysqli->query("SELECT Name FROM City", MYSQLI_USE_RESULT);

if ($result) {
    while ($row = $result->fetch_assoc()) {
```

```

        echo $row['Name'] . PHP_EOL;
    }
}
$uresult->close();
?>

```

**Example 21.16. Unbuffered query example: pdo\_mysql**

```

<?php
$pdo = new PDO("mysql:host=localhost;dbname=world", 'my_user', 'my_pass');
$pdo->setAttribute(PDO::MYSQL_ATTR_USE_BUFFERED_QUERY, false);

$uresult = $pdo->query("SELECT Name FROM City");
if ($uresult) {
    while ($row = $uresult->fetch(PDO::FETCH_ASSOC)) {
        echo $row['Name'] . PHP_EOL;
    }
}
?>

```

**Example 21.17. Unbuffered query example: mysql**

```

<?php
$conn = mysql_connect("localhost", "my_user", "my_pass");
$db    = mysql_select_db("world");

$uresult = mysql_unbuffered_query("SELECT Name FROM City");
if ($uresult) {
    while ($row = mysql_fetch_assoc($uresult)) {
        echo $row['Name'] . PHP_EOL;
    }
}
?>

```

**21.9.1.5.2. Character sets**

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Ideally a proper character set will be set at the server level, and doing this is described within the [Character Set Configuration](#) section of the MySQL Server manual. Alternatively, each MySQL API offers a method to set the character set at runtime.



**The character set and character escaping**

The character set should be understood and defined, as it has an affect on every action, and includes security implications. For example, the escaping mechanism (e.g., [mysqli\\_real\\_escape\\_string](#) for mysqli, [mysql\\_real\\_escape\\_string](#) for mysql, and [PDO::quote](#) for PDO\_MySQL) will adhere to this setting. It is important to realize that these functions will not use the character set that is defined with a query, so for example the following will not have an effect on them:

**Example 21.18. Problems with setting the character set with SQL**

```

<?php

$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

// Will not affect $mysqli->real_escape_string();
$mysqli->query("SET NAMES utf8");

```

```
// Will not affect $mysqli->real_escape_string();
$mysqli->query("SET CHARACTER SET utf8");

// But, this will affect $mysqli->real_escape_string();
$mysqli->set_charset('utf8');

?>
```

Below are examples that demonstrate how to properly alter the character set at runtime using each API.

### Example 21.19. Setting the character set example: [mysqli](#)

```
<?php

$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

if (!$mysqli->set_charset('utf8')) {
    printf("Error loading character set utf8: %s\n", $mysqli->error);
} else {
    printf("Current character set: %s\n", $mysqli->character_set_name());
}

print_r( $mysqli->get_charset() );

?>
```

### Example 21.20. Setting the character set example: [pdo\\_mysql](#)

Note: This only works as of PHP 5.3.6.

```
<?php
$pdo = new PDO("mysql:host=localhost;dbname=world;charset=utf8", 'my_user', 'my_pass');
?>
```

### Example 21.21. Setting the character set example: [mysql](#)

```
<?php

$conn = mysql_connect("localhost", "my_user", "my_pass");
$db   = mysql_select_db("world");

if (!mysql_set_charset('utf8', $conn)) {
    echo "Error: Unable to set the character set.\n";
    exit;
}

echo 'Your current character set is: ' . mysql_client_encoding($conn);

?>
```

## 21.9.2. Original MySQL API ([Mysql](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

This extension is deprecated as of PHP 5.5.0, and is not recommended for writing new code as it will be removed in the future. Instead, either the [mysqli](#) or [PDO\\_MySQL](#) extension should be used. See also the [MySQL API Overview](#) for further help while choosing a MySQL API.



These functions allow you to access MySQL database servers. More information about MySQL can be found at <http://www.mysql.com/>.

Documentation for MySQL can be found at <http://dev.mysql.com/doc/>.

### 21.9.2.1. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.2.1.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

In order to have these functions available, you must compile PHP with MySQL support.

#### 21.9.2.1.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

For compiling, simply use the `--with-mysql[=DIR]` configuration option where the optional `[DIR]` points to the MySQL installation directory.

Although this MySQL extension is compatible with MySQL 4.1.0 and greater, it doesn't support the extra functionality that these versions provide. For that, use the [MySQLi](#) extension.

If you would like to install the `mysql` extension along with the `mysqli` extension you have to use the same client library to avoid any conflicts.

##### 21.9.2.1.2.1. Installation on Linux Systems

Copyright 1997-2012 the PHP Documentation Group. [2906]

Note: `[DIR]` is the path to the MySQL client library files (*headers and libraries*), which can be downloaded from MySQL.

**Table 21.34. ext/mysql compile time support matrix**

PHP Version	Default	Configure Options: <a href="#">mysqlnd</a>	Configure Options: <a href="#">libmysqlclient</a>	Changelog
4.x.x	<code>libmysqlclient</code>	Not Available	<code>--without-mysql</code> to disable	MySQL enabled by default, MySQL client libraries are bundled
5.0.x, 5.1.x, 5.2.x	<code>libmysqlclient</code>	Not Available	<code>--with-mysql=[DIR]</code>	MySQL is no longer enabled by default, and the MySQL client libraries are <b>no longer bundled</b>
5.3.x	<code>libmysqlclient</code>	<code>--with-mysql=mysqlnd</code>	<code>--with-mysql=[DIR]</code>	<code>mysqlnd</code> is now available
5.4.x	<code>mysqlnd</code>	<code>--with-mysql</code>	<code>--with-mysql=[DIR]</code>	<code>mysqlnd</code> is now the default

##### 21.9.2.1.2.2. Installation on Windows Systems

Copyright 1997-2012 the PHP Documentation Group. [2906]

## 21.9.2.1.2.2.1. PHP 4

Copyright 1997-2012 the PHP Documentation Group. [2906]

The PHP MySQL extension is compiled into PHP.

## 21.9.2.1.2.2.2. PHP 5.0.x, 5.1.x, 5.2.x

Copyright 1997-2012 the PHP Documentation Group. [2906]

MySQL is no longer enabled by default, so the `php_mysql.dll` DLL must be enabled inside of `php.ini`. Also, PHP needs access to the MySQL client library. A file named `libmysql.dll` is included in the Windows PHP distribution and in order for PHP to talk to MySQL this file needs to be available to the Windows systems `PATH`. See the FAQ titled "How do I add my PHP directory to the `PATH` on Windows" for information on how to do this. Although copying `libmysql.dll` to the Windows system directory also works (because the system directory is by default in the system's `PATH`), it's not recommended.

As with enabling any PHP extension (such as `php_mysql.dll`), the PHP directive `extension_dir` should be set to the directory where the PHP extensions are located. See also the [Manual Windows Installation Instructions](#). An example `extension_dir` value for PHP 5 is `c:\php\ext`

**Note**

If when starting the web server an error similar to the following occurs: "Unable to load dynamic library './php\_mysql.dll'", this is because `php_mysql.dll` and/or `libmysql.dll` cannot be found by the system.

## 21.9.2.1.2.2.3. PHP 5.3.0+

Copyright 1997-2012 the PHP Documentation Group. [2906]

The [MySQL Native Driver](#) is enabled by default. Include `php_mysql.dll`, but `libmysql.dll` is no longer required or used.

## 21.9.2.1.2.3. MySQL Installation Notes

Copyright 1997-2012 the PHP Documentation Group. [2906]

**Warning**

Crashes and startup problems of PHP may be encountered when loading this extension in conjunction with the `recode` extension. See the [recode](#) extension for more information.

**Note**

If you need charsets other than *latin* (default), you have to install external (not bundled) `libmysqlclient` with compiled charset support.

## 21.9.2.1.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.35. MySQL Configuration Options**

Name	Default	Changeable	Changelog
<a href="#">mysql.allow_local_infile</a>	"1"	PHP_INI_SYSTEM	
<a href="#">mysql.allow_persistent</a>	"1"	PHP_INI_SYSTEM	
<a href="#">mysql.max_persistent</a>	"-1"	PHP_INI_SYSTEM	

Name	Default	Changeable	Changelog
<a href="#">mysql.max_links</a>	"-1"	PHP_INI_SYSTEM	
<a href="#">mysql.trace_mode</a>	"0"	PHP_INI_ALL	Available since PHP 4.3.0.
<a href="#">mysql.default_port</a>	NULL	PHP_INI_ALL	
<a href="#">mysql.default_socket</a>	NULL	PHP_INI_ALL	Available since PHP 4.0.1.
<a href="#">mysql.default_host</a>	NULL	PHP_INI_ALL	
<a href="#">mysql.default_user</a>	NULL	PHP_INI_ALL	
<a href="#">mysql.default_password</a>	NULL	PHP_INI_ALL	
<a href="#">mysql.connect_timeout</a>	"60"	PHP_INI_ALL	PHP_INI_SYSTEM in PHP <= 4.3.2. Available since PHP 4.3.0.

For further details and definitions of the PHP\_INI\_\* modes, see the <http://www.php.net/manual/en/configuration.changes.modes>.

Here's a short explanation of the configuration directives.

<a href="#">mysql.allow_local_infile</a> integer	Allow accessing, from PHP's perspective, local files with LOAD DATA statements
<a href="#">mysql.allow_persistent</a> boolean	Whether to allow <a href="#">persistent connections</a> to MySQL.
<a href="#">mysql.max_persistent</a> integer	The maximum number of persistent MySQL connections per process.
<a href="#">mysql.max_links</a> integer	The maximum number of MySQL connections per process, including persistent connections.
<a href="#">mysql.trace_mode</a> boolean	Trace mode. When <a href="#">mysql.trace_mode</a> is enabled, warnings for table/index scans, non free result sets, and SQL-Errors will be displayed. (Introduced in PHP 4.3.0)
<a href="#">mysql.default_port</a> string	The default TCP port number to use when connecting to the database server if no other port is specified. If no default is specified, the port will be obtained from the <code>MYSQL_TCP_PORT</code> environment variable, the <code>mysql-tcp</code> entry in <code>/etc/services</code> or the compile-time <code>MYSQL_PORT</code> constant, in that order. Win32 will only use the <code>MYSQL_PORT</code> constant.
<a href="#">mysql.default_socket</a> string	The default socket name to use when connecting to a local database server if no other socket name is specified.
<a href="#">mysql.default_host</a> string	The default server host to use when connecting to the database server if no other host is specified. Doesn't apply in <a href="#">SQL safe mode</a> .
<a href="#">mysql.default_user</a> string	The default user name to use when connecting to the database server if no other name is specified. Doesn't apply in <a href="#">SQL safe mode</a> .
<a href="#">mysql.default_password</a> string	The default password to use when connecting to the database server if no other password is specified. Doesn't apply in <a href="#">SQL safe mode</a> .
<a href="#">mysql.connect_timeout</a> integer	Connect timeout in seconds. On Linux this timeout is also used for waiting for the first answer from the server.

#### 21.9.2.1.4. Resource Types

Copyright 1997-2012 the PHP Documentation Group. [2906]

There are two resource types used in the MySQL module. The first one is the link identifier for a database connection, the second a resource which holds the result of a query.

#### 21.9.2.2. Changelog

Copyright 1997-2012 the PHP Documentation Group. [2906]

The following changes have been made to classes/functions/methods of this extension.

##### General Changelog for the `ext/mysql` extension

This changelog references the `ext/mysql` extension.

##### Changes to existing functions

The following list is a compilation of changelog entries from the `ext/mysql` functions.

##### Global `ext/mysql` changes

The following is a list of changes to the entire `ext/mysql` extension.

Version	Description
5.5.0	This extension has been deprecated. Connecting to a MySQL database via <code>mysql_connect</code> , <code>mysql_pconnect</code> or an implicit connection via any other <code>mysql_*</code> function will generate an <code>E_DEPRECATED</code> error.
5.5.0	All of the old deprecated functions and aliases now emit <code>E_DEPRECATED</code> errors. These functions are:  <code>mysql()</code> , <code>mysql_fieldname()</code> , <code>mysql_fieldtable()</code> , <code>mysql_fieldlen()</code> , <code>mysql_fieldtype()</code> , <code>mysql_fieldflags()</code> , <code>mysql_selectdb()</code> , <code>mysql_createdb()</code> , <code>mysql_dropdb()</code> , <code>mysql_freeresult()</code> , <code>mysql_numfields()</code> , <code>mysql_numrows()</code> , <code>mysql_listdbs()</code> , <code>mysql_listtables()</code> , <code>mysql_listfields()</code> , <code>mysql_db_name()</code> , <code>mysql_dbname()</code> , <code>mysql_tablename()</code> , and <code>mysql_table_name()</code> .

#### 21.9.2.3. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

Since PHP 4.3.0 it is possible to specify additional client flags for the `mysql_connect` and `mysql_pconnect` functions. The following constants are defined:

**Table 21.36. MySQL client constants**

Constant	Description
<code>MYSQL_CLIENT_COMPRESS</code>	Use compression protocol
<code>MYSQL_CLIENT_IGNORE_SPACE</code>	Allow space after function names

Constant	Description
<code>MYSQL_CLIENT_INTERACTIVE</code>	Allow <code>interactive_timeout</code> seconds (instead of <code>wait_timeout</code> ) of inactivity before closing the connection.
<code>MYSQL_CLIENT_SSL</code>	Use SSL encryption. This flag is only available with version 4.x of the MySQL client library or newer. Version 3.23.x is bundled both with PHP 4 and Windows binaries of PHP 5.

The function `mysql_fetch_array` uses a constant for the different types of result arrays. The following constants are defined:

**Table 21.37. MySQL fetch constants**

Constant	Description
<code>MYSQL_ASSOC</code>	Columns are returned into the array having the fieldname as the array index.
<code>MYSQL_BOTH</code>	Columns are returned into the array having both a numerical index and the fieldname as the array index.
<code>MYSQL_NUM</code>	Columns are returned into the array having a numerical index to the fields. This index starts with 0, the first field in the result.

## 21.9.2.4. Examples

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

### 21.9.2.4.1. MySQL extension overview example

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

This simple example shows how to connect, execute a query, print resulting rows and disconnect from a MySQL database.

#### Example 21.22. MySQL extension overview example

```
<?php
// Connecting, selecting database
$link = mysql_connect('mysql_host', 'mysql_user', 'mysql_password')
    or die('Could not connect: ' . mysql_error());
echo 'Connected successfully';
mysql_select_db('my_database') or die('Could not select database');

// Performing SQL query
$query = 'SELECT * FROM my_table';
$result = mysql_query($query) or die('Query failed: ' . mysql_error());

// Printing results in HTML
echo "<table>\n";
while ($line = mysql_fetch_array($result, MYSQL_ASSOC)) {
    echo "\t<tr>\n";
    foreach ($line as $col_value) {
        echo "\t\t<td>$col_value</td>\n";
    }
    echo "\t</tr>\n";
}
echo "</table>\n";

// Free resultset
mysql_free_result($result);
```

```
// Closing connection
mysql_close($link);
?>
```

## 21.9.2.5. MySQL Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]



### Note

Most MySQL functions accept *link\_identifier* as the last optional parameter. If it is not provided, last opened connection is used. If it doesn't exist, connection is tried to establish with default parameters defined in `php.ini`. If it is not successful, functions return `FALSE`.

### 21.9.2.5.1. `mysql_affected_rows`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_affected_rows`

Get number of affected rows in previous MySQL operation



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_affected_rows
PDOStatement::rowCount
```

### Description

```
int mysql_affected_rows(
    resource link_identifier
    = =NULL);
```

Get the number of affected rows by the last INSERT, UPDATE, REPLACE or DELETE query associated with *link\_identifier*.

### Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

### Return Values

Returns the number of affected rows on success, and -1 if the last query failed.

If the last query was a DELETE query with no WHERE clause, all of the records will have been deleted from the table but this function will return zero with MySQL versions prior to 4.1.2.

When using UPDATE, MySQL will not update columns where the new value is the same as the old value. This creates the possibility that `mysql_affected_rows` may not actually equal the number of rows matched, only the number of rows that were literally affected by the query.

The REPLACE statement first deletes the record with the same primary key and then inserts the new record. This function returns the number of deleted records plus the number of inserted records.

In the case of "INSERT ... ON DUPLICATE KEY UPDATE" queries, the return value will be [1](#) if an insert was performed, or [2](#) for an update of an existing row.

## Examples

### Example 21.23. [mysql\\_affected\\_rows](#) example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
mysql_select_db('mydb');

/* this should return the correct numbers of deleted records */
mysql_query('DELETE FROM mytable WHERE id < 10');
printf("Records deleted: %d\n", mysql_affected_rows());

/* with a where clause that is never true, it should return 0 */
mysql_query('DELETE FROM mytable WHERE 0');
printf("Records deleted: %d\n", mysql_affected_rows());
?>
```

The above example will output something similar to:

```
Records deleted: 10
Records deleted: 0
```

### Example 21.24. [mysql\\_affected\\_rows](#) example using transactions

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
mysql_select_db('mydb');

/* Update records */
mysql_query("UPDATE mytable SET used=1 WHERE id < 10");
printf ("Updated records: %d\n", mysql_affected_rows());
mysql_query("COMMIT");
?>
```

The above example will output something similar to:

```
Updated Records: 10
```

## Notes



### Transactions

If you are using transactions, you need to call [mysql\\_affected\\_rows](#) after your INSERT, UPDATE, or DELETE query, not after the COMMIT.

**SELECT Statements**

To retrieve the number of rows returned by a SELECT, it is possible to use `mysql_num_rows`.

**Cascaded Foreign Keys**

`mysql_affected_rows` does not count rows affected implicitly through the use of ON DELETE CASCADE and/or ON UPDATE CASCADE in foreign key constraints.

**See Also**

`mysql_num_rows`  
`mysql_info`

**21.9.2.5.2. `mysql_client_encoding`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_client_encoding`

Returns the name of the character set

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_character_set_name`

**Description**

```
string mysql_client_encoding(
    resource link_identifier
    = =NULL);
```

Retrieves the `character_set` variable from MySQL.

**Parameters**

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

Returns the default character set name for the current connection.

**Examples****Example 21.25. `mysql_client_encoding` example**

```
<?php
$link    = mysql_connect('localhost', 'mysql_user', 'mysql_password');
$charset = mysql_client_encoding($link);

echo "The current character set is: $charset\n";
```



```
?>
```

The above example will output something similar to:

```
The current character set is: latin1
```

### See Also

[mysql\\_set\\_charset](#)  
[mysql\\_real\\_escape\\_string](#)

### 21.9.2.5.3. *mysql\_close*

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_close](#)

Close MySQL connection



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_close](#)  
 PDO: Assign the value of `NULL` to the PDO object

### Description

```
bool mysql_close(
    resource link_identifier
    = =NULL);
```

*mysql\_close* closes the non-persistent connection to the MySQL server that's associated with the specified link identifier. If *link\_identifier* isn't specified, the last opened link is used.

Using *mysql\_close* isn't usually necessary, as non-persistent open links are automatically closed at the end of the script's execution. See also [freeing resources](#).

### Parameters

*link\_identifier* The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Examples

#### Example 21.26. *mysql\_close* example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
```

```
die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);
?>
```

The above example will output:

```
Connected successfully
```

## Notes



### Note

`mysql_close` will not close persistent links created by `mysql_pconnect`.

## See Also

[mysql\\_connect](#)  
[mysql\\_free\\_result](#)

### 21.9.2.5.4. [mysql\\_connect](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_connect](#)

Open a connection to a MySQL Server



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_connect](#)  
[PDO::\\_\\_construct](#)

## Description

```
resource mysql_connect(
    string server
        = =ini_get("mysql.default_host"),
    string username
        = =ini_get("mysql.default_user"),
    string password
        = =ini_get("mysql.default_password"),
    bool new_link
        = =false,
    int client_flags
        = =0);
```

Opens or reuses a connection to a MySQL server.

## Parameters

*server*

The MySQL server. It can also include a port number. e.g. "hostname:port" or a path to a local socket e.g. ":/path/to/socket" for the localhost.

If the PHP directive `mysql.default_host` is undefined (default), then the default value is 'localhost:3306'. In [SQL safe mode](#), this parameter is ignored and value 'localhost:3306' is always used.

*username*

The username. Default value is defined by `mysql.default_user`. In [SQL safe mode](#), this parameter is ignored and the name of the user that owns the server process is used.

*password*

The password. Default value is defined by `mysql.default_password`. In [SQL safe mode](#), this parameter is ignored and empty password is used.

*new\_link*

If a second call is made to `mysql_connect` with the same arguments, no new link will be established, but instead, the link identifier of the already opened link will be returned. The *new\_link* parameter modifies this behavior and makes `mysql_connect` always open a new link, even if `mysql_connect` was called before with the same parameters. In [SQL safe mode](#), this parameter is ignored.

*client\_flags*

The *client\_flags* parameter can be a combination of the following constants: 128 (enable `LOAD DATA LOCAL` handling), `MYSQL_CLIENT_SSL`, `MYSQL_CLIENT_COMPRESS`, `MYSQL_CLIENT_IGNORE_SPACE` or `MYSQL_CLIENT_INTERACTIVE`. Read the section about [Table 21.36, “MySQL client constants”](#) for further information. In [SQL safe mode](#), this parameter is ignored.

## Return Values

Returns a MySQL link identifier on success or `FALSE` on failure.

## Changelog

Version	Description
5.5.0	This function will generate an <code>E_DEPRECATED</code> error.
4.3.0	Added the <i>client_flags</i> parameter.
4.2.0	Added the <i>new_link</i> parameter.

## Examples

### Example 21.27. `mysql_connect` example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);
?>
```

### Example 21.28. `mysql_connect` example using `hostname:port` syntax

```
<?php
// we connect to example.com and port 3307
$link = mysql_connect('example.com:3307', 'mysql_user', 'mysql_password');
```

```

if (!$link) {
    die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);

// we connect to localhost at port 3307
$link = mysql_connect('127.0.0.1:3307', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);
?>

```

### Example 21.29. `mysql_connect` example using `"/path/to/socket"` syntax

```

<?php
// we connect to localhost and socket e.g. /tmp/mysql.sock

// variant 1: omit localhost
$link = mysql_connect('/:tmp/mysql', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);

// variant 2: with localhost
$link = mysql_connect('localhost:/tmp/mysql.sock', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
echo 'Connected successfully';
mysql_close($link);
?>

```

### Notes



#### Note

Whenever you specify "localhost" or "localhost:port" as server, the MySQL client library will override this and try to connect to a local socket (named pipe on Windows). If you want to use TCP/IP, use "127.0.0.1" instead of "localhost". If the MySQL client library tries to connect to the wrong local socket, you should set the correct path as `mysql.default_host string` in your PHP configuration and leave the server field blank.



#### Note

The link to the server will be closed as soon as the execution of the script ends, unless it's closed earlier by explicitly calling `mysql_close`.



#### Note

You can suppress the error message on failure by prepending a `@` to the function name.



#### Note

Error "Can't create TCP/IP socket (10106)" usually means that the `variables_order` configure directive doesn't contain character `E`. On Windows, if

the environment is not copied the `SYSTEMROOT` environment variable won't be available and PHP will have problems loading Winsock.

### See Also

`mysql_pconnect`  
`mysql_close`

#### 21.9.2.5.5. `mysql_create_db`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_create_db`

Create a MySQL database



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_query`  
`PDO::query`

### Description

```
bool mysql_create_db(
    string database_name,
    resource link_identifier
    = =NULL);
```

`mysql_create_db` attempts to create a new database on the server associated with the specified link identifier.

### Parameters

*database\_name*

The name of the database being created.

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Examples

#### Example 21.30. `mysql_create_db` alternative example

The function `mysql_create_db` is deprecated. It is preferable to use `mysql_query` to issue an `sql CREATE DATABASE` statement instead.

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
```

```
$sql = 'CREATE DATABASE my_db';
if (mysql_query($sql, $link)) {
    echo "Database my_db created successfully\n";
} else {
    echo 'Error creating database: ' . mysql_error() . "\n";
}
?>
```

The above example will output something similar to:

```
Database my_db created successfully
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_createdb](#)



### Note

This function will not be available if the MySQL extension was built against a MySQL 4.x client library.

## See Also

[mysql\\_query](#)  
[mysql\\_select\\_db](#)

### 21.9.2.5.6. [mysql\\_data\\_seek](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_data\\_seek](#)

Move internal result pointer



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_data\\_seek](#)  
[PDO::FETCH\\_ORI\\_ABS](#)

## Description

```
bool mysql_data_seek(
    resource result,
    int row_number);
```

[mysql\\_data\\_seek](#) moves the internal row pointer of the MySQL result associated with the specified result identifier to point to the specified row number. The next call to a MySQL fetch function, such as [mysql\\_fetch\\_assoc](#), would return that row.

*row\_number* starts at 0. The *row\_number* should be a value in the range from 0 to [mysql\\_num\\_rows](#) - 1. However if the result set is empty ([mysql\\_num\\_rows](#) == 0), a seek to 0 will fail with a [E\\_WARNING](#) and [mysql\\_data\\_seek](#) will return [FALSE](#).

## Parameters

- result* The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).
- row\_number* The desired row number of the new result pointer.

## Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

## Examples

### Example 21.31. [mysql\\_data\\_seek](#) example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
$db_selected = mysql_select_db('sample_db');
if (!$db_selected) {
    die('Could not select database: ' . mysql_error());
}
$query = 'SELECT last_name, first_name FROM friends';
$result = mysql_query($query);
if (!$result) {
    die('Query failed: ' . mysql_error());
}
/* fetch rows in reverse order */
for ($i = mysql_num_rows($result) - 1; $i >= 0; $i--) {
    if (!mysql_data_seek($result, $i)) {
        echo "Cannot seek to row $i: " . mysql_error() . "\n";
        continue;
    }

    if (!$row = mysql_fetch_assoc($result)) {
        continue;
    }

    echo $row['last_name'] . ' ' . $row['first_name'] . "<br />\n";
}

mysql_free_result($result);
?>
```

## Notes



### Note

The function [mysql\\_data\\_seek](#) can be used in conjunction only with [mysql\\_query](#), not with [mysql\\_unbuffered\\_query](#).

## See Also

[mysql\\_query](#)  
[mysql\\_num\\_rows](#)  
[mysql\\_fetch\\_row](#)  
[mysql\\_fetch\\_assoc](#)  
[mysql\\_fetch\\_array](#)  
[mysql\\_fetch\\_object](#)

### 21.9.2.5.7. [mysql\\_db\\_name](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_db\\_name](#)

Retrieves database name from the call to [mysql\\_list\\_dbs](#)



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

Query: `SELECT DATABASE()`

### Description

```
string mysql_db_name(
    resource result,
    int row,
    mixed field
    = =NULL);
```

Retrieve the database name from a call to [mysql\\_list\\_dbs](#).

### Parameters

<i>result</i>	The result pointer from a call to <a href="#">mysql_list_dbs</a> .
<i>row</i>	The index into the result set.
<i>field</i>	The field name.

### Return Values

Returns the database name on success, and `FALSE` on failure. If `FALSE` is returned, use [mysql\\_error](#) to determine the nature of the error.

### Changelog

Version	Description
5.5.0	The <a href="#">mysql_list_dbs</a> function is deprecated, and emits an <code>E_DEPRECATED</code> level error.

### Examples

#### Example 21.32. [mysql\\_db\\_name](#) example

```
<?php
error_reporting(E_ALL);

$link = mysql_connect('dbhost', 'username', 'password');
$db_list = mysql_list_dbs($link);

$i = 0;
$cnt = mysql_num_rows($db_list);
while ($i < $cnt) {
    echo mysql_db_name($db_list, $i) . "\n";
    $i++;
}
?>
```

### Notes



**Note**

For backward compatibility, the following deprecated alias may be used:  
`mysql_dbname`

**See Also**

`mysql_list_dbs`  
`mysql_tablename`

**21.9.2.5.8. `mysql_db_query`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_db_query`

Selects a database and executes a query on it

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_select_db` then the query  
`PDO::__construct`

**Description**

```
resource mysql_db_query(
    string database,
    string query,
    resource link_identifier
    = =NULL);
```

`mysql_db_query` selects a database, and executes a query on it.

**Warning**

This function has been *DEPRECATED* as of PHP 5.3.0. Relying on this feature is highly discouraged.

**Parameters**

<code>database</code>	The name of the database that will be selected.
<code>query</code>	The MySQL query. Data inside the query should be <a href="#">properly escaped</a> .
<code>link_identifier</code>	The MySQL connection. If the link identifier is not specified, the last link opened by <code>mysql_connect</code> is assumed. If no such link is found, it will try to create one as if <code>mysql_connect</code> was called with no arguments. If no connection is found or established, an <code>E_WARNING</code> level error is generated.

**Return Values**

Returns a positive MySQL result resource to the query result, or `FALSE` on error. The function also returns `TRUE` / `FALSE` for `INSERT/UPDATE/DELETE` queries to indicate success/failure.

**Changelog**

Version	Description
5.3.0	This function now throws an E_DEPRECATED notice.
4.0.6	This function is deprecated, do not use this function. Use <a href="#">mysql_select_db</a> and <a href="#">mysql_query</a> instead.

## Examples

### Example 21.33. [mysql\\_db\\_query](#) alternative example

```
<?php

if (!$link = mysql_connect('mysql_host', 'mysql_user', 'mysql_password')) {
    echo 'Could not connect to mysql';
    exit;
}

if (!mysql_select_db('mysql_dbname', $link)) {
    echo 'Could not select database';
    exit;
}

$sql      = 'SELECT foo FROM bar WHERE id = 42';
$result   = mysql_query($sql, $link);

if (!$result) {
    echo "DB Error, could not query the database\n";
    echo 'MySQL Error: ' . mysql_error();
    exit;
}

while ($row = mysql_fetch_assoc($result)) {
    echo $row['foo'];
}

mysql_free_result($result);

?>
```

## Notes



### Note

Be aware that this function does *NOT* switch back to the database you were connected to before. In other words, you can't use this function to *temporarily* run a sql query on another database, you would have to manually switch back. Users are strongly encouraged to use the [database.table](#) syntax in their sql queries or [mysql\\_select\\_db](#) instead of this function.

## See Also

[mysql\\_query](#)  
[mysql\\_select\\_db](#)

### 21.9.2.5.9. [mysql\\_drop\\_db](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_drop\\_db](#)

Drop (delete) a MySQL database

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

Execute a `DROP DATABASE` query

**Description**

```
bool mysql_drop_db(
    string database_name,
    resource link_identifier
    = =NULL);
```

`mysql_drop_db` attempts to drop (remove) an entire database from the server associated with the specified link identifier. This function is deprecated, it is preferable to use `mysql_query` to issue an `sql DROP DATABASE` statement instead.

**Parameters**

*database\_name*

The name of the database that will be deleted.

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**Examples****Example 21.34. `mysql_drop_db` alternative example**

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}

$sql = 'DROP DATABASE my_db';
if (mysql_query($sql, $link)) {
    echo "Database my_db was successfully dropped\n";
} else {
    echo 'Error dropping database: ' . mysql_error() . "\n";
}
?>
```

**Notes****Warning**

This function will not be available if the MySQL extension was built against a MySQL 4.x client library.

**Note**

For backward compatibility, the following deprecated alias may be used:  
`mysql_dropdb`

**See Also**[mysql\\_query](#)**21.9.2.5.10. [mysql\\_errno](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_errno](#)

Returns the numerical value of the error message from previous MySQL operation

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_errno](#)  
[PDO::errorCode](#)

**Description**

```
int mysql_errno(
    resource link_identifier
    = =NULL);
```

Returns the error number from the last MySQL function.

Errors coming back from the MySQL database backend no longer issue warnings. Instead, use [mysql\\_errno](#) to retrieve the error code. Note that this function only returns the error code from the most recently executed MySQL function (not including [mysql\\_error](#) and [mysql\\_errno](#)), so if you want to use it, make sure you check the value before calling another MySQL function.

**Parameters***link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

**Return Values**

Returns the error number from the last MySQL function, or 0 (zero) if no error occurred.

**Examples****Example 21.35. [mysql\\_errno](#) example**

```
<?php
$link = mysql_connect("localhost", "mysql_user", "mysql_password");

if (!mysql_select_db("nonexistentdb", $link)) {
    echo mysql_errno($link) . ": " . mysql_error($link) . "\n";
}

mysql_select_db("kossu", $link);
if (!mysql_query("SELECT * FROM nonexistenttable", $link)) {
    echo mysql_errno($link) . ": " . mysql_error($link) . "\n";
}
?>
```

The above example will output something similar to:

```
1049: Unknown database 'nonexistentdb'
1146: Table 'kossu.nonexistenttable' doesn't exist
```

### See Also

[mysql\\_error](#)  
[MySQL error codes](#)

#### 21.9.2.5.11. [mysql\\_error](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_error](#)

Returns the text of the error message from previous MySQL operation



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_error
PDO::errorInfo
```

### Description

```
string mysql_error(
    resource link_identifier
    = =NULL);
```

Returns the error text from the last MySQL function. Errors coming back from the MySQL database backend no longer issue warnings. Instead, use [mysql\\_error](#) to retrieve the error text. Note that this function only returns the error text from the most recently executed MySQL function (not including [mysql\\_error](#) and [mysql\\_errno](#)), so if you want to use it, make sure you check the value before calling another MySQL function.

### Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

### Return Values

Returns the error text from the last MySQL function, or '' (empty string) if no error occurred.

### Examples

#### Example 21.36. [mysql\\_error](#) example

```
<?php
$link = mysql_connect("localhost", "mysql_user", "mysql_password");
```

```
mysql_select_db("nonexistentdb", $link);
echo mysql_errno($link) . ": " . mysql_error($link) . "\n";

mysql_select_db("kossu", $link);
mysql_query("SELECT * FROM nonexistenttable", $link);
echo mysql_errno($link) . ": " . mysql_error($link) . "\n";
?>
```

The above example will output something similar to:

```
1049: Unknown database 'nonexistentdb'
1146: Table 'kossu.nonexistenttable' doesn't exist
```

### See Also

[mysql\\_errno](#)  
[MySQL error codes](#)

#### 21.9.2.5.12. [mysql\\_escape\\_string](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_escape\\_string](#)

Escapes a string for use in a `mysql_query`



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_escape_string
PDO::quote
```

### Description

```
string mysql_escape_string(
    string unescaped_string);
```

This function will escape the *unescaped\_string*, so that it is safe to place it in a `mysql_query`. This function is deprecated.

This function is identical to [mysql\\_real\\_escape\\_string](#) except that [mysql\\_real\\_escape\\_string](#) takes a connection handler and escapes the string according to the current character set. [mysql\\_escape\\_string](#) does not take a connection argument and does not respect the current charset setting.



#### Warning

This function has been *DEPRECATED* as of PHP 5.3.0. Relying on this feature is highly discouraged.

### Parameters

*unescaped\_string*                      The string that is to be escaped.

### Return Values

Returns the escaped string.

### Changelog

Version	Description
5.3.0	This function now throws an E_DEPRECATED notice.
4.3.0	This function became deprecated, do not use this function. Instead, use <a href="#">mysql_real_escape_string</a> .

### Examples

#### Example 21.37. [mysql\\_escape\\_string](#) example

```
<?php
$item = "Zak's Laptop";
$escaped_item = mysql_escape_string($item);
printf("Escaped string: %s\n", $escaped_item);
?>
```

The above example will output:

```
Escaped string: Zak\'s Laptop
```

### Notes



#### Note

[mysql\\_escape\\_string](#) does not escape `%` and `_`.

### See Also

[mysql\\_real\\_escape\\_string](#)

[addslashes](#)

The [magic\\_quotes\\_gpc](#) directive.

#### 21.9.2.5.13. [mysql\\_fetch\\_array](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_fetch\\_array](#)

Fetch a result row as an associative array, a numeric array, or both



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_fetch\\_array](#)

`PDOStatement::fetch`

### Description

```
array mysql_fetch_array(
    resource result,
    int result_type
    = =MYSQL_BOTH);
```

Returns an array that corresponds to the fetched row and moves the internal data pointer ahead.

### Parameters

*result*                      The result resource that is being evaluated. This result comes from a call to `mysql_query`.

*result\_type*                The type of array that is to be fetched. It's a constant and can take the following values: `MYSQL_ASSOC` , `MYSQL_NUM` , and `MYSQL_BOTH` .

### Return Values

Returns an array of strings that corresponds to the fetched row, or `FALSE` if there are no more rows. The type of returned array depends on how *result\_type* is defined. By using `MYSQL_BOTH` (default), you'll get an array with both associative and number indices. Using `MYSQL_ASSOC` , you only get associative indices (as `mysql_fetch_assoc` works), using `MYSQL_NUM` , you only get number indices (as `mysql_fetch_row` works).

If two or more columns of the result have the same field names, the last column will take precedence. To access the other column(s) of the same name, you must use the numeric index of the column or make an alias for the column. For aliased columns, you cannot access the contents with the original column name.

### Examples

#### Example 21.38. Query with aliased duplicate field names

```
SELECT table1.field AS foo, table2.field AS bar FROM table1, table2
```

#### Example 21.39. `mysql_fetch_array` with `MYSQL_NUM`

```
<?php
mysql_connect("localhost", "mysql_user", "mysql_password") or
    die("Could not connect: " . mysql_error());
mysql_select_db("mydb");

$result = mysql_query("SELECT id, name FROM mytable");

while ($row = mysql_fetch_array($result, MYSQL_NUM)) {
    printf("ID: %s Name: %s", $row[0], $row[1]);
}

mysql_free_result($result);
?>
```

#### Example 21.40. `mysql_fetch_array` with `MYSQL_ASSOC`

```
<?php
mysql_connect("localhost", "mysql_user", "mysql_password") or
    die("Could not connect: " . mysql_error());
mysql_select_db("mydb");

$result = mysql_query("SELECT id, name FROM mytable");
```



```
while ($row = mysql_fetch_array($result, MYSQL_ASSOC)) {
    printf("ID: %s Name: %s", $row["id"], $row["name"]);
}

mysql_free_result($result);
?>
```

**Example 21.41. `mysql_fetch_array` with `MYSQL_BOTH`**

```
<?php
mysql_connect("localhost", "mysql_user", "mysql_password") or
    die("Could not connect: " . mysql_error());
mysql_select_db("mydb");

$result = mysql_query("SELECT id, name FROM mytable");

while ($row = mysql_fetch_array($result, MYSQL_BOTH)) {
    printf("ID: %s Name: %s", $row[0], $row["name"]);
}

mysql_free_result($result);
?>
```

**Notes****Performance**

An important thing to note is that using `mysql_fetch_array` is *not significantly* slower than using `mysql_fetch_row`, while it provides a significant added value.

**Note**

Field names returned by this function are *case-sensitive*.

**Note**

This function sets NULL fields to the PHP `NULL` value.

**See Also**

[mysql\\_fetch\\_row](#)  
[mysql\\_fetch\\_assoc](#)  
[mysql\\_data\\_seek](#)  
[mysql\\_query](#)

**21.9.2.5.14. `mysql_fetch_assoc`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_fetch\\_assoc](#)

Fetch a result row as an associative array

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_fetch_assoc  
PDOStatement::fetch(PDO::FETCH_ASSOC)
```

## Description

```
array mysqli_fetch_assoc(  
    resource result);
```

Returns an associative array that corresponds to the fetched row and moves the internal data pointer ahead. `mysqli_fetch_assoc` is equivalent to calling `mysqli_fetch_array` with `MYSQLI_ASSOC` for the optional second parameter. It only returns an associative array.

## Parameters

*result* The result resource that is being evaluated. This result comes from a call to `mysqli_query`.

## Return Values

Returns an associative array of strings that corresponds to the fetched row, or `FALSE` if there are no more rows.

If two or more columns of the result have the same field names, the last column will take precedence. To access the other column(s) of the same name, you either need to access the result with numeric indices by using `mysqli_fetch_row` or add alias names. See the example at the `mysqli_fetch_array` description about aliases.

## Examples

### Example 21.42. An expanded `mysqli_fetch_assoc` example

```
<?php  
  
$conn = mysqli_connect("localhost", "mysql_user", "mysql_password");  
  
if (!$conn) {  
    echo "Unable to connect to DB: " . mysqli_error();  
    exit;  
}  
  
if (!mysqli_select_db("mydbname")) {  
    echo "Unable to select mydbname: " . mysqli_error();  
    exit;  
}  
  
$sql = "SELECT id as userid, fullname, userstatus  
        FROM   sometable  
        WHERE  userstatus = 1";  
  
$result = mysqli_query($sql);  
  
if (!$result) {  
    echo "Could not successfully run query ($sql) from DB: " . mysqli_error();  
    exit;  
}  
  
if (mysqli_num_rows($result) == 0) {  
    echo "No rows found, nothing to print so am exiting";  
    exit;  
}  
  
// While a row of data exists, put that row in $row as an associative array  
// Note: If you're expecting just one row, no need to use a loop  
// Note: If you put extract($row); inside the following loop, you'll  
//       then create $userid, $fullname, and $userstatus
```

```
while ($row = mysql_fetch_assoc($result)) {
    echo $row["userid"];
    echo $row["fullname"];
    echo $row["userstatus"];
}

mysql_free_result($result);

?>
```

## Notes



### Performance

An important thing to note is that using `mysql_fetch_assoc` is *not significantly* slower than using `mysql_fetch_row`, while it provides a significant added value.



### Note

Field names returned by this function are *case-sensitive*.



### Note

This function sets NULL fields to the PHP [NULL](#) value.

## See Also

[mysql\\_fetch\\_row](#)  
[mysql\\_fetch\\_array](#)  
[mysql\\_data\\_seek](#)  
[mysql\\_query](#)  
[mysql\\_error](#)

### 21.9.2.5.15. [mysql\\_fetch\\_field](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_fetch\\_field](#)

Get column information from a result and return as an object



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_fetch\\_field](#)  
`PDOStatement::getColumnMeta`

## Description

```
object mysql_fetch_field(
    resource result,
    int field_offset
    = 0);
```

Returns an object containing field information. This function can be used to obtain information about fields in the provided query result.

## Parameters

- result* The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).
- field\_offset* The numerical field offset. If the field offset is not specified, the next field that was not yet retrieved by this function is retrieved. The [field\\_offset](#) starts at 0.

## Return Values

Returns an object containing field information. The properties of the object are:

- name - column name
- table - name of the table the column belongs to, which is the alias name if one is defined
- max\_length - maximum length of the column
- not\_null - 1 if the column cannot be [NULL](#)
- primary\_key - 1 if the column is a primary key
- unique\_key - 1 if the column is a unique key
- multiple\_key - 1 if the column is a non-unique key
- numeric - 1 if the column is numeric
- blob - 1 if the column is a BLOB
- type - the type of the column
- unsigned - 1 if the column is unsigned
- zerofill - 1 if the column is zero-filled

## Examples

### Example 21.43. [mysql\\_fetch\\_field](#) example

```
<?php
$conn = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$conn) {
    die('Could not connect: ' . mysql_error());
}
mysql_select_db('database');
$result = mysql_query('select * from table');
if (!$result) {
    die('Query failed: ' . mysql_error());
}
/* get column metadata */
$i = 0;
while ($i < mysql_num_fields($result)) {
    echo "Information for column $i:<br />\n";
    $meta = mysql_fetch_field($result, $i);
    if (!$meta) {
        echo "No information available<br />\n";
    }
    echo "<pre>
blob:          $meta->blob
max_length:    $meta->max_length
multiple_key:  $meta->multiple_key
name:          $meta->name
not_null:      $meta->not_null
numeric:       $meta->numeric
```

```

primary_key: $meta->primary_key
table:       $meta->table
type:       $meta->type
unique_key: $meta->unique_key
unsigned:   $meta->unsigned
zerofill:   $meta->zerofill
</pre>";
    $i++;
}
mysql_free_result($result);
?>

```

## Notes



### Note

Field names returned by this function are *case-sensitive*.



### Note

If field or tablename are aliased in the SQL query the aliased name will be returned. The original name can be retrieved for instance by using [mysqli\\_result::fetch\\_field](#).

## See Also

[mysql\\_field\\_seek](#)

### 21.9.2.5.16. [mysql\\_fetch\\_lengths](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_fetch\\_lengths](#)

Get the length of each output in a result



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_fetch\\_lengths](#)  
[PDOStatement::getColumnMeta](#)

## Description

```

array mysql_fetch_lengths(
    resource result);

```

Returns an array that corresponds to the lengths of each field in the last row fetched by MySQL.

[mysql\\_fetch\\_lengths](#) stores the lengths of each result column in the last row returned by [mysql\\_fetch\\_row](#), [mysql\\_fetch\\_assoc](#), [mysql\\_fetch\\_array](#), and [mysql\\_fetch\\_object](#) in an array, starting at offset 0.

## Parameters

*result*

The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).

## Return Values

An array of lengths on success or `FALSE` on failure.

## Examples

### Example 21.44. A `mysql_fetch_lengths` example

```
<?php
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");
if (!$result) {
    echo 'Could not run query: ' . mysql_error();
    exit;
}
$row      = mysql_fetch_assoc($result);
$lengths = mysql_fetch_lengths($result);

print_r($row);
print_r($lengths);
?>
```

The above example will output something similar to:

```
Array
(
    [id] => 42
    [email] => user@example.com
)
Array
(
    [0] => 2
    [1] => 16
)
```

## See Also

`mysql_field_len`  
`mysql_fetch_row`  
`strlen`

### 21.9.2.5.17. `mysql_fetch_object`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_fetch_object`

Fetch a result row as an object



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_fetch_object`  
`PDOStatement::fetch(PDO::FETCH_OBJ)`

## Description

```
object mysql_fetch_object(
```

```
resource result,
string class_name,
array params);
```

Returns an object with properties that correspond to the fetched row and moves the internal data pointer ahead.

### Parameters

<i>result</i>	The result resource that is being evaluated. This result comes from a call to <code>mysql_query</code> .
<i>class_name</i>	The name of the class to instantiate, set the properties of and return. If not specified, a <code>stdClass</code> object is returned.
<i>params</i>	An optional array of parameters to pass to the constructor for <i>class_name</i> objects.

### Return Values

Returns an object with string properties that correspond to the fetched row, or `FALSE` if there are no more rows.

### Changelog

Version	Description
5.0.0	Added the ability to return as a different object.

### Examples

#### Example 21.45. `mysql_fetch_object` example

```
<?php
mysql_connect("hostname", "user", "password");
mysql_select_db("mydb");
$result = mysql_query("select * from mytable");
while ($row = mysql_fetch_object($result)) {
    echo $row->user_id;
    echo $row->fullname;
}
mysql_free_result($result);
?>
```

#### Example 21.46. `mysql_fetch_object` example

```
<?php
class foo {
    public $name;
}

mysql_connect("hostname", "user", "password");
mysql_select_db("mydb");

$result = mysql_query("select name from mytable limit 1");
$obj = mysql_fetch_object($result, 'foo');
var_dump($obj);
?>
```

### Notes

**Performance**

Speed-wise, the function is identical to `mysql_fetch_array`, and almost as quick as `mysql_fetch_row` (the difference is insignificant).

**Note**

`mysql_fetch_object` is similar to `mysql_fetch_array`, with one difference - an object is returned, instead of an array. Indirectly, that means that you can only access the data by the field names, and not by their offsets (numbers are illegal property names).

**Note**

Field names returned by this function are *case-sensitive*.

**Note**

This function sets NULL fields to the PHP `NULL` value.

**See Also**

`mysql_fetch_array`  
`mysql_fetch_assoc`  
`mysql_fetch_row`  
`mysql_data_seek`  
`mysql_query`

**21.9.2.5.18. `mysql_fetch_row`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_fetch_row`

Get a result row as an enumerated array

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_fetch_row  
PDOStatement::fetch(PDO::FETCH_NUM)
```

**Description**

```
array mysql_fetch_row(  
    resource result);
```

Returns a numerical array that corresponds to the fetched row and moves the internal data pointer ahead.

**Parameters**

*result*

The result resource that is being evaluated. This result comes from a call to `mysql_query`.

**Return Values**



Returns an numerical array of strings that corresponds to the fetched row, or `FALSE` if there are no more rows.

`mysql_fetch_row` fetches one row of data from the result associated with the specified result identifier. The row is returned as an array. Each result column is stored in an array offset, starting at offset 0.

## Examples

### Example 21.47. Fetching one row with `mysql_fetch_row`

```
<?php
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");
if (!$result) {
    echo 'Could not run query: ' . mysql_error();
    exit;
}
$row = mysql_fetch_row($result);

echo $row[0]; // 42
echo $row[1]; // the email value
?>
```

## Notes



### Note

This function sets NULL fields to the PHP `NULL` value.

## See Also

`mysql_fetch_array`  
`mysql_fetch_assoc`  
`mysql_fetch_object`  
`mysql_data_seek`  
`mysql_fetch_lengths`  
`mysql_result`

### 21.9.2.5.19. `mysql_field_flags`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_field_flags`

Get the flags associated with the specified field in a result



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_fetch_fields` [flags]  
`PDOStatement::getColumnMeta` [flags]

## Description

```
string mysql_field_flags(
    resource result,
```

```
int field_offset);
```

`mysql_field_flags` returns the field flags of the specified field. The flags are reported as a single word per flag separated by a single space, so that you can split the returned value using `explode`.

### Parameters

*result*

The result resource that is being evaluated. This result comes from a call to `mysql_query`.

*field\_offset*

The numerical field offset. The *field\_offset* starts at 0. If *field\_offset* does not exist, an error of level `E_WARNING` is also issued.

### Return Values

Returns a string of flags associated with the result or `FALSE` on failure.

The following flags are reported, if your version of MySQL is current enough to support them: `"not_null"`, `"primary_key"`, `"unique_key"`, `"multiple_key"`, `"blob"`, `"unsigned"`, `"zerofill"`, `"binary"`, `"enum"`, `"auto_increment"` and `"timestamp"`.

### Examples

#### Example 21.48. A `mysql_field_flags` example

```
<?php
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");
if (!$result) {
    echo 'Could not run query: ' . mysql_error();
    exit;
}
$flags = mysql_field_flags($result, 0);

echo $flags;
print_r(explode(' ', $flags));
?>
```

The above example will output something similar to:

```
not_null primary_key auto_increment
Array
(
    [0] => not_null
    [1] => primary_key
    [2] => auto_increment
)
```

### Notes



#### Note

For backward compatibility, the following deprecated alias may be used: `mysql_fieldflags`

### See Also

`mysql_field_type`  
`mysql_field_len`

### 21.9.2.5.20. `mysql_field_len`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_field_len`

Returns the length of the specified field



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_fetch_field_direct [length]  
PDOStatement::getColumnMeta [len]
```

#### Description

```
int mysql_field_len(  
    resource result,  
    int field_offset);
```

`mysql_field_len` returns the length of the specified field.

#### Parameters

*result*

The result resource that is being evaluated. This result comes from a call to `mysql_query`.

*field\_offset*

The numerical field offset. The *field\_offset* starts at 0. If *field\_offset* does not exist, an error of level `E_WARNING` is also issued.

#### Return Values

The length of the specified field index on success or `FALSE` on failure.

#### Examples

##### Example 21.49. `mysql_field_len` example

```
<?php  
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");  
if (!$result) {  
    echo 'Could not run query: ' . mysql_error();  
    exit;  
}  
  
// Will get the length of the id field as specified in the database  
// schema.  
$length = mysql_field_len($result, 0);  
echo $length;  
?>
```

#### Notes



#### Note

For backward compatibility, the following deprecated alias may be used:  
`mysql_fieldlen`

**See Also**

[mysql\\_fetch\\_lengths](#)  
[strlen](#)

**21.9.2.5.21. [mysql\\_field\\_name](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_field\\_name](#)

Get the name of the specified field in a result

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_fetch\\_field\\_direct](#) [*name*] or [*orgname*]  
[PDOStatement::getColumnMeta](#) [*name*]

**Description**

```
string mysql_field_name(
    resource result,
    int field_offset);
```

[mysql\\_field\\_name](#) returns the name of the specified field index.

**Parameters**

*result*

The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).

*field\_offset*

The numerical field offset. The *field\_offset* starts at 0. If *field\_offset* does not exist, an error of level [E\\_WARNING](#) is also issued.

**Return Values**

The name of the specified field index on success or [FALSE](#) on failure.

**Examples****Example 21.50. [mysql\\_field\\_name](#) example**

```
<?php
/* The users table consists of three fields:
 *   user_id
 *   username
 *   password.
 */
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect to MySQL server: ' . mysql_error());
}
$dbname = 'mydb';
$db_selected = mysql_select_db($dbname, $link);
if (!$db_selected) {
    die("Could not set $dbname: " . mysql_error());
}
```

```
$res = mysql_query('select * from users', $link);

echo mysql_field_name($res, 0) . "\n";
echo mysql_field_name($res, 2);
?>
```

The above example will output:

```
user_id
password
```

## Notes



### Note

Field names returned by this function are *case-sensitive*.



### Note

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_fieldname](#)

## See Also

[mysql\\_field\\_type](#)  
[mysql\\_field\\_len](#)

### 21.9.2.5.22. [mysql\\_field\\_seek](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_field\\_seek](#)

Set result pointer to a specified field offset



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_field\\_seek](#)  
[PDOStatement::fetch](#) using optional parameters

## Description

```
bool mysql_field_seek(
    resource result,
    int field_offset);
```

Seeks to the specified field offset. If the next call to [mysql\\_fetch\\_field](#) doesn't include a field offset, the field offset specified in [mysql\\_field\\_seek](#) will be returned.

## Parameters

*result*

The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).

`field_offset`

The numerical field offset. The `field_offset` starts at 0. If `field_offset` does not exist, an error of level `E_WARNING` is also issued.

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**See Also**`mysql_fetch_field`**21.9.2.5.23. `mysql_field_table`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_field_table`

Get name of the table the specified field is in

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_fetch_field_direct` [`table`] or [`orgtable`]  
`PDOStatement::getColumnMeta` [`table`]

**Description**

```
string mysql_field_table(
    resource result,
    int field_offset);
```

Returns the name of the table that the specified field is in.

**Parameters**`result`

The result resource that is being evaluated. This result comes from a call to `mysql_query`.

`field_offset`

The numerical field offset. The `field_offset` starts at 0. If `field_offset` does not exist, an error of level `E_WARNING` is also issued.

**Return Values**

The name of the table on success.

**Examples****Example 21.51. A `mysql_field_table` example**

```
<?php
$query = "SELECT account.*, country.* FROM account, country WHERE country.name = 'Portugal' AND account.co
// get the result from the DB
$result = mysql_query($query);
// Lists the table name and then the field name
```

```
for ($i = 0; $i < mysql_num_fields($result); ++$i) {
    $table = mysql_field_table($result, $i);
    $field = mysql_field_name($result, $i);

    echo "$table: $field\n";
}
?>
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_fieldtable](#)

## See Also

[mysql\\_list\\_tables](#)

### 21.9.2.5.24. [mysql\\_field\\_type](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_field\\_type](#)

Get the type of the specified field in a result



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_fetch\\_field\\_direct](#) [type]  
[PDOStatement::getColumnMeta](#) [driver:decl\_type] or [pdo\_type]

## Description

```
string mysql_field_type(
    resource result,
    int field_offset);
```

[mysql\\_field\\_type](#) is similar to the [mysql\\_field\\_name](#) function. The arguments are identical, but the field type is returned instead.

## Parameters

*result*

The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).

*field\_offset*

The numerical field offset. The *field\_offset* starts at 0. If *field\_offset* does not exist, an error of level [E\\_WARNING](#) is also issued.

## Return Values

The returned field type will be one of "int", "real", "string", "blob", and others as detailed in the [MySQL documentation](#).

## Examples

**Example 21.52. `mysql_field_type` example**

```
<?php
mysql_connect("localhost", "mysql_username", "mysql_password");
mysql_select_db("mysql");
$result = mysql_query("SELECT * FROM func");
$fields = mysql_num_fields($result);
$rows   = mysql_num_rows($result);
$table  = mysql_field_table($result, 0);
echo "Your '" . $table . "' table has " . $fields . " fields and " . $rows . " record(s)\n";
echo "The table has the following fields:\n";
for ($i=0; $i < $fields; $i++) {
    $type = mysql_field_type($result, $i);
    $name = mysql_field_name($result, $i);
    $len  = mysql_field_len($result, $i);
    $flags = mysql_field_flags($result, $i);
    echo $type . " " . $name . " " . $len . " " . $flags . "\n";
}
mysql_free_result($result);
mysql_close();
?>
```

The above example will output something similar to:

```
Your 'func' table has 4 fields and 1 record(s)
The table has the following fields:
string name 64 not_null primary_key binary
int ret 1 not_null
string dl 128 not_null
string type 9 not_null enum
```

**Notes****Note**

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_fieldtype](#)

**See Also**

[mysql\\_field\\_name](#)  
[mysql\\_field\\_len](#)

**21.9.2.5.25. `mysql_free_result`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_free\\_result](#)

Free result memory

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_free\\_result](#)

Assign the value of `NULL` to the PDO object, or  
[PDOStatement::closeCursor](#)



## Description

```
bool mysql_free_result(  
    resource result);
```

`mysql_free_result` will free all memory associated with the result identifier `result`.

`mysql_free_result` only needs to be called if you are concerned about how much memory is being used for queries that return large result sets. All associated result memory is automatically freed at the end of the script's execution.

## Parameters

`result` The result resource that is being evaluated. This result comes from a call to `mysql_query`.

## Return Values

Returns `TRUE` on success or `FALSE` on failure.

If a non-resource is used for the `result`, an error of level `E_WARNING` will be emitted. It's worth noting that `mysql_query` only returns a resource for `SELECT`, `SHOW`, `EXPLAIN`, and `DESCRIBE` queries.

## Examples

### Example 21.53. A `mysql_free_result` example

```
<?php  
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");  
if (!$result) {  
    echo 'Could not run query: ' . mysql_error();  
    exit;  
}  
/* Use the result, assuming we're done with it afterwards */  
$row = mysql_fetch_assoc($result);  
  
/* Now we free up the result and continue on with our script */  
mysql_free_result($result);  
  
echo $row['id'];  
echo $row['email'];  
?>
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
`mysql_freeresult`

## See Also

[mysql\\_query](#)  
[is\\_resource](#)

### 21.9.2.5.26. `mysql_get_client_info`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_get\\_client\\_info](#)

## Get MySQL client info

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_get_client_info  
PDO::getAttribute(PDO::ATTR_CLIENT_VERSION)
```

**Description**

```
string mysql_get_client_info();
```

`mysql_get_client_info` returns a string that represents the client library version.

**Return Values**

The MySQL client version.

**Examples****Example 21.54. `mysql_get_client_info` example**

```
<?php  
printf("MySQL client info: %s\n", mysql_get_client_info());  
?>
```

The above example will output something similar to:

```
MySQL client info: 3.23.39
```

**See Also**

```
mysql_get_host_info  
mysql_get_proto_info  
mysql_get_server_info
```

**21.9.2.5.27. `mysql_get_host_info`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_get\\_host\\_info](#)

## Get MySQL host info

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_get_host_info  
PDO::getAttribute(PDO::ATTR_CONNECTION_STATUS)
```

## Description

```
string mysql_get_host_info(  
    resource link_identifier  
    = =NULL);
```

Describes the type of connection in use for the connection, including the server host name.

## Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

## Return Values

Returns a string describing the type of MySQL connection in use for the connection or [FALSE](#) on failure.

## Examples

### Example 21.55. [mysql\\_get\\_host\\_info](#) example

```
<?php  
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');  
if (!$link) {  
    die('Could not connect: ' . mysql_error());  
}  
printf("MySQL host info: %s\n", mysql_get_host_info());  
?>
```

The above example will output something similar to:

```
MySQL host info: Localhost via UNIX socket
```

## See Also

[mysql\\_get\\_client\\_info](#)  
[mysql\\_get\\_proto\\_info](#)  
[mysql\\_get\\_server\\_info](#)

### 21.9.2.5.28. [mysql\\_get\\_proto\\_info](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_get\\_proto\\_info](#)

Get MySQL protocol info



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_get\\_proto\\_info](#)

## Description

```
int mysql_get_proto_info(
    resource link_identifier
    = =NULL);
```

Retrieves the MySQL protocol.

## Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

## Return Values

Returns the MySQL protocol on success or [FALSE](#) on failure.

## Examples

### Example 21.56. [mysql\\_get\\_proto\\_info](#) example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
printf("MySQL protocol version: %s\n", mysql_get_proto_info());
?>
```

The above example will output something similar to:

```
MySQL protocol version: 10
```

## See Also

[mysql\\_get\\_client\\_info](#)  
[mysql\\_get\\_host\\_info](#)  
[mysql\\_get\\_server\\_info](#)

### 21.9.2.5.29. [mysql\\_get\\_server\\_info](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_get\\_server\\_info](#)

Get MySQL server info



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_get\\_server\\_info](#)

```
mysql::getAttribute(PDO::ATTR_SERVER_VERSION)
```

## Description

```
string mysql_get_server_info(
    resource link_identifier
    = =NULL);
```

Retrieves the MySQL server version.

## Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

## Return Values

Returns the MySQL server version on success or `FALSE` on failure.

## Examples

### Example 21.57. `mysql_get_server_info` example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
printf("MySQL server version: %s\n", mysql_get_server_info());
?>
```

The above example will output something similar to:

```
MySQL server version: 4.0.1-alpha
```

## See Also

`mysql_get_client_info`  
`mysql_get_host_info`  
`mysql_get_proto_info`  
`phpversion`

### 21.9.2.5.30. `mysql_info`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_info`

Get information about the most recent query



## Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

**mysql\_info****Description**

```
string mysql_info(
    resource link_identifier
    = =NULL);
```

Returns detailed information about the last query.

**Parameters**

`link_identifier`

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

Returns information about the statement on success, or `FALSE` on failure. See the example below for which statements provide information, and what the returned value may look like. Statements that are not listed will return `FALSE`.

**Examples****Example 21.58. Relevant MySQL Statements**

Statements that return string values. The numbers are only for illustrating purpose; their values will correspond to the query.

```
INSERT INTO ... SELECT ...
String format: Records: 23 Duplicates: 0 Warnings: 0
INSERT INTO ... VALUES (...),(...),(...)...
String format: Records: 37 Duplicates: 0 Warnings: 0
LOAD DATA INFILE ...
String format: Records: 42 Deleted: 0 Skipped: 0 Warnings: 0
ALTER TABLE
String format: Records: 60 Duplicates: 0 Warnings: 0
UPDATE
String format: Rows matched: 65 Changed: 65 Warnings: 0
```

**Notes****Note**

`mysql_info` returns a non-`FALSE` value for the `INSERT ... VALUES` statement only if multiple value lists are specified in the statement.

**See Also**

`mysql_affected_rows`  
`mysql_insert_id`  
`mysql_stat`

**21.9.2.5.31. mysql\_insert\_id**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_insert_id`

Get the ID generated in the last query



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_insert_id
PDO::lastInsertId
```

### Description

```
int mysql_insert_id(
    resource link_identifier
    = =NULL);
```

Retrieves the ID generated for an AUTO\_INCREMENT column by the previous query (usually INSERT).

### Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

### Return Values

The ID generated for an AUTO\_INCREMENT column by the previous query on success, `0` if the previous query does not generate an AUTO\_INCREMENT value, or `FALSE` if no MySQL connection was established.

### Examples

#### Example 21.59. [mysql\\_insert\\_id](#) example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
mysql_select_db('mydb');

mysql_query("INSERT INTO mytable (product) values ('kossu')");
printf("Last inserted record has id %d\n", mysql_insert_id());
?>
```

### Notes



### Caution

[mysql\\_insert\\_id](#) will convert the return type of the native MySQL C API function [mysql\\_insert\\_id\(\)](#) to a type of `long` (named `int` in PHP). If your AUTO\_INCREMENT column has a column type of BIGINT (64 bits) the conversion may result in an incorrect value. Instead, use the internal MySQL SQL function `LAST_INSERT_ID()` in a SQL query. For more information about PHP's maximum integer values, please see the [integer](#) documentation.

**Note**

Because `mysql_insert_id` acts on the last performed query, be sure to call `mysql_insert_id` immediately after the query that generates the value.

**Note**

The value of the MySQL SQL function `LAST_INSERT_ID()` always contains the most recently generated `AUTO_INCREMENT` value, and is not reset between queries.

**See Also**

`mysql_query`  
`mysql_info`

**21.9.2.5.32. `mysql_list_dbs`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_list_dbs`

List databases available on a MySQL server

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

SQL Query: `SHOW DATABASES`

**Description**

```
resource mysql_list_dbs(
    resource link_identifier
    = =NULL);
```

Returns a result pointer containing the databases available from the current `mysql` daemon.

**Warning**

This function has been *DEPRECATED* as of PHP 5.4.0. Relying on this function is highly discouraged.

**Parameters**

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

Returns a result pointer resource on success, or `FALSE` on failure. Use the `mysql_tablename` function to traverse this result pointer, or any function for result tables, such as `mysql_fetch_array`.

**Examples**

**Example 21.60. `mysql_list_dbs` example**



```
<?php
// Usage without mysql_list_dbs()
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
$res = mysql_query("SHOW DATABASES");

while ($row = mysql_fetch_assoc($res)) {
    echo $row['Database'] . "\n";
}

// Deprecated as of PHP 5.4.0
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
$db_list = mysql_list_dbs($link);

while ($row = mysql_fetch_object($db_list)) {
    echo $row->Database . "\n";
}
?>
```

The above example will output something similar to:

```
database1
database2
database3
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:

`mysql_listdbs`

## See Also

[mysql\\_db\\_name](#)  
[mysql\\_select\\_db](#)

### 21.9.2.5.33. [mysql\\_list\\_fields](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_list\\_fields](#)

List MySQL table fields



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

SQL Query: `SHOW COLUMNS FROM sometable`

## Description

```
resource mysql_list_fields(
    string database_name,
    string table_name,
    resource link_identifier
    = =NULL);
```

Retrieves information about the given table name.

This function is deprecated. It is preferable to use `mysql_query` to issue an SQL `SHOW COLUMNS FROM table [LIKE 'name']` statement instead.

### Parameters

<i>database_name</i>	The name of the database that's being queried.
<i>table_name</i>	The name of the table that's being queried.
<i>link_identifier</i>	The MySQL connection. If the link identifier is not specified, the last link opened by <code>mysql_connect</code> is assumed. If no such link is found, it will try to create one as if <code>mysql_connect</code> was called with no arguments. If no connection is found or established, an <code>E_WARNING</code> level error is generated.

### Return Values

A result pointer resource on success, or `FALSE` on failure.

The returned result can be used with `mysql_field_flags`, `mysql_field_len`, `mysql_field_name` and `mysql_field_type`.

### Examples

#### Example 21.61. Alternate to deprecated `mysql_list_fields`

```
<?php
$result = mysql_query("SHOW COLUMNS FROM sometable");
if (!$result) {
    echo 'Could not run query: ' . mysql_error();
    exit;
}
if (mysql_num_rows($result) > 0) {
    while ($row = mysql_fetch_assoc($result)) {
        print_r($row);
    }
}
?>
```

The above example will output something similar to:

```
Array
(
    [Field] => id
    [Type] => int(7)
    [Null] =>
    [Key] => PRI
    [Default] =>
    [Extra] => auto_increment
)
Array
(
    [Field] => email
    [Type] => varchar(100)
    [Null] =>
    [Key] =>
    [Default] =>
    [Extra] =>
)
```

### Notes

**Note**

For backward compatibility, the following deprecated alias may be used:  
`mysql_listfields`

**See Also**

`mysql_field_flags`  
`mysql_info`

**21.9.2.5.34. `mysql_list_processes`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_list_processes`

List MySQL processes

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_thread_id`

**Description**

```
resource mysql_list_processes(
    resource link_identifier
    = =NULL);
```

Retrieves the current MySQL server threads.

**Parameters**

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

A result pointer resource on success or `FALSE` on failure.

**Examples****Example 21.62. `mysql_list_processes` example**

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');

$result = mysql_list_processes($link);
while ($row = mysql_fetch_assoc($result)){
    printf("%s %s %s %s %s\n", $row["Id"], $row["Host"], $row["db"],
        $row["Command"], $row["Time"]);
}
mysql_free_result($result);
?>
```

The above example will output something similar to:

```
1 localhost test Processlist 0
4 localhost mysql sleep 5
```

### See Also

[mysql\\_thread\\_id](#)  
[mysql\\_stat](#)

#### 21.9.2.5.35. [mysql\\_list\\_tables](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_list\\_tables](#)

List tables in a MySQL database



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

SQL Query: `SHOW TABLES FROM sometable`

### Description

```
resource mysql_list_tables(
    string database,
    resource link_identifier
    = =NULL);
```

Retrieves a list of table names from a MySQL database.

This function is deprecated. It is preferable to use [mysql\\_query](#) to issue an SQL `SHOW TABLES [FROM db_name] [LIKE 'pattern']` statement instead.

### Parameters

<i>database</i>	The name of the database
<i>link_identifier</i>	The MySQL connection. If the link identifier is not specified, the last link opened by <a href="#">mysql_connect</a> is assumed. If no such link is found, it will try to create one as if <a href="#">mysql_connect</a> was called with no arguments. If no connection is found or established, an <a href="#">E_WARNING</a> level error is generated.

### Return Values

A result pointer resource on success or `FALSE` on failure.

Use the [mysql\\_tablename](#) function to traverse this result pointer, or any function for result tables, such as [mysql\\_fetch\\_array](#).

### Changelog

Version	Description
4.3.7	This function became deprecated.

## Examples

### Example 21.63. [mysql\\_list\\_tables](#) alternative example

```
<?php
$dbname = 'mysql_dbname';

if (!mysql_connect('mysql_host', 'mysql_user', 'mysql_password')) {
    echo 'Could not connect to mysql';
    exit;
}

$sql = "SHOW TABLES FROM $dbname";
$result = mysql_query($sql);

if (!$result) {
    echo "DB Error, could not list tables\n";
    echo 'MySQL Error: ' . mysql_error();
    exit;
}

while ($row = mysql_fetch_row($result)) {
    echo "Table: {$row[0]}\n";
}

mysql_free_result($result);
?>
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_listtables](#)

## See Also

[mysql\\_list\\_dbs](#)  
[mysql\\_tablename](#)

### 21.9.2.5.36. [mysql\\_num\\_fields](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_num\\_fields](#)

Get number of fields in result



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_field\\_count](#)  
[PDOStatement::columnCount](#)

## Description

```
int mysql_num_fields(
    resource result);
```

Retrieves the number of fields from a query.

## Parameters

`result` The result resource that is being evaluated. This result comes from a call to `mysql_query`.

## Return Values

Returns the number of fields in the result set resource on success or `FALSE` on failure.

## Examples

### Example 21.64. A `mysql_num_fields` example

```
<?php
$result = mysql_query("SELECT id,email FROM people WHERE id = '42'");
if (!$result) {
    echo 'Could not run query: ' . mysql_error();
    exit;
}

/* returns 2 because id,email == two fields */
echo mysql_num_fields($result);
?>
```

## Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
`mysql_numfields`

## See Also

`mysql_select_db`  
`mysql_query`  
`mysql_fetch_field`  
`mysql_num_rows`

### 21.9.2.5.37. `mysql_num_rows`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_num_rows`

Get number of rows in result



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the `MySQLi` or `PDO_MySQL` extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_stmt_num_rows`  
`PDOStatement::rowCount`

## Description

```
int mysql_num_rows(
    resource result);
```

Retrieves the number of rows from a result set. This command is only valid for statements like `SELECT` or `SHOW` that return an actual result set. To retrieve the number of rows affected by a `INSERT`, `UPDATE`, `REPLACE` or `DELETE` query, use [mysql\\_affected\\_rows](#).

### Parameters

*result* The result resource that is being evaluated. This result comes from a call to [mysql\\_query](#).

### Return Values

The number of rows in a result set on success or `FALSE` on failure.

### Examples

#### Example 21.65. [mysql\\_num\\_rows](#) example

```
<?php
$link = mysql_connect("localhost", "mysql_user", "mysql_password");
mysql_select_db("database", $link);

$result = mysql_query("SELECT * FROM table1", $link);
$num_rows = mysql_num_rows($result);

echo "$num_rows Rows\n";

?>
```

### Notes



#### Note

If you use [mysql\\_unbuffered\\_query](#), [mysql\\_num\\_rows](#) will not return the correct value until all the rows in the result set have been retrieved.



#### Note

For backward compatibility, the following deprecated alias may be used:  
[mysql\\_numrows](#)

### See Also

[mysql\\_affected\\_rows](#)  
[mysql\\_connect](#)  
[mysql\\_data\\_seek](#)  
[mysql\\_select\\_db](#)  
[mysql\\_query](#)

#### 21.9.2.5.38. [mysql\\_pconnect](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_pconnect](#)

Open a persistent connection to a MySQL server



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See

also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_connect` with `p`: host prefix  
`PDO::__construct` with `PDO::ATTR_PERSISTENT` as a driver option

## Description

```
resource mysql_pconnect(
    string server
        = =ini_get("mysql.default_host"),
    string username
        = =ini_get("mysql.default_user"),
    string password
        = =ini_get("mysql.default_password"),
    int client_flags
        = =0);
```

Establishes a persistent connection to a MySQL server.

`mysql_pconnect` acts very much like `mysql_connect` with two major differences.

First, when connecting, the function would first try to find a (persistent) link that's already open with the same host, username and password. If one is found, an identifier for it will be returned instead of opening a new connection.

Second, the connection to the SQL server will not be closed when the execution of the script ends. Instead, the link will remain open for future use (`mysql_close` will not close links established by `mysql_pconnect`).

This type of link is therefore called 'persistent'.

## Parameters

<i>server</i>	The MySQL server. It can also include a port number. e.g. "hostname:port" or a path to a local socket e.g. ":/path/to/socket" for the localhost.  If the PHP directive <code>mysql.default_host</code> is undefined (default), then the default value is 'localhost:3306'
<i>username</i>	The username. Default value is the name of the user that owns the server process.
<i>password</i>	The password. Default value is an empty password.
<i>client_flags</i>	The <code>client_flags</code> parameter can be a combination of the following constants: 128 (enable <code>LOAD DATA LOCAL</code> handling), <code>MYSQL_CLIENT_SSL</code> , <code>MYSQL_CLIENT_COMPRESS</code> , <code>MYSQL_CLIENT_IGNORE_SPACE</code> or <code>MYSQL_CLIENT_INTERACTIVE</code> .

## Return Values

Returns a MySQL persistent link identifier on success, or `FALSE` on failure.

## Changelog

Version	Description
5.5.0	This function will generate an <code>E_DEPRECATED</code> error.
4.3.0	Added the <code>client_flags</code> parameter.



## Notes



### Note

Note, that these kind of links only work if you are using a module version of PHP. See the [Persistent Database Connections](#) section for more information.



### Warning

Using persistent connections can require a bit of tuning of your Apache and MySQL configurations to ensure that you do not exceed the number of connections allowed by MySQL.



### Note

You can suppress the error message on failure by prepending a `@` to the function name.

## See Also

[mysql\\_connect](#)  
[Persistent Database Connections](#)

### 21.9.2.5.39. [mysql\\_ping](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_ping](#)

Ping a server connection or reconnect if there is no connection



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_ping](#)

## Description

```
bool mysql_ping(  
    resource link_identifier  
    = =NULL);
```

Checks whether or not the connection to the server is working. If it has gone down, an automatic reconnection is attempted. This function can be used by scripts that remain idle for a long while, to check whether or not the server has closed the connection and reconnect if necessary.



### Note

Automatic reconnection is disabled by default in versions of MySQL  $\geq$  5.0.3.

## Parameters

[link\\_identifier](#)

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

## Return Values

Returns [TRUE](#) if the connection to the server MySQL server is working, otherwise [FALSE](#) .

## Examples

### Example 21.66. A [mysql\\_ping](#) example

```

<?php
set_time_limit(0);

$conn = mysql_connect('localhost', 'mysqluser', 'mypass');
$db   = mysql_select_db('mydb');

/* Assuming this query will take a long time */
$result = mysql_query($sql);
if (!$result) {
    echo 'Query #1 failed, exiting.';
    exit;
}

/* Make sure the connection is still alive, if not, try to reconnect */
if (!mysql_ping($conn)) {
    echo 'Lost connection, exiting after query #1';
    exit;
}
mysql_free_result($result);

/* So the connection is still alive, let's run another query */
$result2 = mysql_query($sql2);
?>

```

## See Also

[mysql\\_thread\\_id](#)  
[mysql\\_list\\_processes](#)

### 21.9.2.5.40. [mysql\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_query](#)

Send a MySQL query



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_query](#)  
[PDO::query](#)

## Description

```

resource mysql_query(
    string query,
    resource link_identifier
    = =NULL);

```

[mysql\\_query](#) sends a unique query (multiple queries are not supported) to the currently active database on the server that's associated with the specified [link\\_identifier](#).

## Parameters

<code>query</code>	An SQL query  The query string should not end with a semicolon. Data inside the query should be <a href="#">properly escaped</a> .
<code>link_identifier</code>	The MySQL connection. If the link identifier is not specified, the last link opened by <code>mysql_connect</code> is assumed. If no such link is found, it will try to create one as if <code>mysql_connect</code> was called with no arguments. If no connection is found or established, an <code>E_WARNING</code> level error is generated.

## Return Values

For SELECT, SHOW, DESCRIBE, EXPLAIN and other statements returning resultset, `mysql_query` returns a resource on success, or `FALSE` on error.

For other type of SQL statements, INSERT, UPDATE, DELETE, DROP, etc, `mysql_query` returns `TRUE` on success or `FALSE` on error.

The returned result resource should be passed to `mysql_fetch_array`, and other functions for dealing with result tables, to access the returned data.

Use `mysql_num_rows` to find out how many rows were returned for a SELECT statement or `mysql_affected_rows` to find out how many rows were affected by a DELETE, INSERT, REPLACE, or UPDATE statement.

`mysql_query` will also fail and return `FALSE` if the user does not have permission to access the table(s) referenced by the query.

## Examples

### Example 21.67. Invalid Query

The following query is syntactically invalid, so `mysql_query` fails and returns `FALSE`.

```
<?php
$result = mysql_query('SELECT * WHERE 1=1');
if (!$result) {
    die('Invalid query: ' . mysql_error());
}
?>
```

### Example 21.68. Valid Query

The following query is valid, so `mysql_query` returns a resource.

```
<?php
// This could be supplied by a user, for example
$firstname = 'fred';
$lastname = 'fox';

// Formulate Query
// This is the best way to perform an SQL query
// For more examples, see mysql_real_escape_string()
$query = sprintf("SELECT firstname, lastname, address, age FROM friends
    WHERE firstname='%s' AND lastname='%s'",
    mysql_real_escape_string($firstname),
    mysql_real_escape_string($lastname));
```

```
// Perform Query
$result = mysql_query($query);

// Check result
// This shows the actual query sent to MySQL, and the error. Useful for debugging.
if (!$result) {
    $message = 'Invalid query: ' . mysql_error() . "\n";
    $message .= 'Whole query: ' . $query;
    die($message);
}

// Use result
// Attempting to print $result won't allow access to information in the resource
// One of the mysql result functions must be used
// See also mysql_result(), mysql_fetch_array(), mysql_fetch_row(), etc.
while ($row = mysql_fetch_assoc($result)) {
    echo $row['firstname'];
    echo $row['lastname'];
    echo $row['address'];
    echo $row['age'];
}

// Free the resources associated with the result set
// This is done automatically at the end of the script
mysql_free_result($result);
?>
```

### See Also

[mysql\\_connect](#)  
[mysql\\_error](#)  
[mysql\\_real\\_escape\\_string](#)  
[mysql\\_result](#)  
[mysql\\_fetch\\_assoc](#)  
[mysql\\_unbuffered\\_query](#)

#### 21.9.2.5.41. [mysql\\_real\\_escape\\_string](#)

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

- [mysql\\_real\\_escape\\_string](#)

Escapes special characters in a string for use in an SQL statement



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_real\\_escape\\_string](#)  
[PDO::quote](#)

#### Description

```
string mysql_real_escape_string(
    string unescaped_string,
    resource link_identifier
    = =NULL);
```

Escapes special characters in the *unescaped\_string*, taking into account the current character set of the connection so that it is safe to place it in a *mysql\_query*. If binary data is to be inserted, this function must be used.

`mysql_real_escape_string` calls MySQL's library function `mysql_real_escape_string`, which prepends backslashes to the following characters: `\x00`, `\n`, `\r`, `\`, `'`, `"` and `\x1a`.

This function must always (with few exceptions) be used to make data safe before sending a query to MySQL.



#### Security: the default character set

The character set must be set either at the server level, or with the API function `mysql_set_charset` for it to affect `mysql_real_escape_string`. See the concepts section on [character sets](#) for more information.

#### Parameters

*unescaped\_string*

The string that is to be escaped.

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

#### Return Values

Returns the escaped string, or `FALSE` on error.

#### Examples

##### Example 21.69. Simple `mysql_real_escape_string` example

```
<?php
// Connect
$link = mysql_connect('mysql_host', 'mysql_user', 'mysql_password')
    OR die(mysql_error());

// Query
$query = sprintf("SELECT * FROM users WHERE user='%s' AND password='%s'",
    mysql_real_escape_string($user),
    mysql_real_escape_string($password));
?>
```

##### Example 21.70. An example SQL Injection Attack

```
<?php
// We didn't check $_POST['password'], it could be anything the user wanted! For example:
$_POST['username'] = 'aidan';
$_POST['password'] = "' OR ''='";

// Query database to check if there are any matching users
$query = "SELECT * FROM users WHERE user='{$_POST['username']}' AND password='{$_POST['password']}'";
mysql_query($query);

// This means the query sent to MySQL would be:
echo $query;
?>
```

The query sent to MySQL:

```
SELECT * FROM users WHERE user='aidan' AND password='' OR ''=''
```

This would allow anyone to log in without a valid password.

## Notes



### Note

A MySQL connection is required before using `mysql_real_escape_string` otherwise an error of level `E_WARNING` is generated, and `FALSE` is returned. If `link_identifier` isn't defined, the last MySQL connection is used.



### Note

If `magic_quotes_gpc` is enabled, first apply `stripslashes` to the data. Using this function on data which has already been escaped will escape the data twice.



### Note

If this function is not used to escape data, the query is vulnerable to [SQL Injection Attacks](#).



### Note

`mysql_real_escape_string` does not escape `%` and `_`. These are wildcards in MySQL if combined with `LIKE`, `GRANT`, or `REVOKE`.

## See Also

`mysql_set_charset`  
`mysql_client_encoding`  
`addslashes`  
`stripslashes`  
The `magic_quotes_gpc` directive  
The `magic_quotes_runtime` directive

### 21.9.2.5.42. `mysql_result`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_result`

Get result data



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_data_seek` in conjunction with `mysqli_field_seek` and `mysqli_fetch_field`  
`PDOStatement::fetchColumn`

## Description

```
string mysql_result(
    resource result,
    int row,
    mixed field
    = =0);
```

Retrieves the contents of one cell from a MySQL result set.

When working on large result sets, you should consider using one of the functions that fetch an entire row (specified below). As these functions return the contents of multiple cells in one function call, they're MUCH quicker than [mysql\\_result](#). Also, note that specifying a numeric offset for the field argument is much quicker than specifying a fieldname or tablename.fieldname argument.

### Parameters

<i>result</i>	The result resource that is being evaluated. This result comes from a call to <a href="#">mysql_query</a> .
<i>row</i>	The row number from the result that's being retrieved. Row numbers start at 0.
<i>field</i>	The name or offset of the field being retrieved.  It can be the field's offset, the field's name, or the field's table dot field name (tablename.fieldname). If the column name has been aliased ('select foo as bar from...'), use the alias instead of the column name. If undefined, the first field is retrieved.

### Return Values

The contents of one cell from a MySQL result set on success, or [FALSE](#) on failure.

### Examples

#### Example 21.71. [mysql\\_result](#) example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Could not connect: ' . mysql_error());
}
if (!mysql_select_db('database_name')) {
    die('Could not select database: ' . mysql_error());
}
$result = mysql_query('SELECT name FROM work.employee');
if (!$result) {
    die('Could not query:' . mysql_error());
}
echo mysql_result($result, 2); // outputs third employee's name

mysql_close($link);
?>
```

### Notes



#### Note

Calls to [mysql\\_result](#) should not be mixed with calls to other functions that deal with the result set.

### See Also

[mysql\\_fetch\\_row](#)  
[mysql\\_fetch\\_array](#)  
[mysql\\_fetch\\_assoc](#)  
[mysql\\_fetch\\_object](#)

#### 21.9.2.5.43. [mysql\\_select\\_db](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_select_db`

Select a MySQL database



### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_select_db`  
 PDO: `::__construct` (part of dsn)

### Description

```
bool mysql_select_db(
    string database_name,
    resource link_identifier
    = =NULL);
```

Sets the current active database on the server that's associated with the specified link identifier. Every subsequent call to `mysql_query` will be made on the active database.

### Parameters

*database\_name*

The name of the database that is to be selected.

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an [E\\_WARNING](#) level error is generated.

### Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

### Examples

#### Example 21.72. `mysql_select_db` example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
if (!$link) {
    die('Not connected : ' . mysql_error());
}

// make foo the current db
$db_selected = mysql_select_db('foo', $link);
if (!$db_selected) {
    die ('Can\'t use foo : ' . mysql_error());
}
?>
```

### Notes



### Note

For backward compatibility, the following deprecated alias may be used:  
`mysql_selectdb`



**See Also**

[mysql\\_connect](#)  
[mysql\\_pconnect](#)  
[mysql\\_query](#)

**21.9.2.5.44. [mysql\\_set\\_charset](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_set\\_charset](#)

Sets the client character set

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

[mysqli\\_character\\_set\\_name](#)

PDO: Add `charset` to the connection string, such as `charset=utf8`

**Description**

```
bool mysql_set_charset(
    string charset,
    resource link_identifier
    = =NULL);
```

Sets the default character set for the current connection.

**Parameters**

*charset*

A valid character set name.

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by [mysql\\_connect](#) is assumed. If no such link is found, it will try to create one as if [mysql\\_connect](#) was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**Notes****Note**

This function requires MySQL 5.0.7 or later.

**Note**

This is the preferred way to change the charset. Using [mysql\\_query](#) to set it (such as `SET NAMES utf8`) is not recommended. See the [MySQL character set concepts](#) section for more information.

**See Also**

[mysql\\_client\\_encoding](#)  
[Setting character sets in MySQL](#)

---

## List of character sets that MySQL supports

### 21.9.2.5.45. `mysql_stat`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysql_stat`

Get current system status



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

```
mysqli_stat  
PDO::getAttribute(PDO::ATTR_SERVER_INFO)
```

#### Description

```
string mysql_stat(  
    resource link_identifier  
    = =NULL);
```

`mysql_stat` returns the current server status.

#### Parameters

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

#### Return Values

Returns a string with the status for uptime, threads, queries, open tables, flush tables and queries per second. For a complete list of other status variables, you have to use the `SHOW STATUS SQL` command. If *link\_identifier* is invalid, `NULL` is returned.

#### Examples

##### Example 21.73. `mysql_stat` example

```
<?php  
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');  
$status = explode(' ', mysql_stat($link));  
print_r($status);  
?>
```

The above example will output something similar to:

```
Array  
(  
    [0] => Uptime: 5380  
    [1] => Threads: 2  
    [2] => Questions: 1321299  
    [3] => Slow queries: 0  
    [4] => Opens: 26  
    [5] => Flush tables: 1
```

```
[6] => Open tables: 17
[7] => Queries per second avg: 245.595
)
```

#### Example 21.74. Alternative `mysql_stat` example

```
<?php
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');
$result = mysql_query('SHOW STATUS', $link);
while ($row = mysql_fetch_assoc($result)) {
    echo $row['Variable_name'] . ' = ' . $row['Value'] . "\n";
}
?>
```

The above example will output something similar to:

```
back_log = 50
basedir = /usr/local/
bdb_cache_size = 8388600
bdb_log_buffer_size = 32768
bdb_home = /var/db/mysql/
bdb_max_lock = 10000
bdb_logdir =
bdb_shared_data = OFF
bdb_tmpdir = /var/tmp/
...
```

#### See Also

[mysql\\_get\\_server\\_info](#)  
[mysql\\_list\\_processes](#)

#### 21.9.2.5.46. `mysql_tablename`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_tablename](#)

Get table name of field



#### Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

SQL Query: `SHOW TABLES`

#### Description

```
string mysql_tablename(
    resource result,
    int i);
```

Retrieves the table name from a *result*.

This function is deprecated. It is preferable to use `mysql_query` to issue an SQL `SHOW TABLES [FROM db_name] [LIKE 'pattern']` statement instead.

**Parameters**

<i>result</i>	A result pointer resource that's returned from <a href="#">mysql_list_tables</a> .
<i>i</i>	The integer index (row/table number)

**Return Values**

The name of the table on success or `FALSE` on failure.

Use the [mysql\\_tablename](#) function to traverse this result pointer, or any function for result tables, such as [mysql\\_fetch\\_array](#).

**Changelog**

Version	Description
5.5.0	The <a href="#">mysql_tablename</a> function is deprecated, and emits an <code>E_DEPRECATED</code> level error.

**Examples****Example 21.75. [mysql\\_tablename](#) example**

```
<?php
mysql_connect("localhost", "mysql_user", "mysql_password");
$result = mysql_list_tables("mydb");
$num_rows = mysql_num_rows($result);
for ($i = 0; $i < $num_rows; $i++) {
    echo "Table: ", mysql_tablename($result, $i), "\n";
}

mysql_free_result($result);
?>
```

**Notes****Note**

The [mysql\\_num\\_rows](#) function may be used to determine the number of tables in the result pointer.

**See Also**

[mysql\\_list\\_tables](#)  
[mysql\\_field\\_table](#)  
[mysql\\_db\\_name](#)

**21.9.2.5.47. [mysql\\_thread\\_id](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_thread\\_id](#)

Return the current thread ID

**Warning**

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See

also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

`mysqli_thread_id`

## Description

```
int mysql_thread_id(  
    resource link_identifier  
    = =NULL);
```

Retrieves the current thread ID. If the connection is lost, and a reconnect with `mysql_ping` is executed, the thread ID will change. This means only retrieve the thread ID when needed.

## Parameters

`link_identifier` The MySQL connection. If the link identifier is not specified, the last link opened by `mysql_connect` is assumed. If no such link is found, it will try to create one as if `mysql_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

## Return Values

The thread ID on success or `FALSE` on failure.

## Examples

### Example 21.76. `mysql_thread_id` example

```
<?php  
$link = mysql_connect('localhost', 'mysql_user', 'mysql_password');  
$thread_id = mysql_thread_id($link);  
if ($thread_id){  
    printf("current thread id is %d\n", $thread_id);  
}  
?>
```

The above example will output something similar to:

```
current thread id is 73
```

## See Also

[mysql\\_ping](#)  
[mysql\\_list\\_processes](#)

### 21.9.2.5.48. `mysql_unbuffered_query`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysql\\_unbuffered\\_query](#)

Send an SQL query to MySQL without fetching and buffering the result rows.



## Warning

This extension is deprecated as of PHP 5.5.0, and will be removed in the future. Instead, the [MySQLi](#) or [PDO\\_MySQL](#) extension should be used. See

also [MySQL: choosing an API](#) guide and [related FAQ](#) for more information. Alternatives to this function include:

See: [Buffered and Unbuffered queries](#)

## Description

```
resource mysqli_unbuffered_query(  
    string query,  
    resource link_identifier  
    = =NULL);
```

`mysqli_unbuffered_query` sends the SQL query *query* to MySQL without automatically fetching and buffering the result rows as `mysqli_query` does. This saves a considerable amount of memory with SQL queries that produce large result sets, and you can start working on the result set immediately after the first row has been retrieved as you don't have to wait until the complete SQL query has been performed. To use `mysqli_unbuffered_query` while multiple database connections are open, you must specify the optional parameter *link\_identifier* to identify which connection you want to use.

## Parameters

*query*

The SQL query to execute.

Data inside the query should be [properly escaped](#).

*link\_identifier*

The MySQL connection. If the link identifier is not specified, the last link opened by `mysqli_connect` is assumed. If no such link is found, it will try to create one as if `mysqli_connect` was called with no arguments. If no connection is found or established, an `E_WARNING` level error is generated.

## Return Values

For SELECT, SHOW, DESCRIBE or EXPLAIN statements, `mysqli_unbuffered_query` returns a resource on success, or `FALSE` on error.

For other type of SQL statements, UPDATE, DELETE, DROP, etc, `mysqli_unbuffered_query` returns `TRUE` on success or `FALSE` on error.

## Notes



### Note

The benefits of `mysqli_unbuffered_query` come at a cost: you cannot use `mysqli_num_rows` and `mysqli_data_seek` on a result set returned from `mysqli_unbuffered_query`, until all rows are fetched. You also have to fetch all result rows from an unbuffered SQL query before you can send a new SQL query to MySQL, using the same *link\_identifier*.

## See Also

[mysqli\\_query](#)

## 21.9.3. MySQL Improved Extension (*Mysqli*)

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqli` extension allows you to access the functionality provided by MySQL 4.1 and above. More information about the MySQL Database server can be found at <http://www.mysql.com/>

An overview of software available for using MySQL from PHP can be found at [Section 21.9.3.2, "Overview"](#)

Documentation for MySQL can be found at <http://dev.mysql.com/doc/>.

Parts of this documentation included from MySQL manual with permissions of Oracle Corporation.

### 21.9.3.1. Examples

Copyright 1997-2012 the PHP Documentation Group. [2906]

All examples in the `mysqli` documentation use the world database. The world database can be found at <http://downloads.mysql.com/docs/world.sql.gz>

### 21.9.3.2. Overview

Copyright 1997-2012 the PHP Documentation Group. [2906]

This section provides an introduction to the options available to you when developing a PHP application that needs to interact with a MySQL database.

*What is an API?*

An Application Programming Interface, or API, defines the classes, methods, functions and variables that your application will need to call in order to carry out its desired task. In the case of PHP applications that need to communicate with databases the necessary APIs are usually exposed via PHP extensions.

APIs can be procedural or object-oriented. With a procedural API you call functions to carry out tasks, with the object-oriented API you instantiate classes and then call methods on the resulting objects. Of the two the latter is usually the preferred interface, as it is more modern and leads to better organized code.

When writing PHP applications that need to connect to the MySQL server there are several API options available. This document discusses what is available and how to select the best solution for your application.

*What is a Connector?*

In the MySQL documentation, the term *connector* refers to a piece of software that allows your application to connect to the MySQL database server. MySQL provides connectors for a variety of languages, including PHP.

If your PHP application needs to communicate with a database server you will need to write PHP code to perform such activities as connecting to the database server, querying the database and other database-related functions. Software is required to provide the API that your PHP application will use, and also handle the communication between your application and the database server, possibly using other intermediate libraries where necessary. This software is known generically as a connector, as it allows your application to *connect* to a database server.

*What is a Driver?*

A driver is a piece of software designed to communicate with a specific type of database server. The driver may also call a library, such as the MySQL Client Library or the MySQL Native Driver. These libraries implement the low-level protocol used to communicate with the MySQL database server.

By way of an example, the [PHP Data Objects \(PDO\) \[2985\]](#) database abstraction layer may use one of several database-specific drivers. One of the drivers it has available is the PDO MYSQL driver, which allows it to interface with the MySQL server.

Sometimes people use the terms connector and driver interchangeably, this can be confusing. In the MySQL-related documentation the term “driver” is reserved for software that provides the database-specific part of a connector package.

### *What is an Extension?*

In the PHP documentation you will come across another term - *extension*. The PHP code consists of a core, with optional extensions to the core functionality. PHP's MySQL-related extensions, such as the [mysqli](#) extension, and the [mysql](#) extension, are implemented using the PHP extension framework.

An extension typically exposes an API to the PHP programmer, to allow its facilities to be used programmatically. However, some extensions which use the PHP extension framework do not expose an API to the PHP programmer.

The PDO MySQL driver extension, for example, does not expose an API to the PHP programmer, but provides an interface to the PDO layer above it.

The terms API and extension should not be taken to mean the same thing, as an extension may not necessarily expose an API to the programmer.

### *What are the main PHP API offerings for using MySQL?*

There are three main API options when considering connecting to a MySQL database server:

- PHP's MySQL Extension
- PHP's [mysqli](#) Extension
- PHP Data Objects (PDO)

Each has its own advantages and disadvantages. The following discussion aims to give a brief introduction to the key aspects of each API.

### *What is PHP's MySQL Extension?*

This is the original extension designed to allow you to develop PHP applications that interact with a MySQL database. The [mysql](#) extension provides a procedural interface and is intended for use only with MySQL versions older than 4.1.3. This extension can be used with versions of MySQL 4.1.3 or newer, but not all of the latest MySQL server features will be available.



#### **Note**

If you are using MySQL versions 4.1.3 or later it is *strongly* recommended that you use the [mysqli](#) extension instead.

The [mysql](#) extension source code is located in the PHP extension directory `ext/mysql`.

For further information on the [mysql](#) extension, see [Section 21.9.2, "Original MySQL API \(MySQL\)"](#).

### *What is PHP's [mysqli](#) Extension?*

The [mysqli](#) extension, or as it is sometimes known, the MySQL *improved* extension, was developed to take advantage of new features found in MySQL systems versions 4.1.3 and newer. The [mysqli](#) extension is included with PHP versions 5 and later.

The [mysqli](#) extension has a number of benefits, the key enhancements over the [mysql](#) extension being:

- Object-oriented interface
- Support for Prepared Statements
- Support for Multiple Statements
- Support for Transactions



- Enhanced debugging capabilities
- Embedded server support

**Note**

If you are using MySQL versions 4.1.3 or later it is *strongly* recommended that you use this extension.

As well as the object-oriented interface the extension also provides a procedural interface.

The `mysqli` extension is built using the PHP extension framework, its source code is located in the directory `ext/mysqli`.

For further information on the `mysqli` extension, see [Section 21.9.3, “MySQL Improved Extension \(Mysqli\)”](#).

*What is PDO?*

PHP Data Objects, or PDO, is a database abstraction layer specifically for PHP applications. PDO provides a consistent API for your PHP application regardless of the type of database server your application will connect to. In theory, if you are using the PDO API, you could switch the database server you used, from say Firebird to MySQL, and only need to make minor changes to your PHP code.

Other examples of database abstraction layers include JDBC for Java applications and DBI for Perl.

While PDO has its advantages, such as a clean, simple, portable API, its main disadvantage is that it doesn't allow you to use all of the advanced features that are available in the latest versions of MySQL server. For example, PDO does not allow you to use MySQL's support for Multiple Statements.

PDO is implemented using the PHP extension framework, its source code is located in the directory `ext/pdo`.

For further information on PDO, see the <http://www.php.net/book.pdo>.

*What is the PDO MYSQL driver?*

The PDO MYSQL driver is not an API as such, at least from the PHP programmer's perspective. In fact the PDO MYSQL driver sits in the layer below PDO itself and provides MySQL-specific functionality. The programmer still calls the PDO API, but PDO uses the PDO MYSQL driver to carry out communication with the MySQL server.

The PDO MYSQL driver is one of several available PDO drivers. Other PDO drivers available include those for the Firebird and PostgreSQL database servers.

The PDO MYSQL driver is implemented using the PHP extension framework. Its source code is located in the directory `ext/pdo_mysql`. It does not expose an API to the PHP programmer.

For further information on the PDO MYSQL driver, see [Section 21.9.4, “MySQL Functions \(PDO\\_MYSQL\) \(MySQL \(PDO\)\)”](#).

*What is PHP's MySQL Native Driver?*

In order to communicate with the MySQL database server the `mysql` extension, `mysqli` and the PDO MYSQL driver each use a low-level library that implements the required protocol. In the past, the only available library was the MySQL Client Library, otherwise known as `libmysqlclient`.

However, the interface presented by `libmysqlclient` was not optimized for communication with PHP applications, as `libmysqlclient` was originally designed with C applications in mind. For this reason the MySQL Native Driver, `mysqlnd`, was developed as an alternative to `libmysqlclient` for PHP applications.

The `mysql` extension, the `mysqli` extension and the PDO MySQL driver can each be individually configured to use either `libmysqlclient` or `mysqlnd`. As `mysqlnd` is designed specifically to be utilised in the PHP system it has numerous memory and speed enhancements over `libmysqlclient`. You are strongly encouraged to take advantage of these improvements.

**Note**

The MySQL Native Driver can only be used with MySQL server versions 4.1.3 and later.

The MySQL Native Driver is implemented using the PHP extension framework. The source code is located in `ext/mysqlnd`. It does not expose an API to the PHP programmer.

*Comparison of Features*

The following table compares the functionality of the three main methods of connecting to MySQL from PHP:

**Table 21.38. Comparison of MySQL API options for PHP**

	PHP's <code>mysqli</code> Extension	PDO (Using PDO MySQL Driver and MySQL Native Driver)	PHP's MySQL Extension
PHP version introduced	5.0	5.0	Prior to 3.0
Included with PHP 5.x	yes	yes	Yes
MySQL development status	Active development	Active development as of PHP 5.3	Maintenance only
Recommended by MySQL for new projects	Yes - preferred option	Yes	No
API supports Charsets	Yes	Yes	No
API supports server-side Prepared Statements	Yes	Yes	No
API supports client-side Prepared Statements	No	Yes	No
API supports Stored Procedures	Yes	Yes	No
API supports Multiple Statements	Yes	Most	No
Supports all MySQL 4.1+ functionality	Yes	Most	No

**21.9.3.3. Quick start guide**

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

This quick start guide will help with choosing and gaining familiarity with the PHP MySQL API.

This quick start gives an overview on the `mysqli` extension. Code examples are provided for all major aspects of the API. Database concepts are explained to the degree needed for presenting concepts specific to MySQL.

Required: A familiarity with the PHP programming language, the SQL language, and basic knowledge of the MySQL server.

**21.9.3.3.1. Dual procedural and object-oriented interface**

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The `mysqli` extension features a dual interface. It supports the procedural and object-oriented programming paradigm.

Users migrating from the old `mysql` extension may prefer the procedural interface. The procedural interface is similar to that of the old `mysql` extension. In many cases, the function names differ only by prefix. Some `mysqli` functions take a connection handle as their first argument, whereas matching functions in the old `mysql` interface take it as an optional last argument.

#### Example 21.77. Easy migration from the old `mysql` extension

```
<?php
$mysqli = mysqli_connect("example.com", "user", "password", "database");
$res = mysqli_query($mysqli, "SELECT 'Please, do not use ' AS _msg FROM DUAL");
$row = mysqli_fetch_assoc($res);
echo $row['_msg'];

$mysql = mysql_connect("localhost", "root", "");
mysql_select_db("test");
$res = mysql_query("SELECT 'the mysql extension for new developments.' AS _msg FROM DUAL", $mysql);
$row = mysql_fetch_assoc($res);
echo $row['_msg'];
?>
```

The above example will output:

```
Please, do not use the mysql extension for new developments.
```

#### *The object-oriented interface*

In addition to the classical procedural interface, users can choose to use the object-oriented interface. The documentation is organized using the object-oriented interface. The object-oriented interface shows functions grouped by their purpose, making it easier to get started. The reference section gives examples for both syntax variants.

There are no significant performance differences between the two interfaces. Users can base their choice on personal preference.

#### Example 21.78. Object-oriented and procedural interface

```
<?php
$mysqli = mysqli_connect("example.com", "user", "password", "database");
if (mysqli_connect_errno($mysqli)) {
    echo "Failed to connect to MySQL: " . mysqli_connect_error();
}

$res = mysqli_query($mysqli, "SELECT 'A world full of ' AS _msg FROM DUAL");
$row = mysqli_fetch_assoc($res);
echo $row['_msg'];

$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: " . $mysqli->connect_error;
}

$res = $mysqli->query("SELECT 'choices to please everybody.' AS _msg FROM DUAL");
$row = $res->fetch_assoc();
echo $row['_msg'];
?>
```

The above example will output:

A world full of choices to please everybody.

The object oriented interface is used for the quickstart because the reference section is organized that way.

### *Mixing styles*

It is possible to switch between styles at any time. Mixing both styles is not recommended for code clarity and coding style reasons.

### **Example 21.79. Bad coding style**

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: " . $mysqli->connect_error;
}

$res = mysqli_query($mysqli, "SELECT 'Possible but bad style.' AS _msg FROM DUAL");
if (!$res) {
    echo "Failed to run query: (" . $mysqli->errno . ") " . $mysqli->error;
}

if ($row = $res->fetch_assoc()) {
    echo $row['_msg'];
}
?>
```

The above example will output:

Possible but bad style.

### *See also*

[mysqli::\\_\\_construct](#)  
[mysqli::query](#)  
[mysqli\\_result::fetch\\_assoc](#)  
[\\$mysqli::connect\\_errno](#)  
[\\$mysqli::connect\\_error](#)  
[\\$mysqli::errno](#)  
[\\$mysqli::error](#)  
[The MySQLi Extension Function Summary](#)

### **21.9.3.3.2. Connections**

Copyright 1997-2012 the PHP Documentation Group. [2906]

The MySQL server supports the use of different transport layers for connections. Connections use TCP/IP, Unix domain sockets or Windows named pipes.

The hostname `localhost` has a special meaning. It is bound to the use of Unix domain sockets. It is not possible to open a TCP/IP connection using the hostname `localhost` you must use `127.0.0.1` instead.

### **Example 21.80. Special meaning of localhost**

```
<?php
$mysqli = new mysqli("localhost", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}
echo $mysqli->host_info . "\n";

$mysqli = new mysqli("127.0.0.1", "user", "password", "database", 3306);
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

echo $mysqli->host_info . "\n";
?>
```

The above example will output:

```
Localhost via UNIX socket
127.0.0.1 via TCP/IP
```

### *Connection parameter defaults*

Depending on the connection function used, assorted parameters can be omitted. If a parameter is not provided, then the extension attempts to use the default values that are set in the PHP configuration file.

### **Example 21.81. Setting defaults**

```
mysqli.default_host=192.168.2.27
mysqli.default_user=root
mysqli.default_pw=""
mysqli.default_port=3306
mysqli.default_socket=/tmp/mysql.sock
```

The resulting parameter values are then passed to the client library that is used by the extension. If the client library detects empty or unset parameters, then it may default to the library built-in values.

### *Built-in connection library defaults*

If the host value is unset or empty, then the client library will default to a Unix socket connection on [localhost](#). If socket is unset or empty, and a Unix socket connection is requested, then a connection to the default socket on [/tmp/mysql.sock](#) is attempted.

On Windows systems, the host name `.` is interpreted by the client library as an attempt to open a Windows named pipe based connection. In this case the socket parameter is interpreted as the pipe name. If not given or empty, then the socket (pipe name) defaults to `\\.pipe\MySQL`.

If neither a Unix domain socket based not a Windows named pipe based connection is to be established and the port parameter value is unset, the library will default to port [3306](#).

The [mysqldb](#) library and the MySQL Client Library (libmysqlclient) implement the same logic for determining defaults.

### *Connection options*

Connection options are available to, for example, set init commands which are executed upon connect, or for requesting use of a certain charset. Connection options must be set before a network connection is established.

For setting a connection option, the connect operation has to be performed in three steps: creating a connection handle with [mysqli\\_init](#), setting the requested options using [mysqli\\_options](#), and establishing the network connection with [mysqli\\_real\\_connect](#).

### *Connection pooling*

The [mysqli](#) extension supports persistent database connections, which are a special kind of pooled connections. By default, every database connection opened by a script is either explicitly closed by the user during runtime or released automatically at the end of the script. A persistent connection is not. Instead it is put into a pool for later reuse, if a connection to the same server using the same username, password, socket, port and default database is opened. Reuse saves connection overhead.

Every PHP process is using its own [mysqli](#) connection pool. Depending on the web server deployment model, a PHP process may serve one or multiple requests. Therefore, a pooled connection may be used by one or more scripts subsequently.

### *Persistent connection*

If a unused persistent connection for a given combination of host, username, password, socket, port and default database can not be found in the connection pool, then [mysqli](#) opens a new connection. The use of persistent connections can be enabled and disabled using the PHP directive [mysqli.allow\\_persistent](#). The total number of connections opened by a script can be limited with [mysqli.max\\_links](#). The maximum number of persistent connections per PHP process can be restricted with [mysqli.max\\_persistent](#). Please note, that the web server may spawn many PHP processes.

A common complain about persistent connections is that their state is not reset before reuse. For example, open and unfinished transactions are not automatically rolled back. But also, authorization changes which happened in the time between putting the connection into the pool and reusing it are not reflected. This may be seen as an unwanted side-effect. On the contrary, the name [persistent](#) may be understood as a promise that the state is persisted.

The [mysqli](#) extension supports both interpretations of a persistent connection: state persisted, and state reset before reuse. The default is reset. Before a persistent connection is reused, the [mysqli](#) extension implicitly calls [mysqli\\_change\\_user](#) to reset the state. The persistent connection appears to the user as if it was just opened. No artifacts from previous usages are visible.

The [mysqli\\_change\\_user](#) function is an expensive operation. For best performance, users may want to recompile the extension with the compile flag [MYSQLI\\_NO\\_CHANGE\\_USER\\_ON\\_PCONNECT](#) being set.

It is left to the user to choose between safe behavior and best performance. Both are valid optimization goals. For ease of use, the safe behavior has been made the default at the expense of maximum performance.

### *See also*

[mysqli::\\_\\_construct](#)  
[mysqli::init](#)  
[mysqli::options](#)  
[mysqli::real\\_connect](#)  
[mysqli::change\\_user](#)  
[\\$mysqli::host\\_info](#)  
[MySQLi Configuration Options](#)  
[Persistent Database Connections](#)

## 21.9.3.3.3. Executing statements

Copyright 1997-2012 the PHP Documentation Group. [2906]

Statements can be executed with the [mysqli\\_query](#), [mysqli\\_real\\_query](#) and [mysqli\\_multi\\_query](#) functions. The [mysqli\\_query](#) function is the most common, and combines

the executing statement with a buffered fetch of its result set, if any, in one call. Calling `mysqli_query` is identical to calling `mysqli_real_query` followed by `mysqli_store_result`.

### Example 21.82. Connecting to MySQL

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}
?>
```

#### *Buffered result sets*

After statement execution results can be retrieved at once to be buffered by the client or by read row by row. Client-side result set buffering allows the server to free resources associated with the statement results as early as possible. Generally speaking, clients are slow consuming result sets. Therefore, it is recommended to use buffered result sets. `mysqli_query` combines statement execution and result set buffering.

PHP applications can navigate freely through buffered results. Navigation is fast because the result sets are held in client memory. Please, keep in mind that it is often easier to scale by client than it is to scale the server.

### Example 21.83. Navigation through buffered results

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$res = $mysqli->query("SELECT id FROM test ORDER BY id ASC");

echo "Reverse order...\n";
for ($row_no = $res->num_rows - 1; $row_no >= 0; $row_no--) {
    $res->data_seek($row_no);
    $row = $res->fetch_assoc();
    echo " id = " . $row['id'] . "\n";
}

echo "Result set order...\n";
$res->data_seek(0);
while ($row = $res->fetch_assoc()) {
    echo " id = " . $row['id'] . "\n";
}
?>
```

The above example will output:

```
Reverse order...
id = 3
id = 2
id = 1
Result set order...
id = 1
id = 2
id = 3
```

### *Unbuffered result sets*

If client memory is a short resource and freeing server resources as early as possible to keep server load low is not needed, unbuffered results can be used. Scrolling through unbuffered results is not possible before all rows have been read.

### **Example 21.84. Navigation through unbuffered results**

```
<?php
$mysqli->real_query("SELECT id FROM test ORDER BY id ASC");
$res = $mysqli->use_result();

echo "Result set order...\n";
while ($row = $res->fetch_assoc()) {
    echo " id = " . $row['id'] . "\n";
}
?>
```

### *Result set values data types*

The `mysqli_query`, `mysqli_real_query` and `mysqli_multi_query` functions are used to execute non-prepared statements. At the level of the MySQL Client Server Protocol, the command `COM_QUERY` and the text protocol are used for statement execution. With the text protocol, the MySQL server converts all data of a result sets into strings before sending. This conversion is done regardless of the SQL result set column data type. The mysql client libraries receive all column values as strings. No further client-side casting is done to convert columns back to their native types. Instead, all values are provided as PHP strings.

### **Example 21.85. Text protocol returns strings by default**

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$res = $mysqli->query("SELECT id, label FROM test WHERE id = 1");
$row = $res->fetch_assoc();

printf("id = %s (%s)\n", $row['id'], gettype($row['id']));
printf("label = %s (%s)\n", $row['label'], gettype($row['label']));
?>
```

The above example will output:



```
id = 1 (string)
label = a (string)
```

It is possible to convert integer and float columns back to PHP numbers by setting the [MYSQLI\\_OPT\\_INT\\_AND\\_FLOAT\\_NATIVE](#) connection option, if using the `mysqlnd` library. If set, the `mysqlnd` library will check the result set meta data column types and convert numeric SQL columns to PHP numbers, if the PHP data type value range allows for it. This way, for example, SQL INT columns are returned as integers.

#### Example 21.86. Native data types with `mysqlnd` and connection option

```
<?php
$mysqli = mysqli_init();
$mysqli->options(MYSQLI_OPT_INT_AND_FLOAT_NATIVE, 1);
$mysqli->real_connect("example.com", "user", "password", "database");

if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$res = $mysqli->query("SELECT id, label FROM test WHERE id = 1");
$row = $res->fetch_assoc();

printf("id = %s (%s)\n", $row['id'], gettype($row['id']));
printf("label = %s (%s)\n", $row['label'], gettype($row['label']));
?>
```

The above example will output:

```
id = 1 (integer)
label = a (string)
```

See also

```
mysqli::__construct
mysqli::init
mysqli::options
mysqli::real_connect
mysqli::query
mysqli::multi_query
mysqli::use_result
mysqli::store_result
mysqli_result::free
```

### 21.9.3.3.4. Prepared Statements

Copyright 1997-2012 the PHP Documentation Group. [2906]

The MySQL database supports prepared statements. A prepared statement or a parameterized statement is used to execute the same statement repeatedly with high efficiency.

*Basic workflow*

The prepared statement execution consists of two stages: prepare and execute. At the prepare stage a statement template is sent to the database server. The server performs a syntax check and initializes server internal resources for later use.

The MySQL server supports using anonymous, positional placeholder with `?`.

#### Example 21.87. First stage: prepare

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

/* Non-prepared statement */
if (!$mysqli->query("DROP TABLE IF EXISTS test") || !$mysqli->query("CREATE TABLE test(id INT)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

/* Prepared statement, stage 1: prepare */
if (!$stmt = $mysqli->prepare("INSERT INTO test(id) VALUES (?)")) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}
?>
```

Prepare is followed by execute. During execute the client binds parameter values and sends them to the server. The server creates a statement from the statement template and the bound values to execute it using the previously created internal resources.

#### Example 21.88. Second stage: bind and execute

```
<?php
/* Prepared statement, stage 2: bind and execute */
$id = 1;
if (!$stmt->bind_param("i", $id)) {
    echo "Binding parameters failed: (" . $stmt->errno . ") " . $stmt->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}
?>
```

#### *Repeated execution*

A prepared statement can be executed repeatedly. Upon every execution the current value of the bound variable is evaluated and sent to the server. The statement is not parsed again. The statement template is not transferred to the server again.

#### Example 21.89. INSERT prepared once, executed multiple times

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

/* Non-prepared statement */
if (!$mysqli->query("DROP TABLE IF EXISTS test") || !$mysqli->query("CREATE TABLE test(id INT)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

/* Prepared statement, stage 1: prepare */
```

```
if (!($stmt = $mysqli->prepare("INSERT INTO test(id) VALUES (?)")) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

/* Prepared statement, stage 2: bind and execute */
$id = 1;
if (!$stmt->bind_param("i", $id)) {
    echo "Binding parameters failed: (" . $stmt->errno . ") " . $stmt->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}

/* Prepared statement: repeated execution, only data transferred from client to server */
for ($id = 2; $id < 5; $id++) {
    if (!$stmt->execute()) {
        echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
    }
}

/* explicit close recommended */
$stmt->close();

/* Non-prepared statement */
$res = $mysqli->query("SELECT id FROM test");
var_dump($res->fetch_all());
?>
```

The above example will output:

```
array(4) {
  [0]=>
  array(1) {
    [0]=>
    string(1) "1"
  }
  [1]=>
  array(1) {
    [0]=>
    string(1) "2"
  }
  [2]=>
  array(1) {
    [0]=>
    string(1) "3"
  }
  [3]=>
  array(1) {
    [0]=>
    string(1) "4"
  }
}
```

Every prepared statement occupies server resources. Statements should be closed explicitly immediately after use. If not done explicitly, the statement will be closed when the statement handle is freed by PHP.

Using a prepared statement is not always the most efficient way of executing a statement. A prepared statement executed only once causes more client-server round-trips than a non-prepared statement. This is why the [SELECT](#) is not run as a prepared statement above.

Also, consider the use of the MySQL multi-INSERT SQL syntax for INSERTs. For the example, multi-INSERT requires less round-trips between the server and client than the prepared statement shown above.

**Example 21.90. Less round trips using multi-INSERT SQL**

```
<?php
if (!$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3), (4)")) {
    echo "Multi-INSERT failed: (" . $mysqli->errno . ") " . $mysqli->error;
}
?>
```

*Result set values data types*

The MySQL Client Server Protocol defines a different data transfer protocol for prepared statements and non-prepared statements. Prepared statements are using the so called binary protocol. The MySQL server sends result set data "as is" in binary format. Results are not serialized into strings before sending. The client libraries do not receive strings only. Instead, they will receive binary data and try to convert the values into appropriate PHP data types. For example, results from an SQL [INT](#) column will be provided as PHP integer variables.

**Example 21.91. Native datatypes**

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$stmt = $mysqli->prepare("SELECT id, label FROM test WHERE id = 1");
$stmt->execute();
$res = $stmt->get_result();
$row = $res->fetch_assoc();

printf("id = %s (%s)\n", $row['id'], gettype($row['id']));
printf("label = %s (%s)\n", $row['label'], gettype($row['label']));
?>
```

The above example will output:

```
id = 1 (integer)
label = a (string)
```

This behavior differs from non-prepared statements. By default, non-prepared statements return all results as strings. This default can be changed using a connection option. If the connection option is used, there are no differences.

*Fetching results using bound variables*

Results from prepared statements can either be retrieved by binding output variables, or by requesting a [mysqli\\_result](#) object.

Output variables must be bound after statement execution. One variable must be bound for every column of the statements result set.

**Example 21.92. Output variable binding**

```

<?php
mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt = $mysqli->prepare("SELECT id, label FROM test")) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$out_id = NULL;
$out_label = NULL;
if (!$stmt->bind_result($out_id, $out_label)) {
    echo "Binding output parameters failed: (" . $stmt->errno . ") " . $stmt->error;
}

while ($stmt->fetch()) {
    printf("id = %s (%s), label = %s (%s)\n", $out_id, gettype($out_id), $out_label, gettype($out_label));
}
?>

```

The above example will output:

```
id = 1 (integer), label = a (string)
```

Prepared statements return unbuffered result sets by default. The results of the statement are not implicitly fetched and transferred from the server to the client for client-side buffering. The result set takes server resources until all results have been fetched by the client. Thus it is recommended to consume results timely. If a client fails to fetch all results or the client closes the statement before having fetched all data, the data has to be fetched implicitly by `mysqli`.

It is also possible to buffer the results of a prepared statement using `mysqli_stmt_store_result`.

#### *Fetching results using `mysqli_result` interface*

Instead of using bound results, results can also be retrieved through the `mysqli_result` interface. `mysqli_stmt_get_result` returns a buffered result set.

#### **Example 21.93. Using `mysqli_result` to fetch results**

```

<?php
mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt = $mysqli->prepare("SELECT id, label FROM test ORDER BY id ASC")) {

```

```

    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}

if (!$res = $stmt->get_result()) {
    echo "Getting result set failed: (" . $stmt->errno . ") " . $stmt->error;
}

var_dump($res->fetch_all());
?>

```

The above example will output:

```

array(1) {
  [0]=>
  array(2) {
    [0]=>
    int(1)
    [1]=>
    string(1) "a"
  }
}

```

Using the [mysqli\\_result](#) interface offers the additional benefit of flexible client-side result set navigation.

#### Example 21.94. Buffered result set for flexible read out

```

<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT, label CHAR(1))") ||
    !$mysqli->query("INSERT INTO test(id, label) VALUES (1, 'a'), (2, 'b'), (3, 'c')")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt = $mysqli->prepare("SELECT id, label FROM test")) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}

if (!$res = $stmt->get_result()) {
    echo "Getting result set failed: (" . $stmt->errno . ") " . $stmt->error;
}

for ($row_no = ($res->num_rows - 1); $row_no >= 0; $row_no--) {
    $res->data_seek($row_no);
    var_dump($res->fetch_assoc());
}
$res->close();
?>

```

The above example will output:

```

array(2) {
  ["id"]=>
  int(3)
  ["label"]=>
  string(1) "c"
}
array(2) {
  ["id"]=>
  int(2)
  ["label"]=>
  string(1) "b"
}
array(2) {
  ["id"]=>
  int(1)
  ["label"]=>
  string(1) "a"
}

```

### *Escaping and SQL injection*

Bound variables will be escaped automatically by the server. The server inserts their escaped values at the appropriate places into the statement template before execution. A hint must be provided to the server for the type of bound variable, to create an appropriate conversion. See the [mysqli\\_stmt\\_bind\\_param](#) function for more information.

The automatic escaping of values within the server is sometimes considered a security feature to prevent SQL injection. The same degree of security can be achieved with non-prepared statements, if input values are escaped correctly.

### *Client-side prepared statement emulation*

The API does not include emulation for client-side prepared statement emulation.

### *Quick prepared - non-prepared statement comparison*

The table below compares server-side prepared and non-prepared statements.

**Table 21.39. Comparison of prepared and non-prepared statements**

	<b>Prepared Statement</b>	<b>Non-prepared statement</b>
Client-server round trips, SELECT, single execution	2	1
Statement string transferred from client to server	1	1
Client-server round trips, SELECT, repeated (n) execution	1 + n	n
Statement string transferred from client to server	1 template, n times bound parameter, if any	n times together with parameter, if any
Input parameter binding API	Yes, automatic input escaping	No, manual input escaping
Output variable binding API	Yes	No
Supports use of mysqli_result API	Yes, use <a href="#">mysqli_stmt_get_result</a>	Yes
Buffered result sets	Yes, use <a href="#">mysqli_stmt_get_result</a> or binding with <a href="#">mysqli_stmt_store_result</a>	Yes, default of <a href="#">mysqli_query</a>

	Prepared Statement	Non-prepared statement
Unbuffered result sets	Yes, use output binding API	Yes, use <code>mysqli_real_query</code> with <code>mysqli_use_result</code>
MySQL Client Server protocol data transfer flavor	Binary protocol	Text protocol
Result set values SQL data types	Preserved when fetching	Converted to string or preserved when fetching
Supports all SQL statements	Recent MySQL versions support most but not all	Yes

See also

```
mysqli::__construct
mysqli::query
mysqli::prepare
mysqli_stmt::prepare
mysqli_stmt::execute
mysqli_stmt::bind_param
mysqli_stmt::bind_result
```

### 21.9.3.3.5. Stored Procedures

Copyright 1997-2012 the PHP Documentation Group. [2906]

The MySQL database supports stored procedures. A stored procedure is a subroutine stored in the database catalog. Applications can call and execute the stored procedure. The `CALL` SQL statement is used to execute a stored procedure.

#### Parameter

Stored procedures can have `IN`, `INOUT` and `OUT` parameters, depending on the MySQL version. The `mysqli` interface has no special notion for the different kinds of parameters.

#### IN parameter

Input parameters are provided with the `CALL` statement. Please, make sure values are escaped correctly.

#### Example 21.95. Calling a stored procedure

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") || !$mysqli->query("CREATE TABLE test(id INT)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->query("DROP PROCEDURE IF EXISTS p") ||
    !$mysqli->query("CREATE PROCEDURE p(IN id_val INT) BEGIN INSERT INTO test(id) VALUES(id_val); END;")) {
    echo "Stored procedure creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->query("CALL p(1)")) {
    echo "CALL failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$res = $mysqli->query("SELECT id FROM test")) {
    echo "SELECT failed: (" . $mysqli->errno . ") " . $mysqli->error;
}
```



```

}

var_dump($res->fetch_assoc());
?>

```

The above example will output:

```

array(1) {
  ["id"]=>
  string(1) "1"
}

```

### *INOUT/OUT parameter*

The values of *INOUT/OUT* parameters are accessed using session variables.

### **Example 21.96. Using session variables**

```

<?php
mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP PROCEDURE IF EXISTS p") ||
    !$mysqli->query('CREATE PROCEDURE p(OUT msg VARCHAR(50)) BEGIN SELECT "Hi!" INTO msg; END;')) {
    echo "Stored procedure creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->query("SET @msg = ''") || !$mysqli->query("CALL p(@msg)")) {
    echo "CALL failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$res = $mysqli->query("SELECT @msg as _p_out")) {
    echo "Fetch failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$row = $res->fetch_assoc();
echo $row['_p_out'];
?>

```

The above example will output:

```

Hi!

```

Application and framework developers may be able to provide a more convenient API using a mix of session variables and databased catalog inspection. However, please note the possible performance impact of a custom solution based on catalog inspection.

### *Handling result sets*

Stored procedures can return result sets. Result sets returned from a stored procedure cannot be fetched correctly using `mysqli_query`. The `mysqli_query` function combines statement execution and fetching the first result set into a buffered result set, if any. However, there are additional stored procedure result sets hidden from the user which cause `mysqli_query` to fail returning the user expected result sets.

Result sets returned from a stored procedure are fetched using `mysqli_real_query` or `mysqli_multi_query`. Both functions allow fetching any number of result sets returned by a statement, such as `CALL`. Failing to fetch all result sets returned by a stored procedure causes an error.

### Example 21.97. Fetching results from stored procedures

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->query("DROP PROCEDURE IF EXISTS p") ||
    !$mysqli->query('CREATE PROCEDURE p() READS SQL DATA BEGIN SELECT id FROM test; SELECT id + 1 FROM test')) {
    echo "Stored procedure creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->multi_query("CALL p()")) {
    echo "CALL failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

do {
    if ($res = $mysqli->store_result()) {
        printf("---\n");
        var_dump($res->fetch_all());
        $res->free();
    } else {
        if ($mysqli->errno) {
            echo "Store failed: (" . $mysqli->errno . ") " . $mysqli->error;
        }
    }
} while ($mysqli->more_results() && $mysqli->next_result());
?>
```

The above example will output:

```
---
array(3) {
  [0]=>
  array(1) {
    [0]=>
    string(1) "1"
  }
  [1]=>
  array(1) {
    [0]=>
    string(1) "2"
  }
  [2]=>
  array(1) {
    [0]=>
    string(1) "3"
  }
}
---
array(3) {
  [0]=>
  array(1) {
    [0]=>
    string(1) "2"
  }
}
```

```

}
[1]=>
array(1) {
  [0]=>
    string(1) "3"
}
[2]=>
array(1) {
  [0]=>
    string(1) "4"
}
}

```

### *Use of prepared statements*

No special handling is required when using the prepared statement interface for fetching results from the same stored procedure as above. The prepared statement and non-prepared statement interfaces are similar. Please note, that not every MySQL server version may support preparing the [CALL SQL](#) statement.

### **Example 21.98. Stored Procedures and Prepared Statements**

```

<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$mysqli->query("DROP PROCEDURE IF EXISTS p") ||
    !$mysqli->query('CREATE PROCEDURE p() READS SQL DATA BEGIN SELECT id FROM test; SELECT id + 1 FROM test;')) {
    echo "Stored procedure creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!(($stmt = $mysqli->prepare("CALL p()"))) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}

do {
    if ($res = $stmt->get_result()) {
        printf("---\n");
        var_dump(mysqli_fetch_all($res));
        mysqli_free_result($res);
    } else {
        if ($stmt->errno) {
            echo "Store failed: (" . $stmt->errno . ") " . $stmt->error;
        }
    }
} while ($stmt->more_results() && $stmt->next_result());
?>

```

Of course, use of the bind API for fetching is supported as well.

### **Example 21.99. Stored Procedures and Prepared Statements using bind API**

```

<?php
if (!$stmt = $mysqli->prepare("CALL p())) {
    echo "Prepare failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

if (!$stmt->execute()) {
    echo "Execute failed: (" . $stmt->errno . ") " . $stmt->error;
}

do {

    $id_out = NULL;
    if (!$stmt->bind_result($id_out)) {
        echo "Bind failed: (" . $stmt->errno . ") " . $stmt->error;
    }

    while ($stmt->fetch()) {
        echo "id = $id_out\n";
    }
} while ($stmt->more_results() && $stmt->next_result());
?>

```

See also

[mysqli::query](#)  
[mysqli::multi\\_query](#)  
[mysqli\\_result::next-result](#)  
[mysqli\\_result::more-results](#)

### 21.9.3.3.6. Multiple Statements

Copyright 1997-2012 the PHP Documentation Group. [2906]

MySQL optionally allows having multiple statements in one statement string. Sending multiple statements at once reduces client-server round trips but requires special handling.

Multiple statements or multi queries must be executed with [mysqli\\_multi\\_query](#). The individual statements of the statement string are separated by semicolon. Then, all result sets returned by the executed statements must be fetched.

The MySQL server allows having statements that do return result sets and statements that do not return result sets in one multiple statement.

#### Example 21.100. Multiple Statements

```

<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

if (!$mysqli->query("DROP TABLE IF EXISTS test") || !$mysqli->query("CREATE TABLE test(id INT)")) {
    echo "Table creation failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

$sql = "SELECT COUNT(*) AS _num FROM test; ";
$sql.= "INSERT INTO test(id) VALUES (1); ";
$sql.= "SELECT COUNT(*) AS _num FROM test; ";

if (!$mysqli->multi_query($sql)) {
    echo "Multi query failed: (" . $mysqli->errno . ") " . $mysqli->error;
}

do {
    if ($res = $mysqli->store_result()) {
        var_dump($res->fetch_all(MYSQLI_ASSOC));
    }
}

```

```
        $res->free();
    }
} while ($mysqli->more_results() && $mysqli->next_result());
?>
```

The above example will output:

```
array(1) {
  [0]=>
  array(1) {
    ["_num"]=>
    string(1) "0"
  }
}
array(1) {
  [0]=>
  array(1) {
    ["_num"]=>
    string(1) "1"
  }
}
```

### *Security considerations*

The API functions [mysqli\\_query](#) and [mysqli\\_real\\_query](#) do not set a connection flag necessary for activating multi queries in the server. An extra API call is used for multiple statements to reduce the likeliness of accidental SQL injection attacks. An attacker may try to add statements such as `;` [DROP DATABASE mysql](#) or `;` [SELECT SLEEP\(999\)](#). If the attacker succeeds in adding SQL to the statement string but [mysqli\\_multi\\_query](#) is not used, the server will not execute the second, injected and malicious SQL statement.

### **Example 21.101. SQL Injection**

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
$res    = $mysqli->query("SELECT 1; DROP TABLE mysql.user");
if (!$res) {
    echo "Error executing query: (" . $mysqli->errno . ") " . $mysqli->error;
}
?>
```

The above example will output:

```
Error executing query: (1064) You have an error in your SQL syntax;
check the manual that corresponds to your MySQL server version for the right syntax
to use near 'DROP TABLE mysql.user' at line 1
```

### *Prepared statements*

Use of the multiple statement with prepared statements is not supported.

### *See also*

[mysqli::query](#)  
[mysqli::multi\\_query](#)  
[mysqli\\_result::next-result](#)  
[mysqli\\_result::more-results](#)

### 21.9.3.3.7. API support for transactions

Copyright 1997-2012 the PHP Documentation Group. [2906]

The MySQL server supports transactions depending on the storage engine used. Since MySQL 5.5, the default storage engine is InnoDB. InnoDB has full ACID transaction support.

Transactions can either be controlled using SQL or API calls. It is recommended to use API calls for enabling and disabling the auto commit mode and for committing and rolling back transactions.

#### Example 21.102. Setting auto commit mode with SQL and through the API

```
<?php
mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

/* Recommended: using API to control transactional settings */
mysqli->autocommit(false);

/* Won't be monitored and recognized by the replication and the load balancing plugin */
if (!$mysqli->query('SET AUTOCOMMIT = 0')) {
    echo "Query failed: (" . $mysqli->errno . ") " . $mysqli->error;
}
?>
```

Optional feature packages, such as the replication and load balancing plugin, can easily monitor API calls. The replication plugin offers transaction aware load balancing, if transactions are controlled with API calls. Transaction aware load balancing is not available if SQL statements are used for setting auto commit mode, committing or rolling back a transaction.

#### Example 21.103. Commit and rollback

```
<?php
mysqli = new mysqli("example.com", "user", "password", "database");
mysqli->autocommit(false);

mysqli->query("INSERT INTO test(id) VALUES (1)");
mysqli->rollback();

mysqli->query("INSERT INTO test(id) VALUES (2)");
mysqli->commit();
?>
```

Please note, that the MySQL server cannot roll back all statements. Some statements cause an implicit commit.

See also

[mysqli::autocommit](#)  
[mysqli\\_result::commit](#)  
[mysqli\\_result::rollback](#)

### 21.9.3.3.8. Metadata

Copyright 1997-2012 the PHP Documentation Group. [2906]

A MySQL result set contains metadata. The metadata describes the columns found in the result set. All metadata sent by MySQL is accessible through the [mysqli](#) interface. The extension performs no or

negligible changes to the information it receives. Differences between MySQL server versions are not aligned.

Meta data is access through the [mysqli\\_result](#) interface.

#### Example 21.104. Accessing result set meta data

```
<?php
$mysqli = new mysqli("example.com", "user", "password", "database");
if ($mysqli->connect_errno) {
    echo "Failed to connect to MySQL: (" . $mysqli->connect_errno . ") " . $mysqli->connect_error;
}

$res = $mysqli->query("SELECT 1 AS _one, 'Hello' AS _two FROM DUAL");
var_dump($res->fetch_fields());
?>
```

The above example will output:

```
array(2) {
  [0]=>
  object(stdClass)#3 (13) {
    ["name"]=>
    string(4) "_one"
    ["orgname"]=>
    string(0) ""
    ["table"]=>
    string(0) ""
    ["orgtable"]=>
    string(0) ""
    ["def"]=>
    string(0) ""
    ["db"]=>
    string(0) ""
    ["catalog"]=>
    string(3) "def"
    ["max_length"]=>
    int(1)
    ["length"]=>
    int(1)
    ["charsetnr"]=>
    int(63)
    ["flags"]=>
    int(32897)
    ["type"]=>
    int(8)
    ["decimals"]=>
    int(0)
  }
  [1]=>
  object(stdClass)#4 (13) {
    ["name"]=>
    string(4) "_two"
    ["orgname"]=>
    string(0) ""
    ["table"]=>
    string(0) ""
    ["orgtable"]=>
    string(0) ""
    ["def"]=>
    string(0) ""
    ["db"]=>
    string(0) ""
    ["catalog"]=>
    string(3) "def"
    ["max_length"]=>
    int(5)
  }
}
```

```
[ "length" ]=>
int(5)
[ "charsetnr" ]=>
int(8)
[ "flags" ]=>
int(1)
[ "type" ]=>
int(253)
[ "decimals" ]=>
int(31)
}
}
```

### Prepared statements

Meta data of result sets created using prepared statements are accessed the same way. A suitable `mysqli_result` handle is returned by `mysqli_stmt_result_metadata`.

### Example 21.105. Prepared statements metadata

```
<?php
$stmt = $mysqli->prepare("SELECT 1 AS _one, 'Hello' AS _two FROM DUAL");
$stmt->execute();
$res = $stmt->result_metadata();
var_dump($res->fetch_fields());
?>
```

See also

`mysqli::query`  
`mysqli_result::fetch_fields`

## 21.9.3.4. Installing/Configuring

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

### 21.9.3.4.1. Requirements

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

In order to have these functions available, you must compile PHP with support for the `mysqli` extension.



#### Note

The `mysqli` extension is designed to work with MySQL version 4.1.13 or newer, or 5.0.7 or newer. For previous versions, please see the [MySQL](#) extension documentation.

### 21.9.3.4.2. Installation

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The `mysqli` extension was introduced with PHP version 5.0.0. The MySQL Native Driver was included in PHP version 5.3.0.

#### 21.9.3.4.2.1. Installation on Linux

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The common Unix distributions include binary versions of PHP that can be installed. Although these binary versions are typically built with support for MySQL extensions enabled, the extension libraries



themselves may need to be installed using an additional package. Check the package manager that comes with your chosen distribution for availability.

Unless your Unix distribution comes with a binary package of PHP with the [mysqli](#) extension available, you will need to build PHP from source code. Building PHP from source allows you to specify the MySQL extensions you want to use, as well as your choice of client library for each extension.

The MySQL Native Driver is the recommended option, as it results in improved performance and gives access to features not available when using the MySQL Client Library. Refer to [What is PHP's MySQL Native Driver? \[2985\]](#) for a brief overview of the advantages of MySQL Native Driver.

The `/path/to/mysqli_config` represents the location of the `mysqli_config` program that comes with MySQL Server.

**Table 21.40. mysqli compile time support matrix**

PHP Version	Default	Configure Options: <a href="#">mysqli</a>	Configure Options: <a href="#">libmysqlclient</a>	Changelog
5.0.x, 5.1.x, 5.2.x	libmysqlclient	Not Available	<code>--with-mysqli=/path/to/mysqli_config</code>	
5.3.x	libmysqlclient	<code>--with-mysqli=mysqli</code>	<code>--with-mysqli=/path/to/mysqli_config</code>	mysqli is now supported
5.4.x	mysqli	<code>--with-mysqli</code>	<code>--with-mysqli=/path/to/mysqli_config</code>	mysqli is now the default

Note that it is possible to freely mix MySQL extensions and client libraries. For example, it is possible to enable the MySQL extension to use the MySQL Client Library (`libmysqlclient`), while configuring the [mysqli](#) extension to use the MySQL Native Driver. However, all permutations of extension and client library are possible.

The following example builds the MySQL extension to use the MySQL Client Library, and the [mysqli](#) and PDO MySQL extensions to use the MySQL Native Driver:

```
./configure --with-mysql=/usr/bin/mysqli_config \
--with-mysqli=mysqli \
--with-pdo-mysql=mysqli \
[other options]
```

#### 21.9.3.4.2.2. Installation on Windows Systems

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

On Windows, PHP is most commonly installed using the binary installer.

##### 21.9.3.4.2.2.1. PHP 5.0, 5.1, 5.2

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Once PHP has been installed, some configuration is required to enable [mysqli](#) and specify the client library you want it to use.

The [mysqli](#) extension is not enabled by default, so the `php_mysqli.dll` DLL must be enabled inside of `php.ini`. In order to do this you need to find the `php.ini` file (typically located

in `c:\php`), and make sure you remove the comment (semi-colon) from the start of the line `extension=php_mysqli.dll`, in the section marked `[PHP_MYSQLI]`.

Also, if you want to use the MySQL Client Library with `mysqli`, you need to make sure PHP can access the client library file. The MySQL Client Library is included as a file named `libmysql.dll` in the Windows PHP distribution. This file needs to be available in the Windows system's `PATH` environment variable, so that it can be successfully loaded. See the FAQ titled "[How do I add my PHP directory to the PATH on Windows](#)" for information on how to do this. Copying `libmysql.dll` to the Windows system directory (typically `c:\Windows\system`) also works, as the system directory is by default in the system's `PATH`. However, this practice is strongly discouraged.

As with enabling any PHP extension (such as `php_mysqli.dll`), the PHP directive `extension_dir` should be set to the directory where the PHP extensions are located. See also the [Manual Windows Installation Instructions](#). An example `extension_dir` value for PHP 5 is `c:\php\ext`.



#### Note

If when starting the web server an error similar to the following occurs: "Unable to load dynamic library './php\_mysqli.dll'", this is because `php_mysqli.dll` and/or `libmysql.dll` cannot be found by the system.

#### 21.9.3.4.2.2. PHP 5.3.0+

Copyright 1997-2012 the PHP Documentation Group. [2906]

On Windows, for PHP versions 5.3 and newer, the `mysqli` extension is enabled and uses the MySQL Native Driver by default. This means you don't need to worry about configuring access to `libmysql.dll`.

#### 21.9.3.4.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.41. MySQLi Configuration Options**

Name	Default	Changeable	Changelog
<code>mysqli.allow_local_infile</code>	"1"	PHP_INI_SYSTEM	Available since PHP 5.2.4.
<code>mysqli.allow_persistent</code>	"1"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqli.max_persistent</code>	"-1"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqli.max_links</code>	"-1"	PHP_INI_SYSTEM	Available since PHP 5.0.0.
<code>mysqli.default_port</code>	"3306"	PHP_INI_ALL	Available since PHP 5.0.0.
<code>mysqli.default_socket</code>	NULL	PHP_INI_ALL	Available since PHP 5.0.0.
<code>mysqli.default_host</code>	NULL	PHP_INI_ALL	Available since PHP 5.0.0.
<code>mysqli.default_user</code>	NULL	PHP_INI_ALL	Available since PHP 5.0.0.
<code>mysqli.default_pw</code>	NULL	PHP_INI_ALL	Available since PHP 5.0.0.

Name	Default	Changeable	Changelog
<a href="#">mysqli.reconnect</a>	"0"	PHP_INI_SYSTEM	Available since PHP 4.3.5.
<a href="#">mysqli.cache_size</a>	"2000"	PHP_INI_SYSTEM	Available since PHP 5.3.0.

For further details and definitions of the preceding `PHP_INI_*` constants, see the chapter on [configuration changes](#).

Here's a short explanation of the configuration directives.

<a href="#">mysqli.allow_local_infile</a> integer	Allow accessing, from PHP's perspective, local files with LOAD DATA statements
<a href="#">mysqli.allow_persistent</a> integer	Enable the ability to create persistent connections using <a href="#">mysqli_connect</a> .
<a href="#">mysqli.max_persistent</a> integer	Maximum of persistent connections that can be made. Set to 0 for unlimited.
<a href="#">mysqli.max_links</a> integer	The maximum number of MySQL connections per process.
<a href="#">mysqli.default_port</a> integer	The default TCP port number to use when connecting to the database server if no other port is specified. If no default is specified, the port will be obtained from the <code>MYSQL_TCP_PORT</code> environment variable, the <code>mysql-tcp</code> entry in <code>/etc/services</code> or the compile-time <code>MYSQL_PORT</code> constant, in that order. Win32 will only use the <code>MYSQL_PORT</code> constant.
<a href="#">mysqli.default_socket</a> string	The default socket name to use when connecting to a local database server if no other socket name is specified.
<a href="#">mysqli.default_host</a> string	The default server host to use when connecting to the database server if no other host is specified. Doesn't apply in <a href="#">safe mode</a> .
<a href="#">mysqli.default_user</a> string	The default user name to use when connecting to the database server if no other name is specified. Doesn't apply in <a href="#">safe mode</a> .
<a href="#">mysqli.default_pw</a> string	The default password to use when connecting to the database server if no other password is specified. Doesn't apply in <a href="#">safe mode</a> .
<a href="#">mysqli.reconnect</a> integer	Automatically reconnect if the connection was lost.
<a href="#">mysqli.cache_size</a> integer	Available only with <a href="#">mysqli</a> .

Users cannot set `MYSQL_OPT_READ_TIMEOUT` through an API call or runtime configuration setting. Note that if it were possible there would be differences between how `libmysqlclient` and streams would interpret the value of `MYSQL_OPT_READ_TIMEOUT`.

#### 21.9.3.4.4. Resource Types

Copyright 1997-2012 the PHP Documentation Group. [2906]

This extension has no resource types defined.

#### 21.9.3.5. The `mysqli` Extension and Persistent Connections

Copyright 1997-2012 the PHP Documentation Group. [2906]

Persistent connection support was introduced in PHP 5.3 for the `mysqli` extension. Support was already present in PDO MySQL and `ext/mysql`. The idea behind persistent connections is that a connection between a client process and a database can be reused by a client process, rather than

being created and destroyed multiple times. This reduces the overhead of creating fresh connections every time one is required, as unused connections are cached and ready to be reused.

Unlike the `mysql` extension, `mysqli` does not provide a separate function for opening persistent connections. To open a persistent connection you must prepend `p:` to the hostname when connecting.

The problem with persistent connections is that they can be left in unpredictable states by clients. For example, a table lock might be activated before a client terminates unexpectedly. A new client process reusing this persistent connection will get the connection “as is”. Any cleanup would need to be done by the new client process before it could make good use of the persistent connection, increasing the burden on the programmer.

The persistent connection of the `mysqli` extension however provides built-in cleanup handling code. The cleanup carried out by `mysqli` includes:

- Rollback active transactions
- Close and drop temporary tables
- Unlock tables
- Reset session variables
- Close prepared statements (always happens with PHP)
- Close handler
- Release locks acquired with `GET_LOCK`

This ensures that persistent connections are in a clean state on return from the connection pool, before the client process uses them.

The `mysqli` extension does this cleanup by automatically calling the C-API function `mysql_change_user()`.

The automatic cleanup feature has advantages and disadvantages though. The advantage is that the programmer no longer needs to worry about adding cleanup code, as it is called automatically. However, the disadvantage is that the code could *potentially* be a little slower, as the code to perform the cleanup needs to run each time a connection is returned from the connection pool.

It is possible to switch off the automatic cleanup code, by compiling PHP with `MYSQLI_NO_CHANGE_USER_ON_PCONNECT` defined.

**Note**

The `mysqli` extension supports persistent connections when using either MySQL Native Driver or MySQL Client Library.

### 21.9.3.6. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

`MYSQLI_READ_DEFAULT_GROUP` Read options from the named group from `my.cnf` or the file specified with `MYSQLI_READ_DEFAULT_FILE`

`MYSQLI_READ_DEFAULT_FILE` Read options from the named option file instead of from `my.cnf`

`MYSQLI_OPT_CONNECT_TIMEOUT` Connect timeout in seconds

`MYSQLI_OPT_LOCAL_INFILE` Enables command `LOAD LOCAL INFILE`

<code>MYSQLI_INIT_COMMAND</code>	Command to execute when connecting to MySQL server. Will automatically be re-executed when reconnecting.
<code>MYSQLI_CLIENT_SSL</code>	Use SSL (encrypted protocol). This option should not be set by application programs; it is set internally in the MySQL client library
<code>MYSQLI_CLIENT_COMPRESS</code>	Use compression protocol
<code>MYSQLI_CLIENT_INTERACTIVE</code>	Allow <code>interactive_timeout</code> seconds (instead of <code>wait_timeout</code> seconds) of inactivity before closing the connection. The client's session <code>wait_timeout</code> variable will be set to the value of the session <code>interactive_timeout</code> variable.
<code>MYSQLI_CLIENT_IGNORE_SPACE</code>	Allow spaces after function names. Makes all functions names reserved words.
<code>MYSQLI_CLIENT_NO_SCHEMA</code>	Don't allow the <code>db_name.tbl_name.col_name</code> syntax.
<code>MYSQLI_CLIENT_MULTI_QUERY</code>	Allows multiple semicolon-delimited queries in a single <code>mysqli_query</code> call.
<code>MYSQLI_STORE_RESULT</code>	For using buffered resultsets
<code>MYSQLI_USE_RESULT</code>	For using unbuffered resultsets
<code>MYSQLI_ASSOC</code>	Columns are returned into the array having the fieldname as the array index.
<code>MYSQLI_NUM</code>	Columns are returned into the array having an enumerated index.
<code>MYSQLI_BOTH</code>	Columns are returned into the array having both a numerical index and the fieldname as the associative index.
<code>MYSQLI_NOT_NULL_FLAG</code>	Indicates that a field is defined as <code>NOT NULL</code>
<code>MYSQLI_PRI_KEY_FLAG</code>	Field is part of a primary index
<code>MYSQLI_UNIQUE_KEY_FLAG</code>	Field is part of a unique index.
<code>MYSQLI_MULTIPLE_KEY_FLAG</code>	Field is part of an index.
<code>MYSQLI_BLOB_FLAG</code>	Field is defined as <code>BLOB</code>
<code>MYSQLI_UNSIGNED_FLAG</code>	Field is defined as <code>UNSIGNED</code>
<code>MYSQLI_ZEROFILL_FLAG</code>	Field is defined as <code>ZEROFILL</code>
<code>MYSQLI_AUTO_INCREMENT_FLAG</code>	Field is defined as <code>AUTO_INCREMENT</code>
<code>MYSQLI_TIMESTAMP_FLAG</code>	Field is defined as <code>TIMESTAMP</code>
<code>MYSQLI_SET_FLAG</code>	Field is defined as <code>SET</code>
<code>MYSQLI_NUM_FLAG</code>	Field is defined as <code>NUMERIC</code>
<code>MYSQLI_PART_KEY_FLAG</code>	Field is part of an multi-index
<code>MYSQLI_GROUP_FLAG</code>	Field is part of <code>GROUP BY</code>
<code>MYSQLI_TYPE_DECIMAL</code>	Field is defined as <code>DECIMAL</code>
<code>MYSQLI_TYPE_NEWDECIMAL</code>	Precision math <code>DECIMAL</code> or <code>NUMERIC</code> field (MySQL 5.0.3 and up)
<code>MYSQLI_TYPE_BIT</code>	Field is defined as <code>BIT</code> (MySQL 5.0.3 and up)

<code>MYSQLI_TYPE_TINY</code>	Field is defined as <code>TINYINT</code>
<code>MYSQLI_TYPE_SHORT</code>	Field is defined as <code>SMALLINT</code>
<code>MYSQLI_TYPE_LONG</code>	Field is defined as <code>INT</code>
<code>MYSQLI_TYPE_FLOAT</code>	Field is defined as <code>FLOAT</code>
<code>MYSQLI_TYPE_DOUBLE</code>	Field is defined as <code>DOUBLE</code>
<code>MYSQLI_TYPE_NULL</code>	Field is defined as <code>DEFAULT NULL</code>
<code>MYSQLI_TYPE_TIMESTAMP</code>	Field is defined as <code>TIMESTAMP</code>
<code>MYSQLI_TYPE_LONGLONG</code>	Field is defined as <code>BIGINT</code>
<code>MYSQLI_TYPE_INT24</code>	Field is defined as <code>MEDIUMINT</code>
<code>MYSQLI_TYPE_DATE</code>	Field is defined as <code>DATE</code>
<code>MYSQLI_TYPE_TIME</code>	Field is defined as <code>TIME</code>
<code>MYSQLI_TYPE_DATETIME</code>	Field is defined as <code>DATETIME</code>
<code>MYSQLI_TYPE_YEAR</code>	Field is defined as <code>YEAR</code>
<code>MYSQLI_TYPE_NEWDATE</code>	Field is defined as <code>DATE</code>
<code>MYSQLI_TYPE_INTERVAL</code>	Field is defined as <code>INTERVAL</code>
<code>MYSQLI_TYPE_ENUM</code>	Field is defined as <code>ENUM</code>
<code>MYSQLI_TYPE_SET</code>	Field is defined as <code>SET</code>
<code>MYSQLI_TYPE_TINY_BLOB</code>	Field is defined as <code>TINYBLOB</code>
<code>MYSQLI_TYPE_MEDIUM_BLOB</code>	Field is defined as <code>MEDIUMBLOB</code>
<code>MYSQLI_TYPE_LONG_BLOB</code>	Field is defined as <code>LOB</code>
<code>MYSQLI_TYPE_BLOB</code>	Field is defined as <code>BLOB</code>
<code>MYSQLI_TYPE_VAR_STRING</code>	Field is defined as <code>VARCHAR</code>
<code>MYSQLI_TYPE_STRING</code>	Field is defined as <code>CHAR</code> or <code>BINARY</code>
<code>MYSQLI_TYPE_CHAR</code>	Field is defined as <code>TINYINT</code> . For <code>CHAR</code> , see <code>MYSQLI_TYPE_STRING</code>
<code>MYSQLI_TYPE_GEOMETRY</code>	Field is defined as <code>GEOMETRY</code>
<code>MYSQLI_NEED_DATA</code>	More data available for bind variable
<code>MYSQLI_NO_DATA</code>	No more data available for bind variable
<code>MYSQLI_DATA_TRUNCATED</code>	Data truncation occurred. Available since PHP 5.1.0 and MySQL 5.0.5.
<code>MYSQLI_ENUM_FLAG</code>	Field is defined as <code>ENUM</code> . Available since PHP 5.3.0.
<code>MYSQLI_CURSOR_TYPE_FOR_UPDATE</code>	
<code>MYSQLI_CURSOR_TYPE_NO_CURSOR</code>	
<code>MYSQLI_CURSOR_TYPE_READ_ONLY</code>	

---

<code>MYSQLI_CURSOR_TYPE_SCROLLABLE</code>	
<code>MYSQLI_STMT_ATTR_CURSOR_TYPE</code>	
<code>MYSQLI_STMT_ATTR_PREFETCH_ROWS</code>	
<code>MYSQLI_STMT_ATTR_UPDATE_MAX_LENGTH</code>	
<code>MYSQLI_SET_CHARSET_NAME</code>	
<code>MYSQLI_REPORT_INDEX</code>	Report if no index or bad index was used in a query.
<code>MYSQLI_REPORT_ERROR</code>	Report errors from <code>mysqli</code> function calls.
<code>MYSQLI_REPORT_STRICT</code>	Throw a <code>mysqli_sql_exception</code> for errors instead of warnings.
<code>MYSQLI_REPORT_ALL</code>	Set all options on (report all).
<code>MYSQLI_REPORT_OFF</code>	Turns reporting off.
<code>MYSQLI_DEBUG_TRACE_ENABLED</code>	Is set to 1 if <code>mysqli_debug</code> functionality is enabled.
<code>MYSQLI_SERVER_QUERY_NO_GOOD_INDEX_USED</code>	
<code>MYSQLI_SERVER_QUERY_NO_INDEX_USED</code>	
<code>MYSQLI_REFRESH_GRANT</code>	Refreshes the grant tables.
<code>MYSQLI_REFRESH_LOG</code>	Flushes the logs, like executing the <code>FLUSH LOGS</code> SQL statement.
<code>MYSQLI_REFRESH_TABLES</code>	Flushes the table cache, like executing the <code>FLUSH TABLES</code> SQL statement.
<code>MYSQLI_REFRESH_HOSTS</code>	Flushes the host cache, like executing the <code>FLUSH HOSTS</code> SQL statement.
<code>MYSQLI_REFRESH_STATUS</code>	Reset the status variables, like executing the <code>FLUSH STATUS</code> SQL statement.
<code>MYSQLI_REFRESH_THREADS</code>	Flushes the thread cache.
<code>MYSQLI_REFRESH_SLAVE</code>	On a slave replication server: resets the master server information, and restarts the slave. Like executing the <code>RESET SLAVE</code> SQL statement.
<code>MYSQLI_REFRESH_MASTER</code>	On a master replication server: removes the binary log files listed in the binary log index, and truncates the index file. Like executing the <code>RESET MASTER</code> SQL statement.

### 21.9.3.7. Notes

Copyright 1997-2012 the PHP Documentation Group. [2906]

Some implementation notes:

1. Support was added for `MYSQL_TYPE_GEOMETRY` to the `MySQLi` extension in PHP 5.3.
2. Note there are different internal implementations within `libmysqlclient` and `mysqlnd` for handling columns of type `MYSQL_TYPE_GEOMETRY`. Generally speaking, `mysqlnd` will allocate significantly less memory. For example, if there is a `POINT` column in a result set, `libmysqlclient` may pre-allocate up to 4GB of RAM although less than 50 bytes are needed for holding a `POINT` column in memory. Memory allocation is much lower, less than 50 bytes, if using `mysqlnd`.

**21.9.3.8. The MySQLi Extension Function Summary**

Copyright 1997-2012 the PHP Documentation Group. [2906]

**Table 21.42. Summary of `mysqli` methods**

<b>mysqli Class</b>			
<b>OOP Interface</b>	<b>Procedural Interface</b>	<b>Alias (Do not use)</b>	<b>Description</b>
<i>Properties</i>			
<a href="#">\$mysqli::affected_rows</a>	<code>mysqli_affected_rows</code>	N/A	Gets the number of affected rows in a previous MySQL operation
<a href="#">\$mysqli::client_info</a>	<code>mysqli_get_client_info</code>	N/A	Returns the MySQL client version as a string
<a href="#">\$mysqli::client_version</a>	<code>mysqli_get_client_version</code>	N/A	Returns MySQL client version info as an integer
<a href="#">\$mysqli::connect_errno</a>	<code>mysqli_connect_errno</code>	N/A	Returns the error code from last connect call
<a href="#">\$mysqli::connect_error</a>	<code>mysqli_connect_error</code>	N/A	Returns a string description of the last connect error
<a href="#">\$mysqli::errno</a>	<code>mysqli_errno</code>	N/A	Returns the error code for the most recent function call
<a href="#">\$mysqli::error</a>	<code>mysqli_error</code>	N/A	Returns a string description of the last error
<a href="#">\$mysqli::field_count</a>	<code>mysqli_field_count</code>	N/A	Returns the number of columns for the most recent query
<a href="#">\$mysqli::host_info</a>	<code>mysqli_get_host_info</code>	N/A	Returns a string representing the type of connection used
<a href="#">\$mysqli::protocol_version</a>	<code>mysqli_get_proto_info</code>	N/A	Returns the version of the MySQL protocol used
<a href="#">\$mysqli::server_info</a>	<code>mysqli_get_server_info</code>	N/A	Returns the version of the MySQL server
<a href="#">\$mysqli::server_version</a>	<code>mysqli_get_server_version</code>	N/A	Returns the version of the MySQL server as an integer
<a href="#">\$mysqli::info</a>	<code>mysqli_info</code>	N/A	Retrieves information about the most recently executed query
<a href="#">\$mysqli::insert_id</a>	<code>mysqli_insert_id</code>	N/A	Returns the auto generated id used in the last query
<a href="#">\$mysqli::sqlstate</a>	<code>mysqli_sqlstate</code>	N/A	Returns the SQLSTATE error from previous MySQL operation



<b>mysqli Class</b>			
<b>OOP Interface</b>	<b>Procedural Interface</b>	<b>Alias (Do not use)</b>	<b>Description</b>
<a href="#">\$mysqli::warning_count</a>	<code>mysqli_warning_count</code>	N/A	Returns the number of warnings from the last query for the given link
<i>Methods</i>			
<code>mysqli::autocommit</code>	<code>mysqli_autocommit</code>	N/A	Turns on or off auto-committing database modifications
<code>mysqli::change_user</code>	<code>mysqli_change_user</code>	N/A	Changes the user of the specified database connection
<code>mysqli::character_set_name</code> <a href="#">mysqli::client_encoding</a>	<code>mysqli_character_set_name</code>	<code>mysqli_client_encoding</code>	Returns the default character set for the database connection
<code>mysqli::close</code>	<code>mysqli_close</code>	N/A	Closes a previously opened database connection
<code>mysqli::commit</code>	<code>mysqli_commit</code>	N/A	Commits the current transaction
<code>mysqli::__construct</code>	<code>mysqli_connect</code>	N/A	Open a new connection to the MySQL server [Note: static (i.e. class) method]
<code>mysqli::debug</code>	<code>mysqli_debug</code>	N/A	Performs debugging operations
<code>mysqli::dump_debug_info</code>	<code>mysqli_dump_debug_info</code>	N/A	Dump debugging information into the log
<code>mysqli::get_charset</code>	<code>mysqli_get_charset</code>	N/A	Returns a character set object
<code>mysqli::get_connection_stats</code>	<code>mysqli_get_connection_stats</code>	N/A	Returns client connection statistics. Available only with <a href="#">mysqliInD</a> .
<code>mysqli::get_client_info</code>	<code>mysqli_get_client_info</code>	N/A	Returns the MySQL client version as a string
<code>mysqli::get_client_statistics</code>	<code>mysqli_get_client_statistics</code>	N/A	Returns client per-process statistics. Available only with <a href="#">mysqliInD</a> .
<code>mysqli::get_cache_statistics</code>	<code>mysqli_get_cache_statistics</code>	N/A	Returns client Zval cache statistics. Available only with <a href="#">mysqliInD</a> .
<code>mysqli::get_server_info</code>	<code>mysqli_get_server_info</code>	N/A	NOT DOCUMENTED
<code>mysqli::get_warnings</code>	<code>mysqli_get_warnings</code>	N/A	NOT DOCUMENTED
<code>mysqli::init</code>	<code>mysqli_init</code>	N/A	Initializes MySQLi and returns a resource for use with <code>mysqli_real_connect</code> .

<b>mysqli Class</b>			
<b>OOP Interface</b>	<b>Procedural Interface</b>	<b>Alias (Do not use)</b>	<b>Description</b>
			[Not called on an object, as it returns a \$mysqli object.]
<code>mysqli::kill</code>	<code>mysqli_kill</code>	N/A	Asks the server to kill a MySQL thread
<code>mysqli::more_results</code>	<code>mysqli_more_results</code>	N/A	Check if there are any more query results from a multi query
<code>mysqli::multi_query</code>	<code>mysqli_multi_query</code>	N/A	Performs a query on the database
<code>mysqli::next_result</code>	<code>mysqli_next_result</code>	N/A	Prepare next result from multi_query
<code>mysqli::options</code>	<code>mysqli_options</code>	<code>mysqli_set_opt</code>	Set options
<code>mysqli::ping</code>	<code>mysqli_ping</code>	N/A	Pings a server connection, or tries to reconnect if the connection has gone down
<code>mysqli::prepare</code>	<code>mysqli_prepare</code>	N/A	Prepare an SQL statement for execution
<code>mysqli::query</code>	<code>mysqli_query</code>	N/A	Performs a query on the database
<code>mysqli::real_connect</code>	<code>mysqli_real_connect</code>	N/A	Opens a connection to a mysql server
<code>mysqli::real_escape_string</code> <code>mysqli::escape_string</code>	<code>mysqli_real_escape_string</code> <code>mysqli_escape_string</code>	<code>mysqli_escape_string</code>	Escapes special characters in a string for use in an SQL statement, taking into account the current charset of the connection
<code>mysqli::real_query</code>	<code>mysqli_real_query</code>	N/A	Execute an SQL query
<code>mysqli::refresh</code>	<code>mysqli_refresh</code>	N/A	Flushes tables or caches, or resets the replication server information
<code>mysqli::rollback</code>	<code>mysqli_rollback</code>	N/A	Rolls back current transaction
<code>mysqli::select_db</code>	<code>mysqli_select_db</code>	N/A	Selects the default database for database queries
<code>mysqli::set_charset</code>	<code>mysqli_set_charset</code>	N/A	Sets the default client character set
<code>mysqli::set_local_infile_handler_default</code>	<code>mysqli_set_local_infile_handler_default</code>	N/A	Unsets user defined handler for load local infile command

<b>mysqli Class</b>			
<b>OOP Interface</b>	<b>Procedural Interface</b>	<b>Alias (Do not use)</b>	<b>Description</b>
<code>mysqli::set_local_infile_handler</code>	<code>mysqli_set_local_infile_handler</code>	N/A	Set callback function for LOAD DATA LOCAL INFILE command
<code>mysqli::ssl_set</code>	<code>mysqli_ssl_set</code>	N/A	Used for establishing secure connections using SSL
<code>mysqli::stat</code>	<code>mysqli_stat</code>	N/A	Gets the current system status
<code>mysqli::stmt_init</code>	<code>mysqli_stmt_init</code>	N/A	Initializes a statement and returns an object for use with <code>mysqli_stmt_prepare</code>
<code>mysqli::store_result</code>	<code>mysqli_store_result</code>	N/A	Transfers a result set from the last query
<code>mysqli::thread_id</code>	<code>mysqli_thread_id</code>	N/A	Returns the thread ID for the current connection
<code>mysqli::thread_safe</code>	<code>mysqli_thread_safe</code>	N/A	Returns whether thread safety is given or not
<code>mysqli::use_result</code>	<code>mysqli_use_result</code>	N/A	Initiate a result set retrieval

**Table 21.43. Summary of `mysqli_stmt` methods**

<b>MySQL_STMT</b>			
<b>OOP Interface</b>	<b>Procedural Interface</b>	<b>Alias (Do not use)</b>	<b>Description</b>
<i>Properties</i>			
<code>\$mysqli_stmt::affected_rows</code>	<code>mysqli_stmt_affected_rows</code>	N/A	Returns the total number of rows changed, deleted, or inserted by the last executed statement
<code>\$mysqli_stmt::errno</code>	<code>mysqli_stmt_errno</code>	N/A	Returns the error code for the most recent statement call
<code>\$mysqli_stmt::error</code>	<code>mysqli_stmt_error</code>	N/A	Returns a string description for last statement error
<code>\$mysqli_stmt::field_count</code>	<code>mysqli_stmt_field_count</code>	N/A	Returns the number of field in the given statement - not documented
<code>\$mysqli_stmt::insert_id</code>	<code>mysqli_stmt_insert_id</code>	N/A	Get the ID generated from the previous INSERT operation
<code>\$mysqli_stmt::num_rows</code>	<code>mysqli_stmt_num_rows</code>	N/A	Return the number of rows in statements result set

MySQL_STMT			
OOP Interface	Procedural Interface	Alias (Do not use)	Description
<code>\$mysqli_stmt::param_count</code>	<code>mysqli_stmt_param_count</code>	<code>mysqli_param_count</code>	Returns the number of parameter for the given statement
<code>\$mysqli_stmt::sqlstate</code>	<code>mysqli_stmt_sqlstate</code>	N/A	Returns SQLSTATE error from previous statement operation
<i>Methods</i>			
<code>mysqli_stmt::attr_get</code>	<code>mysqli_stmt_attr_get</code>	N/A	Used to get the current value of a statement attribute
<code>mysqli_stmt::attr_set</code>	<code>mysqli_stmt_attr_set</code>	N/A	Used to modify the behavior of a prepared statement
<code>mysqli_stmt::bind_param</code>	<code>mysqli_stmt_bind_param</code>	<code>mysqli_bind_param</code>	Binds variables to a prepared statement as parameters
<code>mysqli_stmt::bind_result</code>	<code>mysqli_stmt_bind_result</code>	<code>mysqli_bind_result</code>	Binds variables to a prepared statement for result storage
<code>mysqli_stmt::close</code>	<code>mysqli_stmt_close</code>	N/A	Closes a prepared statement
<code>mysqli_stmt::data_seek</code>	<code>mysqli_stmt_data_seek</code>	N/A	Seeks to an arbitrary row in statement result set
<code>mysqli_stmt::execute</code>	<code>mysqli_stmt_execute</code>	<code>mysqli_execute</code>	Executes a prepared Query
<code>mysqli_stmt::fetch</code>	<code>mysqli_stmt_fetch</code>	<code>mysqli_fetch</code>	Fetch results from a prepared statement into the bound variables
<code>mysqli_stmt::free_result</code>	<code>mysqli_stmt_free_result</code>	N/A	Frees stored result memory for the given statement handle
<code>mysqli_stmt::get_result</code>	<code>mysqli_stmt_get_result</code>	N/A	Gets a result set from a prepared statement. Available only with <a href="#">mysqli</a> .
<code>mysqli_stmt::get_warnings</code>	<code>mysqli_stmt_get_warnings</code>	N/A	NOT DOCUMENTED
<code>\$mysqli_stmt::more_results()</code>	<code>mysqli_stmt_more_results()</code>	N/A	NOT DOCUMENTED Available only with <a href="#">mysqli</a> .
<code>\$mysqli_stmt::next_result()</code>	<code>mysqli_stmt_next_result()</code>	N/A	NOT DOCUMENTED Available only with <a href="#">mysqli</a> .
<code>mysqli_stmt::num_rows</code>	<code>mysqli_stmt_num_rows</code>	N/A	See also property <a href="#">\$mysqli_stmt::num_rows</a>
<code>mysqli_stmt::prepare</code>	<code>mysqli_stmt_prepare</code>	N/A	Prepare an SQL statement for execution

MySQL_STMT			
OOP Interface	Procedural Interface	Alias (Do not use)	Description
<code>mysqli_stmt::reset</code>	<code>mysqli_stmt_reset</code>	N/A	Resets a prepared statement
<code>mysqli_stmt::result_metadata</code>	<code>mysqli_stmt_result_metadata</code>	<code>mysqli_stmt_get_metadata</code>	Returns result set metadata from a prepared statement
<code>mysqli_stmt::send_long_data</code>	<code>mysqli_stmt_send_long_data</code>	<code>mysqli_stmt_send_long_data</code>	Send data in blocks
<code>mysqli_stmt::store_result</code>	<code>mysqli_stmt_store_result</code>	N/A	Transfers a result set from a prepared statement

Table 21.44. Summary of `mysqli_result` methods

mysqli_result			
OOP Interface	Procedural Interface	Alias (Do not use)	Description
<i>Properties</i>			
<code>\$mysqli_result::current_field</code>	<code>mysqli_field_tell</code>	N/A	Get current field offset of a result pointer
<code>\$mysqli_result::field_count</code>	<code>mysqli_num_fields</code>	N/A	Get the number of fields in a result
<code>\$mysqli_result::lengths</code>	<code>mysqli_fetch_lengths</code>	N/A	Returns the lengths of the columns of the current row in the result set
<code>\$mysqli_result::num_rows</code>	<code>mysqli_num_rows</code>	N/A	Gets the number of rows in a result
<i>Methods</i>			
<code>mysqli_result::data_seek</code>	<code>mysqli_data_seek</code>	N/A	Adjusts the result pointer to an arbitrary row in the result
<code>mysqli_result::fetch</code>	<code>mysqli_fetch_all</code>	N/A	Fetches all result rows and returns the result set as an associative array, a numeric array, or both. Available only with <code>mysqli</code> .
<code>mysqli_result::fetch_array</code>	<code>mysqli_fetch_array</code>	N/A	Fetch a result row as an associative, a numeric array, or both
<code>mysqli_result::fetch_assoc</code>	<code>mysqli_fetch_assoc</code>	N/A	Fetch a result row as an associative array
<code>mysqli_result::fetch_direct_field</code>	<code>mysqli_fetch_direct_field</code>	N/A	Fetch meta-data for a single field
<code>mysqli_result::fetch_field</code>	<code>mysqli_fetch_field</code>	N/A	Returns the next field in the result set
<code>mysqli_result::fetch_fields</code>	<code>mysqli_fetch_fields</code>	N/A	Returns an array of objects representing the fields in a result set

<code>mysqli_result</code>			
OOP Interface	Procedural Interface	Alias (Do not use)	Description
<code>mysqli_result::fetch_object</code>	<code>mysqli_fetch_object</code>	N/A	Returns the current row of a result set as an object
<code>mysqli_result::fetch_row</code>	<code>mysqli_fetch_row</code>	N/A	Get a result row as an enumerated array
<code>mysqli_result::field_seek</code>	<code>mysqli_field_seek</code>	N/A	Set result pointer to a specified field offset
<code>mysqli_result::free</code> , <code>mysqli_result::close</code> , <code>mysqli_result::free_result</code>	<code>mysqli_free_result</code>	N/A	Frees the memory associated with a result

**Table 21.45. Summary of `mysqli_driver` methods**

MySQL_Driver			
OOP Interface	Procedural Interface	Alias (Do not use)	Description
<i>Properties</i>			
N/A			
<i>Methods</i>			
<code>mysqli_driver::embedded_server_end</code>	<code>mysqli_embedded_server_end</code>	N/A	NOT DOCUMENTED
<code>mysqli_driver::embedded_server_start</code>	<code>mysqli_embedded_server_start</code>	N/A	NOT DOCUMENTED



**Note**

Alias functions are provided for backward compatibility purposes only. Do not use them in new projects.

### 21.9.3.9. The `mysqli` class (`mysqli`)

Copyright 1997-2012 the PHP Documentation Group. [2906]

Represents a connection between PHP and a MySQL database.

```

mysqli {
    mysqli

    Properties

    int
        mysqli->affected_rows ;

    string
        mysqli->client_info ;

    int
        mysqli->client_version ;

    string
        mysqli->connect_errno ;

    string
        mysqli->connect_error ;

    int
        mysqli->errno ;

    array

```

```
mysqli->error_list ;

string
mysqli->error ;

int
mysqli->field_count ;

int
mysqli->client_version ;

string
mysqli->host_info ;

string
mysqli->protocol_version ;

string
mysqli->server_info ;

int
mysqli->server_version ;

string
mysqli->info ;

mixed
mysqli->insert_id ;

string
mysqli->sqlstate ;

int
mysqli->thread_id ;

int
mysqli->warning_count ;

Methods

mysqli::__construct(
    string host
        = =ini_get("mysqli.default_host"),
    string username
        = =ini_get("mysqli.default_user"),
    string passwd
        = =ini_get("mysqli.default_pw"),
    string dbname
        = "=",
    int port
        = =ini_get("mysqli.default_port"),
    string socket
        = =ini_get("mysqli.default_socket"));

bool mysqli::autocommit(
    bool mode);

bool mysqli::change_user(
    string user,
    string password,
    string database);

string mysqli::character_set_name();

bool mysqli::close();

bool mysqli::commit();

bool mysqli::debug(
    string message);

bool mysqli::dump_debug_info();
```

```
object mysqli::get_charset();

string mysqli::get_client_info();

bool mysqli::get_connection_stats();

mysqli_warning mysqli::get_warnings();

mysqli mysqli::init();

bool mysqli::kill(
    int processid);

bool mysqli::more_results();

bool mysqli::multi_query(
    string query);

bool mysqli::next_result();

bool mysqli::options(
    int option,
    mixed value);

bool mysqli::ping();

public int mysqli::poll(
    array read,
    array error,
    array reject,
    int sec,
    int usec);

mysqli_stmt mysqli::prepare(
    string query);

mixed mysqli::query(
    string query,
    int resultmode
        = =MYSQLI_STORE_RESULT);

bool mysqli::real_connect(
    string host,
    string username,
    string passwd,
    string dbname,
    int port,
    string socket,
    int flags);

string mysqli::escape_string(
    string escapestr);

bool mysqli::real_query(
    string query);

public mysqli_result mysqli::reap_async_query();

public bool mysqli::refresh(
    int options);

bool mysqli::rollback();

int mysqli::rpl_query_type(
    string query);

bool mysqli::select_db(
    string dbname);

bool mysqli::send_query(
    string query);
```



```
bool mysqli::set_charset(
    string charset);

bool mysqli::set_local_infile_handler(
    mysqli link,
    callable read_func);

bool mysqli::ssl_set(
    string key,
    string cert,
    string ca,
    string capath,
    string cipher);

string mysqli::stat();

mysqli_stmt mysqli::stmt_init();

mysqli_result mysqli::store_result();

mysqli_result mysqli::use_result();
}
```

### 21.9.3.9.1. `mysqli::$affected_rows`, `mysqli_affected_rows`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$affected_rows`

`mysqli_affected_rows`

Gets the number of affected rows in a previous MySQL operation

#### Description

Object oriented style

```
int
mysqli->affected_rows ;
```

Procedural style

```
int mysqli_affected_rows(
    mysqli link);
```

Returns the number of rows affected by the last `INSERT`, `UPDATE`, `REPLACE` or `DELETE` query.

For `SELECT` statements `mysqli_affected_rows` works like `mysqli_num_rows`.

#### Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

#### Return Values

An integer greater than zero indicates the number of rows affected or retrieved. Zero indicates that no records were updated for an `UPDATE` statement, no rows matched the `WHERE` clause in the query or that no query has yet been executed. -1 indicates that the query returned an error.



#### Note

If the number of affected rows is greater than maximal int value, the number of affected rows will be returned as a string.

#### Examples

**Example 21.106. \$mysqli->affected\_rows example**

## Object oriented style

```

<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Insert rows */
$mysqli->query("CREATE TABLE Language SELECT * from CountryLanguage");
printf("Affected rows (INSERT): %d\n", $mysqli->affected_rows);

$mysqli->query("ALTER TABLE Language ADD Status int default 0");

/* update rows */
$mysqli->query("UPDATE Language SET Status=1 WHERE Percentage > 50");
printf("Affected rows (UPDATE): %d\n", $mysqli->affected_rows);

/* delete rows */
$mysqli->query("DELETE FROM Language WHERE Percentage < 50");
printf("Affected rows (DELETE): %d\n", $mysqli->affected_rows);

/* select all rows */
$result = $mysqli->query("SELECT CountryCode FROM Language");
printf("Affected rows (SELECT): %d\n", $mysqli->affected_rows);

$result->close();

/* Delete table Language */
$mysqli->query("DROP TABLE Language");

/* close connection */
$mysqli->close();
?>

```

## Procedural style

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

if (!$link) {
    printf("Can't connect to localhost. Error: %s\n", mysqli_connect_error());
    exit();
}

/* Insert rows */
mysqli_query($link, "CREATE TABLE Language SELECT * from CountryLanguage");
printf("Affected rows (INSERT): %d\n", mysqli_affected_rows($link));

mysqli_query($link, "ALTER TABLE Language ADD Status int default 0");

/* update rows */
mysqli_query($link, "UPDATE Language SET Status=1 WHERE Percentage > 50");
printf("Affected rows (UPDATE): %d\n", mysqli_affected_rows($link));

/* delete rows */
mysqli_query($link, "DELETE FROM Language WHERE Percentage < 50");
printf("Affected rows (DELETE): %d\n", mysqli_affected_rows($link));

/* select all rows */
$result = mysqli_query($link, "SELECT CountryCode FROM Language");
printf("Affected rows (SELECT): %d\n", mysqli_affected_rows($link));

```

```
mysqli_free_result($result);

/* Delete table Language */
mysqli_query($link, "DROP TABLE Language");

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Affected rows (INSERT): 984
Affected rows (UPDATE): 168
Affected rows (DELETE): 815
Affected rows (SELECT): 169
```

### See Also

[mysqli\\_num\\_rows](#)  
[mysqli\\_info](#)

#### 21.9.3.9.2. [mysqli::autocommit](#), [mysqli\\_autocommit](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::autocommit](#)

[mysqli\\_autocommit](#)

Turns on or off auto-committing database modifications

### Description

Object oriented style

```
bool mysqli::autocommit(
    bool mode);
```

Procedural style

```
bool mysqli_autocommit(
    mysqli link,
    bool mode);
```

Turns on or off auto-commit mode on queries for the database connection.

To determine the current state of autocommit use the SQL command `SELECT @@autocommit`.

### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

*mode* Whether to turn on auto-commit or not.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Notes

**Note**

This function doesn't work with non transactional table types (like MyISAM or ISAM).

**Examples****Example 21.107. `mysqli::autocommit` example**

## Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* turn autocommit on */
$mysqli->autocommit(TRUE);

if ($result = $mysqli->query("SELECT @@autocommit")) {
    $row = $result->fetch_row();
    printf("Autocommit is %s\n", $row[0]);
    $result->free();
}

/* close connection */
$mysqli->close();
?>
```

## Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

if (!$link) {
    printf("Can't connect to localhost. Error: %s\n", mysqli_connect_error());
    exit();
}

/* turn autocommit on */
mysqli_autocommit($link, TRUE);

if ($result = mysqli_query($link, "SELECT @@autocommit")) {
    $row = mysqli_fetch_row($result);
    printf("Autocommit is %s\n", $row[0]);
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Autocommit is 1
```

**See Also**

```
mysqli_commit
mysqli_rollback
```

### 21.9.3.9.3. *mysqli::change\_user*, *mysqli\_change\_user*

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli::change\_user*

```
mysqli_change_user
```

Changes the user of the specified database connection

#### Description

Object oriented style

```
bool mysqli::change_user(
    string user,
    string password,
    string database);
```

Procedural style

```
bool mysqli_change_user(
    mysqli link,
    string user,
    string password,
    string database);
```

Changes the user of the specified database connection and sets the current database.

In order to successfully change users a valid *username* and *password* parameters must be provided and that user must have sufficient permissions to access the desired database. If for any reason authorization fails, the current user authentication will remain.

#### Parameters

<i>link</i>	Procedural style only: A link identifier returned by <i>mysqli_connect</i> or <i>mysqli_init</i>
<i>user</i>	The MySQL user name.
<i>password</i>	The MySQL password.
<i>database</i>	The database to change to.

If desired, the `NULL` value may be passed resulting in only changing the user and not selecting a database. To select a database in this case use the *mysqli\_select\_db* function.

#### Return Values

Returns `TRUE` on success or `FALSE` on failure.

#### Notes



#### Note

Using this command will always cause the current database connection to behave as if was a completely new database connection, regardless of if the operation was completed successfully. This reset includes performing a rollback on any active transactions, closing all temporary tables, and unlocking all locked tables.

## Examples

### Example 21.108. `mysqli::change_user` example

#### Object oriented style

```
<?php
/* connect database test */
$mysqli = new mysqli("localhost", "my_user", "my_password", "test");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Set Variable a */
$mysqli->query("SET @a:=1");

/* reset all and select a new database */
$mysqli->change_user("my_user", "my_password", "world");

if ($result = $mysqli->query("SELECT DATABASE()")) {
    $row = $result->fetch_row();
    printf("Default database: %s\n", $row[0]);
    $result->close();
}

if ($result = $mysqli->query("SELECT @a")) {
    $row = $result->fetch_row();
    if ($row[0] === NULL) {
        printf("Value of variable a is NULL\n");
    }
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

#### Procedural style

```
<?php
/* connect database test */
$link = mysqli_connect("localhost", "my_user", "my_password", "test");

/* check connection */
if (!$link) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Set Variable a */
mysqli_query($link, "SET @a:=1");

/* reset all and select a new database */
mysqli_change_user($link, "my_user", "my_password", "world");

if ($result = mysqli_query($link, "SELECT DATABASE()")) {
    $row = mysqli_fetch_row($result);
    printf("Default database: %s\n", $row[0]);
    mysqli_free_result($result);
}

if ($result = mysqli_query($link, "SELECT @a")) {
    $row = mysqli_fetch_row($result);
}
```

```
if ($row[0] === NULL) {
    printf("Value of variable a is NULL\n");
}
mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Default database: world
Value of variable a is NULL
```

### See Also

[mysqli\\_connect](#)  
[mysqli\\_select\\_db](#)

#### 21.9.3.9.4. [mysqli::character\\_set\\_name](#), [mysqli\\_character\\_set\\_name](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::character\\_set\\_name](#)  
[mysqli\\_character\\_set\\_name](#)

Returns the default character set for the database connection

### Description

Object oriented style

```
string mysqli::character_set_name();
```

Procedural style

```
string mysqli_character_set_name(
    mysqli link);
```

Returns the current character set for the database connection.

### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

### Return Values

The default character set for the current connection

### Examples

#### Example 21.109. [mysqli::character\\_set\\_name](#) example

Object oriented style

```
<?php
```

```
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Print current character set */
$charset = $mysqli->character_set_name();
printf("Current character set is %s\n", $charset);

$mysqli->close();
?>
```

### Procedural style

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (!$link) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Print current character set */
$charset = mysqli_character_set_name($link);
printf("Current character set is %s\n", $charset);

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Current character set is latin1_swedish_ci
```

### See Also

[mysqli\\_set\\_charset](#)  
[mysqli\\_client\\_encoding](#)  
[mysqli\\_real\\_escape\\_string](#)

### 21.9.3.9.5. [mysqli::\\$client\\_info](#), [mysqli\\_get\\_client\\_info](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$client\\_info](#)  
[mysqli\\_get\\_client\\_info](#)

Get MySQL client info

### Description

Object oriented style

```
string
```



```
mysqli->client_info ;
```

### Procedural style

```
string mysqli_get_client_info(  
    mysqli link);
```

Returns a string that represents the MySQL client library version.

### Return Values

A string that represents the MySQL client library version

### Examples

#### Example 21.110. `mysqli_get_client_info`

```
<?php  
  
/* We don't need a connection to determine  
   the version of mysql client library */  
  
printf("Client library version: %s\n", mysqli_get_client_info());  
?>
```

### See Also

[mysqli\\_get\\_client\\_version](#)  
[mysqli\\_get\\_server\\_info](#)  
[mysqli\\_get\\_server\\_version](#)

#### 21.9.3.9.6. `mysqli::$client_version`, `mysqli_get_client_version`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$client_version`  
`mysqli_get_client_version`

Returns the MySQL client version as a string

### Description

Object oriented style

```
int  
mysqli->client_version ;
```

### Procedural style

```
int mysqli_get_client_version(  
    mysqli link);
```

Returns client version number as an integer.

### Return Values

A number that represents the MySQL client library version in format: `main_version*10000 + minor_version *100 + sub_version`. For example, 4.1.0 is returned as 40100.

This is useful to quickly determine the version of the client library to know if some capability exists.

### Examples

**Example 21.111. `mysqli_get_client_version`**

```
<?php
/* We don't need a connection to determine
   the version of mysql client library */
printf("Client library version: %d\n", mysqli_get_client_version());
?>
```

**See Also**

[mysqli\\_get\\_client\\_info](#)  
[mysqli\\_get\\_server\\_info](#)  
[mysqli\\_get\\_server\\_version](#)

**21.9.3.9.7. `mysqli::close`, `mysqli_close`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::close](#)

[mysqli\\_close](#)

Closes a previously opened database connection

**Description**

Object oriented style

```
bool mysqli::close();
```

Procedural style

```
bool mysqli_close(
    mysqli link);
```

Closes a previously opened database connection.

**Parameters**

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

**Return Values**

Returns [TRUE](#) on success or [FALSE](#) on failure.

**Examples**

See [mysqli\\_connect](#).

**See Also**

[mysqli::\\_\\_construct](#)  
[mysqli\\_init](#)  
[mysqli\\_real\\_connect](#)

**21.9.3.9.8. `mysqli::commit`, `mysqli_commit`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::commit`

`mysqli_commit`

Commits the current transaction

## Description

Object oriented style

```
bool mysqli::commit();
```

Procedural style

```
bool mysqli_commit(
    mysqli link);
```

Commits the current transaction for the database connection.

## Parameters

*link*

Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

## Return Values

Returns `TRUE` on success or `FALSE` on failure.

## Examples

### Example 21.112. `mysqli::commit` example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE Language LIKE CountryLanguage");

/* set autocommit to off */
$mysqli->autocommit(FALSE);

/* Insert some values */
$mysqli->query("INSERT INTO Language VALUES ('DEU', 'Bavarian', 'F', 11.2)");
$mysqli->query("INSERT INTO Language VALUES ('DEU', 'Swabian', 'F', 9.4)");

/* commit transaction */
$mysqli->commit();

/* drop table */
$mysqli->query("DROP TABLE Language");

/* close connection */
$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "test");

/* check connection */
if (!$link) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* set autocommit to off */
mysqli_autocommit($link, FALSE);

mysqli_query($link, "CREATE TABLE Language LIKE CountryLanguage");

/* Insert some values */
mysqli_query($link, "INSERT INTO Language VALUES ('DEU', 'Bavarian', 'F', 11.2)");
mysqli_query($link, "INSERT INTO Language VALUES ('DEU', 'Swabian', 'F', 9.4)");

/* commit transaction */
mysqli_commit($link);

/* close connection */
mysqli_close($link);
?>
```

### See Also

[mysqli\\_autocommit](#)  
[mysqli\\_rollback](#)

#### 21.9.3.9.9. [mysqli::\\$connect\\_errno](#), [mysqli\\_connect\\_errno](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$connect\\_errno](#)

[mysqli\\_connect\\_errno](#)

Returns the error code from last connect call

### Description

Object oriented style

```
string
mysqli->connect_errno ;
```

Procedural style

```
int mysqli_connect_errno();
```

Returns the last error code number from the last call to [mysqli\\_connect](#).



### Note

Client error message numbers are listed in the MySQL [errmsg.h](#) header file, server error message numbers are listed in [mysqld\\_error.h](#). In the MySQL source distribution you can find a complete list of error messages and error numbers in the file [Docs/mysqld\\_error.txt](#).

### Return Values

An error code value for the last call to [mysqli\\_connect](#), if it failed. zero means no error occurred.

### Examples

**Example 21.113. `$mysqli->connect_errno` example**

Object oriented style

```
<?php
$mysqli = @new mysqli('localhost', 'fake_user', 'my_password', 'my_db');

if ($mysqli->connect_errno) {
    die('Connect Error: ' . $mysqli->connect_errno);
}
?>
```

Procedural style

```
<?php
$link = @mysqli_connect('localhost', 'fake_user', 'my_password', 'my_db');

if (!$link) {
    die('Connect Error: ' . mysqli_connect_errno());
}
?>
```

The above examples will output:

```
Connect Error: 1045
```

**See Also**

[mysqli\\_connect](#)  
[mysqli\\_connect\\_error](#)  
[mysqli\\_errno](#)  
[mysqli\\_error](#)  
[mysqli\\_sqlstate](#)

**21.9.3.9.10. `mysqli::$connect_error`, `mysqli_connect_error`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$connect_error`

`mysqli_connect_error`

Returns a string description of the last connect error

**Description**

Object oriented style

```
string
mysqli->connect_error ;
```

Procedural style

```
string mysqli_connect_error();
```

Returns the last error message string from the last call to `mysqli_connect`.

**Return Values**

A string that describes the error. [NULL](#) is returned if no error occurred.

## Examples

### Example 21.114. `$mysqli->connect_error` example

Object oriented style

```
<?php
$mysqli = @new mysqli('localhost', 'fake_user', 'my_password', 'my_db');

// Works as of PHP 5.2.9 and 5.3.0.
if ($mysqli->connect_error) {
    die('Connect Error: ' . $mysqli->connect_error);
}
?>
```

Procedural style

```
<?php
$link = @mysqli_connect('localhost', 'fake_user', 'my_password', 'my_db');

if (!$link) {
    die('Connect Error: ' . mysqli_connect_error());
}
?>
```

The above examples will output:

```
Connect Error: Access denied for user 'fake_user'@'localhost' (using password: YES)
```

## Notes



### Warning

The `mysqli->connect_error` property only works properly as of PHP versions 5.2.9 and 5.3.0. Use the [mysqli\\_connect\\_error](#) function if compatibility with earlier PHP versions is required.

## See Also

[mysqli\\_connect](#)  
[mysqli\\_connect\\_errno](#)  
[mysqli\\_errno](#)  
[mysqli\\_error](#)  
[mysqli\\_sqlstate](#)

### 21.9.3.9.11. [mysqli::\\_\\_construct](#), [mysqli\\_connect](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\_\\_construct](#)

[mysqli\\_connect](#)

Open a new connection to the MySQL server

## Description

## Object oriented style

```
mysqli::__construct(
    string host
        = =ini_get("mysqli.default_host"),
    string username
        = =ini_get("mysqli.default_user"),
    string passwd
        = =ini_get("mysqli.default_pw"),
    string dbname
        = "",
    int port
        = =ini_get("mysqli.default_port"),
    string socket
        = =ini_get("mysqli.default_socket"));
```

## Procedural style

```
mysqli mysqli_connect(
    string host
        = =ini_get("mysqli.default_host"),
    string username
        = =ini_get("mysqli.default_user"),
    string passwd
        = =ini_get("mysqli.default_pw"),
    string dbname
        = "",
    int port
        = =ini_get("mysqli.default_port"),
    string socket
        = =ini_get("mysqli.default_socket"));
```

Opens a connection to the MySQL Server running on.

## Parameters

<i>host</i>	Can be either a host name or an IP address. Passing the <b>NULL</b> value or the string "localhost" to this parameter, the local host is assumed. When possible, pipes will be used instead of the TCP/IP protocol.  Prepending <i>host</i> by <b>p:</b> opens a persistent connection. <i>mysqli_change_user</i> is automatically called on connections opened from the connection pool.
<i>username</i>	The MySQL user name.
<i>passwd</i>	If not provided or <b>NULL</b> , the MySQL server will attempt to authenticate the user against those user records which have no password only. This allows one username to be used with different permissions (depending on if a password as provided or not).
<i>dbname</i>	If provided will specify the default database to be used when performing queries.
<i>port</i>	Specifies the port number to attempt to connect to the MySQL server.
<i>socket</i>	Specifies the socket or named pipe that should be used.

**Note**

Specifying the *socket* parameter will not explicitly determine the type of connection to be used when connecting to the MySQL server. How the connection is made to the

MySQL database is determined by the *host* parameter.

## Return Values

Returns an object which represents the connection to a MySQL Server.

## Changelog

Version	Description
5.3.0	Added the ability of persistent connections.

## Examples

### Example 21.115. *mysqli::\_\_construct* example

Object oriented style

```
<?php
$mysqli = new mysqli('localhost', 'my_user', 'my_password', 'my_db');

/*
 * This is the "official" OO way to do it,
 * BUT $connect_error was broken until PHP 5.2.9 and 5.3.0.
 */
if ($mysqli->connect_error) {
    die('Connect Error (' . $mysqli->connect_errno . ') '
        . $mysqli->connect_error);
}

/*
 * Use this instead of $connect_error if you need to ensure
 * compatibility with PHP versions prior to 5.2.9 and 5.3.0.
 */
if (mysqli_connect_error()) {
    die('Connect Error (' . mysqli_connect_errno() . ') '
        . mysqli_connect_error());
}

echo 'Success... ' . $mysqli->host_info . "\n";

$mysqli->close();
?>
```

Object oriented style when extending *mysqli* class

```
<?php

class foo_mysqli extends mysqli {
    public function __construct($host, $user, $pass, $db) {
        parent::__construct($host, $user, $pass, $db);

        if (mysqli_connect_error()) {
            die('Connect Error (' . mysqli_connect_errno() . ') '
                . mysqli_connect_error());
        }
    }
}

$db = new foo_mysqli('localhost', 'my_user', 'my_password', 'my_db');

echo 'Success... ' . $db->host_info . "\n";

$db->close();
?>
```



## Procedural style

```
<?php
$link = mysqli_connect('localhost', 'my_user', 'my_password', 'my_db');

if (!$link) {
    die('Connect Error (' . mysqli_connect_errno() . ') '
        . mysqli_connect_error());
}

echo 'Success... ' . mysqli_get_host_info($link) . "\n";

mysqli_close($link);
?>
```

The above examples will output:

```
Success... MySQL host info: localhost via TCP/IP
```

## Notes



### Note

MySQLnd always assumes the server default charset. This charset is sent during connection hand-shake/authentication, which MySQLnd will use.

Libmysqlclient uses the default charset set in the `my.cnf` or by an explicit call to `mysqli_options` prior to calling `mysqli_real_connect`, but after `mysqli_init`.



### Note

OO syntax only: If a connection fails an object is still returned. To check if the connection failed then use either the `mysqli_connect_error` function or the `mysqli->connect_error` property as in the preceding examples.



### Note

If it is necessary to set options, such as the connection timeout, `mysqli_real_connect` must be used instead.



### Note

Calling the constructor with no parameters is the same as calling `mysqli_init`.



### Note

Error "Can't create TCP/IP socket (10106)" usually means that the `variables_order` configure directive doesn't contain character `E`. On Windows, if the environment is not copied the `SYSTEMROOT` environment variable won't be available and PHP will have problems loading Winsock.

## See Also

[mysqli\\_real\\_connect](#)

```
mysqli_options  
mysqli_connect_errno  
mysqli_connect_error  
mysqli_close
```

### 21.9.3.9.12. [mysqli::debug](#), [mysqli\\_debug](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::debug](#)

[mysqli\\_debug](#)

Performs debugging operations

#### Description

Object oriented style

```
bool mysqli::debug(  
    string message);
```

Procedural style

```
bool mysqli_debug(  
    string message);
```

Performs debugging operations using the Fred Fish debugging library.

#### Parameters

*message* A string representing the debugging operation to perform

#### Return Values

Returns [TRUE](#) .

#### Notes



#### Note

To use the [mysqli\\_debug](#) function you must compile the MySQL client library to support debugging.

#### Examples

##### Example 21.116. Generating a Trace File

```
<?php  
  
/* Create a trace file in '/tmp/client.trace' on the local (client) machine: */  
mysqli_debug("d:t:o,/tmp/client.trace");  
  
?>
```

#### See Also

[mysqli\\_dump\\_debug\\_info](#)  
[mysqli\\_report](#)

### 21.9.3.9.13. [mysqli::dump\\_debug\\_info](#), [mysqli\\_dump\\_debug\\_info](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::dump_debug_info`

`mysqli_dump_debug_info`

Dump debugging information into the log

### Description

Object oriented style

```
bool mysqli::dump_debug_info();
```

Procedural style

```
bool mysqli_dump_debug_info(  
    mysqli link);
```

This function is designed to be executed by an user with the SUPER privilege and is used to dump debugging information into the log for the MySQL Server relating to the connection.

### Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### See Also

`mysqli_debug`

#### 21.9.3.9.14. `mysqli::$errno`, `mysqli_errno`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$errno`

`mysqli_errno`

Returns the error code for the most recent function call

### Description

Object oriented style

```
int  
mysqli->errno ;
```

Procedural style

```
int mysqli_errno(  
    mysqli link);
```

Returns the last error code for the most recent MySQLi function call that can succeed or fail.

Client error message numbers are listed in the MySQL `errmsg.h` header file, server error message numbers are listed in `mysqld_error.h`. In the MySQL source distribution you can find a complete list of error messages and error numbers in the file `Docs/mysqld_error.txt`.

### Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

## Return Values

An error code value for the last call, if it failed. zero means no error occurred.

## Examples

### Example 21.117. `$mysqli->errno` example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if ($mysqli->connect_errno) {
    printf("Connect failed: %s\n", $mysqli->connect_error);
    exit();
}

if (!$mysqli->query("SET a=1")) {
    printf("Errorcode: %d\n", $mysqli->errno);
}

/* close connection */
$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if (!mysqli_query($link, "SET a=1")) {
    printf("Errorcode: %d\n", mysqli_errno($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Errorcode: 1193
```

## See Also

[mysqli\\_connect\\_errno](#)  
[mysqli\\_connect\\_error](#)  
[mysqli\\_error](#)  
[mysqli\\_sqlstate](#)

### 21.9.3.9.15. `mysqli::$error_list`, `mysqli_error_list`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$error_list`

`mysqli_error_list`

Returns a list of errors from the last command executed

## Description

Object oriented style

```
array
mysqli->error_list ;
```

Procedural style

```
array mysqli_error_list(
mysqli link);
```

Returns a array of errors for the most recent MySQLi function call that can succeed or fail.

## Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

## Return Values

A list of errors, each as an associative array containing the errno, error, and sqlstate.

## Examples

### Example 21.118. `$mysqli->error_list` example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "nobody", "");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if (!$mysqli->query("SET a=1")) {
    print_r($mysqli->error_list);
}

/* close connection */
$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if (!mysqli_query($link, "SET a=1")) {
```

```
print_r(mysqli_error_list($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Array
(
    [0] => Array
        (
            [errno] => 1193
            [sqlstate] => HY000
            [error] => Unknown system variable 'a'
        )
)
```

### See Also

[mysqli\\_connect\\_errno](#)  
[mysqli\\_connect\\_error](#)  
[mysqli\\_error](#)  
[mysqli\\_sqlstate](#)

#### 21.9.3.9.16. [mysqli::\\$error](#), [mysqli\\_error](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$error](#)

[mysqli\\_error](#)

Returns a string description of the last error

### Description

Object oriented style

```
string
mysqli->error ;
```

Procedural style

```
string mysqli_error(
    mysqli link);
```

Returns the last error message for the most recent MySQLi function call that can succeed or fail.

### Parameters

[link](#) Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

### Return Values

A string that describes the error. An empty string if no error occurred.

### Examples

**Example 21.119. `$mysqli->error` example**

## Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if ($mysqli->connect_errno) {
    printf("Connect failed: %s\n", $mysqli->connect_error);
    exit();
}

if (!$mysqli->query("SET a=1")) {
    printf("Errormessage: %s\n", $mysqli->error);
}

/* close connection */
$mysqli->close();
?>
```

## Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if (!mysqli_query($link, "SET a=1")) {
    printf("Errormessage: %s\n", mysqli_error($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Errormessage: Unknown system variable 'a'
```

**See Also**

[mysqli\\_connect\\_errno](#)  
[mysqli\\_connect\\_error](#)  
[mysqli\\_errno](#)  
[mysqli\\_sqlstate](#)

**21.9.3.9.17. `mysqli::$field_count`, `mysqli_field_count`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$field\\_count](#)  
[mysqli\\_field\\_count](#)

Returns the number of columns for the most recent query

## Description

Object oriented style

```
int
mysqli->field_count ;
```

Procedural style

```
int mysqli_field_count(
mysqli link);
```

Returns the number of columns for the most recent query on the connection represented by the [link](#) parameter. This function can be useful when using the [mysqli\\_store\\_result](#) function to determine if the query should have produced a non-empty result set or not without knowing the nature of the query.

## Parameters

[link](#) Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

An integer representing the number of fields in a result set.

## Examples

### Example 21.120. [\\$mysqli->field\\_count](#) example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "test");

$mysqli->query( "DROP TABLE IF EXISTS friends");
$mysqli->query( "CREATE TABLE friends (id int, name varchar(20))");

$mysqli->query( "INSERT INTO friends VALUES (1,'Hartmut'), (2, 'Ulf')");

$mysqli->real_query("SELECT * FROM friends");

if ($mysqli->field_count) {
    /* this was a select/show or describe query */
    $result = $mysqli->store_result();

    /* process resultset */
    $row = $result->fetch_row();

    /* free resultset */
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "test");
```



```
mysqli_query($link, "DROP TABLE IF EXISTS friends");
mysqli_query($link, "CREATE TABLE friends (id int, name varchar(20))");

mysqli_query($link, "INSERT INTO friends VALUES (1,'Hartmut'), (2, 'Ulf')");

mysqli_real_query($link, "SELECT * FROM friends");

if (mysqli_field_count($link)) {
    /* this was a select/show or describe query */
    $result = mysqli_store_result($link);

    /* process resultset */
    $row = mysqli_fetch_row($result);

    /* free resultset */
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

### 21.9.3.9.18. [mysqli::get\\_charset](#), [mysqli\\_get\\_charset](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::get\\_charset](#)  
[mysqli\\_get\\_charset](#)

Returns a character set object

#### Description

Object oriented style

```
object mysqli::get_charset();
```

Procedural style

```
object mysqli_get_charset(
    mysqli link);
```

Returns a character set object providing several properties of the current active character set.

#### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

#### Return Values

The function returns a character set object with the following properties:

<i>charset</i>	Character set name
<i>collation</i>	Collation name
<i>dir</i>	Directory the charset description was fetched from (?) or "" for built-in character sets
<i>min_length</i>	Minimum character length in bytes
<i>max_length</i>	Maximum character length in bytes

---

<code>number</code>	Internal character set number
<code>state</code>	Character set status (?)

## Examples

### Example 21.121. `mysqli::get_charset` example

Object oriented style

```
<?php
$db = mysqli_init();
$db->real_connect("localhost","root","","test");
var_dump($db->get_charset());
?>
```

Procedural style

```
<?php
$db = mysqli_init();
mysqli_real_connect($db, "localhost","root","","test");
var_dump(mysqli_get_charset($db));
?>
```

The above examples will output:

```
object(stdClass)#2 (7) {
  ["charset"]=>
  string(6) "latin1"
  ["collation"]=>
  string(17) "latin1_swedish_ci"
  ["dir"]=>
  string(0) ""
  ["min_length"]=>
  int(1)
  ["max_length"]=>
  int(1)
  ["number"]=>
  int(8)
  ["state"]=>
  int(801)
}
```

## See Also

[mysqli\\_character\\_set\\_name](#)  
[mysqli\\_set\\_charset](#)

### 21.9.3.9.19. `mysqli::get_client_info`, `mysqli_get_client_info`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::get\\_client\\_info](#)

[mysqli\\_get\\_client\\_info](#)

Get MySQL client info

## Description

### Object oriented style

```
string mysqli::get_client_info();
```

### Procedural style

```
string mysqli_get_client_info(  
    mysqli link);
```

Returns a string that represents the MySQL client library version.

### Return Values

A string that represents the MySQL client library version

### Examples

#### Example 21.122. `mysqli_get_client_info`

```
<?php  
  
/* We don't need a connection to determine  
   the version of mysql client library */  
  
printf("Client library version: %s\n", mysqli_get_client_info());  
?>
```

### See Also

[mysqli\\_get\\_client\\_version](#)  
[mysqli\\_get\\_server\\_info](#)  
[mysqli\\_get\\_server\\_version](#)

#### 21.9.3.9.20. `mysqli_get_client_stats`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_get\\_client\\_stats](#)

Returns client per-process statistics

### Description

```
array mysqli_get_client_stats();
```

Returns client per-process statistics. Available only with [mysqlnd](#).

### Parameters

### Return Values

Returns an array with client stats if success, [FALSE](#) otherwise.

### Examples

#### Example 21.123. A `mysqli_get_client_stats` example

```
<?php  
$link = mysqli_connect();  
print_r(mysqli_get_client_stats());  
?>
```

The above example will output something similar to:

```
Array
(
    [bytes_sent] => 43
    [bytes_received] => 80
    [packets_sent] => 1
    [packets_received] => 2
    [protocol_overhead_in] => 8
    [protocol_overhead_out] => 4
    [bytes_received_ok_packet] => 11
    [bytes_received_eof_packet] => 0
    [bytes_received_rset_header_packet] => 0
    [bytes_received_rset_field_meta_packet] => 0
    [bytes_received_rset_row_packet] => 0
    [bytes_received_prepare_response_packet] => 0
    [bytes_received_change_user_packet] => 0
    [packets_sent_command] => 0
    [packets_received_ok] => 1
    [packets_received_eof] => 0
    [packets_received_rset_header] => 0
    [packets_received_rset_field_meta] => 0
    [packets_received_rset_row] => 0
    [packets_received_prepare_response] => 0
    [packets_received_change_user] => 0
    [result_set_queries] => 0
    [non_result_set_queries] => 0
    [no_index_used] => 0
    [bad_index_used] => 0
    [slow_queries] => 0
    [buffered_sets] => 0
    [unbuffered_sets] => 0
    [ps_buffered_sets] => 0
    [ps_unbuffered_sets] => 0
    [flushed_normal_sets] => 0
    [flushed_ps_sets] => 0
    [ps_prepared_never_executed] => 0
    [ps_prepared_once_executed] => 0
    [rows_fetched_from_server_normal] => 0
    [rows_fetched_from_server_ps] => 0
    [rows_buffered_from_client_normal] => 0
    [rows_buffered_from_client_ps] => 0
    [rows_fetched_from_client_normal_buffered] => 0
    [rows_fetched_from_client_normal_unbuffered] => 0
    [rows_fetched_from_client_ps_buffered] => 0
    [rows_fetched_from_client_ps_unbuffered] => 0
    [rows_fetched_from_client_ps_cursor] => 0
    [rows_skipped_normal] => 0
    [rows_skipped_ps] => 0
    [copy_on_write_saved] => 0
    [copy_on_write_performed] => 0
    [command_buffer_too_small] => 0
    [connect_success] => 1
    [connect_failure] => 0
    [connection_reused] => 0
    [reconnect] => 0
    [pconnect_success] => 0
    [active_connections] => 1
    [active_persistent_connections] => 0
    [explicit_close] => 0
    [implicit_close] => 0
    [disconnect_close] => 0
    [in_middle_of_command_close] => 0
    [explicit_free_result] => 0
    [implicit_free_result] => 0
    [explicit_stmt_close] => 0
    [implicit_stmt_close] => 0
    [mem_emalloc_count] => 0
    [mem_emalloc_ammount] => 0
)
```

```
[mem_ecalloc_count] => 0
[mem_ecalloc_ammount] => 0
[mem_erealloc_count] => 0
[mem_erealloc_ammount] => 0
[mem_efree_count] => 0
[mem_malloc_count] => 0
[mem_malloc_ammount] => 0
[mem_calloc_count] => 0
[mem_calloc_ammount] => 0
[mem_realloc_count] => 0
[mem_realloc_ammount] => 0
[mem_free_count] => 0
[proto_text_fetched_null] => 0
[proto_text_fetched_bit] => 0
[proto_text_fetched_tinyint] => 0
[proto_text_fetched_short] => 0
[proto_text_fetched_int24] => 0
[proto_text_fetched_int] => 0
[proto_text_fetched_bigint] => 0
[proto_text_fetched_decimal] => 0
[proto_text_fetched_float] => 0
[proto_text_fetched_double] => 0
[proto_text_fetched_date] => 0
[proto_text_fetched_year] => 0
[proto_text_fetched_time] => 0
[proto_text_fetched_datetime] => 0
[proto_text_fetched_timestamp] => 0
[proto_text_fetched_string] => 0
[proto_text_fetched_blob] => 0
[proto_text_fetched_enum] => 0
[proto_text_fetched_set] => 0
[proto_text_fetched_geometry] => 0
[proto_text_fetched_other] => 0
[proto_binary_fetched_null] => 0
[proto_binary_fetched_bit] => 0
[proto_binary_fetched_tinyint] => 0
[proto_binary_fetched_short] => 0
[proto_binary_fetched_int24] => 0
[proto_binary_fetched_int] => 0
[proto_binary_fetched_bigint] => 0
[proto_binary_fetched_decimal] => 0
[proto_binary_fetched_float] => 0
[proto_binary_fetched_double] => 0
[proto_binary_fetched_date] => 0
[proto_binary_fetched_year] => 0
[proto_binary_fetched_time] => 0
[proto_binary_fetched_datetime] => 0
[proto_binary_fetched_timestamp] => 0
[proto_binary_fetched_string] => 0
[proto_binary_fetched_blob] => 0
[proto_binary_fetched_enum] => 0
[proto_binary_fetched_set] => 0
[proto_binary_fetched_geometry] => 0
[proto_binary_fetched_other] => 0
)
```

### See Also

[Stats description](#)

#### 21.9.3.9.21. `mysqli_get_client_version`, `mysqli::$client_version`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_get_client_version`

`mysqli::$client_version`

Returns the MySQL client version as a string

## Description

Object oriented style

```
int
mysqli->client_version ;
```

Procedural style

```
int mysqli_get_client_version(
mysqli link);
```

Returns client version number as an integer.

## Return Values

A number that represents the MySQL client library version in format: `main_version*10000 + minor_version *100 + sub_version`. For example, 4.1.0 is returned as 40100.

This is useful to quickly determine the version of the client library to know if some capability exists.

## Examples

### Example 21.124. `mysqli_get_client_version`

```
<?php
/* We don't need a connection to determine
   the version of mysql client library */

printf("Client library version: %d\n", mysqli_get_client_version());
?>
```

## See Also

[mysqli\\_get\\_client\\_info](#)  
[mysqli\\_get\\_server\\_info](#)  
[mysqli\\_get\\_server\\_version](#)

### 21.9.3.9.22. [mysqli::get\\_connection\\_stats](#), [mysqli\\_get\\_connection\\_stats](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::get\\_connection\\_stats](#)

[mysqli\\_get\\_connection\\_stats](#)

Returns statistics about the client connection

## Description

Object oriented style

```
bool mysqli::get_connection_stats();
```

Procedural style

```
array mysqli_get_connection_stats(
mysqli link);
```

Returns statistics about the client connection. Available only with [mysqlnd](#).

## Parameters

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

Returns an array with connection stats if success, `FALSE` otherwise.

## Examples

### Example 21.125. A `mysqli_get_connection_stats` example

```
<?php
$link = mysqli_connect();
print_r(mysqli_get_connection_stats($link));
?>
```

The above example will output something similar to:

```
Array
(
    [bytes_sent] => 43
    [bytes_received] => 80
    [packets_sent] => 1
    [packets_received] => 2
    [protocol_overhead_in] => 8
    [protocol_overhead_out] => 4
    [bytes_received_ok_packet] => 11
    [bytes_received_eof_packet] => 0
    [bytes_received_rset_header_packet] => 0
    [bytes_received_rset_field_meta_packet] => 0
    [bytes_received_rset_row_packet] => 0
    [bytes_received_prepare_response_packet] => 0
    [bytes_received_change_user_packet] => 0
    [packets_sent_command] => 0
    [packets_received_ok] => 1
    [packets_received_eof] => 0
    [packets_received_rset_header] => 0
    [packets_received_rset_field_meta] => 0
    [packets_received_rset_row] => 0
    [packets_received_prepare_response] => 0
    [packets_received_change_user] => 0
    [result_set_queries] => 0
    [non_result_set_queries] => 0
    [no_index_used] => 0
    [bad_index_used] => 0
    [slow_queries] => 0
    [buffered_sets] => 0
    [unbuffered_sets] => 0
    [ps_buffered_sets] => 0
    [ps_unbuffered_sets] => 0
    [flushed_normal_sets] => 0
    [flushed_ps_sets] => 0
    [ps_prepared_never_executed] => 0
    [ps_prepared_once_executed] => 0
    [rows_fetched_from_server_normal] => 0
    [rows_fetched_from_server_ps] => 0
    [rows_buffered_from_client_normal] => 0
    [rows_buffered_from_client_ps] => 0
    [rows_fetched_from_client_normal_buffered] => 0
    [rows_fetched_from_client_normal_unbuffered] => 0
    [rows_fetched_from_client_ps_buffered] => 0
    [rows_fetched_from_client_ps_unbuffered] => 0
    [rows_fetched_from_client_ps_cursor] => 0
    [rows_skipped_normal] => 0
    [rows_skipped_ps] => 0
    [copy_on_write_saved] => 0
    [copy_on_write_performed] => 0
    [command_buffer_too_small] => 0
    [connect_success] => 1
    [connect_failure] => 0
    [connection_reused] => 0
)
```

```
[reconnect] => 0
[pconnect_success] => 0
[active_connections] => 1
[active_persistent_connections] => 0
[explicit_close] => 0
[implicit_close] => 0
[disconnect_close] => 0
[in_middle_of_command_close] => 0
[explicit_free_result] => 0
[implicit_free_result] => 0
[explicit_stmt_close] => 0
[implicit_stmt_close] => 0
[mem_emalloc_count] => 0
[mem_emalloc_ammount] => 0
[mem_ecalloc_count] => 0
[mem_ecalloc_ammount] => 0
[mem_erealloc_count] => 0
[mem_erealloc_ammount] => 0
[mem_efree_count] => 0
[mem_malloc_count] => 0
[mem_malloc_ammount] => 0
[mem_calloc_count] => 0
[mem_calloc_ammount] => 0
[mem_realloc_count] => 0
[mem_realloc_ammount] => 0
[mem_free_count] => 0
[proto_text_fetched_null] => 0
[proto_text_fetched_bit] => 0
[proto_text_fetched_tinyint] => 0
[proto_text_fetched_short] => 0
[proto_text_fetched_int24] => 0
[proto_text_fetched_int] => 0
[proto_text_fetched_bigint] => 0
[proto_text_fetched_decimal] => 0
[proto_text_fetched_float] => 0
[proto_text_fetched_double] => 0
[proto_text_fetched_date] => 0
[proto_text_fetched_year] => 0
[proto_text_fetched_time] => 0
[proto_text_fetched_datetime] => 0
[proto_text_fetched_timestamp] => 0
[proto_text_fetched_string] => 0
[proto_text_fetched_blob] => 0
[proto_text_fetched_enum] => 0
[proto_text_fetched_set] => 0
[proto_text_fetched_geometry] => 0
[proto_text_fetched_other] => 0
[proto_binary_fetched_null] => 0
[proto_binary_fetched_bit] => 0
[proto_binary_fetched_tinyint] => 0
[proto_binary_fetched_short] => 0
[proto_binary_fetched_int24] => 0
[proto_binary_fetched_int] => 0
[proto_binary_fetched_bigint] => 0
[proto_binary_fetched_decimal] => 0
[proto_binary_fetched_float] => 0
[proto_binary_fetched_double] => 0
[proto_binary_fetched_date] => 0
[proto_binary_fetched_year] => 0
[proto_binary_fetched_time] => 0
[proto_binary_fetched_datetime] => 0
[proto_binary_fetched_timestamp] => 0
[proto_binary_fetched_string] => 0
[proto_binary_fetched_blob] => 0
[proto_binary_fetched_enum] => 0
[proto_binary_fetched_set] => 0
[proto_binary_fetched_geometry] => 0
[proto_binary_fetched_other] => 0
)
```



**See Also**[Stats description](#)**21.9.3.9.23. [mysqli::\\$host\\_info](#), [mysqli\\_get\\_host\\_info](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$host\\_info](#)

[mysqli\\_get\\_host\\_info](#)

Returns a string representing the type of connection used

**Description**

Object oriented style

```
string  
mysqli->host_info ;
```

Procedural style

```
string mysqli_get_host_info(  
mysqli link);
```

Returns a string describing the connection represented by the [link](#) parameter (including the server host name).**Parameters**

[link](#) Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

**Return Values**

A character string representing the server hostname and the connection type.

**Examples****Example 21.126. [\\$mysqli->host\\_info](#) example**

Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* print host information */  
printf("Host info: %s\n", $mysqli->host_info);  
  
/* close connection */  
$mysqli->close();  
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* print host information */
printf("Host info: %s\n", mysqli_get_host_info($link));

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Host info: Localhost via UNIX socket
```

### See Also

[mysqli\\_get\\_proto\\_info](#)

### 21.9.3.9.24. [mysqli::\\$protocol\\_version](#), [mysqli\\_get\\_proto\\_info](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$protocol\\_version](#)

[mysqli\\_get\\_proto\\_info](#)

Returns the version of the MySQL protocol used

### Description

Object oriented style

```
string
mysqli->protocol_version ;
```

Procedural style

```
int mysqli_get_proto_info(
    mysqli link);
```

Returns an integer representing the MySQL protocol version used by the connection represented by the *link* parameter.

### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

### Return Values

Returns an integer representing the protocol version.

### Examples

**Example 21.127. `$mysqli->protocol_version` example**

## Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* print protocol version */
printf("Protocol version: %d\n", $mysqli->protocol_version);

/* close connection */
$mysqli->close();
?>
```

## Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* print protocol version */
printf("Protocol version: %d\n", mysqli_get_proto_info($link));

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Protocol version: 10
```

**See Also**

[mysqli\\_get\\_host\\_info](#)

**21.9.3.9.25. `mysqli::$server_info`, `mysqli_get_server_info`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$server\\_info](#)

[mysqli\\_get\\_server\\_info](#)

Returns the version of the MySQL server

**Description**

Object oriented style

```
string  
mysqli->server_info ;
```

### Procedural style

```
string mysqli_get_server_info(  
    mysqli link);
```

Returns a string representing the version of the MySQL server that the MySQLi extension is connected to.

### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

### Return Values

A character string representing the server version.

### Examples

#### Example 21.128. `$mysqli->server_info` example

##### Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* print server version */  
printf("Server version: %s\n", $mysqli->server_info);  
  
/* close connection */  
$mysqli->close();  
?>
```

##### Procedural style

```
<?php  
$link = mysqli_connect("localhost", "my_user", "my_password");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* print server version */  
printf("Server version: %s\n", mysqli_get_server_info($link));  
  
/* close connection */  
mysqli_close($link);  
?>
```

The above examples will output:

```
Server version: 4.1.2-alpha-debug
```

### See Also

```
mysqli_get_client_info  
mysqli_get_client_version  
mysqli_get_server_version
```

#### 21.9.3.9.26. `mysqli::$server_version`, `mysqli_get_server_version`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$server_version`  
`mysqli_get_server_version`

Returns the version of the MySQL server as an integer

### Description

Object oriented style

```
int  
mysqli->server_version ;
```

Procedural style

```
int mysqli_get_server_version(  
mysqli link);
```

The `mysqli_get_server_version` function returns the version of the server connected to (represented by the *link* parameter) as an integer.

### Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

### Return Values

An integer representing the server version.

The form of this version number is `main_version * 10000 + minor_version * 100 + sub_version` (i.e. version 4.1.0 is 40100).

### Examples

#### Example 21.129. `$mysqli->server_version` example

Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* print server version */
```

```
printf("Server version: %d\n", $mysqli->server_version);

/* close connection */
$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* print server version */
printf("Server version: %d\n", mysqli_get_server_version($link));

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Server version: 40102
```

### See Also

[mysqli\\_get\\_client\\_info](#)  
[mysqli\\_get\\_client\\_version](#)  
[mysqli\\_get\\_server\\_info](#)

### 21.9.3.9.27. [mysqli::get\\_warnings](#), [mysqli\\_get\\_warnings](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::get\\_warnings](#)

[mysqli\\_get\\_warnings](#)

Get result of SHOW WARNINGS

### Description

Object oriented style

```
mysqli_warning mysqli::get_warnings();
```

Procedural style

```
mysqli_warning mysqli_get_warnings(
    mysqli link);
```



### Warning

This function is currently not documented; only its argument list is available.

**21.9.3.9.28. *mysqli::\$info*, *mysqli\_info***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli::\$info*

*mysqli\_info*

Retrieves information about the most recently executed query

**Description**

Object oriented style

```
string
mysqli->info ;
```

Procedural style

```
string mysqli_info(
    mysqli link);
```

The *mysqli\_info* function returns a string providing information about the last query executed. The nature of this string is provided below:

**Table 21.46. Possible *mysqli\_info* return values**

Query type	Example result string
INSERT INTO...SELECT...	Records: 100 Duplicates: 0 Warnings: 0
INSERT INTO...VALUES (...),(...),(...)	Records: 3 Duplicates: 0 Warnings: 0
LOAD DATA INFILE ...	Records: 1 Deleted: 0 Skipped: 0 Warnings: 0
ALTER TABLE ...	Records: 3 Duplicates: 0 Warnings: 0
UPDATE ...	Rows matched: 40 Changed: 40 Warnings: 0

**Note**

Queries which do not fall into one of the preceding formats are not supported. In these situations, *mysqli\_info* will return an empty string.

**Parameters**

*link*

Procedural style only: A link identifier returned by *mysqli\_connect* or *mysqli\_init*

**Return Values**

A character string representing additional information about the most recently executed query.

**Examples****Example 21.130. *\$mysqli->info* example**

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
```

```
printf("Connect failed: %s\n", mysqli_connect_error());
exit();
}

$mysqli->query("CREATE TEMPORARY TABLE t1 LIKE City");

/* INSERT INTO .. SELECT */
$mysqli->query("INSERT INTO t1 SELECT * FROM City ORDER BY ID LIMIT 150");
printf("%s\n", $mysqli->info);

/* close connection */
$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TEMPORARY TABLE t1 LIKE City");

/* INSERT INTO .. SELECT */
mysqli_query($link, "INSERT INTO t1 SELECT * FROM City ORDER BY ID LIMIT 150");
printf("%s\n", mysqli_info($link));

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Records: 150 Duplicates: 0 Warnings: 0
```

### See Also

[mysqli\\_affected\\_rows](#)  
[mysqli\\_warning\\_count](#)  
[mysqli\\_num\\_rows](#)

### 21.9.3.9.29. [mysqli::init](#), [mysqli\\_init](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::init](#)

[mysqli\\_init](#)

Initializes MySQLi and returns a resource for use with [mysqli\\_real\\_connect\(\)](#)

### Description

Object oriented style

```
mysqli mysqli::init();
```



## Procedural style

```
mysqli mysqli_init();
```

Allocates or initializes a MySQL object suitable for [mysqli\\_options](#) and [mysqli\\_real\\_connect](#).



### Note

Any subsequent calls to any mysqli function (except [mysqli\\_options](#)) will fail until [mysqli\\_real\\_connect](#) was called.

## Return Values

Returns an object.

## Examples

See [mysqli\\_real\\_connect](#).

## See Also

```
mysqli_options  
mysqli_close  
mysqli_real_connect  
mysqli_connect
```

### 21.9.3.9.30. [mysqli::\\$insert\\_id](#), [mysqli\\_insert\\_id](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::\\$insert\\_id](#)

```
mysqli_insert_id
```

Returns the auto generated id used in the last query

## Description

Object oriented style

```
mixed  
mysqli->insert_id ;
```

Procedural style

```
mixed mysqli_insert_id(  
mysqli link);
```

The [mysqli\\_insert\\_id](#) function returns the ID generated by a query on a table with a column having the AUTO\_INCREMENT attribute. If the last query wasn't an INSERT or UPDATE statement or if the modified table does not have a column with the AUTO\_INCREMENT attribute, this function will return zero.



### Note

Performing an INSERT or UPDATE statement using the [LAST\\_INSERT\\_ID\(\)](#) function will also modify the value returned by the [mysqli\\_insert\\_id](#) function.

## Parameters

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

The value of the [AUTO\\_INCREMENT](#) field that was updated by the previous query. Returns zero if there was no previous query on the connection or if the query did not update an [AUTO\\_INCREMENT](#) value.



### Note

If the number is greater than maximal int value, `mysqli_insert_id` will return a string.

## Examples

### Example 21.131. `$mysqli->insert_id` example

#### Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCity LIKE City");

$query = "INSERT INTO myCity VALUES (NULL, 'Stuttgart', 'DEU', 'Stuttgart', 617000)";
$mysqli->query($query);

printf ("New Record has id %d.\n", $mysqli->insert_id);

/* drop table */
$mysqli->query("DROP TABLE myCity");

/* close connection */
$mysqli->close();
?>
```

#### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCity LIKE City");

$query = "INSERT INTO myCity VALUES (NULL, 'Stuttgart', 'DEU', 'Stuttgart', 617000)";
mysqli_query($link, $query);

printf ("New Record has id %d.\n", mysqli_insert_id($link));

/* drop table */
mysqli_query($link, "DROP TABLE myCity");

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
New Record has id 1.
```

### 21.9.3.9.31. [mysqli::kill](#), [mysqli\\_kill](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::kill](#)

[mysqli\\_kill](#)

Asks the server to kill a MySQL thread

#### Description

Object oriented style

```
bool mysqli::kill(  
    int processid);
```

Procedural style

```
bool mysqli_kill(  
    mysqli link,  
    int processid);
```

This function is used to ask the server to kill a MySQL thread specified by the *processid* parameter. This value must be retrieved by calling the [mysqli\\_thread\\_id](#) function.

To stop a running query you should use the SQL command `KILL QUERY processid`.

#### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

#### Return Values

Returns `TRUE` on success or `FALSE` on failure.

#### Examples

##### Example 21.132. [mysqli::kill](#) example

Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* determine our thread id */  
$thread_id = $mysqli->thread_id;  
  
/* Kill connection */
```

```
$mysqli->kill($thread_id);

/* This should produce an error */
if (!$mysqli->query("CREATE TABLE myCity LIKE City")) {
    printf("Error: %s\n", $mysqli->error);
    exit;
}

/* close connection */
$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* determine our thread id */
$thread_id = mysqli_thread_id($link);

/* Kill connection */
mysqli_kill($link, $thread_id);

/* This should produce an error */
if (!mysqli_query($link, "CREATE TABLE myCity LIKE City")) {
    printf("Error: %s\n", mysqli_error($link));
    exit;
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Error: MySQL server has gone away
```

### See Also

[mysqli\\_thread\\_id](#)

### 21.9.3.9.32. [mysqli::more\\_results](#), [mysqli\\_more\\_results](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::more\\_results](#)

[mysqli\\_more\\_results](#)

Check if there are any more query results from a multi query

### Description

Object oriented style

```
bool mysqli::more_results();
```

Procedural style

```
bool mysqli_more_results(  
    mysqli link);
```

Indicates if one or more result sets are available from a previous call to [mysqli\\_multi\\_query](#).

#### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

#### Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

#### Examples

See [mysqli\\_multi\\_query](#).

#### See Also

[mysqli\\_multi\\_query](#)  
[mysqli\\_next\\_result](#)  
[mysqli\\_store\\_result](#)  
[mysqli\\_use\\_result](#)

### 21.9.3.9.33. [mysqli::multi\\_query](#), [mysqli\\_multi\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::multi\\_query](#)

[mysqli\\_multi\\_query](#)

Performs a query on the database

#### Description

Object oriented style

```
bool mysqli::multi_query(  
    string query);
```

Procedural style

```
bool mysqli_multi_query(  
    mysqli link,  
    string query);
```

Executes one or multiple queries which are concatenated by a semicolon.

To retrieve the resultset from the first query you can use [mysqli\\_use\\_result](#) or [mysqli\\_store\\_result](#). All subsequent query results can be processed using [mysqli\\_more\\_results](#) and [mysqli\\_next\\_result](#).

#### Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

*query*

The query, as a string.

Data inside the query should be [properly escaped](#).

## Return Values

Returns [FALSE](#) if the first statement failed. To retrieve subsequent errors from other statements you have to call [mysqli\\_next\\_result](#) first.

## Examples

### Example 21.133. [mysqli::multi\\_query](#) example

#### Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT CURRENT_USER()";
$query .= "SELECT Name FROM City ORDER BY ID LIMIT 20, 5";

/* execute multi query */
if ($mysqli->multi_query($query)) {
    do {
        /* store first result set */
        if ($result = $mysqli->store_result()) {
            while ($row = $result->fetch_row()) {
                printf("%s\n", $row[0]);
            }
            $result->free();
        }
        /* print divider */
        if ($mysqli->more_results()) {
            printf("-----\n");
        }
    } while ($mysqli->next_result());
}

/* close connection */
$mysqli->close();
?>
```

#### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT CURRENT_USER()";
$query .= "SELECT Name FROM City ORDER BY ID LIMIT 20, 5";

/* execute multi query */
if (mysqli_multi_query($link, $query)) {
    do {
```

```
/* store first result set */
if ($result = mysqli_store_result($link)) {
    while ($row = mysqli_fetch_row($result)) {
        printf("%s\n", $row[0]);
    }
    mysqli_free_result($result);
}
/* print divider */
if (mysqli_more_results($link)) {
    printf("-----\n");
}
} while (mysqli_next_result($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output something similar to:

```
my_user@localhost
-----
Amersfoort
Maastricht
Dordrecht
Leiden
Haarlemmermeer
```

### See Also

[mysqli\\_query](#)  
[mysqli\\_use\\_result](#)  
[mysqli\\_store\\_result](#)  
[mysqli\\_next\\_result](#)  
[mysqli\\_more\\_results](#)

#### 21.9.3.9.34. [mysqli::next\\_result](#), [mysqli\\_next\\_result](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::next\\_result](#)

[mysqli\\_next\\_result](#)

Prepare next result from `multi_query`

### Description

Object oriented style

```
bool mysqli::next_result();
```

Procedural style

```
bool mysqli_next_result(
    mysqli link);
```

Prepares next result set from a previous call to [mysqli\\_multi\\_query](#) which can be retrieved by [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#).

### Parameters

*link*Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`**Return Values**Returns `TRUE` on success or `FALSE` on failure.**Examples**See `mysqli_multi_query`.**See Also**

`mysqli_multi_query`  
`mysqli_more_results`  
`mysqli_store_result`  
`mysqli_use_result`

**21.9.3.9.35. `mysqli::options`, `mysqli_options`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::options`

`mysqli_options`

Set options

**Description**

Object oriented style

```
bool mysqli::options(
    int option,
    mixed value);
```

Procedural style

```
bool mysqli_options(
    mysqli link,
    int option,
    mixed value);
```

Used to set extra connect options and affect behavior for a connection.

This function may be called multiple times to set several options.

`mysqli_options` should be called after `mysqli_init` and before `mysqli_real_connect`.**Parameters***link*Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`*option*

The option that you want to set. It can be one of the following values:

**Table 21.47. Valid options**

Name	Description
<code>MYSQLI_OPT_CONNECT_TIMEOUT</code>	connection timeout in seconds (supported on Windows with TCP/IP since PHP 5.3.1)



Name	Description
<code>MYSQLI_OPT_LOCAL_INFILE</code>	enable/disable use of <code>LOAD LOCAL INFILE</code>
<code>MYSQLI_INIT_COMMAND</code>	command to execute after when connecting to MySQL server
<code>MYSQLI_READ_DEFAULT_FILE</code>	Read options from named option file instead of <code>my.cnf</code>
<code>MYSQLI_READ_DEFAULT_GROUP</code>	Read options from the named group from <code>my.cnf</code> or the file specified with <code>MYSQLI_READ_DEFAULT_FILE</code> .
<code>MYSQLI_SERVER_PUBLIC_KEY</code>	RSA public key file used with the SHA-256 based authentication.

*value* The value for the option.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Changelog

Version	Description
5.5.0	The <code>MYSQLI_SERVER_PUBLIC_KEY</code> option was added.

### Examples

See `mysqli_real_connect`.

### Notes



#### Note

MySQLnd always assumes the server default charset. This charset is sent during connection hand-shake/authentication, which mysqlnd will use.

Libmysqlclient uses the default charset set in the `my.cnf` or by an explicit call to `mysqli_options` prior to calling `mysqli_real_connect`, but after `mysqli_init`.

### See Also

`mysqli_init`  
`mysqli_real_connect`

#### 21.9.3.9.36. `mysqli::ping`, `mysqli_ping`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::ping`

`mysqli_ping`

Pings a server connection, or tries to reconnect if the connection has gone down

### Description

## Object oriented style

```
bool mysqli::ping();
```

## Procedural style

```
bool mysqli_ping(  
    mysqli link);
```

Checks whether the connection to the server is working. If it has gone down, and global option [mysqli.reconnect](#) is enabled an automatic reconnection is attempted.

This function can be used by clients that remain idle for a long while, to check whether the server has closed the connection and reconnect if necessary.

## Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

## Examples

### Example 21.134. [mysqli::ping](#) example

#### Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* check if server is alive */  
if ($mysqli->ping()) {  
    printf ("Our connection is ok!\n");  
} else {  
    printf ("Error: %s\n", $mysqli->error);  
}  
  
/* close connection */  
$mysqli->close();  
?>
```

#### Procedural style

```
<?php  
$link = mysqli_connect("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* check if server is alive */  
if (mysqli_ping($link)) {
```

```
    printf ("Our connection is ok!\n");
} else {
    printf ("Error: %s\n", mysqli_error($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Our connection is ok!
```

### 21.9.3.9.37. [mysqli::poll](#), [mysqli\\_poll](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::poll](#)

[mysqli\\_poll](#)

Poll connections

#### Description

Object oriented style

```
public int mysqli::poll(
    array read,
    array error,
    array reject,
    int sec,
    int usec);
```

Procedural style

```
int mysqli_poll(
    array read,
    array error,
    array reject,
    int sec,
    int usec);
```



#### Warning

This function is currently not documented; only its argument list is available.

Poll connections. Available only with [mysqlind](#). The method can be used as [static](#).

#### Parameters

*read*

*error*

*reject*

*sec*

Number of seconds to wait, must be non-negative.

*usec*

Number of microseconds to wait, must be non-negative.

## Return Values

Returns number of ready connections upon success, [FALSE](#) otherwise.

## Examples

### Example 21.135. A `mysqli_poll` example

```
<?php
$link1 = mysqli_connect();
$link1->query("SELECT 'test'", MYSQLI_ASYNC);
$all_links = array($link1);
$processed = 0;
do {
    $links = $errors = $reject = array();
    foreach ($all_links as $link) {
        $links[] = $errors[] = $reject[] = $link;
    }
    if (!mysqli_poll($links, $errors, $reject, 1)) {
        continue;
    }
    foreach ($links as $link) {
        if ($result = $link->reap_async_query()) {
            print_r($result->fetch_row());
            if (is_object($result))
                mysqli_free_result($result);
        } else die(sprintf("MySQLi Error: %s", mysqli_error($link)));
        $processed++;
    }
} while ($processed < count($all_links));
?>
```

The above example will output:

```
Array
(
    [0] => test
)
```

## See Also

[mysqli\\_query](#)  
[mysqli\\_reap\\_async\\_query](#)

### 21.9.3.9.38. `mysqli::prepare`, `mysqli_prepare`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::prepare](#)

[mysqli\\_prepare](#)

Prepare an SQL statement for execution

## Description

Object oriented style

```
mysqli_stmt mysqli::prepare(
    string query);
```

## Procedural style

```
mysqli_stmt mysqli_prepare(  
    mysqli link,  
    string query);
```

Prepares the SQL query, and returns a statement handle to be used for further operations on the statement. The query must consist of a single SQL statement.

The parameter markers must be bound to application variables using [mysqli\\_stmt\\_bind\\_param](#) and/or [mysqli\\_stmt\\_bind\\_result](#) before executing the statement or fetching rows.

### Parameters

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

*query*

The query, as a string.



#### Note

You should not add a terminating semicolon or `\g` to the statement.

This parameter can include one or more parameter markers in the SQL statement by embedding question mark (?) characters at the appropriate positions.



#### Note

The markers are legal only in certain places in SQL statements. For example, they are allowed in the `VALUES ( )` list of an `INSERT` statement (to specify column values for a row), or in a comparison with a column in a `WHERE` clause to specify a comparison value.

However, they are not allowed for identifiers (such as table or column names), in the select list that names the columns to be returned by a `SELECT` statement, or to specify both operands of a binary operator such as the `=` equal sign. The latter restriction is necessary because it would be impossible to determine the parameter type. It's not allowed to compare marker with `NULL` by `? IS NULL` too. In general, parameters are legal only in Data Manipulation Language (DML) statements, and not in Data Definition Language (DDL) statements.

### Return Values

[mysqli\\_prepare](#) returns a statement object or `FALSE` if an error occurred.

### Examples

#### Example 21.136. `mysqli::prepare` example

Object oriented style

```

<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$city = "Amersfoort";

/* create a prepared statement */
if ($stmt = $mysqli->prepare("SELECT District FROM City WHERE Name=?")) {

    /* bind parameters for markers */
    $stmt->bind_param("s", $city);

    /* execute query */
    $stmt->execute();

    /* bind result variables */
    $stmt->bind_result($district);

    /* fetch value */
    $stmt->fetch();

    printf("%s is in district %s\n", $city, $district);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>

```

### Procedural style

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$city = "Amersfoort";

/* create a prepared statement */
if ($stmt = mysqli_prepare($link, "SELECT District FROM City WHERE Name=?")) {

    /* bind parameters for markers */
    mysqli_stmt_bind_param($stmt, "s", $city);

    /* execute query */
    mysqli_stmt_execute($stmt);

    /* bind result variables */
    mysqli_stmt_bind_result($stmt, $district);

    /* fetch value */
    mysqli_stmt_fetch($stmt);

    printf("%s is in district %s\n", $city, $district);

    /* close statement */
    mysqli_stmt_close($stmt);
}

```

```
}  
  
/* close connection */  
mysqli_close($link);  
?>
```

The above examples will output:

```
Amersfoort is in district Utrecht
```

### See Also

```
mysqli_stmt_execute  
mysqli_stmt_fetch  
mysqli_stmt_bind_param  
mysqli_stmt_bind_result  
mysqli_stmt_close
```

### 21.9.3.9.39. [mysqli::query](#), [mysqli\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::query](#)

[mysqli\\_query](#)

Performs a query on the database

### Description

Object oriented style

```
mixed mysqli::query(  
    string query,  
    int resultmode  
    = MYSQLI_STORE_RESULT);
```

Procedural style

```
mixed mysqli_query(  
    mysqli link,  
    string query,  
    int resultmode  
    = MYSQLI_STORE_RESULT);
```

Performs a *query* against the database.

For non-DML queries (not INSERT, UPDATE or DELETE), this function is similar to calling [mysqli\\_real\\_query](#) followed by either [mysqli\\_use\\_result](#) or [mysqli\\_store\\_result](#).



### Note

In the case where you pass a statement to [mysqli\\_query](#) that is longer than [max\\_allowed\\_packet](#) of the server, the returned error codes are different depending on whether you are using MySQL Native Driver ([mysqlnd](#)) or MySQL Client Library ([libmysqlclient](#)). The behavior is as follows:

- [mysqlnd](#) on Linux returns an error code of 1153. The error message means “got a packet bigger than [max\\_allowed\\_packet](#) bytes”.

- [mysqlnd](#) on Windows returns an error code 2006. This error message means “server has gone away”.
- [libmysqlclient](#) on all platforms returns an error code 2006. This error message means “server has gone away”.

## Parameters

<i>link</i>	Procedural style only: A link identifier returned by <a href="#">mysqli_connect</a> or <a href="#">mysqli_init</a>
<i>query</i>	The query string.  Data inside the query should be <a href="#">properly escaped</a> .
<i>resultmode</i>	Either the constant <a href="#">MYSQLI_USE_RESULT</a> or <a href="#">MYSQLI_STORE_RESULT</a> depending on the desired behavior. By default, <a href="#">MYSQLI_STORE_RESULT</a> is used.  If you use <a href="#">MYSQLI_USE_RESULT</a> all subsequent calls will return error <a href="#">Commands out of sync</a> unless you call <a href="#">mysqli_free_result</a>  With <a href="#">MYSQLI_ASYNC</a> (available with <a href="#">mysqlnd</a> ), it is possible to perform query asynchronously. <a href="#">mysqli_poll</a> is then used to get results from such queries.

## Return Values

Returns [FALSE](#) on failure. For successful [SELECT](#), [SHOW](#), [DESCRIBE](#) or [EXPLAIN](#) queries [mysqli\\_query](#) will return a [mysqli\\_result](#) object. For other successful queries [mysqli\\_query](#) will return [TRUE](#) .

## Changelog

Version	Description
5.3.0	Added the ability of async queries.

## Examples

### Example 21.137. [mysqli::query](#) example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if ($mysqli->connect_errno) {
    printf("Connect failed: %s\n", $mysqli->connect_error);
    exit();
}

/* Create table doesn't return a resultset */
if ($mysqli->query("CREATE TEMPORARY TABLE myCity LIKE City") === TRUE) {
    printf("Table myCity successfully created.\n");
}

/* Select queries return a resultset */
if ($result = $mysqli->query("SELECT Name FROM City LIMIT 10")) {
    printf("Select returned %d rows.\n", $result->num_rows);
}
```



```

    /* free result set */
    $result->close();
}

/* If we have to retrieve large amount of data we use MYSQLI_USE_RESULT */
if ($result = $mysqli->query("SELECT * FROM City", MYSQLI_USE_RESULT)) {

    /* Note, that we can't execute any functions which interact with the
       server until result set was closed. All calls will return an
       'out of sync' error */
    if (!$mysqli->query("SET @a='this will not work'")) {
        printf("Error: %s\n", $mysqli->error);
    }
    $result->close();
}

$mysqli->close();
?>

```

### Procedural style

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Create table doesn't return a resultset */
if (mysqli_query($link, "CREATE TEMPORARY TABLE myCity LIKE City") === TRUE) {
    printf("Table myCity successfully created.\n");
}

/* Select queries return a resultset */
if ($result = mysqli_query($link, "SELECT Name FROM City LIMIT 10")) {
    printf("Select returned %d rows.\n", mysqli_num_rows($result));

    /* free result set */
    mysqli_free_result($result);
}

/* If we have to retrieve large amount of data we use MYSQLI_USE_RESULT */
if ($result = mysqli_query($link, "SELECT * FROM City", MYSQLI_USE_RESULT)) {

    /* Note, that we can't execute any functions which interact with the
       server until result set was closed. All calls will return an
       'out of sync' error */
    if (!mysqli_query($link, "SET @a='this will not work'")) {
        printf("Error: %s\n", mysqli_error($link));
    }
    mysqli_free_result($result);
}

mysqli_close($link);
?>

```

The above examples will output:

```

Table myCity successfully created.
Select returned 10 rows.
Error: Commands out of sync; You can't run this command now

```

---

**See Also**

[mysqli\\_real\\_query](#)  
[mysqli\\_multi\\_query](#)  
[mysqli\\_free\\_result](#)

**21.9.3.9.40. [mysqli::real\\_connect](#), [mysqli\\_real\\_connect](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::real\\_connect](#)

[mysqli\\_real\\_connect](#)

Opens a connection to a mysql server

**Description**

Object oriented style

```
bool mysqli::real_connect(  
    string host,  
    string username,  
    string passwd,  
    string dbname,  
    int port,  
    string socket,  
    int flags);
```

Procedural style

```
bool mysqli_real_connect(  
    mysqli link,  
    string host,  
    string username,  
    string passwd,  
    string dbname,  
    int port,  
    string socket,  
    int flags);
```

Establish a connection to a MySQL database engine.

This function differs from [mysqli\\_connect](#):

- [mysqli\\_real\\_connect](#) needs a valid object which has to be created by function [mysqli\\_init](#).
- With the [mysqli\\_options](#) function you can set various options for connection.
- There is a *flags* parameter.

**Parameters**

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

*host*

Can be either a host name or an IP address. Passing the [NULL](#) value or the string "localhost" to this parameter, the local host is assumed. When possible, pipes will be used instead of the TCP/IP protocol.

*username*

The MySQL user name.

*passwd*

If provided or [NULL](#), the MySQL server will attempt to authenticate the user against those user records which have no password only.

This allows one username to be used with different permissions (depending on if a password as provided or not).

*dbname*

If provided will specify the default database to be used when performing queries.

*port*

Specifies the port number to attempt to connect to the MySQL server.

*socket*

Specifies the socket or named pipe that should be used.



#### Note

Specifying the *socket* parameter will not explicitly determine the type of connection to be used when connecting to the MySQL server. How the connection is made to the MySQL database is determined by the *host* parameter.

*flags*

With the parameter *flags* you can set different connection options:

**Table 21.48. Supported flags**

Name	Description
<code>MYSQLI_CLIENT_COMPRESS</code>	Use compression protocol
<code>MYSQLI_CLIENT_FOUND_ROWS</code>	return number of matched rows, not the number of affected rows
<code>MYSQLI_CLIENT_IGNORE_SPACE</code>	Allow spaces after function names. Makes all function names reserved words.
<code>MYSQLI_CLIENT_INTERACTIVE</code>	Allow <code>interactive_timeout</code> seconds (instead of <code>wait_timeout</code> seconds) of inactivity before closing the connection
<code>MYSQLI_CLIENT_SSL</code>	Use SSL (encryption)



#### Note

For security reasons the `MULTI_STATEMENT` flag is not supported in PHP. If you want to execute multiple queries use the `mysqli_multi_query` function.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Examples

#### Example 21.138. `mysqli::real_connect` example

Object oriented style

```
<?php
$mysqli = mysqli_init();
if (!$mysqli) {
```

```
die('mysqli_init failed');
}

if (!$mysqli->options(MYSQLI_INIT_COMMAND, 'SET AUTOCOMMIT = 0')) {
    die('Setting MYSQLI_INIT_COMMAND failed');
}

if (!$mysqli->options(MYSQLI_OPT_CONNECT_TIMEOUT, 5)) {
    die('Setting MYSQLI_OPT_CONNECT_TIMEOUT failed');
}

if (!$mysqli->real_connect('localhost', 'my_user', 'my_password', 'my_db')) {
    die('Connect Error (' . mysqli_connect_errno() . ') '
        . mysqli_connect_error());
}

echo 'Success... ' . $mysqli->host_info . "\n";

$mysqli->close();
?>
```

### Object oriented style when extending mysqli class

```
<?php

class foo_mysqli extends mysqli {
    public function __construct($host, $user, $pass, $db) {
        parent::init();

        if (!$parent::options(MYSQLI_INIT_COMMAND, 'SET AUTOCOMMIT = 0')) {
            die('Setting MYSQLI_INIT_COMMAND failed');
        }

        if (!$parent::options(MYSQLI_OPT_CONNECT_TIMEOUT, 5)) {
            die('Setting MYSQLI_OPT_CONNECT_TIMEOUT failed');
        }

        if (!$parent::real_connect($host, $user, $pass, $db)) {
            die('Connect Error (' . mysqli_connect_errno() . ') '
                . mysqli_connect_error());
        }
    }
}

$db = new foo_mysqli('localhost', 'my_user', 'my_password', 'my_db');

echo 'Success... ' . $db->host_info . "\n";

$db->close();
?>
```

### Procedural style

```
<?php

$link = mysqli_init();
if (!$link) {
    die('mysqli_init failed');
}

if (!mysqli_options($link, MYSQLI_INIT_COMMAND, 'SET AUTOCOMMIT = 0')) {
    die('Setting MYSQLI_INIT_COMMAND failed');
}

if (!mysqli_options($link, MYSQLI_OPT_CONNECT_TIMEOUT, 5)) {
    die('Setting MYSQLI_OPT_CONNECT_TIMEOUT failed');
}
```

```

}

if (!mysqli_real_connect($link, 'localhost', 'my_user', 'my_password', 'my_db')) {
    die('Connect Error (' . mysqli_connect_errno() . ') '
        . mysqli_connect_error());
}

echo 'Success... ' . mysqli_get_host_info($link) . "\n";

mysqli_close($link);
?>

```

The above examples will output:

```
Success... MySQL host info: localhost via TCP/IP
```

## Notes



### Note

MySQLnd always assumes the server default charset. This charset is sent during connection hand-shake/authentication, which mysqlnd will use.

Libmysqlclient uses the default charset set in the `my.cnf` or by an explicit call to `mysqli_options` prior to calling `mysqli_real_connect`, but after `mysqli_init`.

## See Also

```

mysqli_connect
mysqli_init
mysqli_options
mysqli_ssl_set
mysqli_close

```

### 21.9.3.9.41. `mysqli::real_escape_string`, `mysqli_real_escape_string`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::real_escape_string`

```
mysqli_real_escape_string
```

Escapes special characters in a string for use in an SQL statement, taking into account the current charset of the connection

## Description

Object oriented style

```
string mysqli::escape_string(
    string escapestr);
```

```
string mysqli::real_escape_string(
    string escapestr);
```

Procedural style

```
string mysqli_real_escape_string(
    mysqli link,
```

```
string escapestr);
```

This function is used to create a legal SQL string that you can use in an SQL statement. The given string is encoded to an escaped SQL string, taking into account the current character set of the connection.



### Security: the default character set

The character set must be set either at the server level, or with the API function [mysqli\\_set\\_charset](#) for it to affect [mysqli\\_real\\_escape\\_string](#). See the concepts section on [character sets](#) for more information.

### Parameters

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

*escapestr*

The string to be escaped.

Characters encoded are NUL (ASCII 0), \n, \r, \, ', ", and Control-Z.

### Return Values

Returns an escaped string.

### Examples

#### Example 21.139. [mysqli::real\\_escape\\_string](#) example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TEMPORARY TABLE myCity LIKE City");

$city = "'s Hertogenbosch";

/* this query will fail, cause we didn't escape $city */
if (!$mysqli->query("INSERT into myCity (Name) VALUES ('$city')")) {
    printf("Error: %s\n", $mysqli->sqlstate);
}

$city = $mysqli->real_escape_string($city);

/* this query with escaped $city will work */
if ($mysqli->query("INSERT into myCity (Name) VALUES ('$city')")) {
    printf("%d Row inserted.\n", $mysqli->affected_rows);
}

$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TEMPORARY TABLE myCity LIKE City");

$city = "'s Hertogenbosch";

/* this query will fail, cause we didn't escape $city */
if (!mysqli_query($link, "INSERT into myCity (Name) VALUES ('$city')")) {
    printf("Error: %s\n", mysqli_sqlstate($link));
}

$city = mysqli_real_escape_string($link, $city);

/* this query with escaped $city will work */
if (mysqli_query($link, "INSERT into myCity (Name) VALUES ('$city')")) {
    printf("%d Row inserted.\n", mysqli_affected_rows($link));
}

mysqli_close($link);
?>
```

The above examples will output:

```
Error: 42000
1 Row inserted.
```

## Notes



### Note

For those accustomed to using `mysql_real_escape_string`, note that the arguments of `mysqli_real_escape_string` differ from what `mysql_real_escape_string` expects. The `link` identifier comes first in `mysqli_real_escape_string`, whereas the string to be escaped comes first in `mysql_real_escape_string`.

## See Also

[mysqli\\_set\\_charset](#)  
[mysqli\\_character\\_set\\_name](#)

### 21.9.3.9.42. [mysqli::real\\_query](#), [mysqli\\_real\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::real\\_query](#)

[mysqli\\_real\\_query](#)

Execute an SQL query

## Description

Object oriented style

```
bool mysqli::real_query(  
    string query);
```

### Procedural style

```
bool mysqli_real_query(  
    mysqli link,  
    string query);
```

Executes a single query against the database whose result can then be retrieved or stored using the [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#) functions.

In order to determine if a given query should return a result set or not, see [mysqli\\_field\\_count](#).

### Parameters

<i>link</i>	Procedural style only: A link identifier returned by <a href="#">mysqli_connect</a> or <a href="#">mysqli_init</a>
<i>query</i>	The query, as a string.  Data inside the query should be <a href="#">properly escaped</a> .

### Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

### See Also

[mysqli\\_query](#)  
[mysqli\\_store\\_result](#)  
[mysqli\\_use\\_result](#)

#### 21.9.3.9.43. [mysqli::reap\\_async\\_query](#), [mysqli\\_reap\\_async\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::reap\\_async\\_query](#)

[mysqli\\_reap\\_async\\_query](#)

Get result from async query

### Description

Object oriented style

```
public mysqli_result mysqli::reap_async_query();
```

### Procedural style

```
mysqli_result mysqli_reap_async_query(  
    mysqli link);
```



#### Warning

This function is currently not documented; only its argument list is available.

Get result from async query. Available only with [mysqlnd](#).

### Parameters



*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

### Return Values

Returns `mysqli_result` in success, `FALSE` otherwise.

### See Also

`mysqli_poll`

#### 21.9.3.9.44. `mysqli::refresh`, `mysqli_refresh`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::refresh`

`mysqli_refresh`

Refreshes

### Description

Object oriented style

```
public bool mysqli::refresh(  
    int options);
```

Procedural style

```
int mysqli_refresh(  
    resource link,  
    int options);
```

Flushes tables or caches, or resets the replication server information.

### Parameters

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

*options* The options to refresh, using the `MYSQLI_REFRESH_*` constants as documented within the [MySQLi constants](#) documentation.

See also the official [MySQL Refresh](#) documentation.

### Return Values

`TRUE` if the refresh was a success, otherwise `FALSE`

### See Also

`mysqli_poll`

#### 21.9.3.9.45. `mysqli::rollback`, `mysqli_rollback`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::rollback`

`mysqli_rollback`

Rolls back current transaction

## Description

Object oriented style

```
bool mysqli::rollback();
```

Procedural style

```
bool mysqli_rollback(  
    mysqli link);
```

Rollbacks the current transaction for the database.

## Parameters

*link* Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

## Examples

### Example 21.140. [mysqli::rollback](#) example

Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
/* disable autocommit */  
$mysqli->autocommit(FALSE);  
  
$mysqli->query("CREATE TABLE myCity LIKE City");  
$mysqli->query("ALTER TABLE myCity Type=InnoDB");  
$mysqli->query("INSERT INTO myCity SELECT * FROM City LIMIT 50");  
  
/* commit insert */  
$mysqli->commit();  
  
/* delete all rows */  
$mysqli->query("DELETE FROM myCity");  
  
if ($result = $mysqli->query("SELECT COUNT(*) FROM myCity")) {  
    $row = $result->fetch_row();  
    printf("%d rows in table myCity.\n", $row[0]);  
    /* Free result */  
    $result->close();  
}  
  
/* Rollback */  
$mysqli->rollback();  
  
if ($result = $mysqli->query("SELECT COUNT(*) FROM myCity")) {  
    $row = $result->fetch_row();  
    printf("%d rows in table myCity (after rollback).\n", $row[0]);
```

```

    /* Free result */
    $result->close();
}

/* Drop table myCity */
mysqli->query("DROP TABLE myCity");

mysqli->close();
?>

```

### Procedural style

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* disable autocommit */
mysqli_autocommit($link, FALSE);

mysqli_query($link, "CREATE TABLE myCity LIKE City");
mysqli_query($link, "ALTER TABLE myCity Type=InnoDB");
mysqli_query($link, "INSERT INTO myCity SELECT * FROM City LIMIT 50");

/* commit insert */
mysqli_commit($link);

/* delete all rows */
mysqli_query($link, "DELETE FROM myCity");

if ($result = mysqli_query($link, "SELECT COUNT(*) FROM myCity")) {
    $row = mysqli_fetch_row($result);
    printf("%d rows in table myCity.\n", $row[0]);
    /* Free result */
    mysqli_free_result($result);
}

/* Rollback */
mysqli_rollback($link);

if ($result = mysqli_query($link, "SELECT COUNT(*) FROM myCity")) {
    $row = mysqli_fetch_row($result);
    printf("%d rows in table myCity (after rollback).\n", $row[0]);
    /* Free result */
    mysqli_free_result($result);
}

/* Drop table myCity */
mysqli_query($link, "DROP TABLE myCity");

mysqli_close($link);
?>

```

The above examples will output:

```

0 rows in table myCity.
50 rows in table myCity (after rollback).

```

### See Also

```
mysqli_commit  
mysqli_autocommit
```

#### 21.9.3.9.46. `mysqli::rpl_query_type`, `mysqli_rpl_query_type`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::rpl_query_type`

```
mysqli_rpl_query_type
```

Returns RPL query type

##### Description

Object oriented style

```
int mysqli::rpl_query_type(  
    string query);
```

Procedural style

```
int mysqli_rpl_query_type(  
    mysqli link,  
    string query);
```

Returns `MYSQLI_RPL_MASTER`, `MYSQLI_RPL_SLAVE` or `MYSQLI_RPL_ADMIN` depending on a query type. `INSERT`, `UPDATE` and similar are *master* queries, `SELECT` is *slave*, and `FLUSH`, `REPAIR` and similar are *admin*.



##### Warning

This function is currently not documented; only its argument list is available.



##### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

#### 21.9.3.9.47. `mysqli::select_db`, `mysqli_select_db`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::select_db`

```
mysqli_select_db
```

Selects the default database for database queries

##### Description

Object oriented style

```
bool mysqli::select_db(  
    string dbname);
```

Procedural style

```
bool mysqli_select_db(  
    mysqli link,  
    string dbname);
```

Selects the default database to be used when performing queries against the database connection.

**Note**

This function should only be used to change the default database for the connection. You can select the default database with 4th parameter in [mysqli\\_connect](#).

**Parameters**

<i>link</i>	Procedural style only: A link identifier returned by <a href="#">mysqli_connect</a> or <a href="#">mysqli_init</a>
<i>dbname</i>	The database name.

**Return Values**

Returns [TRUE](#) on success or [FALSE](#) on failure.

**Examples****Example 21.141. [mysqli::select\\_db](#) example**

## Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "test");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* return name of current default database */
if ($result = $mysqli->query("SELECT DATABASE())) {
    $row = $result->fetch_row();
    printf("Default database is %s.\n", $row[0]);
    $result->close();
}

/* change db to world db */
$mysqli->select_db("world");

/* return name of current default database */
if ($result = $mysqli->query("SELECT DATABASE())) {
    $row = $result->fetch_row();
    printf("Default database is %s.\n", $row[0]);
    $result->close();
}

$mysqli->close();
?>
```

## Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "test");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}
```

```
/* return name of current default database */
if ($result = mysqli_query($link, "SELECT DATABASE()")) {
    $row = mysqli_fetch_row($result);
    printf("Default database is %s.\n", $row[0]);
    mysqli_free_result($result);
}

/* change db to world db */
mysqli_select_db($link, "world");

/* return name of current default database */
if ($result = mysqli_query($link, "SELECT DATABASE()")) {
    $row = mysqli_fetch_row($result);
    printf("Default database is %s.\n", $row[0]);
    mysqli_free_result($result);
}

mysqli_close($link);
?>
```

The above examples will output:

```
Default database is test.
Default database is world.
```

### See Also

[mysqli\\_connect](#)  
[mysqli\\_real\\_connect](#)

#### 21.9.3.9.48. [mysqli::send\\_query](#), [mysqli\\_send\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::send\\_query](#)

[mysqli\\_send\\_query](#)

Send the query and return

### Description

Object oriented style

```
bool mysqli::send_query(
    string query);
```

Procedural style

```
bool mysqli_send_query(
    mysqli link,
    string query);
```



#### Warning

This function is currently not documented; only its argument list is available.



#### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

**21.9.3.9.49. *mysqli::set\_charset*, *mysqli\_set\_charset***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli::set\_charset*

*mysqli\_set\_charset*

Sets the default client character set

**Description**

Object oriented style

```
bool mysqli::set_charset(
    string charset);
```

Procedural style

```
bool mysqli_set_charset(
    mysqli link,
    string charset);
```

Sets the default character set to be used when sending data from and to the database server.

**Parameters***link*Procedural style only: A link identifier returned by *mysqli\_connect* or *mysqli\_init**charset*

The charset to be set as default.

**Return Values**Returns **TRUE** on success or **FALSE** on failure.**Notes****Note**

To use this function on a Windows platform you need MySQL client library version 4.1.11 or above (for MySQL 5.0 you need 5.0.6 or above).

**Note**This is the preferred way to change the charset. Using *mysqli\_query* to set it (such as **SET NAMES utf8**) is not recommended. See the [MySQL character set concepts](#) section for more information.**Examples****Example 21.142. *mysqli::set\_charset* example**

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "test");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}
```

```
/* change character set to utf8 */
if (!$mysqli->set_charset("utf8")) {
    printf("Error loading character set utf8: %s\n", $mysqli->error);
} else {
    printf("Current character set: %s\n", $mysqli->character_set_name());
}

$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect('localhost', 'my_user', 'my_password', 'test');

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* change character set to utf8 */
if (!mysqli_set_charset($link, "utf8")) {
    printf("Error loading character set utf8: %s\n", mysqli_error($link));
} else {
    printf("Current character set: %s\n", mysqli_character_set_name($link));
}

mysqli_close($link);
?>
```

The above examples will output:

```
Current character set: utf8
```

### See Also

[mysqli\\_character\\_set\\_name](#)  
[mysqli\\_real\\_escape\\_string](#)  
[List of character sets that MySQL supports](#)

#### 21.9.3.9.50. [mysqli::set\\_local\\_infile\\_default](#), [mysqli\\_set\\_local\\_infile\\_default](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::set\\_local\\_infile\\_default](#)  
[mysqli\\_set\\_local\\_infile\\_default](#)

Unsets user defined handler for load local infile command

### Description

```
void mysqli_set_local_infile_default(
    mysqli link);
```

Deactivates a [LOAD DATA INFILE LOCAL](#) handler previously set with [mysqli\\_set\\_local\\_infile\\_handler](#).

### Parameters



*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

**Return Values**

No value is returned.

**Examples**

See `mysqli_set_local_infile_handler` examples

**See Also**

`mysqli_set_local_infile_handler`

**21.9.3.9.51. `mysqli::set_local_infile_handler`, `mysqli_set_local_infile_handler`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::set_local_infile_handler`

`mysqli_set_local_infile_handler`

Set callback function for LOAD DATA LOCAL INFILE command

**Description**

Object oriented style

```
bool mysqli::set_local_infile_handler(
    mysqli link,
    callable read_func);
```

Procedural style

```
bool mysqli_set_local_infile_handler(
    mysqli link,
    callable read_func);
```

Set callback function for LOAD DATA LOCAL INFILE command

The callbacks task is to read input from the file specified in the `LOAD DATA LOCAL INFILE` and to reformat it into the format understood by `LOAD DATA INFILE`.

The returned data needs to match the format specified in the `LOAD DATA`

**Parameters**

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

*read\_func* A callback function or object method taking the following parameters:

- stream* A PHP stream associated with the SQL commands INFILE
- &buffer* A string buffer to store the rewritten input into
- buflen* The maximum number of characters to be stored in the buffer

*&errmsg*

If an error occurs you can store an error message in here

The callback function should return the number of characters stored in the *buffer* or a negative value if an error occurred.

### Return Values

Returns **TRUE** on success or **FALSE** on failure.

### Examples

#### Example 21.143. *mysqli::set\_local\_infile\_handler* example

Object oriented style

```
<?php
$db = mysqli_init();
$db->real_connect("localhost","root","","test");

function callme($stream, &$buffer, $buflen, &$errmsg)
{
    $buffer = fgets($stream);

    echo $buffer;

    // convert to upper case and replace "," delimiter with [TAB]
    $buffer = strtoupper(str_replace(",","\\t", $buffer));

    return strlen($buffer);
}

echo "Input:\\n";

$db->set_local_infile_handler("callme");
$db->query("LOAD DATA LOCAL INFILE 'input.txt' INTO TABLE t1");
$db->set_local_infile_default();

$res = $db->query("SELECT * FROM t1");

echo "\\nResult:\\n";
while ($row = $res->fetch_assoc()) {
    echo join(",", $row). "\\n";
}
?>
```

Procedural style

```
<?php
$db = mysqli_init();
mysqli_real_connect($db, "localhost","root","","test");

function callme($stream, &$buffer, $buflen, &$errmsg)
{
    $buffer = fgets($stream);

    echo $buffer;

    // convert to upper case and replace "," delimiter with [TAB]
    $buffer = strtoupper(str_replace(",","\\t", $buffer));

    return strlen($buffer);
}
```

```

echo "Input:\n";

mysqli_set_local_infile_handler($db, "callme");
mysqli_query($db, "LOAD DATA LOCAL INFILE 'input.txt' INTO TABLE t1");
mysqli_set_local_infile_default($db);

$res = mysqli_query($db, "SELECT * FROM t1");

echo "\nResult:\n";
while ($row = mysqli_fetch_assoc($res)) {
    echo join(", ", $row)."\n";
}
?>

```

The above examples will output:

```

Input :
23,foo
42,bar

Output:
23,FOO
42,BAR

```

### See Also

`mysqli_set_local_infile_default`

#### 21.9.3.9.52. `mysqli::$sqlstate`, `mysqli_sqlstate`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$sqlstate`

`mysqli_sqlstate`

Returns the SQLSTATE error from previous MySQL operation

### Description

Object oriented style

```

string
mysqli->sqlstate ;

```

Procedural style

```

string mysqli_sqlstate(
    mysqli link);

```

Returns a string containing the SQLSTATE error code for the last error. The error code consists of five characters. '00000' means no error. The values are specified by ANSI SQL and ODBC. For a list of possible values, see <http://dev.mysql.com/doc/mysql/en/error-handling.html>.



### Note

Note that not all MySQL errors are yet mapped to SQLSTATE's. The value `HY000` (general error) is used for unmapped errors.

## Parameters

[link](#) Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

## Return Values

Returns a string containing the SQLSTATE error code for the last error. The error code consists of five characters. '00000' means no error.

## Examples

### Example 21.144. [\\$mysqli->sqlstate](#) example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Table City already exists, so we should get an error */
if (!$mysqli->query("CREATE TABLE City (ID INT, Name VARCHAR(30))")) {
    printf("Error - SQLSTATE %s.\n", $mysqli->sqlstate);
}

$mysqli->close();
?>
```

Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* Table City already exists, so we should get an error */
if (!mysqli_query($link, "CREATE TABLE City (ID INT, Name VARCHAR(30))")) {
    printf("Error - SQLSTATE %s.\n", mysqli_sqlstate($link));
}

mysqli_close($link);
?>
```

The above examples will output:

```
Error - SQLSTATE 42S01.
```

## See Also

[mysqli\\_errno](#)

`mysqli_error`

### 21.9.3.9.53. `mysqli::ssl_set`, `mysqli_ssl_set`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::ssl_set`

`mysqli_ssl_set`

Used for establishing secure connections using SSL

#### Description

Object oriented style

```
bool mysqli::ssl_set(  
    string key,  
    string cert,  
    string ca,  
    string capath,  
    string cipher);
```

Procedural style

```
bool mysqli_ssl_set(  
    mysqli link,  
    string key,  
    string cert,  
    string ca,  
    string capath,  
    string cipher);
```

Used for establishing secure connections using SSL. It must be called before `mysqli_real_connect`. This function does nothing unless OpenSSL support is enabled.

Note that MySQL Native Driver does not support SSL before PHP 5.3.3, so calling this function when using MySQL Native Driver will result in an error. MySQL Native Driver is enabled by default on Microsoft Windows from PHP version 5.3 onwards.

#### Parameters

<i>link</i>	Procedural style only: A link identifier returned by <code>mysqli_connect</code> or <code>mysqli_init</code>
<i>key</i>	The path name to the key file.
<i>cert</i>	The path name to the certificate file.
<i>ca</i>	The path name to the certificate authority file.
<i>capath</i>	The pathname to a directory that contains trusted SSL CA certificates in PEM format.
<i>cipher</i>	A list of allowable ciphers to use for SSL encryption.

Any unused SSL parameters may be given as `NULL`

#### Return Values

This function always returns `TRUE` value. If SSL setup is incorrect `mysqli_real_connect` will return an error when you attempt to connect.

#### See Also

```
mysqli_options  
mysqli_real_connect
```

### 21.9.3.9.54. *mysqli::stat*, *mysqli\_stat*

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli::stat*

```
mysqli_stat
```

Gets the current system status

#### Description

Object oriented style

```
string mysqli::stat();
```

Procedural style

```
string mysqli_stat(  
    mysqli link);
```

*mysqli\_stat* returns a string containing information similar to that provided by the 'mysqladmin status' command. This includes uptime in seconds and the number of running threads, questions, reloads, and open tables.

#### Parameters

*link*

Procedural style only: A link identifier returned by *mysqli\_connect* or *mysqli\_init*

#### Return Values

A string describing the server status. *FALSE* if an error occurred.

#### Examples

##### Example 21.145. *mysqli::stat* example

Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
printf ("System status: %s\n", $mysqli->stat());  
  
$mysqli->close();  
?>
```

Procedural style

```
<?php  
$link = mysqli_connect("localhost", "my_user", "my_password", "world");
```

```
/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

printf("System status: %s\n", mysqli_stat($link));

mysqli_close($link);
?>
```

The above examples will output:

```
System status: Uptime: 272  Threads: 1  Questions: 5340  Slow queries: 0
Opens: 13  Flush tables: 1  Open tables: 0  Queries per second avg: 19.632
Memory in use: 8496K  Max memory used: 8560K
```

### See Also

[mysqli\\_get\\_server\\_info](#)

#### 21.9.3.9.55. [mysqli::stmt\\_init](#), [mysqli\\_stmt\\_init](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::stmt\\_init](#)

[mysqli\\_stmt\\_init](#)

Initializes a statement and returns an object for use with [mysqli\\_stmt\\_prepare](#)

### Description

Object oriented style

```
mysqli_stmt $stmt = mysqli::stmt_init();
```

Procedural style

```
mysqli_stmt $stmt = mysqli_stmt_init(
    mysqli $link);
```

Allocates and initializes a statement object suitable for [mysqli\\_stmt\\_prepare](#).



#### Note

Any subsequent calls to any [mysqli\\_stmt](#) function will fail until [mysqli\\_stmt\\_prepare](#) was called.

### Parameters

*link*

Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

### Return Values

Returns an object.

### See Also

```
mysqli_stmt_prepare
```

### 21.9.3.9.56. [mysqli::store\\_result](#), [mysqli\\_store\\_result](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::store\\_result](#)

```
mysqli_store_result
```

Transfers a result set from the last query

#### Description

Object oriented style

```
mysqli_result mysqli::store_result();
```

Procedural style

```
mysqli_result mysqli_store_result(  
    mysqli link);
```

Transfers the result set from the last query on the database connection represented by the [link](#) parameter to be used with the [mysqli\\_data\\_seek](#) function.

#### Parameters

[link](#) Procedural style only: A link identifier returned by [mysqli\\_connect](#) or [mysqli\\_init](#)

#### Return Values

Returns a buffered result object or [FALSE](#) if an error occurred.



#### Note

[mysqli\\_store\\_result](#) returns [FALSE](#) in case the query didn't return a result set (if the query was, for example an INSERT statement). This function also returns [FALSE](#) if the reading of the result set failed. You can check if you have got an error by checking if [mysqli\\_error](#) doesn't return an empty string, if [mysqli\\_errno](#) returns a non zero value, or if [mysqli\\_field\\_count](#) returns a non zero value. Also possible reason for this function returning [FALSE](#) after successful call to [mysqli\\_query](#) can be too large result set (memory for it cannot be allocated). If [mysqli\\_field\\_count](#) returns a non-zero value, the statement should have produced a non-empty result set.

#### Notes



#### Note

Although it is always good practice to free the memory used by the result of a query using the [mysqli\\_free\\_result](#) function, when transferring large result sets using the [mysqli\\_store\\_result](#) this becomes particularly important.

#### Examples

See [mysqli\\_multi\\_query](#).

#### See Also

[mysqli\\_real\\_query](#)



`mysqli_use_result`

### 21.9.3.9.57. `mysqli::$thread_id, mysqli_thread_id`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$thread_id`

`mysqli_thread_id`

Returns the thread ID for the current connection

#### Description

Object oriented style

```
int
mysqli->thread_id ;
```

Procedural style

```
int mysqli_thread_id(
mysqli link);
```

The `mysqli_thread_id` function returns the thread ID for the current connection which can then be killed using the `mysqli_kill` function. If the connection is lost and you reconnect with `mysqli_ping`, the thread ID will be other. Therefore you should get the thread ID only when you need it.



#### Note

The thread ID is assigned on a connection-by-connection basis. Hence, if the connection is broken and then re-established a new thread ID will be assigned.

To kill a running query you can use the SQL command `KILL QUERY processid`.

#### Parameters

`link` Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

#### Return Values

Returns the Thread ID for the current connection.

#### Examples

##### Example 21.146. `$mysqli->thread_id` example

Object oriented style

```
<?php
mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* determine our thread id */
$thread_id = $mysqli->thread_id;
```

```
/* Kill connection */
$mysqli->kill($thread_id);

/* This should produce an error */
if (!$mysqli->query("CREATE TABLE myCity LIKE City")) {
    printf("Error: %s\n", $mysqli->error);
    exit;
}

/* close connection */
$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* determine our thread id */
$thread_id = mysqli_thread_id($link);

/* Kill connection */
mysqli_kill($link, $thread_id);

/* This should produce an error */
if (!mysqli_query($link, "CREATE TABLE myCity LIKE City")) {
    printf("Error: %s\n", mysqli_error($link));
    exit;
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Error: MySQL server has gone away
```

### See Also

[mysqli\\_kill](#)

#### 21.9.3.9.58. [mysqli::thread\\_safe](#), [mysqli\\_thread\\_safe](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli::thread\\_safe](#)

[mysqli\\_thread\\_safe](#)

Returns whether thread safety is given or not

### Description

Procedural style

```
bool mysqli_thread_safe();
```

Tells whether the client library is compiled as thread-safe.

### Return Values

`TRUE` if the client library is thread-safe, otherwise `FALSE`.

#### 21.9.3.9.59. `mysqli::use_result`, `mysqli_use_result`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::use_result`

`mysqli_use_result`

Initiate a result set retrieval

### Description

Object oriented style

```
mysqli_result mysqli::use_result();
```

Procedural style

```
mysqli_result mysqli_use_result(
    mysqli link);
```

Used to initiate the retrieval of a result set from the last query executed using the `mysqli_real_query` function on the database connection.

Either this or the `mysqli_store_result` function must be called before the results of a query can be retrieved, and one or the other must be called to prevent the next query on that database connection from failing.



#### Note

The `mysqli_use_result` function does not transfer the entire result set from the database and hence cannot be used functions such as `mysqli_data_seek` to move to a particular row within the set. To use this functionality, the result set must be stored using `mysqli_store_result`. One should not use `mysqli_use_result` if a lot of processing on the client side is performed, since this will tie up the server and prevent other threads from updating any tables from which the data is being fetched.

### Return Values

Returns an unbuffered result object or `FALSE` if an error occurred.

### Examples

#### Example 21.147. `mysqli::use_result` example

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
```

```
printf("Connect failed: %s\n", mysqli_connect_error());
exit();
}

$query = "SELECT CURRENT_USER()";
$query .= "SELECT Name FROM City ORDER BY ID LIMIT 20, 5";

/* execute multi query */
if ($mysqli->multi_query($query)) {
    do {
        /* store first result set */
        if ($result = $mysqli->use_result()) {
            while ($row = $result->fetch_row()) {
                printf("%s\n", $row[0]);
            }
            $result->close();
        }
        /* print divider */
        if ($mysqli->more_results()) {
            printf("-----\n");
        }
    } while ($mysqli->next_result());
}

/* close connection */
$mysqli->close();
?>
```

### Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT CURRENT_USER()";
$query .= "SELECT Name FROM City ORDER BY ID LIMIT 20, 5";

/* execute multi query */
if (mysqli_multi_query($link, $query)) {
    do {
        /* store first result set */
        if ($result = mysqli_use_result($link)) {
            while ($row = mysqli_fetch_row($result)) {
                printf("%s\n", $row[0]);
            }
            mysqli_free_result($result);
        }
        /* print divider */
        if (mysqli_more_results($link)) {
            printf("-----\n");
        }
    } while (mysqli_next_result($link));
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
my_user@localhost
```

```
-----
Amersfoort
Maastricht
Dordrecht
Leiden
Haarlemmermeer
```

**See Also**

`mysqli_real_query`  
`mysqli_store_result`

**21.9.3.9.60. `mysqli::$warning_count`, `mysqli_warning_count`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::$warning_count`

`mysqli_warning_count`

Returns the number of warnings from the last query for the given link

**Description**

Object oriented style

```
int
mysqli->warning_count ;
```

Procedural style

```
int mysqli_warning_count(
    mysqli link);
```

Returns the number of warnings from the last query in the connection.

**Note**

For retrieving warning messages you can use the SQL command `SHOW WARNINGS [limit row_count]`.

**Parameters**

*link* Procedural style only: A link identifier returned by `mysqli_connect` or `mysqli_init`

**Return Values**

Number of warnings or zero if there are no warnings.

**Examples****Example 21.148. `$mysqli->warning_count` example**

Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
```

```

if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCity LIKE City");

/* a remarkable city in Wales */
$query = "INSERT INTO myCity (CountryCode, Name) VALUES('GBR',
    'Llanfairpwllgwyngyllgogerychwyrndrobwlllllantysiliogogoch')";

$mysqli->query($query);

if ($mysqli->warning_count) {
    if ($result = $mysqli->query("SHOW WARNINGS")) {
        $row = $result->fetch_row();
        printf("%s (%d): %s\n", $row[0], $row[1], $row[2]);
        $result->close();
    }
}

/* close connection */
$mysqli->close();
?>

```

### Procedural style

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCity LIKE City");

/* a remarkable long city name in Wales */
$query = "INSERT INTO myCity (CountryCode, Name) VALUES('GBR',
    'Llanfairpwllgwyngyllgogerychwyrndrobwlllllantysiliogogoch')";

mysqli_query($link, $query);

if (mysqli_warning_count($link)) {
    if ($result = mysqli_query($link, "SHOW WARNINGS")) {
        $row = mysqli_fetch_row($result);
        printf("%s (%d): %s\n", $row[0], $row[1], $row[2]);
        mysqli_free_result($result);
    }
}

/* close connection */
mysqli_close($link);
?>

```

The above examples will output:

```
Warning (1264): Data truncated for column 'Name' at row 1
```

### See Also

[mysqli\\_errno](#)

```
mysqli_error  
mysqli_sqlstate
```

### 21.9.3.10. The `mysqli_stmt` class (`mysqli_stmt`)

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Represents a prepared statement.

```
mysqli_stmt {  
mysqli_stmt  
  
    Properties  
  
    int  
        mysqli_stmt->affected_rows ;  
  
    int  
        mysqli_stmt->errno ;  
  
    array  
        mysqli_stmt->error_list ;  
  
    string  
        mysqli_stmt->error ;  
  
    int  
        mysqli_stmt->field_count ;  
  
    int  
        mysqli_stmt->insert_id ;  
  
    int  
        mysqli_stmt->num_rows ;  
  
    int  
        mysqli_stmt->param_count ;  
  
    string  
        mysqli_stmt->sqlstate ;  
  
Methods  
  
    int mysqli_stmt::attr_get(  
        int attr);  
  
    bool mysqli_stmt::attr_set(  
        int attr,  
        int mode);  
  
    bool mysqli_stmt::bind_param(  
        string types,  
        mixed var1,  
        mixed ...);  
  
    bool mysqli_stmt::bind_result(  
        mixed var1,  
        mixed ...);  
  
    bool mysqli_stmt::close();  
  
    void mysqli_stmt::data_seek(  
        int offset);  
  
    bool mysqli_stmt::execute();  
  
    bool mysqli_stmt::fetch();
```

```
void mysqli_stmt::free_result();

mysqli_result mysqli_stmt::get_result();

object mysqli_stmt::get_warnings(
    mysqli_stmt stmt);

mixed mysqli_stmt::prepare(
    string query);

bool mysqli_stmt::reset();

mysqli_result mysqli_stmt::result_metadata();

bool mysqli_stmt::send_long_data(
    int param_nr,
    string data);

bool mysqli_stmt::store_result();
}
```

### 21.9.3.10.1. [mysqli\\_stmt::\\$affected\\_rows](#), [mysqli\\_stmt\\_affected\\_rows](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$affected\\_rows](#)

[mysqli\\_stmt\\_affected\\_rows](#)

Returns the total number of rows changed, deleted, or inserted by the last executed statement

#### Description

Object oriented style

```
int
mysqli_stmt->affected_rows ;
```

Procedural style

```
int mysqli_stmt_affected_rows(
    mysqli_stmt stmt);
```

Returns the number of rows affected by [INSERT](#), [UPDATE](#), or [DELETE](#) query.

This function only works with queries which update a table. In order to get the number of rows from a [SELECT](#) query, use [mysqli\\_stmt\\_num\\_rows](#) instead.

#### Parameters

*stmt*

Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

#### Return Values

An integer greater than zero indicates the number of rows affected or retrieved. Zero indicates that no records were updated for an [UPDATE](#)/[DELETE](#) statement, no rows matched the [WHERE](#) clause in the query or that no query has yet been executed. -1 indicates that the query has returned an error. [NULL](#) indicates an invalid argument was supplied to the function.



#### Note

If the number of affected rows is greater than maximal PHP int value, the number of affected rows will be returned as a string value.



## Examples

### Example 21.149. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* create temp table */
$mysqli->query("CREATE TEMPORARY TABLE myCountry LIKE Country");

$query = "INSERT INTO myCountry SELECT * FROM Country WHERE Code LIKE ?";

/* prepare statement */
if ($stmt = $mysqli->prepare($query)) {

    /* Bind variable for placeholder */
    $code = 'A%';
    $stmt->bind_param("s", $code);

    /* execute statement */
    $stmt->execute();

    printf("rows inserted: %d\n", $stmt->affected_rows);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.150. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* create temp table */
mysqli_query($link, "CREATE TEMPORARY TABLE myCountry LIKE Country");

$query = "INSERT INTO myCountry SELECT * FROM Country WHERE Code LIKE ?";

/* prepare statement */
if ($stmt = mysqli_prepare($link, $query)) {

    /* Bind variable for placeholder */
    $code = 'A%';
    mysqli_stmt_bind_param($stmt, "s", $code);

    /* execute statement */
    mysqli_stmt_execute($stmt);

    printf("rows inserted: %d\n", mysqli_stmt_affected_rows($stmt));

    /* close statement */
```

```
mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
rows inserted: 17
```

### See Also

[mysqli\\_stmt\\_num\\_rows](#)  
[mysqli\\_prepare](#)

#### 21.9.3.10.2. [mysqli\\_stmt::attr\\_get](#), [mysqli\\_stmt\\_attr\\_get](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::attr\\_get](#)

[mysqli\\_stmt\\_attr\\_get](#)

Used to get the current value of a statement attribute

### Description

Object oriented style

```
int mysqli_stmt::attr_get(
    int attr);
```

Procedural style

```
int mysqli_stmt_attr_get(
    mysqli_stmt stmt,
    int attr);
```

Gets the current value of a statement attribute.

### Parameters

*stmt* Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

*attr* The attribute that you want to get.

### Return Values

Returns [FALSE](#) if the attribute is not found, otherwise returns the value of the attribute.

#### 21.9.3.10.3. [mysqli\\_stmt::attr\\_set](#), [mysqli\\_stmt\\_attr\\_set](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::attr\\_set](#)

[mysqli\\_stmt\\_attr\\_set](#)

Used to modify the behavior of a prepared statement

**Description**

Object oriented style

```
bool mysqli_stmt::attr_set(
    int attr,
    int mode);
```

Procedural style

```
bool mysqli_stmt_attr_set(
    mysqli_stmt stmt,
    int attr,
    int mode);
```

Used to modify the behavior of a prepared statement. This function may be called multiple times to set several attributes.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by *mysqli\_stmt\_init*.

*attr* The attribute that you want to set. It can have one of the following values:

**Table 21.49. Attribute values**

Character	Description
MYSQLI_STMT_ATTR_UPDATE_MAX_LENGTH	Use to enable <i>mysqli_stmt_store_result</i> to update the metadata <i>MYSQL_FIELD-&gt;max_length</i> value.
MYSQLI_STMT_ATTR_CURSOR_TYPE	Type of cursor to open for statement when <i>mysqli_stmt_execute</i> is invoked. <i>mode</i> can be <i>MYSQLI_CURSOR_TYPE_NO_CURSOR</i> (the default) or <i>MYSQLI_CURSOR_TYPE_READ_ONLY</i> .
MYSQLI_STMT_ATTR_PREFETCH_ROWS	Number of rows to fetch from server at a time when using a cursor. <i>mode</i> can be in the range from 1 to the maximum value of unsigned long. The default is 1.

If you use the *MYSQLI\_STMT\_ATTR\_CURSOR\_TYPE* option with *MYSQLI\_CURSOR\_TYPE\_READ\_ONLY*, a cursor is opened for the statement when you invoke *mysqli\_stmt\_execute*. If there is already an open cursor from a previous *mysqli\_stmt\_execute* call, it closes the cursor before opening a new one. *mysqli\_stmt\_reset* also closes any open cursor before preparing the statement for re-execution. *mysqli\_stmt\_free\_result* closes any open cursor.

If you open a cursor for a prepared statement, *mysqli\_stmt\_store\_result* is unnecessary.

*mode*

The value to assign to the attribute.

**21.9.3.10.4. *mysqli\_stmt::bind\_param*, *mysqli\_stmt\_bind\_param***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli\_stmt::bind\_param*

*mysqli\_stmt\_bind\_param*

Binds variables to a prepared statement as parameters

**Description**

Object oriented style

```
bool mysqli_stmt::bind_param(
    string types,
    mixed var1,
    mixed ...);
```

Procedural style

```
bool mysqli_stmt_bind_param(
    mysqli_stmt stmt,
    string types,
    mixed var1,
    mixed ...);
```

Bind variables for the parameter markers in the SQL statement that was passed to *mysqli\_prepare*.**Note**

If data size of a variable exceeds max. allowed packet size (*max\_allowed\_packet*), you have to specify *b* in *types* and use *mysqli\_stmt\_send\_long\_data* to send the data in packets.

**Note**

Care must be taken when using *mysqli\_stmt\_bind\_param* in conjunction with *call\_user\_func\_array*. Note that *mysqli\_stmt\_bind\_param* requires parameters to be passed by reference, whereas *call\_user\_func\_array* can accept as a parameter a list of variables that can represent references or values.

**Parameters***stmt*Procedural style only: A statement identifier returned by *mysqli\_stmt\_init*.*types*

A string that contains one or more characters which specify the types for the corresponding bind variables:

**Table 21.50. Type specification chars**

Character	Description
i	corresponding variable has type integer
d	corresponding variable has type double
s	corresponding variable has type string

Character	Description
b	corresponding variable is a blob and will be sent in packets

*var1*

The number of variables and length of string *types* must match the parameters in the statement.

## Return Values

Returns **TRUE** on success or **FALSE** on failure.

## Examples

### Example 21.151. Object oriented style

```
<?php
$mysqli = new mysqli('localhost', 'my_user', 'my_password', 'world');

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$stmt = $mysqli->prepare("INSERT INTO CountryLanguage VALUES (?, ?, ?, ?)");
$stmt->bind_param('sssd', $code, $language, $official, $percent);

$code = 'DEU';
$language = 'Bavarian';
$official = "F";
$percent = 11.2;

/* execute prepared statement */
$stmt->execute();

printf("%d Row inserted.\n", $stmt->affected_rows);

/* close statement and connection */
$stmt->close();

/* Clean up table CountryLanguage */
$mysqli->query("DELETE FROM CountryLanguage WHERE Language='Bavarian'");
printf("%d Row deleted.\n", $mysqli->affected_rows);

/* close connection */
$mysqli->close();
?>
```

### Example 21.152. Procedural style

```
<?php
$link = mysqli_connect('localhost', 'my_user', 'my_password', 'world');

/* check connection */
if (!$link) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$stmt = mysqli_prepare($link, "INSERT INTO CountryLanguage VALUES (?, ?, ?, ?)");
mysqli_stmt_bind_param($stmt, 'sssd', $code, $language, $official, $percent);

$code = 'DEU';
```

```
$language = 'Bavarian';
$official = "F";
$percent = 11.2;

/* execute prepared statement */
mysqli_stmt_execute($stmt);

printf("%d Row inserted.\n", mysqli_stmt_affected_rows($stmt));

/* close statement and connection */
mysqli_stmt_close($stmt);

/* Clean up table CountryLanguage */
mysqli_query($link, "DELETE FROM CountryLanguage WHERE Language='Bavarian'");
printf("%d Row deleted.\n", mysqli_affected_rows($link));

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
1 Row inserted.
1 Row deleted.
```

### See Also

[mysqli\\_stmt\\_bind\\_result](#)  
[mysqli\\_stmt\\_execute](#)  
[mysqli\\_stmt\\_fetch](#)  
[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_send\\_long\\_data](#)  
[mysqli\\_stmt\\_errno](#)  
[mysqli\\_stmt\\_error](#)

### 21.9.3.10.5. [mysqli\\_stmt::bind\\_result](#), [mysqli\\_stmt\\_bind\\_result](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::bind\\_result](#)

[mysqli\\_stmt\\_bind\\_result](#)

Binds variables to a prepared statement for result storage

#### Description

Object oriented style

```
bool mysqli_stmt::bind_result(
    mixed var1,
    mixed ...);
```

Procedural style

```
bool mysqli_stmt_bind_result(
    mysqli_stmt stmt,
    mixed var1,
    mixed ...);
```

Binds columns in the result set to variables.

When `mysqli_stmt_fetch` is called to fetch data, the MySQL client/server protocol places the data for the bound columns into the specified variables `var1`, ....



### Note

Note that all columns must be bound after `mysqli_stmt_execute` and prior to calling `mysqli_stmt_fetch`. Depending on column types bound variables can silently change to the corresponding PHP type.

A column can be bound or rebound at any time, even after a result set has been partially retrieved. The new binding takes effect the next time `mysqli_stmt_fetch` is called.

### Parameters

<code>stmt</code>	Procedural style only: A statement identifier returned by <code>mysqli_stmt_init</code> .
<code>var1</code>	The variable to be bound.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Examples

#### Example 21.153. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* prepare statement */
if ($stmt = $mysqli->prepare("SELECT Code, Name FROM Country ORDER BY Name LIMIT 5")) {
    $stmt->execute();

    /* bind variables to prepared statement */
    $stmt->bind_result($col1, $col2);

    /* fetch values */
    while ($stmt->fetch()) {
        printf("%s %s\n", $col1, $col2);
    }

    /* close statement */
    $stmt->close();
}
/* close connection */
$mysqli->close();

?>
```

#### Example 21.154. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (!$link) {
```

```
printf("Connect failed: %s\n", mysqli_connect_error());
exit();
}

/* prepare statement */
if ($stmt = mysqli_prepare($link, "SELECT Code, Name FROM Country ORDER BY Name LIMIT 5")) {
    mysqli_stmt_execute($stmt);

    /* bind variables to prepared statement */
    mysqli_stmt_bind_result($stmt, $col1, $col2);

    /* fetch values */
    while (mysqli_stmt_fetch($stmt)) {
        printf("%s %s\n", $col1, $col2);
    }

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
AFG Afghanistan
ALB Albania
DZA Algeria
ASM American Samoa
AND Andorra
```

### See Also

[mysqli\\_stmt\\_bind\\_param](#)  
[mysqli\\_stmt\\_execute](#)  
[mysqli\\_stmt\\_fetch](#)  
[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_prepare](#)  
[mysqli\\_stmt\\_init](#)  
[mysqli\\_stmt\\_errno](#)  
[mysqli\\_stmt\\_error](#)

#### 21.9.3.10.6. [mysqli\\_stmt::close](#), [mysqli\\_stmt\\_close](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::close](#)

[mysqli\\_stmt\\_close](#)

Closes a prepared statement

### Description

Object oriented style

```
bool mysqli_stmt::close();
```

Procedural style

```
bool mysqli_stmt_close(
```



```
mysqli_stmt stmt);
```

Closes a prepared statement. [mysqli\\_stmt\\_close](#) also deallocates the statement handle. If the current statement has pending or unread results, this function cancels them so that the next query can be executed.

### Parameters

*stmt* Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

### Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

### See Also

[mysqli\\_prepare](#)

#### 21.9.3.10.7. [mysqli\\_stmt::data\\_seek](#), [mysqli\\_stmt\\_data\\_seek](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::data\\_seek](#)

[mysqli\\_stmt\\_data\\_seek](#)

Seeks to an arbitrary row in statement result set

### Description

Object oriented style

```
void mysqli_stmt::data_seek(  
    int offset);
```

Procedural style

```
void mysqli_stmt_data_seek(  
    mysqli_stmt stmt,  
    int offset);
```

Seeks to an arbitrary result pointer in the statement result set.

[mysqli\\_stmt\\_store\\_result](#) must be called prior to [mysqli\\_stmt\\_data\\_seek](#).

### Parameters

*stmt* Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

*offset* Must be between zero and the total number of rows minus one (0.. [mysqli\\_stmt\\_num\\_rows](#) - 1).

### Return Values

No value is returned.

### Examples

#### Example 21.155. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name";
if ($stmt = $mysqli->prepare($query)) {

    /* execute query */
    $stmt->execute();

    /* bind result variables */
    $stmt->bind_result($name, $code);

    /* store result */
    $stmt->store_result();

    /* seek to row no. 400 */
    $stmt->data_seek(399);

    /* fetch values */
    $stmt->fetch();

    printf ("City: %s Countrycode: %s\n", $name, $code);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.156. Procedural style

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name";
if ($stmt = mysqli_prepare($link, $query)) {

    /* execute query */
    mysqli_stmt_execute($stmt);

    /* bind result variables */
    mysqli_stmt_bind_result($stmt, $name, $code);

    /* store result */
    mysqli_stmt_store_result($stmt);

    /* seek to row no. 400 */
    mysqli_stmt_data_seek($stmt, 399);

    /* fetch values */
    mysqli_stmt_fetch($stmt);

    printf ("City: %s Countrycode: %s\n", $name, $code);
}
```

```
    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
City: Benin City Countrycode: NGA
```

### See Also

[mysqli\\_prepare](#)

#### 21.9.3.10.8. [mysqli\\_stmt::\\$errno](#), [mysqli\\_stmt\\_errno](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$errno](#)

[mysqli\\_stmt\\_errno](#)

Returns the error code for the most recent statement call

### Description

Object oriented style

```
int
mysqli_stmt->errno ;
```

Procedural style

```
int mysqli_stmt_errno(
mysqli_stmt stmt);
```

Returns the error code for the most recently invoked statement function that can succeed or fail.

Client error message numbers are listed in the MySQL [errmsg.h](#) header file, server error message numbers are listed in [mysqld\\_error.h](#). In the MySQL source distribution you can find a complete list of error messages and error numbers in the file [Docs/mysqld\\_error.txt](#).

### Parameters

*stmt*

Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

### Return Values

An error code value. Zero means no error occurred.

### Examples

#### Example 21.157. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCountry LIKE Country");
$mysqli->query("INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = $mysqli->prepare($query)) {

    /* drop table */
    $mysqli->query("DROP TABLE myCountry");

    /* execute query */
    $stmt->execute();

    printf("Error: %d.\n", $stmt->errno);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.158. Procedural style

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCountry LIKE Country");
mysqli_query($link, "INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = mysqli_prepare($link, $query)) {

    /* drop table */
    mysqli_query($link, "DROP TABLE myCountry");

    /* execute query */
    mysqli_stmt_execute($stmt);

    printf("Error: %d.\n", mysqli_stmt_errno($stmt));

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Error: 1146.
```

### See Also

[mysqli\\_stmt\\_error](#)  
[mysqli\\_stmt\\_sqlstate](#)

#### 21.9.3.10.9. [mysqli\\_stmt::\\$error\\_list](#), [mysqli\\_stmt\\_error\\_list](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$error\\_list](#)

[mysqli\\_stmt\\_error\\_list](#)

Returns a list of errors from the last statement executed

### Description

Object oriented style

```
array  
mysqli_stmt->error_list ;
```

Procedural style

```
array mysqli_stmt_error_list(  
    mysqli_stmt stmt);
```

Returns an array of errors for the most recently invoked statement function that can succeed or fail.

### Parameters

*stmt* Procedural style only: A statement identifier returned by [mysqli\\_stmt\\_init](#).

### Return Values

A list of errors, each as an associative array containing the errno, error, and sqlstate.

### Examples

#### Example 21.159. Object oriented style

```
<?php  
/* Open a connection */  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}  
  
$mysqli->query("CREATE TABLE myCountry LIKE Country");  
$mysqli->query("INSERT INTO myCountry SELECT * FROM Country");  
  
$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
```

```

if ($stmt = $mysqli->prepare($query)) {

    /* drop table */
    $mysqli->query("DROP TABLE myCountry");

    /* execute query */
    $stmt->execute();

    echo "Error:\n";
    print_r($stmt->error_list);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>

```

**Example 21.160. Procedural style**

```

<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCountry LIKE Country");
mysqli_query($link, "INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = mysqli_prepare($link, $query)) {

    /* drop table */
    mysqli_query($link, "DROP TABLE myCountry");

    /* execute query */
    mysqli_stmt_execute($stmt);

    echo "Error:\n";
    print_r(mysqli_stmt_error_list($stmt));

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>

```

The above examples will output:

```

Array
(
    [0] => Array
        (
            [errno] => 1146
            [sqlstate] => 42S02
            [error] => Table 'world.myCountry' doesn't exist
        )
)

```

)

**See Also**

`mysqli_stmt_error`  
`mysqli_stmt_errno`  
`mysqli_stmt_sqlstate`

**21.9.3.10.10. `mysqli_stmt::$error`, `mysqli_stmt_error`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::$error`

`mysqli_stmt_error`

Returns a string description for last statement error

**Description**

Object oriented style

```
string
mysqli_stmt->error ;
```

Procedural style

```
string mysqli_stmt_error(
    mysqli_stmt stmt);
```

Returns a containing the error message for the most recently invoked statement function that can succeed or fail.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

**Return Values**

A string that describes the error. An empty string if no error occurred.

**Examples****Example 21.161. Object oriented style**

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCountry LIKE Country");
$mysqli->query("INSERT INTO myCountry SELECT * FROM Country");
```

```

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = $mysqli->prepare($query)) {

    /* drop table */
    $mysqli->query("DROP TABLE myCountry");

    /* execute query */
    $stmt->execute();

    printf("Error: %s.\n", $stmt->error);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>

```

**Example 21.162. Procedural style**

```

<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCountry LIKE Country");
mysqli_query($link, "INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = mysqli_prepare($link, $query)) {

    /* drop table */
    mysqli_query($link, "DROP TABLE myCountry");

    /* execute query */
    mysqli_stmt_execute($stmt);

    printf("Error: %s.\n", mysqli_stmt_error($stmt));

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>

```

The above examples will output:

```
Error: Table 'world.myCountry' doesn't exist.
```

**See Also**

[mysqli\\_stmt\\_errno](#)  
[mysqli\\_stmt\\_sqlstate](#)



**21.9.3.10.11. *mysqli\_stmt::execute*, *mysqli\_stmt\_execute***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli\_stmt::execute*

*mysqli\_stmt\_execute*

Executes a prepared Query

**Description**

Object oriented style

```
bool mysqli_stmt::execute();
```

Procedural style

```
bool mysqli_stmt_execute(
    mysqli_stmt stmt);
```

Executes a query that has been previously prepared using the *mysqli\_prepare* function. When executed any parameter markers which exist will automatically be replaced with the appropriate data.

If the statement is *UPDATE*, *DELETE*, or *INSERT*, the total number of affected rows can be determined by using the *mysqli\_stmt\_affected\_rows* function. Likewise, if the query yields a result set the *mysqli\_stmt\_fetch* function is used.

**Note**

When using *mysqli\_stmt\_execute*, the *mysqli\_stmt\_fetch* function must be used to fetch the data prior to performing any additional queries.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by *mysqli\_stmt\_init*.

**Return Values**

Returns *TRUE* on success or *FALSE* on failure.

**Examples****Example 21.163. Object oriented style**

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCity LIKE City");

/* Prepare an insert statement */
$query = "INSERT INTO myCity (Name, CountryCode, District) VALUES (?, ?, ?)";
$stmt = $mysqli->prepare($query);

$stmt->bind_param("sss", $val1, $val2, $val3);

$val1 = 'Stuttgart';
```

```

$val2 = 'DEU';
$val3 = 'Baden-Wuerttemberg';

/* Execute the statement */
$stmt->execute();

$val1 = 'Bordeaux';
$val2 = 'FRA';
$val3 = 'Aquitaine';

/* Execute the statement */
$stmt->execute();

/* close statement */
$stmt->close();

/* retrieve all rows from myCity */
$query = "SELECT Name, CountryCode, District FROM myCity";
if ($result = $mysqli->query($query)) {
    while ($row = $result->fetch_row()) {
        printf("%s (%s,%s)\n", $row[0], $row[1], $row[2]);
    }
    /* free result set */
    $result->close();
}

/* remove table */
$mysqli->query("DROP TABLE myCity");

/* close connection */
$mysqli->close();
?>

```

**Example 21.164. Procedural style**

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCity LIKE City");

/* Prepare an insert statement */
$query = "INSERT INTO myCity (Name, CountryCode, District) VALUES (?, ?, ?)";
$stmt = mysqli_prepare($link, $query);

mysqli_stmt_bind_param($stmt, "sss", $val1, $val2, $val3);

$val1 = 'Stuttgart';
$val2 = 'DEU';
$val3 = 'Baden-Wuerttemberg';

/* Execute the statement */
mysqli_stmt_execute($stmt);

$val1 = 'Bordeaux';
$val2 = 'FRA';
$val3 = 'Aquitaine';

/* Execute the statement */
mysqli_stmt_execute($stmt);

/* close statement */
mysqli_stmt_close($stmt);

```

```
/* retrieve all rows from myCity */
$query = "SELECT Name, CountryCode, District FROM myCity";
if ($result = mysqli_query($link, $query)) {
    while ($row = mysqli_fetch_row($result)) {
        printf("%s (%s,%s)\n", $row[0], $row[1], $row[2]);
    }
    /* free result set */
    mysqli_free_result($result);
}

/* remove table */
mysqli_query($link, "DROP TABLE myCity");

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Stuttgart (DEU,Baden-Wuerttemberg)
Bordeaux (FRA,Aquitaine)
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_bind\\_param](#)

#### 21.9.3.10.12. [mysqli\\_stmt::fetch](#), [mysqli\\_stmt\\_fetch](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::fetch](#)

[mysqli\\_stmt\\_fetch](#)

Fetch results from a prepared statement into the bound variables

### Description

Object oriented style

```
bool mysqli_stmt::fetch();
```

Procedural style

```
bool mysqli_stmt_fetch(
    mysqli_stmt stmt);
```

Fetch the result from a prepared statement into the variables bound by [mysqli\\_stmt\\_bind\\_result](#).



#### Note

Note that all columns must be bound by the application before calling [mysqli\\_stmt\\_fetch](#).



#### Note

Data are transferred unbuffered without calling [mysqli\\_stmt\\_store\\_result](#) which can decrease performance (but reduces memory cost).

### Parameters

*stmt*Procedural style only: A statement identifier returned by `mysqli_stmt_init`.**Return Values****Table 21.51. Return Values**

Value	Description
TRUE	Success. Data has been fetched
FALSE	Error occurred
NULL	No more rows/data exists or data truncation occurred

**Examples****Example 21.165. Object oriented style**

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 150,5";

if ($stmt = $mysqli->prepare($query)) {

    /* execute statement */
    $stmt->execute();

    /* bind result variables */
    $stmt->bind_result($name, $code);

    /* fetch values */
    while ($stmt->fetch()) {
        printf ("%s (%s)\n", $name, $code);
    }

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

**Example 21.166. Procedural style**

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 150,5";
```

```
if ($stmt = mysqli_prepare($link, $query)) {  
  
    /* execute statement */  
    mysqli_stmt_execute($stmt);  
  
    /* bind result variables */  
    mysqli_stmt_bind_result($stmt, $name, $code);  
  
    /* fetch values */  
    while (mysqli_stmt_fetch($stmt)) {  
        printf ("%s (%s)\n", $name, $code);  
    }  
  
    /* close statement */  
    mysqli_stmt_close($stmt);  
}  
  
/* close connection */  
mysqli_close($link);  
?>
```

The above examples will output:

```
Rockford (USA)  
Tallahassee (USA)  
Salinas (USA)  
Santa Clarita (USA)  
Springfield (USA)
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_errno](#)  
[mysqli\\_stmt\\_error](#)  
[mysqli\\_stmt\\_bind\\_result](#)

### 21.9.3.10.13. [mysqli\\_stmt::\\$field\\_count](#), [mysqli\\_stmt\\_field\\_count](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$field\\_count](#)

[mysqli\\_stmt\\_field\\_count](#)

Returns the number of field in the given statement

### Description

Object oriented style

```
int  
mysqli_stmt->field_count ;
```

Procedural style

```
int mysqli_stmt_field_count(  
    mysqli_stmt stmt);
```



### Warning

This function is currently not documented; only its argument list is available.

**21.9.3.10.14. `mysqli_stmt::free_result`, `mysqli_stmt_free_result`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::free_result`

`mysqli_stmt_free_result`

Frees stored result memory for the given statement handle

**Description**

Object oriented style

```
void mysqli_stmt::free_result();
```

Procedural style

```
void mysqli_stmt_free_result(  
    mysqli_stmt stmt);
```

Frees the result memory associated with the statement, which was allocated by `mysqli_stmt_store_result`.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

**Return Values**

No value is returned.

**See Also**

`mysqli_stmt_store_result`

**21.9.3.10.15. `mysqli_stmt::get_result`, `mysqli_stmt_get_result`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::get_result`

`mysqli_stmt_get_result`

Gets a result set from a prepared statement

**Description**

Object oriented style

```
mysqli_result mysqli_stmt::get_result();
```

Procedural style

```
mysqli_result mysqli_stmt_get_result(  
    mysqli_stmt stmt);
```

Call to return a result set from a prepared statement query.

**Parameters**

*stmt*Procedural style only: A statement identifier returned by *mysqli\_stmt\_init*.**Return Values**Returns a resultset or **FALSE** on failure.**MySQL Native Driver Only**Available only with *mysqlnd*.**Examples****Example 21.167. Object oriented style**

```
<?php
$mysqli = new mysqli("127.0.0.1", "user", "password", "world");

if($mysqli->connect_error)
{
    die("$mysqli->connect_errno: $mysqli->connect_error");
}

$query = "SELECT Name, Population, Continent FROM Country WHERE Continent=? ORDER BY Name LIMIT 1";

$stmt = $mysqli->stmt_init();
if(!$stmt->prepare($query))
{
    print "Failed to prepare statement\n";
}
else
{
    $stmt->bind_param("s", $continent);

    $continent_array = array('Europe', 'Africa', 'Asia', 'North America');

    foreach($continent_array as $continent)
    {
        $stmt->execute();
        $result = $stmt->get_result();
        while ($row = $result->fetch_array(MYSQLI_NUM))
        {
            foreach ($row as $r)
            {
                print "$r ";
            }
            print "\n";
        }
    }

    $stmt->close();
    $mysqli->close();
?>
```

**Example 21.168. Procedural style**

```
<?php

$link = mysqli_connect("127.0.0.1", "user", "password", "world");

if (!$link)
{
```

```
$error = mysqli_connect_error();
$errorno = mysqli_connect_errno();
print "$errorno: $error\n";
exit();
}

$query = "SELECT Name, Population, Continent FROM Country WHERE Continent=? ORDER BY Name LIMIT 1";

$stmt = mysqli_stmt_init($link);
if(!mysqli_stmt_prepare($stmt, $query))
{
    print "Failed to prepare statement\n";
}
else
{
    mysqli_stmt_bind_param($stmt, "s", $continent);

    $continent_array = array('Europe','Africa','Asia','North America');

    foreach($continent_array as $continent)
    {
        mysqli_stmt_execute($stmt);
        $result = mysqli_stmt_get_result($stmt);
        while ($row = mysqli_fetch_array($result, MYSQLI_NUM))
        {
            foreach ($row as $r)
            {
                print "$r ";
            }
            print "\n";
        }
    }
}
mysqli_stmt_close($stmt);
mysqli_close($link);
?>
```

The above examples will output:

```
Albania 3401200 Europe
Algeria 31471000 Africa
Afghanistan 22720000 Asia
Anguilla 8000 North America
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_result\\_metadata](#)  
[mysqli\\_stmt\\_fetch](#)  
[mysqli\\_fetch\\_array](#)  
[mysqli\\_stmt\\_store\\_result](#)

#### 21.9.3.10.16. [mysqli\\_stmt::get\\_warnings](#), [mysqli\\_stmt\\_get\\_warnings](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::get\\_warnings](#)

[mysqli\\_stmt\\_get\\_warnings](#)

Get result of SHOW WARNINGS

### Description



## Object oriented style

```
object mysqli_stmt::get_warnings(
    mysqli_stmt $stmt);
```

## Procedural style

```
object mysqli_stmt_get_warnings(
    mysqli_stmt $stmt);
```

**Warning**

This function is currently not documented; only its argument list is available.

**21.9.3.10.17. *mysqli\_stmt::\$insert\_id*, *mysqli\_stmt\_insert\_id***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli\_stmt::\$insert\_id*  
*mysqli\_stmt\_insert\_id*

Get the ID generated from the previous INSERT operation

**Description**

## Object oriented style

```
int
mysqli_stmt->insert_id ;
```

## Procedural style

```
mixed mysqli_stmt_insert_id(
    mysqli_stmt $stmt);
```

**Warning**

This function is currently not documented; only its argument list is available.

**21.9.3.10.18. *mysqli\_stmt::\$more\_results*, *mysqli\_stmt\_more\_results***

Copyright 1997-2012 the PHP Documentation Group. [2906]

- *mysqli\_stmt::\$more\_results*  
*mysqli\_stmt\_more\_results*

Check if there are more query results from a multiple query

**Description**

## Object oriented style (method):

```
public bool mysqli_stmt::$more_results();
```

## Procedural style:

```
bool mysqli_stmt_more_results(
    mysqli_stmt $stmt);
```

Checks if there are more query results from a multiple query.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

**Return Values**

Returns `TRUE` if more results exist, otherwise `FALSE` .

**See Also**

`mysqli_stmt::next_result`  
`mysqli::multi_query`

**21.9.3.10.19. `mysqli_stmt::next_result`, `mysqli_stmt_next_result`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::next_result`

`mysqli_stmt_next_result`

Reads the next result from a multiple query

**Description**

Object oriented style (method):

```
public bool mysqli_stmt::next_result();
```

Procedural style:

```
bool mysqli_stmt_next_result(
    mysqli_stmt stmt);
```

Reads the next result from a multiple query.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**Errors/Exceptions**

Emits an `E_STRICT` level error if a result set does not exist, and suggests using `mysqli_stmt::more_results` in these cases, before calling `mysqli_stmt::next_result`.

**See Also**

`mysqli_stmt::more_results`  
`mysqli::multi_query`

**21.9.3.10.20. `mysqli_stmt::$num_rows`, `mysqli_stmt_num_rows`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::$num_rows`

`mysqli_stmt_num_rows`

Return the number of rows in statements result set

## Description

Object oriented style

```
int
mysqli_stmt->num_rows ;
```

Procedural style

```
int mysqli_stmt_num_rows(
mysqli_stmt stmt);
```

Returns the number of rows in the result set. The use of `mysqli_stmt_num_rows` depends on whether or not you used `mysqli_stmt_store_result` to buffer the entire result set in the statement handle.

If you use `mysqli_stmt_store_result`, `mysqli_stmt_num_rows` may be called immediately.

## Parameters

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

## Return Values

An integer representing the number of rows in result set.

## Examples

### Example 21.169. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name LIMIT 20";
if ($stmt = $mysqli->prepare($query)) {

    /* execute query */
    $stmt->execute();

    /* store result */
    $stmt->store_result();

    printf("Number of rows: %d.\n", $stmt->num_rows);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

**Example 21.170. Procedural style**

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name LIMIT 20";
if ($stmt = mysqli_prepare($link, $query)) {

    /* execute query */
    mysqli_stmt_execute($stmt);

    /* store result */
    mysqli_stmt_store_result($stmt);

    printf("Number of rows: %d.\n", mysqli_stmt_num_rows($stmt));

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Number of rows: 20.
```

**See Also**

[mysqli\\_stmt\\_affected\\_rows](#)  
[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_store\\_result](#)

**21.9.3.10.21. [mysqli\\_stmt::\\$param\\_count](#), [mysqli\\_stmt\\_param\\_count](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$param\\_count](#)

[mysqli\\_stmt\\_param\\_count](#)

Returns the number of parameter for the given statement

**Description**

Object oriented style

```
int
mysqli_stmt->param_count ;
```

Procedural style

```
int mysqli_stmt_param_count(
    mysqli_stmt stmt);
```

Returns the number of parameter markers present in the prepared statement.

### Parameters

*stmt* Procedural style only: A statement identifier returned by *mysqli\_stmt\_init*.

### Return Values

Returns an integer representing the number of parameters.

### Examples

#### Example 21.171. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($stmt = $mysqli->prepare("SELECT Name FROM Country WHERE Name=? OR Code=?")) {

    $marker = $stmt->param_count;
    printf("Statement has %d markers.\n", $marker);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

#### Example 21.172. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($stmt = mysqli_prepare($link, "SELECT Name FROM Country WHERE Name=? OR Code=?")) {

    $marker = mysqli_stmt_param_count($stmt);
    printf("Statement has %d markers.\n", $marker);

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Statement has 2 markers.
```

### See Also

`mysqli_prepare`

#### 21.9.3.10.22. `mysqli_stmt::prepare`, `mysqli_stmt_prepare`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::prepare`

`mysqli_stmt_prepare`

Prepare an SQL statement for execution

### Description

Object oriented style

```
mixed mysqli_stmt::prepare(
    string query);
```

Procedural style

```
bool mysqli_stmt_prepare(
    mysqli_stmt stmt,
    string query);
```

Prepares the SQL query pointed to by the null-terminated string query.

The parameter markers must be bound to application variables using `mysqli_stmt_bind_param` and/or `mysqli_stmt_bind_result` before executing the statement or fetching rows.



#### Note

In the case where you pass a statement to `mysqli_stmt_prepare` that is longer than `max_allowed_packet` of the server, the returned error codes are different depending on whether you are using MySQL Native Driver (`mysqlnd`) or MySQL Client Library (`libmysqlclient`). The behavior is as follows:

- `mysqlnd` on Linux returns an error code of 1153. The error message means “got a packet bigger than `max_allowed_packet` bytes”.
- `mysqlnd` on Windows returns an error code 2006. This error message means “server has gone away”.
- `libmysqlclient` on all platforms returns an error code 2006. This error message means “server has gone away”.

### Parameters

`stmt`

Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

`query`

The query, as a string. It must consist of a single SQL statement.

You can include one or more parameter markers in the SQL statement by embedding question mark (?) characters at the appropriate positions.

**Note**

You should not add a terminating semicolon or `\g` to the statement.

**Note**

The markers are legal only in certain places in SQL statements. For example, they are allowed in the `VALUES()` list of an `INSERT` statement (to specify column values for a row), or in a comparison with a column in a `WHERE` clause to specify a comparison value.

However, they are not allowed for identifiers (such as table or column names), in the select list that names the columns to be returned by a `SELECT` statement), or to specify both operands of a binary operator such as the `=` equal sign. The latter restriction is necessary because it would be impossible to determine the parameter type. In general, parameters are legal only in Data Manipulation Language (DML) statements, and not in Data Definition Language (DDL) statements.

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**Examples****Example 21.173. Object oriented style**

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$city = "Amersfoort";

/* create a prepared statement */
$stmt = $mysqli->stmt_init();
if ($stmt->prepare("SELECT District FROM City WHERE Name=?")) {

    /* bind parameters for markers */
    $stmt->bind_param("s", $city);

    /* execute query */
    $stmt->execute();

    /* bind result variables */
```

```
$stmt->bind_result($district);

/* fetch value */
$stmt->fetch();

printf("%s is in district %s\n", $city, $district);

/* close statement */
$stmt->close();
}

/* close connection */
mysqli->close();
?>
```

### Example 21.174. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$city = "Amersfoort";

/* create a prepared statement */
$stmt = mysqli_stmt_init($link);
if (mysqli_stmt_prepare($stmt, 'SELECT District FROM City WHERE Name=?')) {

    /* bind parameters for markers */
    mysqli_stmt_bind_param($stmt, "s", $city);

    /* execute query */
    mysqli_stmt_execute($stmt);

    /* bind result variables */
    mysqli_stmt_bind_result($stmt, $district);

    /* fetch value */
    mysqli_stmt_fetch($stmt);

    printf("%s is in district %s\n", $city, $district);

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Amersfoort is in district Utrecht
```

### See Also

[mysqli\\_stmt\\_init](#)  
[mysqli\\_stmt\\_execute](#)



```
mysqli_stmt_fetch  
mysqli_stmt_bind_param  
mysqli_stmt_bind_result  
mysqli_stmt_close
```

### 21.9.3.10.23. `mysqli_stmt::reset`, `mysqli_stmt_reset`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::reset`

```
mysqli_stmt_reset
```

Resets a prepared statement

#### Description

Object oriented style

```
bool mysqli_stmt::reset();
```

Procedural style

```
bool mysqli_stmt_reset(  
    mysqli_stmt stmt);
```

Resets a prepared statement on client and server to state after prepare.

It resets the statement on the server, data sent using `mysqli_stmt_send_long_data`, unbuffered result sets and current errors. It does not clear bindings or stored result sets. Stored result sets will be cleared when executing the prepared statement (or closing it).

To prepare a statement with another query use function `mysqli_stmt_prepare`.

#### Parameters

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

#### Return Values

Returns `TRUE` on success or `FALSE` on failure.

#### See Also

`mysqli_prepare`

### 21.9.3.10.24. `mysqli_stmt::result_metadata`, `mysqli_stmt_result_metadata`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_stmt::result_metadata`

```
mysqli_stmt_result_metadata
```

Returns result set metadata from a prepared statement

#### Description

Object oriented style

```
mysqli_result mysqli_stmt::result_metadata();
```

## Procedural style

```
mysqli_result mysqli_stmt_result_metadata(
    mysqli_stmt stmt);
```

If a statement passed to `mysqli_prepare` is one that produces a result set, `mysqli_stmt_result_metadata` returns the result object that can be used to process the meta information such as total number of fields and individual field information.

**Note**

This result set pointer can be passed as an argument to any of the field-based functions that process result set metadata, such as:

- `mysqli_num_fields`
- `mysqli_fetch_field`
- `mysqli_fetch_field_direct`
- `mysqli_fetch_fields`
- `mysqli_field_count`
- `mysqli_field_seek`
- `mysqli_field_tell`
- `mysqli_free_result`

The result set structure should be freed when you are done with it, which you can do by passing it to `mysqli_free_result`

**Note**

The result set returned by `mysqli_stmt_result_metadata` contains only metadata. It does not contain any row results. The rows are obtained by using the statement handle with `mysqli_stmt_fetch`.

**Parameters**

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

**Return Values**

Returns a result object or `FALSE` if an error occurred.

**Examples****Example 21.175. Object oriented style**

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "test");

$mysqli->query("DROP TABLE IF EXISTS friends");
$mysqli->query("CREATE TABLE friends (id int, name varchar(20))");

$mysqli->query("INSERT INTO friends VALUES (1, 'Hartmut'), (2, 'Ulf')");

$stmt = $mysqli->prepare("SELECT id, name FROM friends");
$stmt->execute();
```

```
/* get resultset for metadata */
$result = $stmt->result_metadata();

/* retrieve field information from metadata result set */
$field = $result->fetch_field();

printf("Fieldname: %s\n", $field->name);

/* close resultset */
$result->close();

/* close connection */
mysqli->close();
?>
```

### Example 21.176. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "test");

mysqli_query($link, "DROP TABLE IF EXISTS friends");
mysqli_query($link, "CREATE TABLE friends (id int, name varchar(20))");
mysqli_query($link, "INSERT INTO friends VALUES (1,'Hartmut'), (2, 'Ulf')");

$stmt = mysqli_prepare($link, "SELECT id, name FROM friends");
mysqli_stmt_execute($stmt);

/* get resultset for metadata */
$result = mysqli_stmt_result_metadata($stmt);

/* retrieve field information from metadata result set */
$field = mysqli_fetch_field($result);

printf("Fieldname: %s\n", $field->name);

/* close resultset */
mysqli_free_result($result);

/* close connection */
mysqli_close($link);
?>
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_free\\_result](#)

### 21.9.3.10.25. [mysqli\\_stmt::send\\_long\\_data](#), [mysqli\\_stmt\\_send\\_long\\_data](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::send\\_long\\_data](#)

[mysqli\\_stmt\\_send\\_long\\_data](#)

Send data in blocks

### Description

Object oriented style

```
bool mysqli_stmt::send_long_data(
    int param_nr,
```

```
string data);
```

### Procedural style

```
bool mysqli_stmt_send_long_data(  
    mysqli_stmt stmt,  
    int param_nr,  
    string data);
```

Allows to send parameter data to the server in pieces (or chunks), e.g. if the size of a blob exceeds the size of [max\\_allowed\\_packet](#). This function can be called multiple times to send the parts of a character or binary data value for a column, which must be one of the TEXT or BLOB datatypes.

### Parameters

<i>stmt</i>	Procedural style only: A statement identifier returned by <a href="#">mysqli_stmt_init</a> .
<i>param_nr</i>	Indicates which parameter to associate the data with. Parameters are numbered beginning with 0.
<i>data</i>	A string containing data to be sent.

### Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

### Examples

#### Example 21.177. Object oriented style

```
<?php  
$stmt = $mysqli->prepare("INSERT INTO messages (message) VALUES (?)");  
$null = NULL;  
$stmt->bind_param("b", $null);  
$fp = fopen("messages.txt", "r");  
while (!feof($fp)) {  
    $stmt->send_long_data(0, fread($fp, 8192));  
}  
fclose($fp);  
$stmt->execute();  
?>
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_bind\\_param](#)

#### 21.9.3.10.26. [mysqli\\_stmt::\\$sqlstate](#), [mysqli\\_stmt\\_sqlstate](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::\\$sqlstate](#)

[mysqli\\_stmt\\_sqlstate](#)

Returns SQLSTATE error from previous statement operation

### Description

Object oriented style

```
string
mysqli_stmt->sqlstate ;
```

### Procedural style

```
string mysqli_stmt_sqlstate(
    mysqli_stmt stmt);
```

Returns a string containing the SQLSTATE error code for the most recently invoked prepared statement function that can succeed or fail. The error code consists of five characters. '00000' means no error. The values are specified by ANSI SQL and ODBC. For a list of possible values, see <http://dev.mysql.com/doc/mysq/en/error-handling.html>.

### Parameters

*stmt* Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

### Return Values

Returns a string containing the SQLSTATE error code for the last error. The error code consists of five characters. '00000' means no error.

### Notes



#### Note

Note that not all MySQL errors are yet mapped to SQLSTATE's. The value `HY000` (general error) is used for unmapped errors.

### Examples

#### Example 21.178. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$mysqli->query("CREATE TABLE myCountry LIKE Country");
$mysqli->query("INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = $mysqli->prepare($query)) {

    /* drop table */
    $mysqli->query("DROP TABLE myCountry");

    /* execute query */
    $stmt->execute();

    printf("Error: %s.\n", $stmt->sqlstate);

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

**Example 21.179. Procedural style**

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

mysqli_query($link, "CREATE TABLE myCountry LIKE Country");
mysqli_query($link, "INSERT INTO myCountry SELECT * FROM Country");

$query = "SELECT Name, Code FROM myCountry ORDER BY Name";
if ($stmt = mysqli_prepare($link, $query)) {

    /* drop table */
    mysqli_query($link, "DROP TABLE myCountry");

    /* execute query */
    mysqli_stmt_execute($stmt);

    printf("Error: %s.\n", mysqli_stmt_sqlstate($stmt));

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Error: 42S02.
```

**See Also**

[mysqli\\_stmt\\_errno](#)  
[mysqli\\_stmt\\_error](#)

**21.9.3.10.27. [mysqli\\_stmt::store\\_result](#), [mysqli\\_stmt\\_store\\_result](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_stmt::store\\_result](#)

[mysqli\\_stmt\\_store\\_result](#)

Transfers a result set from a prepared statement

**Description**

Object oriented style

```
bool mysqli_stmt::store_result();
```

## Procedural style

```
bool mysqli_stmt_store_result(
    mysqli_stmt stmt);
```

You must call `mysqli_stmt_store_result` for every query that successfully produces a result set (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`), and only if you want to buffer the complete result set by the client, so that the subsequent `mysqli_stmt_fetch` call returns buffered data.



### Note

It is unnecessary to call `mysqli_stmt_store_result` for other queries, but if you do, it will not harm or cause any notable performance in all cases. You can detect whether the query produced a result set by checking if `mysqli_stmt_result_metadata` returns `NULL`.

## Parameters

*stmt*

Procedural style only: A statement identifier returned by `mysqli_stmt_init`.

## Return Values

Returns `TRUE` on success or `FALSE` on failure.

## Examples

### Example 21.180. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name LIMIT 20";
if ($stmt = $mysqli->prepare($query)) {

    /* execute query */
    $stmt->execute();

    /* store result */
    $stmt->store_result();

    printf("Number of rows: %d.\n", $stmt->num_rows);

    /* free result */
    $stmt->free_result();

    /* close statement */
    $stmt->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.181. Procedural style

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name LIMIT 20";
if ($stmt = mysqli_prepare($link, $query)) {

    /* execute query */
    mysqli_stmt_execute($stmt);

    /* store result */
    mysqli_stmt_store_result($stmt);

    printf("Number of rows: %d.\n", mysqli_stmt_num_rows($stmt));

    /* free result */
    mysqli_stmt_free_result($stmt);

    /* close statement */
    mysqli_stmt_close($stmt);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Number of rows: 20.
```

### See Also

[mysqli\\_prepare](#)  
[mysqli\\_stmt\\_result\\_metadata](#)  
[mysqli\\_stmt\\_fetch](#)

### 21.9.3.11. The `mysqli_result` class ([mysqli\\_result](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

Represents the result set obtained from a query against the database.

*Changelog*

**Table 21.52. Changelog**

Version	Description
5.4.0	<a href="#">Iterator</a> support was added, as <a href="#">mysqli_result</a> now implements <a href="#">Traversable</a> .

```
mysqli_result {
    mysqli_result
```



```

    Traversable

    Properties

    int
        mysqli_result->current_field ;

    int
        mysqli_result->field_count ;

    array
        mysqli_result->lengths ;

    int
        mysqli_result->num_rows ;

Methods

bool mysqli_result::data_seek(
    int offset);

mixed mysqli_result::fetch_all(
    int resulttype
        = =MYSQLI_NUM);

mixed mysqli_result::fetch_array(
    int resulttype
        = =MYSQLI_BOTH);

array mysqli_result::fetch_assoc();

object mysqli_result::fetch_field_direct(
    int fieldnr);

object mysqli_result::fetch_field();

array mysqli_result::fetch_fields();

object mysqli_result::fetch_object(
    string class_name,
    array params);

mixed mysqli_result::fetch_row();

bool mysqli_result::field_seek(
    int fieldnr);

void mysqli_result::free();
}

```

#### 21.9.3.11.1. `mysqli_result::$current_field`, `mysqli_field_tell`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::$current_field`  
`mysqli_field_tell`

Get current field offset of a result pointer

#### Description

Object oriented style

```

int
    mysqli_result->current_field ;

```

## Procedural style

```
int mysqli_field_tell(
    mysqli_result result);
```

Returns the position of the field cursor used for the last `mysqli_fetch_field` call. This value can be used as an argument to `mysqli_field_seek`.

## Parameters

*result* Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

## Return Values

Returns current offset of field cursor.

## Examples

### Example 21.182. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = $mysqli->query($query)) {

    /* Get field information for all columns */
    while ($finfo = $result->fetch_field()) {

        /* get fieldpointer offset */
        $currentfield = $result->current_field;

        printf("Column %d:\n", $currentfield);
        printf("Name:      %s\n", $finfo->name);
        printf("Table:     %s\n", $finfo->table);
        printf("max. Len:  %d\n", $finfo->max_length);
        printf("Flags:     %d\n", $finfo->flags);
        printf("Type:      %d\n\n", $finfo->type);
    }
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.183. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
```

```
        exit();
    }

    $query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

    if ($result = mysqli_query($link, $query)) {

        /* Get field information for all fields */
        while ($finfo = mysqli_fetch_field($result)) {

            /* get fieldpointer offset */
            $currentfield = mysqli_field_tell($result);

            printf("Column %d:\n", $currentfield);
            printf("Name:      %s\n", $finfo->name);
            printf("Table:     %s\n", $finfo->table);
            printf("max. Len:  %d\n", $finfo->max_length);
            printf("Flags:     %d\n", $finfo->flags);
            printf("Type:      %d\n\n", $finfo->type);
        }
        mysqli_free_result($result);
    }

    /* close connection */
    mysqli_close($link);
?>
```

The above examples will output:

```
Column 1:
Name:      Name
Table:     Country
max. Len:  11
Flags:     1
Type:      254

Column 2:
Name:      SurfaceArea
Table:     Country
max. Len:  10
Flags:     32769
Type:      4
```

## See Also

[mysqli\\_fetch\\_field](#)  
[mysqli\\_field\\_seek](#)

### 21.9.3.11.2. [mysqli\\_result::data\\_seek](#), [mysqli\\_data\\_seek](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::data\\_seek](#)

[mysqli\\_data\\_seek](#)

Adjusts the result pointer to an arbitrary row in the result

## Description

Object oriented style

```
bool mysqli_result::data_seek(
```

```
int offset);
```

### Procedural style

```
bool mysqli_data_seek(
    mysqli_result result,
    int offset);
```

The `mysqli_data_seek` function seeks to an arbitrary result pointer specified by the `offset` in the result set.

### Parameters

*result*

Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

*offset*

The field offset. Must be between zero and the total number of rows minus one (`0..mysqli_num_rows - 1`).

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Notes



#### Note

This function can only be used with buffered results attained from the use of the `mysqli_store_result` or `mysqli_query` functions.

### Examples

#### Example 21.184. Object oriented style

```
<?php
/* Open a connection */
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name";
if ($result = $mysqli->query( $query)) {

    /* seek to row no. 400 */
    $result->data_seek(399);

    /* fetch row */
    $row = $result->fetch_row();

    printf ("City: %s Countrycode: %s\n", $row[0], $row[1]);

    /* free result set*/
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

**Example 21.185. Procedural style**

```
<?php
/* Open a connection */
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (!$link) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER BY Name";

if ($result = mysqli_query($link, $query)) {

    /* seek to row no. 400 */
    mysqli_data_seek($result, 399);

    /* fetch row */
    $row = mysqli_fetch_row($result);

    printf ("City: %s Countrycode: %s\n", $row[0], $row[1]);

    /* free result set*/
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
City: Benin City Countrycode: NGA
```

**See Also**

[mysqli\\_store\\_result](#)  
[mysqli\\_fetch\\_row](#)  
[mysqli\\_fetch\\_array](#)  
[mysqli\\_fetch\\_assoc](#)  
[mysqli\\_fetch\\_object](#)  
[mysqli\\_query](#)  
[mysqli\\_num\\_rows](#)

**21.9.3.11.3. [mysqli\\_result::fetch\\_all](#), [mysqli\\_fetch\\_all](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_all](#)

[mysqli\\_fetch\\_all](#)

Fetches all result rows as an associative array, a numeric array, or both

**Description**

Object oriented style

```
mixed mysqli_result::fetch_all(
```

```
int resulttype
= =MYSQLI_NUM);
```

### Procedural style

```
mixed mysqli_fetch_all(
    mysqli_result result,
    int resulttype
    = =MYSQLI_NUM);
```

[mysqli\\_fetch\\_all](#) fetches all result rows and returns the result set as an associative array, a numeric array, or both.

### Parameters

*result*

Procedural style only: A result set identifier returned by [mysqli\\_query](#), [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#).

*resulttype*

This optional parameter is a constant indicating what type of array should be produced from the current row data. The possible values for this parameter are the constants [MYSQLI\\_ASSOC](#) , [MYSQLI\\_NUM](#) , or [MYSQLI\\_BOTH](#) .

### Return Values

Returns an array of associative or numeric arrays holding result rows.

### MySQL Native Driver Only

Available only with [mysqlnd](#).

As [mysqli\\_fetch\\_all](#) returns all the rows as an array in a single step, it may consume more memory than some similar functions such as [mysqli\\_fetch\\_array](#), which only returns one row at a time from the result set. Further, if you need to iterate over the result set, you will need a looping construct that will further impact performance. For these reasons [mysqli\\_fetch\\_all](#) should only be used in those situations where the fetched result set will be sent to another layer for processing.

### See Also

[mysqli\\_fetch\\_array](#)  
[mysqli\\_query](#)

#### 21.9.3.11.4. [mysqli\\_result::fetch\\_array](#), [mysqli\\_fetch\\_array](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_array](#)

[mysqli\\_fetch\\_array](#)

Fetch a result row as an associative, a numeric array, or both

### Description

Object oriented style

```
mixed mysqli_result::fetch_array(
    int resulttype
    = =MYSQLI_BOTH);
```

Procedural style

```
mixed mysqli_fetch_array(
    mysqli_result result,
    int resulttype
    = =MYSQLI_BOTH);
```

Returns an array that corresponds to the fetched row or `NULL` if there are no more rows for the resultset represented by the `result` parameter.

`mysqli_fetch_array` is an extended version of the `mysqli_fetch_row` function. In addition to storing the data in the numeric indices of the result array, the `mysqli_fetch_array` function can also store the data in associative indices, using the field names of the result set as keys.



#### Note

Field names returned by this function are *case-sensitive*.



#### Note

This function sets `NULL` fields to the PHP `NULL` value.

If two or more columns of the result have the same field names, the last column will take precedence and overwrite the earlier data. In order to access multiple columns with the same name, the numerically indexed version of the row must be used.

### Parameters

`result`

Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

`resulttype`

This optional parameter is a constant indicating what type of array should be produced from the current row data. The possible values for this parameter are the constants `MYSQLI_ASSOC`, `MYSQLI_NUM`, or `MYSQLI_BOTH`.

By using the `MYSQLI_ASSOC` constant this function will behave identically to the `mysqli_fetch_assoc`, while `MYSQLI_NUM` will behave identically to the `mysqli_fetch_row` function. The final option `MYSQLI_BOTH` will create a single array with the attributes of both.

### Return Values

Returns an array of strings that corresponds to the fetched row or `NULL` if there are no more rows in resultset.

### Examples

#### Example 21.186. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID LIMIT 3";
```

```

$result = $mysqli->query($query);

/* numeric array */
$row = $result->fetch_array(MYSQLI_NUM);
printf ("%s (%s)\n", $row[0], $row[1]);

/* associative array */
$row = $result->fetch_array(MYSQLI_ASSOC);
printf ("%s (%s)\n", $row["Name"], $row["CountryCode"]);

/* associative and numeric array */
$row = $result->fetch_array(MYSQLI_BOTH);
printf ("%s (%s)\n", $row[0], $row["CountryCode"]);

/* free result set */
$result->free();

/* close connection */
$mysqli->close();
?>

```

**Example 21.187. Procedural style**

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID LIMIT 3";
$result = mysqli_query($link, $query);

/* numeric array */
$row = mysqli_fetch_array($result, MYSQLI_NUM);
printf ("%s (%s)\n", $row[0], $row[1]);

/* associative array */
$row = mysqli_fetch_array($result, MYSQLI_ASSOC);
printf ("%s (%s)\n", $row["Name"], $row["CountryCode"]);

/* associative and numeric array */
$row = mysqli_fetch_array($result, MYSQLI_BOTH);
printf ("%s (%s)\n", $row[0], $row["CountryCode"]);

/* free result set */
mysqli_free_result($result);

/* close connection */
mysqli_close($link);
?>

```

The above examples will output:

```

Kabul (AFG)
Qandahar (AFG)
Herat (AFG)

```

**See Also**

[mysqli\\_fetch\\_assoc](#)



```
mysqli_fetch_row
mysqli_fetch_object
mysqli_query
mysqli_data_seek
```

### 21.9.3.11.5. `mysqli_result::fetch_assoc`, `mysqli_fetch_assoc`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::fetch_assoc`

```
mysqli_fetch_assoc
```

Fetch a result row as an associative array

#### Description

Object oriented style

```
array mysqli_result::fetch_assoc();
```

Procedural style

```
array mysqli_fetch_assoc(
    mysqli_result result);
```

Returns an associative array that corresponds to the fetched row or `NULL` if there are no more rows.



#### Note

Field names returned by this function are *case-sensitive*.



#### Note

This function sets `NULL` fields to the PHP `NULL` value.

#### Parameters

*result* Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

#### Return Values

Returns an associative array of strings representing the fetched row in the result set, where each key in the array represents the name of one of the result set's columns or `NULL` if there are no more rows in resultset.

If two or more columns of the result have the same field names, the last column will take precedence. To access the other column(s) of the same name, you either need to access the result with numeric indices by using `mysqli_fetch_row` or add alias names.

#### Examples

##### Example 21.188. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");
```

```
/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";

if ($result = $mysqli->query($query)) {

    /* fetch associative array */
    while ($row = $result->fetch_assoc()) {
        printf ("%s (%s)\n", $row["Name"], $row["CountryCode"]);
    }

    /* free result set */
    $result->free();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.189. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";

if ($result = mysqli_query($link, $query)) {

    /* fetch associative array */
    while ($row = mysqli_fetch_assoc($result)) {
        printf ("%s (%s)\n", $row["Name"], $row["CountryCode"]);
    }

    /* free result set */
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Pueblo (USA)
Arvada (USA)
Cape Coral (USA)
Green Bay (USA)
Santa Clara (USA)
```

### Example 21.190. A `mysqli_result` example comparing `iterator` usage

```

<?php
$c = mysqli_connect('127.0.0.1','user', 'pass');

// Using iterators (support was added with PHP 5.4)
foreach ( $c->query('SELECT user,host FROM mysql.user') as $row ) {
    printf("%s'@'%s'\n", $row['user'], $row['host']);
}

echo "\n=====\n";

// Not using iterators
$result = $c->query('SELECT user,host FROM mysql.user');
while ($row = $result->fetch_assoc()) {
    printf("%s'@'%s'\n", $row['user'], $row['host']);
}

?>

```

The above example will output something similar to:

```

'root'@'192.168.1.1'
'root'@'127.0.0.1'
'dude'@'localhost'
'lebowski'@'localhost'

=====

'root'@'192.168.1.1'
'root'@'127.0.0.1'
'dude'@'localhost'
'lebowski'@'localhost'

```

### See Also

[mysqli\\_fetch\\_array](#)  
[mysqli\\_fetch\\_row](#)  
[mysqli\\_fetch\\_object](#)  
[mysqli\\_query](#)  
[mysqli\\_data\\_seek](#)

#### 21.9.3.11.6. `mysqli_result::fetch_field_direct`, `mysqli_fetch_field_direct`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::fetch_field_direct`

`mysqli_fetch_field_direct`

Fetch meta-data for a single field

### Description

Object oriented style

```
object mysqli_result::fetch_field_direct(
    int fieldnr);
```

Procedural style

```
object mysqli_fetch_field_direct(
    mysqli_result result,
    int fieldnr);
```

Returns an object which contains field definition information from the specified result set.

### Parameters

*result*

Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

*fieldnr*

The field number. This value must be in the range from 0 to `number of fields - 1`.

### Return Values

Returns an object which contains field definition information or `FALSE` if no field information for specified `fieldnr` is available.

**Table 21.53. Object attributes**

Attribute	Description
name	The name of the column
orgname	Original column name if an alias was specified
table	The name of the table this field belongs to (if not calculated)
orgtable	Original table name if an alias was specified
def	The default value for this field, represented as a string
max_length	The maximum width of the field for the result set.
length	The width of the field, as specified in the table definition.
charsetnr	The character set number for the field.
flags	An integer representing the bit-flags for the field.
type	The data type used for this field
decimals	The number of decimals used (for integer fields)

### Examples

#### Example 21.191. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Name LIMIT 5";

if ($result = $mysqli->query($query)) {

    /* Get field information for column 'SurfaceArea' */
    $finfo = $result->fetch_field_direct(1);

    printf("Name:      %s\n", $finfo->name);
    printf("Table:     %s\n", $finfo->table);
    printf("max. Len:  %d\n", $finfo->max_length);
    printf("Flags:     %d\n", $finfo->flags);
}
```

```
    printf("Type:      %d\n", $finfo->type);

    $result->close();
}

/* close connection */
mysqli->close();
?>
```

### Example 21.192. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Name LIMIT 5";

if ($result = mysqli_query($link, $query)) {

    /* Get field information for column 'SurfaceArea' */
    $finfo = mysqli_fetch_field_direct($result, 1);

    printf("Name:      %s\n", $finfo->name);
    printf("Table:      %s\n", $finfo->table);
    printf("max. Len:  %d\n", $finfo->max_length);
    printf("Flags:      %d\n", $finfo->flags);
    printf("Type:      %d\n", $finfo->type);

    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Name:      SurfaceArea
Table:      Country
max. Len:  10
Flags:      32769
Type:      4
```

### See Also

[mysqli\\_num\\_fields](#)  
[mysqli\\_fetch\\_field](#)  
[mysqli\\_fetch\\_fields](#)

#### 21.9.3.11.7. [mysqli\\_result::fetch\\_field](#), [mysqli\\_fetch\\_field](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_field](#)  
[mysqli\\_fetch\\_field](#)

Returns the next field in the result set

## Description

Object oriented style

```
object mysqli_result::fetch_field();
```

Procedural style

```
object mysqli_fetch_field(
    mysqli_result result);
```

Returns the definition of one column of a result set as an object. Call this function repeatedly to retrieve information about all columns in the result set.

## Parameters

*result* Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

## Return Values

Returns an object which contains field definition information or `FALSE` if no field information is available.

**Table 21.54. Object properties**

Property	Description
name	The name of the column
orgname	Original column name if an alias was specified
table	The name of the table this field belongs to (if not calculated)
orgtable	Original table name if an alias was specified
def	Reserved for default value, currently always ""
db	Database (since PHP 5.3.6)
catalog	The catalog name, always "def" (since PHP 5.3.6)
max_length	The maximum width of the field for the result set.
length	The width of the field, as specified in the table definition.
charsetnr	The character set number for the field.
flags	An integer representing the bit-flags for the field.
type	The data type used for this field
decimals	The number of decimals used (for integer fields)

## Examples

### Example 21.193. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
```

```

    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = $mysqli->query($query)) {

    /* Get field information for all columns */
    while ($finfo = $result->fetch_field()) {

        printf("Name:      %s\n", $finfo->name);
        printf("Table:     %s\n", $finfo->table);
        printf("max. Len:  %d\n", $finfo->max_length);
        printf("Flags:     %d\n", $finfo->flags);
        printf("Type:      %d\n\n", $finfo->type);
    }
    $result->close();
}

/* close connection */
$mysqli->close();
?>

```

**Example 21.194. Procedural style**

```

<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = mysqli_query($link, $query)) {

    /* Get field information for all fields */
    while ($finfo = mysqli_fetch_field($result)) {

        printf("Name:      %s\n", $finfo->name);
        printf("Table:     %s\n", $finfo->table);
        printf("max. Len:  %d\n", $finfo->max_length);
        printf("Flags:     %d\n", $finfo->flags);
        printf("Type:      %d\n\n", $finfo->type);
    }
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>

```

The above examples will output:

```

Name:      Name
Table:     Country
max. Len:  11
Flags:     1
Type:      254

Name:      SurfaceArea
Table:     Country
max. Len:  10

```

```
Flags: 32769
Type: 4
```

**See Also**

```
mysqli_num_fields
mysqli_fetch_field_direct
mysqli_fetch_fields
mysqli_field_seek
```

**21.9.3.11.8. [mysqli\\_result::fetch\\_fields](#), [mysqli\\_fetch\\_fields](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_fields](#)

```
mysqli_fetch_fields
```

Returns an array of objects representing the fields in a result set

**Description**

Object oriented style

```
array mysqli_result::fetch_fields();
```

Procedural style

```
array mysqli_fetch_fields(
    mysqli_result result);
```

This function serves an identical purpose to the [mysqli\\_fetch\\_field](#) function with the single difference that, instead of returning one object at a time for each field, the columns are returned as an array of objects.

**Parameters**

*result* Procedural style only: A result set identifier returned by [mysqli\\_query](#), [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#).

**Return Values**

Returns an array of objects which contains field definition information or [FALSE](#) if no field information is available.

**Table 21.55. Object properties**

Property	Description
name	The name of the column
orgname	Original column name if an alias was specified
table	The name of the table this field belongs to (if not calculated)
orgtable	Original table name if an alias was specified
max_length	The maximum width of the field for the result set.
length	The width of the field, as specified in the table definition.



Property	Description
charsetnr	The character set number for the field.
flags	An integer representing the bit-flags for the field.
type	The data type used for this field
decimals	The number of decimals used (for integer fields)

## Examples

### Example 21.195. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = $mysqli->query($query)) {

    /* Get field information for all columns */
    $finfo = $result->fetch_fields();

    foreach ($finfo as $val) {
        printf("Name:      %s\n", $val->name);
        printf("Table:      %s\n", $val->table);
        printf("max. Len: %d\n", $val->max_length);
        printf("Flags:      %d\n", $val->flags);
        printf("Type:       %d\n\n", $val->type);
    }
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.196. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = mysqli_query($link, $query)) {

    /* Get field information for all columns */
    $finfo = mysqli_fetch_fields($result);

    foreach ($finfo as $val) {
        printf("Name:      %s\n", $val->name);
        printf("Table:      %s\n", $val->table);
        printf("max. Len: %d\n", $val->max_length);
        printf("Flags:      %d\n", $val->flags);
        printf("Type:       %d\n\n", $val->type);
    }
}
```

```
}
mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Name:      Name
Table:     Country
max. Len:  11
Flags:     1
Type:      254

Name:      SurfaceArea
Table:     Country
max. Len:  10
Flags:     32769
Type:      4
```

### See Also

[mysqli\\_num\\_fields](#)  
[mysqli\\_fetch\\_field\\_direct](#)  
[mysqli\\_fetch\\_field](#)

#### 21.9.3.11.9. [mysqli\\_result::fetch\\_object](#), [mysqli\\_fetch\\_object](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_object](#)

[mysqli\\_fetch\\_object](#)

Returns the current row of a result set as an object

### Description

Object oriented style

```
object mysqli_result::fetch_object(
    string class_name,
    array params);
```

Procedural style

```
object mysqli_fetch_object(
    mysqli_result result,
    string class_name,
    array params);
```

The [mysqli\\_fetch\\_object](#) will return the current row result set as an object where the attributes of the object represent the names of the fields found within the result set.

Note that [mysqli\\_fetch\\_object](#) sets the properties of the object before calling the object constructor.

### Parameters

<i>result</i>	Procedural style only: A result set identifier returned by <code>mysqli_query</code> , <code>mysqli_store_result</code> or <code>mysqli_use_result</code> .
<i>class_name</i>	The name of the class to instantiate, set the properties of and return. If not specified, a <code>stdClass</code> object is returned.
<i>params</i>	An optional array of parameters to pass to the constructor for <i>class_name</i> objects.

## Return Values

Returns an object with string properties that corresponds to the fetched row or `NULL` if there are no more rows in resultset.



### Note

Field names returned by this function are *case-sensitive*.



### Note

This function sets `NULL` fields to the PHP `NULL` value.

## Changelog

Version	Description
5.0.0	Added the ability to return as a different object.

## Examples

### Example 21.197. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";

if ($result = $mysqli->query($query)) {

    /* fetch object array */
    while ($obj = $result->fetch_object()) {
        printf ("%s (%s)\n", $obj->Name, $obj->CountryCode);
    }

    /* free result set */
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.198. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";

if ($result = mysqli_query($link, $query)) {

    /* fetch associative array */
    while ($obj = mysqli_fetch_object($result)) {
        printf ("%s (%s)\n", $obj->Name, $obj->CountryCode);
    }

    /* free result set */
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Pueblo (USA)
Arvada (USA)
Cape Coral (USA)
Green Bay (USA)
Santa Clara (USA)
```

### See Also

[mysqli\\_fetch\\_array](#)  
[mysqli\\_fetch\\_assoc](#)  
[mysqli\\_fetch\\_row](#)  
[mysqli\\_query](#)  
[mysqli\\_data\\_seek](#)

#### 21.9.3.11.10. [mysqli\\_result::fetch\\_row](#), [mysqli\\_fetch\\_row](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::fetch\\_row](#)

[mysqli\\_fetch\\_row](#)

Get a result row as an enumerated array

### Description

Object oriented style

```
mixed mysqli_result::fetch_row();
```

Procedural style

```
mixed mysqli_fetch_row(
    mysqli_result result);
```

Fetches one row of data from the result set and returns it as an enumerated array, where each column is stored in an array offset starting from 0 (zero). Each subsequent call to this function will return the next row within the result set, or `NULL` if there are no more rows.

### Parameters

*result* Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

### Return Values

`mysqli_fetch_row` returns an array of strings that corresponds to the fetched row or `NULL` if there are no more rows in result set.



#### Note

This function sets NULL fields to the PHP `NULL` value.

### Examples

#### Example 21.199. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";

if ($result = $mysqli->query($query)) {

    /* fetch object array */
    while ($row = $result->fetch_row()) {
        printf ("%s (%s)\n", $row[0], $row[1]);
    }

    /* free result set */
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

#### Example 21.200. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, CountryCode FROM City ORDER by ID DESC LIMIT 50,5";
```

```
if ($result = mysqli_query($link, $query)) {  
  
    /* fetch associative array */  
    while ($row = mysqli_fetch_row($result)) {  
        printf ("%s (%s)\n", $row[0], $row[1]);  
    }  
  
    /* free result set */  
    mysqli_free_result($result);  
}  
  
/* close connection */  
mysqli_close($link);  
?>
```

The above examples will output:

```
Pueblo (USA)  
Arvada (USA)  
Cape Coral (USA)  
Green Bay (USA)  
Santa Clara (USA)
```

### See Also

[mysqli\\_fetch\\_array](#)  
[mysqli\\_fetch\\_assoc](#)  
[mysqli\\_fetch\\_object](#)  
[mysqli\\_query](#)  
[mysqli\\_data\\_seek](#)

### 21.9.3.11.11. [mysqli\\_result::\\$field\\_count](#), [mysqli\\_num\\_fields](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::\\$field\\_count](#)

[mysqli\\_num\\_fields](#)

Get the number of fields in a result

#### Description

Object oriented style

```
int  
mysqli_result->field_count ;
```

Procedural style

```
int mysqli_num_fields(  
    mysqli_result result);
```

Returns the number of fields from specified result set.

#### Parameters

*result*

Procedural style only: A result set identifier returned by [mysqli\\_query](#), [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#).

## Return Values

The number of fields from a result set.

## Examples

### Example 21.201. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($result = $mysqli->query("SELECT * FROM City ORDER BY ID LIMIT 1")) {

    /* determine number of fields in result set */
    $field_cnt = $result->field_count;

    printf("Result set has %d fields.\n", $field_cnt);

    /* close result set */
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.202. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($result = mysqli_query($link, "SELECT * FROM City ORDER BY ID LIMIT 1")) {

    /* determine number of fields in result set */
    $field_cnt = mysqli_num_fields($result);

    printf("Result set has %d fields.\n", $field_cnt);

    /* close result set */
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Result set has 5 fields.
```

## See Also

[mysqli\\_fetch\\_field](#)

### 21.9.3.11.12. [mysqli\\_result::field\\_seek](#), [mysqli\\_field\\_seek](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_result::field\\_seek](#)

[mysqli\\_field\\_seek](#)

Set result pointer to a specified field offset

## Description

Object oriented style

```
bool mysqli_result::field_seek(  
    int fieldnr);
```

Procedural style

```
bool mysqli_field_seek(  
    mysqli_result result,  
    int fieldnr);
```

Sets the field cursor to the given offset. The next call to [mysqli\\_fetch\\_field](#) will retrieve the field definition of the column associated with that offset.



### Note

To seek to the beginning of a row, pass an offset value of zero.

## Parameters

*result*

Procedural style only: A result set identifier returned by [mysqli\\_query](#), [mysqli\\_store\\_result](#) or [mysqli\\_use\\_result](#).

*fieldnr*

The field number. This value must be in the range from 0 to [number of fields](#) - 1.

## Return Values

Returns [TRUE](#) on success or [FALSE](#) on failure.

## Examples

### Example 21.203. Object oriented style

```
<?php  
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");  
  
/* check connection */  
if (mysqli_connect_errno()) {  
    printf("Connect failed: %s\n", mysqli_connect_error());  
    exit();  
}
```



```
$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = $mysqli->query($query)) {

    /* Get field information for 2nd column */
    $result->field_seek(1);
    $finfo = $result->fetch_field();

    printf("Name:      %s\n", $finfo->name);
    printf("Table:     %s\n", $finfo->table);
    printf("max. Len:  %d\n", $finfo->max_length);
    printf("Flags:     %d\n", $finfo->flags);
    printf("Type:      %d\n\n", $finfo->type);

    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.204. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT Name, SurfaceArea from Country ORDER BY Code LIMIT 5";

if ($result = mysqli_query($link, $query)) {

    /* Get field information for 2nd column */
    mysqli_field_seek($result, 1);
    $finfo = mysqli_fetch_field($result);

    printf("Name:      %s\n", $finfo->name);
    printf("Table:     %s\n", $finfo->table);
    printf("max. Len:  %d\n", $finfo->max_length);
    printf("Flags:     %d\n", $finfo->flags);
    printf("Type:      %d\n\n", $finfo->type);

    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Name:      SurfaceArea
Table:     Country
max. Len:  10
Flags:     32769
Type:      4
```

**See Also**

---

`mysqli_fetch_field`

### 21.9.3.11.13. `mysqli_result::free`, `mysqli_free_result`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::free`

`mysqli_free_result`

Frees the memory associated with a result

#### Description

Object oriented style

```
void mysqli_result::free();
```

```
void mysqli_result::close();
```

```
void mysqli_result::free_result();
```

Procedural style

```
void mysqli_free_result(  
    mysqli_result result);
```

Frees the memory associated with the result.



#### Note

You should always free your result with `mysqli_free_result`, when your result object is not needed anymore.

#### Parameters

*result*

Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

#### Return Values

No value is returned.

#### See Also

`mysqli_query`  
`mysqli_stmt_store_result`  
`mysqli_store_result`  
`mysqli_use_result`

### 21.9.3.11.14. `mysqli_result::$lengths`, `mysqli_fetch_lengths`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::$lengths`

`mysqli_fetch_lengths`

Returns the lengths of the columns of the current row in the result set

#### Description

## Object oriented style

```
array
mysqli_result->lengths ;
```

## Procedural style

```
array mysqli_fetch_lengths(
mysqli_result result);
```

The `mysqli_fetch_lengths` function returns an array containing the lengths of every column of the current row within the result set.

## Parameters

*result* Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

## Return Values

An array of integers representing the size of each column (not including any terminating null characters). `FALSE` if an error occurred.

`mysqli_fetch_lengths` is valid only for the current row of the result set. It returns `FALSE` if you call it before calling `mysqli_fetch_row/array/object` or after retrieving all rows in the result.

## Examples

### Example 21.205. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT * from Country ORDER BY Code LIMIT 1";

if ($result = $mysqli->query($query)) {

    $row = $result->fetch_row();

    /* display column lengths */
    foreach ($result->lengths as $i => $val) {
        printf("Field %2d has Length %2d\n", $i+1, $val);
    }
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.206. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");
```

```
/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$query = "SELECT * from Country ORDER BY Code LIMIT 1";

if ($result = mysqli_query($link, $query)) {

    $row = mysqli_fetch_row($result);

    /* display column lengths */
    foreach (mysqli_fetch_lengths($result) as $i => $val) {
        printf("Field %2d has Length %2d\n", $i+1, $val);
    }
    mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Field 1 has Length 3
Field 2 has Length 5
Field 3 has Length 13
Field 4 has Length 9
Field 5 has Length 6
Field 6 has Length 1
Field 7 has Length 6
Field 8 has Length 4
Field 9 has Length 6
Field 10 has Length 6
Field 11 has Length 5
Field 12 has Length 44
Field 13 has Length 7
Field 14 has Length 3
Field 15 has Length 2
```

#### 21.9.3.11.15. `mysqli_result::$num_rows, mysqli_num_rows`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_result::$num_rows`

`mysqli_num_rows`

Gets the number of rows in a result

#### Description

Object oriented style

```
int
mysqli_result->num_rows ;
```

Procedural style

```
int mysqli_num_rows(
    mysqli_result result);
```

Returns the number of rows in the result set.

The behaviour of `mysqli_num_rows` depends on whether buffered or unbuffered result sets are being used. For unbuffered result sets, `mysqli_num_rows` will not return the correct number of rows until all the rows in the result have been retrieved.

## Parameters

`result` Procedural style only: A result set identifier returned by `mysqli_query`, `mysqli_store_result` or `mysqli_use_result`.

## Return Values

Returns number of rows in the result set.



### Note

If the number of rows is greater than `MAXINT`, the number will be returned as a string.

## Examples

### Example 21.207. Object oriented style

```
<?php
$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($result = $mysqli->query("SELECT Code, Name FROM Country ORDER BY Name")) {

    /* determine number of rows result set */
    $row_cnt = $result->num_rows;

    printf("Result set has %d rows.\n", $row_cnt);

    /* close result set */
    $result->close();
}

/* close connection */
$mysqli->close();
?>
```

### Example 21.208. Procedural style

```
<?php
$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

if ($result = mysqli_query($link, "SELECT Code, Name FROM Country ORDER BY Name")) {

    /* determine number of rows result set */
    $row_cnt = mysqli_num_rows($result);
}
```

```
printf("Result set has %d rows.\n", $row_cnt);

/* close result set */
mysqli_free_result($result);
}

/* close connection */
mysqli_close($link);
?>
```

The above examples will output:

```
Result set has 239 rows.
```

### See Also

[mysqli\\_affected\\_rows](#)  
[mysqli\\_store\\_result](#)  
[mysqli\\_use\\_result](#)  
[mysqli\\_query](#)

### 21.9.3.12. The `mysqli_driver` class ([mysqli\\_driver](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

MySQLi Driver.

```
mysqli_driver {
    mysqli_driver

    Properties

    public readonly string
        client_info ;

    public readonly string
        client_version ;

    public readonly string
        driver_version ;

    public readonly string
        embedded ;

    public bool
        reconnect ;

    public int
        report_mode ;

    Methods

    void mysqli_driver::embedded_server_end();

    bool mysqli_driver::embedded_server_start(
        bool start,
        array arguments,
        array groups);
}
```

<code>client_info</code>	The Client API header version
<code>client_version</code>	The Client version
<code>driver_version</code>	The MySQLi Driver version
<code>embedded</code>	Whether MySQLi Embedded support is enabled
<code>reconnect</code>	Allow or prevent reconnect (see the <code>mysqli.reconnect</code> INI directive)
<code>report_mode</code>	Set to <code>MYSQLI_REPORT_OFF</code> , <code>MYSQLI_REPORT_ALL</code> or any combination of <code>MYSQLI_REPORT_STRICT</code> (throw Exceptions for errors), <code>MYSQLI_REPORT_ERROR</code> (report errors) and <code>MYSQLI_REPORT_INDEX</code> (errors regarding indexes). See also <code>mysqli_report</code> .

#### 21.9.3.12.1. `mysqli_driver::embedded_server_end`, `mysqli_embedded_server_end`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_driver::embedded_server_end`

`mysqli_embedded_server_end`

Stop embedded server

#### Description

Object oriented style

```
void mysqli_driver::embedded_server_end();
```

Procedural style

```
void mysqli_embedded_server_end();
```



#### Warning

This function is currently not documented; only its argument list is available.

#### 21.9.3.12.2. `mysqli_driver::embedded_server_start`, `mysqli_embedded_server_start`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_driver::embedded_server_start`

`mysqli_embedded_server_start`

Initialize and start embedded server

#### Description

Object oriented style

```
bool mysqli_driver::embedded_server_start(
    bool start,
    array arguments,
    array groups);
```

Procedural style

```
bool mysqli_embedded_server_start(
    bool start,
```

```
array arguments,
array groups);
```

**Warning**

This function is currently not documented; only its argument list is available.

**21.9.3.12.3. `mysqli_driver::$report_mode`, `mysqli_report`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_driver::$report_mode`

`mysqli_report`

Enables or disables internal report functions

**Description**

Object oriented style

```
int
mysqli_driver->report_mode ;
```

Procedural style

```
bool mysqli_report(
int flags);
```

A function helpful in improving queries during code development and testing. Depending on the flags, it reports errors from `mysqli` function calls or queries that don't use an index (or use a bad index).

**Parameters**

*flags*

**Table 21.56. Supported flags**

Name	Description
<code>MYSQLE_REPORT_OFF</code>	Turns reporting off
<code>MYSQLE_REPORT_ERROR</code>	Report errors from <code>mysqli</code> function calls
<code>MYSQLE_REPORT_STRICT</code>	Throw <code>mysqli_sql_exception</code> for errors instead of warnings
<code>MYSQLE_REPORT_INDEX</code>	Report if no index or bad index was used in a query
<code>MYSQLE_REPORT_ALL</code>	Set all options (report all)

**Return Values**

Returns `TRUE` on success or `FALSE` on failure.

**Changelog**

Version	Description
5.3.4	Changing the reporting mode is now be per-request, rather than per-process.
5.2.15	Changing the reporting mode is now be per-request, rather than per-process.



## Examples

### Example 21.209. Object oriented style

```
<?php

$mysqli = new mysqli("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* activate reporting */
$driver = new mysqli_driver();
$driver->report_mode = MYSQLI_REPORT_ALL;

try {

    /* this query should report an error */
    $result = $mysqli->query("SELECT Name FROM Nonexistingtable WHERE population > 50000");

    /* this query should report a bad index */
    $result = $mysqli->query("SELECT Name FROM City WHERE population > 50000");

    $result->close();

    $mysqli->close();

} catch (mysqli_sql_exception $e) {

    echo $e->__toString();
}
?>
```

### Example 21.210. Procedural style

```
<?php
/* activate reporting */
mysqli_report(MYSQLI_REPORT_ALL);

$link = mysqli_connect("localhost", "my_user", "my_password", "world");

/* check connection */
if (mysqli_connect_errno()) {
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

/* this query should report an error */
$result = mysqli_query("SELECT Name FROM Nonexistingtable WHERE population > 50000");

/* this query should report a bad index */
$result = mysqli_query("SELECT Name FROM City WHERE population > 50000");

mysqli_free_result($result);

mysqli_close($link);
?>
```

## See Also

[mysqli\\_debug](#)

```
mysqli_dump_debug_info  
mysqli_sql_exception  
set_exception_handler  
error_reporting
```

### 21.9.3.13. The `mysqli_warning` class (`mysqli_warning`)

Copyright 1997-2012 the PHP Documentation Group. [2906]

Represents a MySQL warning.

```
mysqli_warning {  
    mysqli_warning  
  
    Properties  
  
    public  
        message ;  
  
    public  
        sqlstate ;  
  
    public  
        errno ;  
Methods  
  
    public mysqli_warning::__construct();  
  
    public void mysqli_warning::next();  
}
```

<code>message</code>	Message string
<code>sqlstate</code>	SQL state
<code>errno</code>	Error number

#### 21.9.3.13.1. `mysqli_warning::__construct`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_warning::__construct`

The `__construct` purpose

#### Description

```
public mysqli_warning::__construct();
```



#### Warning

This function is currently not documented; only its argument list is available.

#### Parameters

This function has no parameters.

#### Return Values

### 21.9.3.13.2. `mysqli_warning::next`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_warning::next`

The next purpose

#### Description

```
public void mysqli_warning::next();
```



#### Warning

This function is currently not documented; only its argument list is available.

#### Parameters

This function has no parameters.

#### Return Values

### 21.9.3.14. The `mysqli_sql_exception` class (`mysqli_sql_exception`)

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqli` exception handling class.

```
mysqli_sql_exception {
mysqli_sql_exception extends RuntimeException

    Properties

    protected
    code ;

    protected
    sqlstate ;
}
```

<code>message</code>	The error message.
<code>file</code>	The file with the error.
<code>line</code>	The line with the error.
<code>code</code>	The code causing the error.
<code>sqlstate</code>	The sql state with the error.

### 21.9.3.15. Aliases and deprecated `Mysqli` Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.3.15.1. `mysqli_bind_param`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_bind\\_param](#)

Alias for [mysqli\\_stmt\\_bind\\_param](#)

#### Description

This function is an alias of [mysqli\\_stmt\\_bind\\_param](#).



#### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

#### See Also

[mysqli\\_stmt\\_bind\\_param](#)

### 21.9.3.15.2. [mysqli\\_bind\\_result](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_bind\\_result](#)

Alias for [mysqli\\_stmt\\_bind\\_result](#)

#### Description

This function is an alias of [mysqli\\_stmt\\_bind\\_result](#).



#### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

#### See Also

[mysqli\\_stmt\\_bind\\_result](#)

### 21.9.3.15.3. [mysqli\\_client\\_encoding](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_client\\_encoding](#)

Alias of [mysqli\\_character\\_set\\_name](#)

#### Description

This function is an alias of [mysqli\\_character\\_set\\_name](#).



#### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

#### See Also

[mysqli\\_real\\_escape\\_string](#)

### 21.9.3.15.4. [mysqli\\_connect](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_connect`

Alias of `mysqli::__construct`

### Description

This function is an alias of: `mysqli::__construct`

#### 21.9.3.15.5. `mysqli::disable_reads_from_master`, `mysqli_disable_reads_from_master`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli::disable_reads_from_master`

`mysqli_disable_reads_from_master`

Disable reads from master

### Description

Object oriented style

```
void mysqli::disable_reads_from_master();
```

Procedural style

```
bool mysqli_disable_reads_from_master(  
    mysqli link);
```



#### Warning

This function is currently not documented; only its argument list is available.



#### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

#### 21.9.3.15.6. `mysqli_disable_rpl_parse`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_disable_rpl_parse`

Disable RPL parse

### Description

```
bool mysqli_disable_rpl_parse(  
    mysqli link);
```



#### Warning

This function is currently not documented; only its argument list is available.



#### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

#### 21.9.3.15.7. `mysqli_enable_reads_from_master`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_enable\\_reads\\_from\\_master](#)

Enable reads from master

#### Description

```
bool mysqli_enable_reads_from_master(  
    mysqli link);
```



#### Warning

This function is currently not documented; only its argument list is available.



#### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

#### 21.9.3.15.8. [mysqli\\_enable\\_rpl\\_parse](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_enable\\_rpl\\_parse](#)

Enable RPL parse

#### Description

```
bool mysqli_enable_rpl_parse(  
    mysqli link);
```



#### Warning

This function is currently not documented; only its argument list is available.



#### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

#### 21.9.3.15.9. [mysqli\\_escape\\_string](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_escape\\_string](#)

Alias of [mysqli\\_real\\_escape\\_string](#)

#### Description

This function is an alias of: [mysqli\\_real\\_escape\\_string](#).

#### 21.9.3.15.10. [mysqli\\_execute](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_execute](#)

Alias for [mysqli\\_stmt\\_execute](#)

#### Description

This function is an alias of [mysqli\\_stmt\\_execute](#).

## Notes



### Note

`mysqli_execute` is deprecated and will be removed.

## See Also

`mysqli_stmt_execute`

### 21.9.3.15.11. `mysqli_fetch`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_fetch`

Alias for `mysqli_stmt_fetch`

## Description

This function is an alias of `mysqli_stmt_fetch`.



### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

## See Also

`mysqli_stmt_fetch`

### 21.9.3.15.12. `mysqli_get_cache_stats`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_get_cache_stats`

Returns client Zval cache statistics

## Description

```
array mysqli_get_cache_stats();
```



### Warning

This function is currently not documented; only its argument list is available.

Returns client Zval cache statistics. Available only with `mysqlnd`.

## Parameters

## Return Values

Returns an array with client Zval cache stats if success, `FALSE` otherwise.

## Examples

### Example 21.211. A `mysqli_get_cache_stats` example

```
<?php
$link = mysqli_connect();
print_r(mysqli_get_cache_stats());
```

```
?>
```

The above example will output something similar to:

```
Array
(
    [bytes_sent] => 43
    [bytes_received] => 80
    [packets_sent] => 1
    [packets_received] => 2
    [protocol_overhead_in] => 8
    [protocol_overhead_out] => 4
    [bytes_received_ok_packet] => 11
    [bytes_received_eof_packet] => 0
    [bytes_received_rset_header_packet] => 0
    [bytes_received_rset_field_meta_packet] => 0
    [bytes_received_rset_row_packet] => 0
    [bytes_received_prepare_response_packet] => 0
    [bytes_received_change_user_packet] => 0
    [packets_sent_command] => 0
    [packets_received_ok] => 1
    [packets_received_eof] => 0
    [packets_received_rset_header] => 0
    [packets_received_rset_field_meta] => 0
    [packets_received_rset_row] => 0
    [packets_received_prepare_response] => 0
    [packets_received_change_user] => 0
    [result_set_queries] => 0
    [non_result_set_queries] => 0
    [no_index_used] => 0
    [bad_index_used] => 0
    [slow_queries] => 0
    [buffered_sets] => 0
    [unbuffered_sets] => 0
    [ps_buffered_sets] => 0
    [ps_unbuffered_sets] => 0
    [flushed_normal_sets] => 0
    [flushed_ps_sets] => 0
    [ps_prepared_never_executed] => 0
    [ps_prepared_once_executed] => 0
    [rows_fetched_from_server_normal] => 0
    [rows_fetched_from_server_ps] => 0
    [rows_buffered_from_client_normal] => 0
    [rows_buffered_from_client_ps] => 0
    [rows_fetched_from_client_normal_buffered] => 0
    [rows_fetched_from_client_normal_unbuffered] => 0
    [rows_fetched_from_client_ps_buffered] => 0
    [rows_fetched_from_client_ps_unbuffered] => 0
    [rows_fetched_from_client_ps_cursor] => 0
    [rows_skipped_normal] => 0
    [rows_skipped_ps] => 0
    [copy_on_write_saved] => 0
    [copy_on_write_performed] => 0
    [command_buffer_too_small] => 0
    [connect_success] => 1
    [connect_failure] => 0
    [connection_reused] => 0
    [reconnect] => 0
    [pconnect_success] => 0
    [active_connections] => 1
    [active_persistent_connections] => 0
    [explicit_close] => 0
    [implicit_close] => 0
    [disconnect_close] => 0
    [in_middle_of_command_close] => 0
    [explicit_free_result] => 0
    [implicit_free_result] => 0
    [explicit_stmt_close] => 0
    [implicit_stmt_close] => 0
)
```



```
[mem_emalloc_count] => 0
[mem_emalloc_ammount] => 0
[mem_ecalloc_count] => 0
[mem_ecalloc_ammount] => 0
[mem_erealloc_count] => 0
[mem_erealloc_ammount] => 0
[mem_efree_count] => 0
[mem_malloc_count] => 0
[mem_malloc_ammount] => 0
[mem_calloc_count] => 0
[mem_calloc_ammount] => 0
[mem_realloc_count] => 0
[mem_realloc_ammount] => 0
[mem_free_count] => 0
[proto_text_fetched_null] => 0
[proto_text_fetched_bit] => 0
[proto_text_fetched_tinyint] => 0
[proto_text_fetched_short] => 0
[proto_text_fetched_int24] => 0
[proto_text_fetched_int] => 0
[proto_text_fetched_bigint] => 0
[proto_text_fetched_decimal] => 0
[proto_text_fetched_float] => 0
[proto_text_fetched_double] => 0
[proto_text_fetched_date] => 0
[proto_text_fetched_year] => 0
[proto_text_fetched_time] => 0
[proto_text_fetched_datetime] => 0
[proto_text_fetched_timestamp] => 0
[proto_text_fetched_string] => 0
[proto_text_fetched_blob] => 0
[proto_text_fetched_enum] => 0
[proto_text_fetched_set] => 0
[proto_text_fetched_geometry] => 0
[proto_text_fetched_other] => 0
[proto_binary_fetched_null] => 0
[proto_binary_fetched_bit] => 0
[proto_binary_fetched_tinyint] => 0
[proto_binary_fetched_short] => 0
[proto_binary_fetched_int24] => 0
[proto_binary_fetched_int] => 0
[proto_binary_fetched_bigint] => 0
[proto_binary_fetched_decimal] => 0
[proto_binary_fetched_float] => 0
[proto_binary_fetched_double] => 0
[proto_binary_fetched_date] => 0
[proto_binary_fetched_year] => 0
[proto_binary_fetched_time] => 0
[proto_binary_fetched_datetime] => 0
[proto_binary_fetched_timestamp] => 0
[proto_binary_fetched_string] => 0
[proto_binary_fetched_blob] => 0
[proto_binary_fetched_enum] => 0
[proto_binary_fetched_set] => 0
[proto_binary_fetched_geometry] => 0
[proto_binary_fetched_other] => 0
)
```

## See Also

[Stats description](#)

### 21.9.3.15.13. [mysqli\\_get\\_metadata](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_get\\_metadata](#)

Alias for [mysqli\\_stmt\\_result\\_metadata](#)

## Description

This function is an alias of [mysqli\\_stmt\\_result\\_metadata](#).



### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

## See Also

[mysqli\\_stmt\\_result\\_metadata](#)

### 21.9.3.15.14. [mysqli\\_master\\_query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_master\\_query](#)

Enforce execution of a query on the master in a master/slave setup

## Description

```
bool mysqli_master_query(  
    mysqli link,  
    string query);
```



### Warning

This function is currently not documented; only its argument list is available.



### Warning

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

### 21.9.3.15.15. [mysqli\\_param\\_count](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_param\\_count](#)

Alias for [mysqli\\_stmt\\_param\\_count](#)

## Description

This function is an alias of [mysqli\\_stmt\\_param\\_count](#).



### Warning

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

## See Also

[mysqli\\_stmt\\_param\\_count](#)

### 21.9.3.15.16. [mysqli\\_report](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqli\\_report](#)

Alias of of [mysqli\\_driver->report\\_mode](#)

**Description**

This function is an alias of: `mysqli_driver->report_mode`

**21.9.3.15.17. `mysqli_rpl_parse_enabled`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_rpl_parse_enabled`

Check if RPL parse is enabled

**Description**

```
int mysqli_rpl_parse_enabled(  
    mysqli link);
```

**Warning**

This function is currently not documented; only its argument list is available.

**Warning**

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

**21.9.3.15.18. `mysqli_rpl_probe`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_rpl_probe`

RPL probe

**Description**

```
bool mysqli_rpl_probe(  
    mysqli link);
```

**Warning**

This function is currently not documented; only its argument list is available.

**Warning**

This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.

**21.9.3.15.19. `mysqli_send_long_data`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqli_send_long_data`

Alias for `mysqli_stmt_send_long_data`

**Description**

This function is an alias of `mysqli_stmt_send_long_data`.

**Warning**

This function has been *DEPRECATED* as of PHP 5.3.0 and *REMOVED* as of PHP 5.4.0.

**See Also**[mysqli\\_stmt\\_send\\_long\\_data](#)**21.9.3.15.20. [mysqli\\_set\\_opt](#)**[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

- [mysqli\\_set\\_opt](#)

Alias of [mysqli\\_options](#)**Description**This function is an alias of [mysqli\\_options](#).**21.9.3.15.21. [mysqli\\_slave\\_query](#)**[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

- [mysqli\\_slave\\_query](#)

Force execution of a query on a slave in a master/slave setup

**Description**

```
bool mysqli_slave_query(
    mysqli link,
    string query);
```

**Warning**

This function is currently not documented; only its argument list is available.

**Warning**This function has been *DEPRECATED* and *REMOVED* as of PHP 5.3.0.**21.9.3.16. Changelog**[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The following changes have been made to classes/functions/methods of this extension.

**21.9.4. MySQL Functions (PDO\_MYSQL) (MySQL (PDO))**[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)PDO\_MYSQL is a driver that implements the [PHP Data Objects \(PDO\) interface](#) to enable access from PHP to MySQL 3.x, 4.x and 5.x databases.

PDO\_MYSQL will take advantage of native prepared statement support present in MySQL 4.1 and higher. If you're using an older version of the mysql client libraries, PDO will emulate them for you.

**Warning**

Beware: Some MySQL table types (storage engines) do not support transactions. When writing transactional database code using a table type that does not support transactions, MySQL will pretend that a transaction was

initiated successfully. In addition, any DDL queries issued will implicitly commit any pending transactions.

Use `--with-pdo-mysql[=DIR]` to install the PDO MySQL extension, where the optional `[=DIR]` is the MySQL base install directory. If `mysqlnd` is passed as `[=DIR]`, then the MySQL native driver will be used.

Optionally, the `--with-mysql-sock[=DIR]` sets to location to the MySQL unix socket pointer for all MySQL extensions, including PDO\_MYSQL. If unspecified, the default locations are searched.

Optionally, the `--with-zlib-dir[=DIR]` is used to set the path to the libz install prefix.

```
$ ./configure --with-pdo-mysql --with-mysql-sock=/var/mysql/mysql.sock
```

SSL support is enabled using the appropriate [PDO\\_MySQL constants \[3197\]](#), which is equivalent to calling the [MySQL C API function `mysql\_ssl\_set\(\)`](#). Also, SSL cannot be enabled with `PDO::setAttribute` because the connection already exists. See also the MySQL documentation about [connecting to MySQL with SSL](#).

**Table 21.57. Changelog**

Version	Description
5.4.0	MySQL client libraries 4.1 and below are no longer supported.
5.3.9	Added SSL support with <code>mysqlnd</code> and OpenSSL.
5.3.7	Added SSL support with <code>libmysqlclient</code> and OpenSSL.

The constants below are defined by this driver, and will only be available when the extension has been either compiled into PHP or dynamically loaded at runtime. In addition, these driver-specific constants should only be used if you are using this driver. Using driver-specific attributes with another driver may result in unexpected behaviour. `PDO::getAttribute` may be used to obtain the `PDO_ATTR_DRIVER_NAME` attribute to check the driver, if your code can run against multiple drivers.

`PDO::MYSQL_ATTR_USE_BUFFERED_QUERY` (integer) If this attribute is set to `TRUE` on a `PDOStatement`, the MySQL driver will use the buffered versions of the MySQL API. If you're writing portable code, you should use `PDOStatement::fetchAll` instead.

**Example 21.212. Forcing queries to be buffered in mysql**

```
<?php
if ($db->getAttribute(PDO::ATTR_DRIVER_NAME) == 'mysql') {
    $stmt = $db->prepare('select * from foo',
        array(PDO::MYSQL_ATTR_USE_BUFFERED_QUERY => true));
} else {
    die("my application only works with mysql; I should use \$stmt->fet
}
?>
```

`PDO::MYSQL_ATTR_LOCAL_INFILE` (integer) Enable `LOAD LOCAL INFILE`.

Note, this constant can only be used in the `driver_options` array when constructing a new database handle.

- `PDO::MYSQL_ATTR_INIT_COMMAND` (integer) Command to execute when connecting to the MySQL server. Will automatically be re-executed when reconnecting.
- Note, this constant can only be used in the `driver_options` array when constructing a new database handle.
- `PDO::MYSQL_ATTR_READ_DEFAULT_GROUP` (integer) Read options from the named option file instead of from `my.cnf`. This option is not available if `mysqlnd` is used, because `mysqlnd` does not read the `mysql` configuration files.
- `PDO::MYSQL_ATTR_READ_DEFAULT_SECTION` (integer) Read options from the named group from `my.cnf` or the file specified with `MYSQL_READ_DEFAULT_FILE`. This option is not available if `mysqlnd` is used, because `mysqlnd` does not read the `mysql` configuration files.
- `PDO::MYSQL_ATTR_MAX_BUFFER_SIZE` (integer) Maximum buffer size. Defaults to 1 MiB. This constant is not supported when compiled against `mysqlnd`.
- `PDO::MYSQL_ATTR_DIRECT_QUERY` (integer) Perform direct queries, don't use prepared statements.
- `PDO::MYSQL_ATTR_FOUND_ROWS` (integer) Return the number of found (matched) rows, not the number of changed rows.
- `PDO::MYSQL_ATTR_IGNORE_SPACE` (integer) Permit spaces after function names. Makes all functions names reserved words.
- `PDO::MYSQL_ATTR_COMPRESS` (integer) Enable network communication compression. This is not supported when compiled against `mysqlnd`.
- `PDO::MYSQL_ATTR_SSL_CA` (integer) The file path to the SSL certificate authority.
- This exists as of PHP 5.3.7.
- `PDO::MYSQL_ATTR_SSL_CAPATH` (integer) The file path to the directory that contains the trusted SSL CA certificates, which are stored in PEM format.
- This exists as of PHP 5.3.7.
- `PDO::MYSQL_ATTR_SSL_CERT` (integer) The file path to the SSL certificate.
- This exists as of PHP 5.3.7.
- `PDO::MYSQL_ATTR_SSL_CIPHER` (integer) A list of one or more permissible ciphers to use for SSL encryption, in a format understood by OpenSSL. For example: `DHE-RSA-AES256-SHA:AES128-SHA`
- This exists as of PHP 5.3.7.
- `PDO::MYSQL_ATTR_SSL_KEY` (integer) The file path to the SSL key.
- This exists as of PHP 5.3.7.

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.58. PDO\_MYSQL Configuration Options**

Name	Default	Changeable
<code>pdo_mysql.default_socket</code>	<code>"/tmp/mysql.sock"</code>	PHP_INI_SYSTEM
<code>pdo_mysql.debug</code>	NULL	PHP_INI_SYSTEM

For further details and definitions of the PHP\_INI\_\* modes, see the <http://www.php.net/manual/en/configuration.changes.modes>.

Here's a short explanation of the configuration directives.

<code>pdo_mysql.default_socket</code> string	Sets a Unix domain socket. This value can either be set at compile time if a domain socket is found at configure. This ini setting is Unix only.
<code>pdo_mysql.debug</code> boolean	Enables debugging for PDO_MYSQL. This setting is only available when PDO_MYSQL is compiled against mysqlnd and in PDO debug mode.

### 21.9.4.1. PDO\_MYSQL DSN

Copyright 1997-2012 the PHP Documentation Group. [2906]

- PDO\_MYSQL DSN

Connecting to MySQL databases

#### Description

The PDO\_MYSQL Data Source Name (DSN) is composed of the following elements:

DSN prefix	The DSN prefix is <code>mysql:</code> .
<code>host</code>	The hostname on which the database server resides.
<code>port</code>	The port number where the database server is listening.
<code>dbname</code>	The name of the database.
<code>unix_socket</code>	The MySQL Unix socket (shouldn't be used with <code>host</code> or <code>port</code> ).
<code>charset</code>	The character set. See the <a href="#">character set</a> concepts documentation for more information.  Prior to PHP 5.3.6, this element was silently ignored. The same behaviour can be partly replicated with the <code>PDO::MYSQL_ATTR_INIT_COMMAND</code> driver option, as the following example shows.



#### Warning

The method in the below example can only be used with character sets that share the same lower 7 bit representation as ASCII, such as ISO-8859-1 and UTF-8. Users using character sets that have different representations (such as UTF-16 or Big5) *must* use the `charset` option provided in PHP 5.3.6 and later versions.

#### Example 21.213. Setting the connection character set to UTF-8 prior to PHP 5.3.6

```
<?php
$dsn = 'mysql:host=localhost;dbname=testdb';
$username = 'username';
$password = 'password';
$options = array(
```

```

PDO::MYSQL_ATTR_INIT_COMMAND => 'SET NAMES utf8',
);

$dbh = new PDO($dsn, $username, $password, $options);
?>

```

## Changelog

Version	Description
5.3.6	Prior to version 5.3.6, <code>charset</code> was ignored.

## Examples

### Example 21.214. PDO\_MYSQL DSN examples

The following example shows a PDO\_MYSQL DSN for connecting to MySQL databases:

```
mysql:host=localhost;dbname=testdb
```

More complete examples:

```
mysql:host=localhost;port=3307;dbname=testdb
mysql:unix_socket=/tmp/mysql.sock;dbname=testdb
```

## Notes



### Unix only:

When the host name is set to "`localhost`", then the connection to the server is made thru a domain socket. If PDO\_MYSQL is compiled against `libmysqlclient` then the location of the socket file is at `libmysqlclient`'s compiled in location. If PDO\_MYSQL is compiled against `mysqlnd` a default socket can be set thru the `pdo_mysql.default_socket` setting.

## 21.9.5. MySQL Native Driver ([Mysqlnd](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

MySQL Native Driver is a replacement for the MySQL Client Library (`libmysqlclient`). MySQL Native Driver is part of the official PHP sources as of PHP 5.3.0.

The MySQL database extensions MySQL extension, `mysqli` and PDO MYSQL all communicate with the MySQL server. In the past, this was done by the extension using the services provided by the MySQL Client Library. The extensions were compiled against the MySQL Client Library in order to use its client-server protocol.

With MySQL Native Driver there is now an alternative, as the MySQL database extensions can be compiled to use MySQL Native Driver instead of the MySQL Client Library.

MySQL Native Driver is written in C as a PHP extension.

### 21.9.5.1. Overview

Copyright 1997-2012 the PHP Documentation Group. [2906]



### *What it is not*

Although MySQL Native Driver is written as a PHP extension, it is important to note that it does not provide a new API to the PHP programmer. The programmer APIs for MySQL database connectivity are provided by the MySQL extension, `mysqli` and PDO MySQL. These extensions can now use the services of MySQL Native Driver to communicate with the MySQL Server. Therefore, you should not think of MySQL Native Driver as an API.

### *Why use it?*

Using the MySQL Native Driver offers a number of advantages over using the MySQL Client Library.

The older MySQL Client Library was written by MySQL AB (now Oracle Corporation) and so was released under the MySQL license. This ultimately led to MySQL support being disabled by default in PHP. However, the MySQL Native Driver has been developed as part of the PHP project, and is therefore released under the PHP license. This removes licensing issues that have been problematic in the past.

Also, in the past, you needed to build the MySQL database extensions against a copy of the MySQL Client Library. This typically meant you needed to have MySQL installed on a machine where you were building the PHP source code. Also, when your PHP application was running, the MySQL database extensions would call down to the MySQL Client library file at run time, so the file needed to be installed on your system. With MySQL Native Driver that is no longer the case as it is included as part of the standard distribution. So you do not need MySQL installed in order to build PHP or run PHP database applications.

Because MySQL Native Driver is written as a PHP extension, it is tightly coupled to the workings of PHP. This leads to gains in efficiency, especially when it comes to memory usage, as the driver uses the PHP memory management system. It also supports the PHP memory limit. Using MySQL Native Driver leads to comparable or better performance than using MySQL Client Library, it always ensures the most efficient use of memory. One example of the memory efficiency is the fact that when using the MySQL Client Library, each row is stored in memory twice, whereas with the MySQL Native Driver each row is only stored once in memory.



#### **Reporting memory usage**

Because MySQL Native Driver uses the PHP memory management system, its memory usage can be tracked with `memory_get_usage`. This is not possible with `libmysqlclient` because it uses the C function `malloc()` instead.

### *Special features*

MySQL Native Driver also provides some special features not available when the MySQL database extensions use MySQL Client Library. These special features are listed below:

- Improved persistent connections
- The special function `mysqli_fetch_all`
- Performance statistics calls: `mysqli_get_cache_stats`, `mysqli_get_client_stats`, `mysqli_get_connection_stats`

The performance statistics facility can prove to be very useful in identifying performance bottlenecks.

MySQL Native Driver also allows for persistent connections when used with the `mysqli` extension.

### *SSL Support*

MySQL Native Driver has supported SSL since PHP version 5.3.3

### *Compressed Protocol Support*

As of PHP 5.3.2 MySQL Native Driver supports the compressed client server protocol. MySQL Native Driver did not support this in 5.3.0 and 5.3.1. Extensions such as [ext/mysql](#), [ext/mysqli](#), that are configured to use MySQL Native Driver, can also take advantage of this feature. Note that [PDO\\_MYSQL](#) does *NOT* support compression when used together with [mysqlnd](#).

#### *Named Pipes Support*

Named pipes support for Windows was added in PHP version 5.4.0.

### 21.9.5.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### *Changelog*

**Table 21.59. Changelog**

Version	Description
5.3.0	The MySQL Native Driver was added, with support for all MySQL extensions (i.e., <a href="#">mysql</a> , <a href="#">mysqli</a> and <a href="#">PDO_MYSQL</a> ). Passing in <a href="#">mysqlnd</a> to the appropriate configure switch enables this support.
5.4.0	The MySQL Native Driver is now the default for all MySQL extensions (i.e., <a href="#">mysql</a> , <a href="#">mysqli</a> and <a href="#">PDO_MYSQL</a> ). Passing in <a href="#">mysqlnd</a> to configure is now optional.
5.5.0	SHA-256 Authentication Plugin support was added

#### *Installation on Unix*

The MySQL database extensions must be configured to use the MySQL Client Library. In order to use the MySQL Native Driver, PHP needs to be built specifying that the MySQL database extensions are compiled with MySQL Native Driver support. This is done through configuration options prior to building the PHP source code.

For example, to build the MySQL extension, [mysqli](#) and PDO MYSQL using the MySQL Native Driver, the following command would be given:

```
./configure --with-mysql=mysqlnd \  
--with-mysqli=mysqlnd \  
--with-pdo-mysql=mysqlnd \  
[other options]
```

#### *Installation on Windows*

In the official PHP Windows distributions from 5.3 onwards, MySQL Native Driver is enabled by default, so no additional configuration is required to use it. All MySQL database extensions will use MySQL Native Driver in this case.

#### *SHA-256 Authentication Plugin support*

The MySQL Native Driver requires the OpenSSL functionality of PHP to be loaded and enabled to connect to MySQL through accounts that use the MySQL SHA-256 Authentication Plugin. For example, PHP could be configured using:

```
./configure --with-mysql=mysqlnd \  
--with-mysqli=mysqlnd \  
--with-openssl=
```

```
--with-pdo-mysql=mysqlnd \
--with-openssl
[other options]
```

### 21.9.5.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.60. MySQL Native Driver Configuration Options**

Name	Default	Changeable	Changelog
<code>mysqlnd.collect_statistics</code>	"1"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.collect_memory_statistics</code>	"0"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.debug</code>	"0"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.log_mask</code>	0	PHP_INI_ALL	Available since PHP 5.3.0
<code>mysqlnd.mempool_default_size</code>	16000	PHP_INI_ALL	Available since PHP 5.3.3
<code>mysqlnd.net_read_timeout</code>	31536000"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.net_cmd_buffer_size</code>	512 - "2048", 5.3.1 - "4096"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.net_read_buffer_size</code>	32768"	PHP_INI_SYSTEM	Available since PHP 5.3.0.
<code>mysqlnd.sha256_server_public_key</code>	"public_key"	PHP_INI_PERDIR	Available since PHP 5.5.0.

For further details and definitions of the `PHP_INI_*` modes, see the <http://www.php.net/manual/en/configuration.changes.modes>.

Here's a short explanation of the configuration directives.

`mysqlnd.collect_statistics` **boolean** Enables the collection of various client statistics which can be accessed through `mysqli_get_client_stats`, `mysqli_get_connection_stats`, `mysqli_get_cache_stats` and are shown in `mysqlnd` section of the output of the `phpinfo` function as well.

This configuration setting enables all [MySQL Native Driver statistics](#) except those relating to memory management.

`mysqlnd.collect_memory_statistics` **boolean** Enable the collection of various memory statistics which can be accessed through `mysqli_get_client_stats`, `mysqli_get_connection_stats`, `mysqli_get_cache_stats` and are shown in `mysqlnd` section of the output of the `phpinfo` function as well.

This configuration setting enables the memory management statistics within the overall set of [MySQL Native Driver statistics](#).

`mysqlnd.debug` **string** Records communication from all extensions using `mysqlnd` to the specified log file.

The format of the directive is `mysqlnd.debug = "option1[,parameter_option1] [:option2[,parameter_option2]]"`.

The options for the format string are as follows:

- `A[,file]` - Appends trace output to specified file. Also ensures that data is written after each write. This is done by closing and reopening the trace file (this is slow). It helps ensure a complete log file should the application crash.
- `a[,file]` - Appends trace output to the specified file.
- `d` - Enables output from `DEBUG_<N>` macros for the current state. May be followed by a list of keywords which selects output only for the `DEBUG` macros with that keyword. An empty list of keywords implies output for all macros.
- `f[,functions]` - Limits debugger actions to the specified list of functions. An empty list of functions implies that all functions are selected.
- `F` - Marks each debugger output line with the name of the source file containing the macro causing the output.
- `i` - Marks each debugger output line with the PID of the current process.
- `L` - Marks each debugger output line with the name of the source file line number of the macro causing the output.
- `n` - Marks each debugger output line with the current function nesting depth
- `o[,file]` - Similar to `a[,file]` but overwrites old file, and does not append.
- `O[,file]` - Similar to `A[,file]` but overwrites old file, and does not append.
- `t[,N]` - Enables function control flow tracing. The maximum nesting depth is specified by `N`, and defaults to 200.
- `x` - This option activates profiling.

Example:

```
d:t:x:0,/tmp/mysqlnd.trace
```



#### Note

This feature is only available with a debug build of PHP. Works on Microsoft Windows if using a debug build of PHP and PHP was built using Microsoft Visual C version 9 and above.

`mysqlnd.log_mask` integer

Defines which queries will be logged. The default 0, which disables logging. Define using an integer, and not with PHP constants. For

example, a value of 48 (16 + 32) will log slow queries which either use 'no good index' (`SERVER_QUERY_NO_GOOD_INDEX_USED = 16`) or no index at all (`SERVER_QUERY_NO_INDEX_USED = 32`). A value of 2043 (1 + 2 + 8 + ... + 1024) will log all slow query types.

The types are as follows: `SERVER_STATUS_IN_TRANS=1`, `SERVER_STATUS_AUTOCOMMIT=2`, `SERVER_MORE_RESULTS_EXISTS=8`, `SERVER_QUERY_NO_GOOD_INDEX_USED=16`, `SERVER_QUERY_NO_INDEX_USED=32`, `SERVER_STATUS_CURSOR_EXISTS=64`, `SERVER_STATUS_LAST_ROW_SENT=128`, `SERVER_STATUS_DB_DROPPED=256`, `SERVER_STATUS_NO_BACKSLASH_ESCAPES=512`, and `SERVER_QUERY_WAS_SLOW=1024`.

[mysqlnd.mempool\\_default\\_size](#) integer Default size of the mysqlnd memory pool, which is used by result sets.

[mysqlnd.net\\_read\\_timeout](#) integer [mysqlnd](#) and the MySQL Client Library, [libmysqlclient](#) use different networking APIs. [mysqlnd](#) uses PHP streams, whereas [libmysqlclient](#) uses its own wrapper around the operating level network calls. PHP, by default, sets a read timeout of 60s for streams. This is set via [php.ini](#), [default\\_socket\\_timeout](#). This default applies to all streams that set no other timeout value. [mysqlnd](#) does not set any other value and therefore connections of long running queries can be disconnected after [default\\_socket\\_timeout](#) seconds resulting in an error message "2006 - MySQL Server has gone away". The MySQL Client Library sets a default timeout of 365 \* 24 \* 3600 seconds (1 year) and waits for other timeouts to occur, such as TCP/IP timeouts. [mysqlnd](#) now uses the same very long timeout. The value is configurable through a new [php.ini](#) setting: [mysqlnd.net\\_read\\_timeout](#). [mysqlnd.net\\_read\\_timeout](#) gets used by any extension ([ext/mysql](#), [ext/mysql\\_i](#), [PDO\\_MySQL](#)) that uses [mysqlnd](#). [mysqlnd](#) tells PHP Streams to use [mysqlnd.net\\_read\\_timeout](#). Please note that there may be subtle differences between [MYSQL\\_OPT\\_READ\\_TIMEOUT](#) from the MySQL Client Library and PHP Streams, for example [MYSQL\\_OPT\\_READ\\_TIMEOUT](#) is documented to work only for TCP/IP connections and, prior to MySQL 5.1.2, only for Windows. PHP streams may not have this limitation. Please check the streams documentation, if in doubt.

[mysqlnd.net\\_cmd\\_buffer\\_size](#) long [mysqlnd](#) allocates an internal command/network buffer of [mysqlnd.net\\_cmd\\_buffer\\_size](#) (in [php.ini](#)) bytes for every connection. If a MySQL Client Server protocol command, for example, `COM_QUERY` ("normal" query), does not fit into the buffer, [mysqlnd](#) will grow the buffer to the size required for sending the command. Whenever the buffer gets extended for one connection, [command\\_buffer\\_too\\_small](#) will be incremented by one.

If [mysqlnd](#) has to grow the buffer beyond its initial size of [mysqlnd.net\\_cmd\\_buffer\\_size](#) bytes for almost every connection, you should consider increasing the default size to avoid re-allocations.

The default buffer size is 2048 bytes in PHP 5.3.0. In later versions the default is 4096 bytes. The default can be changed either through the `php.ini` setting `mysqlnd.net_cmd_buffer_size` or using `mysqli_options(MYSQLI_OPT_NET_CMD_BUFFER_SIZE, int size)`.

It is recommended that the buffer size be set to no less than 4096 bytes because `mysqlnd` also uses it when reading certain communication packets from MySQL. In PHP 5.3.0, `mysqlnd` will not grow the buffer if MySQL sends a packet that is larger than the current size of the buffer. As a consequence, `mysqlnd` is unable to decode the packet and the client application will get an error. There are only two situations when the packet can be larger than the 2048 bytes default of `mysqlnd.net_cmd_buffer_size` in PHP 5.3.0: the packet transports a very long error message, or the packet holds column meta data from `COM_LIST_FIELD` (`mysql_list_fields()`) and the meta data come from a string column with a very long default value (>1900 bytes).

As of PHP 5.3.2 `mysqlnd` does not allow setting buffers smaller than 4096 bytes.

The value can also be set using `mysqli_options(link, MYSQLI_OPT_NET_CMD_BUFFER_SIZE, size)`.

`mysqlnd.net_read_buffer_size`  
long

Maximum read chunk size in bytes when reading the body of a MySQL command packet. The MySQL client server protocol encapsulates all its commands in packets. The packets consist of a small header and a body with the actual payload. The size of the body is encoded in the header. `mysqlnd` reads the body in chunks of `MIN(header.size, mysqlnd.net_read_buffer_size)` bytes. If a packet body is larger than `mysqlnd.net_read_buffer_size` bytes, `mysqlnd` has to call `read()` multiple times.

The value can also be set using `mysqli_options(link, MYSQLI_OPT_NET_READ_BUFFER_SIZE, size)`.

`mysqlnd.sha256_server_public_key`  
string

SHA-256 Authentication Plugin related. File with the MySQL server public RSA key.

Clients can either omit setting a public RSA key, specify the key through this PHP configuration setting or set the key at runtime using `mysqli_options`. If not public RSA key file is given by the client, then the key will be exchanged as part of the standard SHA-256 Authentication Plugin authentication procedure.

#### 21.9.5.4. Persistent Connections

Copyright 1997-2012 the PHP Documentation Group. [2906]

##### *Using Persistent Connections*

If `mysqli` is used with `mysqlnd`, when a persistent connection is created it generates a `COM_CHANGE_USER` (`mysql_change_user()`) call on the server. This ensures that re-authentication of the connection takes place.

As there is some overhead associated with the `COM_CHANGE_USER` call, it is possible to switch this off at compile time. Reusing a persistent connection will then generate a `COM_PING` (`mysql_ping`) call to simply test the connection is reusable.

Generation of `COM_CHANGE_USER` can be switched off with the compile flag `MYSQLI_NO_CHANGE_USER_ON_PCONNECT`. For example:

```
shell# CFLAGS="-DMYSQLI_NO_CHANGE_USER_ON_PCONNECT" ./configure --with-mysql=/usr/local/mysql/ --with-m
```

Or alternatively:

```
shell# export CFLAGS="-DMYSQLI_NO_CHANGE_USER_ON_PCONNECT"
shell# configure --whatever-option
shell# make clean
shell# make
```

Note that only `mysqli` on `mysqlnd` uses `COM_CHANGE_USER`. Other extension-driver combinations use `COM_PING` on initial use of a persistent connection.

### 21.9.5.5. Statistics

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### *Using Statistical Data*

MySQL Native Driver contains support for gathering statistics on the communication between the client and the server. The statistics gathered are of three main types:

- Client statistics
- Connection statistics
- Zval cache statistics

If you are using the `mysqli` extension, these statistics can be obtained through three API calls:

- `mysqli_get_client_stats`
- `mysqli_get_connection_stats`
- `mysqli_get_cache_stats`



#### **Note**

Statistics are aggregated among all extensions that use MySQL Native Driver. For example, when compiling both `ext/mysql` and `ext/mysqli` against MySQL Native Driver, both function calls of `ext/mysql` and `ext/mysqli` will change the statistics. There is no way to find out how much a certain API call of any extension that has been compiled against MySQL Native Driver has impacted a certain statistic. You can configure the PDO MySQL Driver, `ext/mysql` and `ext/mysqli` to optionally use the MySQL Native Driver. When doing so, all three extensions will change the statistics.

#### *Accessing Client Statistics*

To access client statistics, you need to call `mysqli_get_client_stats`. The function call does not require any parameters.

The function returns an associative array that contains the name of the statistic as the key and the statistical data as the value.

Client statistics can also be accessed by calling the `phpinfo` function.

#### *Accessing Connection Statistics*

To access connection statistics call `mysqli_get_connection_stats`. This takes the database connection handle as the parameter.



The function returns an associative array that contains the name of the statistic as the key and the statistical data as the value.

*Accessing Zval Cache Statistics*

The MySQL Native Driver also collects statistics from its internal Zval cache. These statistics can be accessed by calling `mysqli_get_cache_stats`.

The Zval cache statistics obtained may lead to a tweaking of `php.ini` settings related to the Zval cache, resulting in better performance.

*Buffered and Unbuffered Result Sets*

Result sets can be buffered or unbuffered. Using default settings, `ext/mysql` and `ext/mysqli` work with buffered result sets for normal (non prepared statement) queries. Buffered result sets are cached on the client. After the query execution all results are fetched from the MySQL Server and stored in a cache on the client. The big advantage of buffered result sets is that they allow the server to free all resources allocated to a result set, once the results have been fetched by the client.

Unbuffered result sets on the other hand are kept much longer on the server. If you want to reduce memory consumption on the client, but increase load on the server, use unbuffered results. If you experience a high server load and the figures for unbuffered result sets are high, you should consider moving the load to the clients. Clients typically scale better than servers. "Load" does not only refer to memory buffers - the server also needs to keep other resources open, for example file handles and threads, before a result set can be freed.

Prepared Statements use unbuffered result sets by default. However, you can use `mysqli_stmt_store_result` to enable buffered result sets.

*Statistics returned by MySQL Native Driver*

The following tables show a list of statistics returned by the `mysqli_get_client_stats`, `mysqli_get_connection_stats` and `mysqli_get_cache_stats` functions.

**Table 21.61. Returned mysqlnd statistics: Network**

Statistic	Scope	Description	Notes
<code>bytes_sent</code>	Connection	Number of bytes sent from PHP to the MySQL server	Can be used to check the efficiency of the compression protocol
<code>bytes_received</code>	Connection	Number of bytes received from MySQL server	Can be used to check the efficiency of the compression protocol
<code>packets_sent</code>	Connection	Number of MySQL Client Server protocol packets sent	Used for debugging Client Server protocol implementation
<code>packets_received</code>	Connection	Number of MySQL Client Server protocol packets received	Used for debugging Client Server protocol implementation
<code>protocol_overhead_in</code>	Connection	MySQL Client Server protocol overhead in bytes for incoming traffic. Currently only the Packet Header (4 bytes) is considered as overhead. <code>protocol_overhead_in = packets_received * 4</code>	Used for debugging Client Server protocol implementation
<code>protocol_overhead_out</code>	Connection	MySQL Client Server protocol overhead in bytes for outgoing traffic. Currently only the Packet Header (4 bytes) is considered as overhead. <code>protocol_overhead_out = packets_sent * 4</code>	Used for debugging Client Server protocol implementation
<code>bytes_received_ok</code>	Connection	Total size of bytes of MySQL Client Server protocol OK packets received.	Used for debugging CS protocol implementation. Note that the total



Statistic	Scope	Description	Notes
		OK packets can contain a status message. The length of the status message can vary and thus the size of an OK packet is not fixed.	size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received</a>	Connection	Number of MySQL Client Server protocol OK packets received.	Used for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received</a>	Connection	Total size in bytes of MySQL Client Server protocol EOF packets received. EOF can vary in size depending on the server version. Also, EOF can transport an error message.	Used for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received</a>	Connection	Number of MySQL Client Server protocol EOF packets. Like with other packet statistics the number of packets will be increased even if PHP does not receive the expected packet but, for example, an error message.	Used for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received</a>	Connection	Total size in bytes of MySQL Client Server protocol result set header packets. The size of the packets varies depending on the payload ( <a href="#">LOAD LOCAL INFILE</a> , <a href="#">INSERT</a> , <a href="#">UPDATE</a> , <a href="#">SELECT</a> , error message).	Used for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received</a>	Connection	Number of MySQL Client Server protocol result set header packets.	Used for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received</a>	Connection	Total size in bytes of MySQL Client Server protocol result set meta data (field information) packets. Of course the size varies with the fields in the result set. The packet may also transport an error or an EOF packet in case of <a href="#">COM_LIST_FIELDS</a> .	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received</a>	Connection	Number of MySQL Client Server protocol result set meta data (field information) packets.	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received</a>	Connection	Total size in bytes of MySQL Client Server protocol result set row data packets. The packet may also transport an error or an EOF packet. You can reverse engineer the number of error and EOF packets by subtracting <a href="#">rows_fetched_from_server_normal</a> and <a href="#">rows_fetched_from_server_ps</a>	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).

Statistic	Scope	Description	Notes
		from <a href="#">bytes_received_rset_row_packet</a> .	
<a href="#">packets_received_rset_row_data</a>	Connection	Number of MySQL Client Server protocol result set row data packets and their total size in bytes.	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received_prepare_response</a>	Connection	Total size in bytes of MySQL Client Server protocol OK for Prepared Statement Initialization packets (prepared statement init packets). The packet may also transport an error. The packet size depends on the MySQL version: 9 bytes with MySQL 4.1 and 12 bytes from MySQL 5.0 on. There is no safe way to know how many errors happened. You may be able to guess that an error has occurred if, for example, you always connect to MySQL 5.0 or newer and, <a href="#">bytes_received_prepare_response_packet</a> != <a href="#">packets_received_prepare_response</a> * 12. See also <a href="#">ps_prepared_never_executed</a> , <a href="#">ps_prepared_once_executed</a> .	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received_prepare_init</a>	Connection	Number of MySQL Client Server protocol OK for Prepared Statement Initialization packets (prepared statement init packets).	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">bytes_received_change_user</a>	Connection	Total size in bytes of MySQL Client Server protocol COM_CHANGE_USER packets. The packet may also transport an error or EOF.	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received_change_user</a>	Connection	Number of MySQL Client Server protocol COM_CHANGE_USER packets	Only useful for debugging CS protocol implementation. Note that the total size in bytes includes the size of the header packet (4 bytes, see protocol overhead).
<a href="#">packets_received_commands</a>	Connection	Number of MySQL Client Server protocol commands sent from PHP to MySQL. There is no way to know which specific commands and how many of them have been sent. At its best you can use it to check if PHP has sent any commands to MySQL to know if you can consider to disable MySQL support in your PHP binary. There is also no way to reverse engineer the number of errors that may have occurred while sending data to MySQL. The only error that is recorded is	Only useful for debugging CS protocol implementation.

Statistic	Scope	Description	Notes
		command_buffer_too_small (see below).	
<a href="#">bytes_received</a>	Connection	Number of bytes of payload fetched by the PHP client from <a href="#">mysqlnd</a> using the text protocol.	<p>This is the size of the actual data contained in result sets that do not originate from prepared statements and which have been fetched by the PHP client. Note that although a full result set may have been pulled from MySQL by <a href="#">mysqlnd</a>, this statistic only counts actual data pulled from <a href="#">mysqlnd</a> by the PHP client. An example of a code sequence that will increase the value is as follows:</p> <pre>\$mysqli = new mysqli(); \$res = \$mysqli-&gt;query("SELECT 'abc'"); \$res-&gt;fetch_assoc(); \$res-&gt;close();</pre> <p>Every fetch operation will increase the value.</p> <p>The statistic will not be increased if the result set is only buffered on the client, but not fetched, such as in the following example:</p> <pre>\$mysqli = new mysqli(); \$res = \$mysqli-&gt;query("SELECT 'abc'"); \$res-&gt;close();</pre> <p>This statistic is available as of PHP version 5.3.4.</p>
<a href="#">bytes_received_real_data</a>	Connection	Number of bytes of the payload fetched by the PHP client from <a href="#">mysqlnd</a> using the prepared statement protocol.	<p>This is the size of the actual data contained in result sets that originate from prepared statements and which has been fetched by the PHP client. The value will not be increased if the result set is not subsequently read by the PHP client. Note that although a full result set may have been pulled from MySQL by <a href="#">mysqlnd</a>, this statistic only counts actual data pulled from <a href="#">mysqlnd</a> by the PHP client. See also <a href="#">bytes_received_real_data_normal</a>. This statistic is available as of PHP version 5.3.4.</p>

*Result Set*

**Table 21.62. Returned mysqlnd statistics: Result Set**

Statistic	Scope	Description	Notes
<a href="#">result_queries</a>	Connection	Number of queries that have generated a result set. Examples of queries that generate a result set: <a href="#">SELECT</a> , <a href="#">SHOW</a> . The statistic will not be incremented if	You may use it as an indirect measure for the number of queries PHP has sent to MySQL, for example, to identify a client that causes a high database load.

Statistic	Scope	Description	Notes
		there is an error reading the result set header packet from the line.	
<a href="#">non_res</a>	Connection	Number of queries that did not generate a result set. Examples of queries that do not generate a result set: <a href="#">INSERT</a> , <a href="#">UPDATE</a> , <a href="#">LOAD DATA</a> , <a href="#">SHOW</a> . The statistic will not be incremented if there is an error reading the result set header packet from the line.	You may use it as an indirect measure for the number of queries PHP has sent to MySQL, for example, to identify a client that causes a high database load.
<a href="#">no_index</a>	Connection	Number of queries that have generated a result set but did not use an index (see also <a href="#">mysqld</a> start option <code>–log-queries-not-using-indexes</code> ). If you want these queries to be reported you can use <code>mysqli_report(MYSQLI_REPORT_INDEX)</code> to make <code>ext/mysqli</code> throw an exception. If you prefer a warning instead of an exception use <code>mysqli_report(MYSQLI_REPORT_INDEX ^ MYSQLI_REPORT_STRICT)</code> .	
<a href="#">bad_index</a>	Connection	Number of queries that have generated a result set and did not use a good index (see also <a href="#">mysqld</a> start option <code>–log-slow-queries</code> ).	If you want these queries to be reported you can use <code>mysqli_report(MYSQLI_REPORT_INDEX)</code> to make <code>ext/mysqli</code> throw an exception. If you prefer a warning instead of an exception use <code>mysqli_report(MYSQLI_REPORT_INDEX ^ MYSQLI_REPORT_STRICT)</code>
<a href="#">slow_query</a>	Connection	SQL statements that took more than <a href="#">long_query_time</a> seconds to execute and required at least <a href="#">min_examined_row_limit</a> rows to be examined.	Not reported through <code>mysqli_report</code>
<a href="#">buffered</a>	Connection	Number of buffered result sets returned by “normal” queries. “Normal” means “not prepared statement” in the following notes.	Examples of API calls that will buffer result sets on the client: <code>mysql_query</code> , <code>mysqli_query</code> , <code>mysqli_store_result</code> , <code>mysqli_stmt_get_result</code> . Buffering result sets on the client ensures that server resources are freed as soon as possible and it makes result set scrolling easier. The downside is the additional memory consumption on the client for buffering data. Note that <code>mysqlnd</code> (unlike the MySQL Client Library) respects the PHP memory limit because it uses PHP internal memory management functions to allocate memory. This is also the reason why <code>memory_get_usage</code> reports a higher memory consumption when using <code>mysqlnd</code> instead of the MySQL Client Library. <code>memory_get_usage</code> does not

Statistic	Scope	Description	Notes
			measure the memory consumption of the MySQL Client Library at all because the MySQL Client Library does not use PHP internal memory management functions monitored by the function!
<a href="#">unbuffered</a>	Connection	Number of unbuffered result sets returned by normal (non prepared statement) queries.	Examples of API calls that will not buffer result sets on the client: <a href="#">mysqli_use_result</a>
<a href="#">ps_buffered</a>	Connection	Number of buffered result sets returned by prepared statements. By default prepared statements are unbuffered.	Examples of API calls that will not buffer result sets on the client: <a href="#">mysqli_stmt_store_result</a>
<a href="#">ps_unbuffered</a>	Connection	Number of unbuffered result sets returned by prepared statements.	By default prepared statements are unbuffered.
<a href="#">flushed_normal</a>	Connection	Number of result sets from normal (non prepared statement) queries with unread data which have been flushed silently for you. Flushing happens only with unbuffered result sets.	Unbuffered result sets must be fetched completely before a new query can be run on the connection otherwise MySQL will throw an error. If the application does not fetch all rows from an unbuffered result set, <a href="#">mysqlnd</a> does implicitly fetch the result set to clear the line. See also <a href="#">rows_skipped_normal</a> , <a href="#">rows_skipped_ps</a> . Some possible causes for an implicit flush: <ul style="list-style-type: none"> <li>• Faulty client application</li> <li>• Client stopped reading after it found what it was looking for but has made MySQL calculate more records than needed</li> <li>• Client application has stopped unexpectedly</li> </ul>
<a href="#">flushed_ps</a>	Connection	Number of result sets from prepared statements with unread data which have been flushed silently for you. Flushing happens only with unbuffered result sets.	Unbuffered result sets must be fetched completely before a new query can be run on the connection otherwise MySQL will throw an error. If the application does not fetch all rows from an unbuffered result set, <a href="#">mysqlnd</a> does implicitly fetch the result set to clear the line. See also <a href="#">rows_skipped_normal</a> , <a href="#">rows_skipped_ps</a> . Some possible causes for an implicit flush: <ul style="list-style-type: none"> <li>• Faulty client application</li> <li>• Client stopped reading after it found what it was looking for but has made MySQL calculate more records than needed</li> <li>• Client application has stopped unexpectedly</li> </ul>

Statistic	Scope	Description	Notes
<a href="#">ps_prepared_statements</a>	Connection	Number of statements prepared but never executed.	Prepared statements occupy server resources. You should not prepare a statement if you do not plan to execute it.
<a href="#">ps_prepared_statements_executed</a>	Connection	Number of prepared statements executed only once.	One of the ideas behind prepared statements is that the same query gets executed over and over again (with different parameters) and some parsing and other preparation work can be saved, if statement execution is split up in separate prepare and execute stages. The idea is to prepare once and “cache” results, for example, the parse tree to be reused during multiple statement executions. If you execute a prepared statement only once the two stage processing can be inefficient compared to “normal” queries because all the caching means extra work and it takes (limited) server resources to hold the cached information. Consequently, prepared statements that are executed only once may cause performance hurts.
<a href="#">rows_fetched_from_result_set</a> <a href="#">rows_fetched_from_result_set_implicit</a>	Connection	Total number of result set rows successfully fetched from MySQL regardless if the client application has consumed them or not. Some of the rows may not have been fetched by the client application but have been flushed implicitly.	See also <a href="#">packets_received_rset_row</a>
<a href="#">rows_buffered_fetching</a> <a href="#">rows_buffered_fetching_implicit</a>	Connection	Total number of successfully buffered rows originating from a “normal” query or a prepared statement. This is the number of rows that have been fetched from MySQL and buffered on client. Note that there are two distinct statistics on rows that have been buffered (MySQL to <i>mysqld</i> internal buffer) and buffered rows that have been fetched by the client application ( <i>mysqld</i> internal buffer to client application). If the number of buffered rows is higher than the number of fetched buffered rows it can mean that the client application runs queries that cause larger result sets than needed resulting in rows not read by the client.	Examples of queries that will buffer results: <a href="#">mysql_query</a> , <a href="#">mysql_store_result</a>
<a href="#">rows_fetched_from_result_set_buffered</a> <a href="#">rows_fetched_from_result_set_buffered_implicit</a>	Connection	Total number of rows fetched by the client from a buffered result set created by a normal query or a prepared statement.	
<a href="#">rows_fetched_from_result_set_unbuffered</a> <a href="#">rows_fetched_from_result_set_unbuffered_implicit</a>	Connection	Total number of rows fetched by the client from a unbuffered result set	

Statistic	Scope	Description	Notes
		created by a "normal" query or a prepared statement.	
<code>rows_fetched</code>	Connection	Total number of rows fetched by the client from a cursor created by a prepared statement.	
<code>rows_skipped</code> <code>rows_skipped_ps</code>	Connection	Reserved for future use (currently not supported)	
<code>copy_on_write_enabled</code> <code>copy_on_write_percent</code>	Process	With mysqlnd, variables returned by the extensions point into mysqlnd internal network result buffers. If you do not change the variables, fetched data will be kept only once in memory. If you change the variables, mysqlnd has to perform a copy-on-write to protect the internal network result buffers from being changed. With the MySQL Client Library you always hold fetched data twice in memory. Once in the internal MySQL Client Library buffers and once in the variables returned by the extensions. In theory mysqlnd can save up to 40% memory. However, note that the memory saving cannot be measured using <code>memory_get_usage</code> .	
<code>explicit_result_set_free</code> <code>implicit_result_set_free</code>	Connection Process (only during prepared statement cleanup)	Total number of freed result sets.	The free is always considered explicit but for result sets created by an init command, for example, <code>mysql_options(MYSQLI_INIT_COMMAND, ...)</code>
<code>proto_text_fetched</code> <code>proto_text_fetched_int</code> <code>proto_text_fetched_short</code> <code>proto_text_fetched_int24</code> <code>proto_text_fetched_int</code> <code>proto_text_fetched_bigint</code> <code>proto_text_fetched_decimal</code> <code>proto_text_fetched_float</code> <code>proto_text_fetched_double</code> <code>proto_text_fetched_date</code> <code>proto_text_fetched_year</code> <code>proto_text_fetched_time</code> <code>proto_text_fetched_datetime</code> <code>proto_text_fetched_timestamp</code> <code>proto_text_fetched_string</code> <code>proto_text_fetched_blob</code> <code>proto_text_fetched_enum</code> <code>proto_text_fetched_set</code> <code>proto_text_fetched_geometry</code> <code>proto_text_fetched_other</code>	Connection	Total number of columns of a certain type fetched from a normal query (MySQL text protocol).	Mapping from C API / MySQL meta data type to statistics name: <ul style="list-style-type: none"> <li><code>MYSQL_TYPE_NULL</code> - <code>proto_text_fetched_null</code></li> <li><code>MYSQL_TYPE_BIT</code> - <code>proto_text_fetched_bit</code></li> <li><code>MYSQL_TYPE_TINY</code> - <code>proto_text_fetched_tinyint</code></li> <li><code>MYSQL_TYPE_SHORT</code> - <code>proto_text_fetched_short</code></li> <li><code>MYSQL_TYPE_INT24</code> - <code>proto_text_fetched_int24</code></li> <li><code>MYSQL_TYPE_LONG</code> - <code>proto_text_fetched_int</code></li> <li><code>MYSQL_TYPE_LONGLONG</code> - <code>proto_text_fetched_bigint</code></li> </ul>



Statistic	Scope	Description	Notes
			<ul style="list-style-type: none"> <li>• <code>MYSQL_TYPE_DECIMAL</code>, <code>MYSQL_TYPE_NEWDECIMAL</code> - <code>proto_text_fetched_decimal</code></li> <li>• <code>MYSQL_TYPE_FLOAT</code> - <code>proto_text_fetched_float</code></li> <li>• <code>MYSQL_TYPE_DOUBLE</code> - <code>proto_text_fetched_double</code></li> <li>• <code>MYSQL_TYPE_DATE</code>, <code>MYSQL_TYPE_NEWDATE</code> - <code>proto_text_fetched_date</code></li> <li>• <code>MYSQL_TYPE_YEAR</code> - <code>proto_text_fetched_year</code></li> <li>• <code>MYSQL_TYPE_TIME</code> - <code>proto_text_fetched_time</code></li> <li>• <code>MYSQL_TYPE_DATETIME</code> - <code>proto_text_fetched_datetime</code></li> <li>• <code>MYSQL_TYPE_TIMESTAMP</code> - <code>proto_text_fetched_timestamp</code></li> <li>• <code>MYSQL_TYPE_STRING</code>, <code>MYSQL_TYPE_VARSTRING</code>, <code>MYSQL_TYPE_VARCHAR</code> - <code>proto_text_fetched_string</code></li> <li>• <code>MYSQL_TYPE_TINY_BLOB</code>, <code>MYSQL_TYPE_MEDIUM_BLOB</code>, <code>MYSQL_TYPE_LONG_BLOB</code>, <code>MYSQL_TYPE_BLOB</code> - <code>proto_text_fetched_blob</code></li> <li>• <code>MYSQL_TYPE_ENUM</code> - <code>proto_text_fetched_enum</code></li> <li>• <code>MYSQL_TYPE_SET</code> - <code>proto_text_fetched_set</code></li> <li>• <code>MYSQL_TYPE_GEOMETRY</code> - <code>proto_text_fetched_geometry</code></li> <li>• Any <code>MYSQL_TYPE_*</code> not listed before (there should be none) - <code>proto_text_fetched_other</code></li> </ul> <p>Note that the <code>MYSQL_*</code>-type constants may not be associated with the very same SQL column types in every version of MySQL.</p>
<code>proto_binary_fetched</code> , <code>proto_binary_fetched_long</code> , <code>proto_binary_fetched_medium</code> , <code>proto_binary_fetched_short</code> ,	<b>Connection</b>	<b>Total number</b> , of columns of a certain type fetched from a prepared statement (MySQL binary protocol).	For type mapping see <code>proto_text_*</code> described in the preceding text.



Statistic	Scope	Description	Notes
proto_binary_fetched_int24, proto_binary_fetched_int, proto_binary_fetched_bigint, proto_binary_fetched_decimal, proto_binary_fetched_float, proto_binary_fetched_double, proto_binary_fetched_date, proto_binary_fetched_year, proto_binary_fetched_time, proto_binary_fetched_datetime, proto_binary_fetched_timestamp, proto_binary_fetched_string, proto_binary_fetched_blob, proto_binary_fetched_enum, proto_binary_fetched_set, proto_binary_fetched_geometry, proto_binary_fetched_other			

**Table 21.63. Returned mysqlnd statistics: Connection**

Statistic	Scope	Description	Notes
<code>connect_success</code> , <code>connect_failure</code>	Connection	Total number of successful / failed connection attempt.	Reused connections and all other kinds of connections are included.
<code>reconnect</code>	Process	Total number of (real_)connect attempts made on an already opened connection handle.	The code sequence <code>\$link = new mysqli(...); \$link-&gt;real_connect(...)</code> will cause a reconnect. But <code>\$link = new mysqli(...); \$link-&gt;connect(...)</code> will not because <code>\$link-&gt;connect(...)</code> will explicitly close the existing connection before a new connection is established.
<code>pconnect_success</code>	Connection	Total number of successful persistent connection attempts.	Note that <code>connect_success</code> holds the sum of successful persistent and non-persistent connection attempts. The number of successful non-persistent connection attempts is <code>connect_success - pconnect_success</code> .
<code>active_connections</code>	Connection	Total number of active persistent and non-persistent connections.	
<code>active_persistent_connections</code>	Connection	Total number of active persistent connections.	The total number of active non-persistent connections is <code>active_connections - active_persistent_connections</code> .
<code>explicit_closes</code>	Connection	Total number of explicitly closed connections (ext/mysqli only).	Examples of code snippets that cause an explicit close :  <pre>\$link = new mysqli(...); \$link-&gt;close(...) \$link = new mysqli(...); \$link-&gt;connect(...)</pre>
<code>implicit_closes</code>	Connection	Total number of implicitly closed connections (ext/mysqli only).	Examples of code snippets that cause an implicit close :

Statistic	Scope	Description	Notes
			<ul style="list-style-type: none"> <li><code>\$link = new mysqli(...);</code> <code>\$link-&gt;real_connect(...)</code></li> <li><code>unset(\$link)</code></li> <li>Persistent connection: pooled connection has been created with <code>real_connect</code> and there may be unknown options set - close implicitly to avoid returning a connection with unknown options</li> <li>Persistent connection: <code>ping/change_user</code> fails and <code>ext/mysqli</code> closes the connection</li> <li>end of script execution: close connections that have not been closed by the user</li> </ul>
<code>disconnect</code>	Connection	Connection failures indicated by the C API call <code>mysql_real_connect</code> during an attempt to establish a connection.	It is called <code>disconnect_close</code> because the connection handle passed to the C API call will be closed.
<code>in_middle</code>	Process	A connection has been closed in the middle of a command execution (outstanding result sets not fetched, after sending a query and before retrieving an answer, while fetching data, while transferring data with LOAD DATA).	Unless you use asynchronous queries this should only happen if your script stops unexpectedly and PHP shuts down the connections for you.
<code>init_command</code>	Connection	Total number of init command executions, for example, <code>mysqli_options(MYSQLI_INIT_COMMAND)</code>	The number of successful executions is <code>init_command_executed_count - init_command_failed_count</code> .
<code>init_command</code>	Connection	Total number of failed init commands.	

**Table 21.64. Returned mysqlnd statistics: COM\_\* Command**

Statistic	Scope	Description	Notes
<code>com_quit</code> , <code>com_init_db</code> , <code>com_query</code> , <code>com_field_list</code> , <code>com_create_db</code> , <code>com_drop_db</code> , <code>com_refresh</code> , <code>com_shutdown</code> , <code>com_statistics</code> , <code>com_process_info</code> , <code>com_connect</code> , <code>com_process_kill</code> , <code>com_debug</code> , <code>com_ping</code> , <code>com_time</code> , <code>com_delayed_insert</code> , <code>com_change_user</code> , <code>com_binlog_dump</code> ,	Connection	Total number of attempts to send a certain COM_* command from PHP to MySQL.	<p>The statistics are incremented after checking the line and immediately before sending the corresponding MySQL client server protocol packet. If <code>mysqlnd</code> fails to send the packet over the wire the statistics will not be decremented. In case of a failure <code>mysqlnd</code> emits a PHP warning "Error while sending %s packet. PID=%d."</p> <p>Usage examples:</p> <ul style="list-style-type: none"> <li>Check if PHP sends certain commands to MySQL, for example, check if a client sends <code>COM_PROCESS_KILL</code></li> </ul>

Statistic	Scope	Description	Notes
<code>com_table_dump</code> , <code>com_connect_out</code> , <code>com_register_slave</code> , <code>com_stmt_prepare</code> , <code>com_stmt_execute</code> , <code>com_stmt_send_long_data</code> , <code>com_stmt_close</code> , <code>com_stmt_reset</code> , <code>com_stmt_set_option</code> , <code>com_stmt_fetch</code> , <code>com_daemon</code>			<ul style="list-style-type: none"> <li>Calculate the average number of prepared statement executions by comparing <code>COM_EXECUTE</code> with <code>COM_PREPARE</code></li> <li>Check if PHP has run any non-prepared SQL statements by checking if <code>COM_QUERY</code> is zero</li> <li>Identify PHP scripts that run an excessive number of SQL statements by checking <code>COM_QUERY</code> and <code>COM_EXECUTE</code></li> </ul>

*Miscellaneous*

**Table 21.65. Returned mysqlnd statistics: Miscellaneous**

Statistic	Scope	Description	Notes
<code>explicit_close</code> , <code>implicit_stmt_close</code>	Process	Total number of close prepared statements.	A close is always considered explicit but for a failed prepare.
<code>mem_emalloc_count</code> , <code>mem_emalloc_ammount</code> , <code>mem_ecalloc_count</code> , <code>mem_ecalloc_ammount</code> , <code>mem_erealloc_count</code> , <code>mem_erealloc_ammount</code> , <code>mem_efree_count</code> , <code>mem_malloc_count</code> , <code>mem_malloc_ammount</code> , <code>mem_calloc_count</code> , <code>mem_calloc_ammount</code> , <code>mem_realloc_count</code> , <code>mem_realloc_ammount</code> , <code>mem_free_count</code>	Process	Memory management calls.	Development only.
<code>command_buffer_too_small</code>	Connection	Number of network command buffer extensions while sending commands from PHP to MySQL.	<p>mysqlnd allocates an internal command/network buffer of <code>mysqlnd.net_cmd_buffer_size</code> (<code>php.ini</code>) bytes for every connection. If a MySQL Client Server protocol command, for example, <code>COM_QUERY</code> (normal query), does not fit into the buffer, mysqlnd will grow the buffer to what is needed for sending the command. Whenever the buffer gets extended for one connection <code>command_buffer_too_small</code> will be incremented by one.</p> <p>If mysqlnd has to grow the buffer beyond its initial size of <code>mysqlnd.net_cmd_buffer_size</code> (<code>php.ini</code>) bytes for almost every connection, you should consider to increase the default size to avoid re-allocations.</p>

Statistic	Scope	Description	Notes
			<p>The default buffer size is 2048 bytes in PHP 5.3.0. In future versions the default will be 4kB or larger. The default can be changed either through the <code>php.ini</code> setting <code>mysqlnd.net_cmd_buffer_size</code> or using <code>mysqli_options(MYSQLI_OPT_NET_CMD_BUFFER_SIZE, int size)</code>.</p> <p>It is recommended to set the buffer size to no less than 4096 bytes because <code>mysqlnd</code> also uses it when reading certain communication packet from MySQL. In PHP 5.3.0, <code>mysqlnd</code> will not grow the buffer if MySQL sends a packet that is larger than the current size of the buffer. As a consequence <code>mysqlnd</code> is unable to decode the packet and the client application will get an error. There are only two situations when the packet can be larger than the 2048 bytes default of <code>mysqlnd.net_cmd_buffer_size</code> in PHP 5.3.0: the packet transports a very long error message or the packet holds column meta data from <code>COM_LIST_FIELD</code> (<code>mysql_list_fields</code>) and the meta data comes from a string column with a very long default value (&gt;1900 bytes). No bug report on this exists - it should happen rarely.</p> <p>As of PHP 5.3.2 <code>mysqlnd</code> does not allow setting buffers smaller than 4096 bytes.</p>
<code>connection_reused</code>			

### 21.9.5.6. Notes

Copyright 1997-2012 the PHP Documentation Group. [2906]

This section provides a collection of miscellaneous notes on MySQL Native Driver usage.

- Using `mysqlnd` means using PHP streams for underlying connectivity. For `mysqlnd`, the PHP streams documentation (<http://www.php.net/manual/en/book.stream>) should be consulted on such details as timeout settings, not the documentation for the MySQL Client Library.

### 21.9.5.7. MySQL Native Driver Plugin API

Copyright 1997-2012 the PHP Documentation Group. [2906]

The MySQL Native Driver Plugin API is a feature of MySQL Native Driver, or `mysqlnd`. `mysqlnd` plugins operate in the layer between PHP applications and the MySQL server. This is comparable to MySQL Proxy. MySQL Proxy operates on a layer between any MySQL client application, for example, a PHP application and, the MySQL server. `mysqlnd` plugins can undertake typical MySQL Proxy tasks such as load balancing, monitoring and performance optimizations. Due to the different architecture

and location, `mysqlnd` plugins do not have some of MySQL Proxy's disadvantages. For example, with plugins, there is no single point of failure, no dedicated proxy server to deploy, and no new programming language to learn (Lua).

A `mysqlnd` plugin can be thought of as an extension to `mysqlnd`. Plugins can intercept the majority of `mysqlnd` functions. The `mysqlnd` functions are called by the PHP MySQL extensions such as `ext/mysql`, `ext/mysqli`, and `PDO_MYSQL`. As a result, it is possible for a `mysqlnd` plugin to intercept all calls made to these extensions from the client application.

Internal `mysqlnd` function calls can also be intercepted, or replaced. There are no restrictions on manipulating `mysqlnd` internal function tables. It is possible to set things up so that when certain `mysqlnd` functions are called by the extensions that use `mysqlnd`, the call is directed to the appropriate function in the `mysqlnd` plugin. The ability to manipulate `mysqlnd` internal function tables in this way allows maximum flexibility for plugins.

`mysqlnd` plugins are in fact PHP Extensions, written in C, that use the `mysqlnd` plugin API (which is built into MySQL Native Driver, `mysqlnd`). Plugins can be made 100% transparent to PHP applications. No application changes are needed because plugins operate on a different layer. The `mysqlnd` plugin can be thought of as operating in a layer below `mysqlnd`.

The following list represents some possible applications of `mysqlnd` plugins.

- Load Balancing
  - Read/Write Splitting. An example of this is the `PECL/mysqlnd_ms` (Master Slave) extension. This extension splits read/write queries for a replication setup.
  - Failover
  - Round-Robin, least loaded
- Monitoring
  - Query Logging
  - Query Analysis
  - Query Auditing. An example of this is the `PECL/mysqlnd_sip` (SQL Injection Protection) extension. This extension inspects queries and executes only those that are allowed according to a ruleset.
- Performance
  - Caching. An example of this is the `PECL/mysqlnd_qc` (Query Cache) extension.
  - Throttling
  - Sharding. An example of this is the `PECL/mysqlnd_mc` (Multi Connect) extension. This extension will attempt to split a `SELECT` statement into `n`-parts, using `SELECT ... LIMIT part_1, SELECT LIMIT part_n`. It sends the queries to distinct MySQL servers and merges the result at the client.

#### *MySQL Native Driver Plugins Available*

There are a number of `mysqlnd` plugins already available. These include:

- `PECL/mysqlnd_mc` - Multi Connect plugin.
- `PECL/mysqlnd_ms` - Master Slave plugin.
- `PECL/mysqlnd_qc` - Query Cache plugin.
- `PECL/mysqlnd_pscache` - Prepared Statement Handle Cache plugin.
- `PECL/mysqlnd_sip` - SQL Injection Protection plugin.

- [PECL/mysqlnd\\_uh](#) - User Handler plugin.

#### 21.9.5.7.1. A comparison of `mysqlnd` plugins with MySQL Proxy

Copyright 1997-2012 the PHP Documentation Group. [2906]

`mysqlnd` plugins and MySQL Proxy are different technologies using different approaches. Both are valid tools for solving a variety of common tasks such as load balancing, monitoring, and performance enhancements. An important difference is that MySQL Proxy works with all MySQL clients, whereas `mysqlnd` plugins are specific to PHP applications.

As a PHP Extension, a `mysqlnd` plugin gets installed on the PHP application server, along with the rest of PHP. MySQL Proxy can either be run on the PHP application server or can be installed on a dedicated machine to handle multiple PHP application servers.

Deploying MySQL Proxy on the application server has two advantages:

1. No single point of failure
2. Easy to scale out (horizontal scale out, scale by client)

MySQL Proxy (and `mysqlnd` plugins) can solve problems easily which otherwise would have required changes to existing applications.

However, MySQL Proxy does have some disadvantages:

- MySQL Proxy is a new component and technology to master and deploy.
- MySQL Proxy requires knowledge of the Lua scripting language.

MySQL Proxy can be customized with C and Lua programming. Lua is the preferred scripting language of MySQL Proxy. For most PHP experts Lua is a new language to learn. A `mysqlnd` plugin can be written in C. It is also possible to write plugins in PHP using [PECL/mysqlnd\\_uh](#).

MySQL Proxy runs as a daemon - a background process. MySQL Proxy can recall earlier decisions, as all state can be retained. However, a `mysqlnd` plugin is bound to the request-based lifecycle of PHP. MySQL Proxy can also share one-time computed results among multiple application servers. A `mysqlnd` plugin would need to store data in a persistent medium to be able to do this. Another daemon would need to be used for this purpose, such as Memcache. This gives MySQL Proxy an advantage in this case.

MySQL Proxy works on top of the wire protocol. With MySQL Proxy you have to parse and reverse engineer the MySQL Client Server Protocol. Actions are limited to those that can be achieved by manipulating the communication protocol. If the wire protocol changes (which happens very rarely) MySQL Proxy scripts would need to be changed as well.

`mysqlnd` plugins work on top of the C API, which mirrors the `libmysqlclient` client and Connector/C APIs. This C API is basically a wrapper around the MySQL Client Server protocol, or wire protocol, as it is sometimes called. You can intercept all C API calls. PHP makes use of the C API, therefore you can hook all PHP calls, without the need to program at the level of the wire protocol.

`mysqlnd` implements the wire protocol. Plugins can therefore parse, reverse engineer, manipulate and even replace the communication protocol. However, this is usually not required.

As plugins allow you to create implementations that use two levels (C API and wire protocol), they have greater flexibility than MySQL Proxy. If a `mysqlnd` plugin is implemented using the C API, any subsequent changes to the wire protocol do not require changes to the plugin itself.

#### 21.9.5.7.2. Obtaining the `mysqlnd` plugin API

Copyright 1997-2012 the PHP Documentation Group. [2906]

The [mysqlnd](#) plugin API is simply part of the MySQL Native Driver PHP extension, [ext/mysqlnd](#). Development started on the [mysqlnd](#) plugin API in December 2009. It is developed as part of the PHP source repository, and as such is available to the public either via Git, or through source snapshot downloads.

The following table shows PHP versions and the corresponding [mysqlnd](#) version contained within.

**Table 21.66. The bundled mysqlnd version per PHP release**

PHP Version	MySQL Native Driver version
5.3.0	5.0.5
5.3.1	5.0.5
5.3.2	5.0.7
5.3.3	5.0.7
5.3.4	5.0.7

Plugin developers can determine the [mysqlnd](#) version through accessing `MYSQLND_VERSION`, which is a string of the format “mysqlnd 5.0.7-dev - 091210 - \$Revision: 300535”, or through `MYSQLND_VERSION_ID`, which is an integer such as 50007. Developers can calculate the version number as follows:

**Table 21.67. MYSQLND\_VERSION\_ID calculation table**

Version (part)	Example
Major*10000	5*10000 = 50000
Minor*100	0*100 = 0
Patch	7 = 7
MYSQLND_VERSION_ID	50007

During development, developers should refer to the [mysqlnd](#) version number for compatibility and version tests, as several iterations of [mysqlnd](#) could occur during the lifetime of a PHP development branch with a single PHP version number.

### 21.9.5.7.3. MySQL Native Driver Plugin Architecture

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

This section provides an overview of the [mysqlnd](#) plugin architecture.

*MySQL Native Driver Overview*

Before developing [mysqlnd](#) plugins, it is useful to know a little of how [mysqlnd](#) itself is organized. [Mysqlnd](#) consists of the following modules:

**Table 21.68. The mysqlnd organization chart, per module**

Modules Statistics	mysqlnd_statistics.c
Connection	mysqlnd.c
Resultset	mysqlnd_result.c
Resultset Metadata	mysqlnd_result_meta.c
Statement	mysqlnd_ps.c
Network	mysqlnd_net.c
Wire protocol	mysqlnd_wireprotocol.c

*C Object Oriented Paradigm*



At the code level, *mysqlnd* uses a C pattern for implementing object orientation.

In C you use a *struct* to represent an object. Members of the struct represent object properties. Struct members pointing to functions represent methods.

Unlike with other languages such as C++ or Java, there are no fixed rules on inheritance in the C object oriented paradigm. However, there are some conventions that need to be followed that will be discussed later.

### *The PHP Life Cycle*

When considering the PHP life cycle there are two basic cycles:

- PHP engine startup and shutdown cycle
- Request cycle

When the PHP engine starts up it will call the module initialization (MINIT) function of each registered extension. This allows each module to setup variables and allocate resources that will exist for the lifetime of the PHP engine process. When the PHP engine shuts down it will call the module shutdown (MSHUTDOWN) function of each extension.

During the lifetime of the PHP engine it will receive a number of requests. Each request constitutes another life cycle. On each request the PHP engine will call the request initialization function of each extension. The extension can perform any variable setup and resource allocation required for request processing. As the request cycle ends the engine calls the request shutdown (RSHUTDOWN) function of each extension so the extension can perform any cleanup required.

### *How a plugin works*

A *mysqlnd* plugin works by intercepting calls made to *mysqlnd* by extensions that use *mysqlnd*. This is achieved by obtaining the *mysqlnd* function table, backing it up, and replacing it by a custom function table, which calls the functions of the plugin as required.

The following code shows how the *mysqlnd* function table is replaced:

```
/* a place to store original function table */
struct st_mysqlnd_conn_methods org_methods;

void minit_register_hooks(TSRMLS_D) {
    /* active function table */
    struct st_mysqlnd_conn_methods * current_methods
        = mysqlnd_conn_get_methods();

    /* backup original function table */
    memcpy(&org_methods, current_methods,
        sizeof(struct st_mysqlnd_conn_methods));

    /* install new methods */
    current_methods->query = MYSQLND_METHOD(my_conn_class, query);
}
```

Connection function table manipulations must be done during Module Initialization (MINIT). The function table is a global shared resource. In an multi-threaded environment, with a TSRM build, the manipulation of a global shared resource during the request processing will almost certainly result in conflicts.



#### **Note**

Do not use any fixed-size logic when manipulating the *mysqlnd* function table: new methods may be added at the end of the function table. The function table may change at any time in the future.



*Calling parent methods*

If the original function table entries are backed up, it is still possible to call the original function table entries - the parent methods.

In some cases, such as for `Connection::stmt_init()`, it is vital to call the parent method prior to any other activity in the derived method.

```
MYSQLND_METHOD(my_conn_class, query)(MYSQLND *conn,
    const char *query, unsigned int query_len TSRMLS_DC) {

    php_printf("my_conn_class::query(query = %s)\n", query);

    query = "SELECT 'query rewritten' FROM DUAL";
    query_len = strlen(query);

    return org_methods.query(conn, query, query_len); /* return with call to parent */
}
```

*Extending properties*

A `mysqlnd` object is represented by a C struct. It is not possible to add a member to a C struct at run time. Users of `mysqlnd` objects cannot simply add properties to the objects.

Arbitrary data (properties) can be added to a `mysqlnd` objects using an appropriate function of the `mysqlnd_plugin_get_plugin_<object>_data()` family. When allocating an object `mysqlnd` reserves space at the end of the object to hold a `void *` pointer to arbitrary data. `mysqlnd` reserves space for one `void *` pointer per plugin.

The following table shows how to calculate the position of the pointer for a specific plugin:

**Table 21.69. Pointer calculations for `mysqlnd`**

Memory address	Contents
0	Beginning of the <code>mysqlnd</code> object C struct
n	End of the <code>mysqlnd</code> object C struct
n + (m x sizeof(void*))	<code>void*</code> to object data of the m-th plugin

If you plan to subclass any of the `mysqlnd` object constructors, which is allowed, you must keep this in mind!

The following code shows extending properties:

```
/* any data we want to associate */
typedef struct my_conn_properties {
    unsigned long query_counter;
} MY_CONN_PROPERTIES;

/* plugin id */
unsigned int my_plugin_id;

void minit_register_hooks(TSRMLS_D) {
    /* obtain unique plugin ID */
    my_plugin_id = mysqlnd_plugin_register();
    /* snip - see Extending Connection: methods */
}

static MY_CONN_PROPERTIES** get_conn_properties(const MYSQLND *conn TSRMLS_DC) {
    MY_CONN_PROPERTIES** props;
    props = (MY_CONN_PROPERTIES**)mysqlnd_plugin_get_plugin_connection_data(
        conn, my_plugin_id);
    if (!props || !(*props)) {
```

```

*props = mnd_pecalloc(1, sizeof(MY_CONN_PROPERTIES), conn->persistent);
(*props)->query_counter = 0;
}
return props;
}

```

The plugin developer is responsible for the management of plugin data memory.

Use of the *mysqlnd* memory allocator is recommended for plugin data. These functions are named using the convention: `mnd_*loc()`. The *mysqlnd* allocator has some useful features, such as the ability to use a debug allocator in a non-debug build.

**Table 21.70. When and how to subclass**

	When to subclass?	Each instance has its own private function table?	How to subclass?
Connection (MYSQLND)	MINIT	No	<code>mysqlnd_conn_get_methods()</code>
Resultset (MYSQLND_RES)	MINIT or later	Yes	<code>mysqlnd_result_get_methods()</code> or object method function table manipulation
Resultset Meta (MYSQLND_RES_METADATA)	MINIT	No	<code>mysqlnd_result_metadata_get_methods()</code>
Statement (MYSQLND_STMT)	MINIT	No	<code>mysqlnd_stmt_get_methods()</code>
Network (MYSQLND_NET)	MINIT or later	Yes	<code>mysqlnd_net_get_methods()</code> or object method function table manipulation
Wire protocol (MYSQLND_PROTOCOL)	MINIT or later	Yes	<code>mysqlnd_protocol_get_methods()</code> or object method function table manipulation

You must not manipulate function tables at any time later than MINIT if it is not allowed according to the above table.

Some classes contain a pointer to the method function table. All instances of such a class will share the same function table. To avoid chaos, in particular in threaded environments, such function tables must only be manipulated during MINIT.

Other classes use copies of a globally shared function table. The class function table copy is created together with the object. Each object uses its own function table. This gives you two options: you can manipulate the default function table of an object at MINIT, and you can additionally refine methods of an object without impacting other instances of the same class.

The advantage of the shared function table approach is performance. There is no need to copy a function table for each and every object.

**Table 21.71. Constructor status**

	Allocation, construction, reset	Can be modified?	Caller
Connection (MYSQLND)	<code>mysqlnd_init()</code>	No	<code>mysqlnd_connect()</code>
Resultset (MYSQLND_RES)	Allocation:	Yes, but call parent!	<ul style="list-style-type: none"> <li>• <code>Connection::list_fields()</code></li> </ul>

	<ul style="list-style-type: none"> <li>• Connection::result_init()</li> </ul> <p>Reset and re-initialized during:</p> <ul style="list-style-type: none"> <li>• Result::use_result()</li> <li>• Result::store_result</li> </ul>		<ul style="list-style-type: none"> <li>• Statement::get_result()</li> <li>• Statement::prepare() (Metadata only)</li> <li>• Statement::resultMetadata()</li> </ul>
Resultset Meta (MYSQLND_RES_METADATA)	Connection::result_metadata()	Yes, but call parent!	Result::read_result_metadata()
Statement (MYSQLND_STMT)	Connection::stmt_init()	Yes, but call parent!	Connection::stmt_init()
Network (MYSQLND_NET)	mysqlnd_net_init()	No	Connection::init()
Wire protocol (MYSQLND_PROTOCOL)	mysqlnd_protocol_init()	No	Connection::init()

It is strongly recommended that you do not entirely replace a constructor. The constructors perform memory allocations. The memory allocations are vital for the *mysqlnd* plugin API and the object logic of *mysqlnd*. If you do not care about warnings and insist on hooking the constructors, you should at least call the parent constructor before doing anything in your constructor.

Regardless of all warnings, it can be useful to subclass constructors. Constructors are the perfect place for modifying the function tables of objects with non-shared object tables, such as Resultset, Network, Wire Protocol.

**Table 21.72. Destruction status**

	Derived method must call parent?	Destructor
Connection	yes, after method execution	free_contents(), end_psession()
Resultset	yes, after method execution	free_result()
Resultset Meta	yes, after method execution	free()
Statement	yes, after method execution	dtor(), free_stmt_content()
Network	yes, after method execution	free()
Wire protocol	yes, after method execution	free()

The destructors are the appropriate place to free properties, `mysqlnd_plugin_get_plugin_<object>_data()`.

The listed destructors may not be equivalent to the actual *mysqlnd* method freeing the object itself. However, they are the best possible place for you to hook in and free your plugin data. As with constructors you may replace the methods entirely but this is not recommended. If multiple methods are listed in the above table you will need to hook all of the listed methods and free your plugin data in whichever method is called first by *mysqlnd*.

The recommended method for plugins is to simply hook the methods, free your memory and call the parent implementation immediately following this.



**Caution**

Due to a bug in PHP versions 5.3.0 to 5.3.3, plugins do not associate plugin data with a persistent connection. This is because `ext/mysql` and `ext/mysqli` do not trigger all the necessary `mysqlnd_end_psession()` method calls and the plugin may therefore leak memory. This has been fixed in PHP 5.3.4.

**21.9.5.7.4. The *mysqlnd* plugin API**

Copyright 1997-2012 the PHP Documentation Group. [2906]

The following is a list of functions provided in the *mysqlnd* plugin API:

- *mysqlnd\_plugin\_register()*
- *mysqlnd\_plugin\_count()*
- *mysqlnd\_plugin\_get\_plugin\_connection\_data()*
- *mysqlnd\_plugin\_get\_plugin\_result\_data()*
- *mysqlnd\_plugin\_get\_plugin\_stmt\_data()*
- *mysqlnd\_plugin\_get\_plugin\_net\_data()*
- *mysqlnd\_plugin\_get\_plugin\_protocol\_data()*
- *mysqlnd\_conn\_get\_methods()*
- *mysqlnd\_result\_get\_methods()*
- *mysqlnd\_result\_meta\_get\_methods()*
- *mysqlnd\_stmt\_get\_methods()*
- *mysqlnd\_net\_get\_methods()*
- *mysqlnd\_protocol\_get\_methods()*

There is no formal definition of what a plugin is and how a plugin mechanism works.

Components often found in plugins mechanisms are:

- A plugin manager
- A plugin API
- Application services (or modules)
- Application service APIs (or module APIs)

The *mysqlnd* plugin concept employs these features, and additionally enjoys an open architecture.

*No Restrictions*

A plugin has full access to the inner workings of *mysqlnd*. There are no security limits or restrictions. Everything can be overwritten to implement friendly or hostile algorithms. It is recommended you only deploy plugins from a trusted source.

As discussed previously, plugins can use pointers freely. These pointers are not restricted in any way, and can point into another plugin's data. Simple offset arithmetic can be used to read another plugin's data.

It is recommended that you write cooperative plugins, and that you always call the parent method. The plugins should always cooperate with *mysqlnd* itself.

**Table 21.73. Issues: an example of chaining and cooperation**

Extension	<i>mysqlnd.query()</i> pointer	call stack if calling parent
ext/ <i>mysqlnd</i>	<i>mysqlnd.query()</i>	<i>mysqlnd.query</i>

<code>ext/mysqlnd_cache</code>	<code>mysqlnd_cache.query()</code>	<ol style="list-style-type: none"> <li>1. <code>mysqlnd_cache.query()</code></li> <li>2. <code>mysqlnd.query</code></li> </ol>
<code>ext/mysqlnd_monitor</code>	<code>mysqlnd_monitor.query()</code>	<ol style="list-style-type: none"> <li>1. <code>mysqlnd_monitor.query()</code></li> <li>2. <code>mysqlnd_cache.query()</code></li> <li>3. <code>mysqlnd.query</code></li> </ol>

In this scenario, a cache (`ext/mysqlnd_cache`) and a monitor (`ext/mysqlnd_monitor`) plugin are loaded. Both subclass `Connection::query()`. Plugin registration happens at `MINIT` using the logic shown previously. PHP calls extensions in alphabetical order by default. Plugins are not aware of each other and do not set extension dependencies.

By default the plugins call the parent implementation of the query method in their derived version of the method.

### PHP Extension Recap

This is a recap of what happens when using an example plugin, `ext/mysqlnd_plugin`, which exposes the `mysqlnd` C plugin API to PHP:

- Any PHP MySQL application tries to establish a connection to 192.168.2.29
- The PHP application will either use `ext/mysql`, `ext/mysqli` or `PDO_MYSQL`. All three PHP MySQL extensions use `mysqlnd` to establish the connection to 192.168.2.29.
- `Mysqlnd` calls its connect method, which has been subclassed by `ext/mysqlnd_plugin`.
- `ext/mysqlnd_plugin` calls the userspace hook `proxy::connect()` registered by the user.
- The userspace hook changes the connection host IP from 192.168.2.29 to 127.0.0.1 and returns the connection established by `parent::connect()`.
- `ext/mysqlnd_plugin` performs the equivalent of `parent::connect(127.0.0.1)` by calling the original `mysqlnd` method for establishing a connection.
- `ext/mysqlnd` establishes a connection and returns to `ext/mysqlnd_plugin`. `ext/mysqlnd_plugin` returns as well.
- Whatever PHP MySQL extension had been used by the application, it receives a connection to 127.0.0.1. The PHP MySQL extension itself returns to the PHP application. The circle is closed.

### 21.9.5.7.5. Getting started building a mysqlnd plugin

Copyright 1997-2012 the PHP Documentation Group. [2906]

It is important to remember that a `mysqlnd` plugin is itself a PHP extension.

The following code shows the basic structure of the `MINIT` function that will be used in the typical `mysqlnd` plugin:

```
/* my_php_mysqlnd_plugin.c */

static PHP_MINIT_FUNCTION(mysqlnd_plugin) {
    /* globals, ini entries, resources, classes */

    /* register mysqlnd plugin */
    mysqlnd_plugin_id = mysqlnd_plugin_register();

    conn_m = mysqlnd_get_conn_methods();
    memcpy(org_conn_m, conn_m,
```

```

sizeof(struct st_mysqlnd_conn_methods));

conn_m->query = MYSQLND_METHOD(mysqlnd_plugin_conn, query);
conn_m->connect = MYSQLND_METHOD(mysqlnd_plugin_conn, connect);
}

```

```

/* my_mysqlnd_plugin.c */

enum_func_status MYSQLND_METHOD(mysqlnd_plugin_conn, query)(/* ... */) {
    /* ... */
}
enum_func_status MYSQLND_METHOD(mysqlnd_plugin_conn, connect)(/* ... */) {
    /* ... */
}

```

### Task analysis: from C to userspace

```

class proxy extends mysqlnd_plugin_connection {
    public function connect($host, ...) { .. }
}
mysqlnd_plugin_set_conn_proxy(new proxy());

```

#### Process:

1. PHP: user registers plugin callback
2. PHP: user calls any PHP MySQL API to connect to MySQL
3. C: ext/\*mysql\* calls mysqlnd method
4. C: mysqlnd ends up in ext/mysqlnd\_plugin
5. C: ext/mysqlnd\_plugin
  - a. Calls userspace callback
  - b. Or original [mysqlnd](#) method, if userspace callback not set

You need to carry out the following:

1. Write a class "mysqlnd\_plugin\_connection" in C
2. Accept and register proxy object through "mysqlnd\_plugin\_set\_conn\_proxy()"
3. Call userspace proxy methods from C (optimization - zend\_interfaces.h)

Userspace object methods can either be called using [call\\_user\\_function\(\)](#) or you can operate at a level closer to the Zend Engine and use [zend\\_call\\_method\(\)](#).

#### Optimization: calling methods from C using [zend\\_call\\_method](#)

The following code snippet shows the prototype for the [zend\\_call\\_method](#) function, taken from [zend\\_interfaces.h](#).

```

ZEND_API zval* zend_call_method(
    zval **object_pp, zend_class_entry *obj_ce,
    zend_function **fn_proxy, char *function_name,
    int function_name_len, zval **retval_ptr_ptr,
    int param_count, zval* arg1, zval* arg2 TSRMLS_DC
);

```

Zend API supports only two arguments. You may need more, for example:

```
enum_func_status (*func_mysqlnd_conn_connect)(
    MYSQLND *conn, const char *host,
    const char * user, const char * passwd,
    unsigned int passwd_len, const char * db,
    unsigned int db_len, unsigned int port,
    const char * socket, unsigned int mysql_flags TSRMLS_DC
);
```

To get around this problem you will need to make a copy of `zend_call_method()` and add a facility for additional parameters. You can do this by creating a set of `MY_ZEND_CALL_METHOD_WRAPPER` macros.

### Calling PHP userspace

This code snippet shows the optimized method for calling a userspace function from C:

```
/* my_mysqlnd_plugin.c */
MYSQLND_METHOD(my_conn_class, connect)(
    MYSQLND *conn, const char *host /* ... */ TSRMLS_DC) {
    enum_func_status ret = FAIL;
    zval * global_user_conn_proxy = fetch_userspace_proxy();
    if (global_user_conn_proxy) {
        /* call userspace proxy */
        ret = MY_ZEND_CALL_METHOD_WRAPPER(global_user_conn_proxy, host, /*...*/);
    } else {
        /* or original mysqlnd method = do nothing, be transparent */
        ret = org_methods.connect(conn, host, user, passwd,
            passwd_len, db, db_len, port,
            socket, mysql_flags TSRMLS_CC);
    }
    return ret;
}
```

### Calling userspace: simple arguments

```
/* my_mysqlnd_plugin.c */
MYSQLND_METHOD(my_conn_class, connect)(
    /* ... */, const char *host, /* ... */) {
    /* ... */
    if (global_user_conn_proxy) {
        /* ... */
        zval* zv_host;
        MAKE_STD_ZVAL(zv_host);
        ZVAL_STRING(zv_host, host, 1);
        MY_ZEND_CALL_METHOD_WRAPPER(global_user_conn_proxy, zv_retval, zv_host /*, ...*/);
        zval_ptr_dtor(&zv_host);
        /* ... */
    }
    /* ... */
}
```

### Calling userspace: structs as arguments

```
/* my_mysqlnd_plugin.c */
MYSQLND_METHOD(my_conn_class, connect)(
    MYSQLND *conn, /* ... */) {
    /* ... */
    if (global_user_conn_proxy) {
        /* ... */
        zval* zv_conn;
        ZEND_REGISTER_RESOURCE(zv_conn, (void *)conn, le_mysqlnd_plugin_conn);
        MY_ZEND_CALL_METHOD_WRAPPER(global_user_conn_proxy, zv_retval, zv_conn, zv_host /*, ...*/);
        zval_ptr_dtor(&zv_conn);
        /* ... */
    }
}
```

```
/* ... */
}
```

The first argument of many [mysqlnd](#) methods is a C "object". For example, the first argument of the `connect()` method is a pointer to `MYSQLND`. The struct `MYSQLND` represents a [mysqlnd](#) connection object.

The [mysqlnd](#) connection object pointer can be compared to a standard I/O file handle. Like a standard I/O file handle a [mysqlnd](#) connection object shall be linked to the userspace using the PHP resource variable type.

*From C to userspace and back*

```
class proxy extends mysqlnd_plugin_connection {
    public function connect($conn, $host, ...) {
        /* "pre" hook */
        printf("Connecting to host = '%s'\n", $host);
        debug_print_backtrace();
        return parent::connect($conn);
    }

    public function query($conn, $query) {
        /* "post" hook */
        $ret = parent::query($conn, $query);
        printf("Query = '%s'\n", $query);
        return $ret;
    }
}
mysqlnd_plugin_set_conn_proxy(new proxy());
```

PHP users must be able to call the parent implementation of an overwritten method.

As a result of subclassing it is possible to refine only selected methods and you can choose to have "pre" or "post" hooks.

*Buildin class: `mysqlnd_plugin_connection::connect()`*

```
/* my_mysqlnd_plugin_classes.c */

PHP_METHOD("mysqlnd_plugin_connection", connect) {
    /* ... simplified! ... */
    zval* mysqlnd_rsrc;
    MYSQLND* conn;
    char* host; int host_len;
    if (zend_parse_parameters(ZEND_NUM_ARGS() TSRMLS_CC, "rs",
        &mysqlnd_rsrc, &host, &host_len) == FAILURE) {
        RETURN_NULL();
    }
    ZEND_FETCH_RESOURCE(conn, MYSQLND* conn, &mysqlnd_rsrc, -1,
        "MySQLnd Connection", le_mysqlnd_plugin_conn);
    if (PASS == org_methods.connect(conn, host, /* simplified! */ TSRMLS_CC))
        RETVAL_TRUE;
    else
        RETVAL_FALSE;
}
```

## 21.9.6. MySQLnd replication and load balancing plugin ([mysqlnd\\_ms](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

The [mysqlnd](#) replication and load balancing plugin ([mysqlnd\\_ms](#)) adds easy to use MySQL replication support to all PHP MySQL extensions that use [mysqlnd](#).

As of version PHP 5.3.3 the MySQL native driver for PHP ([mysqlnd](#)) features an internal plugin C API. C plugins, such as the replication and load balancing plugin, can extend the functionality of [mysqlnd](#).



The MySQL native driver for PHP is a C library that ships together with PHP as of PHP 5.3.0. It serves as a drop-in replacement for the MySQL Client Library (libmysqlclient). Using `mysqld` has several advantages: no extra downloads are required because it's bundled with PHP, it's under the PHP license, there is lower memory consumption in certain cases, and it contains new functionality such as asynchronous queries.

Mysqld plugins like `mysqld_ms` operate, for the most part, transparently from a user perspective. The replication and load balancing plugin supports all PHP applications, and all MySQL PHP extensions. It does not change existing APIs. Therefore, it can easily be used with existing PHP applications.

### 21.9.6.1. Key Features

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The key features of PECL/`mysqld_ms` are as follows.

- Transparent and therefore easy to use.
  - Supports all of the PHP MySQL extensions.
  - SSL support.
  - A consistent API.
  - Little to no application changes required, dependent on the required usage scenario.
  - Lazy connections: connections to master and slave servers are not opened before a SQL statement is executed.
  - Optional: automatic use of master after the first write in a web request, to lower the possible impact of replication lag.
- Can be used with any MySQL clustering solution.
  - MySQL Replication: Read-write splitting is done by the plugin. Primary focus of the plugin.
  - MySQL Cluster: Read-write splitting can be disabled. Configuration of multiple masters possible
  - Third-party solutions: the plugin is optimized for MySQL Replication but can be used with any other kind of MySQL clustering solution.
- Featured read-write split strategies
  - Automatic detection of SELECT.
  - Supports SQL hints to overrule automatism.
  - User-defined.
  - Can be disabled for, for example, when using synchronous clusters such as MySQL Cluster.
- Featured load balancing strategies
  - Round Robin: choose a different slave in round-robin fashion for every slave request.
  - Random: choose a random slave for every slave request.
  - Random once (sticky): choose a random slave once to run all slave requests for the duration of a web request.
  - User-defined. The application can register callbacks with `mysqld_ms`.

- PHP 5.4.0 or newer: transaction aware when using API calls only to control transactions.
- Weighted load balancing: servers can be assigned different priorities, for example, to direct more requests to a powerful machine than to another less powerful machine. Or, to prefer nearby machines to reduce latency.
- Global transaction ID
- Client-side emulation. Makes manual master server failover and slave promotion easier with asynchronous clusters, such as MySQL Replication.
- Support for built-in global transaction identifier feature of MySQL 5.6.5 or newer.
- Supports using transaction ids to identify up-to-date asynchronous slaves for reading when session consistency is required.
- Throttling: optionally, the plugin can wait for a slave to become "synchronous" before continuing.
- Service and consistency levels
  - Applications can request eventual, session and strong consistency service levels for connections. Appropriate cluster nodes will be searched automatically.
  - Eventual consistent MySQL Replication slave accesses can be replaced with fast local cache accesses transparently to reduce server load.
- Partitioning and sharding
  - Servers of a replication cluster can be organized into groups. SQL hints can be used to manually direct queries to a specific group. Grouping can be used to partition (shard) the data, or to cure the issue of hotspots with updates.
  - MySQL Replication filters are supported through the table filter.

### 21.9.6.2. Limitations

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The built-in read-write-split mechanism is very basic. Every query which starts with `SELECT` is considered a read request to be sent to a MySQL slave server. All other queries (such as `SHOW` statements) are considered as write requests that are sent to the MySQL master server. The build-in behavior can be overruled using [SQL hints](#), or a user-defined [callback function](#).

The read-write splitter is not aware of multi-statements. Multi-statements are considered as one statement. The decision of where to run the statement will be based on the beginning of the statement string. For example, if using `mysqli_multi_query` to execute the multi-statement `SELECT id FROM test ; INSERT INTO test(id) VALUES (1)`, the statement will be redirected to a slave server because it begins with `SELECT`. The `INSERT` statement, which is also part of the multi-statement, will not be redirected to a master server.



#### Note

Applications must be aware of the consequences of connection switches that are performed for load balancing purposes. Please check the documentation on [connection pooling and switching](#), [transaction handling](#), [failover load balancing](#) and [read-write splitting](#).

### 21.9.6.3. On the name

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The shortcut `mysqlnd_ms` stands for `mysqlnd master slave plugin`. The name was chosen for a quick-and-dirty proof-of-concept. In the beginning the developers did not expect to continue using the code base.

#### 21.9.6.4. Quickstart and Examples

Copyright 1997-2012 the PHP Documentation Group. [2906]

The mysqlnd replication load balancing plugin is easy to use. This quickstart will demo typical use-cases, and provide practical advice on getting started.

It is strongly recommended to read the reference sections in addition to the quickstart. The quickstart tries to avoid discussing theoretical concepts and limitations. Instead, it will link to the reference sections. It is safe to begin with the quickstart. However, before using the plugin in mission critical environments we urge you to read additionally the background information from the reference sections.

The focus is on using PECL `mysqlnd_ms` for work with an asynchronous MySQL cluster, namely MySQL replication. Generally speaking an asynchronous cluster is more difficult to use than a synchronous one. Thus, users of, for example, MySQL Cluster will find more information than needed.

##### 21.9.6.4.1. Setup

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin is implemented as a PHP extension. See also the [installation instructions](#) to install the `PECL/mysqlnd_ms` extension.

Compile or configure the PHP MySQL extension (API) (`mysqli`, `PDO_MYSQL`, `mysql`) that you plan to use with support for the `mysqlnd` library. `PECL/mysqlnd_ms` is a plugin for the `mysqlnd` library. To use the plugin with any of the PHP MySQL extensions, the extension has to use the `mysqlnd` library.

Then, load the extension into PHP and activate the plugin in the PHP configuration file using the PHP configuration directive named `mysqlnd_ms.enable`.

##### Example 21.215. Enabling the plugin (`php.ini`)

```
mysqlnd_ms.enable=1
mysqlnd_ms.config_file=/path/to/mysqlnd_ms_plugin.ini
```

The plugin uses its own configuration file. Use the PHP configuration directive `mysqlnd_ms.config_file` to set the full file path to the plugin-specific configuration file. This file must be readable by PHP (e.g., the web server user). Please note, the configuration directive `mysqlnd_ms.config_file` supersedes `mysqlnd_ms.ini_file` since 1.4.0. It is a common pitfall to use the old, no longer available configuration directive.

Create a plugin-specific configuration file. Save the file to the path set by the PHP configuration directive `mysqlnd_ms.config_file`.

The plugin's [configuration file](#) is JSON based. It is divided into one or more sections. Each section has a name, for example, `myapp`. Every section makes its own set of configuration settings.

A section must, at a minimum, list the MySQL replication master server, and set a list of slaves. The plugin supports using only one master server per section. Multi-master MySQL replication setups are not yet fully supported. Use the configuration setting `master` to set the hostname, and the port or socket of the MySQL master server. MySQL slave servers are configured using the `slave` keyword.

##### Example 21.216. Minimal plugin-specific configuration file (`mysqlnd_ms_plugin.ini`)

```
{
  "myapp": {
```

```

"master": {
  "master_0": {
    "host": "localhost"
  }
},
"slave": [
]
}

```

Configuring a MySQL slave server list is required, although it may contain an empty list. It is recommended to always configure at least one slave server.

Server lists can use [anonymous or non-anonymous syntax \[3280\]](#). Non-anonymous lists include alias names for the servers, such as `master_0` for the master in the above example. The quickstart uses the more verbose non-anonymous syntax.

#### Example 21.217. Recommended minimal plugin-specific config (`mysqlnd_ms_plugin.ini`)

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\\tmp\\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27",
        "port": "3306"
      }
    }
  }
}

```

If there are at least two servers in total, the plugin can start to load balance and switch connections. Switching connections is not always transparent and can cause issues in certain cases. The reference sections about [connection pooling and switching](#), [transaction handling](#), [fail over load balancing](#) and [read-write splitting](#) all provide more details. And potential pitfalls are described later in this guide.

It is the responsibility of the application to handle potential issues caused by connection switches, by configuring a master with at least one slave server, which allows switching to work therefore related problems can be found.

The MySQL master and MySQL slave servers, which you configure, do not need to be part of MySQL replication setup. For testing purpose you can use single MySQL server and make it known to the plugin as a master and slave server as shown below. This could help you to detect many potential issues with connection switches. However, such a setup will not be prone to the issues caused by replication lag.

#### Example 21.218. Using one server as a master and as a slave (testing only!)

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\\tmp\\mysql.sock"
      }
    },
  },
}

```

```

    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    }
  }
}

```

The plugin attempts to notify you of invalid configurations. Since 1.5.0 it will throw a warning during PHP startup if the configuration file cannot be read, is empty or parsing the JSON failed. Depending on your PHP settings those errors may appear in some log files only. Further validation is done when a connection is to be established and the configuration file is searched for valid sections. Setting [mysqlnd\\_ms.force\\_config\\_usage](#) may help debugging a faulty setup. Please, see also [configuration file debugging notes \[3282\]](#).

#### 21.9.6.4.2. Running statements

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin can be used with any PHP MySQL extension (`mysqli`, `mysql`, and `PDO_MYSQL`) that is compiled to use the `mysqlnd` library. PECL/`mysqlnd_ms` plugs into the `mysqlnd` library. It does not change the API or behavior of those extensions.

Whenever a connection to MySQL is being opened, the plugin compares the host parameter value of the connect call, with the section names from the plugin specific configuration file. If, for example, the plugin specific configuration file has a section `myapp` then the section should be referenced by opening a MySQL connection to the host `myapp`

#### Example 21.219. Plugin specific configuration file (`mysqlnd_ms_plugin.ini`)

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27",
        "port": "3306"
      }
    }
  }
}

```

#### Example 21.220. Opening a load balanced connection

```

<?php
/* Load balanced following "myapp" section rules from the plugins config file */
$mysqli = new mysqli("myapp", "username", "password", "database");
$pdo = new PDO('mysql:host=myapp;dbname=database', 'username', 'password');
$mysql = mysql_connect("myapp", "username", "password");
?>

```

The connection examples above will be load balanced. The plugin will send read-only statements to the MySQL slave server with the IP `192.168.2.27` and will listen on port `3306` for the MySQL client connection. All other statements will be directed to the MySQL master server running on the host

[localhost](#). If on Unix like operating systems, the master on [localhost](#) will be accepting MySQL client connections on the Unix domain socket `/tmp/mysql.sock`, while TCP/IP is the default port on Windows. The plugin will use the user name [username](#) and the password [password](#) to connect to any of the MySQL servers listed in the section [myapp](#) of the plugins configuration file. Upon connect, the plugin will select [database](#) as the current schemata.

The username, password and schema name are taken from the connect API calls and used for all servers. In other words: you must use the same username and password for every MySQL server listed in a plugin configuration file section. This is not a general limitation. As of PECL/mysqlnd\_ms 1.1.0, it is possible to set the [username \[3284\]](#) and [password \[3284\]](#) for any server in the plugins configuration file, to be used instead of the credentials passed to the API call.

The plugin does not change the API for running statements. [Read-write splitting](#) works out of the box. The following example assumes that there is no significant replication lag between the master and the slave.

### Example 21.221. Executing statements

```
<?php
/* Load balanced following "myapp" section rules from the plugins config file */
$mysqli = new mysqli("myapp", "username", "password", "database");
if (mysqli_connect_errno())
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Statements will be run on the master */
if (!$mysqli->query("DROP TABLE IF EXISTS test")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
if (!$mysqli->query("CREATE TABLE test(id INT)")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
if (!$mysqli->query("INSERT INTO test(id) VALUES (1)")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}

/* read-only: statement will be run on a slave */
if (!$res = $mysqli->query("SELECT id FROM test")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
} else {
    $row = $res->fetch_assoc();
    $res->close();
    printf("Slave returns id = '%s'\n", $row['id']);
}
$mysqli->close();
?>
```

The above example will output something similar to:

```
Slave returns id = '1'
```

#### 21.9.6.4.3. Connection state

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin changes the semantics of a PHP MySQL connection handle. A new connection handle represents a connection pool, instead of a single MySQL client-server network connection. The connection pool consists of a master connection, and optionally any number of slave connections.

Every connection from the connection pool has its own state. For example, SQL user variables, temporary tables and transactions are part of the state. For a complete list of items that belong to the

state of a connection, see the [connection pooling and switching](#) concepts documentation. If the plugin decides to switch connections for load balancing, the application could be given a connection which has a different state. Applications must be made aware of this.

#### Example 21.222. Plugin config with one slave and one master

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27",
        "port": "3306"
      }
    }
  }
}
```

#### Example 21.223. Pitfall: connection state and SQL user variables

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Connection 1, connection bound SQL user variable, no SELECT thus run on master */
if (!$mysqli->query("SET @myrole='master'")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}

/* Connection 2, run on slave because SELECT */
if (!($res = $mysqli->query("SELECT @myrole AS _role"))) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
} else {
    $row = $res->fetch_assoc();
    $res->close();
    printf("@myrole = '%s'\n", $row['_role']);
}
$mysqli->close();
?>
```

The above example will output:

```
@myrole = ''
```

The example opens a load balanced connection and executes two statements. The first statement `SET @myrole='master'` does not begin with the string `SELECT`. Therefore the plugin does not recognize it as a read-only query which shall be run on a slave. The plugin runs the statement on the connection to the master. The statement sets a SQL user variable which is bound to the master connection. The state of the master connection has been changed.

The next statement is `SELECT @myrole AS _role`. The plugin does recognize it as a read-only query and sends it to the slave. The statement is run on a connection to the slave. This second connection does not have any SQL user variables bound to it. It has a different state than the first

connection to the master. The requested SQL user variable is not set. The example script prints `@myrole = ''`.

It is the responsibility of the application developer to take care of the connection state. The plugin does not monitor all connection state changing activities. Monitoring all possible cases would be a very CPU intensive task, if it could be done at all.

The pitfalls can easily be worked around using SQL hints.

#### 21.9.6.4.4. SQL Hints

Copyright 1997-2012 the PHP Documentation Group. [2906]

SQL hints can force a query to choose a specific server from the connection pool. It gives the plugin a hint to use a designated server, which can solve issues caused by connection switches and connection state.

SQL hints are standard compliant SQL comments. Because SQL comments are supposed to be ignored by SQL processing systems, they do not interfere with other programs such as the MySQL Server, the MySQL Proxy, or a firewall.

Three SQL hints are supported by the plugin: The `MYSQLND_MS_MASTER_SWITCH` hint makes the plugin run a statement on the master, `MYSQLND_MS_SLAVE_SWITCH` enforces the use of the slave, and `MYSQLND_MS_LAST_USED_SWITCH` will run a statement on the same server that was used for the previous statement.

The plugin scans the beginning of a statement for the existence of an SQL hint. SQL hints are only recognized if they appear at the beginning of the statement.

#### Example 21.224. Plugin config with one slave and one master

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27",
        "port": "3306"
      }
    }
  }
}
```

#### Example 21.225. SQL hints to prevent connection switches

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (mysqli_connect_errno())
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Connection 1, connection bound SQL user variable, no SELECT thus run on master */
if (!$mysqli->query("SET @myrole='master'")) {
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}

/* Connection 1, run on master because of SQL hint */
if (!$res = $mysqli->query(sprintf("/*%s*/SELECT @myrole AS _role", MYSQLND_MS_LAST_USED_SWITCH))) {
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
```



```

} else {
    $row = $res->fetch_assoc();
    $res->close();
    printf("@myrole = '%s'\n", $row['_role']);
}
$mysqli->close();
?>

```

The above example will output:

```
@myrole = 'master'
```

In the above example, using `MYSQLND_MS_LAST_USED_SWITCH` prevents session switching from the master to a slave when running the `SELECT` statement.

SQL hints can also be used to run `SELECT` statements on the MySQL master server. This may be desired if the MySQL slave servers are typically behind the master, but you need current data from the cluster.

In version 1.2.0 the concept of a service level has been introduced to address cases when current data is required. Using a service level requires less attention and removes the need of using SQL hints for this use case. Please, find more information below in the service level and consistency section.

#### Example 21.226. Fighting replication lag

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Force use of master, master has always fresh and current data */
if (!$mysqli->query(sprintf("/*%s*/SELECT critical_data FROM important_table", MYSQLND_MS_MASTER_SWITCH)))
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
?>

```

A use case may include the creation of tables on a slave. If an SQL hint is not given, then the plugin will send `CREATE` and `INSERT` statements to the master. Use the SQL hint `MYSQLND_MS_SLAVE_SWITCH` if you want to run any such statement on a slave, for example, to build temporary reporting tables.

#### Example 21.227. Table creation on a slave

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Force use of slave */
if (!$mysqli->query(sprintf("/*%s*/CREATE TABLE slave_reporting(id INT)", MYSQLND_MS_SLAVE_SWITCH))) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
/* Continue using this particular slave connection */
if (!$mysqli->query(sprintf("/*%s*/INSERT INTO slave_reporting(id) VALUES (1), (2), (3)", MYSQLND_MS_SLAVE_SWITCH)))
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
/* Don't use MYSQLND_MS_SLAVE_SWITCH which would allow switching to another slave! */
if ($res = $mysqli->query(sprintf("/*%s*/SELECT COUNT(*) AS _num FROM slave_reporting", MYSQLND_MS_SLAVE_SWITCH)))

```

```

$row = $res->fetch_assoc();
$res->close();
printf("There are %d rows in the table 'slave_reporting'", $row['_num']);
} else {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}
$mysqli->close();
?>

```

The SQL hint `MYSQLND_MS_LAST_USED` forbids switching a connection, and forces use of the previously used connection.

#### 21.9.6.4.5. Transactions

Copyright 1997-2012 the PHP Documentation Group. [2906]

The current version of the plugin is not transaction safe by default, because it is not aware of running transactions in all cases. SQL transactions are units of work to be run on a single server. The plugin does not always know when the unit of work starts and when it ends. Therefore, the plugin may decide to switch connections in the middle of a transaction.

No kind of MySQL load balancer can detect transaction boundaries without any kind of hint from the application.

You can either use SQL hints to work around this limitation. Alternatively, you can activate transaction API call monitoring. In the latter case you must use API calls only to control transactions, see below.

#### Example 21.228. Plugin config with one slave and one master

```

[myapp]
{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\tmp\mysql.sock"
            }
        },
        "slave": {
            "slave_0": {
                "host": "192.168.2.27",
                "port": "3306"
            }
        }
    }
}

```

#### Example 21.229. Using SQL hints for transactions

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Not a SELECT, will use master */
if (!$mysqli->query("START TRANSACTION")) {
    /* Please use better error handling in your code */
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* Prevent connection switch! */
if (!$mysqli->query(sprintf("/*%s*/INSERT INTO test(id) VALUES (1)", MYSQLND_MS_LAST_USED_SWITCH))) {

```

```

/* Please do proper ROLLBACK in your code, don't just die */
die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if ($res = $mysqli->query(sprintf("/*%s*/SELECT COUNT(*) AS _num FROM test", MYSQLND_MS_LAST_USED_SWITCH))) {
    $row = $res->fetch_assoc();
    $res->close();
    if ($row['_num'] > 1000) {
        if (!$mysqli->query(sprintf("/*%s*/INSERT INTO events(task) VALUES ('cleanup'", MYSQLND_MS_LAST_USED_SWITCH))) {
            die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
        }
    }
} else {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if (!$mysqli->query(sprintf("/*%s*/UPDATE log SET last_update = NOW()", MYSQLND_MS_LAST_USED_SWITCH))) {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if (!$mysqli->query(sprintf("/*%s*/COMMIT", MYSQLND_MS_LAST_USED_SWITCH))) {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

$mysqli->close();
?>

```

Starting with PHP 5.4.0, the `mysqlnd` library allows the plugin to monitor the status of the `autocommit` mode, if the mode is set by API calls instead of using SQL statements such as `SET AUTOCOMMIT=0`. This makes it possible for the plugin to become transaction aware. In this case, you do not need to use SQL hints.

If using PHP 5.4.0 or newer, API calls that enable `autocommit` mode, and when setting the plugin configuration option `trx_stickiness=master`, the plugin can automatically disable load balancing and connection switches for SQL transactions. In this configuration, the plugin stops load balancing if `autocommit` is disabled and directs all statements to the master. This prevents connection switches in the middle of a transaction. Once `autocommit` is re-enabled, the plugin starts to load balance statements again.

API based transaction boundary detection has been improved with PHP 5.5.0 and PECL/`mysqlnd_ms` 1.5.0 to cover not only calls to `mysqli_autocommit` but also `mysqli_begin`, `mysqli_commit` and `mysqli_rollback`.

#### Example 21.230. Transaction aware load balancing: `trx_stickiness` setting

```

{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\tmp\mysql.sock"
            }
        },
        "slave": {
            "slave_0": {
                "host": "127.0.0.1",
                "port": "3306"
            }
        },
        "trx_stickiness": "master"
    }
}

```

#### Example 21.231. Transaction aware

```
<?php
```

```

$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Disable autocommit, plugin will run all statements on the master */
$mysqli->autocommit(FALSE);

if (!$mysqli->query("INSERT INTO test(id) VALUES (1)") {
    /* Please do proper ROLLBACK in your code, don't just die */
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if ($res = $mysqli->query("SELECT COUNT(*) AS _num FROM test")) {
    $row = $res->fetch_assoc();
    $res->close();
    if ($row['_num'] > 1000) {
        if (!$mysqli->query("INSERT INTO events(task) VALUES ('cleanup')")) {
            die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
        }
    }
} else {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if (!$mysqli->query("UPDATE log SET last_update = NOW()") {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}
if (!$mysqli->commit()) {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* Plugin assumes that the transaction has ended and starts load balancing again */
$mysqli->autocommit(TRUE);
$mysqli->close();
?>

```



#### Version requirement

The plugin configuration option `trx_stickiness=master` requires PHP 5.4.0 or newer.

Please note the restrictions outlined in the [transaction handling](#) concepts section.

#### 21.9.6.4.6. Service level and consistency

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Version requirement

Service levels have been introduced in PECL `mysqlnd_ms` version 1.2.0-alpha. `mysqlnd_ms_set_qos` is available with PHP 5.4.0 or newer.

Different types of MySQL cluster solutions offer different service and data consistency levels to their users. An asynchronous MySQL replication cluster offers eventual consistency by default. A read executed on an asynchronous slave may return current, stale or no data at all, depending on whether the slave has replayed all changesets from the master or not.

Applications using an MySQL replication cluster need to be designed to work correctly with eventual consistent data. In some cases, however, stale data is not acceptable. In those cases only certain slaves or even only master accesses are allowed to achieve the required quality of service from the cluster.

As of PECL `mysqlnd_ms` 1.2.0 the plugin is capable of selecting MySQL replication nodes automatically that deliver session consistency or strong consistency. Session consistency means that one client can read its writes. Other clients may or may not see the clients' write. Strong consistency means that all clients will see all writes from the client.

**Example 21.232. Session consistency: read your writes**

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    }
  }
}
```

**Example 21.233. Requesting session consistency**

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* read-write splitting: master used */
if (!$mysqli->query("INSERT INTO orders(order_id, item) VALUES (1, 'christmas tree, 1.8m')")) {
    /* Please use better error handling in your code */
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* Request session consistency: read your writes */
if (!mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_SESSION))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Plugin picks a node which has the changes, here: master */
if (!$res = $mysqli->query("SELECT item FROM orders WHERE order_id = 1"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

var_dump($res->fetch_assoc());

/* Back to eventual consistency: stale data allowed */
if (!mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Plugin picks any slave, stale data is allowed */
if (!$res = $mysqli->query("SELECT item, price FROM specials"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
?>
```

Service levels can be set in the plugins configuration file and at runtime using [mysqlnd\\_ms\\_set\\_qos](#). In the example the function is used to enforce session consistency (read your writes) for all future statements until further notice. The `SELECT` statement on the `orders` table is run on the master to ensure the previous write can be seen by the client. Read-write splitting logic has been adapted to fulfill the service level.

After the application has read its changes from the `orders` table it returns to the default service level, which is eventual consistency. Eventual consistency puts no restrictions on choosing a node for statement execution. Thus, the `SELECT` statement on the `specials` table is executed on a slave.

The new functionality supersedes the use of SQL hints and the [master\\_on\\_write](#) configuration option. In many cases [mysqlnd\\_ms\\_set\\_qos](#) is easier to use, more powerful improves portability.

### Example 21.234. Maximum age/slave lag

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    },
    "failover" : "master"
  }
}
```

### Example 21.235. Limiting slave lag

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Read from slaves lagging no more than four seconds */
$ret = mysqlnd_ms_set_qos($mysqli,
    MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL,
    MYSQLND_MS_QOS_OPTION_AGE, 4);

if (!$ret)
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Plugin picks any slave, which may or may not have the changes */
if (!$res = $mysqli->query("SELECT item, price FROM daytrade"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Back to default: use of all slaves and masters permitted */
if (!mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
?>
```

The eventual consistency service level can be used with an optional parameter to set a maximum slave lag for choosing slaves. If set, the plugin checks `SHOW SLAVE STATUS` for all configured slaves. In case of the example, only slaves for which `Slave_IO_Running=Yes`, `Slave_SQL_Running=Yes` and `Seconds_Behind_Master <= 4` is true are considered for executing the statement `SELECT item, price FROM daytrade`.

Checking `SHOW SLAVE STATUS` is done transparently from an applications perspective. Errors, if any, are reported as warnings. No error will be set on the connection handle. Even if all `SHOW SLAVE STATUS` SQL statements executed by the plugin fail, the execution of the users statement is not stopped, given that master fail over is enabled. Thus, no application changes are required.



#### Expensive and slow operation

Checking `SHOW SLAVE STATUS` for all slaves adds overhead to the application. It is an expensive and slow background operation. Try to minimize the use of it. Unfortunately, a MySQL replication cluster does not give clients the possibility to

request a list of candidates from a central instance. Thus, a more efficient way of checking the slaves lag is not available.

Please, note the limitations and properties of `SHOW SLAVE STATUS` as explained in the MySQL reference manual.

To prevent `mysqLnd_ms` from emitting a warning if no slaves can be found that lag no more than the defined number of seconds behind the master, it is necessary to enable master fail over in the plugins configuration file. If no slaves can be found and fail over is turned on, the plugin picks a master for executing the statement.

If no slave can be found and fail over is turned off, the plugin emits a warning, it does not execute the statement and it sets an error on the connection.

#### Example 21.236. Fail over not set

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    }
  }
}
```

#### Example 21.237. No slave within time limit

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Read from slaves lagging no more than four seconds */
$ret = mysqlnd_ms_set_qos($mysqli,
  MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL,
  MYSQLND_MS_QOS_OPTION_AGE, 4);

if (!$ret)
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Plugin picks any slave, which may or may not have the changes */
if (!$res = $mysqli->query("SELECT item, price FROM daytrade"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Back to default: use of all slaves and masters permitted */
if (!mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
?>
```

The above example will output:

```
PHP Warning: mysqli::query(): (mysqlnd_ms) Couldn't find the appropriate slave connection. 0 slaves to
```

```
PHP Warning: mysqli::query(): (mysqlnd_ms) No connection selected by the last filter in %s on line %d
[2000] (mysqlnd_ms) No connection selected by the last filter
```

### 21.9.6.4.7. Global transaction IDs

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Version requirement

A client-side global transaction ID injection has been introduced in `mysqlnd_ms` version 1.2.0-alpha. The feature is not required for synchronous clusters, such as MySQL Cluster. Use it with asynchronous clusters such as classical MySQL replication.

As of MySQL 5.6.5-m8 the MySQL server features built-in global transaction identifiers. The MySQL built-in global transaction ID feature is supported by PECL/`mysqlnd_ms` 1.3.0-alpha or later.

PECL/`mysqlnd_ms` can either use its own global transaction ID emulation or the global transaction ID feature built-in to MySQL 5.6.5-m8 or later. From a developer perspective the client-side and server-side approach offer the same features with regards to service levels provided by PECL/`mysqlnd_ms`. Their differences are discussed in the [concepts section](#).

The quickstart first demonstrates the use of the client-side global transaction ID emulation built-in to PECL/`mysqlnd_ms` before it shows how to use the server-side counterpart. The order ensures that the underlying idea is discussed first.

#### *Idea and client-side emulation*

In its most basic form a global transaction ID (GTID) is a counter in a table on the master. The counter is incremented whenever a transaction is committed on the master. Slaves replicate the table. The counter serves two purposes. In case of a master failure, it helps the database administrator to identify the most recent slave for promoting it to the new master. The most recent slave is the one with the highest counter value. Applications can use the global transaction ID to search for slaves which have replicated a certain write (identified by a global transaction ID) already.

PECL/`mysqlnd_ms` can inject SQL for every committed transaction to increment a GTID counter. The so created GTID is accessible by the application to identify an applications write operation. This enables the plugin to deliver session consistency (read your writes) service level by not only querying masters but also slaves which have replicated the change already. Read load is taken away from the master.

Client-side global transaction ID emulation has some limitations. Please, read the [concepts section](#) carefully to fully understand the principles and ideas behind it, before using in production environments. The background knowledge is not required to continue with the quickstart.

First, create a counter table on your master server and insert a record into it. The plugin does not assist creating the table. Database administrators must make sure it exists. Depending on the error reporting mode, the plugin will silently ignore the lack of the table or bail out.

#### **Example 21.238. Create counter table on master**

```
CREATE TABLE `trx` (
  `trx_id` int(11) DEFAULT NULL,
  `last_update` timestamp NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
) ENGINE=InnoDB DEFAULT CHARSET=latin1
INSERT INTO `trx`(`trx_id`) VALUES (1);
```



In the plugins configuration file set the SQL to update the global transaction ID table using `on_commit` from the `global_transaction_id_injection` section. Make sure the table name used for the `UPDATE` statement is fully qualified. In the example, `test.trx` is used to refer to table `trx` in the schema (database) `test`. Use the table that was created in the previous step. It is important to set the fully qualified table name because the connection on which the injection is done may use a different default database. Make sure the user that opens the connection is allowed to execute the `UPDATE`.

Enable reporting of errors that may occur when `mysqld_ms` does global transaction ID injection.

### Example 21.239. Plugin config: SQL for client-side GTID injection

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\/tmp\/mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    },
    "global_transaction_id_injection": {
      "on_commit": "UPDATE test.trx SET trx_id = trx_id + 1",
      "report_error": true
    }
  }
}
```

### Example 21.240. Transparent global transaction ID injection

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("DROP TABLE IF EXISTS test"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("CREATE TABLE test(id INT)"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("INSERT INTO test(id) VALUES (1)"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* auto commit mode, read on slave, no increment */
if (!($res = $mysqli->query("SELECT id FROM test")))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

var_dump($res->fetch_assoc());
?>
```

The above example will output:

```
array(1) {
  ["id"]=>
```

```
string(1) "1"
}
```

The example runs three statements in auto commit mode on the master, causing three transactions on the master. For every such statement, the plugin will inject the configured `UPDATE` transparently before executing the users SQL statement. When the script ends the global transaction ID counter on the master has been incremented by three.

The fourth SQL statement executed in the example, a `SELECT`, does not trigger an increment. Only transactions (writes) executed on a master shall increment the GTID counter.



#### SQL for global transaction ID: efficient solution wanted!

The SQL used for the client-side global transaction ID emulation is inefficient. It is optimized for clarity not for performance. Do not use it for production environments. Please, help finding an efficient solution for inclusion in the manual. We appreciate your input.

#### Example 21.241. Plugin config: SQL for fetching GTID

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    },
    "global_transaction_id_injection": {
      "on_commit": "UPDATE test.trx SET trx_id = trx_id + 1",
      "fetch_last_gtid" : "SELECT MAX(trx_id) FROM test.trx",
      "report_error": true
    }
  }
}
```

#### Example 21.242. Obtaining GTID after injection

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("DROP TABLE IF EXISTS test"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

printf("GTID after transaction %s\n", mysqlnd_ms_get_last_gtid($mysqli));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("CREATE TABLE test(id INT)"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

printf("GTID after transaction %s\n", mysqlnd_ms_get_last_gtid($mysqli));
?>
```

The above example will output:

```
GTID after transaction 7
GTID after transaction 8
```

Applications can ask PECL `mysqlnd_ms` for a global transaction ID which belongs to the last write operation performed by the application. The function `mysqlnd_ms_get_last_gtid` returns the GTID obtained when executing the SQL statement from the `fetch_last_gtid` entry of the `global_transaction_id_injection` section from the plugins configuration file. The function may be called after the GTID has been incremented.

Applications are advised not to run the SQL statement themselves as this bares the risk of accidentally causing an implicit GTID increment. Also, if the function is used, it is easy to migrate an application from one SQL statement for fetching a transaction ID to another, for example, if any MySQL server ever features built-in global transaction ID support.

The quickstart shows a SQL statement which will return a GTID equal or greater to that created for the previous statement. It is exactly the GTID created for the previous statement if no other clients have incremented the GTID in the time span between the statement execution and the `SELECT` to fetch the GTID. Otherwise, it is greater.

#### Example 21.243. Plugin config: Checking for a certain GTID

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": "3306"
      }
    },
    "global_transaction_id_injection":{
      "on_commit":"UPDATE test.trx SET trx_id = trx_id + 1",
      "fetch_last_gtid" : "SELECT MAX(trx_id) FROM test.trx",
      "check_for_gtid" : "SELECT trx_id FROM test.trx WHERE trx_id >= #GTID",
      "report_error":true
    }
  }
}
```

#### Example 21.244. Session consistency service level and GTID combined

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1)"))
  die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
```

```

/* GTID as an identifier for the last write */
$gtid = mysqlnd_ms_get_last_gtid($mysqli);

/* Session consistency (read your writes): try to read from slaves not only master */
if (false == mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_SESSION, MYSQLND_MS_QOS_OPTION_GTID, $gtid)) {
    die(sprintf("[006] [%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* Either run on master or a slave which has replicated the INSERT */
if (!($res = $mysqli->query("SELECT id FROM test"))) {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

var_dump($res->fetch_assoc());
?>

```

A GTID returned from `mysqlnd_ms_get_last_gtid` can be used as an option for the session consistency service level. Session consistency delivers read your writes. Session consistency can be requested by calling `mysqlnd_ms_set_qos`. In the example, the plugin will execute the `SELECT` statement either on the master or on a slave which has replicated the previous `INSERT` already.

PECL `mysqlnd_ms` will transparently check every configured slave if it has replicated the `INSERT` by checking the slaves GTID table. The check is done running the SQL set with the `check_for_gtid` option from the `global_transaction_id_injection` section of the plugins configuration file. Please note, that this is a slow and expensive procedure. Applications should try to use it sparsely and only if read load on the master becomes to high otherwise.

#### *Use of the server-side global transaction ID feature*

Starting with MySQL 5.6.5-m8 the MySQL Replication system features server-side global transaction IDs. Transaction identifiers are automatically generated and maintained by the server. Users do not need to take care of maintaining them. There is no need to setup any tables in advance, or for setting `on_commit`. A client-side emulation is no longer needed.

Clients can continue to use global transaction identifier to achieve session consistency when reading from MySQL Replication slaves. The algorithm works as described above. Different SQL statements must be configured for `fetch_last_gtid` and `check_for_gtid`. The statements are given below. Please note, MySQL 5.6.5-m8 is a development version. Details of the server implementation may change in the future and require adoption of the SQL statements shown.

Using the following configuration any of the above described functionality can be used together with the server-side global transaction ID feature. `mysqlnd_ms_get_last_gtid` and `mysqlnd_ms_set_qos` continue to work as described above. The only difference is that the server does not use a simple sequence number but a string containing of a server identifier and a sequence number. Thus, users cannot easily derive an order from GTIDs returned by `mysqlnd_ms_get_last_gtid`.

#### **Example 21.245. Plugin config: using MySQL 5.6.5-m8 built-in GTID feature**

```

{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\\tmp\\mysql.sock"
            }
        },
        "slave": {
            "slave_0": {
                "host": "127.0.0.1",
                "port": "3306"
            }
        },
        "global_transaction_id_injection":{

```

```

        "fetch_last_gtid" : "SELECT @@GLOBAL.GTID_DONE AS trx_id FROM DUAL",
        "check_for_gtid"  : "SELECT GTID_SUBSET('#GTID', @@GLOBAL.GTID_DONE) AS trx_id FROM DUAL",
        "report_error":true
    }
}

```

### 21.9.6.4.8. Cache integration

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Version requirement, dependencies and status

Please, find more about version requirements, extension load order dependencies and the current status in the [concepts section](#)!

Databases clusters can deliver different levels of consistency. As of PECL/mysqlnd\_ms 1.2.0 it is possible to advice the plugin to consider only cluster nodes that can deliver the consistency level requested. For example, if using asynchronous MySQL Replication with its cluster-wide eventual consistency, it is possible to request session consistency (read your writes) at any time using `mysqlnd_ms_set_qos`. Please, see also the [service level and consistency](#) introduction.

#### Example 21.246. Recap: quality of service to request read your writes

```

/* Request session consistency: read your writes */
if (!mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_SESSION))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

```

Assuming PECL/mysqlnd has been explicitly told to deliver no consistency level higher than eventual consistency, it is possible to replace a database node read access with a client-side cache using time-to-live (TTL) as its invalidation strategy. Both the database node and the cache may or may not serve current data as this is what eventual consistency defines.

Replacing a database node read access with a local cache access can improve overall performance and lower the database load. If the cache entry is every reused by other clients than the one creating the cache entry, a database access is saved and thus database load is lowered. Furthermore, system performance can become better if computation and delivery of a database query is slower than a local cache access.

#### Example 21.247. Plugin config: no special entries for caching

```

{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\/tmp\/mysql.sock"
            }
        },
        "slave": {
            "slave_0": {
                "host": "127.0.0.1",
                "port": "3306"
            }
        }
    }
}

```

#### Example 21.248. Caching a slave request

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1)"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Explicitly allow eventual consistency and caching (TTL <= 60 seconds) */
if (false == mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL, MYSQLND_MS_QOS_OPTION_CACHE,
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* To make this example work, we must wait for a slave to catch up. Brute force style. */
$attempts = 0;
do {
    /* check if slave has the table */
    if ($res = $mysqli->query("SELECT id FROM test")) {
        break;
    } else if ($mysqli->errno) {
        die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
    }
    /* wait for slave to catch up */
    usleep(200000);
} while ($attempts++ < 10);

/* Query has been run on a slave, result is in the cache */
assert($res);
var_dump($res->fetch_assoc());

/* Served from cache */
$res = $mysqli->query("SELECT id FROM test");
?>
    
```

The example shows how to use the cache feature. First, you have to set the quality of service to eventual consistency and explicitly allow for caching. This is done by calling `mysqlnd_ms_set_qos`. Then, the result set of every read-only statement is cached for upto that many seconds as allowed with `mysqlnd_ms_set_qos`.

The actual TTL is lower or equal to the value set with `mysqlnd_ms_set_qos`. The value passed to the function sets the maximum age (seconds) of the data delivered. To calculate the actual TTL value the replication lag on a slave is checked and subtracted from the given value. If, for example, the maximum age is set to 60 seconds and the slave reports a lag of 10 seconds the resulting TTL is 50 seconds. The TTL is calculated individually for every cached query.

#### Example 21.249. Read your writes and caching combined

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

if (!$mysqli->query("DROP TABLE IF EXISTS test") ||
    !$mysqli->query("CREATE TABLE test(id INT)") ||
    !$mysqli->query("INSERT INTO test(id) VALUES (1)"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Explicitly allow eventual consistency and caching (TTL <= 60 seconds) */
if (false == mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL, MYSQLND_MS_QOS_OPTION_CACHE,
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* To make this example work, we must wait for a slave to catch up. Brute force style. */
    
```

```

$attempts = 0;
do {
    /* check if slave has the table */
    if ($res = $mysqli->query("SELECT id FROM test")) {
        break;
    } else if ($mysqli->errno) {
        die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
    }
    /* wait for slave to catch up */
    usleep(200000);
} while ($attempts++ < 10);

assert($res);

/* Query has been run on a slave, result is in the cache */
var_dump($res->fetch_assoc());

/* Served from cache */
if (!($res = $mysqli->query("SELECT id FROM test")))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
var_dump($res->fetch_assoc());

/* Update on master */
if (!$mysqli->query("UPDATE test SET id = 2"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Read your writes */
if (false == mysqlnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_SESSION)) {
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
}

/* Fetch latest data */
if (!($res = $mysqli->query("SELECT id FROM test")))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
var_dump($res->fetch_assoc());
?>
    
```

The quality of service can be changed at any time to avoid further cache usage. If needed, you can switch to read your writes (session consistency). In that case, the cache will not be used and fresh data is read.

#### 21.9.6.4.9. Failover

Copyright 1997-2012 the PHP Documentation Group. [2906]

By default, the plugin does not attempt to fail over if connecting to a host fails. This prevents pitfalls related to [connection state](#). It is recommended to manually handle connection errors in a way similar to a failed transaction. You should catch the error, rebuild the connection state and rerun your query as shown below.

If connection state is no issue to you, you can alternatively enable automatic and silent failover. Depending on the configuration, the automatic and silent failover will either attempt to fail over to the master before issuing an error or, try to connect to other slaves, given the query allows for it, before attempting to connect to a master. Because [automatic failover](#) is not fool-proof, it is not discussed in the quickstart. Instead, details are given in the [concepts](#) section below.

#### Example 21.250. Manual failover, automatic optional

```

{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\tmp\mysql.sock"
            }
        }
    },
}
    
```

```

    "slave": {
      "slave_0": {
        "host": "simulate_slave_failure",
        "port": "0"
      },
      "slave_1": {
        "host": "127.0.0.1",
        "port": 3311
      }
    },
    "filters": { "roundrobin": [] }
  }
}

```

### Example 21.251. Manual failover

```

<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

$sql = "SELECT 1 FROM DUAL";

/* error handling as it should be done regardless of the plugin */
if (!$res = $link->query($sql)) {
    /* plugin specific: check for connection error */
    switch ($link->errno) {
        case 2002:
        case 2003:
        case 2005:
            printf("Connection error - trying next slave!\n");
            /* load balancer will pick next slave */
            $res = $link->query($sql);
            break;
        default:
            /* no connection error, failover is unlikely to help */
            die(sprintf("SQL error: [%d] %s", $link->errno, $link->error));
            break;
    }
}
if ($res) {
    var_dump($res->fetch_assoc());
}
?>

```

## 21.9.6.4.10. Partitioning and Sharding

Copyright 1997-2012 the PHP Documentation Group. [2906]

Database clustering is done for various reasons. Clusters can improve availability, fault tolerance, and increase performance by applying a divide and conquer approach as work is distributed over many machines. Clustering is sometimes combined with partitioning and sharding to further break up a large complex task into smaller, more manageable units.

The `mysqlnd_ms` plugin aims to support a wide variety of MySQL database clusters. Some flavors of MySQL database clusters have built-in methods for partitioning and sharding, which could be transparent to use. The plugin supports the two most common approaches: MySQL Replication table filtering, and Sharding (application based partitioning).

MySQL Replication supports partitioning as filters that allow you to create slaves that replicate all or specific databases of the master, or tables. It is then in the responsibility of the application to choose a slave according to the filter rules. You can either use the `mysqlnd_ms` `node_groups` filter to manually support this, or use the experimental table filter.



Manual partitioning or sharding is supported through the node grouping filter, and SQL hints as of 1.5.0. The `node_groups` filter lets you assign a symbolic name to a group of master and slave servers. In the example, the master `master_0` and `slave_0` form a group with the name `Partition_A`. It is entirely up to you to decide what makes up a group. For example, you may use node groups for sharding, and use the group names to address shards like `Shard_A_Range_0_100`.

### Example 21.252. Cluster node groups

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\/tmp\/mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "simulate_slave_failure",
        "port": "0"
      },
      "slave_1": {
        "host": "127.0.0.1",
        "port": 3311
      }
    },
    "filters": {
      "node_groups": {
        "Partition_A": {
          "master": ["master_0"],
          "slave": ["slave_0"]
        }
      },
      "roundrobin": []
    }
  }
}
```

### Example 21.253. Manual partitioning using SQL hints

```
<?php
function select($mysqli, $msg, $hint = '') {
  /* Note: weak test, two connections to two servers may have the same thread id */
  $sql = sprintf("SELECT CONNECTION_ID() AS _thread, '%s' AS _hint FROM DUAL", $msg);
  if ($hint) {
    $sql = $hint . $sql;
  }
  if (!$res = $mysqli->query($sql)) {
    printf("[%d] %s", $mysqli->errno, $mysqli->error);
    return false;
  }
  $row = $res->fetch_assoc();
  printf("%d - %s - %s\n", $row['_thread'], $row['_hint'], $sql);
  return true;
}

$mysqli = new mysqli("myapp", "user", "password", "database");
if (!$mysqli)
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* All slaves allowed */
select($mysqli, "slave_0");
select($mysqli, "slave_1");

/* only servers of node group "Partition_A" allowed */
select($mysqli, "slave_1", "/*Partition_A*/");
```

```
select($mysqli, "slave_1", "/*Partition_A*/");
?>
```

```
6804 - slave_0 - SELECT CONNECTION_ID() AS _thread, 'slave1' AS _hint FROM DUAL
2442 - slave_1 - SELECT CONNECTION_ID() AS _thread, 'slave2' AS _hint FROM DUAL
6804 - slave_0 - /*Partition_A*/SELECT CONNECTION_ID() AS _thread, 'slave1' AS _hint FROM DUAL
6804 - slave_0 - /*Partition_A*/SELECT CONNECTION_ID() AS _thread, 'slave1' AS _hint FROM DUAL
```

By default, the plugin will use all configured master and slave servers for query execution. But if a query begins with a SQL hint like `/*node_group*/`, the plugin will only consider the servers listed in the `node_group` for query execution. Thus, `SELECT` queries prefixed with `/*Partition_A*/` will only be executed on `slave_0`.

### 21.9.6.5. Concepts

Copyright 1997-2012 the PHP Documentation Group. [2906]

This explains the architecture and related concepts for this plugin, and describes the impact that MySQL replication and this plugin have on developmental tasks while using a database cluster. Reading and understanding these concepts is required, in order to use this plugin with success.

#### 21.9.6.5.1. Architecture

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqlnd` replication and load balancing plugin is implemented as a PHP extension. It is written in C and operates under the hood of PHP. During the startup of the PHP interpreter, in the module init phase of the PHP engine, it gets registered as a `mysqlnd` plugin to replace selected `mysqlnd C` methods.

At PHP runtime, it inspects queries sent from `mysqlnd` (PHP) to the MySQL server. If a query is recognized as read-only, it will be sent to one of the configured slave servers. Statements are considered read-only if they either start with `SELECT`, the SQL hint `/*ms=slave*/` or a slave had been chosen for running the previous query, and the query started with the SQL hint `/*ms=last_used*/`. In all other cases, the query will be sent to the MySQL replication master server.

For better portability, applications should use the `MYSQLND_MS_MASTER_SWITCH`, `MYSQLND_MS_SLAVE_SWITCH`, and `MYSQLND_MS_LAST_USED_SWITCH` predefined `mysqlnd_ms constants`, instead of their literal values, such as `/*ms=slave*/`.

The plugin handles the opening and closing of database connections to both master and slave servers. From an application point of view, there continues to be only one connection handle. However, internally, this one public connection handle represents a pool of network connections that are managed by the plugin. The plugin proxies queries to the master server, and to the slaves using multiple connections.

Database connections have a state consisting of, for example, transaction status, transaction settings, character set settings, and temporary tables. The plugin will try to maintain the same state among all internal connections, whenever this can be done in an automatic and transparent way. In cases where it is not easily possible to maintain state among all connections, such as when using `BEGIN TRANSACTION`, the plugin leaves it to the user to handle.

#### 21.9.6.5.2. Connection pooling and switching

Copyright 1997-2012 the PHP Documentation Group. [2906]

The replication and load balancing plugin changes the semantics of a PHP MySQL connection handle. The existing API of the PHP MySQL extensions (`mysqli`, `mysql`, and `PDO_MYSQL`) are not changed

in a way that functions are added or removed. But their behaviour changes when using the plugin. Existing applications do not need to be adapted to a new API, but they may need to be modified because of the behaviour changes.

The plugin breaks the one-by-one relationship between a [mysqli](#), [mysql](#), and [PDO\\_MYSQL](#) connection handle and a MySQL network connection. And a [mysqli](#), [mysql](#), and [PDO\\_MYSQL](#) connection handle represents a local pool of connections to the configured MySQL replication master and MySQL replication slave servers. The plugin redirects queries to the master and slave servers. At some point in time one and the same PHP connection handle may point to the MySQL master server. Later on, it may point to one of the slave servers or still the master. Manipulating and replacing the network connection referenced by a PHP MySQL connection handle is not a transparent operation.

Every MySQL connection has a state. The state of the connections in the connection pool of the plugin can differ. Whenever the plugin switches from one wire connection to another, the current state of the user connection may change. The applications must be aware of this.

The following list shows what the connection state consists of. The list may not be complete.

- Transaction status
- Temporary tables
- Table locks
- Session system variables and session user variables
- The current database set using [USE](#) and other state chaining SQL commands
- Prepared statements
- [HANDLER](#) variables
- Locks acquired with [GET\\_LOCK\( \)](#)

Connection switches happen right before queries are executed. The plugin does not switch the current connection until the next statement is executed.



#### Replication issues

See also the MySQL reference manual chapter about [replication features](#) and related issues. Some restrictions may not be related to the PHP plugin, but are properties of the MySQL replication system.

#### Broadcasted messages

The plugins philosophy is to align the state of connections in the pool only if the state is under full control of the plugin, or if it is necessary for security reasons. Just a few actions that change the state of the connection fall into this category.

The following is a list of connection client library calls that change state, and are broadcasted to all open connections in the connection pool.

If any of the listed calls below are to be executed, the plugin loops over all open master and slave connections. The loop continues until all servers have been contacted, and the loop does not break if a server indicates a failure. If possible, the failure will propagate to the called user API function, which may be detected depending on which underlying library function was triggered.

Library call	Notes	Version
<a href="#">change_user()</a>	Called by the <a href="#">mysqli_change_user</a> user API call. Also triggered upon reuse of a persistent <a href="#">mysqli</a> connection.	Since 1.0.0.

Library call	Notes	Version
<code>select_db</code>	Called by the following user API calls: <code>mysql_select_db</code> , <code>mysql_list_tables</code> , <code>mysql_db_query</code> , <code>mysql_list_fields</code> , <code>mysqli_select_db</code> . Note, that SQL <code>USE</code> is not monitored.	Since 1.0.0.
<code>set_charset</code>	Called by the following user API calls: <code>mysql_set_charset</code> , <code>mysqli_set_charset</code> . Note, that SQL <code>SET NAMES</code> is not monitored.	Since 1.0.0.
<code>set_server_option</code>	Called by the following user API calls: <code>mysqli_multi_query</code> , <code>mysqli_real_query</code> , <code>mysqli_query</code> , <code>mysql_query</code> .	Since 1.0.0.
<code>set_client_option</code>	Called by the following user API calls: <code>mysqli_options</code> , <code>mysqli_ssl_set</code> , <code>mysqli_connect</code> , <code>mysql_connect</code> , <code>mysql_pconnect</code> .	Since 1.0.0.
<code>set_autocommit</code>	Called by the following user API calls: <code>mysqli_autocommit</code> , <code>PDO::setAttribute(PDO::ATTR_AUTOCOMMIT)</code> .	Since 1.0.0. PHP >= 5.4.0.
<code>ssl_set</code>	Called by the following user API calls: <code>mysqli_ssl_set</code> .	Since 1.1.0.

### Broadcasting and lazy connections

The plugin does not proxy or “remember” all settings to apply them on connections opened in the future. This is important to remember, if using [lazy connections](#). Lazy connections are connections which are not opened before the client sends the first connection. Use of lazy connections is the default plugin action.

The following connection library calls each changed state, and their execution is recorded for later use when lazy connections are opened. This helps ensure that the connection state of all connections in the connection pool are comparable.

Library call	Notes	Version
<code>change_user</code>	User, password and database recorded for future use.	Since 1.1.0.
<code>select_db</code>	Database recorded for future use.	Since 1.1.0.
<code>set_charset</code>	Calls <code>(set_client_option(MYSQL_SET_CHARSET_NAME, charset))</code> on lazy connection to ensure <code>charset</code> will be used upon opening the lazy connection.	Since 1.1.0.
<code>set_autocommit</code>	Adds <code>SET AUTOCOMMIT=0 1</code> to the list of init commands of a lazy connection using <code>set_client_option(MYSQL_INIT_COMMAND, "SET AUTOCOMMIT=...%quot;)</code> .	Since 1.1.0. PHP >= 5.4.0.



### Connection state

The connection state is not only changed by API calls. Thus, even if PECL `mysqlnd_ms` monitors all API calls, the application must still be aware. Ultimately, it is the applications responsibility to maintain the connection state, if needed.

### Charsets and string escaping

Due to the use of lazy connections, which are a default, it can happen that an application tries to escape a string for use within SQL statements before a connection has been established. In this case string escaping is not possible. The string escape function does not know what charset to use before a connection has been established.

To overcome the problem a new configuration setting `server_charset` has been introduced in version 1.4.0.

Attention has to be paid on escaping strings with a certain charset but using the result on a connection that uses a different charset. Please note, that PECL/mysqlnd\_ms manipulates connections and one application level connection represents a pool of multiple connections that all may have different default charsets. It is recommended to configure the servers involved to use the same default charsets. The configuration setting `server_charset` does help with this situation as well. If using `server_charset`, the plugin will set the given charset on all newly opened connections.

### 21.9.6.5.3. Transaction handling

Copyright 1997-2012 the PHP Documentation Group. [2906]

Transaction handling is fundamentally changed. An SQL transaction is a unit of work that is run on one database server. The unit of work consists of one or more SQL statements.

By default the plugin is not aware of SQL transactions. The plugin may switch connections for load balancing at any point in time. Connection switches may happen in the middle of a transaction. This is against the nature of an SQL transaction. By default, the plugin is not transaction safe.

Any kind of MySQL load balancer must be hinted about the begin and end of a transaction. Hinting can either be done implicitly by monitoring API calls or using SQL hints. Both options are supported by the plugin, depending on your PHP version. API monitoring requires PHP 5.4.0 or newer. The plugin, like any other MySQL load balancer, cannot detect transaction boundaries based on the MySQL Client Server Protocol. Thus, entirely transparent transaction aware load balancing is not possible. The least intrusive option is API monitoring, which requires little to no application changes, depending on your application.

Please, find examples of using SQL hints or the API monitoring in the [examples section](#). The details behind the API monitoring, which makes the plugin transaction aware, are described below.

Beginning with PHP 5.4.0, the [mysqlnd](#) library allows this plugin to subclass the library C API call `set_autocommit()`, to detect the status of `autocommit` mode.

The PHP MySQL extensions either issue a query (such as `SET AUTOCOMMIT=0|1`), or use the `mysqlnd` library call `set_autocommit()` to control the `autocommit` setting. If an extension makes use of `set_autocommit()`, the plugin can be made transaction aware. Transaction awareness cannot be achieved if using SQL to set the `autocommit` mode. The library function `set_autocommit()` is called by the `mysqli_autocommit` and `PDO::setAttribute(PDO::ATTR_AUTOCOMMIT)` user API calls.

The plugin configuration option `trx_stickiness=master` can be used to make the plugin transactional aware. In this mode, the plugin stops load balancing if `autocommit` becomes disabled, and directs all statements to the master until `autocommit` gets enabled.

An application that does not want to set SQL hints for transactions but wants to use the transparent API monitoring to avoid application changes must make sure that the `autocommit` settings is changed exclusively through the listed API calls.

API based transaction boundary detection has been improved with PHP 5.5.0 and PECL/mysqlnd\_ms 1.5.0 to cover not only calls to `mysqli_autocommit` but also `mysqli_begin`, `mysqli_commit` and `mysqli_rollback`.

### 21.9.6.5.4. Error handling

Copyright 1997-2012 the PHP Documentation Group. [2906]

Applications using PECL/mysqlnd\_ms should implement proper error handling for all user API calls. And because the plugin changes the semantics of a connection handle, API calls may return unexpected errors. If using the plugin on a connection handle that no longer represents an individual network connection, but a connection pool, an error code and error message will be set on the connection handle whenever an error occurs on any of the network connections behind.

If using lazy connections, which is the default, connections are not opened until they are needed for query execution. Therefore, an API call for a statement execution may return a connection error. In the example below, an error is provoked when trying to run a statement on a slave. Opening a slave connection fails because the plugin configuration file lists an invalid host name for the slave.

**Example 21.254. Provoking a connection error**

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost",
        "socket": "\tmp\mysql.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "invalid_host_name",
      }
    },
    "lazy_connections": 1
  }
}
```

The explicit activation of lazy connections is for demonstration purpose only.

**Example 21.255. Connection error on query execution**

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (mysqli_connect_errno())
  /* Of course, your error handling is nicer... */
  die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Connection 1, connection bound SQL user variable, no SELECT thus run on master */
if (!$mysqli->query("SET @myrole='master'")) {
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}

/* Connection 2, run on slave because SELECT, provoke connection error */
if (!$res = $mysqli->query("SELECT @myrole AS _role")) {
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
} else {
  $row = $res->fetch_assoc();
  $res->close();
  printf("@myrole = '%s'\n", $row['_role']);
}
$mysqli->close();
?>
```

The above example will output something similar to:

```
PHP Warning: mysqli::query(): php_network_getaddresses: getaddrinfo failed: Name or service not known in %
PHP Warning: mysqli::query(): [2002] php_network_getaddresses: getaddrinfo failed: Name or service not kno
[2002] php_network_getaddresses: getaddrinfo failed: Name or service not known
```

Applications are expected to handle possible connection errors by implementing proper error handling.

Depending on the use case, applications may want to handle connection errors differently from other errors. Typical connection errors are `2002 (CR_CONNECTION_ERROR) - Can't connect to local MySQL server through socket '%s' (%d)`, `2003 (CR_CONN_HOST_ERROR)`

- Can't connect to MySQL server on '%s' (%d) and 2005 (CR\_UNKNOWN\_HOST) - Unknown MySQL server host '%s' (%d). For example, the application may test for the error codes and manually perform a fail over. The plugins philosophy is not to offer automatic fail over, beyond master fail over, because fail over is not a transparent operation.

### Example 21.256. Provoking a connection error

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "invalid_host_name"
      },
      "slave_1": {
        "host": "192.168.78.136"
      }
    },
    "lazy_connections": 1,
    "filters": {
      "roundrobin": [
        ]
      ]
    }
  }
}
```

Explicitly activating lazy connections is done for demonstration purposes, as is round robin load balancing as opposed to the default `random once` type.

### Example 21.257. Most basic failover

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
if (mysqli_connect_errno())
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Connection 1, connection bound SQL user variable, no SELECT thus run on master */
if (!$mysqli->query("SET @myrole='master'")) {
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
}

/* Connection 2, first slave */
$res = $mysqli->query("SELECT VERSION() AS _version");
/* Hackish manual fail over */
if (2002 == $mysqli->errno || 2003 == $mysqli->errno || 2004 == $mysqli->errno) {
    /* Connection 3, first slave connection failed, trying next slave */
    $res = $mysqli->query("SELECT VERSION() AS _version");
}

if (!$res) {
    printf("ERROR, [%d] '%s'\n", $mysqli->errno, $mysqli->error);
} else {
    /* Error messages are taken from connection 3, thus no error */
    printf("SUCCESS, [%d] '%s'\n", $mysqli->errno, $mysqli->error);
    $row = $res->fetch_assoc();
    $res->close();
    printf("version = %s\n", $row['_version']);
}
$mysqli->close();
?>
```



The above example will output something similar to:

```
[1045] Access denied for user 'username'@'localhost' (using password: YES)
PHP Warning: mysqli::query(): php_network_getaddresses: getaddrinfo failed: Name or service not known in 3
PHP Warning: mysqli::query(): [2002] php_network_getaddresses: getaddrinfo failed: Name or service not known
SUCCESS, [0] ''
version = 5.6.2-m5-log
```

In some cases, it may not be easily possible to retrieve all errors that occur on all network connections through a connection handle. For example, let's assume a connection handle represents a pool of three open connections. One connection to a master and two connections to the slaves. The application changes the current database using the user API call `mysqli_select_db`, which then calls the `mysqld` library function to change the schemata. `mysqld_ms` monitors the function, and tries to change the current database on all connections to harmonize their state. Now, assume the master succeeds in changing the database, and both slaves fail. Upon the initial error from the first slave, the plugin will set an appropriate error on the connection handle. The same is done when the second slave fails to change the database. The error message from the first slave is lost.

Such cases can be debugged by either checking for errors of the type `E_WARNING` (see above) or, if no other option, investigation of the [mysqld\\_ms debug and trace log](#).

### 21.9.6.5.5. Failover

Copyright 1997-2012 the PHP Documentation Group. [2906]

By default, connection failover handling is left to the user. The application is responsible for checking return values of the database functions it calls and reacting to possible errors. If, for example, the plugin recognizes a query as a read-only query to be sent to the slave servers, and the slave server selected by the plugin is not available, the plugin will raise an error after not executing the statement.

*Default: manual failover*

It is up to the application to handle the error and, if required, re-issue the query to trigger the selection of another slave server for statement execution. The plugin will make no attempts to failover automatically, because the plugin cannot ensure that an automatic failover will not change the state of the connection. For example, the application may have issued a query which depends on SQL user variables which are bound to a specific connection. Such a query might return incorrect results if the plugin would switch the connection implicitly as part of automatic failover. To ensure correct results, the application must take care of the failover, and rebuild the required connection state. Therefore, by default, no automatic failover is performed by the plugin.

A user that does not change the connection state after opening a connection may activate automatic failover. Please note, that automatic failover logic is limited to connection attempts. Automatic failover is not used for already established connections. There is no way to instruct the plugin to attempt failover on a connection that has been connected to MySQL already in the past.

*Automatic failover*

The failover policy is configured in the plugins configuration file, by using the [failover](#) configuration directive.

Automatic and silent failover can be enabled through the [failover](#) configuration directive. Automatic failover can either be configured to try exactly one master after a slave failure or, alternatively, loop over slaves and masters before returning an error to the user. The number of connection attempts can be limited and failed hosts can be excluded from future load balancing attempts. Limiting the number of retries and remembering failed hosts are considered experimental features, albeit being reasonable stable. Syntax and semantics may change in future versions.



Please note, since version 1.5.0 automatic failover is disabled for the duration of a transaction if transaction stickiness is enabled and transaction boundaries have been detected. The plugin will not switch connections for the duration of a transaction. It will also not perform automatic and silent failover. Instead an error will be thrown. It is then left to the user to handle the failure of the transaction. Please check, the [trx\\_stickiness](#) documentation how to do this.

A basic manual failover example is provided within the [error handling](#) section.

#### *Standby servers*

Using [weighted load balancing](#), introduced in PECL/mysqld 1.4.0, it is possible to configure standby servers that are sparsely used during normal operations. A standby server that is primarily used as a worst-case standby failover target can be assigned a very low weight/priority in relation to all other servers. As long as all servers are up and running the majority of the workload is assigned to the servers which have high weight values. Few requests will be directed to the standby system which has a very low weight value.

Upon failure of the servers with a high priority, you can still failover to the standby, which has been given a low load balancing priority by assigning a low weight to it. Failover can be done manually or automatically. If done automatically, you may want to combine it with the [remember\\_failed](#) option.

At this point, it is not possible to instruct the load balancer to direct no requests at all to a standby. This may not be much of a limitation given that the highest weight you can assign to a server is 65535. Given two slaves, of which one shall act as a standby and has been assigned a weight of 1, the standby will have to handle far less than one percent of the overall workload.

#### *Failover and primary copy*

Please note, if using a primary copy cluster, such as MySQL Replication, it is difficult to do connection failover in case of a master failure. At any time there is only one master in the cluster for a given dataset. The master is a single point of failure. If the master fails, clients have no target to fail over write requests. In case of a master outage the database administrator must take care of the situation and update the client configurations, if need be.

### **21.9.6.5.6. Load balancing**

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Four load balancing strategies are supported to distribute statements over the configured MySQL slave servers:

random	Chooses a random server whenever a statement is executed.
random once (default)	Chooses a random server after the first statement is executed, and uses the decision for the rest of the PHP request.  It is the default, and the lowest impact on the connection state.
round robin	Iterates over the list of configured servers.
user-defined via callback	Is used to implement any other strategy.

The load balancing policy is configured in the plugins configuration file using the [random](#), [roundrobin](#), and [user filters](#).

Servers can be prioritized assigning a weight. A server that has been given a weight of two will get twice as many requests as a server that has been given the default weight of one. Prioritization can be handy in heterogenous environments. For example, you may want to assign more requests to a powerful machine than to a less powerful. Or, you may have configured servers that are close or far from the client, thus expose different latencies.

### 21.9.6.5.7. Read-write splitting

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin executes read-only statements on the configured MySQL slaves, and all other queries on the MySQL master. Statements are considered read-only if they either start with `SELECT`, the SQL hint `/*ms=slave*/`, or if a slave had been chosen for running the previous query and the query starts with the SQL hint `/*ms=last_used*/`. In all other cases, the query will be sent to the MySQL replication master server. It is recommended to use the constants `MYSQLND_MS_SLAVE_SWITCH`, `MYSQLND_MS_MASTER_SWITCH` and `MYSQLND_MS_LAST_USED_SWITCH` instead of `/*ms=slave*/`. See also the [list of `mysqld\_ms` constants](#).

SQL hints are a special kind of standard compliant SQL comments. The plugin does check every statement for certain SQL hints. The SQL hints are described within the [`mysqld\_ms` constants](#) documentation, constants that are exported by the extension. Other systems involved with the statement processing, such as the MySQL server, SQL firewalls, and SQL proxies, are unaffected by the SQL hints, because those systems are designed to ignore SQL comments.

The built-in read-write splitter can be replaced by a user-defined filter, see also the [user filter](#) documentation.

A user-defined read-write splitter can request the built-in logic to send a statement to a specific location, by invoking `mysqld_ms_is_select`.



#### Note

The built-in read-write splitter is not aware of multi-statements. Multi-statements are seen as one statement. The splitter will check the beginning of the statement to decide where to run the statement. If, for example, a multi-statement begins with `SELECT 1 FROM DUAL; INSERT INTO test(id) VALUES (1); ...` the plugin will run it on a slave although the statement is not read-only.

### 21.9.6.5.8. Filter

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Version requirement

Filters exist as of `mysqld_ms` version 1.1.0-beta.

**filters.** PHP applications that implement a MySQL replication cluster must first identify a group of servers in the cluster which could execute a statement before the statement is executed by one of the candidates. In other words: a defined list of servers must be filtered until only one server is available.

The process of filtering may include using one or more filters, and filters can be chained. And they are executed in the order they are defined in the plugins configuration file.



#### Explanation: comparing filter chaining to pipes

The concept of chained filters can be compared to using pipes to connect command line utilities on an operating system command shell. For example, an input stream is passed to a processor, filtered, and then transferred to be output. Then, the output is passed as input to the next command, which is connected to the previous using the pipe operator.

Available filters:

- Load balancing filters: [random](#) and [roundrobin](#).
- Selection filter: [user](#), [user\\_multi](#), [quality\\_of\\_service](#).

The `random` filter implements the 'random' and 'random once' load balancing policies. The 'round robin' load balancing can be configured through the `roundrobin` filter. Setting a 'user defined callback' for server selection is possible with the `user` filter. The `quality_of_service` filter finds cluster nodes capable of delivering a certain service, for example, read-your-writes or, not lagging more seconds behind the master than allowed.

Filters can accept parameters to change their behaviour. The `random` filter accepts an optional `sticky` parameter. If set to true, the filter changes load balancing from random to random once. Random picks a random server every time a statement is to be executed. Random once picks a random server when the first statement is to be executed and uses the same server for the rest of the PHP request.

One of the biggest strength of the filter concept is the possibility to chain filters. This strength does not become immediately visible because the `random`, `roundrobin` and `user` filters are supposed to output no more than one server. If a filter reduces the list of candidates for running a statement to only one server, it makes little sense to use that one server as input for another filter for further reduction of the list of candidates.

An example filter sequence that will fail:

- Statement to be executed: `SELECT 1 FROM DUAL`. Passed to all filters.
- All configured nodes are passed as input to the first filter. Master nodes: `master_0`. Slave nodes: `slave_0`, `slave_1`
- Filter: `random`, argument `sticky=1`. Picks a random slave once to be used for the rest of the PHP request. Output: `slave_0`.
- Output of `slave_0` and the statement to be executed is passed as input to the next filter. Here: `roundrobin`, server list passed to filter is: `slave_0`.
- Filter: `roundrobin`. Server list consists of one server only, round robin will always return the same server.

If trying to use such a filter sequence, the plugin may emit a warning like `(mysqld_ms) Error while creating filter '%s' . Non-multi filter '%s' already created. Stopping in %s on line %d`. Furthermore, an appropriate error on the connection handle may be set.

A second type of filter exists: multi filter. A multi filter emits zero, one or multiple servers after processing. The `quality_of_service` filter is an example. If the service quality requested sets an upper limit for the slave lag and more than one slave is lagging behind less than the allowed number of seconds, the filter returns more than one cluster node. A multi filter must be followed by other to further reduce the list of candidates for statement execution until a candidate is found.

A filter sequence with the `quality_of_service` multi filter followed by a load balancing filter.

- Statement to be executed: `SELECT sum(price) FROM orders WHERE order_id = 1`. Passed to all filters.
- All configured nodes are passed as input to the first filter. Master nodes: `master_0`. Slave nodes: `slave_0`, `slave_1`, `slave_2`, `slave_3`
- Filter: `quality_of_service`, rule set: `session_consistency (read-your-writes)` Output: `master_0`
- Output of `master_0` and the statement to be executed is passed as input to the next filter, which is `roundrobin`.
- Filter: `roundrobin`. Server list consists of one server. Round robin selects `master_0`.

A filter sequence must not end with a multi filter. If trying to use a filter sequence which ends with a multi filter the plugin may emit a warning like `(mysqld_ms) Error in configuration. Last`

`filter` is multi filter. Needs to be non-multi one. Stopping in %s on line %d. Furthermore, an appropriate error on the connection handle may be set.



### Speculation towards the future: MySQL replication filtering

In future versions, there may be additional multi filters. For example, there may be a `table` filter to support MySQL replication filtering. This would allow you to define rules for which database or table is to be replicated to which node of a replication cluster. Assume your replication cluster consists of four slaves (`slave_0`, `slave_1`, `slave_2`, `slave_3`) two of which replicate a database named `sales` (`slave_0`, `slave_1`). If the application queries the database `sales`, the hypothetical `table` filter reduces the list of possible servers to `slave_0` and `slave_1`. Because the output and list of candidates consists of more than one server, it is necessary and possible to add additional filters to the candidate list, for example, using a load balancing filter to identify a server for statement execution.

#### 21.9.6.5.9. Service level and consistency

Copyright 1997-2012 the PHP Documentation Group. [2906]



### Version requirement

Service levels have been introduced in `mysqlnd_ms` version 1.2.0-alpha. `mysqlnd_ms_set_qos` requires PHP 5.4.0 or newer.

The plugin can be used with different kinds of MySQL database clusters. Different clusters can deliver different levels of service to applications. The service levels can be grouped by the data consistency levels that can be achieved. The plugin knows about:

- eventual consistency
- session consistency
- strong consistency

Depending how a cluster is used it may be possible to achieve higher service levels than the default one. For example, a read from an asynchronous MySQL replication slave is eventual consistent. Thus, one may say the default consistency level of a MySQL replication cluster is eventual consistency. However, if the master only is used by a client for reading and writing during a session, session consistency (read your writes) is given. PECL `mysqlnd` 1.2.0 abstracts the details of choosing an appropriate node for any of the above service levels from the user.

Service levels can be set through the qualify-of-service filter in the [plugins configuration file](#) and at runtime using the function `mysqlnd_ms_set_qos`.

The plugin defines the different service levels as follows.

Eventual consistency is the default service provided by an asynchronous cluster, such as classical MySQL replication. A read operation executed on an arbitrary node may or may not return stale data. The applications view of the data is eventual consistent.

Session consistency is given if a client can always read its own writes. An asynchronous MySQL replication cluster can deliver session consistency if clients always use the master after the first write or never query a slave which has not yet replicated the clients write operation.

The plugins understanding of strong consistency is that all clients always see the committed writes of all other clients. This is the default when using MySQL Cluster or any other cluster offering synchronous data distribution.

*Service level parameters*

Eventual consistency and session consistency service level accept parameters.

Eventual consistency is the service provided by classical MySQL replication. By default, all nodes qualify for read requests. An optional `age` parameter can be given to filter out nodes which lag more than a certain number of seconds behind the master. The plugin is using `SHOW SLAVE STATUS` to measure the lag. Please, see the MySQL reference manual to learn about accuracy and reliability of the `SHOW SLAVE STATUS` command.

Session consistency (read your writes) accepts an optional `GTID` parameter to consider reading not only from the master but also from slaves which already have replicated a certain write described by its transaction identifier. This way, when using asynchronous MySQL replication, read requests may be load balanced over slaves while still ensuring session consistency.

The latter requires the use of [client-side global transaction id injection](#).

#### *Advantages of the new approach*

The new approach supersedes the use of SQL hints and the configuration option `master_on_write` in some respects. If an application running on top of an asynchronous MySQL replication cluster cannot accept stale data for certain reads, it is easier to tell the plugin to choose appropriate nodes than prefixing all read statements in question with the SQL hint to enforce the use of the master. Furthermore, the plugin may be able to use selected slaves for reading.

The `master_on_write` configuration option makes the plugin use the master after the first write (session consistency, read your writes). In some cases, session consistency may not be needed for the rest of the session but only for some, few read operations. Thus, `master_on_write` may result in more read load on the master than necessary. In those cases it is better to request a higher than default service level only for those reads that actually need it. Once the reads are done, the application can return to default service level. Switching between service levels is only possible using `mysqLnd_ms_set_qos`.

#### *Performance considerations*

A MySQL replication cluster cannot tell clients which slaves are capable of delivering which level of service. Thus, in some cases, clients need to query the slaves to check their status. PECL `mysqLnd_ms` transparently runs the necessary SQL in the background. However, this is an expensive and slow operation. SQL statements are run if eventual consistency is combined with an `age` (slave lag) limit and if session consistency is combined with a global transaction ID.

If eventual consistency is combined with an maximum `age` (slave lag), the plugin selects candidates for statement execution and load balancing for each statement as follows. If the statement is a write all masters are considered as candidates. Slaves are not checked and not considered as candidates. If the statement is a read, the plugin transparently executes `SHOW SLAVE STATUS` on every slaves connection. It will loop over all connections, send the statement and then start checking for results. Usually, this is slightly faster than a loop over all connections in which for every connection a query is send and the plugin waits for its results. A slave is considered a candidate if `SHOW SLAVE STATUS` reports `Slave_IO_Running=Yes`, `Slave_SQL_Running=Yes` and `Seconds_Behind_Master` is less or equal than the allowed maximum `age`. In case of an SQL error, the plugin emits a warning but does not set an error on the connection. The error is not set to make it possible to use the plugin as a drop-in.

If session consistency is combined with a global transaction ID, the plugin executes the SQL statement set with the `fetch_last_gtid` entry of the `global_transaction_id_injection` section from the plugins configuration file. Further details are identical to those described above.

In version 1.2.0 no additional optimizations are done for executing background queries. Future versions may contain optimizations, depending on user demand.

If no parameters and options are set, no SQL is needed. In that case, the plugin consider all nodes of the type shown below.

- eventual consistency, no further options set: all masters, all slaves
- session consistency, no further options set: all masters
- strong consistency (no options allowed): all masters

#### *Throttling*

The quality of service filter can be combined with [Global transaction IDs](#) to throttle clients. Throttling does reduce the write load on the master by slowing down clients. If session consistency is requested and global transactions identifier are used to check the status of a slave, the check can be done in two ways. By default a slave is checked and skipped immediately if it does not match the criteria for session consistency. Alternatively, the plugin can wait for a slave to catch up to the master until session consistency is possible. To enable the throttling, you have to set [wait\\_for\\_gtid\\_timeout](#) configuration option.

### 21.9.6.5.10. Global transaction IDs

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### **Version requirement**

Client side global transaction ID injection exists as of `mysqLnd_ms` version 1.2.0-alpha. Transaction boundaries are detected by monitoring API calls. This is possible as of PHP 5.4.0. Please, see also [Transaction handling](#).

As of MySQL 5.6.5-m8 the MySQL server features built-in global transaction identifiers. The MySQL built-in global transaction ID feature is supported by PECL/`mysqLnd_ms` 1.3.0-alpha or later. Neither are client-side transaction boundary monitoring nor any setup activities required if using the server feature.

#### *Idea and client-side emulation*

PECL/`mysqLnd_ms` can do client-side transparent global transaction ID injection. In its most basic form, a global transaction identifier is a counter which is incremented for every transaction executed on the master. The counter is held in a table on the master. Slaves replicate the counter table.

In case of a master failure a database administrator can easily identify the most recent slave for promoting it as a new master. The most recent slave has the highest transaction identifier.

Application developers can ask the plugin for the global transaction identifier (GTID) for their last successful write operation. The plugin will return an identifier that refers to a transaction no older than that of the clients last write operation. Then, the GTID can be passed as a parameter to the quality of service (QoS) filter as an option for session consistency. Session consistency ensures read your writes. The filter ensures that all reads are either directed to a master or a slave which has replicated the write referenced by the GTID.

#### *When injection is done*

The plugin transparently maintains the GTID table on the master. In autocommit mode the plugin injects an `UPDATE` statement before executing the users statement for every master use. In manual transaction mode, the injection is done before the application calls `commit()` to close a transaction. The configuration option `report_error` of the GTID section in the plugins configuration file is used to control whether a failed injection shall abort the current operation or be ignored silently (default).

Please note, the PHP version requirements for [transaction boundary monitoring](#) and their limits.

#### *Limitations*

Client-side global transaction ID injection has shortcomings. The potential issues are not specific to PECL/`mysqLnd_ms` but are rather of general nature.



- Global transaction ID tables must be deployed on all masters and replicas.
- The GTID can have holes. Only PHP clients using the plugin will maintain the table. Other clients will not.
- Client-side transaction boundary detection is based on API calls only.
- Client-side transaction boundary detection does not take implicit commit into account. Some MySQL SQL statements cause an implicit commit and cannot be rolled back.

*Using server-side global transaction identifier*

Starting with PECL/mysqlnd\_ms 1.3.0-alpha the MySQL 5.6.5-m8 or newer built-in global transaction identifier feature is supported. Use of the server feature lifts all of the above listed limitations. Please, see the MySQL Reference Manual for limitations and preconditions for using server built-in global transaction identifiers.

Whether to use the client-side emulation or the server built-in functionality is a question not directly related to the plugin, thus it is not discussed in depth. There are no plans to remove the client-side emulation and you can continue to use it, if the server-side solution is no option. This may be the case in heterogenous environments with old MySQL server or, if any of the server-side solution limitations is not acceptable.

From an applications perspective there is hardly a difference in using one or the other approach. The following properties differ.

- Client-side emulation, as shown in the manual, is using an easy to compare sequence number for global transactions. Multi-master is not handled to keep the manual examples easy.

Server-side built-in feature is using a combination of a server identifier and a sequence number as a global transaction identifier. Comparison cannot use numeric algebra. Instead a SQL function must be used. Please, see the MySQL Reference Manual for details.

- Plugin global transaction ID statistics are only available with client-side emulation because they monitor the emulation.



**Global transaction identifiers in distributed systems**

Global transaction identifiers can serve multiple purposes in the context of distributed systems, such as a database cluster. Global transaction identifiers can be used for, for example, system wide identification of transactions, global ordering of transactions, heartbeat mechanism and for checking the replication status of replicas. PECL/mysqlnd\_ms, a clientside driver based software, does focus on using GTIDs for tasks that can be handled at the client, such as checking the replication status of replicas for asynchronous replication setups.

**21.9.6.5.11. Cache integration**

Copyright 1997-2012 the PHP Documentation Group. [2906]



**Version requirement**

The feature requires use of PECL/mysqlnd\_ms 1.3.0-beta or later, and PECL/mysqlnd\_qc 1.1.0-alpha or newer. PECL/mysqlnd\_ms must be compiled to support the feature. PHP 5.4.0 or newer is required.



**Setup: extension load order**

PECL/mysqlnd\_ms must be loaded before PECL/mysqlnd\_qc, when using shared extensions.



### Feature stability

The cache integration is of beta quality.



### Suitable MySQL clusters

The feature is targeted for use with MySQL Replication (primary copy). Currently, no other kinds of MySQL clusters are supported. Users of such cluster must control PECL/[mysqld\\_qc](#) manually if they are interested in client-side query caching.

Support for MySQL replication clusters (asynchronous primary copy) is the main focus of PECL/[mysqld\\_ms](#). The slaves of a MySQL replication cluster may or may not reflect the latest updates from the master. Slaves are asynchronous and can lag behind the master. A read from a slave is eventual consistent from a cluster-wide perspective.

The same level of consistency is offered by a local cache using time-to-live (TTL) invalidation strategy. Current data or stale data may be served. Eventually, data searched for in the cache is not available and the source of the cache needs to be accessed.

Given that both a MySQL Replication slave (asynchronous secondary) and a local TTL-driven cache deliver the same level of service it is possible to transparently replace a remote database access with a local cache access to gain better possibility.

As of PECL/[mysqld\\_ms](#) 1.3.0-beta the plugin is capable of transparently controlling PECL/[mysqld\\_ms](#) 1.1.0-alpha or newer to cache a read-only query if explicitly allowed by setting an appropriate quality of service through [mysqld\\_ms\\_set\\_qos](#). Please, see the [quickstart](#) for a code example. Both plugins must be installed, PECL/[mysqld\\_ms](#) must be compiled to support the cache feature and PHP 5.4.0 or newer has to be used.

Applications have full control of cache usage and can request fresh data at any time, if need be. The cache usage can be enabled and disabled time during the execution of a script. The cache will be used if [mysqld\\_ms\\_set\\_qos](#) sets the quality of service to eventual consistency and enables cache usage. Cache usage is disabled by requesting higher consistency levels, for example, session consistency (read your writes). Once the quality of service has been relaxed to eventual consistency the cache can be used again.

If caching is enabled for a read-only statement, PECL/[mysqld\\_ms](#) may inject [SQL hints to control caching](#) by PECL/[mysqld\\_qc](#). It may modify the SQL statement it got from the application. Subsequent SQL processors are supposed to ignore the SQL hints. A SQL hint is a SQL comment. Comments must not be ignored, for example, by the database server.

The TTL of a cache entry is computed on a per statement basis. Applications set an maximum age for the data they want to retrieve using [mysqld\\_ms\\_set\\_qos](#). The age sets an approximate upper limit of how many seconds the data returned may lag behind the master.

The following logic is used to compute the actual TTL if caching is enabled. The logic takes the estimated slave lag into account for choosing a TTL. If, for example, there are two slaves lagging 5 and 10 seconds behind and the maximum age allowed is 60 seconds, the TTL is set to 50 seconds. Please note, the age setting is no more than an estimated guess.

- Check whether the statement is read-only. If not, don't cache.
- If caching is enabled, check the slave lag of all configured slaves. Establish slave connections if none exist so far and lazy connections are used.
- Send [SHOW SLAVE STATUS](#) to all slaves. Do not wait for the first slave to reply before sending to the second slave. Clients often wait long for replies, thus we send out all requests in a burst before fetching in a second stage.



- Loop over all slaves. For every slave wait for its reply. Do not start checking another slave before the currently waited for slave has replied. Check for [Slave\\_IO\\_Running=Yes](#) and [Slave\\_SQL\\_Running=Yes](#). If both conditions hold true, fetch the value of [Seconds\\_Behind\\_Master](#). In case of any errors or if conditions fail, set an error on the slave connection. Skip any such slave connection for the rest of connection filtering.
- Search for the maximum value of [Seconds\\_Behind\\_Master](#) from all slaves that passed the previous conditions. Subtract the value from the maximum age provided by the user with [mysqLnd\\_ms\\_set\\_qos](#). Use the result as a TTL.
- The filtering may sort out all slaves. If so, the maximum age is used as TTL, because the maximum lag found equals zero. It is perfectly valid to sort out all slaves. In the following it is up to subsequent filter to decide what to do. The built-in load balancing filter will pick the master.
- Inject the appropriate SQL hints to enable caching by PECL/mysqLnd\_qc.
- Proceed with the connection filtering, e.g. apply load balancing rules to pick a slave.
- PECL/mysqLnd\_qc is loaded after PECL/mysqLnd\_ms by PHP. Thus, it will see all query modifications of PECL/mysqLnd\_ms and cache the query if instructed to do so.

The algorithm may seem expensive. [SHOW SLAVE STATUS](#) is a very fast operation. Given a sufficient number of requests and cache hits per second the cost of checking the slaves lag can easily outweigh the costs of the cache decision.

Suggestions on a better algorithm are always welcome.

#### 21.9.6.5.12. Supported clusters

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Any application using any kind of MySQL cluster is faced with the same tasks:

- Identify nodes capable of executing a given statement with the required service level
- Load balance requests within the list of candidates
- Automatic fail over within candidates, if needed

The plugin is optimized for fulfilling these tasks in the context of a classical asynchronous MySQL replication cluster consisting of a single master and many slaves (primary copy). When using classical, asynchronous MySQL replication all of the above listed tasks need to be mastered at the client side.

Other types of MySQL cluster may have lower requirements on the application side. For example, if all nodes in the cluster can answer read and write requests, no read-write splitting needs to be done (multi-master, update-all). If all nodes in the cluster are synchronous, they automatically provide the highest possible quality of service which makes choosing a node easier. In this case, the plugin may serve the application after some reconfiguration to disable certain features, such as built-in read-write splitting.



#### Documentation focus

The documentation focusses describing the use of the plugin with classical asynchronous MySQL replication clusters (primary copy). Support for this kind of cluster has been the original development goal. Use of other clusters is briefly described below. Please note, that this is still work in progress.

*Primary copy (MySQL Replication)*

This is the primary use case of the plugin. Follow the hints given in the descriptions of each feature.

- Configure one master and one or more slaves. [Server configuration details \[3280\]](#) are given in the setup section.
- Use random load balancing policy together with the `sticky` flag.
- If you do not plan to use the `service level` API calls, add the `master on write` flag.
- Please, make yourself aware of the properties of automatic failover before adding a `failover` directive.
- Consider the use of `trx_stickiness` to execute transactions on the primary only. Please, read carefully how it works before you rely on it.

#### Example 21.258. Enabling the plugin (`php.ini`)

```
mysqlnd_ms.enable=1
mysqlnd_ms.config_file=/path/to/mysqlnd_ms_plugin.ini
```

#### Example 21.259. Basic plugin configuration (`mysqlnd_ms_plugin.ini`) for MySQL Replication

```
{
  "myapp": {
    "master": {
      "master_1": {
        "host": "localhost",
        "socket": "\tmp\mysql57.sock"
      }
    },
    "slave": {
      "slave_0": {
        "host": "127.0.0.1",
        "port": 3308
      },
      "slave_1": {
        "host": "192.168.2.28",
        "port": 3306
      }
    },
    "filters": {
      "random": {
        "sticky": "1"
      }
    }
  }
}
```

#### *Primary copy with multi primaries (MMM - MySQL Multi Master)*

MySQL Replication allows you to create cluster topologies with multiple masters (primaries). Write-write conflicts are not handled by the replication system. This is no update anywhere setup. Thus, data must be partitioned manually and clients must be redirected in accordance to the partitioning rules. The recommended setup is equal to the sharding setup below.

#### *Manual sharding, possibly combined with primary copy and multiple primaries*

Use SQL hints and the node group filter for clusters that use data partitioning but leave query redirection to the client. The example configuration shows a multi master setup with two shards.

#### Example 21.260. Multiple primaries - multi master (`php.ini`)

```
mysqlnd_ms.enable=1
```

```
mysqlnd_ms.config_file=/path/to/mysqlnd_ms_plugin.ini
mysqlnd_ms.multi_master=1
```

### Example 21.261. Primary copy with multiple primaries and partitioning

```
{
  "myapp": {
    "master": {
      "master_1": {
        "host": "localhost",
        "socket": "\tmp\mysql57.sock"
      }
      "master_2": {
        "host": "192.168.2.27",
        "socket": "3306"
      }
    },
    "slave": {
      "slave_1": {
        "host": "127.0.0.1",
        "port": 3308
      },
      "slave_2": {
        "host": "192.168.2.28",
        "port": 3306
      }
    },
    "filters": {
      "node_groups": {
        "Partition_A": {
          "master": ["master_1"],
          "slave": ["slave_1"]
        },
        "Partition_B": {
          "master": ["master_2"],
          "slave": ["slave_2"]
        }
      }
    },
    "roundrobin": []
  }
}
```

The plugin can also be used with a loose collection of unrelated shards. For such a cluster, configure masters only and disable read write splitting. The nodes of such a cluster are called masters in the plugin configuration as they accept both reads and writes for their partition.

#### *Using synchronous update everywhere clusters such as MySQL Cluster*

MySQL Cluster is a synchronous cluster solution. All cluster nodes accept read and write requests. In the context of the plugin, all nodes shall be considered as masters.

Use the load balancing and fail over features only.

- Disable the plugins [built-in read-write splitting](#).
- Configure masters only.
- Consider random once load balancing strategy, which is the plugins default. If random once is used, only masters are configured and no SQL hints are used to force using a certain node, no connection switches will happen for the duration of a web request. Thus, no special handling is required for transactions. The plugin will pick one master at the beginning of the PHP script and use it until the script terminates.

- Do not set the quality of service. All nodes have all the data. This automatically gives you the highest possible service quality (strong consistency).
- Do not enable client-side global transaction injection. It is neither required to help with server-side fail over nor to assist the quality of service filter choosing an appropriate node.

Disabling built-in read-write splitting.

- Set `mysqld_ms.disable_rw_split=1`
- Do not use [SQL hints](#) to enforce the use of slaves

Configure masters only.

- Set `mysqld_ms.multi_master=1`.
- Do not configure any slaves.
- Set `failover=loop_before_master` in the plugins configuration file to avoid warnings about the empty slave list and to make the failover logic loop over all configured masters before emitting an error.

Please, note the warnings about automatic failover given in the previous sections.

#### Example 21.262. Multiple primaries - multi master (`php.ini`)

```
mysqld_ms.enable=1
mysqld_ms.config_file=/path/to/mysqld_ms_plugin.ini
mysqld_ms.multi_master=1
mysqld_ms.disable_rw_split=1
```

#### Example 21.263. Synchronous update anywhere cluster

```
"myapp": {
  "master": {
    "master_1": {
      "host": "localhost",
      "socket": "\\tmp\\mysql57.sock"
    },
    "master_2": {
      "host": "192.168.2.28",
      "port": 3306
    }
  },
  "slave": {
  },
  "filters": {
    "roundrobin": {
    }
  },
  "failover": {
    "strategy": "loop_before_master",
    "remember_failed": true
  }
}
```

If running an update everywhere cluster that has no built-in partitioning to avoid hot spots and high collision rates, consider using the node groups filter to keep updates on a frequently accessed table on one of the nodes. This may help to reduce collision rates and thus improve performance.

## 21.9.6.6. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

### 21.9.6.6.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

PHP 5.3.6 or newer. Some advanced functionality requires PHP 5.4.0 or newer.

The [mysqlnd\\_ms](#) replication and load balancing plugin supports all PHP applications and all available PHP MySQL extensions ([mysqli](#), [mysql](#), [PDO\\_MYSQL](#)). The PHP MySQL extension must be configured to use [mysqlnd](#) in order to be able to use the [mysqlnd\\_ms](#) plugin for [mysqlnd](#).

### 21.9.6.6.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

This [PECL](#) extension is not bundled with PHP.

Information for installing this PECL extension may be found in the manual chapter titled [Installation of PECL extensions](#). Additional information such as new releases, downloads, source files, maintainer information, and a CHANGELOG, can be located here: [http://pecl.php.net/package/mysqlnd\\_ms](http://pecl.php.net/package/mysqlnd_ms)

A DLL for this PECL extension is currently unavailable. See also the [building on Windows](#) section.

### 21.9.6.6.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in [php.ini](#).

**Table 21.74. Mysqlnd\_ms Configure Options**

Name	Default	Changeable	Changelog
<a href="#">mysqlnd_ms.enable</a>	0	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.force_config_usage</a>	0	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.ini_file</a>	""	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.config_file</a>	""	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.collect_statistics</a>	0	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.multi_master</a>	0	PHP_INI_SYSTEM	
<a href="#">mysqlnd_ms.disable_rw_split</a>	0	PHP_INI_SYSTEM	

Here's a short explanation of the configuration directives.

[mysqlnd\\_ms.enable](#) integer Enables or disables the plugin. If disabled, the extension will not plug into [mysqlnd](#) to proxy internal [mysqlnd](#) C API calls.

[mysqlnd\\_ms.force\\_config\\_usage](#) integer If enabled, the plugin checks if the host (server) parameters value of any MySQL connection attempt, matches a section name from the plugin configuration file. If not, the connection attempt is blocked.

This setting is not only useful to restrict PHP to certain servers but also to debug configuration file problems. The configuration file validity is checked at two different stages. The first check is performed when PHP begins to handle a web request. At this point the plugin reads and decodes the configuration file. Errors thrown

at this early stage in an extensions life cycle may not be shown properly to the user. Thus, the plugin buffers the errors, if any, and additionally displays them when establishing a connection to MySQL. By default a buffered startup error will emit an error of type `E_WARNING`. If `force_config_usage` is set, the error type used is `E_RECOVERABLE_ERROR`.

Please, see also [configuration file debugging notes \[3282\]](#).

`mysqlnd_ms.ini_file` string Plugin specific configuration file. This setting has been renamed to `mysqlnd_ms.config_file` in version 1.4.0.

`mysqlnd_ms.config_file` string Plugin specific configuration file. This setting supersedes `mysqlnd_ms.ini_file` since 1.4.0.

`mysqlnd_ms.collect_statistics` integer Enables or disables the collection of statistics. The collection of statistics is disabled by default for performance reasons. Statistics are returned by the function `mysqlnd_ms_get_stats`.

`mysqlnd_ms.multi_master` integer Enables or disables support of MySQL multi master replication setups. Please, see also [supported clusters](#).

`mysqlnd_ms.disable_rw_splitting` integer Enables or disables built-in read write splitting.

Controls whether load balancing and lazy connection functionality can be used independently of read write splitting. If read write splitting is disabled, only servers from the master list will be used for statement execution. All configured slave servers will be ignored.

The SQL hint `MYSQLND_MS_USE_SLAVE` will not be recognized. If found, the statement will be redirected to a master.

Disabling read write splitting impacts the return value of `mysqlnd_ms_query_is_select`. The function will no longer propose query execution on slave servers.



#### Multiple master servers

Setting `mysqlnd_ms.multi_master=1` allows the plugin to use multiple master servers, instead of only the first master server of the master list.

Please, see also [supported clusters](#).

#### 21.9.6.6.4. Plugin configuration file (>=1.1.x)

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Changelog: Feature was added in PECL/mysqlnd\_ms 1.1.0-beta

The below description applies to PECL/mysqlnd\_ms >= 1.1.0-beta. It is not valid for prior versions.

The plugin uses its own configuration file. The configuration file holds information about the MySQL replication master server, the MySQL replication slave servers, the server pick (load balancing) policy, the failover strategy, and the use of lazy connections.

The plugin loads its configuration file at the beginning of a web request. It is then cached in memory and used for the duration of the web request. This way, there is no need to restart PHP after deploying the configuration file. Configuration file changes will become active almost instantly.

The PHP configuration directive `mysqld_ms.config_file` is used to set the plugins configuration file. Please note, that the PHP configuration directive may not be evaluated for every web request. Therefore, changing the plugins configuration file name or location may require a PHP restart. However, no restart is required to read changes if an already existing plugin configuration file is updated.

Using and parsing JSON is efficient, and using JSON makes it easier to express hierarchical data structures than the standard `php.ini` format.

#### Example 21.264. Converting a PHP array (hash) into JSON format

Or alternatively, a developer may be more familiar with the PHP array syntax, and prefer it. This example demonstrates how a developer might convert a PHP array to JSON.

```
<?php
$config = array(
    "myapp" => array(
        "master" => array(
            "master_0" => array(
                "host" => "localhost",
                "socket" => "/tmp/mysql.sock",
            ),
        ),
        "slave" => array(),
    ),
);

file_put_contents("mysqld_ms.ini", json_encode($config, JSON_PRETTY_PRINT));
printf("mysqld_ms.ini file created...\n");
printf("Dumping file contents...\n");
printf("%s\n", str_repeat("-", 80));
echo file_get_contents("mysqld_ms.ini");
printf("\n%s\n", str_repeat("-", 80));
?>
```

The above example will output:

```
mysqld_ms.ini file created...
Dumping file contents...
-----
{
    "myapp": {
        "master": {
            "master_0": {
                "host": "localhost",
                "socket": "\\tmp\\mysql.sock"
            }
        },
        "slave": [
        ]
    }
}
-----
```

A plugin configuration file consists of one or more sections. Sections are represented by the top-level object properties of the object encoded in the JSON file. Sections could also be called *configuration names*.

Applications reference sections by their name. Applications use section names as the host (server) parameter to the various connect methods of the [mysqli](#), [mysql](#) and [PDO\\_MYSQL](#) extensions. Upon connect, the [mysqld](#) plugin compares the hostname with all of the section names from the plugin

configuration file. If the hostname and section name match, then the plugin will load the settings for that section.

### Example 21.265. Using section names example

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27"
      },
      "slave_1": {
        "host": "192.168.2.27",
        "port": 3306
      }
    }
  },
  "localhost": {
    "master": [
      {
        "host": "localhost",
        "socket": "\path\to\mysql.sock"
      }
    ],
    "slave": [
      {
        "host": "192.168.3.24",
        "port": "3305"
      },
      {
        "host": "192.168.3.65",
        "port": "3309"
      }
    ]
  }
}
```

```
<?php
/* All of the following connections will be load balanced */
$mysqli = new mysqli("myapp", "username", "password", "database");
$pdo = new PDO('mysql:host=myapp;dbname=database', 'username', 'password');
$mysql = mysql_connect("myapp", "username", "password");

$mysqli = new mysqli("localhost", "username", "password", "database");
?>
```

Section names are strings. It is valid to use a section name such as `192.168.2.1`, `127.0.0.1` or `localhost`. If, for example, an application connects to `localhost` and a plugin configuration section `localhost` exists, the semantics of the connect operation are changed. The application will no longer only use the MySQL server running on the host `localhost`, but the plugin will start to load balance MySQL queries following the rules from the `localhost` configuration section. This way you can load balance queries from an application without changing the applications source code. Please keep in mind, that such a configuration may not contribute to overall readability of your applications source code. Using section names that can be mixed up with host names should be seen as a last resort.

Each configuration section contains, at a minimum, a list of master servers and a list of slave servers. The master list is configured with the keyword `master`, while the slave list is configured with the `slave` keyword. Failing to provide a slave list will result in a fatal `E_ERROR` level error, although a slave list



may be empty. It is possible to allow no slaves. However, this is only recommended with synchronous clusters, please see also [supported clusters](#). The main part of the documentation focusses on the use of asynchronous MySQL replication clusters.

The master and slave server lists can be optionally indexed by symbolic names for the servers they describe. Alternatively, an array of descriptions for slave and master servers may be used.

#### Example 21.266. List of anonymous slaves

```
"slave": [
  {
    "host": "192.168.3.24",
    "port": "3305"
  },
  {
    "host": "192.168.3.65",
    "port": "3309"
  }
]
```

An anonymous server list is encoded by the [JSON array](#) type. Optionally, symbolic names may be used for indexing the slave or master servers of a server list, and done so using the [JSON object](#) type.

#### Example 21.267. Master list using symbolic names

```
"master": {
  "master_0": {
    "host": "localhost"
  }
}
```

It is recommended to index the server lists with symbolic server names. The alias names will be shown in error messages.

The order of servers is preserved and taken into account by `mysqld_ms`. If, for example, you configure round robin load balancing strategy, the first `SELECT` statement will be executed on the slave that appears first in the slave server list.

A configured server can be described with the `host`, `port`, `socket`, `db`, `user`, `password` and `connect_flags`. It is mandatory to set the database server host using the `host` keyword. All other settings are optional.

#### Example 21.268. Keywords to configure a server

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "db_server_host",
        "port": "db_server_port",
        "socket": "db_server_socket",
        "db": "database_resp_schema",
        "user": "user",
        "password": "password",
        "connect_flags": 0
      }
    },
    "slave": {
      "slave_0": {
        "host": "db_server_host",
```

```

        "port": "db_server_port",
        "socket": "db_server_socket"
    }
}
}
}

```

If a setting is omitted, the plugin will use the value provided by the user API call used to open a connection. Please, see the [using section names example \[3280\]](#) above.

The configuration file format has been changed in version 1.1.0-beta to allow for chained filters. Filters are responsible for filtering the configured list of servers to identify a server for execution of a given statement. Filters are configured with the `filter` keyword. Filters are executed by `mysqlnd_ms` in the order of their appearance. Defining filters is optional. A configuration section in the plugins configuration file does not need to have a `filters` entry.

Filters replace the `pick[]` setting from prior versions. The new `random` and `roundrobin` provide the same functionality.

**Example 21.269. New `roundrobin` filter, old functionality**

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      },
      "slave_1": {
        "host": "192.168.78.137",
        "port": "3306"
      }
    },
    "filters": {
      "roundrobin": [
      ]
    }
  }
}

```

The function `mysqlnd_ms_set_user_pick_server` has been removed. Setting a callback is now done with the `user` filter. Some filters accept parameters. The `user` filter requires and accepts a mandatory `callback` parameter to set the callback previously set through the function `mysqlnd_ms_set_user_pick_server`.

**Example 21.270. The `user` filter replaces `mysqlnd_ms_set_user_pick_server`**

```

"filters": {
  "user": {
    "callback": "pick_server"
  }
}

```

The validity of the configuration file is checked both when reading the configuration file and later when establishing a connection. The configuration file is read during PHP request startup. At this early stage

a PHP extension may not display error messages properly. In the worst case, no error is shown and a connection attempt fails without an adequate error message. This problem has been cured in version 1.5.0.

**Example 21.271. Common error message in case of configuration file issues (upto version 1.5.0)**

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
?>
```

The above example will output:

```
Warning: mysqli::mysqli(): (mysqlnd_ms) (mysqlnd_ms) Failed to parse config file [s1.json]. Please, ver
Warning: mysqli::mysqli(): (HY000/2002): php_network_getaddresses: getaddrinfo failed: Name or service
Warning: mysqli::query(): Couldn't fetch mysqli in Command line code on line 1
Fatal error: Call to a member function fetch_assoc() on a non-object in Command line code on line 1
```

Since version 1.5.0 startup errors are additionally buffered and emitted when a connection attempt is made. Use the configuration directive [mysqlnd\\_ms.force\\_config\\_usage](#) to set the error type used to display buffered errors. By default an error of type [E\\_WARNING](#) will be emitted.

**Example 21.272. Improved configuration file validation since 1.5.0**

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
?>
```

The above example will output:

```
Warning: mysqli::mysqli(): (mysqlnd_ms) (mysqlnd_ms) Failed to parse config file [s1.json]. Please, ver
```

It can be useful to set [mysqlnd\\_ms.force\\_config\\_usage = 1](#) when debugging potential configuration file errors. This will not only turn the type of buffered startup errors into [E\\_RECOVERABLE\\_ERROR](#) but also help detecting misspelled section names.

**Example 21.273. Possibly more precise error due to [mysqlnd\\_ms.force\\_config\\_usage=1](#)**

```
mysqlnd_ms.force_config_usage=1
```

```
<?php
$mysqli = new mysqli("invalid_section", "username", "password", "database");
?>
```

The above example will output:

```
Warning: mysqli::mysqli(): (mysqlnd_ms) Exclusive usage of configuration enforced but did not find the
```

Here is a short explanation of the configuration directives that can be used.

`master` array or object

List of MySQL replication master servers. The list of either of the `JSON type array` to declare an anonymous list of servers or of the `JSON type object`. Please, see [above \[3280\]](#) for examples.

Setting at least one master server is mandatory. The plugin will issue an error of type `E_ERROR` if the user has failed to provide a master server list for a configuration section. The fatal error may read `(mysqlnd_ms) Section [master] doesn't exist for host [name_of_a_config_section] in %s on line %d`.

A server is described with the `host`, `port`, `socket`, `db`, `user`, `password` and `connect_flags`. It is mandatory to provide at a value for `host`. If any of the other values is not given, it will be taken from the user API connect call, please, see also: [using section names example \[3280\]](#).

Table of server configuration keywords.

Keyword	Description	Version
<code>host</code>	Database server host. This is a mandatory setting. Failing to provide, will cause an error of type <code>E_RECOVERABLE_ERROR</code> when the plugin tries to connect to the server. The error message may read <code>(mysqlnd_ms) Cannot find [host] in [%s] section in config in %s on line %d</code> .	Since 1.1.0.
<code>port</code>	Database server TCP/IP port.	Since 1.1.0.
<code>socket</code>	Database server Unix domain socket.	Since 1.1.0.
<code>db</code>	Database (schemata).	Since 1.1.0.
<code>user</code>	MySQL database user.	Since 1.1.0.
<code>password</code>	MySQL database user password.	Since 1.1.0.
<code>connect_flags</code>	Connection flags.	Since 1.1.0.

The plugin supports using only one master server. An experimental setting exists to enable multi-master support. The details are not documented. The setting is meant for development only.

`slave` array or object

List of one or more MySQL replication slave servers. The syntax is identical to setting master servers, please, see `master` above for details.

The plugin supports using one or more slave servers.

Setting a list of slave servers is mandatory. The plugin will report an error of the type `E_ERROR` if `slave` is not given for a configuration section. The fatal error message may read `(mysqlnd_ms) Section [slave] doesn't exist for host [%s] in %s on line %d`. Note, that it is valid to use an empty slave server list. The error has been introduced to prevent accidentally setting no slaves by forgetting about the `slave` setting. A master-only setup is still possible using an empty slave server list.

If an empty slave list is configured and an attempt is made to execute a statement on a slave the plugin may emit a warning like `(mysqld_ms) Couldn't find the appropriate slave connection. 0 slaves to choose from.` upon statement execution. It is possible that another warning follows such as `(mysqld_ms) No connection selected by the last filter.`

`global_transaction_id_injection` Global transaction identifier configuration related to both the use of the server built-in global transaction ID feature and the client-side emulation.

Keyword	Description	Version
<code>fetch_global_transaction_id</code>	SQL statement for accessing the latest global transaction identifier. The SQL statement is run if the plugin needs to know the most recent global transaction identifier. This can be the case, for example, when checking MySQL Replication slave status. Also used with <code>mysqld_ms_get_last_gtid</code> .	Since 1.2.0.
<code>check_replica_consistency</code>	SQL statement for checking if a replica has replicated all transactions up to and including ones searched for. The SQL statement is run when searching for replicas which can offer a higher level of consistency than eventual consistency. The statement must contain a placeholder <code>#GTID</code> which is to be replaced with the global transaction identifier searched for by the plugin. Please, check the <a href="#">quickstart</a> for examples.	Since 1.2.0.
<code>report_errors</code>	Whether to emit an error of type warning if an issue occurs while executing any of the configured SQL statements.	Since 1.2.0.
<code>on_completion</code>	Client-side global transaction ID emulation only. SQL statement to run when a transaction finished to update the global transaction identifier sequence number on the master. Please, see the <a href="#">quickstart</a> for examples.	Since 1.2.0.
<code>wait_for_slave</code>	Instructs the plugin to wait up to <code>wait_for_gtid_timeout</code> seconds for a slave to catch up when searching for slaves that can deliver session consistency. The setting limits the time spend for polling the slave status. If polling the status takes very long, the total clock time spend waiting may exceed <code>wait_for_gtid_timeout</code> . The plugin calls <code>sleep(1)</code> to sleep one second between each two polls.  The setting can be used both with the plugins client-side emulation and the server-side global transaction identifier feature of MySQL 5.6.  Waiting for a slave to replicate a certain GTID needed for session consistency also means	Since 1.4.0.

Keyword	Description	Version
	throttling the client. By throttling the client the write load on the master is reduced indirectly. A primary copy based replication system, such as MySQL Replication, is given more time to reach a consistent state. This can be desired, for example, to increase the number of data copies for high availability considerations or to prevent the master from being overloaded.	

`filters` object

List of filters. A filter is responsible to filter the list of available servers for executing a given statement. Filters can be chained. The `random` and `roundrobin` filter replace the `pick[ ]` directive used in prior version to select a load balancing policy. The `user` filter replaces the `mysqlnd_ms_set_user_pick_server` function.

Filters may accept parameters to refine their actions.

If no load balancing policy is set, the plugin will default to `random_once`. The `random_once` policy picks a random slave server when running the first read-only statement. The slave server will be used for all read-only statements until the PHP script execution ends. No load balancing policy is set and thus, defaulting takes place, if neither the `random` nor the `roundrobin` are part of a configuration section.

If a filter chain is configured so that a filter which output no more than once server is used as input for a filter which should be given more than one server as input, the plugin may emit a warning upon opening a connection. The warning may read: (`mysqlnd_ms`) `Error while creating filter '%s' . Non-multi filter '%s' already created. Stopping in %s on line %d`. Furthermore, an error of the error code `2000`, the sql state `HY000` and an error message similar to the warning may be set on the connection handle.

**Example 21.274. Invalid filter sequence**

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "filters": [
      "roundrobin",
      "random"
    ]
  }
}

```

```
<?php
```

```
$link = new mysqli("myapp", "root", "", "test");
printf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error());
$link->query("SELECT 1 FROM DUAL");
?>
```

The above example will output:

```
PHP Warning: mysqli::mysqli(): (HY000/2000): (mysqlnd_ms) Error while
[2000] (mysqlnd_ms) Error while creating filter 'random' . Non-multi fi
PHP Warning: mysqli::query(): Couldn't fetch mysqli in filter_warning.
```

Filter: `random` object

The `random` filter features the `random` and `random once` load balancing policies, set through the `pick[]` directive in older versions.

The `random` policy will pick a random server whenever a read-only statement is to be executed. The `random once` strategy picks a random slave server once and continues using the slave for the rest of the PHP web request. `Random once` is a default, if load balancing is not configured through a filter.

If the `random` filter is not given any arguments, it stands for random load balancing policy.

#### Example 21.275. Random load balancing with `random` filter

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      },
      "slave_1": {
        "host": "192.168.78.137",
        "port": "3306"
      }
    }
  },
  "filters": [
    "random"
  ]
}
```

Optionally, the `sticky` argument can be passed to the filter. If the parameter `sticky` is set to the string `1`, the filter follows the `random once` load balancing strategy.

#### Example 21.276. Random once load balancing with `random` filter

```
{
  "filters": {
```

```

        "random": {
            "sticky": "1"
        }
    }
}

```

Both the `random` and `roundrobin` filters support setting a priority, a weight for a server, since PECL/mysqlnd\_ms 1.4.0. If the `weight` argument is passed to the filter, it must assign a weight for all servers. Servers must be given an alias name in the `slave` respectively `master` server lists. The alias must be used to reference servers for assigning a priority with `weight`.

#### Example 21.277. Referencing error

```
[E_RECOVERABLE_ERROR] mysqli_real_connect(): (mysqlnd_ms) Unknown server 's
```

Using a wrong alias name with `weight` may result in an error similar to the shown above.

If `weight` is omitted, the default weight of all servers is one.

#### Example 21.278. Assigning a `weight` for load balancing

```

{
    "myapp": {
        "master": {
            "master1": {
                "host": "localhost",
                "socket": "\/var\/run\/mysql\/mysql.sock"
            }
        },
        "slave": {
            "slave1": {
                "host": "192.168.2.28",
                "port": 3306
            },
            "slave2": {
                "host": "192.168.2.29",
                "port": 3306
            },
            "slave3": {
                "host": "192.0.43.10",
                "port": 3306
            }
        },
        "filters": {
            "random": {
                "weights": {
                    "slave1": 8,
                    "slave2": 4,
                    "slave3": 1,
                    "master1": 1
                }
            }
        }
    }
}

```

At the average a server assigned a weight of two will be selected twice as often as a server assigned a weight of one. Different



weights can be assigned to reflect differently sized machines, to prefer co-located slaves which have a low network latency or, to configure a standby failover server. In the latter case, you may want to assign the standby server a very low weight in relation to the other servers. For example, given the configuration above `slave3` will get only some eight percent of the requests in the average. As long as `slave1` and `slave2` are running, it will be used sparsely, similar to a standby failover server. Upon failure of `slave1` and `slave2`, the usage of `slave3` increases. Please, check the notes on failover before using `weight` this way.

Valid weight values range from 1 to 65535.

Unknown arguments are ignored by the filter. No warning or error is given.

The filter expects one or more servers as input. Outputs one server. A filter sequence such as `random`, `roundrobin` may cause a warning and an error message to be set on the connection handle when executing a statement.

List of filter arguments.

Keyword	Description	Version
<code>sticky</code>	Enables or disabled random once load balancing policy. See above.	Since 1.2.0.
<code>weight</code>	Assigns a load balancing weight/priority to a server. Please, see above for a description.	Since 1.4.0.

Filter: `roundrobin` object

If using the `roundrobin` filter, the plugin iterates over the list of configured slave servers to pick a server for statement execution. If the plugin reaches the end of the list, it wraps around to the beginning of the list and picks the first configured slave server.

**Example 21.279. `roundrobin` filter**

```

{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "filters": [
      "roundrobin"
    ]
  }
}

```

Expects one or more servers as input. Outputs one server. A filter sequence such as `roundrobin`, `random` may cause a warning and an error message to be set on the connection handle when executing a statement.

List of filter arguments.

Keyword	Description	Version
<code>weight</code>	Assigns a load balancing weight/priority to a server. Please, find a description <a href="#">above</a> .	Since 1.4.0.

Filter: `user` object

The `user` replaces `mysqld_ms_set_user_pick_server` function, which was removed in 1.1.0-beta. The filter sets a callback for user-defined read/write splitting and server selection.

The plugins built-in read/write query split mechanism decisions can be overwritten in two ways. The easiest way is to prepend a query string with the SQL hints `MYSQLND_MS_MASTER_SWITCH`, `MYSQLND_MS_SLAVE_SWITCH` or `MYSQLND_MS_LAST_USED_SWITCH`. Using SQL hints one can control, for example, whether a query shall be send to the MySQL replication master server or one of the slave servers. By help of SQL hints it is not possible to pick a certain slave server for query execution.

Full control on server selection can be gained using a callback function. Use of a callback is recommended to expert users only because the callback has to cover all cases otherwise handled by the plugin.

The plugin will invoke the callback function for selecting a server from the lists of configured master and slave servers. The callback function inspects the query to run and picks a server for query execution by returning the hosts URI, as found in the master and slave list.

If the lazy connections are enabled and the callback chooses a slave server for which no connection has been established so far and establishing the connection to the slave fails, the plugin will return an error upon the next action on the failed connection, for example, when running a query. It is the responsibility of the application developer to handle the error. For example, the application can re-run the query to trigger a new server selection and callback invocation. If so, the callback must make sure to select a different slave, or check slave availability, before returning to the plugin to prevent an endless loop.

#### Example 21.280. Setting a callback

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "filters": {
      "user": {
        "callback": "pick_server"
      }
    }
  }
}
```

```

    }
  }
}

```

The callback is supposed to return a host to run the query on. The host URI is to be taken from the master and slave connection lists passed to the callback function. If callback returns a value neither found in the master nor in the slave connection lists the plugin will emit an error of the type `E_RECOVERABLE_ERROR`. The error may read like `(mysqld_ms) User filter callback has returned an unknown server. The server 'server that is not in master or slave list' can neither be found in the master list nor in the slave list.` If the application catches the error to ignore it, follow up errors may be set on the connection handle, for example, `(mysqld_ms) No connection selected by the last filter with the error code 2000 and the sqlstate HY000.` Furthermore a warning may be emitted.

Referencing a non-existing function as a callback will result in any error of the type `E_RECOVERABLE_ERROR` whenever the plugin tries to callback function. The error message may reads like: `(mysqld_ms) Specified callback (pick_server) is not a valid callback.` If the application catches the error to ignore it, follow up errors may be set on the connection handle, for example, `(mysqld_ms) Specified callback (pick_server) is not a valid callback with the error code 2000 and the sqlstate HY000.` Furthermore a warning may be emitted.

The following parameters are passed from the plugin to the callback.

Parameter	Description	Version
<code>connection</code>	URI of the currently connected database server.	Since 1.1.0.
<code>query</code>	Query string of the statement for which a server needs to be picked.	Since 1.1.0.
<code>master</code>	List of master servers to choose from. Note, that the list of master servers may not be identical to the list of configured master servers if the filter is not the first in the filter chain. Previously run filters may have reduced the master list already.	Since 1.1.0.
<code>slave</code>	List of slave servers to choose from. Note, that the list of master servers may not be identical to the list of configured master servers if the filter is not the first in the filter chain. Previously run filters may have reduced the master list already.	Since 1.1.0.
<code>last_server</code>	URI of the server of the connection used to execute the previous statement on.	Since 1.1.0.
<code>in_transaction</code>	Boolean flag indicating whether the statement is part of an open transaction. If autocommit	Since 1.1.0.

Parameter	Description	Version
	<p>mode is turned off, this will be set to <code>TRUE</code> . Otherwise it is set to <code>FALSE</code> .</p> <p>Transaction detection is based on monitoring the <code>mysqlnd</code> library call <code>set_autocommit</code>. Monitoring is not possible before PHP 5.4.0. Please, see <a href="#">connection pooling and switching</a> concepts discussion for further details.</p>	

**Example 21.281. Using a callback**

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.2.27",
        "port": "3306"
      },
      "slave_1": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "filters": {
      "user": {
        "callback": "pick_server"
      }
    }
  }
}
```

```
<?php
function pick_server($connected, $query, $masters, $slaves, $last_used_conn)
{
  static $slave_idx = 0;
  static $num_slaves = NULL;
  if (is_null($num_slaves))
    $num_slaves = count($slaves);

  /* default: fallback to the plugins build-in logic */
  $ret = NULL;

  printf("User has connected to '%s'...\n", $connected);
  printf("... deciding where to run '%s'\n", $query);

  $where = mysqlnd_ms_query_is_select($query);
  switch ($where)
  {
    case MYSQLND_MS_QUERY_USE_MASTER:
      printf("... using master\n");
      $ret = $masters[0];
      break;
    case MYSQLND_MS_QUERY_USE_SLAVE:
      /* SELECT or SQL hint for using slave */
      if (strstr($query, "FROM table_on_slave_a_only"))
      {
        /* a table which is only on the first configured slave */
        printf("... access to table available only on slave A detected\n");
        $ret = $slaves[0];
      }
    }
  }
}
```

```

    }
    else
    {
        /* round robin */
        printf("... some read-only query for a slave\n");
        $ret = $slaves[$slave_idx++ % $num_slaves];
    }
    break;
case MYSQLND_MS_QUERY_LAST_USED:
    printf("... using last used server\n");
    $ret = $last_used_connection;
    break;
}

printf("... ret = '%s'\n", $ret);
return $ret;
}

$mysqli = new mysqli("myapp", "root", "", "test");

if (!$res = $mysqli->query("SELECT 1 FROM DUAL"))
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
    $res->close();

if (!$res = $mysqli->query("SELECT 2 FROM DUAL"))
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
    $res->close();

if (!$res = $mysqli->query("SELECT * FROM table_on_slave_a_only"))
    printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
    $res->close();

$mysqli->close();
?>

```

The above example will output:

```

User has connected to 'myapp'...
... deciding where to run 'SELECT 1 FROM DUAL'
... some read-only query for a slave
... ret = 'tcp://192.168.2.27:3306'
User has connected to 'myapp'...
... deciding where to run 'SELECT 2 FROM DUAL'
... some read-only query for a slave
... ret = 'tcp://192.168.78.136:3306'
User has connected to 'myapp'...
... deciding where to run 'SELECT * FROM table_on_slave_a_only'
... access to table available only on slave A detected
... ret = 'tcp://192.168.2.27:3306'

```

Filter: [user\\_multi](#) object

The [user\\_multi](#) differs from the [user](#) only in one aspect. Otherwise, their syntax is identical. The [user](#) filter must pick and return exactly one node for statement execution. A filter chain usually ends with a filter that emits only one node. The filter chain shall reduce the list of candidates for statement execution down to one. This, only one node left, is the case after the [user](#) filter has been run.

The [user\\_multi](#) filter is a multi filter. It returns a list of slave and a list of master servers. This list needs further filtering to identify

exactly one node for statement execution. A multi filter is typically placed at the top of the filter chain. The `quality_of_service` filter is another example of a multi filter.

The return value of the callback set for `user_multi` must be an array with two elements. The first element holds a list of selected master servers. The second element contains a list of selected slave servers. The lists shall contain the keys of the slave and master servers as found in the slave and master lists passed to the callback. The below example returns random master and slave lists extracted from the functions input.

### Example 21.282. Returning random masters and slaves

```
<?php
function pick_server($connected, $query, $masters, $slaves, $last_used_conn)
{
    $picked_masters = array()
    foreach ($masters as $key => $value) {
        if (mt_rand(0, 2) > 1)
            $picked_masters[] = $key;
    }
    $picked_slaves = array()
    foreach ($slaves as $key => $value) {
        if (mt_rand(0, 2) > 1)
            $picked_slaves[] = $key;
    }
    return array($picked_masters, $picked_slaves);
}
?>
```

The plugin will issue an error of type `E_RECOVERABLE` if the callback fails to return a server list. The error may read `(mysqlnd_ms) User multi filter callback has not returned a list of servers to use. The callback must return an array in %s on line %d`. In case the server list is not empty but has invalid servers key/ids in it, an error of type `E_RECOVERABLE` will be thrown with an error message like `(mysqlnd_ms) User multi filter callback has returned an invalid list of servers to use. Server id is negative in %s on line %d, or similar`.

Whether an error is emitted in case of an empty slave or master list depends on the configuration. If an empty master list is returned for a write operation, it is likely that the plugin will emit a warning that may read `(mysqlnd_ms) Couldn't find the appropriate master connection. 0 masters to choose from. Something is wrong in %s on line %d`. Typically a follow up error of type `E_ERROR` will happen. In case of a read operation and an empty slave list the behavior depends on the fail over configuration. If fail over to master is enabled, no error should appear. If fail over to master is deactivated the plugin will emit a warning that may read `(mysqlnd_ms) Couldn't find the appropriate slave connection. 0 slaves to choose from. Something is wrong in %s on line %d`.

Filter: `node_groups` object

The `node_groups` filter lets you group cluster nodes and query selected groups, for example, to support data partitioning. Data partitioning can be required for manual sharding, primary copy based clusters running multiple masters, or to avoid hot spots in

update everywhere clusters that have no built-in partitioning. The filter is a multi filter which returns zero, one or multiple of its input servers. Thus, it must be followed by other filters to reduce the number of candidates down to one for statement execution.

Keyword	Description	Version
<code>user defined node group name</code>	<p>One or more node groups must be defined. A node group can have an arbitrary user defined name. The name is used in combination with a SQL hint to restrict query execution to the nodes listed for the node group. To run a query on any of the servers of a node group, the query must begin with the SQL hint <code>/*user defined node group name*/</code>. Please note, no white space is allowed around <code>user defined node group name</code>. Because <code>user defined node group name</code> is used as-is as part of a SQL hint, you should choose the name that is compliant with the SQL language.</p> <p>Each node group entry must contain a list of <code>master</code> servers. Additional <code>slave</code> servers are allowed. Failing to provide a list of <code>master</code> for a node group <code>name_of_group</code> may cause an error of type <code>E_RECOVERABLE_ERROR</code> like <code>(mysqld_ms) No masters configured in node group 'name_of_group' for 'node_groups' filter</code>.</p> <p>The list of master and slave servers must reference corresponding entries in the <code>global master</code> respectively <code>slave</code> server list. Referencing an unknown server in either of the both server lists may cause an <code>E_RECOVERABLE_ERROR</code> error like <code>(mysqld_ms) Unknown master 'server_alias_name' (section 'name_of_group') in 'node_groups' filter configuration</code>.</p> <p><b>Example 21.283. Manual partitioning</b></p> <pre> {   "myapp": {     "master": {       "master_0": {         "host": "localhost",         "socket": "\\tmp\\mysql.sock"       },     },     "slave": {       "slave_0": {         "host": "192.168.2.28",         "port": 3306       },       "slave_1": {         "host": "127.0.0.1",         "port": 3311       }     }   }, } </pre>	Since 1.5.0.

Keyword	Description	Version
	<pre> filters: {   node_groups: {     Partition_A: {       master: ["master_0"],       slave: ["slave_0"]     },     roundrobin: []   } } </pre> <p>Please note, if a filter chain generates an empty slave list and the PHP configuration directive <code>mysqlnd_ms.multi_master=0</code> is used, the plugin may emit a warning.</p>	

Filter: `quality_of_service` object

The `quality_of_service` identifies cluster nodes capable of delivering a certain quality of service. It is a multi filter which returns zero, one or multiple of its input servers. Thus, it must be followed by other filters to reduce the number of candidates down to one for statement execution.

The `quality_of_service` filter has been introduced in 1.2.0-alpha. In the 1.2 series the filters focus is on the consistency aspect of service quality. Different types of clusters offer different default data consistencies. For example, an asynchronous MySQL replication slave offers eventual consistency. The slave may not be able to deliver requested data because it has not replicated the write, it may serve stale database because its lagging behind or it may serve current information. Often, this is acceptable. In some cases higher consistency levels are needed for the application to work correct. In those cases, the `quality_of_service` can filter out cluster nodes which cannot deliver the necessary quality of service.

The `quality_of_service` filter can be replaced or created at runtime. A successful call to `mysqlnd_ms_set_qos` removes all existing `qos` filter entries from the filter list and installs a new one at the very beginning. All settings that can be made through `mysqlnd_ms_set_qos` can also be in the plugins configuration file. However, use of the function is by far the most common use case. Instead of setting session consistency and strong consistency service levels in the plugins configuration file it is recommended to define only masters and no slaves. Both service levels will force the use of masters only. Using an empty slave list shortens the configuration file, thus improving readability. The only service level for which there is a case of defining in the plugins configuration file is the combination of eventual consistency and maximum slave lag.

Keyword	Description	Version
<code>eventual</code>	<p>Request eventual consistency. Allows the use of all master and slave servers. Data returned may or may not be current.</p> <p>Eventual consistency accepts an optional <code>age</code> parameter. If <code>age</code> is given the plugin</p>	Since 1.2.0.



Keyword	Description	Version
	<p>considers only slaves for reading for which MySQL replication reports a slave lag less or equal to <code>age</code>. The replication lag is measure using <code>SHOW SLAVE STATUS</code>. If the plugin fails to fetch the replication lag, the slave tested is skipped. Implementation details and tips are given in the <a href="#">quality of service concepts section</a>.</p> <p>Please note, if a filter chain generates an empty slave list and the PHP configuration directive <code>mysqLnd_ms.multi_master=0</code> is used, the plugin may emit a warning.</p> <p><b>Example 21.284. Global limit on slave lag</b></p> <pre data-bbox="767 725 1332 1473"> {   "myapp": {     "master": {       "master_0": {         "host": "localhost"       }     },     "slave": {       "slave_0": {         "host": "192.168.2.27",         "port": "3306"       },       "slave_1": {         "host": "192.168.78.136",         "port": "3306"       }     },     "filters": {       "quality_of_service": {         "eventual_consistency": {           "age": 123         }       }     }   } } </pre>	
<a href="#">session consistency</a>	<p><b>Request session consistency</b> (read your writes). Allows use of all masters and all slaves which are in sync with the master. If no further parameters are given slaves are filtered out as there is no reliable way to test if a slave has caught up to the master or is lagging behind. Please note, if a filter chain generates an empty slave list and the PHP configuration directive <code>mysqLnd_ms.multi_master=0</code> is used, the plugin may emit a warning.</p> <p>Session consistency temporarily requested using <code>mysqLnd_ms_set_qos</code> is a valuable alternative to using <code>master_on_write</code>. <code>master_on_write</code> is likely to send more statements to the master than needed. The application may be able to continue operation</p>	<p>Since 1.1.0.</p>

Keyword	Description	Version
	at a lower consistency level after it has done some critical reads.	
<code>strong</code>	Requests strong consistency. Only masters will be used.	Since 1.2.0.

`failover` Up to and including 1.3.x: string. Since 1.4.0: object.

Failover policy. Supported policies: `disabled` (default), `master`, `loop_before_master` (Since 1.4.0).

If no failover policy is set, the plugin will not do any automatic failover (`failover=disabled`). Whenever the plugin fails to connect a server it will emit a warning and set the connections error code and message. Thereafter it is up to the application to handle the error and, for example, resend the last statement to trigger the selection of another server.

Please note, the automatic failover logic is applied when opening connections only. Once a connection has been opened no automatic attempts are made to reopen it in case of an error. If, for example, the server a connection is connected to is shut down and the user attempts to run a statement on the connection, no automatic failover will be tried. Instead, an error will be reported.

If using `failover=master` the plugin will implicitly failover to a master, if available. Please check the concepts documentation to learn about potential pitfalls and risks of using `failover=master`.

**Example 21.285. Optional master failover when failing to connect to slave (PECL/mysqld\_ms < 1.4.0)**

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "failover": "master"
  }
}
```

Since PECL/mysqld\_ms 1.4.0 the failover configuration keyword refers to an object.

**Example 21.286. New syntax since 1.4.0**

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
```

```

        "slave_0": {
            "host": "192.168.78.136",
            "port": "3306"
        },
        "failover": {"strategy": "master" }
    }
}

```

Keyword	Description	Version
<code>strategy</code>	<p>Failover policy. Possible values: <code>disabled</code> (default), <code>master</code>, <code>loop_before_master</code></p> <p>A value of <code>disabled</code> disables automatic failover.</p> <p>Setting <code>master</code> instructs the plugin to try to connect to a master in case of a slave connection error. If the master connection attempt fails, the plugin enters the failover loop and returns an error to the user.</p> <p>If using <code>loop_before_master</code> and a slave request is made, the plugin tries to connect to other slaves before failing over to a master. If multiple masters are given and multi master is enabled, the plugin also loops over the list of masters and attempts to connect before returning an error to the user.</p>	Since 1.4.0.
<code>remember_failed</code>	<p>Remember failures for the duration of a web request. Default: <code>false</code>.</p> <p>If set to <code>true</code> the plugin will remember failed hosts and skip the hosts in all future load balancing made for the duration of the current web request.</p>	Since 1.4.0. The feature is only available together with the <code>random</code> and <code>roundrobin</code> load balancing filter. Use of the setting is recommended.
<code>max_retries</code>	<p>Maximum number of connection attempts before skipping host. Default: <code>0</code> (no limit).</p> <p>The setting is used to prevent hosts from being dropped of the host list upon the first failure. If set to <code>n &gt; 0</code>, the plugin will keep the node in the node list even after a failed connection attempt. The node will not be removed immediately from the slave respectively master lists after the first connection failure but instead be tried to connect to up to <code>n</code> times in future load balancing rounds before being removed.</p>	Since 1.4.0. The feature is only available together with the <code>random</code> and <code>roundrobin</code> load balancing filter.

[lazy\\_connections](#) bool

Setting [failover](#) to any other value but [disabled](#), [master](#) or [loop\\_before\\_master](#) will not emit any warning or error.

Controls the use of lazy connections. Lazy connections are connections which are not opened before the client sends the first connection. Lazy connections are a default.

It is strongly recommended to use lazy connections. Lazy connections help to keep the number of open connections low. If you disable lazy connections and, for example, configure one MySQL replication master server and two MySQL replication slaves, the plugin will open three connections upon the first call to a connect function although the application might use the master connection only.

Lazy connections bare a risk if you make heavy use of actions which change the state of a connection. The plugin does not dispatch all state changing actions to all connections from the connection pool. The few dispatched actions are applied to already opened connections only. Lazy connections opened in the future are not affected. Only some settings are "remembered" and applied when lazy connections are opened.

#### Example 21.287. Disabling lazy connection

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "lazy_connections": 0
  }
}
```

Please, see also [server\\_charset](#) to overcome potential problems with string escaping and servers using different default charsets.

[server\\_charset](#) string

The setting has been introduced in 1.4.0. It is recommended to set it if using lazy connections.

The [server\\_charset](#) setting serves two purposes. It acts as a fallback charset to be used for string escaping done before a connection has been established and it helps to avoid escaping pitfalls in heterogeneous environments which servers using different default charsets.

String escaping takes a connections charset into account. String escaping is not possible before a connection has been opened and the connections charset is known. The use of lazy connections delays the actual opening of connections until a statement is send.

An application using lazy connections may attempt to escape a string before sending a statement. In fact, this should be a common case as the statement string may contain the string that is to be escaped. However, due to the lazy connection feature no connection has been opened yet and escaping fails. The plugin may report an error of the type `E_WARNING` and a message like `(mysqlnd_ms) string escaping doesn't work without established connection. Possible solution is to add server_charset to your configuration` to inform you of the pitfall.

Setting `server_charset` makes the plugin use the given charset for string escaping done on lazy connection handles before establishing a network connection to MySQL. Furthermore, the plugin will enforce the use of the charset when the connection is established.

Enforcing the use of the configured charset used for escaping is done to prevent tapping into the pitfall of using a different charset for escaping than used later for the connection. This has the additional benefit of removing the need to align the charset configuration of all servers used. No matter what the default charset on any of the servers is, the plugin will set the configured one as a default.

The plugin does not stop the user from changing the charset at any time using the `set_charset` call or corresponding SQL statements. Please, note that the use of SQL is not recommended as it cannot be monitored by the plugin. The user can, for example, change the charset on a lazy connection handle after escaping a string and before the actual connection is opened. The charset set by the user will be used for any subsequent escaping before the connection is established. The connection will be established using the configured charset, no matter what the server charset is or what the user has set before. Once a connection has been opened, `set_charset` is of no meaning anymore.

#### Example 21.288. String escaping on a lazy connection handle

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "lazy_connections": 1,
    "server_charset" : "utf8"
  }
}
```

```
<?php
$mysqli = new mysqli("myapp", "username", "password", "database");
$mysqli->real_escape("this will be escaped using the server_charset set
```

```
$mysqlcli->set_charset("latin1");
$mysqlcli->real_escape("this will be escaped using latin1");
/* server_charset implicitly set - utf8 connection */
$mysqlcli->query("SELECT 'This connection will be set to server_charset upon
/* latin1 used from now on */
$mysqlcli->set_charset("latin1");
?>
```

`master_on_write` bool

If set, the plugin will use the master server only after the first statement has been executed on the master. Applications can still send statements to the slaves using SQL hints to overrule the automatic decision.

The setting may help with replication lag. If an application runs an `INSERT` the plugin will, by default, use the master to execute all following statements, including `SELECT` statements. This helps to avoid problems with reads from slaves which have not replicated the `INSERT` yet.

#### Example 21.289. Master on write for consistent reads

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "master_on_write": 1
  }
}
```

Please, note the `quality_of_service` filter introduced in version 1.2.0-alpha. It gives finer control, for example, for achieving read-your-writes and, it offers additional functionality introducing `service levels`.

All `transaction stickiness` settings, including `trx_stickiness=on`, are overruled by `master_on_write=1`.

`trx_stickiness` string

Transaction stickiness policy. Supported policies: `disabled` (default), `master`.

The setting requires 5.4.0 or newer. If used with PHP older than 5.4.0, the plugin will emit a warning like `(mysqlnd_ms) trx_stickiness strategy is not supported before PHP 5.3.99`.

If no transaction stickiness policy is set or, if setting `trx_stickiness=disabled`, the plugin is not transaction aware. Thus, the plugin may load balance connections and switch connections in the middle of a transaction. The plugin is not transaction safe. SQL hints must be used avoid connection switches during a transaction.

As of PHP 5.4.0 the `mysqlnd` library allows the plugin to monitor the `autocommit` mode set by calls to the libraries `set_autocommit()` function. If setting `set_stickiness=master` and `autocommit` gets disabled by a PHP MySQL extension invoking the `mysqlnd` library internal function call `set_autocommit()`, the plugin is made aware of the begin of a transaction. Then, the plugin stops load balancing and directs all statements to the master server until `autocommit` is enabled. Thus, no SQL hints are required.

An example of a PHP MySQL API function calling the `mysqlnd` library internal function call `set_autocommit()` is `mysqli_autocommit`.

Although setting `trx_stickiness=master`, the plugin cannot be made aware of `autocommit` mode changes caused by SQL statements such as `SET AUTOCOMMIT=0` or `BEGIN`.

As of PHP 5.5.0, the `mysqlnd` library features additional C API calls to control transactions. The level of control matches the one offered by SQL statements. The `mysqli` API has been modified to use these calls. Since version 1.5.0, PECL/`mysqlnd_ms` can monitor not only `mysqli_autocommit`, but also `mysqli_begin`, `mysqli_commit` and `mysqli_rollback` to detect transaction boundaries and stop load balancing for the duration of a transaction.

#### Example 21.290. Using master to execute transactions

```
{
  "myapp": {
    "master": {
      "master_0": {
        "host": "localhost"
      }
    },
    "slave": {
      "slave_0": {
        "host": "192.168.78.136",
        "port": "3306"
      }
    },
    "trx_stickiness": "master"
  }
}
```

Since version 1.5.0 automatic and silent failover is disabled for the duration of a transaction. If the boundaries of a transaction have been properly detected, transaction stickiness is enabled and a server fails, the plugin will not attempt to fail over to the next server, if any, regardless of the failover policy configured. The user must handle the error manually. Depending on the configuration, the plugin may emit an error of type `E_WARNING` reading like `(mysqlnd_ms) Automatic failover is not permitted in the middle of a transaction`. This error may then be overwritten by follow up errors such as `(mysqlnd_ms) No connection selected by the last filter`. Those errors will be generated by the failing query function.

### Example 21.291. No automatic failover, error handling pitfall

```
<?php
/* assumption: automatic failover configured */
$mysqli = new mysqli("myapp", "username", "password", "database");

/* sets plugin internal state in_trx = 1 */
$mysqli->autocommit(false);

/* assumption: server fails */
if (!$res = $mysqli->query("SELECT 'Assume this query fails' AS _msg FROM
/* handle failure of transaction, plugin internal state is still in_trx =
printf("[%d] %s", $mysqli->errno, $mysqli->error);
/*
If using autocommit() based transaction detection it is a
MUST to call autocommit(true). Otherwise the plugin assumes
the current transaction continues and connection
changes remain forbidden.
*/
$mysqli->autocommit(true);
/* Likewise, you'll want to start a new transaction */
$mysqli->autocommit(false);
}
/* latin1 used from now on */
$mysqli->set_charset("latin1");
?>
```

If a server fails in the middle of a transaction the plugin continues to refuse to switch connections until the current transaction has been finished. Recall that the plugin monitors API calls to detect transaction boundaries. Thus, you have to, for example, enable auto commit mode to end the current transaction before the plugin continues load balancing and switches the server. Likewise, you will want to start a new transaction immediately thereafter and disable auto commit mode again.

Not handling failed queries and not ending a failed transaction using API calls may cause all following commands emit errors such as [Commands out of sync; you can't run this command now](#). Thus, it is important to handle all errors.

#### 21.9.6.6.5. Plugin configuration file (<= 1.0.x)

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Note

The below description applies to PECL/mysqld\_ms < 1.1.0-beta. It is not valid for later versions.

The plugin is using its own configuration file. The configuration file holds information on the MySQL replication master server, the MySQL replication slave servers, the server pick (load balancing) policy, the failover strategy and the use of lazy connections.

The PHP configuration directive [mysqld\\_ms.ini\\_file](#) is used to set the plugins configuration file.

The configuration file mimics standard the [php.ini](#) format. It consists of one or more sections. Every section defines its own unit of settings. There is no global section for setting defaults.

Applications reference sections by their name. Applications use section names as the host (server) parameter to the various connect methods of the [mysqli](#), [mysql](#) and [PDO\\_MYSQL](#) extensions.



Upon connect the `mysqld` plugin compares the hostname with all section names from the plugin configuration file. If hostname and section name match, the plugin will load the sections settings.

### Example 21.292. Using section names example

```
[myapp]
master[] = localhost
slave[] = 192.168.2.27
slave[] = 192.168.2.28:3306
[localhost]
master[] = localhost:/tmp/mysql/mysql.sock
slave[] = 192.168.3.24:3305
slave[] = 192.168.3.65:3309
```

```
<?php
/* All of the following connections will be load balanced */
$mysqli = new mysqli("myapp", "username", "password", "database");
$pdo = new PDO('mysql:host=myapp;dbname=database', 'username', 'password');
$mysql = mysql_connect("myapp", "username", "password");

$mysqli = new mysqli("localhost", "username", "password", "database");
?>
```

Section names are strings. It is valid to use a section name such as `192.168.2.1`, `127.0.0.1` or `localhost`. If, for example, an application connects to `localhost` and a plugin configuration section `[localhost]` exists, the semantics of the connect operation are changed. The application will no longer only use the MySQL server running on the host `localhost` but the plugin will start to load balance MySQL queries following the rules from the `[localhost]` configuration section. This way you can load balance queries from an application without changing the applications source code.

The `master[]`, `slave[]` and `pick[]` configuration directives use a list-like syntax. Configuration directives supporting list-like syntax may appear multiple times in a configuration section. The plugin maintains the order in which entries appear when interpreting them. For example, the below example shows two `slave[]` configuration directives in the configuration section `[myapp]`. If doing round-robin load balancing for read-only queries, the plugin will send the first read-only query to the MySQL server `mysql_slave_1` because it is the first in the list. The second read-only query will be send to the MySQL server `mysql_slave_2` because it is the second in the list. Configuration directives supporting list-like syntax result are ordered from top to bottom in accordance to their appearance within a configuration section.

### Example 21.293. List-like syntax

```
[myapp]
master[] = mysql_master_server
slave[] = mysql_slave_1
slave[] = mysql_slave_2
```

Here is a short explanation of the configuration directives that can be used.

`master[]` string

URI of a MySQL replication master server. The URI follows the syntax `hostname[:port|unix_domain_socket]`.

The plugin supports using only one master server.

Setting a master server is mandatory. The plugin will report a warning upon connect if the user has failed to provide a master server for a configuration section. The warning may read (`mysqld_ms`) `Cannot find master section in config.`

	<p>Furthermore the plugin may set an error code for the connection handle such as <code>HY000/2000 (CR_UNKNOWN_ERROR)</code>. The corresponding error message depends on your language settings.</p>
<code>slave[]</code> string	<p>URI of one or more MySQL replication slave servers. The URI follows the syntax <code>hostname[:port unix_domain_socket]</code>.</p> <p>The plugin supports using one or more slave servers.</p> <p>Setting a slave server is mandatory. The plugin will report a warning upon connect if the user has failed to provide at least one slave server for a configuration section. The warning may read <code>(mysqld_ms) Cannot find slaves section in config</code>. Furthermore the plugin may set an error code for the connection handle such as <code>HY000/2000 (CR_UNKNOWN_ERROR)</code>. The corresponding error message depends on your language settings.</p>
<code>pick[]</code> string	<p>Load balancing (server picking) policy. Supported policies: <code>random</code>, <code>random_once</code> (default), <code>roundrobin</code>, <code>user</code>.</p> <p>If no load balancing policy is set, the plugin will default to <code>random_once</code>. The <code>random_once</code> policy picks a random slave server when running the first read-only statement. The slave server will be used for all read-only statements until the PHP script execution ends.</p> <p>The <code>random</code> policy will pick a random server whenever a read-only statement is to be executed.</p> <p>If using <code>roundrobin</code> the plugin iterates over the list of configured slave servers to pick a server for statement execution. If the plugin reaches the end of the list, it wraps around to the beginning of the list and picks the first configured slave server.</p> <p>Setting more than one load balancing policy for a configuration section makes only sense in conjunction with <code>user</code> and <code>mysqld_ms_set_user_pick_server</code>. If the user defined callback fails to pick a server, the plugin falls back to the second configured load balancing policy.</p>
<code>failover</code> string	<p>Failover policy. Supported policies: <code>disabled</code> (default), <code>master</code>.</p> <p>If no failover policy is set, the plugin will not do any automatic failover (<code>failover=disabled</code>). Whenever the plugin fails to connect a server it will emit a warning and set the connections error code and message. Thereafter it is up to the application to handle the error and, for example, resent the last statement to trigger the selection of another server.</p> <p>If using <code>failover=master</code> the plugin will implicitly failover to a slave, if available. Please check the concepts documentation to learn about potential pitfalls and risks of using <code>failover=master</code>.</p>
<code>lazy_connections</code> bool	<p>Controls the use of lazy connections. Lazy connections are connections which are not opened before the client sends the first connection.</p> <p>It is strongly recommended to use lazy connections. Lazy connections help to keep the number of open connections low. If you disable lazy connections and, for example, configure one</p>

MySQL replication master server and two MySQL replication slaves, the plugin will open three connections upon the first call to a connect function although the application might use the master connection only.

Lazy connections bare a risk if you make heavy use of actions which change the state of a connection. The plugin does not dispatch all state changing actions to all connections from the connection pool. The few dispatched actions are applied to already opened connections only. Lazy connections opened in the future are not affected. If, for example, the connection character set is changed using a PHP MySQL API call, the plugin will change the character set of all currently opened connection. It will not remember the character set change to apply it on lazy connections opened in the future. As a result the internal connection pool would hold connections using different character sets. This is not desired. Remember that character sets are taken into account for escaping.

`master_on_write` bool

If set, the plugin will use the master server only after the first statement has been executed on the master. Applications can still send statements to the slaves using SQL hints to overrule the automatic decision.

The setting may help with replication lag. If an application runs an `INSERT` the plugin will, by default, use the master to execute all following statements, including `SELECT` statements. This helps to avoid problems with reads from slaves which have not replicated the `INSERT` yet.

`trx_stickiness` string

Transaction stickiness policy. Supported policies: `disabled` (default), `master`.

Experimental feature.

The setting requires 5.4.0 or newer. If used with PHP older than 5.4.0, the plugin will emit a warning like `(mysqlnd_ms) trx_stickiness strategy is not supported before PHP 5.3.99`.

If no transaction stickiness policy is set or, if setting `trx_stickiness=disabled`, the plugin is not transaction aware. Thus, the plugin may load balance connections and switch connections in the middle of a transaction. The plugin is not transaction safe. SQL hints must be used avoid connection switches during a transaction.

As of PHP 5.4.0 the `mysqlnd` library allows the plugin to monitor the `autocommit` mode set by calls to the libraries `trx_autocommit()` function. If setting `trx_stickiness=master` and `autocommit` gets disabled by a PHP MySQL extension invoking the `mysqlnd` library internal function call `trx_autocommit()`, the plugin is made aware of the begin of a transaction. Then, the plugin stops load balancing and directs all statements to the master server until `autocommit` is enabled. Thus, no SQL hints are required.

An example of a PHP MySQL API function calling the `mysqlnd` library internal function call `trx_autocommit()` is `mysqli_autocommit`.

Although setting `trx_stickiness=master`, the plugin cannot be made aware of `autocommit` mode changes caused by SQL statements such as `SET AUTOCOMMIT=0`.

### 21.9.6.6. Testing

Copyright 1997-2012 the PHP Documentation Group. [2906]



#### Note

The section applies to `mysqlnd_ms` 1.1.0 or newer, not the 1.0 series.

The PECL/`mysqlnd_ms` test suite is in the `tests/` directory of the source distribution. The test suite consists of standard phpt tests, which are described on the PHP Quality Assurance Teams website.

Running the tests requires setting up one to four MySQL servers. Some tests don't connect to MySQL at all. Others require one server for testing. Some require two distinct servers. In some cases two servers are used to emulate a replication setup. In other cases a master and a slave of an existing MySQL replication setup are required for testing. The tests will try to detect how many servers and what kind of servers are given. If the required servers are not found, the test will be skipped automatically.

Before running the tests, edit `tests/config.inc` to configure the MySQL servers to be used for testing.

The most basic configuration is as follows.

```
putenv("MYSQL_TEST_HOST=localhost");
putenv("MYSQL_TEST_PORT=3306");
putenv("MYSQL_TEST_USER=root");
putenv("MYSQL_TEST_PASSWD=");
putenv("MYSQL_TEST_DB=test");
putenv("MYSQL_TEST_ENGINE=MyISAM");
putenv("MYSQL_TEST_SOCKET=");

putenv("MYSQL_TEST_SKIP_CONNECT_FAILURE=1");
putenv("MYSQL_TEST_CONNECT_FLAGS=0");
putenv("MYSQL_TEST_EXPERIMENTAL=0");

/* replication cluster emulation */
putenv("MYSQL_TEST_EMULATED_MASTER_HOST=". getenv("MYSQL_TEST_HOST"));
putenv("MYSQL_TEST_EMULATED_SLAVE_HOST=". getenv("MYSQL_TEST_HOST"));

/* real replication cluster */
putenv("MYSQL_TEST_MASTER_HOST=". getenv("MYSQL_TEST_EMULATED_MASTER_HOST"));
putenv("MYSQL_TEST_SLAVE_HOST=". getenv("MYSQL_TEST_EMULATED_SLAVE_HOST"));
```

`MYSQL_TEST_HOST`, `MYSQL_TEST_PORT` and `MYSQL_TEST_SOCKET` define the hostname, TCP/IP port and Unix domain socket of the default database server. `MYSQL_TEST_USER` and `MYSQL_TEST_PASSWD` contain the user and password needed to connect to the database/schema configured with `MYSQL_TEST_DB`. All configured servers must have the same database user configured to give access to the test database.

Using `host`, `host:port` or `host:/path/to/socket` syntax one can set an alternate host, host and port or host and socket for any of the servers.

```
putenv("MYSQL_TEST_SLAVE_HOST=192.168.78.136:3307");
putenv("MYSQL_TEST_MASTER_HOST=myserver_hostname:/path/to/socket");
```



```
| | | | | info : Freeing memory of members
| | | | | info : scheme=unix:///tmp/mysql.sock
| | | | | >mysqlnd_error_list_pdtor
| | | | | <mysqlnd_error_list_pdtor
| | | | | <mysqlnd_conn::free_contents
| | | | | <mysqlnd_conn::connect
[...]
```

The trace log can also be used to verify correct behaviour of PECL/`mysqlnd_ms` itself, for example, to check which server has been selected for query execution and why.

Example excerpt from the debug log (plugin decision):

```
[...]
>mysqlnd_ms::query
| info : query=DROP TABLE IF EXISTS test
| >_mysqlnd_plugin_get_plugin_connection_data
| info : plugin_id=5
| <_mysqlnd_plugin_get_plugin_connection_data
| >mysqlnd_ms_pick_server_ex
| | info : conn_data=0x7fb6a7d3e5a0 *conn_data=0x7fb6a7d410d0
| | >mysqlnd_ms_select_servers_all
| | <mysqlnd_ms_select_servers_all
| | >mysqlnd_ms_choose_connection_rr
| | >mysqlnd_ms_query_is_select
[...]
```

```
| | <mysqlnd_ms_query_is_select
[...]
```

```
| | info : Init the master context
| | info : list(0x7fb6a7d3f598) has 1
| | info : Using master connection
| | >mysqlnd_ms_advanced_connect
| | >mysqlnd_conn::connect
| | | info : host=localhost user=root db=test port=3306 flags=131072 persistent=0 state=0
```

In this case the statement `DROP TABLE IF EXISTS test` has been executed. Note that the statement string is shown in the log file. You may want to take measures to restrict access to the log for security considerations.

The statement has been load balanced using round robin policy, as you can easily guess from the functions name `>mysqlnd_ms_choose_connection_rr`. It has been sent to a master server running on `host=localhost user=root db=test port=3306 flags=131072 persistent=0 state=0`.

### 21.9.6.6.8. Monitoring

Copyright 1997-2012 the PHP Documentation Group. [2906]

Plugin activity can be monitored using the `mysqlnd` trace log, `mysqlnd` statistics, `mysqlnd_ms` plugin statistics and external PHP debugging tools. Use of the trace log should be limited to debugging. It is recommended to use the plugins statistics for monitoring.

Writing a trace log is a slow operation. If using an external PHP debugging tool, please refer to the vendors manual about its performance impact and the type of information collected. In many cases, external debugging tools will provide call stacks. Often, a call stack or a trace log is more difficult to interpret than the statistics provided by the plugin.

Plugin statistics tell how often which kind of cluster node has been used (slave or master), why the node was used, if lazy connections have been used and if global transaction ID injection has been performed. The monitoring information provided enables user to verify plugin decisions and to plan their cluster resources based on usage pattern. The function `mysqlnd_ms_get_stats` is used to access the statistics. Please, see the functions description for a list of available statistics.

Statistics are collected on a per PHP process basis. Their scope is a PHP process. Depending on the PHP deployment model a process may serve one or multiple web requests. If using CGI model, a PHP process serves one web request. If using FastCGI or pre-fork web server models, a PHP process usually serves multiple web requests. The same is the case with a threaded web server. Please, note that threads running in parallel can update the statistics in parallel. Thus, if using a threaded PHP deployment model, statistics can be changed by more than one script at a time. A script cannot rely on the fact that it sees only its own changes to statistics.

**Example 21.294. Verify plugin activity in a non-threaded deployment model**

```
mysqlnd_ms.enable=1
mysqlnd_ms.collect_statistics=1
```

```
<?php
/* Load balanced following "myapp" section rules from the plugins config file (not shown) */
$mysqli = new mysqli("myapp", "username", "password", "database");
if (mysqli_connect_errno())
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

$stats_before = mysqlnd_ms_get_stats();
if ($res = $mysqli->query("SELECT 'Read request' FROM DUAL")) {
    var_dump($res->fetch_all());
}
$stats_after = mysqlnd_ms_get_stats();
if ($stats_after['use_slave'] <= $stats_before['use_slave']) {
    echo "According to the statistics the read request has not been run on a slave!";
}
?>
```

Statistics are aggregated for all plugin activities and all connections handled by the plugin. It is not possible to tell how much a certain connection handle has contributed to the overall statistics.

Utilizing PHPs [register\\_shutdown\\_function](#) function or the [auto\\_append\\_file](#) PHP configuration directive it is easily possible to dump statistics into, for example, a log file when a script finishes. Instead of using a log file it is also possible to send the statistics to an external monitoring tool for recording and display.

**Example 21.295. Recording statistics during shutdown**

```
mysqlnd_ms.enable=1
mysqlnd_ms.collect_statistics=1
error_log=/tmp/php_errors.log
```

```
<?php
function check_stats() {
    $msg = str_repeat("-", 80) . "\n";
    $msg .= var_export(mysqlnd_ms_get_stats(), true) . "\n";
    $msg .= str_repeat("-", 80) . "\n";
    error_log($msg);
}
register_shutdown_function("check_stats");
?>
```

**21.9.6.7. Predefined Constants**

Copyright 1997-2012 the PHP Documentation Group. [2906]



The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

*SQL hint related*

**Example 21.296. Example demonstrating the usage of [mysqLnd\\_ms](#) constants**

The [mysqLnd](#) replication and load balancing plugin ([mysqLnd\\_ms](#)) performs read/write splitting. This directs write queries to a MySQL master server, and read-only queries to the MySQL slave servers. The plugin has a built-in read/write split logic. All queries which start with [SELECT](#) are considered read-only queries, which are then sent to a MySQL slave server that is listed in the plugin configuration file. All other queries are directed to the MySQL master server that is also specified in the plugin configuration file.

User supplied SQL hints can be used to overrule automatic read/write splitting, to gain full control on the process. SQL hints are standards compliant SQL comments. The plugin will scan the beginning of a query string for an SQL comment for certain commands, which then control query redirection. Other systems involved in the query processing are unaffected by the SQL hints because other systems will ignore the SQL comments.

The plugin supports three SQL hints to direct queries to either the MySQL slave servers, the MySQL master server, or the last used MySQL server. SQL hints must be placed at the beginning of a query to be recognized by the plugin.

For better portability, it is recommended to use the string constants [MYSQLND\\_MS\\_MASTER\\_SWITCH](#) , [MYSQLND\\_MS\\_SLAVE\\_SWITCH](#) and [MYSQLND\\_MS\\_LAST\\_USED\\_SWITCH](#) instead of their literal values.

```
<?php
/* Use constants for maximum portability */
$master_query = "/*" . MYSQLND_MS_MASTER_SWITCH . "*/SELECT id FROM test";

/* Valid but less portable: using literal instead of constant */
$slave_query = "/*ms=slave*/SHOW TABLES";

printf("master_query = '%s'\n", $master_query);
printf("slave_query = '%s'\n", $slave_query);
?>
```

The above examples will output:

```
master_query = /*ms=master*/SELECT id FROM test
slave_query = /*ms=slave*/SHOW TABLES
```

[MYSQLND\\_MS\\_MASTER\\_SWITCH](#) (string) SQL hint used to send a query to the MySQL replication master server.

[MYSQLND\\_MS\\_SLAVE\\_SWITCH](#) (string) SQL hint used to send a query to one of the MySQL replication slave servers.

[MYSQLND\\_MS\\_LAST\\_USED\\_SWITCH](#) (string) SQL hint used to send a query to the last used MySQL server. The last used MySQL server can either be a master or a slave server in a MySQL replication setup.

[mysqLnd\\_ms\\_query\\_is\\_select](#) related

[MYSQLND\\_MS\\_QUERY\\_USE\\_MASTER](#) (integer) **if** [mysqLnd\\_ms\\_is\\_select](#) returns [MYSQLND\\_MS\\_QUERY\\_USE\\_MASTER](#) for a given query, the built-in read/write split mechanism recommends sending the query to a MySQL replication master server.



- `MYSQLND_MS_QUERY_USE_SLAVE` (integer) If `mysqlnd_ms_is_select` returns `MYSQLND_MS_QUERY_USE_SLAVE` for a given query, the built-in read/write split mechanism recommends sending the query to a MySQL replication slave server.
- `MYSQLND_MS_QUERY_USE_LAST_USED` (integer) If `mysqlnd_ms_is_select` returns `MYSQLND_MS_QUERY_USE_LAST_USED` for a given query, the built-in read/write split mechanism recommends sending the query to the last used server.
- `mysqlnd_ms_set_qos`, quality of service filter and service level related
- `MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL` (integer) Use to request the service level eventual consistency from the `mysqlnd_ms_set_qos`. Eventual consistency is the default quality of service when reading from an asynchronous MySQL replication slave. Data returned in this service level may or may not be stale, depending on whether the selected slaves happen to have replicated the latest changes from the MySQL replication master or not.
- `MYSQLND_MS_QOS_CONSISTENCY_SESSION` (integer) Use to request the service level session consistency from the `mysqlnd_ms_set_qos`. Session consistency is defined as read your writes. The client is guaranteed to see his latest changes.
- `MYSQLND_MS_QOS_CONSISTENCY_STRONG` (integer) Use to request the service level strong consistency from the `mysqlnd_ms_set_qos`. Strong consistency is used to ensure all clients see each others changes.
- `MYSQLND_MS_QOS_OPTION_GTID` (integer) Used as a service level option with `mysqlnd_ms_set_qos` to parameterize session consistency.
- `MYSQLND_MS_QOS_OPTION_AGE` (integer) Used as a service level option with `mysqlnd_ms_set_qos` to parameterize eventual consistency.

*Other*

The plugins version number can be obtained using `MYSQLND_MS_VERSION` or `MYSQLND_MS_VERSION_ID`. `MYSQLND_MS_VERSION` is the string representation of the numerical version number `MYSQLND_MS_VERSION_ID`, which is an integer such as 10000. Developers can calculate the version number as follows.

Version (part)	Example
Major*10000	1*10000 = 10000
Minor*100	0*100 = 0
Patch	0 = 0
<code>MYSQLND_MS_VERSION_ID</code>	10000

`MYSQLND_MS_VERSION` (string) Plugin version string, for example, "1.0.0-prototype".

`MYSQLND_MS_VERSION_ID` (integer) Plugin version number, for example, 10000.

### 21.9.6.8. MySQLnd\_ms Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.6.8.1. mysqlnd\_ms\_get\_last\_gtid

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_ms_get_last_gtid`

Returns the latest global transaction ID

## Description

```
string mysqlnd_ms_get_last_gtid(
    mixed connection);
```

Returns a global transaction identifier which belongs to a write operation no older than the last write performed by the client. It is not guaranteed that the global transaction identifier is identical to that one created for the last write transaction performed by the client.

## Parameters

*connection*

A PECL/mysqlnd\_ms connection handle to a MySQL server of the type [PDO\\_MYSQL](#), [mysqli](#) or [ext/mysql](#). The connection handle is obtained when opening a connection with a host name that matches a [mysqlnd\\_ms](#) configuration file entry using any of the above three MySQL driver extensions.

## Return Values

Returns a global transaction ID (GTID) on success. Otherwise, returns [FALSE](#).

The function [mysqlnd\\_ms\\_get\\_last\\_gtid](#) returns the GTID obtained when executing the SQL statement from the [fetch\\_last\\_gtid](#) entry of the [global\\_transaction\\_id\\_injection](#) section from the plugins configuration file.

The function may be called after the GTID has been incremented.

## Notes



### Note

[mysqlnd\\_ms\\_get\\_last\\_gtid](#) requires PHP >= 5.4.0 and PECL [mysqlnd\\_ms](#) >= 1.2.0. Internally, it is using a [mysqlnd](#) library C functionality not available with PHP 5.3.

## Examples

### Example 21.297. [mysqlnd\\_ms\\_get\\_last\\_gtid](#) example

```
<?php
/* Open mysqlnd_ms connection using mysqli, PDO_MySQL or mysql extension */
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("DROP TABLE IF EXISTS test"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

printf("GTID after transaction %s\n", mysqlnd_ms_get_last_gtid($mysqli));

/* auto commit mode, transaction on master, GTID must be incremented */
if (!$mysqli->query("CREATE TABLE test(id INT)"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

printf("GTID after transaction %s\n", mysqlnd_ms_get_last_gtid($mysqli));
?>
```

## See Also

[Global Transaction IDs](#)

### 21.9.6.8.2. [mysqlnd\\_ms\\_get\\_last\\_used\\_connection](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_ms\\_get\\_last\\_used\\_connection](#)

Returns an array which describes the last used connection

#### Description

```
array mysqlnd_ms_get_last_used_connection(
    mixed connection);
```

Returns an array which describes the last used connection from the plugins connection pool currently pointed to by the user connection handle. If using the plugin, a user connection handle represents a pool of database connections. It is not possible to tell from the user connection handles properties to which database server from the pool the user connection handle points.

The function can be used to debug or monitor PECL [mysqlnd\\_ms](#).

#### Parameters

*connection* A MySQL connection handle obtained from any of the connect functions of the [mysqli](#), [mysql](#) or [PDO\\_MYSQL](#) extensions.

#### Return Values

[FALSE](#) on error. Otherwise, an array which describes the connection used to execute the last statement on.

Array which describes the connection.

Property	Description	Version
<a href="#">scheme</a>	Connection scheme. Either <code>tcp://host:port</code> or <code>unix://host:socket</code> . If you want to distinguish connections from each other use a combination of <a href="#">scheme</a> and <a href="#">thread_id</a> as a unique key. Neither <a href="#">scheme</a> nor <a href="#">thread_id</a> alone are sufficient to distinguish two connections from each other. Two servers may assign the same <a href="#">thread_id</a> to two different connections. Thus, connections in the pool may have the same <a href="#">thread_id</a> . Also, do not rely on uniqueness of <a href="#">scheme</a> in a pool. Your QA engineers may use the same MySQL server instance for two distinct logical roles and add it multiple times to the pool. This hack is used, for example, in the test suite.	Since 1.1.0.
<a href="#">host</a>	Database server host used with the connection. The host is only set with TCP/IP connections. It is empty with Unix domain or Windows named pipe connections,	Since 1.1.0.
<a href="#">host_info</a>	A character string representing the server hostname and the connection type.	Since 1.1.2.
<a href="#">port</a>	Database server port used with the connection.	Since 1.1.0.
<a href="#">socket_connection</a>	Unix domain socket or Windows named pipe used with the connection. The value is empty for TCP/IP connections.	Since 1.1.2.
<a href="#">thread_id</a>	Connection thread id.	Since 1.1.0.
<a href="#">last_message</a>	Info message obtained from the MySQL C API function <code>mysql_info()</code> . Please, see <a href="#">mysqli_info</a> for a description.	Since 1.1.0.
<a href="#">errno</a>	Error code.	Since 1.1.0.
<a href="#">error</a>	Error message.	Since 1.1.0.
<a href="#">sqlstate</a>	Error SQLstate code.	Since 1.1.0.

#### Notes



**Note**

`mysqlnd_ms_get_last_used_connection` requires PHP >= 5.4.0 and PECL `mysqlnd_ms` >> 1.1.0. Internally, it is using a `mysqlnd` library C call not available with PHP 5.3.

**Examples**

The example assumes that `myapp` refers to a plugin configuration file section and represents a connection pool.

**Example 21.298. `mysqlnd_ms_get_last_used_connection` example**

```
<?php
$link = new mysqli("myapp", "user", "password", "database");
$res = $link->query("SELECT 1 FROM DUAL");
var_dump(mysqlnd_ms_get_last_used_connection($link));
?>
```

The above example will output:

```
array(10) {
  ["scheme"]=>
  string(22) "unix:///tmp/mysql.sock"
  ["host_info"]=>
  string(25) "Localhost via UNIX socket"
  ["host"]=>
  string(0) ""
  ["port"]=>
  int(3306)
  ["socket_or_pipe"]=>
  string(15) "/tmp/mysql.sock"
  ["thread_id"]=>
  int(46253)
  ["last_message"]=>
  string(0) ""
  ["errno"]=>
  int(0)
  ["error"]=>
  string(0) ""
  ["sqlstate"]=>
  string(5) "00000"
}
```

**21.9.6.8.3. `mysqlnd_ms_get_stats`**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_ms_get_stats`

Returns query distribution and connection statistics

**Description**

```
array mysqlnd_ms_get_stats();
```

Returns an array of statistics collected by the replication and load balancing plugin.

The PHP configuration setting `mysqlnd_ms.collect_statistics` controls the collection of statistics. The collection of statistics is disabled by default for performance reasons.

The scope of the statistics is the PHP process. Depending on your deployment model a PHP process may handle one or multiple requests.

Statistics are aggregated for all connections and all storage handler. It is not possible to tell how much queries originating from `mysqli`, `PDO_MySQL` or `mysql` API calls have contributed to the aggregated data values.

**Parameters**

This function has no parameters.

**Return Values**

Returns `NULL` if the PHP configuration directive `mysqlnd_ms.enable` has disabled the plugin. Otherwise, returns array of statistics.

Array of statistics

Statistic	Description	Version
<code>use_slave</code>	<p>The semantics of this statistic has changed between 1.0.1 - 1.1.0.</p> <p>The meaning for version 1.0.1 is as follows. Number of statements considered as read-only by the built-in query analyzer. Neither statements which begin with a SQL hint to force use of slave nor statements directed to a slave by an user-defined callback are included. The total number of statements sent to the slaves is <code>use_slave + use_slave_sql_hint + use_slave_callback</code>.</p> <p>PECL/mysqlnd_ms 1.1.0 introduces a new concept of chained filters. The statistics is now set by the internal load balancing filter. With version 1.1.0 the load balancing filter is always the last in the filter chain, if used. In future versions a load balancing filter may be followed by other filters causing another change in the meaning of the statistic. If, in the future, a load balancing filter is followed by another filter it is no longer guaranteed that the statement, which increments <code>use_slave</code>, will be executed on the slaves.</p> <p>The meaning for version 1.1.0 is as follows. Number of statements sent to the slaves. Statements directed to a slave by the user filter (an user-defined callback) are not included. The latter are counted by <code>use_slave_callback</code>.</p>	Since 1.0.0.
<code>use_master</code>	<p>The semantics of this statistic has changed between 1.0.1 - 1.1.0.</p> <p>The meaning for version 1.0.1 is as follows. Number of statements not considered as read-only by the built-in query analyzer. Neither statements which begin with a SQL hint to force use of master nor statements directed to a master by an user-defined callback are included. The total number of statements sent to the master is <code>use_master + use_master_sql_hint + use_master_callback</code>.</p> <p>PECL/mysqlnd_ms 1.1.0 introduces a new concept of chained filters. The statistics is now set by the internal load balancing filter. With version 1.1.0 the load balancing filter is always the last in the filter chain, if used. In future versions a load balancing filter may be followed by other filters causing another change in the meaning of the statistic. If, in the future, a load balancing filter is followed by another filter it is no longer guaranteed that the statement, which increments <code>use_master</code>, will be executed on the slaves.</p> <p>The meaning for version 1.1.0 is as follows. Number of statements sent to the masters. Statements directed to a master by the user filter (an user-defined callback) are not included. The latter are counted by <code>use_master_callback</code>.</p>	Since 1.0.0.

Statistic	Description	Version
<code>use_slave_hint</code>	Number of statements the built-in query analyzer recommends sending to a slave because they contain no SQL hint to force use of a certain server. The recommendation may be overruled in the following. It is not guaranteed whether the statement will be executed on a slave or not. This is how often the internal <code>is_select</code> function has guessed that a slave shall be used. Please, see also the user space function <code>mysqlnd_ms_query_is_select</code> .	Since 1.1.0.
<code>use_master_hint</code>	Number of statements the built-in query analyzer recommends sending to a master because they contain no SQL hint to force use of a certain server. The recommendation may be overruled in the following. It is not guaranteed whether the statement will be executed on a slave or not. This is how often the internal <code>is_select</code> function has guessed that a master shall be used. Please, see also the user space function <code>mysqlnd_ms_query_is_select</code> .	Since 1.1.0.
<code>use_slave_hint_force</code>	Number of statements sent to a slave because statement begins with the SQL hint to force use of slave.	Since 1.0.0.
<code>use_master_hint_force</code>	Number of statements sent to a master because statement begins with the SQL hint to force use of master.	Since 1.0.0.
<code>use_last_used_server</code>	Number of statements sent to server which has run the previous statement, because statement begins with the SQL hint to force use of previously used server.	Since 1.0.0.
<code>use_slave_callback</code>	Number of statements sent to a slave because an user-defined callback has chosen a slave server for statement execution.	Since 1.0.0.
<code>use_master_callback</code>	Number of statements sent to a master because an user-defined callback has chosen a master server for statement execution.	Since 1.0.0.
<code>non_lazy_connections_slave_success</code>	Number of successfully opened slave connections from configurations not using <code>lazy connections</code> . The total number of successfully opened slave connections is <code>non_lazy_connections_slave_success + lazy_connections_slave_success</code>	Since 1.0.0.
<code>non_lazy_connections_slave_failure</code>	Number of failed slave connection attempts from configurations not using <code>lazy connections</code> . The total number of failed slave connection attempts is <code>non_lazy_connections_slave_failure + lazy_connections_slave_failure</code>	Since 1.0.0.
<code>non_lazy_connections_master_success</code>	Number of successfully opened master connections from configurations not using <code>lazy connections</code> . The total number of successfully opened master connections is <code>non_lazy_connections_master_success + lazy_connections_master_success</code>	Since 1.0.0.
<code>non_lazy_connections_master_failure</code>	Number of failed master connection attempts from configurations not using <code>lazy connections</code> . The total number of failed master connection attempts is <code>non_lazy_connections_master_failure + lazy_connections_master_failure</code>	Since 1.0.0.
<code>lazy_connections_slave_success</code>	Number of successfully opened slave connections from configurations using <code>lazy connections</code> .	Since 1.0.0.
<code>lazy_connections_slave_failure</code>	Number of failed slave connection attempts from configurations using <code>lazy connections</code> .	Since 1.0.0.
<code>lazy_connections_master_success</code>	Number of successfully opened master connections from configurations using <code>lazy connections</code> .	Since 1.0.0.
<code>lazy_connections_master_failure</code>	Number of failed master connection attempts from configurations using <code>lazy connections</code> .	Since 1.0.0.

Statistic	Description	Version
<code>trx_autocommit</code>	Number of <code>autocommit</code> mode activations via API calls. This figure may be used to monitor activity related to the plugin configuration setting <code>trx_stickiness</code> . If, for example, you want to know if a certain API call invokes the <code>mysqlnd</code> library function <code>trx_autocommit()</code> , which is a requirement for <code>trx_stickiness</code> , you may call the user API function in question and check if the statistic has changed. The statistic is modified only by the plugins internal subclassed <code>trx_autocommit()</code> method.	Since 1.0.0.
<code>trx_autocommit_deact</code>	Number of <code>autocommit</code> mode deactivations via API calls.	Since 1.0.0.
<code>trx_master</code>	Number of statements redirected to the master while <code>trx_stickiness=master</code> and <code>autocommit</code> mode is disabled.	Since 1.0.0.
<code>gtid_autocommit</code>	Number of successful SQL injections in <code>autocommit</code> mode as part of the plugins client-side <code>global transaction id emulation</code> .	Since 1.2.0.
<code>gtid_autocommit_fail</code>	Number of failed SQL injections in <code>autocommit</code> mode as part of the plugins client-side <code>global transaction id emulation</code> .	Since 1.2.0.
<code>gtid_commit</code>	Number of successful SQL injections in <code>commit</code> mode as part of the plugins client-side <code>global transaction id emulation</code> .	Since 1.2.0.
<code>gtid_commit_fail</code>	Number of failed SQL injections in <code>commit</code> mode as part of the plugins client-side <code>global transaction id emulation</code> .	Since 1.2.0.
<code>gtid_implicit</code>	Number of successful SQL injections when implicit commit is detected as part of the plugins client-side <code>global transaction id emulation</code> . Implicit commit happens, for example, when <code>autocommit</code> has been turned off, a query is executed and <code>autocommit</code> is enabled again. In that case, the statement will be committed by the server and SQL to maintain is injected before the <code>autocommit</code> is re-enabled. Another sequence causing an an implicit commit is <code>begin()</code> , <code>query()</code> , <code>begin()</code> . The second call to <code>begin()</code> will implicitly commit the transaction started by the first call to <code>begin()</code> . <code>begin()</code> refers to internal library calls not actual PHP user API calls.	Since 1.2.0.
<code>gtid_implicit_fail</code>	Number of failed SQL injections when implicit commit is detected as part of the plugins client-side <code>global transaction id emulation</code> . Implicit commit happens, for example, when <code>autocommit</code> has been turned off, a query is executed and <code>autocommit</code> is enabled again. In that case, the statement will be committed by the server and SQL to maintain is injected before the <code>autocommit</code> is re-enabled.	Since 1.2.0.

## Examples

### Example 21.299. `mysqlnd_ms_get_stats` example

```
<?php
printf("mysqlnd_ms.enable = %d\n", ini_get("mysqlnd_ms.enable"));
printf("mysqlnd_ms.collect_statistics = %d\n", ini_get("mysqlnd_ms.collect_statistics"));
var_dump(mysqlnd_ms_get_stats());
?>
```

The above example will output:

```
mysqlnd_ms.enable = 1
mysqlnd_ms.collect_statistics = 1
array(26) {
  ["use_slave"]=>
    string(1) "0"
```

```
[ "use_master" ]=>
string(1) "0"
[ "use_slave_guess" ]=>
string(1) "0"
[ "use_master_guess" ]=>
string(1) "0"
[ "use_slave_sql_hint" ]=>
string(1) "0"
[ "use_master_sql_hint" ]=>
string(1) "0"
[ "use_last_used_sql_hint" ]=>
string(1) "0"
[ "use_slave_callback" ]=>
string(1) "0"
[ "use_master_callback" ]=>
string(1) "0"
[ "non_lazy_connections_slave_success" ]=>
string(1) "0"
[ "non_lazy_connections_slave_failure" ]=>
string(1) "0"
[ "non_lazy_connections_master_success" ]=>
string(1) "0"
[ "non_lazy_connections_master_failure" ]=>
string(1) "0"
[ "lazy_connections_slave_success" ]=>
string(1) "0"
[ "lazy_connections_slave_failure" ]=>
string(1) "0"
[ "lazy_connections_master_success" ]=>
string(1) "0"
[ "lazy_connections_master_failure" ]=>
string(1) "0"
[ "trx_autocommit_on" ]=>
string(1) "0"
[ "trx_autocommit_off" ]=>
string(1) "0"
[ "trx_master_forced" ]=>
string(1) "0"
[ "gtid_autocommit_injections_success" ]=>
string(1) "0"
[ "gtid_autocommit_injections_failure" ]=>
string(1) "0"
[ "gtid_commit_injections_success" ]=>
string(1) "0"
[ "gtid_commit_injections_failure" ]=>
string(1) "0"
[ "gtid_implicit_commit_injections_success" ]=>
string(1) "0"
[ "gtid_implicit_commit_injections_failure" ]=>
string(1) "0"
}
```

### See Also

[Runtime configuration](#)  
[mysqld\\_ms.collect\\_statistics](#)  
[mysqld\\_ms.enable](#)  
[Monitoring](#)

#### 21.9.6.8.4. `mysqld_ms_match_wild`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqld_ms_match_wild`

Finds whether a table name matches a wildcard pattern or not

### Description



```
bool mysqlnd_ms_match_wild(
    string table_name,
    string wildcard);
```

Finds whether a table name matches a wildcard pattern or not.

This function is not of much practical relevance with PECL [mysqlnd\\_ms](#) 1.1.0 because the plugin does not support MySQL replication table filtering yet.

### Parameters

<i>table_name</i>	The table name to check if it is matched by the wildcard.
<i>wildcard</i>	The wildcard pattern to check against the table name. The wildcard pattern supports the same placeholders as MySQL replication filters do.  MySQL replication filters can be configured by using the MySQL Server configuration options <code>--replicate-wild-do-table</code> and <code>--replicate-wild-do-db</code> . Please, consult the MySQL Reference Manual to learn more about this MySQL Server feature.  The supported placeholders are: <ul style="list-style-type: none"> <li>• <code>%</code> - zero or more literals</li> <li>• <code>%</code> - one literal</li> </ul> Placeholders can be escaped using <code>\</code> .

### Return Values

Returns `TRUE` *table\_name* is matched by *wildcard*. Otherwise, returns `FALSE`

### Examples

#### Example 21.300. [mysqlnd\\_ms\\_match\\_wild](#) example

```
<?php
var_dump(mysqlnd_ms_match_wild("schema_name.table_name", "schema%"));
var_dump(mysqlnd_ms_match_wild("abc", "_"));
var_dump(mysqlnd_ms_match_wild("table1", "table_"));
var_dump(mysqlnd_ms_match_wild("asia_customers", "%customers"));
var_dump(mysqlnd_ms_match_wild("funny%table", "funny\\%table"));
var_dump(mysqlnd_ms_match_wild("funnytable", "funny%table"));
?>
```

The above example will output:

```
bool(true)
bool(false)
bool(true)
bool(true)
bool(true)
bool(true)
bool(true)
```

#### 21.9.6.8.5. [mysqlnd\\_ms\\_query\\_is\\_select](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_ms\\_query\\_is\\_select](#)

Find whether to send the query to the master, the slave or the last used MySQL server

## Description

```
int mysqlnd_ms_query_is_select(
    string query);
```

Finds whether to send the query to the master, the slave or the last used MySQL server.

The plugin's built-in read/write split mechanism will be used to analyze the query string to make a recommendation where to send the query. The built-in read/write split mechanism is very basic and simple. The plugin will recommend sending all queries to the MySQL replication master server but those which begin with `SELECT`, or begin with a SQL hint which enforces sending the query to a slave server. Due to the basic but fast algorithm the plugin may propose to run some read-only statements such as `SHOW TABLES` on the replication master.

## Parameters

*query* Query string to test.

## Return Values

A return value of `MYSQLND_MS_QUERY_USE_MASTER` indicates that the query should be sent to the MySQL replication master server. The function returns a value of `MYSQLND_MS_QUERY_USE_SLAVE` if the query can be run on a slave because it is considered read-only. A value of `MYSQLND_MS_QUERY_USE_LAST_USED` is returned to recommend running the query on the last used server. This can either be a MySQL replication master server or a MySQL replication slave server.

If read write splitting has been disabled by setting `mysqlnd_ms.disable_rw_split`, the function will always return `MYSQLND_MS_QUERY_USE_MASTER` or `MYSQLND_MS_QUERY_USE_LAST_USED`.

## Examples

### Example 21.301. `mysqlnd_ms_query_is_select` example

```
<?php
function is_select($query)
{
    switch (mysqlnd_ms_query_is_select($query))
    {
        case MYSQLND_MS_QUERY_USE_MASTER:
            printf("%s' should be run on the master.\n", $query);
            break;
        case MYSQLND_MS_QUERY_USE_SLAVE:
            printf("%s' should be run on a slave.\n", $query);
            break;
        case MYSQLND_MS_QUERY_USE_LAST_USED:
            printf("%s' should be run on the server that has run the previous query\n", $query);
            break;
        default:
            printf("No suggestion where to run the '%s', fallback to master recommended\n", $query);
            break;
    }
}

is_select("INSERT INTO test(id) VALUES (1)");
is_select("SELECT 1 FROM DUAL");
is_select("/* . MYSQLND_MS_LAST_USED_SWITCH . */SELECT 2 FROM DUAL");
?>
```

The above example will output:

```
INSERT INTO test(id) VALUES (1) should be run on the master.
```

```
SELECT 1 FROM DUAL should be run on a slave.
/*ms=last_used*/SELECT 2 FROM DUAL should be run on the server that has run the previous query
```

### See Also

#### Predefined Constants

[user filter](#)

#### Runtime configuration

[mysqlnd\\_ms.disable\\_rw\\_split](#)

[mysqlnd\\_ms.enable](#)

### 21.9.6.8.6. [mysqlnd\\_ms\\_set\\_qos](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_ms\\_set\\_qos](#)

Sets the quality of service needed from the cluster

### Description

```
bool mysqlnd_ms_set_qos(
    mixed connection,
    int service_level,
    int service_level_option,
    mixed option_value);
```

Sets the quality of service needed from the cluster. A database cluster delivers a certain quality of service to the user depending on its architecture. A major aspect of the quality of service is the consistency level the cluster can offer. An asynchronous MySQL replication cluster defaults to eventual consistency for slave reads: a slave may serve stale data, current data, or it may have not the requested data at all, because it is not synchronous to the master. In a MySQL replication cluster, only master accesses can give strong consistency, which promises that all clients see each others changes.

PECL/mysqlnd\_ms hides the complexity of choosing appropriate nodes to achieve a certain level of service from the cluster. The "Quality of Service" filter implements the necessary logic. The filter can either be configured in the plugins configuration file, or at runtime using [mysqlnd\\_ms\\_set\\_qos](#).

Similar results can be achieved with PECL mysqlnd\_ms < 1.2.0, if using SQL hints to force the use of a certain type of node or using the [master\\_on\\_write](#) plugin configuration option. The first requires more code and causes more work on the application side. The latter is less refined than using the quality of service filter. Settings made through the function call can be reversed, as shown in the example below. The example temporarily switches to a higher service level (session consistency, read your writes) and returns back to the clusters default after it has performed all operations that require the better service. This way, read load on the master can be minimized compared to using [master\\_on\\_write](#), which would continue using the master after the first write.

Since 1.5.0 calls will fail when done in the middle of a transaction if [transaction stickiness](#) is enabled and transaction boundaries have been detected. properly.

### Parameters

*connection*

A PECL/mysqlnd\_ms connection handle to a MySQL server of the type [PDO\\_MYSQL](#), [mysqli](#) or [ext/mysqli](#) for which a service level is to be set. The connection handle is obtained when opening a connection with a host name that matches a [mysqlnd\\_ms](#) configuration file entry using any of the above three MySQL driver extensions.

*service\_level*

The requested service level:

[MYSQLND\\_MS\\_QOS\\_CONSISTENCY\\_EVENTUAL](#) ,

`MYSQLND_MS_QOS_CONSISTENCY_SESSION` or  
`MYSQLND_MS_QOS_CONSISTENCY_STRONG` .

*service\_level\_option*

An option to parameterize the requested service level. The option can either be `MYSQLND_MS_QOS_OPTION_GTID` or `MYSQLND_MS_QOS_OPTION_AGE` .

The option `MYSQLND_MS_QOS_OPTION_GTID` can be used to refine the service level `MYSQLND_MS_QOS_CONSISTENCY_SESSION` . It must be combined with a fourth function parameter, the *option\_value*. The *option\_value* shall be a global transaction ID obtained from `mysqlnd_ms_get_last_gtid`. If set, the plugin considers both master servers and asynchronous slaves for session consistency (read your writes). Otherwise, only masters are used to achieve session consistency. A slave is considered up-to-date and checked if it has already replicated the global transaction ID from *option\_value*. Please note, searching appropriate slaves is an expensive and slow operation. Use the feature sparsely, if the master cannot handle the read load alone.

The `MYSQLND_MS_QOS_OPTION_AGE` option can be combined with the `MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL` service level, to filter out asynchronous slaves that lag more seconds behind the master than *option\_value*. If set, the plugin will only consider slaves for reading if `SHOW SLAVE STATUS` reports `Slave_IO_Running=Yes`, `Slave_SQL_Running=Yes` and `Seconds_Behind_Master <= option_value`. Please note, searching appropriate slaves is an expensive and slow operation. Use the feature sparsely in version 1.2.0. Future versions may improve the algorithm used to identify candidates. Please, see the MySQL reference manual about the precision, accuracy and limitations of the MySQL administrative command `SHOW SLAVE STATUS`.

*option\_value*

Parameter value for the service level option. See also the *service\_level\_option* parameter.

## Return Values

Returns `TRUE` if the connections service level has been switched to the requested. Otherwise, returns `FALSE`

## Notes



### Note

`mysqlnd_ms_set_qos` requires PHP  $\geq$  5.4.0 and PECL `mysqlnd_ms`  $\geq$  1.2.0. Internally, it is using a `mysqlnd` library C functionality not available with PHP 5.3.

## Examples

### Example 21.302. `mysqlnd_ms_set_qos` example

```
<?php
/* Open mysqlnd_ms connection using mysqli, PDO_MySQL or mysql extension */
$mysqli = new mysqli("myapp", "username", "password", "database");
if (!$mysqli)
    /* Of course, your error handling is nicer... */
    die(sprintf("[%d] %s\n", mysqli_connect_errno(), mysqli_connect_error()));

/* Session consistency: read your writes */
```

```
$ret = mySQLnd_ms_set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_SESSION);
if (!$ret)
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Will use master and return fresh data, client can see his last write */
if (!$res = $mysqli->query("SELECT item, price FROM orders WHERE order_id = 1"))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

/* Back to default: use of all slaves and masters permitted, stale data can happen */
if (!$mysqli->set_qos($mysqli, MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL))
    die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));
?>
```

### See Also

[mysqLnd\\_ms\\_get\\_last\\_gtid](#)  
[Service level and consistency concept](#)  
[Filter concept](#)

### 21.9.6.8.7. [mysqLnd\\_ms\\_set\\_user\\_pick\\_server](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqLnd\\_ms\\_set\\_user\\_pick\\_server](#)

Sets a callback for user-defined read/write splitting

### Description

```
bool mySQLnd_ms_set_user_pick_server(
    string function);
```

Sets a callback for user-defined read/write splitting. The plugin will call the callback only if `pick[]=user` is the default rule for server picking in the relevant section of the plugins configuration file.

The plugins built-in read/write query split mechanism decisions can be overwritten in two ways. The easiest way is to prepend the query string with the SQL hints [MYSQLND\\_MS\\_MASTER\\_SWITCH](#) , [MYSQLND\\_MS\\_SLAVE\\_SWITCH](#) or [MYSQLND\\_MS\\_LAST\\_USED\\_SWITCH](#) . Using SQL hints one can control, for example, whether a query shall be send to the MySQL replication master server or one of the slave servers. By help of SQL hints it is not possible to pick a certain slave server for query execution.

Full control on server selection can be gained using a callback function. Use of a callback is recommended to expert users only because the callback has to cover all cases otherwise handled by the plugin.

The plugin will invoke the callback function for selecting a server from the lists of configured master and slave servers. The callback function inspects the query to run and picks a server for query execution by returning the hosts URI, as found in the master and slave list.

If the lazy connections are enabled and the callback chooses a slave server for which no connection has been established so far and establishing the connection to the slave fails, the plugin will return an error upon the next action on the failed connection, for example, when running a query. It is the responsibility of the application developer to handle the error. For example, the application can re-run the query to trigger a new server selection and callback invocation. If so, the callback must make sure to select a different slave, or check slave availability, before returning to the plugin to prevent an endless loop.

### Parameters

*function*                      The function to be called. Class methods may also be invoked statically using this function by passing `array($classname,`

[\\$methodname](#)) to this parameter. Additionally class methods of an object instance may be called by passing [array\(\\$objectinstance, \\$methodname\)](#) to this parameter.

## Return Values

Host to run the query on. The host URI is to be taken from the master and slave connection lists passed to the callback function. If callback returns a value neither found in the master nor in the slave connection lists the plugin will fallback to the second pick method configured via the [pick\[\]](#) setting in the plugin configuration file. If not second pick method is given, the plugin falls back to the build-in default pick method for server selection.

## Notes



### Note

[mysqlnd\\_ms\\_set\\_user\\_pick\\_server](#) is available with PECL [mysqlnd\\_ms](#) < 1.1.0. It has been replaced by the [user](#) filter. Please, check the [Change History](#) for upgrade notes.

## Examples

### Example 21.303. [mysqlnd\\_ms\\_set\\_user\\_pick\\_server](#) example

```
[myapp]
master[] = localhost
slave[] = 192.168.2.27:3306
slave[] = 192.168.78.136:3306
pick[] = user
```

```
<?php

function pick_server($connected, $query, $master, $slaves, $last_used)
{
    static $slave_idx = 0;
    static $num_slaves = NULL;
    if (is_null($num_slaves))
        $num_slaves = count($slaves);

    /* default: fallback to the plugins build-in logic */
    $ret = NULL;

    printf("User has connected to '%s'...\n", $connected);
    printf("... deciding where to run '%s'\n", $query);

    $where = mysqlnd_ms_query_is_select($query);
    switch ($where)
    {
        case MYSQLND_MS_QUERY_USE_MASTER:
            printf("... using master\n");
            $ret = $master[0];
            break;
        case MYSQLND_MS_QUERY_USE_SLAVE:
            /* SELECT or SQL hint for using slave */
            if (striestr($query, "FROM table_on_slave_a_only"))
            {
                /* a table which is only on the first configured slave */
                printf("... access to table available only on slave A detected\n");
                $ret = $slaves[0];
            }
            else
            {
                /* round robin */
                printf("... some read-only query for a slave\n");
            }
    }
}
```

```

    $ret = $slaves[$slave_idx++ % $num_slaves];
  }
  break;
case MYSQLND_MS_QUERY_LAST_USED:
  printf("... using last used server\n");
  $ret = $last_used;
  break;
}

printf("... ret = '%s'\n", $ret);
return $ret;
}

mysqld_ms_set_user_pick_server("pick_server");

$mysqli = new mysqli("myapp", "root", "root", "test");

if (!$res = $mysqli->query("SELECT 1 FROM DUAL"))
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
  $res->close();

if (!$res = $mysqli->query("SELECT 2 FROM DUAL"))
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
  $res->close();

if (!$res = $mysqli->query("SELECT * FROM table_on_slave_a_only"))
  printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
else
  $res->close();

$mysqli->close();
?>

```

The above example will output:

```

User has connected to 'myapp'...
... deciding where to run 'SELECT 1 FROM DUAL'
... some read-only query for a slave
... ret = 'tcp://192.168.2.27:3306'
User has connected to 'myapp'...
... deciding where to run 'SELECT 2 FROM DUAL'
... some read-only query for a slave
... ret = 'tcp://192.168.78.136:3306'
User has connected to 'myapp'...
... deciding where to run 'SELECT * FROM table_on_slave_a_only'
... access to table available only on slave A detected
... ret = 'tcp://192.168.2.27:3306'

```

### See Also

[mysqld\\_ms\\_query\\_is\\_select](#)  
[Filter concept](#)  
[user filter](#)

## 21.9.6.9. Change History

Copyright 1997-2012 the PHP Documentation Group. [2906]

This change history is a high level summary of selected changes that may impact applications and/or break backwards compatibility.

See also the [CHANGES](#) file in the source distribution for a complete list of changes.

### 21.9.6.9.1. PECL/mysqlnd\_ms 1.5 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

1.5.0-alpha

- Release date: 03/2013
- Motto/theme: Sharding support, improved transaction support



#### Note

This is the current development series. All features are at an early stage. Changes may happen at any time without prior notice. Please, do not use this version in production environments.

The documentation may not reflect all changes yet.

#### Bug fixes

- Fixed #60605 PHP segmentation fault when `mysqlnd_ms` is enabled.
- Setting transaction stickiness disables all load balancing, including automatic failover, for the duration of a transaction. So far connection switches could have happened in the middle of a transaction in multi-master configurations and during automatic failover although transaction monitoring had detected transaction boundaries properly.
- BC break and bug fix. SQL hints enforcing the use of a specific kind of server ( `MYSQLND_MS_MASTER_SWITCH` , `MYSQLND_MS_SLAVE_SWITCH` , `MYSQLND_MS_LAST_USED_SWITCH` ) are ignored for the duration of a transaction of transaction stickiness is enabled and transaction boundaries have been detected properly.

This is a change in behaviour. However, it is also a bug fix and a step to align behaviour. If, in previous versions, transaction stickiness, one of the above listed SQL hints and the quality of service filtering was combined it could happen that the SQL hints got ignored. In some cases the SQL hints did work, in other cases they did not. The new behaviour is more consistent. SQL hints will always be ignored for the duration of a transaction, if [transaction stickiness](#) is enabled.

Please note, transaction boundary detection continues to be based on API call monitoring. SQL commands controlling transactions are not monitored.

- BC break and bug fix. Calls to `mysqlnd_ms_set_qos` will fail when done in the middle of a transaction if [transaction stickiness](#) is enabled. Connection switches are not allowed for the duration of a transaction. Changing the quality of service likely results on a different set of servers qualifying for query execution, possibly making it necessary to switch connections. Thus, the call is not allowed in during an active transaction. The quality of server can, however, be changed in between transactions.

#### Feature changes

- Introduced the `node_group` filter. The filter lets you organize servers (master and slaves) into groups. Queries can be directed to a certain group of servers by prefixing the query statement with a SQL hint/comment that contains the groups configured name. Grouping can be used for partitioning and sharding, and also to optimize for local caching. In the case of sharding, a group name can be thought of like a shard key. All queries for a given shard key will be executed on the configured shard. Note: both the client and server must support sharding for sharding to function with `mysqlnd_ms`.
- Extended configuration file validation during PHP startup (RINIT). An `E_WARNING` level error will be thrown if the configuration file can not be read (permissions), is empty, or the file (JSON) could not be parsed. Warnings may appear in log files, which depending on how PHP is configured.



Distributions that aim to provide a pre-configured setup, including a configuration file stub, are asked to put `{ }` into the configuration file to prevent this warning about an invalid configuration file.

Further configuration file validation is done when parsing sections upon opening a connection. Please, note that there may still be situations when an invalid plugin configuration file does not lead to proper error messages but a failure to connect.

- As of PHP 5.5.0, improved support for transaction boundaries detection was added for [mysqli](#). The [mysqli](#) extension has been modified to use the new C API calls of the [mysqlnd](#) library to begin, commit, and rollback a transaction or savepoint. If [trx\\_stickiness](#) is used to enable transaction aware load balancing, the [mysqli\\_begin](#), [mysqli\\_commit](#) and [mysqli\\_rollback](#) functions will now be monitored by the plugin, to go along with the [mysqli\\_autocommit](#) function that was already supported. All SQL features to control transactions are also available through the improved [mysqli](#) transaction control related functions. This means that it is not required to issue SQL statements instead of using API calls. Applications using the appropriate API calls can be load balanced by PECL/[mysqlnd\\_ms](#) in a completely transaction-aware way.

Please note, [PDO\\_MYSQL](#) has not been updated yet to utilize the new [mysqlnd](#) API calls. Thus, transaction boundary detection with [PDO\\_MYSQL](#) continues to be limited to the monitoring by passing in `PDO::ATTR_AUTOCOMMIT` to `PDO::setAttribute`.

- Introduced `trx_stickiness=on`. This [trx\\_stickiness](#) option differs from `trx_stickiness=master` as it tries to execute a read-only transaction on a slave, if quality of service (consistency level) allows the use of a slave. Read-only transactions were introduced in MySQL 5.6, and they offer performance gains.
- Query cache support is considered beta if used with the [mysqli](#) API. It should work fine with primary copy based clusters. For all other APIs, this feature continues to be called experimental.
- The code examples in the [mysqlnd\\_ms](#) source were updated.

#### 21.9.6.9.2. PECL/[mysqlnd\\_ms](#) 1.4 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

##### 1.4.2-stable

- Release date: 08/2012
- Motto/theme: Tweaking based on user feedback

##### 1.4.1-beta

- Release date: 08/2012
- Motto/theme: Tweaking based on user feedback

##### Bug fixes

- Fixed build with PHP 5.5

##### 1.4.0-alpha

- Release date: 07/2012
- Motto/theme: Tweaking based on user feedback

##### Feature changes

- BC break: Renamed plugin configuration setting `ini_file` to `config_file`. In early versions the plugin configuration file used ini style. Back then the configuration setting was named accordingly. It has now been renamed to reflect the newer file format and to distinguish it from PHP's own ini file (configuration directives file).

- Introduced new default charset setting [server\\_charset](#) to allow proper escaping before a connection is opened. This is most useful when using lazy connections, which are a default.
- Introduced [wait\\_for\\_gtid\\_timeout](#) setting to throttle slave reads that need session consistency. If global transaction identifier are used and the service level is set to session consistency, the plugin tries to find up-to-date slaves. The slave status check is done by a SQL statement. If nothing else is set, the slave status is checked only one can the search for more up-to-date slaves continues immediately thereafter. Setting [wait\\_for\\_gtid\\_timeout](#) instructs the plugin to poll a slaves status for [wait\\_for\\_gtid\\_timeout](#) seconds if the first execution of the SQL statement has shown that the slave is not up-to-date yet. The poll will be done once per second. This way, the plugin will wait for slaves to catch up and throttle the client.
- New failover strategy [loop\\_before\\_master](#). By default the plugin does no failover. It is possible to enable automatic failover if a connection attempt fails. Upto version 1.3 only [master](#) strategy existed to failover to a master if a slave connection fails. [loop\\_before\\_master](#) is similar but tries all other slaves before attempting to connect to the master if a slave connection fails.

The number of attempts can be limited using the [max\\_retries](#) option. Failed hosts can be remembered and skipped in load balancing for the rest of the web request. [max\\_retries](#) and [remember\\_failed](#) are considered experimental although decent stability is given. Syntax and semantics may change in the future without prior notice.

### 21.9.6.9.3. PECL/mysqLnd\_ms 1.3 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

#### 1.3.2-stable

- Release date: 04/2012
- Motto/theme: see 1.3.0-alpha

#### Bug fixes

- Fixed problem with multi-master where although in a transaction the queries to the master weren't sticky and were spread all over the masters (RR). Still not sticky for Random. [Random\\_once](#) is not affected.

#### 1.3.1-beta

- Release date: 04/2012
- Motto/theme: see 1.3.0-alpha

#### Bug fixes

- Fixed problem with building together with QC.

#### 1.3.0-alpha

- Release date: 04/2012
- Motto/theme: Query caching through quality-of-service concept

The 1.3 series aims to improve the performance of applications and the overall load of an asynchronous MySQL cluster, for example, a MySQL cluster using MySQL Replication. This is done by transparently replacing a slave access with a local cache access, if the application allows it by setting an appropriate quality of service flag. When using MySQL replication a slave can serve stale data. An application using MySQL replication must continue to work correctly with stale data. Given that the application is know to work correctly with stale data, the slave access can transparently be replace with a local cache access.

[PECL/mysqlnd\\_qc](#) serves as a cache backend. PECL/mysqlnd\_qc supports use of various storage locations, among others main memory, [APC](#) and [MEMCACHE](#).

#### Feature changes

- Added cache option to quality-of-service (QoS) filter.
  - New configure option `enable-mysqlnd-ms-cache-support`
  - New constant `MYSQLND_MS_HAVE_CACHE_SUPPORT`.
  - New constant `MYSQLND_MS_QOS_OPTION_CACHE` to be used with `mysqlnd_ms_set_qos`.
- Support for built-in global transaction identifier feature of MySQL 5.6.5-m8 or newer.

#### 21.9.6.9.4. PECL/mysqlnd\_ms 1.2 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

##### 1.2.1-beta

- Release date: 01/2012
- Motto/theme: see 1.2.0-alpha

Minor test changes.

##### 1.2.0-alpha

- Release date: 11/2011
- Motto/theme: Global Transaction ID injection and quality-of-service concept

In version 1.2 the focus continues to be on supporting MySQL database clusters with asynchronous replication. The plugin tries to make using the cluster introducing a quality-of-service filter which applications can use to define what service quality they need from the cluster. Service levels provided are eventual consistency with optional maximum age/slave slag, session consistency and strong consistency.

Additionally the plugin can do client-side global transaction id injection to make manual master failover easier.

#### Feature changes

- Introduced quality-of-service (QoS) filter. Service levels provided by QoS filter:
  - eventual consistency, optional option slave lag
  - session consistency, optional option GTID
  - strong consistency
- Added the `mysqlnd_ms_set_qos` function to set the required connection quality at runtime. The new constants related to `mysqlnd_ms_set_qos` are:
  - `MYSQLND_MS_QOS_CONSISTENCY_STRONG`
  - `MYSQLND_MS_QOS_CONSISTENCY_SESSION`
  - `MYSQLND_MS_QOS_CONSISTENCY_EVENTUAL`
  - `MYSQLND_MS_QOS_OPTION_GTID`
  - `MYSQLND_MS_QOS_OPTION_AGE`

- Added client-side global transaction id injection (GTID).
- New statistics related to GTID:
  - `gtid_autocommit_injections_success`
  - `gtid_autocommit_injections_failure`
  - `gtid_commit_injections_success`
  - `gtid_commit_injections_failure`
  - `gtid_implicit_commit_injections_success`
  - `gtid_implicit_commit_injections_failure`
- Added `mysqlnd_ms_get_last_gtid` to fetch the last global transaction id.
- Enabled support for multi master without slaves.

#### 21.9.6.9.5. PECL/mysqlnd\_ms 1.1 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

1.1.0

- Release date: 09/2011
- Motto/theme: Cover replication basics with production quality

The 1.1 and 1.0 series expose a similar feature set. Internally, the 1.1 series has been refactored to plan for future feature additions. A new configuration file format has been introduced, and limitations have been lifted. And the code quality and quality assurance has been improved.

Feature changes

- Added the (chainable) [filter concept](#):
  - BC break: `mysqlnd_ms_set_user_pick_server` has been removed. The [http://svn.php.net/viewvc/pecl/mysqlnd\\_ms/trunk/user](http://svn.php.net/viewvc/pecl/mysqlnd_ms/trunk/user) filter has been introduced to replace it. The filter offers similar functionality, but see below for an explanation of the differences.
- New powerful JSON based configuration syntax.
- [Lazy connections improved](#): security relevant, and state changing commands are covered.
- Support for (native) prepared statements.
- New statistics: `use_master_guess`, `use_slave_guess`.
  - BC break: Semantics of statistics changed for `use_slave`, `use_master`. Future changes are likely. Please see, `mysqlnd_ms_get_stats`.
- List of broadcasted messages extended by `ssl_set`.
- Library calls now monitored to remember settings for lazy connections: `change_user`, `select_db`, `set_charset`, `set_autocommit`.
- Introduced `mysqlnd_ms.disable_rw_split`. The configuration setting allows using the load balancing and lazy connection functionality independently of read write splitting.

Bug fixes

- Fixed PECL #22724 - Server switching (`mysqlnd_ms_query_is_select()` case sensitive)

- Fixed PECL #22784 - Using `mysql_connect` and `mysql_select_db` did not work
- Fixed PECL #59982 - Unusable extension with `--enable-mysqlnd-ms-table-filter`. Use of the option is NOT supported. You must not use it. Added note to m4.
- Fixed Bug #60119 - `host="localhost"` lost in `mysqlnd_ms_get_last_used_connection()`

The `mysqlnd_ms_set_user_pick_server` function was removed, and replaced in favor of a new `user` filter. You can no longer set a callback function using `mysqlnd_ms_set_user_pick_server` at runtime, but instead have to configure it in the plugins configuration file. The `user` filter will pass the same arguments to the callback as before. Therefore, you can continue to use the same procedural function as a `callback.callback`. It is no longer possible to use static class methods, or class methods of an object instance, as a callback. Doing so will cause the function executing a statement handled by the plugin to emit an `E_RECOVERABLE_ERROR` level error, which might look like: `"(mysqlnd_ms) Specified callback (picker) is not a valid callback."` Note: this may halt your application.

#### 21.9.6.9.6. PECL/mysqlnd\_ms 1.0 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

1.0.1-alpha

- Release date: 04/2011
- Motto/theme: bug fix release

1.0.0-alpha

- Release date: 04/2011
- Motto/theme: Cover replication basics to test user feedback

The first release of practical use. It features basic automatic read-write splitting, SQL hints to overrule automatic redirection, load balancing of slave requests, lazy connections, and optional, automatic use of the master after the first write.

The public feature set is close to that of the 1.1 release.

1.0.0-pre-alpha

- Release date: 09/2010
- Motto/theme: Proof of concept

Initial check-in. Essentially a demo of the [mysqlnd](#) plugin API.

#### 21.9.7. Mysqlnd query result cache plugin ([mysqlnd\\_qc](#))

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The `mysqlnd` query result cache plugin adds easy to use client-side query caching to all PHP MySQL extensions using [mysqlnd](#).

As of version PHP 5.3.3 the MySQL native driver for PHP ([mysqlnd](#)) features an internal plugin C API. C plugins, such as the query cache plugin, can extend the functionality of [mysqlnd](#).

`mysqlnd` plugins such as the query cache plugin operate transparent from a user perspective. The cache plugin supports all PHP applications and all PHP MySQL extensions ([mysqli](#), [mysql](#), [PDO\\_MYSQL](#)). It does not change existing APIs.

No significant application changes are required to cache a query. The cache has two operation modes. It will either cache all queries (not recommended) or only those queries marked with a certain SQL hint (recommended).

### 21.9.7.1. Key Features

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

- Transparent and therefore easy to use
  - supports all PHP MySQL extensions
  - no API changes
  - very little application changes required
- Flexible invalidation strategy
  - Time-to-Live (TTL)
  - user-defined
- Storage with different scope and life-span
  - Default (Hash, process memory)
  - [APC](#)
  - MEMCACHE
  - sqlite
  - user-defined
- Built-in slam defense to prevent cache stampeding.

### 21.9.7.2. Limitations

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The current 1.0.1 release of PECL `mysqlnd_qc` does not support PHP 5.4. Version 1.1.0-alpha lifts this limitation.

Prepared statements and unbuffered queries are fully supported. Thus, the plugin is capable of caching all statements issued with `mysqli` or `PDO_MySQL`, which are the only two PHP MySQL APIs to offer prepared statement support.

### 21.9.7.3. On the name

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The shortcut `mysqlnd_qc` stands for `mysqlnd query cache plugin`. The name was chosen for a quick-and-dirty proof-of-concept. In the beginning the developers did not expect to continue using the code base. Sometimes PECL/`mysqlnd_qc` has also been called `client-side query result set cache`.

### 21.9.7.4. Quickstart and Examples

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The `mysqlnd` query cache plugin is easy to use. This quickstart will demo typical use-cases, and provide practical advice on getting started.

It is strongly recommended to read the reference sections in addition to the quickstart. It is safe to begin with the quickstart. However, before using the plugin in mission critical environments we urge you to read additionally the background information from the reference sections.

Most of the examples use the [mysqli](#) extension because it is the most feature complete PHP MySQL extension. However, the plugin can be used with any PHP MySQL extension that is using the [mysqlnd](#) library.

#### 21.9.7.4.1. Architecture and Concepts

Copyright 1997-2012 the PHP Documentation Group. [2906]

The query cache plugin is implemented as a PHP extension. It is written in C and operates under the hood of PHP. During the startup of the PHP interpreter, it gets registered as a [mysqlnd](#) plugin to replace selected [mysqlnd](#) C methods. Hereby, it can change the behaviour of any PHP MySQL extension ([mysqli](#), [PDO\\_MYSQL](#), [mysql](#)) compiled to use the [mysqlnd](#) library without changing the extensions API. This makes the plugin compatible with each and every PHP MySQL application. Because existing APIs are not changed, it is almost transparent to use. Please, see the [mysqlnd plugin API description](#) for a discussion of the advantages of the plugin architecture and a comparison with proxy based solutions.

##### *Transparent to use*

At PHP run time PECL/[mysqlnd\\_qc](#) can proxy queries send from PHP ([mysqlnd](#)) to the MySQL server. It then inspects the statement string to find whether it shall cache its results. If so, result set is cached using a storage handler and further executions of the statement are served from the cache for a user-defined period. The Time to Live (TTL) of the cache entry can either be set globally or on a per statement basis.

A statement is either cached if the plugin is instructed to cache all statements globally using a or, if the query string starts with the SQL hint (`/*qc=on*/`). The plugin is capable of caching any query issued by calling appropriate API calls of any of the existing PHP MySQL extensions.

##### *Flexible storage: various storage handler*

Various storage handler are supported to offer different scopes for cache entries. Different scopes allow for different degrees in sharing cache entries among clients.

- [default](#) (built-in): process memory, scope: process, one or more web requests depending on PHP deployment model used
- [APC](#): shared memory, scope: single server, multiple web requests
- [SQLite](#): memory or file, scope: single server, multiple web requests
- [MEMCACHE](#): main memory, scope: single or multiple server, multiple web requests
- [user](#) (built-in): user-defined - any, scope: user-defined - any

Support for the [APC](#), [SQLite](#) and [MEMCACHE](#) storage handler has to be enabled at compile time. The [default](#) and [user](#) handler are built-in. It is possible to switch between compiled-in storage handlers on a per query basis at run time. However, it is recommended to pick one storage handler and use it for all cache entries.

##### *Built-in slam defense to avoid overloading*

To avoid overload situations the cache plugin has a built-in slam defense mechanism. If a popular cache entries expires many clients using the cache entries will try to refresh the cache entry. For the duration of the refresh many clients may access the database server concurrently. In the worst case, the database server becomes overloaded and it takes more and more time to refresh the cache entry, which in turn lets more and more clients try to refresh the cache entry. To prevent this from happening the plugin has a slam defense mechanism. If slam defense is enabled and the plugin detects an expired cache entry it extends the life time of the cache entry before it refreshes the cache entry. This way other concurrent accesses to the expired cache entry are still served from the cache for a certain time. The other concurrent accesses do not trigger a concurrent refresh. Ideally, the cache entry gets

refreshed by the client which extended the cache entries lifespan before other clients try to refresh the cache and potentially cause an overload situation.

#### *Unique approach to caching*

PECL/mysqlnd\_qc has a unique approach to caching result sets that is superior to application based cache solutions. Application based solutions first fetch a result set into PHP variables. Then, the PHP variables are serialized for storage in a persistent cache, and then unserialized when fetching. The mysqlnd query cache stores the raw wire protocol data sent from MySQL to PHP in its cache and replays it, if still valid, on a cache hit. This way, it saves an extra serialization step for a cache put that all application based solutions have to do. It can store the raw wire protocol data in the cache without having to serialize into a PHP variable first and deserializing the PHP variable for storing in the cache again.

### 21.9.7.4.2. Setup

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin is implemented as a PHP extension. See also the [installation instructions](#) to install the PECL/mysqlnd\_qc extension.

Compile or configure the PHP MySQL extension (`mysqli`, `PDO_MYSQL`, `mysql`) that you plan to use with support for the `mysqlnd` library. PECL/mysqlnd\_qc is a plugin for the `mysqlnd` library. To use the plugin with any of the existing PHP MySQL extensions (APIs), the extension has to use the `mysqlnd` library.

Then, load the extension into PHP and activate the plugin in the PHP configuration file using the PHP configuration directive named `mysqlnd_qc.enable_qc`.

#### **Example 21.304. Enabling the plugin (php.ini)**

```
mysqlnd_qc.enable_qc=1
```

### 21.9.7.4.3. Caching queries

Copyright 1997-2012 the PHP Documentation Group. [2906]

There are four ways to trigger caching of a query.

- Use of SQL hints on a per query basis
- User supplied callbacks to decide on a per query basis, for example, using `mysqlnd_qc_is_select`
- `mysqlnd_set_cache_condition` for rule based automatic per query decisions
- `mysqlnd_qc.cache_by_default = 1` to cache all queries blindly

Use of SQL hints and `mysqlnd_qc.cache_by_default = 1` are explained below. Please, refer to the function reference on `mysqlnd_qc_is_select` for a description of using a callback and, `mysqlnd_qc_set_cache_condition` on how to set rules for automatic caching.

A SQL hint is a SQL standards compliant comment. As a SQL comment it is ignored by the database. A statement is considered eligible for caching if it either begins with the SQL hint enabling caching or it is a `SELECT` statement.

An individual query which shall be cached must begin with the SQL hint `/*qc=on*/`. It is recommended to use the PHP constant `MYSQLND_QC_ENABLE_SWITCH` instead of using the string value.

- not eligible for caching and not cached: `INSERT INTO test(id) VALUES (1)`



- not eligible for caching and not cached: `SHOW ENGINES`
- eligible for caching but uncached: `SELECT id FROM test`
- eligible for caching and cached: `/*qc=on*/SELECT id FROM test`

The examples `SELECT` statement string is prefixed with the `MYSQLND_QC_ENABLE_SWITCH` SQL hint to enable caching of the statement. The SQL hint must be given at the very beginning of the statement string to enable caching.

**Example 21.305. Using the `MYSQLND_QC_ENABLE_SWITCH` SQL hint**

```
mysqlnd_qc.enable_qc=1

<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

/* Will be cached because of the SQL hint */
$start = microtime(true);
$res    = $mysqli->query("/* . MYSQLND_QC_ENABLE_SWITCH . */" . "SELECT id FROM test WHERE id = 1");

var_dump($res->fetch_assoc());
$res->free();

printf("Total time uncached query: %.6fs\n", microtime(true) - $start);

/* Cache hit */
$start = microtime(true);
$res    = $mysqli->query("/* . MYSQLND_QC_ENABLE_SWITCH . */" . "SELECT id FROM test WHERE id = 1");

var_dump($res->fetch_assoc());
$res->free();

printf("Total time cached query: %.6fs\n", microtime(true) - $start);
?>
```

The above examples will output something similar to:

```
array(1) {
  ["id"]=>
  string(1) "1"
}
Total time uncached query: 0.000740s
array(1) {
  ["id"]=>
  string(1) "1"
}
Total time cached query: 0.000098s
```

If nothing else is configured, as it is the case in the quickstart example, the plugin will use the built-in `default` storage handler. The `default` storage handler uses process memory to hold a cache entry. Depending on the PHP deployment model, a PHP process may serve one or more web requests. Please, consult the web server manual for details. Details make no difference for the examples given in the quickstart.

The query cache plugin will cache all queries regardless if the query string begins with the SQL hint which enables caching or not, if the PHP configuration directive `mysqlnd_qc.cache_by_default` is

set to 1. The setting `mysqlnd_qc.cache_by_default` is evaluated by the core of the query cache plugins. Neither the built-in nor user-defined storage handler can overrule the setting.

The SQL hint `/*qc=off*/` can be used to disable caching of individual queries if `mysqlnd_qc.cache_by_default = 1`. It is recommended to use the PHP constant `MYSQLND_QC_DISABLE_SWITCH` instead of using the string value.

**Example 21.306. Using the `MYSQLND_QC_DISABLE_SWITCH` SQL hint**

```
mysqlnd_qc.enable_qc=1
mysqlnd_qc.cache_by_default=1
```

```
<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

/* Will be cached although no SQL hint is present because of mysqlnd_qc.cache_by_default = 1*/
$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

$mysqli->query("DELETE FROM test WHERE id = 1");

/* Cache hit - no automatic invalidation and still valid! */
$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

/* Cache miss - query must not be cached because of the SQL hint */
$res = $mysqli->query("/*" . MYSQLND_QC_DISABLE_SWITCH . "*/SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();
?>
```

The above examples will output:

```
array(1) {
  ["id"]=>
  string(1) "1"
}
array(1) {
  ["id"]=>
  string(1) "1"
}
NULL
```

PECL/`mysqlnd_qc` forbids caching of statements for which at least one column from the statements result set shows no table name in its meta data by default. This is usually the case for columns originating from SQL functions such as `NOW()` or `LAST_INSERT_ID()`. The policy aims to prevent pitfalls if caching by default is used.

**Example 21.307. Example showing which type of statements are not cached**

```
mysqlnd_qc.enable_qc=1
mysqlnd_qc.cache_by_default=1
```

```

<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1)");

for ($i = 0; $i < 3; $i++) {

    $start = microtime(true);

    /* Note: statement will not be cached because of NOW() use */
    $res = $mysqli->query("SELECT id, NOW() AS _time FROM test");
    $row = $res->fetch_assoc();

    /* dump results */
    var_dump($row);

    printf("Total time: %.6fs\n", microtime(true) - $start);

    /* pause one second */
    sleep(1);
}
?>

```

The above examples will output something similar to:

```

array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:43:10"
}
Total time: 0.000540s
array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:43:11"
}
Total time: 0.000555s
array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:43:12"
}
Total time: 0.000549s

```

It is possible to enable caching for all statements including those which has columns in their result set for which MySQL reports no table, such as the statement from the example. Set [mysqlnd\\_qc.cache\\_no\\_table = 1](#) to enable caching of such statements. Please, note the difference in the measured times for the above and below examples.

**Example 21.308. Enabling caching for all statements using the [mysqlnd\\_qc.cache\\_no\\_table](#) ini setting**

```

mysqlnd_qc.enable_qc=1
mysqlnd_qc.cache_by_default=1
mysqlnd_qc.cache_no_table=1

```

```
<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1)");

for ($i = 0; $i < 3; $i++) {

    $start = microtime(true);

    /* Note: statement will not be cached because of NOW() use */
    $res = $mysqli->query("SELECT id, NOW() AS _time FROM test");
    $row = $res->fetch_assoc();

    /* dump results */
    var_dump($row);

    printf("Total time: %.6fs\n", microtime(true) - $start);

    /* pause one second */
    sleep(1);
}
?>
```

The above examples will output something similar to:

```
array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:47:45"
}
Total time: 0.000546s
array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:47:45"
}
Total time: 0.000187s
array(2) {
  ["id"]=>
  string(1) "1"
  ["_time"]=>
  string(19) "2012-01-11 15:47:45"
}
Total time: 0.000167s
```



#### Note

Although `mysqlnd_qc.cache_no_table = 1` has been created for use with `mysqlnd_qc.cache_by_default = 1` it is bound it. The plugin will evaluate the `mysqlnd_qc.cache_no_table` whenever a query is to be cached, no matter whether caching has been enabled using a SQL hint or any other measure.

#### 21.9.7.4.4. Setting the TTL

Copyright 1997-2012 the PHP Documentation Group. [2906]

The default invalidation strategy of the query cache plugin is Time to Live (TTL). The built-in storage handlers will use the default TTL defined by the PHP configuration value `mysqlnd_qc.ttl` unless

the query string contains a hint for setting a different `TTL`. The `TTL` is specified in seconds. By default cache entries expire after 30 seconds

The example sets `mysqlnd_qc.ttl=3` to cache statements for three seconds by default. Every second it updates a database table record to hold the current time and executes a `SELECT` statement to fetch the record from the database. The `SELECT` statement is cached for three seconds because it is prefixed with the SQL hint enabling caching. The output verifies that the query results are taken from the cache for the duration of three seconds before they are refreshed.

### Example 21.309. Setting the TTL with the `mysqlnd_qc.ttl` ini setting

```
mysqlnd_qc.enable_qc=1
mysqlnd_qc.ttl=3
```

```
<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id VARCHAR(255))");

for ($i = 0; $i < 7; $i++) {

    /* update DB row */
    if (!$mysqli->query("DELETE FROM test") ||
        !$mysqli->query("INSERT INTO test(id) VALUES (NOW())"))
        /* Of course, a real-life script should do better error handling */
        die(sprintf("[%d] %s\n", $mysqli->errno, $mysqli->error));

    /* select latest row but cache results */
    $query = "/*" . MYSQLND_QC_ENABLE_SWITCH . "*/";
    $query .= "SELECT id AS _time FROM test";
    if (!$res = $mysqli->query($query) ||
        !($row = $res->fetch_assoc()))
    {
        printf("[%d] %s\n", $mysqli->errno, $mysqli->error);
    }
    $res->free();
    printf("Wall time %s - DB row time %s\n", date("H:i:s"), $row['_time']);

    /* pause one second */
    sleep(1);
}
?>
```

The above examples will output something similar to:

```
Wall time 14:55:59 - DB row time 2012-01-11 14:55:59
Wall time 14:56:00 - DB row time 2012-01-11 14:55:59
Wall time 14:56:01 - DB row time 2012-01-11 14:55:59
Wall time 14:56:02 - DB row time 2012-01-11 14:56:02
Wall time 14:56:03 - DB row time 2012-01-11 14:56:02
Wall time 14:56:04 - DB row time 2012-01-11 14:56:02
Wall time 14:56:05 - DB row time 2012-01-11 14:56:05
```

As can be seen from the example, any `TTL` based cache can serve stale data. Cache entries are not automatically invalidated, if underlying data changes. Applications using the default `TTL` invalidation strategy must be able to work correctly with stale data.

A user-defined cache storage handler can implement any invalidation strategy to work around this limitation.

The default `TTL` can be overruled using the SQL hint `/*qc_tt=seconds*/`. The SQL hint must be appear immediately after the SQL hint which enables caching. It is recommended to use the PHP constant `MYSQLND_QC_TTL_SWITCH` instead of using the string value.

### Example 21.310. Setting TTL with SQL hints

```
<?php
$start = microtime(true);

/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

printf("Default TTL\t: %d seconds\n", ini_get("mysqlnd_qc.ttl"));

/* Will be cached for 2 seconds */
$sql = sprintf("/*s*//*s%d*/SELECT id FROM test WHERE id = 1", MYSQLND_QC_ENABLE_SWITCH, MYSQLND_QC_TTL_S);
$res = $mysqli->query($sql);

var_dump($res->fetch_assoc());
$res->free();

$mysqli->query("DELETE FROM test WHERE id = 1");
sleep(1);

/* Cache hit - no automatic invalidation and still valid! */
$res = $mysqli->query($sql);
var_dump($res->fetch_assoc());
$res->free();

sleep(2);

/* Cache miss - cache entry has expired */
$res = $mysqli->query($sql);
var_dump($res->fetch_assoc());
$res->free();

printf("Script runtime\t: %d seconds\n", microtime(true) - $start);
?>
```

The above examples will output something similar to:

```
Default TTL      : 30 seconds
array(1) {
  ["id"]=>
    string(1) "1"
}
array(1) {
  ["id"]=>
    string(1) "1"
}
NULL
Script runtime   : 3 seconds
```

#### 21.9.7.4.5. Pattern based caching

Copyright 1997-2012 the PHP Documentation Group. [2906]

An application has three options for telling PECL/`mysqlnd_qc` whether a particular statement shall be used. The most basic approach is to cache all statements by setting `mysqlnd_qc.cache_by_default = 1`. This approach is often of little practical value. But it enables

users to make a quick estimation about the maximum performance gains from caching. An application designed to use a cache may be able to prefix selected statements with the appropriate SQL hints. However, altering an applications source code may not always be possible or desired, for example, to avoid problems with software updates. Therefore, PECL/mysqlnd\_qc allows setting a callback which decides if a query is to be cached.

The callback is installed with the `mysqlnd_qc_set_is_select` function. The callback is given the statement string of every statement inspected by the plugin. Then, the callback can decide whether to cache the function. The callback is supposed to return `FALSE` if the statement shall not be cached. A return value of `TRUE` makes the plugin try to add the statement into the cache. The cache entry will be given the default TTL ( `mysqlnd_qc.ttl`). If the callback returns a numerical value it is used as the TTL instead of the global default.

### Example 21.311. Setting a callback with `mysqlnd_qc_set_is_select`

```
mysqlnd_qc.enable_qc=1
mysqlnd_qc.collect_statistics=1
```

```
<?php
/* callback which decides if query is cached */
function is_select($query) {
    static $patterns = array(
        /* true - use default from mysqlnd_qc.ttl */
        "@SELECT\s+.*\s+FROM\s+test@ismU" => true,
        /* 3 - use TTL = 3 seconds */
        "@SELECT\s+.*\s+FROM\s+news@ismU" => 3
    );

    /* check if query does match pattern */
    foreach ($patterns as $pattern => $ttl) {
        if (preg_match($pattern, $query)) {
            printf("is_select(%45s): cache\n", $query);
            return $ttl;
        }
    }
    printf("is_select(%45s): do not cache\n", $query);
    return false;
}
/* install callback */
mysqlnd_qc_set_is_select("is_select");

/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* cache put */
$mysqli->query("SELECT id FROM test WHERE id = 1");
/* cache hit */
$mysqli->query("SELECT id FROM test WHERE id = 1");
/* cache put */
$mysqli->query("SELECT * FROM test");

$stats = mysqlnd_qc_get_core_stats();
printf("Cache put: %d\n", $stats['cache_put']);
printf("Cache hit: %d\n", $stats['cache_hit']);
?>
```

The above examples will output something similar to:

```
is_select(          DROP TABLE IF EXISTS test): do not cache
```

```
is_select(          CREATE TABLE test(id INT)): do not cache
is_select(    INSERT INTO test(id) VALUES (1), (2), (3)): do not cache
is_select(          SELECT id FROM test WHERE id = 1): cache
is_select(          SELECT id FROM test WHERE id = 1): cache
is_select(          SELECT * FROM test): cache
Cache put: 2
Cache hit: 1
```

The examples callback tests if a statement string matches a pattern. If this is the case, it either returns `TRUE` to cache the statement using the global default TTL or an alternative TTL.

To minimize application changes the callback can put into and registered in an auto prepend file.

#### 21.9.7.4.6. Slam defense

Copyright 1997-2012 the PHP Documentation Group. [2906]

A badly designed cache can do more harm than good. In the worst case a cache can increase database server load instead of minimizing it. An overload situation can occur if a highly shared cache entry expires (cache stampeding).

Cache entries are shared and reused to a different degree depending on the storage used. The default storage handler stores cache entries in process memory. Thus, a cache entry can be reused for the life-span of a process. Other PHP processes cannot access it. If Memcache is used, a cache entry can be shared among multiple PHP processes and even among multiple machines, depending on the set up being used.

If a highly shared cache entry stored, for example, in Memcache expires, many clients gets a cache miss. Many client requests can no longer be served from the cache but try to run the underlying query on the database server. Until the cache entry is refreshed, more and more clients contact the database server. In the worst case, a total lost of service is the result.

The overload can be avoided using a storage handler which limits the reuse of cache entries to few clients. Then, at the average, its likely that only a limited number of clients will try to refresh a cache entry concurrently.

Additionally, the built-in slam defense mechanism can and should be used. If slam defense is activated an expired cache entry is given an extended life time. The first client getting a cache miss for the expired cache entry tries to refresh the cache entry within the extended life time. All other clients requesting the cache entry are temporarily served from the cache although the original `TTL` of the cache entry has expired. The other clients will not experience a cache miss before the extended life time is over.

#### Example 21.312. Enabling the slam defense mechanism

```
mysqlnd_qc.slam_defense=1
mysqlnd_qc.slam_defense_ttl=1
```

The slam defense mechanism is enabled with the PHP configuration directive `mysqlnd_qc.slam_defense`. The extended life time of a cache entry is set with `mysqlnd_qc.slam_defense_ttl`.

The function `mysqlnd_qc_get_core_stats` returns an array of statistics. The statistics `slam_stale_refresh` and `slam_stale_hit` are incremented if slam defense takes place.

It is not possible to give a one-fits-all recommendation on the slam defense configuration. Users are advised to monitor and test their setup and derive settings accordingly.



### 21.9.7.4.7. Finding cache candidates

Copyright 1997-2012 the PHP Documentation Group. [2906]

A statement should be considered for caching if it is executed often and has a long run time. Cache candidates are found by creating a list of statements sorted by the product of the number of executions multiplied by the statements run time. The function `mysqlnd_qc_get_query_trace_log` returns a query log which help with the task.

Collecting a query trace is a slow operation. Thus, it is disabled by default. The PHP configuration directive `mysqlnd_qc.collect_query_trace` is used to enable it. The functions trace contains one entry for every query issued before the function is called.

#### Example 21.313. Collecting a query trace

```
mysqlnd_qc.enable_qc=1
mysqlnd_qc.collect_query_trace=1
```

```
<?php
/* connect to MySQL */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");

/* dummy queries to fill the query trace */
for ($i = 0; $i < 2; $i++) {
    $res = $mysqli->query("SELECT 1 AS _one FROM DUAL");
    $res->free();
}

/* dump trace */
var_dump(mysqlnd_qc_get_query_trace_log());
?>
```

The above examples will output:

```
array(2) {
  [0]=>
  array(8) {
    ["query"]=>
    string(26) "SELECT 1 AS _one FROM DUAL"
    ["origin"]=>
    string(102) "#0 qc.php(7): mysqli->query('SELECT 1 AS _on...')
#1 {main}"
    ["run_time"]=>
    int(0)
    ["store_time"]=>
    int(25)
    ["eligible_for_caching"]=>
    bool(false)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(false)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [1]=>
  array(8) {
    ["query"]=>
    string(26) "SELECT 1 AS _one FROM DUAL"
    ["origin"]=>
    string(102) "#0 qc.php(7): mysqli->query('SELECT 1 AS _on...')
#1 {main}"
```

```

["run_time"]=>
int(0)
["store_time"]=>
int(8)
["eligible_for_caching"]=>
bool(false)
["no_table"]=>
bool(false)
["was_added"]=>
bool(false)
["was_already_in_cache"]=>
bool(false)
}
}

```

Assorted information is given in the trace. Among them timings and the origin of the query call. The origin property holds a code backtrace to identify the source of the query. The depth of the backtrace can be limited with the PHP configuration directive `mysqlnd_qc.query_trace_bt_depth`. The default depth is 3.

**Example 21.314. Setting the backtrace depth with the `mysqlnd_qc.query_trace_bt_depth` ini setting**

```

mysqlnd_qc.enable_qc=1
mysqlnd_qc.collect_query_trace=1

```

```

<?php
/* connect to MySQL */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* dummy queries to fill the query trace */
for ($i = 0; $i < 3; $i++) {
    $res = $mysqli->query("SELECT id FROM test WHERE id = " . $mysqli->real_escape_string($i));
    $res->free();
}

$trace = mysqlnd_qc_get_query_trace_log();
$summary = array();
foreach ($trace as $entry) {
    if (!isset($summary[$entry['query']])) {
        $summary[$entry['query']] = array(
            "executions" => 1,
            "time"       => $entry['run_time'] + $entry['store_time'],
        );
    } else {
        $summary[$entry['query']]['executions']++;
        $summary[$entry['query']]['time'] += $entry['run_time'] + $entry['store_time'];
    }
}

foreach ($summary as $query => $details) {
    printf("%45s: %5dms (%dx)\n",
        $query, $details['time'], $details['executions']);
}
?>

```

The above examples will output something similar to:

```

        DROP TABLE IF EXISTS test:      0ms (1x)
        CREATE TABLE test(id INT):      0ms (1x)
INSERT INTO test(id) VALUES (1), (2), (3): 0ms (1x)
        SELECT id FROM test WHERE id = 0: 25ms (1x)
        SELECT id FROM test WHERE id = 1: 10ms (1x)
        SELECT id FROM test WHERE id = 2:  9ms (1x)

```

#### 21.9.7.4.8. Measuring cache efficiency

Copyright 1997-2012 the PHP Documentation Group. [2906]

PECL/`mysqlnd_qc` offers three ways to measure the cache efficiency. The function `mysqlnd_qc_get_normalized_query_trace_log` returns statistics aggregated by the normalized query string, `mysqlnd_qc_get_cache_info` gives storage handler specific information which includes a list of all cached items, depending on the storage handler. Additionally, the core of PECL/`mysqlnd_qc` collects high-level summary statistics aggregated per PHP process. The high-level statistics are returned by `mysqlnd_qc_get_core_stats`.

The functions `mysqlnd_qc_get_normalized_query_trace_log` and `mysqlnd_qc_get_core_stats` will not collect data unless data collection has been enabled through their corresponding PHP configuration directives. Data collection is disabled by default for performance considerations. It is configurable with the `mysqlnd_qc.time_statistics` option, which determines if timing information should be collected. Collection of time statistics is enabled by default but only performed if data collection as such has been enabled. Recording time statistics causes extra system calls. In most cases, the benefit of the monitoring outweighs any potential performance penalty of the additional system calls.

**Example 21.315. Collecting statistics data with the `mysqlnd_qc.time_statistics` ini setting**

```

mysqlnd_qc.enable_qc=1
mysqlnd_qc.collect_statistics=1

```

```

<?php
/* connect to MySQL */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* dummy queries */
for ($i = 1; $i <= 4; $i++) {
    $query = sprintf("/*%s*/SELECT id FROM test WHERE id = %d", MYSQLND_QC_ENABLE_SWITCH, $i % 2);
    $res    = $mysqli->query($query);

    $res->free();
}

var_dump(mysqlnd_qc_get_core_stats());
?>

```

The above examples will output something similar to:

```

array(26) {
  ["cache_hit"]=>
  string(1) "2"
  ["cache_miss"]=>
  string(1) "2"
  ["cache_put"]=>

```

```

string(1) "2"
["query_should_cache"]=>
string(1) "4"
["query_should_not_cache"]=>
string(1) "3"
["query_not_cached"]=>
string(1) "3"
["query_could_cache"]=>
string(1) "4"
["query_found_in_cache"]=>
string(1) "2"
["query_uncached_other"]=>
string(1) "0"
["query_uncached_no_table"]=>
string(1) "0"
["query_uncached_no_result"]=>
string(1) "0"
["query_uncached_use_result"]=>
string(1) "0"
["query_aggr_run_time_cache_hit"]=>
string(2) "28"
["query_aggr_run_time_cache_put"]=>
string(3) "900"
["query_aggr_run_time_total"]=>
string(3) "928"
["query_aggr_store_time_cache_hit"]=>
string(2) "14"
["query_aggr_store_time_cache_put"]=>
string(2) "40"
["query_aggr_store_time_total"]=>
string(2) "54"
["receive_bytes_recorded"]=>
string(3) "136"
["receive_bytes_replayed"]=>
string(3) "136"
["send_bytes_recorded"]=>
string(2) "84"
["send_bytes_replayed"]=>
string(2) "84"
["slam_stale_refresh"]=>
string(1) "0"
["slam_stale_hit"]=>
string(1) "0"
["request_counter"]=>
int(1)
["process_hash"]=>
int(1929695233)
}

```

For a quick overview, call `mysqlnd_qc_get_core_stats`. It delivers cache usage, cache timing and traffic related statistics. Values are aggregated on a per process basis for all queries issued by any PHP MySQL API call.

Some storage handler, such as the default handler, can report cache entries, statistics related to the entries and meta data for the underlying query through the `mysqlnd_qc_get_cache_info` function. Please note, that the information returned depends on the storage handler. Values are aggregated on a per process basis.

#### Example 21.316. Example `mysqlnd_qc_get_cache_info` usage

```
mysqlnd_qc.enable_qc=1
```

```
<?php
/* connect to MySQL */
```

```

$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* dummy queries to fill the query trace */
for ($i = 1; $i <= 4; $i++) {
    $query = sprintf("/*%s*/SELECT id FROM test WHERE id = %d", MYSQLND_QC_ENABLE_SWITCH, $i % 2);
    $res = $mysqli->query($query);

    $res->free();
}

var_dump(mysqlnd_qc_get_cache_info());
?>

```

The above examples will output something similar to:

```

array(4) {
  ["num_entries"]=>
  int(2)
  ["handler"]=>
  string(7) "default"
  ["handler_version"]=>
  string(5) "1.0.0"
  ["data"]=>
  array(2) {
    ["Localhost via UNIX socket
3306
root
test|/*qc=on*/SELECT id FROM test WHERE id = 1"]=>
    array(2) {
      ["statistics"]=>
      array(11) {
        ["rows"]=>
        int(1)
        ["stored_size"]=>
        int(71)
        ["cache_hits"]=>
        int(1)
        ["run_time"]=>
        int(391)
        ["store_time"]=>
        int(27)
        ["min_run_time"]=>
        int(16)
        ["max_run_time"]=>
        int(16)
        ["min_store_time"]=>
        int(8)
        ["max_store_time"]=>
        int(8)
        ["avg_run_time"]=>
        int(8)
        ["avg_store_time"]=>
        int(4)
      }
      ["metadata"]=>
      array(1) {
        [0]=>
        array(8) {
          ["name"]=>
          string(2) "id"
          ["orig_name"]=>
          string(2) "id"
          ["table"]=>
          string(4) "test"
          ["orig_table"]=>
          string(4) "test"
        }
      }
    }
  }
}

```

```

    ["db"]=>
    string(4) "test"
    ["max_length"]=>
    int(1)
    ["length"]=>
    int(11)
    ["type"]=>
    int(3)
  }
}
["Localhost via UNIX socket
3306
root
test|/*qc=on*/SELECT id FROM test WHERE id = 0"]=>
array(2) {
  ["statistics"]=>
  array(11) {
    ["rows"]=>
    int(0)
    ["stored_size"]=>
    int(65)
    ["cache_hits"]=>
    int(1)
    ["run_time"]=>
    int(299)
    ["store_time"]=>
    int(13)
    ["min_run_time"]=>
    int(11)
    ["max_run_time"]=>
    int(11)
    ["min_store_time"]=>
    int(6)
    ["max_store_time"]=>
    int(6)
    ["avg_run_time"]=>
    int(5)
    ["avg_store_time"]=>
    int(3)
  }
  ["metadata"]=>
  array(1) {
    [0]=>
    array(8) {
      ["name"]=>
      string(2) "id"
      ["orig_name"]=>
      string(2) "id"
      ["table"]=>
      string(4) "test"
      ["orig_table"]=>
      string(4) "test"
      ["db"]=>
      string(4) "test"
      ["max_length"]=>
      int(0)
      ["length"]=>
      int(11)
      ["type"]=>
      int(3)
    }
  }
}
}
}
}

```

It is possible to further break down the granularity of statistics to the level of the normalized statement string. The normalized statement string is the statements string with all parameters replaced with question marks. For example, the two statements `SELECT id FROM test WHERE id = 0` and

`SELECT id FROM test WHERE id = 1` are normalized into `SELECT id FROM test WHERE id = ?`. Their both statistics are aggregated into one entry for `SELECT id FROM test WHERE id = ?`.

**Example 21.317. Example `mysqld_qc_get_normalized_query_trace_log` usage**

```
mysqld_qc.enable_qc=1
mysqld_qc.collect_normalized_query_trace=1
```

```
<?php
/* connect to MySQL */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* dummy queries to fill the query trace */
for ($i = 1; $i <= 4; $i++) {
    $query = sprintf("/*%s*/SELECT id FROM test WHERE id = %d", MYSQLND_QC_ENABLE_SWITCH, $i % 2);
    $res = $mysqli->query($query);

    $res->free();
}

var_dump(mysqld_qc_get_normalized_query_trace_log());
?>
```

The above examples will output something similar to:

```
array(4) {
  [0]=>
  array(9) {
    ["query"]=>
    string(25) "DROP TABLE IF EXISTS test"
    ["occurrences"]=>
    int(0)
    ["eligible_for_caching"]=>
    bool(false)
    ["avg_run_time"]=>
    int(0)
    ["min_run_time"]=>
    int(0)
    ["max_run_time"]=>
    int(0)
    ["avg_store_time"]=>
    int(0)
    ["min_store_time"]=>
    int(0)
    ["max_store_time"]=>
    int(0)
  }
  [1]=>
  array(9) {
    ["query"]=>
    string(27) "CREATE TABLE test (id INT )"
    ["occurrences"]=>
    int(0)
    ["eligible_for_caching"]=>
    bool(false)
    ["avg_run_time"]=>
    int(0)
    ["min_run_time"]=>
    int(0)
    ["max_run_time"]=>
    int(0)
    ["avg_store_time"]=>
```

```

int(0)
["min_store_time"]=>
int(0)
["max_store_time"]=>
int(0)
}
[2]=>
array(9) {
  ["query"]=>
string(46) "INSERT INTO test (id ) VALUES ( ? ), ( ? ), ( ? )"
  ["occurrences"]=>
int(0)
  ["eligible_for_caching"]=>
bool(false)
  ["avg_run_time"]=>
int(0)
  ["min_run_time"]=>
int(0)
  ["max_run_time"]=>
int(0)
  ["avg_store_time"]=>
int(0)
  ["min_store_time"]=>
int(0)
  ["max_store_time"]=>
int(0)
}
[3]=>
array(9) {
  ["query"]=>
string(31) "SELECT id FROM test WHERE id =?"
  ["occurrences"]=>
int(4)
  ["eligible_for_caching"]=>
bool(true)
  ["avg_run_time"]=>
int(179)
  ["min_run_time"]=>
int(11)
  ["max_run_time"]=>
int(393)
  ["avg_store_time"]=>
int(12)
  ["min_store_time"]=>
int(7)
  ["max_store_time"]=>
int(25)
}
}

```

The source distribution of PECL/mysqlnd\_qc contains a directory [web/](#) in which web based monitoring scripts can be found which give an example how to write a cache monitor. Please, follow the instructions given in the source.

Since PECL/mysqlnd\_qc 1.1.0 it is possible to write statistics into a log file. Please, see [mysqlnd\\_qc.collect\\_statistics\\_log\\_file](#).

#### 21.9.7.4.9. Beyond TTL: user-defined storage

Copyright 1997-2012 the PHP Documentation Group. [2906]

The query cache plugin supports the use of user-defined storage handler. User-defined storage handler can use arbitrarily complex invalidation algorithms and support arbitrary storage media.

All user-defined storage handlers have to provide a certain interface. The functions of the user-defined storage handler will be called by the core of the cache plugin. The necessary interface consists of seven public functions. Both procedural and object oriented user-defined storage handler must implement the same set of functions.



**Example 21.318. Using a user-defined storage handler**

```

<?php
/* Enable default caching of all statements */
ini_set("mysqld_qc.cache_by_default", 1);

/* Procedural user defined storage handler functions */

$__cache = array();

function get_hash($host_info, $port, $user, $db, $query) {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    return md5(sprintf("%s%s%s%s", $host_info, $port, $user, $db, $query));
}

function find_query_in_cache($key) {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    if (isset($__cache[$key])) {
        $tmp = $__cache[$key];
        if ($tmp["valid_until"] < time()) {
            unset($__cache[$key]);
            $ret = NULL;
        } else {
            $ret = $__cache[$key]["data"];
        }
    } else {
        $ret = NULL;
    }

    return $ret;
}

function return_to_cache($key) {
    /*
     * Called on cache hit after cached data has been processed,
     * may be used for reference counting
     */
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());
}

function add_query_to_cache_if_not_exists($key, $data, $ttl, $run_time, $store_time, $row_count) {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    $__cache[$key] = array(
        "data" => $data,
        "row_count" => $row_count,
        "valid_until" => time() + $ttl,
        "hits" => 0,
        "run_time" => $run_time,
        "store_time" => $store_time,
        "cached_run_times" => array(),
        "cached_store_times" => array(),
    );

    return TRUE;
}

function query_is_select($query) {
    printf("\t%s('%s'):", __FUNCTION__, $query);

    $ret = FALSE;
    if (strstr($query, "SELECT") !== FALSE) {
        /* cache for 5 seconds */
        $ret = 5;
    }
}

```

```
printf("%s\n", (FALSE === $ret) ? "FALSE" : $ret);
return $ret;
}

function update_query_run_time_stats($key, $run_time, $store_time) {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    if (isset($__cache[$key])) {
        $__cache[$key]['hits']++;
        $__cache[$key]['cached_run_times'][] = $run_time;
        $__cache[$key]['cached_store_times'][] = $store_time;
    }
}

function get_stats($key = NULL) {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    if ($key && isset($__cache[$key])) {
        $stats = $__cache[$key];
    } else {
        $stats = array();
        foreach ($__cache as $key => $details) {
            $stats[$key] = array(
                'hits' => $details['hits'],
                'bytes' => strlen($details['data']),
                'uncached_run_time' => $details['run_time'],
                'cached_run_time' => (count($details['cached_run_times']))
                    ? array_sum($details['cached_run_times']) / count($details['cached_run_times'])
                    : 0,
            );
        }
    }

    return $stats;
}

function clear_cache() {
    global $__cache;
    printf("\t%s(%d)\n", __FUNCTION__, func_num_args());

    $__cache = array();
    return TRUE;
}

/* Install procedural user-defined storage handler */
if (!mysqlnd_qc_set_user_handlers("get_hash", "find_query_in_cache",
    "return_to_cache", "add_query_to_cache_if_not_exists",
    "query_is_select", "update_query_run_time_stats", "get_stats", "clear_cache")) {

    printf("Failed to install user-defined storage handler\n");
}

/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

printf("\nCache put/cache miss\n");

$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

/* Delete record to verify we get our data from the cache */
$mysqli->query("DELETE FROM test WHERE id = 1");

printf("\nCache hit\n");
```

```

$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

printf("\nDisplay cache statistics\n");
var_dump(mysqlnd_qc_get_cache_info());

printf("\nFlushing cache, cache put/cache miss");
var_dump(mysqlnd_qc_clear_cache());

$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();
?>

```

The above examples will output something similar to:

```

query_is_select('DROP TABLE IF EXISTS test'): FALSE
query_is_select('CREATE TABLE test(id INT)'): FALSE
query_is_select('INSERT INTO test(id) VALUES (1), (2)'): FALSE

Cache put/cache miss
query_is_select('SELECT id FROM test WHERE id = 1'): 5
get_hash(5)
find_query_in_cache(1)
add_query_to_cache_if_not_exists(6)
array(1) {
  ["id"]=>
  string(1) "1"
}
query_is_select('DELETE FROM test WHERE id = 1'): FALSE

Cache hit
query_is_select('SELECT id FROM test WHERE id = 1'): 5
get_hash(5)
find_query_in_cache(1)
return_to_cache(1)
update_query_run_time_stats(3)
array(1) {
  ["id"]=>
  string(1) "1"
}

Display cache statistics
get_stats(0)
array(4) {
  ["num_entries"]=>
  int(1)
  ["handler"]=>
  string(4) "user"
  ["handler_version"]=>
  string(5) "1.0.0"
  ["data"]=>
  array(1) {
    ["18683c177dc89bb352b29965d112fdaa"]=>
    array(4) {
      ["hits"]=>
      int(1)
      ["bytes"]=>
      int(71)
      ["uncached_run_time"]=>
      int(398)
      ["cached_run_time"]=>
      int(4)
    }
  }
}

Flushing cache, cache put/cache miss clear_cache(0)

```

```
bool(true)
  query_is_select('SELECT id FROM test WHERE id = 1'): 5
  get_hash(5)
  find_query_in_cache(1)
  add_query_to_cache_if_not_exists(6)
NULL
```

## 21.9.7.5. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

### 21.9.7.5.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

PHP 5.3.3 or a newer version of PHP.

PECL/`mysqlnd_qc` is a `mysqlnd` plugin. It plugs into the `mysqlnd` library. To use you this plugin with a PHP MySQL extension, the extension (`mysqli`, `mysql`, or `PDO_MYSQL`) must enable the `mysqlnd` library.

For using the `APC` storage handler with PECL/`mysqlnd_qc` 1.0 `APC 3.1.3p1-beta` or newer. PECL/`mysqlnd_qc` 1.2 has been tested with `APC 3.1.13-beta`. The `APC` storage handler cannot be used with a shared build. You cannot use the PHP configuration directive `extension` to load the `APC` and PECL/`mysqlnd_qc` extensions if PECL/`mysqlnd_qc` will use `APC` as a storage handler. For using the `APC` storage handler, you have to statically compile PHP with `APC` and PECL/`mysqlnd_qc` support into PHP.

For using `MEMCACHE` storage handler: Use `libmemcache 0.38` or newer. PECL/`mysqlnd_qc` 1.2 has been tested with `libmemcache 1.4.0`.

For using `sqlite` storage handler: Use the `sqlite3` extension that bundled with PHP.

### 21.9.7.5.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

This `PECL` extension is not bundled with PHP.

Information for installing this `PECL` extension may be found in the manual chapter titled [Installation of PECL extensions](#). Additional information such as new releases, downloads, source files, maintainer information, and a `CHANGELOG`, can be located here: [http://pecl.php.net/package/mysqlnd\\_qc](http://pecl.php.net/package/mysqlnd_qc)

A `DLL` for this `PECL` extension is currently unavailable. See also the [building on Windows](#) section.

### 21.9.7.5.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.75. `mysqlnd_qc` Configure Options**

Name	Default	Changeable	Changelog
<code>mysqlnd_qc.enable_qc</code>	1	PHP_INI_SYSTEM	
<code>mysqlnd_qc.ttl</code>	30	PHP_INI_ALL	
<code>mysqlnd_qc.cache_by_default</code>	0	PHP_INI_ALL	
<code>mysqlnd_qc.cache_no_table</code>	0	PHP_INI_ALL	
<code>mysqlnd_qc.use_request_time</code>	0	PHP_INI_ALL	
<code>mysqlnd_qc.time_statistics</code>	1	PHP_INI_ALL	

Name	Default	Changeable	Changelog
<code>mysqlnd_qc.collect_statistics</code>	0	PHP_INI_ALL	
<code>mysqlnd_qc.collect_statistics_log_file</code>	<code>mysqlnd_qc.stats</code>	PHP_INI_SYSTEM	
<code>mysqlnd_qc.collect_query_trace</code>	0	PHP_INI_SYSTEM	
<code>mysqlnd_qc.query_trace_bt_depth</code>	8	PHP_INI_SYSTEM	
<code>mysqlnd_qc.collect_normalized_query_trace</code>	0	PHP_INI_SYSTEM	
<code>mysqlnd_qc.ignore_sql_comments</code>	1	PHP_INI_ALL	
<code>mysqlnd_qc.slam_defense</code>	0	PHP_INI_SYSTEM	
<code>mysqlnd_qc.slam_defense_ttl</code>	30	PHP_INI_SYSTEM	
<code>mysqlnd_qc.std_data_copy</code>	0	PHP_INI_SYSTEM	
<code>mysqlnd_qc.apc_prefix</code>	<code>qc_</code>	PHP_INI_ALL	
<code>mysqlnd_qc.memc_server</code>	127.0.0.1	PHP_INI_ALL	
<code>mysqlnd_qc.memc_port</code>	11211	PHP_INI_ALL	
<code>mysqlnd_qc.sqlite_data_file</code>	memory:	PHP_INI_ALL	

Here's a short explanation of the configuration directives.

`mysqlnd_qc.enable_qc` integer Enables or disables the plugin. If disabled the extension will not plug into `mysqlnd` to proxy internal `mysqlnd` C API calls.

`mysqlnd_qc.ttl` integer Default Time-to-Live (TTL) for cache entries in seconds.

`mysqlnd_qc.cache_by_default` integer Cache all queries regardless if they begin with the SQL hint that enables caching of a query or not. Storage handler cannot overrule the setting. It is evaluated by the core of the plugin.

`mysqlnd_qc.cache_no_table` integer Whether to cache queries with no table name in any of columns meta data of their result set, for example, `SELECT SLEEP(1)`, `SELECT NOW()`, `SELECT SUBSTRING()`.

`mysqlnd_qc.use_request_time` integer Use PHP global request time to avoid `gettimeofday()` system calls? If using `APC` storage handler it should be set to the value of `apc.use_request_time`, if not warnings will be generated.

`mysqlnd_qc.time_statistics` integer Collect run time and store time statistics using `gettimeofday()` system call? Data will be collected only if you also set `mysqlnd_qc.collect_statistics = 1`,

`mysqlnd_qc.collect_statistics_log_file` integer Collect statistics for `mysqlnd_qc_get_core_stats`? Does not influence storage handler statistics! Handler statistics can be an integral part of the handler internal storage format. Therefore, collection of some handler statistics cannot be disabled.

`mysqlnd_qc.collect_statistics_log_file` integer If `mysqlnd_qc.collect_statistics` and `mysqlnd_qc.collect_statistics_log_file` are set, the plugin will dump statistics into the specified log file at every 10th web request during PHP request shutdown. The log file needs to be writable by the web server user.

Since 1.1.0.

`mysqlnd_qc.collect_query_trace` integer Collect query back traces?

`mysqlnd_qc.query_trace_bt_depth` integer Maximum depth/level of a query code backtrace.

<code>mysqlnd_qc.ignore_sql_comments</code> integer	Whether to remove SQL comments from a query string before hashing it to generate a cache key. Disable if you do not want two statements such as <code>SELECT /*my_source_ip=123*/ id FROM test</code> and <code>SELECT /*my_source_ip=456*/ id FROM test</code> to refer to the same cache entry.  Since 1.1.0.
<code>mysqlnd_qc.slam_defense</code> integer	Activates handler based slam defense (cache stampeding protection) if available. Supported by <code>Default</code> and <code>APC</code> storage handler
<code>mysqlnd_qc.slam_defense_ttl</code> integer	TTL for stale cache entries which are served while another client updates the entries. Supported by <code>APC</code> storage handler.
<code>mysqlnd_qc.collect_normalized_queries</code> integer	Collect aggregated normalized query traces? The setting has no effect by default. You compile the extension using the define <code>NORM_QUERY_TRACE_LOG</code> to make use of the setting.
<code>mysqlnd_qc.std_data_copy</code> integer	Default storage handler: copy cached wire data? EXPERIMENTAL – use default setting!
<code>mysqlnd_qc.apc_prefix</code> string	The <code>APC</code> storage handler stores data in the <code>APC</code> user cache. The setting sets a prefix to be used for cache entries.
<code>mysqlnd_qc.memc_server</code> string	<code>MEMCACHE</code> storage handler: memcache server host.
<code>mysqlnd_qc.memc_port</code> integer	<code>MEMCACHE</code> storage handler: memcached server port.
<code>mysqlnd_qc.sqlite_data_file</code> string	<code>sqlite</code> storage handler: data file. Any setting but <code>:memory:</code> may be of little practical value.

### 21.9.7.6. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

*SQL hint related*

#### Example 21.319. Using SQL hint constants

The query cache is controlled by SQL hints. SQL hints are used to enable and disable caching. SQL hints can be used to set the `TTL` of a query.

The SQL hints recognized by the query cache can be manually changed at compile time. This makes it possible to use `mysqlnd_qc` in environments in which the default SQL hints are already taken and interpreted by other systems. Therefore it is recommended to use the SQL hint string constants instead of manually adding the default SQL hints to the query string.

```
<?php
/* Use constants for maximum portability */
$query = "/*" . MYSQLND_QC_ENABLE_SWITCH . "*/SELECT id FROM test";

/* Valid but less portable: default TTL */
$query = "/*qc=on*/SELECT id FROM test";

/* Valid but less portable: per statement TTL */
$query = "/*qc=on/*qc_ttl=5*/SELECT id FROM test";

printf("MYSQLND_QC_ENABLE_SWITCH: %s\n", MYSQLND_QC_ENABLE_SWITCH);
printf("MYSQLND_QC_DISABLE_SWITCH: %s\n", MYSQLND_QC_DISABLE_SWITCH);
```

```
printf("MYSQLND_QC_TTL_SWITCH: %s\n", MYSQLND_QC_TTL_SWITCH);
?>
```

The above examples will output:

```
MYSQLND_QC_ENABLE_SWITCH: qc=on
MYSQLND_QC_DISABLE_SWITCH: qc=off
MYSQLND_QC_TTL_SWITCH: qc_ttl=
```

`MYSQLND_QC_ENABLE_SWITCH` SQL hint used to enable caching of a query.  
(string)

`MYSQLND_QC_DISABLE_SWITCH` SQL hint used to disable caching of a query if  
(string) `mysqLnd_qc.cache_by_default = 1`.

`MYSQLND_QC_TTL_SWITCH` SQL hint used to set the TTL of a result set.  
(string)

`MYSQLND_QC_SERVER_ID_SWITCH` This SQL hint should not be used in general.  
(string)

It is needed by [PECL/mysqLnd\\_ms](#) to group cache entries for one statement but originating from different physical connections. If the hint is used connection settings such as user, hostname and charset are not considered for generating a cache key of a query. Instead the given value and the query string are used as input to the hashing function that generates the key.

`PECL/mysqLnd_ms` may, if instructed, cache results from MySQL Replication slaves. Because it can hold many connections to the slave the cache key shall not be formed from the user, hostname or other settings that may vary for the various slave connections. Instead, `PECL/mysqLnd_ms` provides an identifier which refers to the group of slave connections that shall be enabled to share cache entries no matter which physical slave connection was to generate the cache entry.

Use of this feature outside of `PECL/mysqLnd_ms` is not recommended.

*[mysqLnd\\_qc\\_set\\_cache\\_condition](#) related*

### Example 21.320. Example `mysqLnd_qc_set_cache_condition` usage

The function `mysqLnd_qc_set_cache_condition` allows setting conditions for automatic caching of statements which don't begin with the SQL hints necessary to manually enable caching.

```
<?php
/* Cache all accesses to tables with the name "new%" in schema/database "db_example" for 1 second */
if (!mysqLnd_qc_set_cache_condition(MYSQLND_QC_CONDITION_META_SCHEMA_PATTERN, "db_example.new%", 1)) {
    die("Failed to set cache condition!");
}

$mysqli = new mysqli("host", "user", "password", "db_example", "port");
/* cached although no SQL hint given */
$mysqli->query("SELECT id, title FROM news");

$pdo_mysql = new PDO("mysql:host=host;dbname=db_example;port=port", "user", "password");
/* not cached: no SQL hint, no pattern match */
$pdo_mysql->query("SELECT id, title FROM latest_news");
/* cached: TTL 1 second, pattern match */
```

```
$pdo_mysql->query("SELECT id, title FROM news");
?>
```

`MYSQLND_QC_CONDITION_META` Used as a parameter of `mysqlnd_qc_set_cache_condition` to set conditions for schema based automatic caching.

(int)

*Other*

The plugin version number can be obtained using either `MYSQLND_QC_VERSION`, which is the string representation of the numerical version number, or `MYSQLND_QC_VERSION_ID`, which is an integer such as 10000. Developers can calculate the version number as follows.

Version (part)	Example
Major*10000	1*10000 = 10000
Minor*100	0*100 = 0
Patch	0 = 0
<code>MYSQLND_QC_VERSION_ID</code>	10000

`MYSQLND_QC_VERSION` (string) Plugin version string, for example, "1.0.0-prototype".

`MYSQLND_QC_VERSION_ID` (int) Plugin version number, for example, 10000.

### 21.9.7.7. `mysqlnd_qc` Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.7.7.1. `mysqlnd_qc_clear_cache`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_qc_clear_cache`

Flush all cache contents

#### Description

```
bool mysqlnd_qc_clear_cache();
```

Flush all cache contents.

Flushing the cache is a storage handler responsibility. All built-in storage handler but the `memcache` storage handler support flushing the cache. The `memcache` storage handler cannot flush its cache contents.

User-defined storage handler may or may not support the operation.

#### Parameters

This function has no parameters.

#### Return Values

Returns `TRUE` on success or `FALSE` on failure.

A return value of `FALSE` indicates that flushing all cache contents has failed or the operation is not supported by the active storage handler. Applications must not expect that calling the function will always flush the cache.

#### 21.9.7.7.2. `mysqlnd_qc_get_available_handlers`

Copyright 1997-2012 the PHP Documentation Group. [2906]



- [mysqlnd\\_qc\\_get\\_available\\_handlers](#)

Returns a list of available storage handler

### Description

```
array mysqlnd_qc_get_available_handlers();
```

Which storage are available depends on the compile time configuration of the query cache plugin. The [default](#) storage handler is always available. All other storage handler must be enabled explicitly when building the extension.

### Parameters

This function has no parameters.

### Return Values

Returns an array of available built-in storage handler. For each storage handler the version number and version string is given.

### Examples

#### Example 21.321. [mysqlnd\\_qc\\_get\\_available\\_handlers](#) example

```
<?php
var_dump(mysqlnd_qc_get_available_handlers());
?>
```

The above examples will output:

```
array(5) {
  ["default"]=>
  array(2) {
    ["version"]=>
    string(5) "1.0.0"
    ["version_number"]=>
    int(100000)
  }
  ["user"]=>
  array(2) {
    ["version"]=>
    string(5) "1.0.0"
    ["version_number"]=>
    int(100000)
  }
  ["APC"]=>
  array(2) {
    ["version"]=>
    string(5) "1.0.0"
    ["version_number"]=>
    int(100000)
  }
  ["MEMCACHE"]=>
  array(2) {
    ["version"]=>
    string(5) "1.0.0"
    ["version_number"]=>
    int(100000)
  }
  ["sqlite"]=>
  array(2) {
    ["version"]=>
    string(5) "1.0.0"
    ["version_number"]=>
    int(100000)
  }
}
```

```
}
}
```

**See Also**[Installation](#)[mysqlnd\\_qc\\_set\\_storage\\_handler](#)**21.9.7.7.3. mysqlnd\_qc\_get\_cache\_info**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_qc\\_get\\_cache\\_info](#)

Returns information on the current handler, the number of cache entries and cache entries, if available

**Description**

```
array mysqlnd_qc_get_cache_info();
```

**Parameters**

This function has no parameters.

**Return Values**

Returns information on the current handler, the number of cache entries and cache entries, if available. If and what data will be returned for the cache entries is subject to the active storage handler. Storage handler are free to return any data. Storage handler are recommended to return at least the data provided by the default handler, if technically possible.

The scope of the information is the PHP process. Depending on the PHP deployment model a process may serve one or more web requests.

Values are aggregated for all cache activities on a per storage handler basis. It is not possible to tell how much queries originating from `mysqli`, `PDO_MySQL` or `mysql`.API calls have contributed to the aggregated data values. Use [mysqlnd\\_qc\\_get\\_core\\_stats](#) to get timing data aggregated for all storage handlers.

Array of cache information

<code>handler</code> string	The active storage handler. All storage handler. Since 1.0.0.
<code>handler_version</code> string	The version of the active storage handler. All storage handler. Since 1.0.0.
<code>num_entries</code> int	The number of cache entries. The value depends on the storage handler in use.  The default, APC and SQLite storage handler provide the actual number of cache entries.  The MEMCACHE storage handler always returns 0. MEMCACHE does not support counting the number of cache entries.  If a user defined handler is used, the number of entries of the <code>data</code> property is reported.

Since 1.0.0.

`data` array

The version of the active storage handler.

Additional storage handler dependent data on the cache entries. Storage handler are requested to provide similar and comparable information. A user defined storage handler is free to return any data.

Since 1.0.0.

The following information is provided by the default storage handler for the `data` property.

The `data` property holds a hash. The hash is indexed by the internal cache entry identifier of the storage handler. The cache entry identifier is human-readable and contains the query string leading to the cache entry. Please, see also the example below. The following data is given for every cache entry.

`statistics` array

Statistics of the cache entry.

Since 1.0.0.

Property	Description	Version
<code>rows</code>	Number of rows of the cached result set.	Since 1.0.0.
<code>storage_size</code>	The size of the cached result set in bytes. This is the size of the payload. The value is not suited for calculating the total memory consumption of all cache entries including the administrative overhead of the cache entries.	Since 1.0.0.
<code>cache_hits</code>	How often the cached entry has been returned.	Since 1.0.0.
<code>runtime</code>	Run time of the statement to which the cache entry belongs. This is the run time of the uncached statement. It is the time between sending the statement to MySQL receiving a reply from MySQL. Run time saved by using the query cache plugin can be calculated like this: <code>cache_hits</code>	Since 1.0.0.

Property	Description	Version
	<code>* ((run_time - avg_run_time) + (store_time - avg_store_time)).</code>	
<code>store_time</code>	Store time of the statements result set to which the cache entry belongs. This is the time it took to fetch and store the results of the uncached statement.	Since 1.0.0.
<code>min_run_time</code>	Minimum run time of the cached statement. How long it took to find the statement in the cache.	Since 1.0.0.
<code>min_store_time</code>	Minimum store time of the cached statement. The time taken for fetching the cached result set from the storage medium and decoding	Since 1.0.0.
<code>avg_run_time</code>	Average run time of the cached statement.	Since 1.0.0.
<code>avg_store_time</code>	Average store time of the cached statement.	Since 1.0.0.
<code>max_run_time</code>	Average run time of the cached statement.	Since 1.0.0.
<code>max_store_time</code>	Average store time of the cached statement.	Since 1.0.0.
<code>validity</code>	Timestamp when the cache entry expires.	Since 1.1.0.

`metadata` array

Metadata of the cache entry. This is the metadata provided by MySQL together with the result set of the statement in question. Different versions of the MySQL server may return different metadata. Unlike with some of the PHP MySQL extensions no attempt is made to hide MySQL server version dependencies and version details from the caller. Please, refer to the MySQL C API documentation that belongs to the MySQL server in use for further details.

The metadata list contains one entry for every column.

Since 1.0.0.

Property	Description	Version
<code>name</code>	The field name. Depending on the MySQL version this may be the fields alias name.	Since 1.0.0.
<code>org_name</code>	The field name.	Since 1.0.0.
<code>table</code>	The table name. If an alias name was used for the table, this usually holds the alias name.	Since 1.0.0.
<code>org_table</code>	The table name.	Since 1.0.0.
<code>db</code>	The database/schema name.	Since 1.0.0.
<code>max_length</code>	The maximum width of the field. Details may vary by MySQL server version.	Since 1.0.0.
<code>length</code>	The width of the field. Details may vary by MySQL server version.	Since 1.0.0.
<code>type</code>	The data type of the field. Details may vary by the MySQL server in use. This is the MySQL C API type constants value. It is recommended to use type constants provided by the <code>mysqli</code> extension to test for its meaning. You should not test for certain type values by comparing with certain numbers.	Since 1.0.0.

The APC storage handler returns the same information for the `data` property but no `metadata`. The `metadata` of a cache entry is set to `NULL`.

The MEMCACHE storage handler does not fill the `data` property. Statistics are not available on a per cache entry basis with the MEMCACHE storage handler.

A user defined storage handler is free to provide any data.

## Examples

### Example 21.322. `mysqld_qc_get_cache_info` example

The example shows the output from the built-in default storage handler. Other storage handler may report different data.

```
<?php
/* Populate the cache, e.g. using mysqli */
$mysqli = new mysqli("host", "user", "password", "schema");
$mysqli->query("/*" . MYSQLND_QC_ENABLE_SWITCH . "*/SELECT id FROM test");

/* Display cache information */
var_dump(mysqld_qc_get_cache_info());
?>
```

The above examples will output:

```
array(4) {
  ["num_entries"]=>
  int(1)
  ["handler"]=>
  string(7) "default"
  ["handler_version"]=>
  string(5) "1.0.0"
  ["data"]=>
  array(1) {
    ["Localhost via UNIX socket 3306 user schema/*qc=on*/SELECT id FROM test"]=>
    array(2) {
      ["statistics"]=>
      array(11) {
        ["rows"]=>
        int(6)
        ["stored_size"]=>
        int(101)
        ["cache_hits"]=>
        int(0)
        ["run_time"]=>
        int(471)
        ["store_time"]=>
        int(27)
        ["min_run_time"]=>
        int(0)
        ["max_run_time"]=>
        int(0)
        ["min_store_time"]=>
        int(0)
        ["max_store_time"]=>
        int(0)
        ["avg_run_time"]=>
        int(0)
        ["avg_store_time"]=>
        int(0)
      }
      ["metadata"]=>
      array(1) {
        [0]=>
        array(8) {
          ["name"]=>
          string(2) "id"
          ["orig_name"]=>
          string(2) "id"
          ["table"]=>
          string(4) "test"
          ["orig_table"]=>
          string(4) "test"
          ["db"]=>
          string(4) "schema"
          ["max_length"]=>
          int(1)
        }
      }
    }
  }
}
```



Statistic	Description	Version
<code>cache_miss</code>	<p>Statement is considered cacheable...</p> <ul style="list-style-type: none"> <li>• ... and has been added to the cache</li> <li>• ... but the <code>PHP</code> configuration directive setting of <code>mysqlnd_qc.cache_no_table = 1</code> has prevented caching.</li> <li>• ... but an unbuffered result set is requested.</li> <li>• ... but a buffered result set was empty.</li> </ul>	Since 1.0.0.
<code>cache_put</code>	<p>Statement is considered cacheable and has been added to the cache. Take care when calculating derived statistics. Storage handler with a storage life time beyond process scope may report <code>cache_put = 0</code> together with <code>cache_hit &gt; 0</code>, if another process has filled the cache. You may want to use <code>num_entries</code> from <code>mysqlnd_qc_get_cache_info</code> if the handler supports it ( <code>default</code>, <code>APC</code>).</p>	Since 1.0.0.
<code>query_should_cache</code>	<p>Statement is considered cacheable based on query string analysis. The statement may or may not be added to the cache. See also <code>cache_put</code>.</p>	Since 1.0.0.
<code>query_should_not_cache</code>	<p>Statement is considered not cacheable based on query string analysis.</p>	Since 1.0.0.
<code>query_not_cached</code>	<p>Statement is considered not cacheable or it is considered cacheable but the storage handler has not returned a hash key for it.</p>	Since 1.0.0.
<code>query_could_cache</code>	<p>Statement is considered cacheable...</p> <ul style="list-style-type: none"> <li>• ... and statement has been run without errors</li> <li>• ... and meta data shows at least one column in the result set</li> </ul> <p>The statement may or may not be in the cache already. It may or</p>	Since 1.0.0.



Statistic	Description	Version
	may not be added to the cache later on.	
<code>query_found_in_cache</code>	Statement is considered cacheable and we have found it in the cache but we have not replayed the cached data yet and we have not send the result set to the client yet. This is not considered a cache hit because the client might not fetch the result or the cached data may be faulty.	Since 1.0.0.
<code>query_uncached_other</code>	Statement is considered cacheable and it may or may not be in the cache already but either replaying cached data has failed, no result set is available or some other error has happened.	
<code>query_uncached_no_table</code>	Statement has not been cached because the result set has at least one column which has no table name in its meta data. An example of such a query is <code>SELECT SLEEP(1)</code> . To cache those statements you have to change default value of the PHP configuration directive <code>mysqld_qc.cache_no_table</code> and set <code>mysqld_qc.cache_no_table = 1</code> . Often, it is not desired to cache such statements.	Since 1.0.0.
<code>query_uncached_use_result</code>	Statement would have been cached if a buffered result set had been used. The situation is also considered as a cache miss and <code>cache_miss</code> will be incremented as well.	Since 1.0.0.
<code>query_aggr_run_time_cache</code>	Aggregated run time (ms) of all cached queries. Cached queries are those which have incremented <code>cache_hit</code> .	Since 1.0.0.
<code>query_aggr_run_time_cache</code>	Aggregated run time (ms) of all uncached queries that have been put into the cache. See also <code>cache_put</code> .	Since 1.0.0.
<code>query_aggr_run_time_total</code>	Aggregated run time (ms) of all uncached and cached queries that have been inspected and executed by the query cache.	Since 1.0.0.
<code>query_aggr_store_time_cache</code>	Aggregated store time (ms) of all cached queries. Cached	Since 1.0.0.

Statistic	Description	Version
	queries are those which have incremented <code>cache_hit</code> .	
<code>query_aggr_store_time_cache</code>	Aggregated store time (ms) of all uncached queries that have been put into the cache. See also <code>cache_put</code> .	Since 1.0.0.
<code>query_aggr_store_time_total</code>	Aggregated store time (ms) of all uncached and cached queries that have been inspected and executed by the query cache.	Since 1.0.0.
<code>receive_bytes_recorded</code>	Recorded incoming network traffic (bytes) send from MySQL to PHP. The traffic may or may not have been added to the cache. The traffic is the total for all queries regardless if cached or not.	Since 1.0.0.
<code>receive_bytes_replayed</code>	Network traffic replayed during cache. This is the total amount of incoming traffic saved because of the usage of the query cache plugin.	Since 1.0.0.
<code>send_bytes_recorded</code>	Recorded outgoing network traffic (bytes) send from MySQL to PHP. The traffic may or may not have been added to the cache. The traffic is the total for all queries regardless if cached or not.	Since 1.0.0.
<code>send_bytes_replayed</code>	Network traffic replayed during cache. This is the total amount of outgoing traffic saved because of the usage of the query cache plugin.	Since 1.0.0.
<code>slam_stale_refresh</code>	Number of cache misses which triggered serving stale data until the client causing the cache miss has refreshed the cache entry.	Since 1.0.0.
<code>slam_stale_hit</code>	Number of cache hits while a stale cache entry gets refreshed.	Since 1.0.0.

## Examples

### Example 21.323. `mysqlnd_qc_get_core_stats` example

```
<?php
/* Enable collection of statistics - default: disabled */
ini_set("mysqlnd_qc.collect_statistics", 1);

/* Enable collection of all timing related statistics -
default: enabled but overruled by mysqlnd_qc.collect_statistics = 0 */
ini_set("mysqlnd_qc.collect_time_statistics", 1);

/* Populate the cache, e.g. using mysqli */
$mysqli = new mysqli('host', 'user', 'password', 'schema');
```

```

/* Cache miss and cache put */
$mysqlcli->query("/*qc=on*/SELECT id FROM test");
/* Cache hit */
$mysqlcli->query("/*qc=on*/SELECT id FROM test");

/* Display core statistics */
var_dump(mysqld_qc_get_core_stats());
?>

```

The above examples will output:

```

array(26) {
  ["cache_hit"]=>
  string(1) "1"
  ["cache_miss"]=>
  string(1) "1"
  ["cache_put"]=>
  string(1) "1"
  ["query_should_cache"]=>
  string(1) "2"
  ["query_should_not_cache"]=>
  string(1) "0"
  ["query_not_cached"]=>
  string(1) "0"
  ["query_could_cache"]=>
  string(1) "2"
  ["query_found_in_cache"]=>
  string(1) "1"
  ["query_uncached_other"]=>
  string(1) "0"
  ["query_uncached_no_table"]=>
  string(1) "0"
  ["query_uncached_no_result"]=>
  string(1) "0"
  ["query_uncached_use_result"]=>
  string(1) "0"
  ["query_aggr_run_time_cache_hit"]=>
  string(1) "4"
  ["query_aggr_run_time_cache_put"]=>
  string(3) "395"
  ["query_aggr_run_time_total"]=>
  string(3) "399"
  ["query_aggr_store_time_cache_hit"]=>
  string(1) "2"
  ["query_aggr_store_time_cache_put"]=>
  string(1) "8"
  ["query_aggr_store_time_total"]=>
  string(2) "10"
  ["receive_bytes_recorded"]=>
  string(2) "65"
  ["receive_bytes_replayed"]=>
  string(2) "65"
  ["send_bytes_recorded"]=>
  string(2) "29"
  ["send_bytes_replayed"]=>
  string(2) "29"
  ["slam_stale_refresh"]=>
  string(1) "0"
  ["slam_stale_hit"]=>
  string(1) "0"
  ["request_counter"]=>
  int(1)
  ["process_hash"]=>
  int(3547549858)
}

```

**See Also**

- [Runtime configuration](#)
- [mysqlnd\\_qc.collect\\_statistics](#)
- [mysqlnd\\_qc.time\\_statistics](#)
- [mysqlnd\\_qc\\_get\\_cache\\_info](#)

**21.9.7.7.5. mysqlnd\_qc\_get\_normalized\_query\_trace\_log**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_qc\\_get\\_normalized\\_query\\_trace\\_log](#)

Returns a normalized query trace log for each query inspected by the query cache

**Description**

```
array mysqlnd_qc_get_normalized_query_trace_log();
```

Returns a normalized query trace log for each query inspected by the query cache. The collection of the trace log is disabled by default. To collect the trace log you have to set the PHP configuration directive `mysqlnd_qc.collect_normalized_query_trace` to `1`

Entries in the trace log are grouped by the normalized query statement. The normalized query statement is the query statement with all statement parameter values being replaced with a question mark. For example, the two statements `SELECT id FROM test WHERE id = 1` and `SELECT id FROM test WHERE id = 2` are normalized as `SELECT id FROM test WHERE id = ?`. Whenever a statement is inspected by the query cache which matches the normalized statement pattern, its statistics are grouped by the normalized statement string.

**Parameters**

This function has no parameters.

**Return Values**

An array of query log. Every list entry contains the normalized query string and further detail information.

Key	Description
<code>query</code>	Normalized statement string.
<code>occurrences</code>	How many statements have matched the normalized statement string in addition to the one which has created the log entry. The value is zero if a statement has been normalized, its normalized representation has been added to the log but no further queries inspected by PECL/mysqlnd_qc have the same normalized statement string.
<code>eligible_for_cache</code>	Whether the statement could be cached. An statement eligible for caching has not necessarily been cached. It not possible to tell for sure if or how many cached statement have contributed to the aggregated normalized statement log entry. However, comparing the minimum and average run time one can make an educated guess.
<code>avg_run_time</code>	The average run time of all queries contributing to the query log entry. The run time is the time between sending the query statement to MySQL and receiving an answer from MySQL.
<code>avg_store_time</code>	The average store time of all queries contributing to the query log entry. The store time is the time needed to fetch a statements result set from the server to the client and, storing it on the client.
<code>min_run_time</code>	The minimum run time of all queries contributing to the query log entry.
<code>min_store_time</code>	The minimum store time of all queries contributing to the query log entry.
<code>max_run_time</code>	The maximum run time of all queries contributing to the query log entry.
<code>max_store_time</code>	The maximum store time of all queries contributing to the query log entry.

## Examples

### Example 21.324. [mysqlnd\\_qc\\_get\\_normalized\\_query\\_trace\\_log](#) example

```
mysqlnd_qc.collect_normalized_query_trace=1
```

```
<?php
/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

/* not cached */
$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

/* cache put */
$res = $mysqli->query("/*" . MYSQLND_QC_ENABLE_SWITCH . "*/" . "SELECT id FROM test WHERE id = 2");
var_dump($res->fetch_assoc());
$res->free();

/* cache hit */
$res = $mysqli->query("/*" . MYSQLND_QC_ENABLE_SWITCH . "*/" . "SELECT id FROM test WHERE id = 2");
var_dump($res->fetch_assoc());
$res->free();

var_dump(mysqlnd_qc_get_normalized_query_trace_log());
?>
```

The above examples will output:

```
array(1) {
  ["id"]=>
  string(1) "1"
}
array(1) {
  ["id"]=>
  string(1) "2"
}
array(1) {
  ["id"]=>
  string(1) "2"
}
array(4) {
  [0]=>
  array(9) {
    ["query"]=>
    string(25) "DROP TABLE IF EXISTS test"
    ["occurrences"]=>
    int(0)
    ["eligible_for_caching"]=>
    bool(false)
    ["avg_run_time"]=>
    int(0)
    ["min_run_time"]=>
    int(0)
    ["max_run_time"]=>
    int(0)
    ["avg_store_time"]=>
    int(0)
    ["min_store_time"]=>
    int(0)
    ["max_store_time"]=>
    int(0)
  }
}
```

```
}
[1]=>
array(9) {
  ["query"]=>
  string(27) "CREATE TABLE test (id INT )"
  ["occurrences"]=>
  int(0)
  ["eligible_for_caching"]=>
  bool(false)
  ["avg_run_time"]=>
  int(0)
  ["min_run_time"]=>
  int(0)
  ["max_run_time"]=>
  int(0)
  ["avg_store_time"]=>
  int(0)
  ["min_store_time"]=>
  int(0)
  ["max_store_time"]=>
  int(0)
}
[2]=>
array(9) {
  ["query"]=>
  string(40) "INSERT INTO test (id ) VALUES ( ? ), ( ? )"
  ["occurrences"]=>
  int(0)
  ["eligible_for_caching"]=>
  bool(false)
  ["avg_run_time"]=>
  int(0)
  ["min_run_time"]=>
  int(0)
  ["max_run_time"]=>
  int(0)
  ["avg_store_time"]=>
  int(0)
  ["min_store_time"]=>
  int(0)
  ["max_store_time"]=>
  int(0)
}
[3]=>
array(9) {
  ["query"]=>
  string(31) "SELECT id FROM test WHERE id =?"
  ["occurrences"]=>
  int(2)
  ["eligible_for_caching"]=>
  bool(true)
  ["avg_run_time"]=>
  int(159)
  ["min_run_time"]=>
  int(12)
  ["max_run_time"]=>
  int(307)
  ["avg_store_time"]=>
  int(10)
  ["min_store_time"]=>
  int(8)
  ["max_store_time"]=>
  int(13)
}
}
```

**See Also**

[Runtime configuration](#)  
[mysqld\\_qc.collect\\_normalized\\_query\\_trace](#)

[mysqlnd\\_qc.time\\_statistics](#)  
[mysqlnd\\_qc\\_get\\_query\\_trace\\_log](#)

### 21.9.7.7.6. `mysqlnd_qc_get_query_trace_log`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_qc_get_query_trace_log`

Returns a backtrace for each query inspected by the query cache

#### Description

```
array mysqlnd_qc_get_query_trace_log();
```

Returns a backtrace for each query inspected by the query cache. The collection of the backtrace is disabled by default. To collect the backtrace you have to set the PHP configuration directive `mysqlnd_qc.collect_query_trace` to `1`

The maximum depth of the backtrace is limited to the depth set with the PHP configuration directive `mysqlnd_qc.query_trace_bt_depth`.

#### Parameters

This function has no parameters.

#### Return Values

An array of query backtrace. Every list entry contains the query string, a backtrace and further detail information.

Key	Description
<code>query</code>	Query string.
<code>origin</code>	Code backtrace.
<code>run_time</code>	Query run time in milliseconds. The collection of all times and the necessary <code>gettimeofday</code> system calls can be disabled by setting the PHP configuration directive <code>mysqlnd_qc.time_statistics</code> to <code>0</code>
<code>store_time</code>	Query result set store time in milliseconds. The collection of all times and the necessary <code>gettimeofday</code> system calls can be disabled by setting the PHP configuration directive <code>mysqlnd_qc.time_statistics</code> to <code>0</code>
<code>eligible_for_cache</code>	<code>TRUE</code> if query is cacheable otherwise <code>FALSE</code> .
<code>no_table_name</code>	<code>TRUE</code> if the query has generated a result set and at least one column from the result set has no table name set in its metadata. This is usually the case with queries which one probably do not want to cache such as <code>SELECT SLEEP(1)</code> . By default any such query will not be added to the cache. See also PHP configuration directive <code>mysqlnd_qc.cache_no_table</code> .
<code>was_added</code>	<code>TRUE</code> if the query result has been put into the cache, otherwise <code>FALSE</code> .
<code>was_already_in_cache</code>	<code>TRUE</code> if the query result would have been added to the cache if it was not already in the cache (cache hit). Otherwise <code>FALSE</code> .

#### Examples

##### Example 21.325. `mysqlnd_qc_get_query_trace_log` example

```
mysqlnd_qc.collect_query_trace=1
```

```
<?php
/* Connect, create and populate test table */
```

```

$mysqli = new mysqli("host", "user", "password", "schema", "port", "socket");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2)");

/* not cached */
$res = $mysqli->query("SELECT id FROM test WHERE id = 1");
var_dump($res->fetch_assoc());
$res->free();

/* cache put */
$res = $mysqli->query("/* . MYSQLND_QC_ENABLE_SWITCH . */" . "SELECT id FROM test WHERE id = 2");
var_dump($res->fetch_assoc());
$res->free();

/* cache hit */
$res = $mysqli->query("/* . MYSQLND_QC_ENABLE_SWITCH . */" . "SELECT id FROM test WHERE id = 2");
var_dump($res->fetch_assoc());
$res->free();

var_dump(mysqlnd_qc_get_query_trace_log());
?>

```

The above examples will output:

```

array(1) {
  ["id"]=>
  string(1) "1"
}
array(1) {
  ["id"]=>
  string(1) "2"
}
array(1) {
  ["id"]=>
  string(1) "2"
}
array(6) {
  [0]=>
  array(8) {
    ["query"]=>
    string(25) "DROP TABLE IF EXISTS test"
    ["origin"]=>
    string(102) "#0 qc.php(4): mysqli->query('DROP TABLE IF E...')
#1 {main}"
    ["run_time"]=>
    int(0)
    ["store_time"]=>
    int(0)
    ["eligible_for_caching"]=>
    bool(false)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(false)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [1]=>
  array(8) {
    ["query"]=>
    string(25) "CREATE TABLE test(id INT)"
    ["origin"]=>
    string(102) "#0 qc.php(5): mysqli->query('CREATE TABLE te...')
#1 {main}"
    ["run_time"]=>
    int(0)
    ["store_time"]=>
    int(0)
  }
}

```



```
    ["eligible_for_caching"]=>
    bool(false)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(false)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [2]=>
  array(8) {
    ["query"]=>
    string(36) "INSERT INTO test(id) VALUES (1), (2)"
    ["origin"]=>
    string(102) "#0 qc.php(6): mysqli->query('INSERT INTO tes...')
#1 {main}"
    ["run_time"]=>
    int(0)
    ["store_time"]=>
    int(0)
    ["eligible_for_caching"]=>
    bool(false)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(false)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [3]=>
  array(8) {
    ["query"]=>
    string(32) "SELECT id FROM test WHERE id = 1"
    ["origin"]=>
    string(102) "#0 qc.php(9): mysqli->query('SELECT id FROM ...')
#1 {main}"
    ["run_time"]=>
    int(0)
    ["store_time"]=>
    int(25)
    ["eligible_for_caching"]=>
    bool(false)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(false)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [4]=>
  array(8) {
    ["query"]=>
    string(41) "/*qc=on*/SELECT id FROM test WHERE id = 2"
    ["origin"]=>
    string(103) "#0 qc.php(14): mysqli->query('/*qc=on*/SELECT...')
#1 {main}"
    ["run_time"]=>
    int(311)
    ["store_time"]=>
    int(13)
    ["eligible_for_caching"]=>
    bool(true)
    ["no_table"]=>
    bool(false)
    ["was_added"]=>
    bool(true)
    ["was_already_in_cache"]=>
    bool(false)
  }
  [5]=>
  array(8) {
    ["query"]=>
```

```
string(41) "/*qc=on*/SELECT id FROM test WHERE id = 2"
["origin"]=>
string(103) "#0 qc.php(19): mysqli->query('/*qc=on*/SELECT...')
#1 {main}"
  ["run_time"]=>
  int(13)
  ["store_time"]=>
  int(8)
  ["eligible_for_caching"]=>
  bool(true)
  ["no_table"]=>
  bool(false)
  ["was_added"]=>
  bool(false)
  ["was_already_in_cache"]=>
  bool(true)
}
```

### See Also

[Runtime configuration](#)  
[mysqlnd\\_qc.collect\\_query\\_trace](#)  
[mysqlnd\\_qc.query\\_trace\\_bt\\_depth](#)  
[mysqlnd\\_qc.time\\_statistics](#)  
[mysqlnd\\_qc.cache\\_no\\_table](#)  
[mysqlnd\\_qc\\_get\\_normalized\\_query\\_trace\\_log](#)

#### 21.9.7.7.7. [mysqlnd\\_qc\\_set\\_cache\\_condition](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_qc\\_set\\_cache\\_condition](#)

Set conditions for automatic caching

### Description

```
bool mysqlnd_qc_set_cache_condition(
    int condition_type,
    mixed condition,
    mixed condition_option);
```

Sets a condition for automatic caching of statements which do not contain the necessary SQL hints to enable caching of them.

### Parameters

*condition\_type*

Type of the condition. The only allowed value is `MYSQLND_QC_CONDITION_META_SCHEMA_PATTERN`.

*condition*

Parameter for the condition set with `condition_type`. Parameter type and structure depend on `condition_type`

If `condition_type` equals `MYSQLND_QC_CONDITION_META_SCHEMA_PATTERN` `condition` must be a string. The string sets a pattern. Statements are cached if table and database meta data entry of their result sets match the pattern. The pattern is checked for a match with the `db` and `org_table` meta data entries provided by the underlying MySQL client server library. Please, check the MySQL Reference manual for details about the two entries. The `db` and `org_table` values are concatenated with a dot (.) before matched against `condition`. Pattern matching supports the wildcards `%` and `_`. The wildcard `%` will

match one or many arbitrary characters. `_` will match one arbitrary character. The escape symbol is backslash.

`condition_option`

Option for `condition`. Type and structure depend on `condition_type`.

If `condition_type` equals `MYSQLND_QC_CONDITION_META_SCHEMA_PATTERN` `condition_options` is the TTL to be used.

## Examples

### Example 21.326. `mysqlnd_qc_set_cache_condition` example

```
<?php
/* Cache all accesses to tables with the name "new%" in schema/database "db_example" for 1 second */
if (!mysqlnd_qc_set_cache_condition(MYSQLND_QC_CONDITION_META_SCHEMA_PATTERN, "db_example.new%", 1)) {
    die("Failed to set cache condition!");
}

$mysqli = new mysqli("host", "user", "password", "db_example", "port");
/* cached although no SQL hint given */
$mysqli->query("SELECT id, title FROM news");

$pdo_mysql = new PDO("mysql:host=host;dbname=db_example;port=port", "user", "password");
/* not cached: no SQL hint, no pattern match */
$pdo_mysql->query("SELECT id, title FROM latest_news");
/* cached: TTL 1 second, pattern match */
$pdo_mysql->query("SELECT id, title FROM news");
?>
```

## Return Values

Returns TRUE on success or FALSE on FAILURE.

## See Also

[Quickstart: pattern based caching](#)

### 21.9.7.7.8. `mysqlnd_qc_set_is_select`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_qc_set_is_select`

Installs a callback which decides whether a statement is cached

## Description

```
mixed mysqlnd_qc_set_is_select(
    string callback);
```

Installs a callback which decides whether a statement is cached.

There are several ways of hinting PECL/`mysqlnd_qc` to cache a query. By default, PECL/`mysqlnd_qc` attempts to cache a if caching of all statements is enabled or the query string begins with a certain SQL hint. The plugin internally calls a function named `is_select()` to find out. This internal function can be replaced with a user-defined callback. Then, the user-defined callback is responsible to decide whether the plugin attempts to cache a statement. Because the internal function is replaced with the callback, the callback gains full control. The callback is free to ignore the configuration setting `mysqlnd_qc.cache_by_default` and SQL hints.

The callback is invoked for every statement inspected by the plugin. It is given the statements string as a parameter. The callback returns `FALSE` if the statement shall not be cached. It returns `TRUE` to

make the plugin attempt to cache the statements result set, if any. A so-created cache entry is given the default TTL set with the PHP configuration directive `mysqlnd_qc.ttl`. If a different TTL shall be used, the callback returns a numeric value to be used as the TTL.

The internal `is_select` function is part of the internal cache storage handler interface. Thus, a user-defined storage handler offers the same capabilities.

### Parameters

This function has no parameters.

### Return Values

Returns `TRUE` on success or `FALSE` on failure.

### Examples

#### Example 21.327. `mysqlnd_qc_set_is_select` example

```
<?php
/* callback which decides if query is cached */
function is_select($query) {
    static $patterns = array(
        /* true - use default from mysqlnd_qc.ttl */
        "@SELECT\s+.*\s+FROM\s+test@ismU" => true,
        /* 3 - use TTL = 3 seconds */
        "@SELECT\s+.*\s+FROM\s+news@ismU" => 3
    );
    /* check if query does match pattern */
    foreach ($patterns as $pattern => $ttl) {
        if (preg_match($pattern, $query)) {
            printf("is_select(%45s): cache\n", $query);
            return $ttl;
        }
    }
    printf("is_select(%45s): do not cache\n", $query);
    return false;
}
mysqlnd_qc_set_is_select("is_select");

/* Connect, create and populate test table */
$mysqli = new mysqli("host", "user", "password", "schema");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1), (2), (3)");

/* cache put */
$mysqli->query("SELECT id FROM test WHERE id = 1");
/* cache hit */
$mysqli->query("SELECT id FROM test WHERE id = 1");
/* cache put */
$mysqli->query("SELECT * FROM test");
?>
```

The above examples will output:

```
is_select(                DROP TABLE IF EXISTS test): do not cache
is_select(                CREATE TABLE test(id INT)): do not cache
is_select(    INSERT INTO test(id) VALUES (1), (2), (3)): do not cache
is_select(                SELECT id FROM test WHERE id = 1): cache
is_select(                SELECT id FROM test WHERE id = 1): cache
is_select(                SELECT * FROM test): cache
```

### See Also

[Runtime configuration](#)  
[mysqlnd\\_qc.ttl](#)  
[mysqlnd\\_qc.cache\\_by\\_default](#)  
[mysqlnd\\_qc.set\\_user\\_handlers](#)

### 21.9.7.7.9. [mysqlnd\\_qc\\_set\\_storage\\_handler](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_qc\\_set\\_storage\\_handler](#)

Change current storage handler

#### Description

```
bool mysqlnd_qc_set_storage_handler(
    string handler);
```

Sets the storage handler used by the query cache. A list of available storage handler can be obtained from [mysqlnd\\_qc\\_get\\_available\\_handlers](#). Which storage are available depends on the compile time configuration of the query cache plugin. The `default` storage handler is always available. All other storage handler must be enabled explicitly when building the extension.

#### Parameters

*handler* Handler can be of type string representing the name of a built-in storage handler or an object of type [mysqlnd\\_qc\\_handler\\_default](#). The names of the built-in storage handler are `default`, `APC`, `MEMCACHE`, `sqlite`.

#### Return Values

Returns `TRUE` on success or `FALSE` on failure.

If changing the storage handler fails a catchable fatal error will be thrown. The query cache cannot operate if the previous storage handler has been shutdown but no new storage handler has been installed.

#### Examples

##### Example 21.328. [mysqlnd\\_qc\\_set\\_storage\\_handler](#) example

The example shows the output from the built-in default storage handler. Other storage handler may report different data.

```
<?php
var_dump(mysqlnd_qc_set_storage_handler("memcache"));

if (true === mysqlnd_qc_set_storage_handler("default"))
    printf("Default storage handler activated");

/* Catchable fatal error */
var_dump(mysqlnd_qc_set_storage_handler("unknown"));
?>
```

The above examples will output:

```
bool(true)
Default storage handler activated
Catchable fatal error: mysqlnd_qc_set_storage_handler(): Unknown handler 'unknown' in (file) on line (1
```

**See Also**[Installation](#)[mysqlnd\\_qc\\_get\\_available\\_handlers](#)**21.9.7.7.10. mysqlnd\_qc\_set\_user\_handlers**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_qc\\_set\\_user\\_handlers](#)

Sets the callback functions for a user-defined procedural storage handler

**Description**

```
bool mysqlnd_qc_set_user_handlers(  
    string get_hash,  
    string find_query_in_cache,  
    string return_to_cache,  
    string add_query_to_cache_if_not_exists,  
    string query_is_select,  
    string update_query_run_time_stats,  
    string get_stats,  
    string clear_cache);
```

Sets the callback functions for a user-defined procedural storage handler.

**Parameters**

<a href="#">get_hash</a>	Name of the user function implementing the storage handler <a href="#">get_hash</a> functionality.
<a href="#">find_query_in_cache</a>	Name of the user function implementing the storage handler <a href="#">find_in_cache</a> functionality.
<a href="#">return_to_cache</a>	Name of the user function implementing the storage handler <a href="#">return_to_cache</a> functionality.
<a href="#">add_query_to_cache_if_not_exists</a>	Name of the user function implementing the storage handler <a href="#">add_query_to_cache_if_not_exists</a> functionality.
<a href="#">query_is_select</a>	Name of the user function implementing the storage handler <a href="#">query_is_select</a> functionality.
<a href="#">update_query_run_time_stats</a>	Name of the user function implementing the storage handler <a href="#">update_query_run_time_stats</a> functionality.
<a href="#">get_stats</a>	Name of the user function implementing the storage handler <a href="#">get_stats</a> functionality.
<a href="#">clear_cache</a>	Name of the user function implementing the storage handler <a href="#">clear_cache</a> functionality.

**Return Values**

Returns TRUE on success or FALSE on FAILURE.

**See Also**[Procedural user-defined storage handler example](#)**21.9.7.8. Change History**

Copyright 1997-2012 the PHP Documentation Group. [2906]

This change history is a high level summary of selected changes that may impact applications and/or break backwards compatibility.

See also the [CHANGES](#) file in the source distribution for a complete list of changes.

#### 21.9.7.8.1. PECL/mysqlnd\_qc 1.2 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

1.2.0 - alpha

- Release date: 03/2013
- Motto/theme: PHP 5.5 compatibility

Feature changes

- Update build for PHP 5.5 (Credits: Remi Collet)
- APC storage handler update
  - Fix build for APC 3.1.13-beta and trunk
- Introduced [MYSQLND\\_QC\\_VERSION](#) and [MYSQLND\\_QC\\_VERSION\\_ID](#).

#### 21.9.7.8.2. PECL/mysqlnd\_qc 1.1 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

1.1.0 - stable

- Release date: 04/2012
- Motto/theme: PHP 5.4 compatibility, schema pattern based caching and mysqlnd\_ms support

1.1.0 - beta

- Release date: 04/2012
- Motto/theme: PHP 5.4 compatibility, schema pattern based caching and mysqlnd\_ms support

1.1.0 - alpha

- Release date: 04/2012
- Motto/theme: PHP 5.4 compatibility, schema pattern based caching and mysqlnd\_ms support

Feature changes

- APC storage handler update
  - Fix build for APC 3.1.9+
  - Note: Use of the APC storage handler is currently not recommended due to stability issues of APC itself.
- New PHP configuration directives
  - [mysqlnd\\_qc.collect\\_statistics\\_log\\_file](#). Aggregated cache statistics log file written after every 10th request served by the PHP process.
  - [mysqlnd\\_qc.ignore\\_sql\\_comments](#). Control whether SQL comments are ignored for cache key hash generation.

- New constants and SQL hints
  - [MYSQLND\\_QC\\_SERVER\\_ID\\_SWITCH](#) allows grouping of cache entries from different physical connections. This is needed by PECL/mysqlnd\_ms.
  - [MYSQLND\\_QC\\_CONDITION\\_META\\_SCHEMA\\_PATTERN](#) to be used with [mysqlnd\\_qc\\_set\\_cache\\_condition](#).
- New function [mysqlnd\\_qc\\_set\\_cache\\_condition](#) for built-in schema pattern based caching. Likely to support a wider range of conditions in the future.
- Report [valid\\_until](#) timestamp for cache entries of the default handler through [mysqlnd\\_qc\\_get\\_cache\\_info](#).
- Include charset number for cache entry hashing. This should prevent serving result sets which have the wrong charset.  
  
API change: [get\\_hash\\_key](#) expects new "charsetnr" (int) parameter after "port".
- API change: changing [is\\_select\(\)](#) signature from bool [is\\_select\(\)](#) to mixed [is\\_select\(\)](#). Mixed can be either boolean or array(long ttl, string server\_id). This is needed by PECL/mysqlnd\_ms.

#### Other

- Support acting as a cache backend for [PECL/mysqlnd\\_ms](#) 1.3.0-beta or later to transparently replace MySQL Replication slave reads with cache accesses, if the user explicitly allows.

#### Bug fixes

- Fixed Bug #59959 (config.m4, wrong library - 64bit memcached handler builds) (Credits: Remi Collet)

### 21.9.7.8.3. PECL/mysqlnd\_qc 1.0 series

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

#### 1.0.1-stable

- Release date: 12/2010
- Motto/theme: Prepared statement support

Added support for Prepared statements and unbuffered queries.

#### 1.0.0-beta

- Release date: 07/2010
- Motto/theme: TTL-based cache with various storage options (Memcache, APC, SQLite, user-defined)

Initial public release of the transparent TTL-based query result cache. Flexible storage of cached results. Various storage media supported.

### 21.9.8. Mysqlnd user handler plugin ([mysqlnd\\_uh](#))

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The mysqlnd user handler plugin ([mysqlnd\\_uh](#)) allows users to set hooks for most internal calls of the MySQL native driver for PHP ([mysqlnd](#)). Mysqlnd and its plugins, including PECL/mysqlnd\_uh, operate on a layer beneath the PHP MySQL extensions. A mysqlnd plugin can be considered as a proxy between the PHP MySQL extensions and the MySQL server as part of the PHP executable on



the client-side. Because the plugins operates on their own layer below the PHP MySQL extensions, they can monitor and change application actions without requiring application changes. If the PHP MySQL extensions (`mysqli`, `mysql`, `PDO_MYSQL`) are compiled to use `mysqlnd` this can be used for:

- Monitoring
  - Queries executed by any of the PHP MySQL extensions
  - Prepared statements executing by any of the PHP MySQL extensions
- Auditing
  - Detection of database usage
  - SQL injection protection using black and white lists
- Assorted
  - Load Balancing connections

The MySQL native driver for PHP (`mysqlnd`) features an internal plugin C API. C plugins, such as the `mysqlnd` user handler plugin, can extend the functionality of `mysqlnd`. PECL/`mysqlnd_uh` makes parts of the internal plugin C API available to the PHP user for plugin development with PHP.



#### Status

The `mysqlnd` user handler plugin is in alpha status. Take appropriate care before using it in production environments.

### 21.9.8.1. Security considerations

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

PECL/`mysqlnd_uh` gives users access to MySQL user names, MySQL password used by any of the PHP MySQL extensions to connect to MySQL. It allows monitoring of all queries and prepared statements exposing the statement string to the user. Therefore, the extension should be installed with care. The `PHP_INI_SYSTEM` configuration setting `mysqlnd_uh.enable` can be used to prevent users from hooking `mysqlnd` calls.

Code obfuscators and similar technologies are not suitable to prevent monitoring of `mysqlnd` library activities if PECL/`mysqlnd_uh` is made available and the user can install a proxy, for example, using [auto\\_prepend\\_file](#).

### 21.9.8.2. Documentation note

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

Many of the `mysqlnd_uh` functions are briefly described because the `mysqli` extension is a thin abstraction layer on top of the MySQL C API that the `mysqlnd` library provides. Therefore, the corresponding `mysqli` documentation (along with the MySQL reference manual) can be consulted to receive more information about a particular function.

### 21.9.8.3. On the name

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The shortcut `mysqlnd_uh` stands for `mysqlnd user handler`, and has been the name since early development.

### 21.9.8.4. Quickstart and Examples

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The `mysqlnd` user handler plugin can be understood as a client-side proxy for all PHP MySQL extensions (`mysqli`, `mysql`, `PDO_MYSQL`), if they are compiled to use the `mysqlnd` library. The extensions use the `mysqlnd` library internally, at the C level, to communicate with the MySQL server. PECL/`mysqlnd_uh` allows it to hook many `mysqlnd` calls. Therefore, most activities of the PHP MySQL extensions can be monitored.

Because monitoring happens at the level of the library, at a layer below the application, it is possible to monitor applications without changing them.

On the C level, the `mysqlnd` library is structured in modules or classes. The extension hooks almost all methods of the `mysqlnd` internal `connection` class and exposes them through the user space class `MysqlndUhConnection`. Some few methods of the `mysqlnd` internal `statement` class are made available to the PHP user with the class `MysqlndUhPreparedStatement`. By subclassing the classes `MysqlndUhConnection` and `MysqlndUhPreparedStatement` users get access to `mysqlnd` internal function calls.

**Note**

The internal `mysqlnd` function calls are not designed to be exposed to the PHP user. Manipulating their activities may cause PHP to crash or leak memory. Often, this is not considered a bug. Please, keep in mind that you are accessing C library functions through PHP which are expected to take certain actions, which you may not be able to emulate in user space. Therefore, it is strongly recommended to always call the parent method implementation when subclassing `MysqlndUhConnection` or `MysqlndUhPreparedStatement`. To prevent the worst case, the extension performs some sanity checks. Please, see also the [Mysqlnd\\_uh Configure Options](#).

#### 21.9.8.4.1. Setup

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugin is implemented as a PHP extension. See the [installation instructions](#) to install the [PECL/mysqlnd\\_uh](#) extension. Then, load the extension into PHP and activate the plugin in the PHP configuration file using the PHP configuration directive named `mysqlnd_uh.enable`. The below example shows the default settings of the extension.

**Example 21.329. Enabling the plugin (php.ini)**

```
mysqlnd_uh.enable=1
mysqlnd_uh.report_wrong_types=1
```

#### 21.9.8.4.2. How it works

Copyright 1997-2012 the PHP Documentation Group. [2906]

This describes the background and inner workings of the `mysqlnd_uh` extension.

Two classes are provided by the extension: `MysqlndUhConnection` and `MysqlndUhPreparedStatement`. `MysqlndUhConnection` lets you access almost all methods of the `mysqlnd` internal `connection` class. The latter exposes some selected methods of the `mysqlnd` internal `statement` class. For example, `MysqlndUhConnection::connect` maps to the `mysqlnd` library C function `mysqlnd_conn__connect`.

As a `mysqlnd` plugin, the `PECL/mysqlnd_uh` extension replaces `mysqlnd` library C functions with its own functions. Whenever a PHP MySQL extension compiled to use `mysqlnd` calls a `mysqlnd` function, the functions installed by the plugin are executed instead of the original `mysqlnd` ones. For example, `mysqli_connect` invokes `mysqlnd_conn__connect`, so the `connect` function installed by `PECL/`

`mysqlnd_uh` will be called. The functions installed by PECL/`mysqlnd_uh` are the methods of the built-in classes.

The built-in PHP classes and their methods do nothing but call their `mysqlnd` C library counterparts, to behave exactly like the original `mysqlnd` function they replace. The code below illustrates in pseudo-code what the extension does.

**Example 21.330. Pseudo-code: what a built-in class does**

```
class MySQLndUhConnection {
    public function connect(($conn, $host, $user, $passwd, $db, $port, $socket, $mysql_flags) {
        MYSQLND* c_mysqlnd_connection = convert_from_php_to_c($conn);
        ...
        return call_c_function(mysqlnd_conn_connect(c_mysqlnd_connection, ...));
    }
}
```

The built-in classes behave like a transparent proxy. It is possible for you to replace the proxy with your own. This is done by subclassing `MySQLndUhConnection` or `MySQLndUhPreparedStatement` to extend the functionality of the proxy, followed by registering a new proxy object. Proxy objects are installed by `mysqlnd_uh_set_connection_proxy` and `mysqlnd_uh_set_statement_proxy`.

**Example 21.331. Installing a proxy**

```
<?php
class proxy extends MySQLndUhConnection {
    public function connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
?>
```

The above example will output:

```
proxy::connect(array (
  0 => NULL,
  1 => 'localhost',
  2 => 'root',
  3 => '',
  4 => 'test',
  5 => 3306,
  6 => NULL,
  7 => 131072,
))
proxy::connect returns true
```

**21.9.8.4.3. Installing a proxy**

Copyright 1997-2012 the PHP Documentation Group. [2906]

The extension provides two built-in classes: `MySQLndUhConnection` and `MySQLndUhPreparedStatement`. The classes are used for hooking `mysqlnd` library calls. Their methods correspond to `mysqlnd` internal functions. By default they act like a transparent proxy and do

nothing but call their `mysqlnd` counterparts. By subclassing the classes you can install your own proxy to monitor `mysqlnd`.

See also the [How it works](#) guide to learn about the inner workings of this extension.

Connection proxies are objects of the type `MySqlndUhConnection`. Connection proxy objects are installed by `mysqlnd_uh_set_connection_proxy`. If you install the built-in class `MySqlndUhConnection` as a proxy, nothing happens. It behaves like a transparent proxy.

#### Example 21.332. Proxy registration, `mysqlnd_uh.enable=1`

```
<?php
mysqlnd_uh_set_connection_proxy(new MySqlndUhConnection());
$mysqli = new mysqli("localhost", "root", "", "test");
?>
```

The `PHP_INI_SYSTEM` configuration setting `mysqlnd_uh.enable` controls whether a proxy may be set. If disabled, the extension will throw errors of type `E_WARNING`

#### Example 21.333. Proxy installation disabled

```
mysqlnd_uh.enable=0
```

```
<?php
mysqlnd_uh_set_connection_proxy(new MySqlndUhConnection());
$mysqli = new mysqli("localhost", "root", "", "test");
?>
```

The above example will output:

```
PHP Warning: MySqlndUhConnection::__construct(): (MySqlnd User Handler) The plugin has been disabled by set
PHP Warning: mysqlnd_uh_set_connection_proxy(): (MySqlnd User Handler) The plugin has been disabled by set
```

To monitor `mysqlnd`, you have to write your own proxy object subclassing `MySqlndUhConnection`. Please, see the function reference for a the list of methods that can be subclassed. Alternatively, you can use reflection to inspect the built-in `MySqlndUhConnection`.

Create a new class `proxy`. Derive it from the built-in class `MySqlndUhConnection`. Replace the `MySqlndUhConnection::connect` method. Print out the host parameter value passed to the method. Make sure that you call the parent implementation of the `connect` method. Failing to do so may give unexpected and undesired results, including memory leaks and crashes.

Register your proxy and open three connections using the PHP MySQL extensions `mysqli`, `mysql`, `PDO_MYSQL`. If the extensions have been compiled to use the `mysqlnd` library, the `proxy::connect` method will be called three times, once for each connection opened.

#### Example 21.334. Connection proxy

```
<?php
class proxy extends MySqlndUhConnection {
    public function connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags) {
        printf("Connection opened to '%s'\n", $host);
        /* Always call the parent implementation! */
        return parent::connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags);
    }
}
```

```

    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysql = mysql_connect("localhost", "root", "");
$pdo = new PDO("mysql:host=localhost;dbname=test", "root", "");
?>

```

The above example will output:

```

Connection opened to 'localhost'
Connection opened to 'localhost'
Connection opened to 'localhost'

```

The use of prepared statement proxies follows the same pattern: create a proxy object of the type `MysqlndUhPreparedStatement` and install the proxy using `mysqlnd_uh_set_statement_proxy`.

#### Example 21.335. Prepared statement proxy

```

<?php
class stmt_proxy extends MysqlndUhPreparedStatement {
    public function prepare($res, $query) {
        printf("%s(%s)\n", __METHOD__, $query);
        return parent::prepare($res, $query);
    }
}
mysqlnd_uh_set_statement_proxy(new stmt_proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$stmt = $mysqli->prepare("SELECT 'mysqlnd hacking made easy' AS _msg FROM DUAL");
?>

```

The above example will output:

```

stmt_proxy::prepare(SELECT 'mysqlnd hacking made easy' AS _msg FROM DUAL)

```

#### 21.9.8.4.4. Basic query monitoring

Copyright 1997-2012 the PHP Documentation Group. [2906]

Basic monitoring of a query statement is easy with PECL/mysqlnd\_uh. Combined with `debug_print_backtrace` it can become a powerful tool, for example, to find the origin of certain statement. This may be desired when searching for slow queries but also after database refactoring to find code still accessing deprecated databases or tables. The latter may be a complicated matter to do otherwise, especially if the application uses auto-generated queries.

#### Example 21.336. Basic Monitoring

```

<?php
class conn_proxy extends MysqlndUhConnection {
    public function query($res, $query) {
        debug_print_backtrace();
        return parent::query($res, $query);
    }
}

```

```
class stmt_proxy extends MysqlndUhPreparedStatement {
    public function prepare($res, $query) {
        debug_print_backtrace();
        return parent::prepare($res, $query);
    }
}
mysqlnd_uh_set_connection_proxy(new conn_proxy());
mysqlnd_uh_set_statement_proxy(new stmt_proxy());

printf("Proxies installed...\n");
$pdo = new PDO("mysql:host=localhost;dbname=test", "root", "");
var_dump($pdo->query("SELECT 1 AS _one FROM DUAL")->fetchAll(PDO::FETCH_ASSOC));

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->prepare("SELECT 1 AS _two FROM DUAL");
?>
```

The above example will output:

```
#0 conn_proxy->query(Resource id #19, SELECT 1 AS _one FROM DUAL)
#1 PDO->query(SELECT 1 AS _one FROM DUAL) called at [example.php:19]
array(1) {
    [0]=>
    array(1) {
        ["_one"]=>
        string(1) "1"
    }
}
#0 stmt_proxy->prepare(Resource id #753, SELECT 1 AS _two FROM DUAL)
#1 mysqli->prepare(SELECT 1 AS _two FROM DUAL) called at [example.php:22]
```

For basic query monitoring you should install a connection and a prepared statement proxy. The connection proxy should subclass `MysqlndUhConnection::query`. All database queries not using native prepared statements will call this method. In the example the `query` function is invoked by a PDO call. By default, `PDO_MySQL` is using prepared statement emulation.

All native prepared statements are prepared with the `prepare` method of `mysqlnd` exported through `MysqlndUhPreparedStatement::prepare`. Subclass `MysqlndUhPreparedStatement` and overwrite `prepare` for native prepared statement monitoring.

### 21.9.8.5. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.8.5.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

[PHP 5.3.3](#) or later. It is recommended to use [PHP 5.4.0](#) or later to get access to the latest `mysqlnd` features.

The `mysqlnd_uh` user handler plugin supports all PHP applications and all available PHP MySQL extensions (`mysqli`, `mysql`, `PDO_MYSQL`). The PHP MySQL extension must be configured to use `mysqlnd` in order to be able to use the `mysqlnd_uh` plugin for `mysqlnd`.

The alpha versions makes use of some `mysqli` features. You must enable `mysqli` to compile the plugin. This requirement may be removed in the future. Note, that this requirement does not restrict you to use the plugin only with `mysqli`. You can use the plugin to monitor `mysql`, `mysqli` and `PDO_MYSQL`.

#### 21.9.8.5.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

Information for installing this PECL extension may be found in the manual chapter titled [Installation of PECL extensions](#). Additional information such as new releases, downloads, source files, maintainer information, and a CHANGELOG, can be located here: <http://pecl.php.net/package/mysqlnd-uh>

PECL/`mysqlnd_uh` is currently not available on Windows. The source code of the extension makes use of `C99` constructs not allowed with PHP Windows builds.

### 21.9.8.5.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.76. MySQLnd\_uh Configure Options**

Name	Default	Changeable	Changelog
<code>mysqlnd_uh.enable</code>	1	PHP_INI_SYSTEM	
<code>mysqlnd_uh.report_wrong1_types</code>		PHP_INI_ALL	

Here's a short explanation of the configuration directives.

`mysqlnd_uh.enable` integer Enables or disables the plugin. If set to disabled, the extension will not allow users to plug into `mysqlnd` to hook `mysqlnd` calls.

`mysqlnd_uh.report_wrong_types` integer Whether to report wrong return value types of user hooks as `E_WARNING` level errors. This is recommended for detecting errors.

### 21.9.8.5.4. Resource Types

Copyright 1997-2012 the PHP Documentation Group. [2906]

This extension has no resource types defined.

### 21.9.8.6. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

Most of the constants refer to details of the MySQL Client Server Protocol. Please, refer to the MySQL reference manual to learn about their meaning. To avoid content duplication, only short descriptions are given.

`MySQLndUhConnection::simpleCommand` related

The following constants can be used to detect what command is to be send through `MySQLndUhConnection::simpleCommand`.

`MYSQLND_UH_MYSQLND_COM_SLEEP` MySQL Client Server protocol command: COM\_SLEEP.  
(integer)

`MYSQLND_UH_MYSQLND_COM_QUIT` MySQL Client Server protocol command: COM\_QUIT.  
(integer)

`MYSQLND_UH_MYSQLND_COM_INIT_DB` MySQL Client Server protocol command: COM\_INIT\_DB.  
(integer)

`MYSQLND_UH_MYSQLND_COM_QUERY` MySQL Client Server protocol command: COM\_QUERY.  
(integer)

`MYSQLND_UH_MYSQLND_COM_FIELD_LIST` MySQL Client Server protocol command: COM\_FIELD\_LIST.  
(integer)

`MYSQLND_UH_MYSQLND_COM_CREATE_DB` MySQL Client Server protocol command: COM\_CREATE\_DB.  
(integer)

`MYSQLND_UH_MYSQLND_COM_DROP_DB` MySQL Client Server protocol command: COM\_DROP\_DB.  
(integer)

`MYSQLND_UH_MYSQLND_COM_REFRESH` MySQL Client Server protocol command: COM\_REFRESH.  
(integer)

`MYSQLND_UH_MYSQLND_COM_SHUTDOWN` MySQL Client Server protocol command: COM\_SHUTDOWN.  
(integer)

`MYSQLND_UH_MYSQLND_COM_STATISTICS` MySQL Client Server protocol command: COM\_STATISTICS.  
(integer)

`MYSQLND_UH_MYSQLND_COM_PROCESS_INFO` MySQL Client Server protocol command: COM\_PROCESS\_INFO.  
(integer)

`MYSQLND_UH_MYSQLND_COM_CONNECT` MySQL Client Server protocol command: COM\_CONNECT.  
(integer)

`MYSQLND_UH_MYSQLND_COM_PROCESS_KILL` MySQL Client Server protocol command: COM\_PROCESS\_KILL.  
(integer)

`MYSQLND_UH_MYSQLND_COM_DEBUG` MySQL Client Server protocol command: COM\_DEBUG.  
(integer)

`MYSQLND_UH_MYSQLND_COM_PING` MySQL Client Server protocol command: COM\_PING.  
(integer)

`MYSQLND_UH_MYSQLND_COM_TIME` MySQL Client Server protocol command: COM\_TIME.  
(integer)

`MYSQLND_UH_MYSQLND_COM_DELAYED_INSERT` MySQL Client Server protocol command:  
(integer) COM\_DELAYED\_INSERT.

`MYSQLND_UH_MYSQLND_COM_CHANGE_USER` MySQL Client Server protocol command: COM\_CHANGE\_USER.  
(integer)

`MYSQLND_UH_MYSQLND_COM_BINLOG_DUMP` MySQL Client Server protocol command: COM\_BINLOG\_DUMP.  
(integer)

`MYSQLND_UH_MYSQLND_COM_TABLE_DUMP` MySQL Client Server protocol command: COM\_TABLE\_DUMP.  
(integer)

`MYSQLND_UH_MYSQLND_COM_CONNECT_OUT` MySQL Client Server protocol command: COM\_CONNECT\_OUT.  
(integer)

`MYSQLND_UH_MYSQLND_COM_REGISTER_SLAVED` MySQL Client Server protocol command:  
(integer) COM\_REGISTER\_SLAVED.

`MYSQLND_UH_MYSQLND_COM_STMT_PREPARE` MySQL Client Server protocol command: COM\_STMT\_PREPARE.  
(integer)

`MYSQLND_UH_MYSQLND_COM_STMT_EXECUTE` MySQL Client Server protocol command: COM\_STMT\_EXECUTE.  
(integer)

`MYSQLND_UH_MYSQLND_COM_STMT_SEND_LONG_DATA` MySQL Client Server protocol command:  
(integer) COM\_STMT\_SEND\_LONG\_DATA.

`MYSQLND_UH_MYSQLND_COM_STMT_CLOSE` MySQL Client Server protocol command: COM\_STMT\_CLOSE.  
(integer)



`MYSQLND_UH_MYSQLND_COM_STMT_RESET` MySQL Client Server protocol command: COM\_STMT\_RESET.  
(integer)

`MYSQLND_UH_MYSQLND_COM_SET_OPTION` MySQL Client Server protocol command: COM\_SET\_OPTION.  
(integer)

`MYSQLND_UH_MYSQLND_COM_STMT_FETCH` MySQL Client Server protocol command: COM\_STMT\_FETCH.  
(integer)

`MYSQLND_UH_MYSQLND_COM_DAEMON` MySQL Client Server protocol command: COM\_DAEMON.  
(integer)

`MYSQLND_UH_MYSQLND_COM_END` MySQL Client Server protocol command: COM\_END.  
(integer)

The following constants can be used to analyze the `ok_packet` argument of `MySQLndUhConnection::simpleCommand`.

`MYSQLND_UH_MYSQLND_PROT_GREETING` MySQL Client Server protocol packet: greeting.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_AUTH` MySQL Client Server protocol packet: authentication.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_OK` MySQL Client Server protocol packet: OK.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_EOF` MySQL Client Server protocol packet: EOF.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_CMD` MySQL Client Server protocol packet: command.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_RS_HEADER` MySQL Client Server protocol packet: result set header.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_RS_FIELD` MySQL Client Server protocol packet: resultset field.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_ROW` MySQL Client Server protocol packet: row.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_STATS` MySQL Client Server protocol packet: stats.  
(integer)

`MYSQLND_UH_MYSQLND_PREPARE` MySQL Client Server protocol packet: prepare response.  
(integer)

`MYSQLND_UH_MYSQLND_CHG_USER` MySQL Client Server protocol packet: change user response.  
(integer)

`MYSQLND_UH_MYSQLND_PROT_LAST` No practical meaning. Last entry marker of internal C data structure list.  
(integer)

`MySQLndUhConnection::close` *related*

The following constants can be used to detect why a connection has been closed through `MySQLndUhConnection::close()`.

`MYSQLND_UH_MYSQLND_CLOSE_USER` User has called `mysqlnd` to close the connection.  
(integer)

`MYSQLND_UH_MYSQLND_CLOSE_IMPLICIT` Implicitly closed, for example, during garbage connection.  
(integer)

`MYSQLND_UH_MYSQLND_CLOSE_ID` Connection error.  
(integer)

`MYSQLND_UH_MYSQLND_CLOSE_ID` No practical meaning. Last entry marker of internal C data structure  
(integer) list.

*MySQLndUhConnection::setServerOption()* related

The following constants can be used to detect which option is set through `MySQLndUhConnection::setServerOption()`.

`MYSQLND_UH_SERVER_OPTION_M` Option: enables multi statement support.  
(integer)

`MYSQLND_UH_SERVER_OPTION_M` Option: disables multi statement support.  
(integer)

*MySQLndUhConnection::setClientOption* related

The following constants can be used to detect which option is set through `MySQLndUhConnection::setClientOption`.

`MYSQLND_UH_MYSQLND_OPTION` Option: connection timeout.  
(integer)

`MYSQLND_UH_MYSQLND_OPTION` Option: whether the MySQL compressed protocol is to be used.  
(integer)

`MYSQLND_UH_MYSQLND_OPTION` Option: named pipe to use for connection (Windows).  
(integer)

`MYSQLND_UH_MYSQLND_OPTION` Option: init command to execute upon connect.  
(integer)

`MYSQLND_UH_MYSQLND_READ_DE` Option: MySQL server default file to read upon connect.  
(integer)

`MYSQLND_UH_MYSQLND_READ_DE` Option: MySQL server default file group to read upon connect.  
(integer)

`MYSQLND_UH_MYSQLND_SET_CHA` Option: charset description files directory.  
(integer)

`MYSQLND_UH_MYSQLND_SET_CHA` Option: charset name.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_LO` Option: Whether to allow `LOAD DATA LOCAL INFILE` use.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_PR` Option: supported protocol version.  
(integer)

`MYSQLND_UH_MYSQLND_SHARED` Option: shared memory base name for shared memory connections.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_RE` Option: connection read timeout.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_W` Option: connection write timeout.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_U` Option: unbuffered result sets.  
(integer)

- `MYSQLND_UH_MYSQLND_OPT_USE_EMBEDDED_SERVER` Embedded server related.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_USE_EMBEDDED_SERVER_BIND` Embedded server related.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_GUEST_CONNECTION` TODO  
(integer)
- `MYSQLND_UH_MYSQLND_SET_CLIENT_IP` TODO  
(integer)
- `MYSQLND_UH_MYSQLND_SECURE_FILE` TODO  
(integer)
- `MYSQLND_UH_MYSQLND_REPORT_TRUNCATION` Option: Whether to report data truncation.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_RECONNECT` Option: Whether to reconnect automatically.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_VERIFY_SERVER_CERT` Option: TODO  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_NETWORK_BUFFER_SIZE_COMMANDS` Option: mysqlnd network buffer size for commands.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_NETWORK_BUFFER_SIZE_READING` Option: mysqlnd network buffer size for reading from the server.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_KEY` Option: SSL key.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_CERTIFICATE` Option: SSL certificate.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_CA` Option: SSL CA.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_PATH_TO_CA` Option: Path to SSL CA.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_CIPHER` Option: SSL cipher.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_SSL_PASSPHRASE` Option: SSL passphrase.  
(integer)
- `MYSQLND_UH_SERVER_OPTION_PLUGIN_DIRECTORY` Option: server plugin directory.  
(integer)
- `MYSQLND_UH_SERVER_OPTION_DEFAULT_AUTHENTICATION_METHOD` Option: default authentication method.  
(integer)
- `MYSQLND_UH_SERVER_OPTION_SECURE_CLIENT_IP` TODO  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_MAXIMUM_ALLOWED_PACKET_SIZE` Option: maximum allowed packet size. Available as of PHP 5.4.0.  
(integer)
- `MYSQLND_UH_MYSQLND_OPT_AUTOMATICALLY_LOAD_EXTENSION` Option: TODO. Available as of PHP 5.4.0.  
(integer)

`MYSQLND_UH_MYSQLND_OPT_INT_INTEGER` (integer) Option: `make mysqLnd` return integer and float columns as long even when using the MySQL Client Server text protocol. Only available with a custom build of `mysqLnd`.

*Other*

The plugins version number can be obtained using `MYSQLND_UH_VERSION` or `MYSQLND_UH_VERSION_ID`. `MYSQLND_UH_VERSION` is the string representation of the numerical version number `MYSQLND_UH_VERSION_ID`, which is an integer such as 10000. Developers can calculate the version number as follows.

Version (part)	Example
Major*10000	1*10000 = 10000
Minor*100	0*100 = 0
Patch	0 = 0
<code>MYSQLND_UH_VERSION_ID</code>	10000

`MYSQLND_UH_VERSION` (string) Plugin version string, for example, "1.0.0-alpha".

`MYSQLND_UH_VERSION_ID` (integer) Plugin version number, for example, 10000.

### 21.9.8.7. The `MysqLndUhConnection` class (`MysqLndUhConnection`)

Copyright 1997-2012 the PHP Documentation Group. [2906]

```
MysqLndUhConnection {
MysqLndUhConnection

    Methods

    public bool MysqLndUhConnection::changeUser(
        mysqLnd_connection connection,
        string user,
        string password,
        string database,
        bool silent,
        int passwd_len);

    public string MysqLndUhConnection::charsetName(
        mysqLnd_connection connection);

    public bool MysqLndUhConnection::close(
        mysqLnd_connection connection,
        int close_type);

    public bool MysqLndUhConnection::connect(
        mysqLnd_connection connection,
        string host,
        string use",
        string password,
        string database,
        int port,
        string socket,
        int mysql_flags);

    public MysqLndUhConnection::__construct();

    public bool MysqLndUhConnection::endPSession(
        mysqLnd_connection connection);

    public string MysqLndUhConnection::escapeString(
        mysqLnd_connection connection,
```

```

    string escape_string);

public int MySQLndUhConnection::getAffectedRows(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getErrorNumber(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getErrorString(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getFieldCount(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getHostInformation(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getLastInsertId(
    mysqlnd_connection connection);

public void MySQLndUhConnection::getLastMessage(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getProtocolInformation(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getServerInformation(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getServerStatistics(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getServerVersion(
    mysqlnd_connection connection);

public string MySQLndUhConnection::getSqlstate(
    mysqlnd_connection connection);

public array MySQLndUhConnection::getStatistics(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getThreadId(
    mysqlnd_connection connection);

public int MySQLndUhConnection::getWarningCount(
    mysqlnd_connection connection);

public bool MySQLndUhConnection::init(
    mysqlnd_connection connection);

public bool MySQLndUhConnection::killConnection(
    mysqlnd_connection connection,
    int pid);

public array MySQLndUhConnection::listFields(
    mysqlnd_connection connection,
    string table,
    string achtung_wild);

public void MySQLndUhConnection::listMethod(
    mysqlnd_connection connection,
    string query,
    string achtung_wild,
    string parl);

public bool MySQLndUhConnection::moreResults(
    mysqlnd_connection connection);

public bool MySQLndUhConnection::nextResult(
    mysqlnd_connection connection);

public bool MySQLndUhConnection::ping(
    mysqlnd_connection connection);

```

```
public bool MysqlndUhConnection::query(
    mysqlnd_connection connection,
    string query);

public bool MysqlndUhConnection::queryReadResultsetHeader(
    mysqlnd_connection connection,
    mysqlnd_statement mysqlnd_stmt);

public bool MysqlndUhConnection::reapQuery(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::refreshServer(
    mysqlnd_connection connection,
    int options);

public bool MysqlndUhConnection::restartPSession(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::selectDb(
    mysqlnd_connection connection,
    string database);

public bool MysqlndUhConnection::sendClose(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::sendQuery(
    mysqlnd_connection connection,
    string query);

public bool MysqlndUhConnection::serverDumpDebugInformation(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::setAutocommit(
    mysqlnd_connection connection,
    int mode);

public bool MysqlndUhConnection::setCharset(
    mysqlnd_connection connection,
    string charset);

public bool MysqlndUhConnection::setClientOption(
    mysqlnd_connection connection,
    int option,
    int value);

public void MysqlndUhConnection::setServerOption(
    mysqlnd_connection connection,
    int option);

public void MysqlndUhConnection::shutdownServer(
    string MYSQLND_UH_RES_MYSQLND_NAME,
    string "level");

public bool MysqlndUhConnection::simpleCommand(
    mysqlnd_connection connection,
    int command,
    string arg,
    int ok_packet,
    bool silent,
    bool ignore_upsert_status);

public bool MysqlndUhConnection::simpleCommandHandleResponse(
    mysqlnd_connection connection,
    int ok_packet,
    bool silent,
    int command,
    bool ignore_upsert_status);

public bool MysqlndUhConnection::sslSet(
    mysqlnd_connection connection,
    string key,
    string cert,
    string ca,
```

```

    string cpath,
    string cipher);

public resource MysqlndUhConnection::stmtInit(
    mysqlnd_connection connection);

public resource MysqlndUhConnection::storeResult(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::txCommit(
    mysqlnd_connection connection);

public bool MysqlndUhConnection::txRollback(
    mysqlnd_connection connection);

public resource MysqlndUhConnection::useResult(
    mysqlnd_connection connection);
}

```

### 21.9.8.7.1. `MysqlndUhConnection::changeUser`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::changeUser`

Changes the user of the specified mysqlnd database connection

#### Description

```

public bool MysqlndUhConnection::changeUser(
    mysqlnd_connection connection,
    string user,
    string password,
    string database,
    bool silent,
    int passwd_len);

```

Changes the user of the specified mysqlnd database connection

#### Parameters

<i>connection</i>	Mysqlnd connection handle. Do not modify!
<i>user</i>	The MySQL user name.
<i>password</i>	The MySQL password.
<i>database</i>	The MySQL database to change to.
<i>silent</i>	Controls if mysqlnd is allowed to emit errors or not.
<i>passwd_len</i>	Length of the MySQL password.

#### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

#### Examples

##### Example 21.337. `MysqlndUhConnection::changeUser` example

```

<?php
class proxy extends MysqlndUhConnection {
    /* Hook mysqlnd's connection::change_user call */
    public function changeUser($res, $user, $passwd, $db, $silent, $passwd_len) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::changeUser($res, $user, $passwd, $db, $silent, $passwd_len);
    }
}

```

```
printf("%s returns %s\n", __METHOD__, var_export($ret, true));
return $ret;
}
}
/* Install proxy/hooks to be used with all future mysqlnd connection */
mysqlnd_uh_set_connection_proxy(new proxy());

/* Create mysqli connection which is using the mysqlnd library */
$mysqli = new mysqli("localhost", "root", "", "test");

/* Example of a user API call which triggers the hooked mysqlnd call */
var_dump($mysqli->change_user("root", "bar", "test"));
?>
```

The above example will output:

```
proxy::changeUser(array (
  0 => NULL,
  1 => 'root',
  2 => 'bar',
  3 => 'test',
  4 => false,
  5 => 3,
))
proxy::changeUser returns false
bool(false)
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_change\\_user](#)

### 21.9.8.7.2. `MysqlndUhConnection::charsetName`

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

- [MysqlndUhConnection::charsetName](#)

Returns the default character set for the database connection

#### Description

```
public string MysqlndUhConnection::charsetName(
    mysqlnd_connection connection);
```

Returns the default character set for the database connection.

#### Parameters

*connection* Mysqlnd connection handle. Do not modify!

#### Return Values

The default character set.

#### Examples

##### Example 21.338. `MysqlndUhConnection::charsetName` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function charsetName($res) {
```



```

printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
$ret = parent::charsetName($res);
printf("%s returns %s\n", __METHOD__, var_export($ret, true));
return $ret;
}
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump(mysqli_character_set_name($mysqli));
?>

```

The above example will output:

```

proxy::charsetName(array (
  0 => NULL,
))
proxy::charsetName returns 'latin1'
string(6) "latin1"

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_character\\_set\\_name](#)

#### 21.9.8.7.3. `MysqlndUhConnection::close`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::close`

Closes a previously opened database connection

### Description

```

public bool MysqlndUhConnection::close(
    mysqlnd_connection connection,
    int close_type);

```

Closes a previously opened database connection.



### Note

Failing to call the parent implementation may cause memory leaks or crash PHP. This is not considered a bug. Please, keep in mind that the `mysqlnd` library functions have never been designed to be exposed to the user space.

### Parameters

*connection*

The connection to be closed. Do not modify!

*close\_type*

Why the connection is to be closed. The value of *close\_type* is one of `MYSQLND_UH_MYSQLND_CLOSE_EXPLICIT`, `MYSQLND_UH_MYSQLND_CLOSE_IMPLICIT`, `MYSQLND_UH_MYSQLND_CLOSE_DISCONNECTED` or `MYSQLND_UH_MYSQLND_CLOSE_LAST`. The latter should never be seen, unless the default behaviour of the `mysqlnd` library has been changed by a plugin.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

## Examples

### Example 21.339. [MysqlndUhConnection::close](#) example

```
<?php
function close_type_to_string($close_type) {
    $mapping = array(
        MYSQLND_UH_MYSQLND_CLOSE_DISCONNECTED => "MYSQLND_UH_MYSQLND_CLOSE_DISCONNECTED",
        MYSQLND_UH_MYSQLND_CLOSE_EXPLICIT => "MYSQLND_UH_MYSQLND_CLOSE_EXPLICIT",
        MYSQLND_UH_MYSQLND_CLOSE_IMPLICIT => "MYSQLND_UH_MYSQLND_CLOSE_IMPLICIT",
        MYSQLND_UH_MYSQLND_CLOSE_LAST => "MYSQLND_UH_MYSQLND_CLOSE_IMPLICIT"
    );
    return (isset($mapping[$close_type])) ? $mapping[$close_type] : 'unknown';
}

class proxy extends MysqlndUhConnection {
    public function close($res, $close_type) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        printf("close_type = %s\n", close_type_to_string($close_type));
        /* WARNING: you must call the parent */
        $ret = parent::close($res, $close_type);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->close();
?>
```

The above example will output:

```
proxy::close(array (
    0 => NULL,
    1 => 0,
))
close_type = MYSQLND_UH_MYSQLND_CLOSE_EXPLICIT
proxy::close returns true
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_close](#)  
[mysql\\_close](#)

### 21.9.8.7.4. [MysqlndUhConnection::connect](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::connect](#)

Open a new connection to the MySQL server

## Description

```
public bool MysqlndUhConnection::connect(
    mysqlnd_connection connection,
    string host,
    string use",
    string password,
    string database,
    int port,
    string socket,
```

```
int mysql_flags);
```

Open a new connection to the MySQL server.

### Parameters

<code>connection</code>	Mysqlnd connection handle. Do not modify!
<code>host</code>	Can be either a host name or an IP address. Passing the <code>NULL</code> value or the string "localhost" to this parameter, the local host is assumed. When possible, pipes will be used instead of the TCP/IP protocol.
<code>user</code>	The MySQL user name.
<code>password</code>	If not provided or <code>NULL</code> , the MySQL server will attempt to authenticate the user against those user records which have no password only. This allows one username to be used with different permissions (depending on if a password as provided or not).
<code>database</code>	If provided will specify the default database to be used when performing queries.
<code>port</code>	Specifies the port number to attempt to connect to the MySQL server.
<code>socket</code>	Specifies the socket or named pipe that should be used. If <code>NULL</code> , mysqlnd will default to <code>/tmp/mysql.sock</code> .
<code>mysql_flags</code>	Connection options.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.340. `MysqlndUhConnection::connect` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::connect($res, $host, $user, $passwd, $db, $port, $socket, $mysql_flags);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
?>
```

The above example will output:

```
proxy::connect(array (
    0 => NULL,
    1 => 'localhost',
    2 => 'root',
    3 => '',
    4 => 'test',
    5 => 3306,
    6 => NULL,
```

```
7 => 131072,
))
proxy::connect returns true
```

### See Also

`mysqlnd_uh_set_connection_proxy`  
`mysqli_connect`  
`mysql_connect`

#### 21.9.8.7.5. `mysqlnd_uh_connection::__construct`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_uh_connection::__construct`

The `__construct` purpose

### Description

```
public MysqlndUhConnection::__construct();
```



### Warning

This function is currently not documented; only its argument list is available.

### Parameters

This function has no parameters.

### Return Values

#### 21.9.8.7.6. `mysqlnd_uh_connection::endPSession`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_uh_connection::endPSession`

End a persistent connection

### Description

```
public bool MysqlndUhConnection::endPSession(
    mysqlnd_connection connection);
```

End a persistent connection



### Warning

This function is currently not documented; only its argument list is available.

### Parameters

`connection` Mysqlnd connection handle. Do not modify!

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

**Example 21.341. [MySQLndUhConnection::endPSession](#) example**

```
<?php
class proxy extends MySQLndUhConnection {
    public function endPSession($conn) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::endPSession($conn);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("p:localhost", "root", "", "test");
$mysqli->close();
?>
```

The above example will output:

```
proxy::endPSession(array (
    0 => NULL,
))
proxy::endPSession returns true
```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

**21.9.8.7.7. [MySQLndUhConnection::escapeString](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::escapeString](#)

Escapes special characters in a string for use in an SQL statement, taking into account the current charset of the connection

**Description**

```
public string MySQLndUhConnection::escapeString(
    mysqlnd_connection connection,
    string escape_string);
```

Escapes special characters in a string for use in an SQL statement, taking into account the current charset of the connection.

**Parameters**

[MYSQLND\\_UH\\_RES\\_MYSQLND\\_NAME](#) MySQLnd connection handle. Do not modify!

[escape\\_string](#) The string to be escaped.

**Return Values**

The escaped string.

**Examples**

**Example 21.342. [MySQLndUhConnection::escapeString](#) example**

```
<?php
```

```
class proxy extends MysqlndUhConnection {
    public function escapeString($res, $string) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::escapeString($res, $string);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->set_charset("latin1");
$mysqli->real_escape_string("test0'test");
?>
```

The above example will output:

```
proxy::escapeString(array (
    0 => NULL,
    1 => 'test0'test',
))
proxy::escapeString returns 'test0\\\'test'
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_real\\_escape\\_string](#)  
[mysql\\_real\\_escape\\_string](#)

### 21.9.8.7.8. [MysqlndUhConnection::getAffectedRows](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getAffectedRows](#)

Gets the number of affected rows in a previous MySQL operation

#### Description

```
public int MysqlndUhConnection::getAffectedRows(
    mysqlnd_connection connection);
```

Gets the number of affected rows in a previous MySQL operation.

#### Parameters

*connection* Mysqlnd connection handle. Do not modify!

#### Return Values

Number of affected rows.

#### Examples

#### Example 21.343. [MysqlndUhConnection::getAffectedRows](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getAffectedRows($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getAffectedRows($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
    }
}
```

```

    return $ret;
}
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT)");
$mysqli->query("INSERT INTO test(id) VALUES (1)");
var_dump($mysqli->affected_rows);
?>

```

The above example will output:

```

proxy::getAffectedRows(array (
  0 => NULL,
))
proxy::getAffectedRows returns 1
int(1)

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_affected\\_rows](#)  
[mysql\\_affected\\_rows](#)

#### 21.9.8.7.9. `MysqlndUhConnection::getErrorNumber`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getErrorNumber](#)

Returns the error code for the most recent function call

### Description

```

public int MysqlndUhConnection::getErrorNumber(
    mysqlnd_connection connection);

```

Returns the error code for the most recent function call.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Error code for the most recent function call.

### Examples

`MysqlndUhConnection::getErrorNumber` is not only executed after the invocation of a user space API call which maps directly to it but also called internally.

#### Example 21.344. `MysqlndUhConnection::getErrorNumber` example

```

<?php
class proxy extends MysqlndUhConnection {
    public function getErrorNumber($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getErrorNumber($res);
    }
}

```

```

    printf("%s returns %s\n", __METHOD__, var_export($ret, true));
    return $ret;
}
}
mysqlnd_uh_set_connection_proxy(new proxy());

printf("connect...\n");
$mysqli = new mysqli("localhost", "root", "", "test");
printf("query...\n");
$mysqli->query("PLEASE_LET_THIS_BE_INVALID_SQL");
printf("errno...\n");
var_dump($mysqli->errno);
printf("close...\n");
$mysqli->close();
?>

```

The above example will output:

```

connect...
proxy::getLineNumber(array (
    0 => NULL,
))
proxy::getLineNumber returns 0
query...
errno...
proxy::getLineNumber(array (
    0 => NULL,
))
proxy::getLineNumber returns 1064
int(1064)
close...

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[MysqlndUhConnection::getErrorString](#)  
[mysqli\\_errno](#)  
[mysql\\_errno](#)

#### 21.9.8.7.10. `MysqlndUhConnection::getErrorString`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::getErrorString`

Returns a string description of the last error

#### Description

```

public string MysqlndUhConnection::getErrorString(
    mysqlnd_connection connection);

```

Returns a string description of the last error.

#### Parameters

*connection* Mysqlnd connection handle. Do not modify!

#### Return Values

Error string for the most recent function call.

#### Examples



[MysqlndUhConnection::getErrorString](#) is not only executed after the invocation of a user space API call which maps directly to it but also called internally.

#### Example 21.345. [MysqlndUhConnection::getErrorString](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getErrorString($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getErrorString($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

printf("connect...\n");
$mysqli = new mysqli("localhost", "root", "", "test");
printf("query...\n");
$mysqli->query("WILL_I_EVER_LEARN_SQL?");
printf("errno...\n");
var_dump($mysqli->error);
printf("close...\n");
$mysqli->close();
?>
```

The above example will output:

```
connect...
proxy::getErrorString(array (
    0 => NULL,
))
proxy::getErrorString returns ''
query...
errno...
proxy::getErrorString(array (
    0 => NULL,
))
proxy::getErrorString returns 'You have an error in your SQL syntax; check the manual that corresponds
string(168) "You have an error in your SQL syntax; check the manual that corresponds to your MySQL serv
close...
```

#### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[MysqlndUhConnection::getErrorNumber](#)  
[mysqli\\_error](#)  
[mysql\\_error](#)

#### 21.9.8.7.11. [MysqlndUhConnection::getFieldCount](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getFieldCount](#)

Returns the number of columns for the most recent query

#### Description

```
public int MysqlndUhConnection::getFieldCount(
    mysqlnd_connection connection);
```

Returns the number of columns for the most recent query.

## Parameters

[connection](#) Mysqlnd connection handle. Do not modify!

## Return Values

Number of columns.

## Examples

[MysqlndUhConnection::getFieldCount](#) is not only executed after the invocation of a user space API call which maps directly to it but also called internally.

### Example 21.346. [MysqlndUhConnection::getFieldCount](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getFieldCount($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getFieldCount($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("WILL_I_EVER_LEARN_SQL?");
var_dump($mysqli->field_count);
$mysqli->query("SELECT 1, 2, 3 FROM DUAL");
var_dump($mysqli->field_count);
?>
```

The above example will output:

```
proxy::getFieldCount(array (
    0 => NULL,
))
proxy::getFieldCount returns 0
int(0)
proxy::getFieldCount(array (
    0 => NULL,
))
proxy::getFieldCount returns 3
proxy::getFieldCount(array (
    0 => NULL,
))
proxy::getFieldCount returns 3
int(3)
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_field\\_count](#)

### 21.9.8.7.12. [MysqlndUhConnection::getHostInformation](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getHostInformation](#)

Returns a string representing the type of connection used

## Description

```
public string MysqlndUhConnection::getHostInformation(
    mysqlnd_connection connection);
```

Returns a string representing the type of connection used.

## Parameters

*connection* Mysqlnd connection handle. Do not modify!

## Return Values

Connection description.

## Examples

### Example 21.347. [MysqlndUhConnection::getHostInformation](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getHostInformation($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getHostInformation($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->host_info);
?>
```

The above example will output:

```
proxy::getHostInformation(array (
    0 => NULL,
))
proxy::getHostInformation returns 'Localhost via UNIX socket'
string(25) "Localhost via UNIX socket"
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_get\\_host\\_info](#)  
[mysql\\_get\\_host\\_info](#)

### 21.9.8.7.13. [MysqlndUhConnection::getLastInsertId](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getLastInsertId](#)

Returns the auto generated id used in the last query.

## Description

```
public int MysqlndUhConnection::getLastInsertId(
    mysqlnd_connection connection);
```

Returns the auto generated id used in the last query.

## Parameters

[connection](#) Mysqlnd connection handle. Do not modify!

## Return Values

Last insert id.

## Examples

### Example 21.348. [MysqlndUhConnection::getLastInsertId](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getLastInsertId($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getLastInsertId($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("DROP TABLE IF EXISTS test");
$mysqli->query("CREATE TABLE test(id INT AUTO_INCREMENT PRIMARY KEY, col VARCHAR(255))");
$mysqli->query("INSERT INTO test(col) VALUES ('a')");
var_dump($mysqli->insert_id);
?>
```

The above example will output:

```
proxy::getLastInsertId(array (
    0 => NULL,
))
proxy::getLastInsertId returns 1
int(1)
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_insert\\_id](#)  
[mysql\\_insert\\_id](#)

### 21.9.8.7.14. [MysqlndUhConnection::getLastMessage](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getLastMessage](#)

Retrieves information about the most recently executed query

## Description

```
public void MysqlndUhConnection::getLastMessage(
    mysqlnd_connection connection);
```

Retrieves information about the most recently executed query.

## Parameters

[connection](#) Mysqlnd connection handle. Do not modify!

## Return Values

Last message. Trying to return a string longer than 511 bytes will cause an error of the type [E\\_WARNING](#) and result in the string being truncated.

## Examples

### Example 21.349. [MysqlndUhConnection::getLastMessage](#) example

```

<?php
class proxy extends MysqlndUhConnection {
    public function getLastMessage($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getLastMessage($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->info);
$mysqli->query("DROP TABLE IF EXISTS test");
var_dump($mysqli->info);
?>

```

The above example will output:

```

proxy::getLastMessage(array (
    0 => NULL,
))
proxy::getLastMessage returns ''
string(0) ""
proxy::getLastMessage(array (
    0 => NULL,
))
proxy::getLastMessage returns ''
string(0) ""

```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_info](#)  
[mysql\\_info](#)

### 21.9.8.7.15. [MysqlndUhConnection::getProtocolInformation](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getProtocolInformation](#)

Returns the version of the MySQL protocol used

## Description

```

public string MysqlndUhConnection::getProtocolInformation(
    mysqlnd_connection connection);

```

Returns the version of the MySQL protocol used.

## Parameters

[connection](#) MySQLnd connection handle. Do not modify!

## Return Values

The protocol version.

## Examples

### Example 21.350. `MysqlndUhConnection::getProtocolInformation` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getProtocolInformation($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getProtocolInformation($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->protocol_version);
?>
```

The above example will output:

```
proxy::getProtocolInformation(array (
  0 => NULL,
))
proxy::getProtocolInformation returns 10
int(10)
```

## See Also

`mysqlnd_uh_set_connection_proxy`  
`mysqli_get_proto_info`  
`mysql_get_proto_info`

### 21.9.8.7.16. `MysqlndUhConnection::getServerInformation`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::getServerInformation`

Returns the version of the MySQL server

## Description

```
public string MysqlndUhConnection::getServerInformation(
    mysqlnd_connection connection);
```

Returns the version of the MySQL server.

## Parameters

*connection* MySQLnd connection handle. Do not modify!

## Return Values

The server version.

## Examples

**Example 21.351. [MySQLndUhConnection::getServerInformation](#) example**

```
<?php
class proxy extends MySQLndUhConnection {
    public function getServerInformation($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getServerInformation($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->server_info);
?>
```

The above example will output:

```
proxy::getServerInformation(array (
    0 => NULL,
))
proxy::getServerInformation returns '5.1.45-debug-log'
string(16) "5.1.45-debug-log"
```

**See Also**

- [mysqlnd\\_uh\\_set\\_connection\\_proxy](#)
- [mysqli\\_get\\_server\\_info](#)
- [mysql\\_get\\_server\\_info](#)

**21.9.8.7.17. [MySQLndUhConnection::getServerStatistics](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::getServerStatistics](#)

Gets the current system status

**Description**

```
public string MySQLndUhConnection::getServerStatistics(
    mysqlnd_connection connection);
```

Gets the current system status.

**Parameters**

*connection* MySQLnd connection handle. Do not modify!

**Return Values**

The system status message.

**Examples**

**Example 21.352. [MySQLndUhConnection::getServerStatistics](#) example**

```
<?php
class proxy extends MySQLndUhConnection {
    public function getServerStatistics($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
```

```

    $ret = parent::getServerStatistics($res);
    printf("%s returns %s\n", __METHOD__, var_export($ret, true));
    return $ret;
}
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump(mysqli_stat($mysqli));
?>

```

The above example will output:

```

proxy::getServerStatistics(array (
    0 => NULL,
))
proxy::getServerStatistics returns 'Uptime: 2059995 Threads: 1 Questions: 126157 Slow queries: 0 Opens: 1
string(140) "Uptime: 2059995 Threads: 1 Questions: 126157 Slow queries: 0 Opens: 6377 Flush tables: 1

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_stat](#)  
[mysql\\_stat](#)

### 21.9.8.7.18. `MySQLndUhConnection::getServerVersion`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MySQLndUhConnection::getServerVersion`

Returns the version of the MySQL server as an integer

#### Description

```

public int MySQLndUhConnection::getServerVersion(
    mysqlnd_connection connection);

```

Returns the version of the MySQL server as an integer.

#### Parameters

*connection* MySQLnd connection handle. Do not modify!

#### Return Values

The MySQL version.

#### Examples

##### Example 21.353. `MySQLndUhConnection::getServerVersion` example

```

<?php
class proxy extends MySQLndUhConnection {
    public function getServerVersion($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getServerVersion($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

```



```
$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->server_version);
?>
```

The above example will output:

```
proxy::getServerVersion(array (
  0 => NULL,
))
proxy::getServerVersion returns 50145
int(50145)
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_get\\_server\\_version](#)  
[mysql\\_get\\_server\\_version](#)

### 21.9.8.7.19. [MySQLndUhConnection::getSqlstate](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::getSqlstate](#)

Returns the SQLSTATE error from previous MySQL operation

#### Description

```
public string MySQLndUhConnection::getSqlstate(
    mysqlnd_connection connection);
```

Returns the SQLSTATE error from previous MySQL operation.

#### Parameters

*connection* MySQLnd connection handle. Do not modify!

#### Return Values

The SQLSTATE code.

#### Examples

##### Example 21.354. [MySQLndUhConnection::getSqlstate](#) example

```
<?php
class proxy extends MySQLndUhConnection {
    public function getSqlstate($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getSqlstate($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->sqlstate);
$mysqli->query("AN_INVALID_REQUEST_TO_PROVOKE_AN_ERROR");
var_dump($mysqli->sqlstate);
?>
```

The above example will output:

```
proxy::getSqlstate(array (
  0 => NULL,
))
proxy::getSqlstate returns '00000'
string(5) "00000"
proxy::getSqlstate(array (
  0 => NULL,
))
proxy::getSqlstate returns '42000'
string(5) "42000"
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_sql\\_state](#)

#### 21.9.8.7.20. [MysqlndUhConnection::getStatistics](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getStatistics](#)

Returns statistics about the client connection.

### Description

```
public array MysqlndUhConnection::getStatistics(
    mysqlnd_connection connection);
```

Returns statistics about the client connection.



### Warning

This function is currently not documented; only its argument list is available.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Connection statistics collected by mysqlnd.

### Examples

#### Example 21.355. [MysqlndUhConnection::getStatistics](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getStatistics($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getStatistics($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->get_connection_stats());
?>
```

The above example will output:

```
proxy::getStatistics(array (
  0 => NULL,
))
proxy::getStatistics returns array (
  'bytes_sent' => '73',
  'bytes_received' => '77',
  'packets_sent' => '2',
  'packets_received' => '2',
  'protocol_overhead_in' => '8',
  'protocol_overhead_out' => '8',
  'bytes_received_ok_packet' => '0',
  'bytes_received_eof_packet' => '0',
  'bytes_received_rset_header_packet' => '0',
  'bytes_received_rset_field_meta_packet' => '0',
  'bytes_received_rset_row_packet' => '0',
  'bytes_received_prepare_response_packet' => '0',
  'bytes_received_change_user_packet' => '0',
  'packets_sent_command' => '0',
  'packets_received_ok' => '0',
  'packets_received_eof' => '0',
  'packets_received_rset_header' => '0',
  'packets_received_rset_field_meta' => '0',
  'packets_received_rset_row' => '0',
  'packets_received_prepare_response' => '0',
  'packets_received_change_user' => '0',
  'result_set_queries' => '0',
  'non_result_set_queries' => '0',
  'no_index_used' => '0',
  'bad_index_used' => '0',
  'slow_queries' => '0',
  'buffered_sets' => '0',
  'unbuffered_sets' => '0',
  'ps_buffered_sets' => '0',
  'ps_unbuffered_sets' => '0',
  'flushed_normal_sets' => '0',
  'flushed_ps_sets' => '0',
  'ps_prepared_never_executed' => '0',
  'ps_prepared_once_executed' => '0',
  'rows_fetched_from_server_normal' => '0',
  'rows_fetched_from_server_ps' => '0',
  'rows_buffered_from_client_normal' => '0',
  'rows_buffered_from_client_ps' => '0',
  'rows_fetched_from_client_normal_buffered' => '0',
  'rows_fetched_from_client_normal_unbuffered' => '0',
  'rows_fetched_from_client_ps_buffered' => '0',
  'rows_fetched_from_client_ps_unbuffered' => '0',
  'rows_fetched_from_client_ps_cursor' => '0',
  'rows_affected_normal' => '0',
  'rows_affected_ps' => '0',
  'rows_skipped_normal' => '0',
  'rows_skipped_ps' => '0',
  'copy_on_write_saved' => '0',
  'copy_on_write_performed' => '0',
  'command_buffer_too_small' => '0',
  'connect_success' => '1',
  'connect_failure' => '0',
  'connection_reused' => '0',
  'reconnect' => '0',
  'pconnect_success' => '0',
  'active_connections' => '1',
  'active_persistent_connections' => '0',
  'explicit_close' => '0',
  'implicit_close' => '0',
  'disconnect_close' => '0',
  'in_middle_of_command_close' => '0',
  'explicit_free_result' => '0',
  'implicit_free_result' => '0',
```

```
'explicit_stmt_close' => '0',
'implicit_stmt_close' => '0',
'mem_emalloc_count' => '0',
'mem_emalloc_amount' => '0',
'mem_ecalloc_count' => '0',
'mem_ecalloc_amount' => '0',
'mem_erealloc_count' => '0',
'mem_erealloc_amount' => '0',
'mem_efree_count' => '0',
'mem_efree_amount' => '0',
'mem_malloc_count' => '0',
'mem_malloc_amount' => '0',
'mem_calloc_count' => '0',
'mem_calloc_amount' => '0',
'mem_realloc_count' => '0',
'mem_realloc_amount' => '0',
'mem_free_count' => '0',
'mem_free_amount' => '0',
'mem_strdup_count' => '0',
'mem_strndup_count' => '0',
'mem_estndup_count' => '0',
'mem_strdup_count' => '0',
'proto_text_fetched_null' => '0',
'proto_text_fetched_bit' => '0',
'proto_text_fetched_tinyint' => '0',
'proto_text_fetched_short' => '0',
'proto_text_fetched_int24' => '0',
'proto_text_fetched_int' => '0',
'proto_text_fetched_bigint' => '0',
'proto_text_fetched_decimal' => '0',
'proto_text_fetched_float' => '0',
'proto_text_fetched_double' => '0',
'proto_text_fetched_date' => '0',
'proto_text_fetched_year' => '0',
'proto_text_fetched_time' => '0',
'proto_text_fetched_datetime' => '0',
'proto_text_fetched_timestamp' => '0',
'proto_text_fetched_string' => '0',
'proto_text_fetched_blob' => '0',
'proto_text_fetched_enum' => '0',
'proto_text_fetched_set' => '0',
'proto_text_fetched_geometry' => '0',
'proto_text_fetched_other' => '0',
'proto_binary_fetched_null' => '0',
'proto_binary_fetched_bit' => '0',
'proto_binary_fetched_tinyint' => '0',
'proto_binary_fetched_short' => '0',
'proto_binary_fetched_int24' => '0',
'proto_binary_fetched_int' => '0',
'proto_binary_fetched_bigint' => '0',
'proto_binary_fetched_decimal' => '0',
'proto_binary_fetched_float' => '0',
'proto_binary_fetched_double' => '0',
'proto_binary_fetched_date' => '0',
'proto_binary_fetched_year' => '0',
'proto_binary_fetched_time' => '0',
'proto_binary_fetched_datetime' => '0',
'proto_binary_fetched_timestamp' => '0',
'proto_binary_fetched_string' => '0',
'proto_binary_fetched_blob' => '0',
'proto_binary_fetched_enum' => '0',
'proto_binary_fetched_set' => '0',
'proto_binary_fetched_geometry' => '0',
'proto_binary_fetched_other' => '0',
'init_command_executed_count' => '0',
'init_command_failed_count' => '0',
'com_quit' => '0',
'com_init_db' => '0',
'com_query' => '0',
'com_field_list' => '0',
'com_create_db' => '0',
'com_drop_db' => '0',
```

```

'com_refresh' => '0',
'com_shutdown' => '0',
'com_statistics' => '0',
'com_process_info' => '0',
'com_connect' => '0',
'com_process_kill' => '0',
'com_debug' => '0',
'com_ping' => '0',
'com_time' => '0',
'com_delayed_insert' => '0',
'com_change_user' => '0',
'com_binlog_dump' => '0',
'com_table_dump' => '0',
'com_connect_out' => '0',
'com_register_slave' => '0',
'com_stmt_prepare' => '0',
'com_stmt_execute' => '0',
'com_stmt_send_long_data' => '0',
'com_stmt_close' => '0',
'com_stmt_reset' => '0',
'com_stmt_set_option' => '0',
'com_stmt_fetch' => '0',
'com_daemon' => '0',
'bytes_received_real_data_normal' => '0',
'bytes_received_real_data_ps' => '0',
)
array(160) {
  ["bytes_sent"]=>
  string(2) "73"
  ["bytes_received"]=>
  string(2) "77"
  ["packets_sent"]=>
  string(1) "2"
  ["packets_received"]=>
  string(1) "2"
  ["protocol_overhead_in"]=>
  string(1) "8"
  ["protocol_overhead_out"]=>
  string(1) "8"
  ["bytes_received_ok_packet"]=>
  string(1) "0"
  ["bytes_received_eof_packet"]=>
  string(1) "0"
  ["bytes_received_rset_header_packet"]=>
  string(1) "0"
  ["bytes_received_rset_field_meta_packet"]=>
  string(1) "0"
  ["bytes_received_rset_row_packet"]=>
  string(1) "0"
  ["bytes_received_prepare_response_packet"]=>
  string(1) "0"
  ["bytes_received_change_user_packet"]=>
  string(1) "0"
  ["packets_sent_command"]=>
  string(1) "0"
  ["packets_received_ok"]=>
  string(1) "0"
  ["packets_received_eof"]=>
  string(1) "0"
  ["packets_received_rset_header"]=>
  string(1) "0"
  ["packets_received_rset_field_meta"]=>
  string(1) "0"
  ["packets_received_rset_row"]=>
  string(1) "0"
  ["packets_received_prepare_response"]=>
  string(1) "0"
  ["packets_received_change_user"]=>
  string(1) "0"
  ["result_set_queries"]=>
  string(1) "0"
  ["non_result_set_queries"]=>

```

```
string(1) "0"  
["no_index_used"]=>  
string(1) "0"  
["bad_index_used"]=>  
string(1) "0"  
["slow_queries"]=>  
string(1) "0"  
["buffered_sets"]=>  
string(1) "0"  
["unbuffered_sets"]=>  
string(1) "0"  
["ps_buffered_sets"]=>  
string(1) "0"  
["ps_unbuffered_sets"]=>  
string(1) "0"  
["flushed_normal_sets"]=>  
string(1) "0"  
["flushed_ps_sets"]=>  
string(1) "0"  
["ps_prepared_never_executed"]=>  
string(1) "0"  
["ps_prepared_once_executed"]=>  
string(1) "0"  
["rows_fetched_from_server_normal"]=>  
string(1) "0"  
["rows_fetched_from_server_ps"]=>  
string(1) "0"  
["rows_buffered_from_client_normal"]=>  
string(1) "0"  
["rows_buffered_from_client_ps"]=>  
string(1) "0"  
["rows_fetched_from_client_normal_buffered"]=>  
string(1) "0"  
["rows_fetched_from_client_normal_unbuffered"]=>  
string(1) "0"  
["rows_fetched_from_client_ps_buffered"]=>  
string(1) "0"  
["rows_fetched_from_client_ps_unbuffered"]=>  
string(1) "0"  
["rows_fetched_from_client_ps_cursor"]=>  
string(1) "0"  
["rows_affected_normal"]=>  
string(1) "0"  
["rows_affected_ps"]=>  
string(1) "0"  
["rows_skipped_normal"]=>  
string(1) "0"  
["rows_skipped_ps"]=>  
string(1) "0"  
["copy_on_write_saved"]=>  
string(1) "0"  
["copy_on_write_performed"]=>  
string(1) "0"  
["command_buffer_too_small"]=>  
string(1) "0"  
["connect_success"]=>  
string(1) "1"  
["connect_failure"]=>  
string(1) "0"  
["connection_reused"]=>  
string(1) "0"  
["reconnect"]=>  
string(1) "0"  
["pconnect_success"]=>  
string(1) "0"  
["active_connections"]=>  
string(1) "1"  
["active_persistent_connections"]=>  
string(1) "0"  
["explicit_close"]=>  
string(1) "0"  
["implicit_close"]=>
```

```
string(1) "0"
["disconnect_close"]=>
string(1) "0"
["in_middle_of_command_close"]=>
string(1) "0"
["explicit_free_result"]=>
string(1) "0"
["implicit_free_result"]=>
string(1) "0"
["explicit_stmt_close"]=>
string(1) "0"
["implicit_stmt_close"]=>
string(1) "0"
["mem_emalloc_count"]=>
string(1) "0"
["mem_emalloc_amount"]=>
string(1) "0"
["mem_ecalloc_count"]=>
string(1) "0"
["mem_ecalloc_amount"]=>
string(1) "0"
["mem_erealloc_count"]=>
string(1) "0"
["mem_erealloc_amount"]=>
string(1) "0"
["mem_efree_count"]=>
string(1) "0"
["mem_efree_amount"]=>
string(1) "0"
["mem_malloc_count"]=>
string(1) "0"
["mem_malloc_amount"]=>
string(1) "0"
["mem_calloc_count"]=>
string(1) "0"
["mem_calloc_amount"]=>
string(1) "0"
["mem_realloc_count"]=>
string(1) "0"
["mem_realloc_amount"]=>
string(1) "0"
["mem_free_count"]=>
string(1) "0"
["mem_free_amount"]=>
string(1) "0"
["mem_strdup_count"]=>
string(1) "0"
["mem_strndup_count"]=>
string(1) "0"
["mem_estndup_count"]=>
string(1) "0"
["mem_strdup_count"]=>
string(1) "0"
["proto_text_fetched_null"]=>
string(1) "0"
["proto_text_fetched_bit"]=>
string(1) "0"
["proto_text_fetched_tinyint"]=>
string(1) "0"
["proto_text_fetched_short"]=>
string(1) "0"
["proto_text_fetched_int24"]=>
string(1) "0"
["proto_text_fetched_int"]=>
string(1) "0"
["proto_text_fetched_bigint"]=>
string(1) "0"
["proto_text_fetched_decimal"]=>
string(1) "0"
["proto_text_fetched_float"]=>
string(1) "0"
["proto_text_fetched_double"]=>
```

```

string(1) "0"
["proto_text_fetched_date"]=>
string(1) "0"
["proto_text_fetched_year"]=>
string(1) "0"
["proto_text_fetched_time"]=>
string(1) "0"
["proto_text_fetched_datetime"]=>
string(1) "0"
["proto_text_fetched_timestamp"]=>
string(1) "0"
["proto_text_fetched_string"]=>
string(1) "0"
["proto_text_fetched_blob"]=>
string(1) "0"
["proto_text_fetched_enum"]=>
string(1) "0"
["proto_text_fetched_set"]=>
string(1) "0"
["proto_text_fetched_geometry"]=>
string(1) "0"
["proto_text_fetched_other"]=>
string(1) "0"
["proto_binary_fetched_null"]=>
string(1) "0"
["proto_binary_fetched_bit"]=>
string(1) "0"
["proto_binary_fetched_tinyint"]=>
string(1) "0"
["proto_binary_fetched_short"]=>
string(1) "0"
["proto_binary_fetched_int24"]=>
string(1) "0"
["proto_binary_fetched_int"]=>
string(1) "0"
["proto_binary_fetched_bigint"]=>
string(1) "0"
["proto_binary_fetched_decimal"]=>
string(1) "0"
["proto_binary_fetched_float"]=>
string(1) "0"
["proto_binary_fetched_double"]=>
string(1) "0"
["proto_binary_fetched_date"]=>
string(1) "0"
["proto_binary_fetched_year"]=>
string(1) "0"
["proto_binary_fetched_time"]=>
string(1) "0"
["proto_binary_fetched_datetime"]=>
string(1) "0"
["proto_binary_fetched_timestamp"]=>
string(1) "0"
["proto_binary_fetched_string"]=>
string(1) "0"
["proto_binary_fetched_blob"]=>
string(1) "0"
["proto_binary_fetched_enum"]=>
string(1) "0"
["proto_binary_fetched_set"]=>
string(1) "0"
["proto_binary_fetched_geometry"]=>
string(1) "0"
["proto_binary_fetched_other"]=>
string(1) "0"
["init_command_executed_count"]=>
string(1) "0"
["init_command_failed_count"]=>
string(1) "0"
["com_quit"]=>
string(1) "0"
["com_init_db"]=>

```



```

string(1) "0"
["com_query"]=>
string(1) "0"
["com_field_list"]=>
string(1) "0"
["com_create_db"]=>
string(1) "0"
["com_drop_db"]=>
string(1) "0"
["com_refresh"]=>
string(1) "0"
["com_shutdown"]=>
string(1) "0"
["com_statistics"]=>
string(1) "0"
["com_process_info"]=>
string(1) "0"
["com_connect"]=>
string(1) "0"
["com_process_kill"]=>
string(1) "0"
["com_debug"]=>
string(1) "0"
["com_ping"]=>
string(1) "0"
["com_time"]=>
string(1) "0"
["com_delayed_insert"]=>
string(1) "0"
["com_change_user"]=>
string(1) "0"
["com_binlog_dump"]=>
string(1) "0"
["com_table_dump"]=>
string(1) "0"
["com_connect_out"]=>
string(1) "0"
["com_register_slave"]=>
string(1) "0"
["com_stmt_prepare"]=>
string(1) "0"
["com_stmt_execute"]=>
string(1) "0"
["com_stmt_send_long_data"]=>
string(1) "0"
["com_stmt_close"]=>
string(1) "0"
["com_stmt_reset"]=>
string(1) "0"
["com_stmt_set_option"]=>
string(1) "0"
["com_stmt_fetch"]=>
string(1) "0"
["com_daemon"]=>
string(1) "0"
["bytes_received_real_data_normal"]=>
string(1) "0"
["bytes_received_real_data_ps"]=>
string(1) "0"
}

```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_get\\_connection\\_stats](#)

**21.9.8.7.21. MySQLndUhConnection::getThreadId**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getThreadId](#)

Returns the thread ID for the current connection

### Description

```
public int MysqlndUhConnection::getThreadId(
    mysqlnd_connection connection);
```

Returns the thread ID for the current connection.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Connection thread id.

### Examples

#### Example 21.356. [MysqlndUhConnection::getThreadId](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getThreadId($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getThreadId($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->thread_id);
?>
```

The above example will output:

```
proxy::getThreadId(array (
    0 => NULL,
))
proxy::getThreadId returns 27646
int(27646)
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_thread\\_id](#)  
[mysql\\_thread\\_id](#)

#### 21.9.8.7.22. [MysqlndUhConnection::getWarningCount](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::getWarningCount](#)

Returns the number of warnings from the last query for the given link

## Description

```
public int MysqlndUhConnection::getWarningCount(
    mysqlnd_connection connection);
```

Returns the number of warnings from the last query for the given link.

## Parameters

[connection](#) Mysqlnd connection handle. Do not modify!

## Return Values

Number of warnings.

## Examples

### Example 21.357. [MysqlndUhConnection::getWarningCount](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function getWarningCount($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::getWarningCount($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
var_dump($mysqli->warning_count);
?>
```

The above example will output:

```
proxy::getWarningCount(array (
    0 => NULL,
))
proxy::getWarningCount returns 0
int(0)
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_warning\\_count](#)

### 21.9.8.7.23. [MysqlndUhConnection::init](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::init](#)

Initialize mysqlnd connection

## Description

```
public bool MysqlndUhConnection::init(
    mysqlnd_connection connection);
```

Initialize mysqlnd connection. This is an mysqlnd internal call to initialize the connection object.

**Note**

Failing to call the parent implementation may cause memory leaks or crash PHP. This is not considered a bug. Please, keep in mind that the [mysqlnd](#) library functions have never been designed to be exposed to the user space.

**Parameters**

[connection](#) Mysqlnd connection handle. Do not modify!

**Return Values**

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

**Examples****Example 21.358. [MysqlndUhConnection::init](#) example**

```
<?php
class proxy extends MysqlndUhConnection {
    public function init($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::init($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
?>
```

The above example will output:

```
proxy::init(array (
    0 => NULL,
))
proxy::init returns true
```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

**21.9.8.7.24. [MysqlndUhConnection::killConnection](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::killConnection](#)

Asks the server to kill a MySQL thread

**Description**

```
public bool MysqlndUhConnection::killConnection(
    mysqlnd_connection connection,
    int pid);
```

Asks the server to kill a MySQL thread.

**Parameters**

`connection` Mysqlnd connection handle. Do not modify!

`pid` Thread Id of the connection to be killed.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.359. `MysqlndUhConnection::kill` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function killConnection($res, $pid) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::killConnection($res, $pid);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->kill($mysqli->thread_id);
?>
```

The above example will output:

```
proxy::killConnection(array (
    0 => NULL,
    1 => 27650,
))
proxy::killConnection returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_kill](#)

#### 21.9.8.7.25. `MysqlndUhConnection::listFields`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::listFields`

List MySQL table fields

### Description

```
public array MysqlndUhConnection::listFields(
    mysqlnd_connection connection,
    string table,
    string achtung_wild);
```

List MySQL table fields.



### Warning

This function is currently not documented; only its argument list is available.

## Parameters

<i>connection</i>	Mysqlnd connection handle. Do not modify!
<i>table</i>	The name of the table that's being queried.
<i>pattern</i>	Name pattern.

## Return Values

## Examples

### Example 21.360. `MysqlndUhConnection::listFields` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function listFields($res, $table, $pattern) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::listFields($res, $table, $pattern);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysql = mysql_connect("localhost", "root", "");
mysql_select_db("test", $mysql);
mysql_query("DROP TABLE IF EXISTS test_a", $mysql);
mysql_query("CREATE TABLE test_a(id INT, coll VARCHAR(255))", $mysql);
$res = mysql_list_fields("test", "test_a", $mysql);
printf("num_rows = %d\n", mysql_num_rows($res));
while ($row = mysql_fetch_assoc($res))
    var_dump($row);
?>
```

The above example will output:

```
proxy::listFields(array (
  0 => NULL,
  1 => 'test_a',
  2 => '',
))
proxy::listFields returns NULL
num_rows = 0
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysql\\_list\\_fields](#)

### 21.9.8.7.26. `MysqlndUhConnection::listMethod`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::listMethod](#)

Wrapper for assorted list commands

## Description

```
public void MysqlndUhConnection::listMethod(
```

```
mysqlnd_connection connection,
string query,
string achtung_wild,
string par1);
```

Wrapper for assorted list commands.



### Warning

This function is currently not documented; only its argument list is available.

### Parameters

<code>connection</code>	Mysqlnd connection handle. Do not modify!
<code>query</code>	SHOW command to be executed.
<code>achtung_wild</code>	
<code>par1</code>	

### Return Values

### Return Values

TODO

### Examples

#### Example 21.361. `MysqlndUhConnection::listMethod` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function listMethod($res, $query, $pattern, $par1) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::listMethod($res, $query, $pattern, $par1);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysql = mysql_connect("localhost", "root", "");
$res = mysql_list_dbs($mysql);
printf("num_rows = %d\n", mysql_num_rows($res));
while ($row = mysql_fetch_assoc($res))
    var_dump($row);
?>
```

The above example will output:

```
proxy::listMethod(array (
    0 => NULL,
    1 => 'SHOW DATABASES',
    2 => '',
    3 => '',
))
proxy::listMethod returns NULL
num_rows = 6
array(1) {
    ["Database"]=>
    string(18) "information_schema"
```

```
}
array(1) {
  ["Database"]=>
  string(5) "mysql"
}
array(1) {
  ["Database"]=>
  string(8) "oxid_new"
}
array(1) {
  ["Database"]=>
  string(7) "phptest"
}
array(1) {
  ["Database"]=>
  string(7) "pushphp"
}
array(1) {
  ["Database"]=>
  string(4) "test"
}
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysql\\_list\\_dbs](#)

#### 21.9.8.7.27. `MysqlndUhConnection::moreResults`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::moreResults`

Check if there are any more query results from a multi query

### Description

```
public bool MysqlndUhConnection::moreResults(
    mysqlnd_connection connection);
```

Check if there are any more query results from a multi query.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.362. `MysqlndUhConnection::moreResults` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function moreResults($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::moreResults($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
```



```

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->multi_query("SELECT 1 AS _one; SELECT 2 AS _two");
do {
    $res = $mysqli->store_result();
    var_dump($res->fetch_assoc());
    printf("%s\n", str_repeat("-", 40));
} while ($mysqli->more_results() && $mysqli->next_result());
?>

```

The above example will output:

```

array(1) {
    ["_one"]=>
        string(1) "1"
}
-----
proxy::moreResults(array (
    0 => NULL,
))
proxy::moreResults returns true
proxy::moreResults(array (
    0 => NULL,
))
proxy::moreResults returns true
array(1) {
    ["_two"]=>
        string(1) "2"
}
-----
proxy::moreResults(array (
    0 => NULL,
))
proxy::moreResults returns false

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_more\\_results](#)

#### 21.9.8.7.28. `MysqlndUhConnection::nextResult`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::nextResult](#)

Prepare next result from `multi_query`

### Description

```

public bool MysqlndUhConnection::nextResult(
    mysqlnd_connection connection);

```

Prepare next result from `multi_query`.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

**Example 21.363. [MySQLndUhConnection::nextResult](#) example**

```
<?php
class proxy extends MySQLndUhConnection {
    public function nextResult($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::nextResult($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->multi_query("SELECT 1 AS _one; SELECT 2 AS _two");
do {
    $res = $mysqli->store_result();
    var_dump($res->fetch_assoc());
    printf("%s\n", str_repeat("-", 40));
} while ($mysqli->more_results() && $mysqli->next_result());
?>
```

The above example will output:

```
array(1) {
    ["_one"]=>
    string(1) "1"
}
-----
proxy::nextResult(array (
    0 => NULL,
))
proxy::nextResult returns true
array(1) {
    ["_two"]=>
    string(1) "2"
}
-----
```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_next\\_result](#)

**21.9.8.7.29. [MySQLndUhConnection::ping](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::ping](#)

Pings a server connection, or tries to reconnect if the connection has gone down

**Description**

```
public bool MySQLndUhConnection::ping(
    mysqlnd_connection connection);
```

Pings a server connection, or tries to reconnect if the connection has gone down.

**Parameters**

*connection* MySQLnd connection handle. Do not modify!

## Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.364. [MysqlndUhConnection::ping](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function ping($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::ping($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->ping();
?>
```

The above example will output:

```
proxy::ping(array (
  0 => NULL,
))
proxy::ping returns true
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

[mysqli\\_ping](#)

[mysql\\_ping](#)

### 21.9.8.7.30. [MysqlndUhConnection::query](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::query](#)

Performs a query on the database

## Description

```
public bool MysqlndUhConnection::query(
    mysqlnd_connection connection,
    string query);
```

Performs a query on the database (COM\_QUERY).

## Parameters

*connection* MySQLnd connection handle. Do not modify!

*query* The query string.

## Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.365. [MysqlndUhConnection::query](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function query($res, $query) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $query = "SELECT 'How about query rewriting?'";
        $ret = parent::query($res, $query);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$res = $mysqli->query("SELECT 'Welcome mysqlnd_uh!' FROM DUAL");
var_dump($res->fetch_assoc());
?>
```

The above example will output:

```
proxy::query(array (
    0 => NULL,
    1 => 'SELECT \'Welcome mysqlnd_uh!\' FROM DUAL',
))
proxy::query returns true
array(1) {
    ["How about query rewriting?"]=>
    string(26) "How about query rewriting?"
}
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_query](#)  
[mysql\\_query](#)

### 21.9.8.7.31. [MysqlndUhConnection::queryReadResultsetHeader](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::queryReadResultsetHeader](#)

Read a result set header

## Description

```
public bool MysqlndUhConnection::queryReadResultsetHeader(
    mysqlnd_connection connection,
    mysqlnd_statement mysqlnd_stmt);
```

Read a result set header.

## Parameters

*connection*

Mysqlnd connection handle. Do not modify!

*mysqlnd\_stmt*

Mysqlnd statement handle. Do not modify! Set to [NULL](#) , if function is not used in the context of a prepared statement.

## Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.366. [MySQLndUhConnection::queryReadResultsetHeader](#) example

```
<?php
class proxy extends MySQLndUhConnection {
    public function queryReadResultsetHeader($res, $stmt) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::queryReadResultsetHeader($res, $stmt);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$res = $mysqli->query("SELECT 'Welcome mysqlnd_uh!' FROM DUAL");
var_dump($res->fetch_assoc());
?>
```

The above example will output:

```
proxy::queryReadResultsetHeader(array (
  0 => NULL,
  1 => NULL,
))
proxy::queryReadResultsetHeader returns true
array(1) {
  ["Welcome mysqlnd_uh!"]=>
    string(19) "Welcome mysqlnd_uh!"
}
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

### 21.9.8.7.32. [MySQLndUhConnection::reapQuery](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::reapQuery](#)

Get result from async query

## Description

```
public bool MySQLndUhConnection::reapQuery(
    mysqlnd_connection connection);
```

Get result from async query.

## Parameters

*connection*

MySQLnd connection handle. Do not modify!

## Return Values

Returns **TRUE** on success. Otherwise, returns **FALSE**

## Examples

### Example 21.367. MysqlndUhConnection::reapQuery example

```
<?php
class proxy extends MysqlndUhConnection {
    public function reapQuery($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::reapQuery($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$conn1 = new mysqli("localhost", "root", "", "test");
$conn2 = new mysqli("localhost", "root", "", "test");

$conn1->query("SELECT 1 as 'one', SLEEP(1) AS _sleep FROM DUAL", MYSQLI_ASYNC | MYSQLI_USE_RESULT);
$conn2->query("SELECT 1.1 as 'one dot one' FROM DUAL", MYSQLI_ASYNC | MYSQLI_USE_RESULT);

$links = array(
    $conn1->thread_id => array('link' => $conn1, 'processed' => false),
    $conn2->thread_id => array('link' => $conn2, 'processed' => false)
);

$saved_errors = array();
do {
    $poll_links = $poll_errors = $poll_reject = array();
    foreach ($links as $thread_id => $link) {
        if (!$link['processed']) {
            $poll_links[] = $link['link'];
            $poll_errors[] = $link['link'];
            $poll_reject[] = $link['link'];
        }
    }
    if (0 == count($poll_links))
        break;

    if (0 == ($num_ready = mysqli_poll($poll_links, $poll_errors, $poll_reject, 0, 200000)))
        continue;

    if (!empty($poll_errors)) {
        die(var_dump($poll_errors));
    }

    foreach ($poll_links as $link) {
        $thread_id = mysqli_thread_id($link);
        $links[$thread_id]['processed'] = true;

        if (is_object($res = mysqli_reap_async_query($link))) {
            // result set object
            while ($row = mysqli_fetch_assoc($res)) {
                // eat up all results
                var_dump($row);
            }
            mysqli_free_result($res);
        } else {
            // either there is no result (no SELECT) or there is an error
            if (mysqli_errno($link) > 0) {
                $saved_errors[$thread_id] = mysqli_errno($link);
                printf("%s' caused %d\n", $links[$thread_id]['query'],    mysqli_errno($link));
            }
        }
    }
} while (true);
?>
```

The above example will output:

```
proxy::reapQuery(array (
  0 => NULL,
))
proxy::reapQuery returns true
array(1) {
  ["one dot one"]=>
  string(3) "1.1"
}
proxy::reapQuery(array (
  0 => NULL,
))
proxy::reapQuery returns true
array(2) {
  ["one"]=>
  string(1) "1"
  ["_sleep"]=>
  string(1) "0"
}
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_real\\_async\\_query](#)

### 21.9.8.7.33. `MySQLndUhConnection::refreshServer`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MySQLndUhConnection::refreshServer`

Flush or reset tables and caches

#### Description

```
public bool MySQLndUhConnection::refreshServer(
    mysqlnd_connection connection,
    int options);
```

Flush or reset tables and caches.



#### Warning

This function is currently not documented; only its argument list is available.

#### Parameters

*connection* MySQLnd connection handle. Do not modify!

*options* What to refresh.

#### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

#### Examples

##### Example 21.368. `MySQLndUhConnection::refreshServer` example

```
<?php
```

```

class proxy extends MysqlndUhConnection {
    public function refreshServer($res, $option) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::refreshServer($res, $option);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
mysqli_refresh($mysqli, 1);
?>

```

The above example will output:

```

proxy::refreshServer(array (
    0 => NULL,
    1 => 1,
))
proxy::refreshServer returns false

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

#### 21.9.8.7.34. [MysqlndUhConnection::restartPSession](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::restartPSession](#)

Restart a persistent mysqlnd connection

### Description

```

public bool MysqlndUhConnection::restartPSession(
    mysqlnd_connection connection);

```

Restart a persistent mysqlnd connection.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

### Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

### Examples

#### Example 21.369. [MysqlndUhConnection::restartPSession](#) example

```

<?php
class proxy extends MysqlndUhConnection {
    public function ping($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::ping($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}

```



```
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->ping();
?>
```

The above example will output:

```
proxy::restartPSession(array (
  0 => NULL,
))
proxy::restartPSession returns true
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

### 21.9.8.7.35. [MysqlndUhConnection::selectDb](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::selectDb](#)

Selects the default database for database queries

## Description

```
public bool MysqlndUhConnection::selectDb(
    mysqlnd_connection connection,
    string database);
```

Selects the default database for database queries.

## Parameters

<i>connection</i>	MySQLnd connection handle. Do not modify!
<i>database</i>	The database name.

## Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.370. [MysqlndUhConnection::selectDb](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function selectDb($res, $database) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::selectDb($res, $database);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->select_db("mysql");
?>
```

The above example will output:

```
proxy::selectDb(array (
  0 => NULL,
  1 => 'mysql',
))
proxy::selectDb returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_select\\_db](#)  
[mysql\\_select\\_db](#)

### 21.9.8.7.36. `MysqlndUhConnection::sendClose`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::sendClose`

Sends a close command to MySQL

### Description

```
public bool MysqlndUhConnection::sendClose(
    mysqlnd_connection connection);
```

Sends a close command to MySQL.

### Parameters

*connection* MySQLnd connection handle. Do not modify!

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.371. `MysqlndUhConnection::sendClose` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function sendClose($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::sendClose($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->close();
?>
```

The above example will output:

```

proxy::sendClose(array (
    0 => NULL,
))
proxy::sendClose returns true
proxy::sendClose(array (
    0 => NULL,
))
proxy::sendClose returns true

```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

**21.9.8.7.37. MysqlndUhConnection::sendQuery**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::sendQuery](#)

Sends a query to MySQL

**Description**

```

public bool MysqlndUhConnection::sendQuery(
    mysqlnd_connection connection,
    string query);

```

Sends a query to MySQL.

**Parameters**

*connection* MySQLnd connection handle. Do not modify!

*query* The query string.

**Return Values**

Returns `TRUE` on success. Otherwise, returns `FALSE`

**Examples****Example 21.372. MysqlndUhConnection::sendQuery example**

```

<?php
class proxy extends MysqlndUhConnection {
    public function sendQuery($res, $query) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::sendQuery($res, $query);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("SELECT 1");
?>

```

The above example will output:

```

proxy::sendQuery(array (
    0 => NULL,

```

```
1 => 'SELECT 1',
))
proxy::sendQuery returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

### 21.9.8.7.38. [MySQLndUhConnection::serverDumpDebugInformation](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::serverDumpDebugInformation](#)

Dump debugging information into the log for the MySQL server

### Description

```
public bool MySQLndUhConnection::serverDumpDebugInformation(
    mysqlnd_connection connection);
```

Dump debugging information into the log for the MySQL server.

### Parameters

*connection* MySQLnd connection handle. Do not modify!

### Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

### Examples

#### Example 21.373. [MySQLndUhConnection::serverDumpDebugInformation](#) example

```
<?php
class proxy extends MySQLndUhConnection {
    public function serverDumpDebugInformation($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::serverDumpDebugInformation($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->dump_debug_info();
?>
```

The above example will output:

```
proxy::serverDumpDebugInformation(array (
    0 => NULL,
))
proxy::serverDumpDebugInformation returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

`mysqli_dump_debug_info`

### 21.9.8.7.39. `MysqlndUhConnection::setAutocommit`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::setAutocommit`

Turns on or off auto-committing database modifications

#### Description

```
public bool MysqlndUhConnection::setAutocommit(
    mysqlnd_connection connection,
    int mode);
```

Turns on or off auto-committing database modifications

#### Parameters

<code>connection</code>	Mysqlnd connection handle. Do not modify!
<code>mode</code>	Whether to turn on auto-commit or not.

#### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

#### Examples

##### Example 21.374. `MysqlndUhConnection::setAutocommit` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function setAutocommit($res, $mode) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::setAutocommit($res, $mode);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->autocommit(false);
$mysqli->autocommit(true);
?>
```

The above example will output:

```
proxy::setAutocommit(array (
    0 => NULL,
    1 => 0,
))
proxy::setAutocommit returns true
proxy::setAutocommit(array (
    0 => NULL,
    1 => 1,
))
proxy::setAutocommit returns true
```

#### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_autocommit](#)

#### 21.9.8.7.40. [MySQLndUhConnection::setCharset](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MySQLndUhConnection::setCharset](#)

Sets the default client character set

#### Description

```
public bool MySQLndUhConnection::setCharset(
    mysqlnd_connection connection,
    string charset);
```

Sets the default client character set.

#### Parameters

<i>connection</i>	MySQLnd connection handle. Do not modify!
<i>charset</i>	The charset to be set as default.

#### Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

#### Examples

##### Example 21.375. [MySQLndUhConnection::setCharset](#) example

```
<?php
class proxy extends MySQLndUhConnection {
    public function setCharset($res, $charset) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::setCharset($res, $charset);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->set_charset("latin1");
?>
```

The above example will output:

```
proxy::setCharset(array (
    0 => NULL,
    1 => 'latin1',
))
proxy::setCharset returns true
```

#### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_set\\_charset](#)

### 21.9.8.7.41. `MysqlndUhConnection::setClientOption`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::setClientOption`

Sets a client option

#### Description

```
public bool MysqlndUhConnection::setClientOption(
    mysqlnd_connection connection,
    int option,
    int value);
```

Sets a client option.

#### Parameters

<code>connection</code>	Mysqlnd connection handle. Do not modify!
<code>option</code>	The option to be set.
<code>value</code>	Optional option value, if required.

#### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

#### Examples

##### Example 21.376. `MysqlndUhConnection::setClientOption` example

```
<?php
function client_option_to_string($option) {
    static $mapping = array(
        MYSQLND_UH_MYSQLND_OPTION_OPT_CONNECT_TIMEOUT => "MYSQLND_UH_MYSQLND_OPTION_OPT_CONNECT_TIMEOUT",
        MYSQLND_UH_MYSQLND_OPTION_OPT_COMPRESS => "MYSQLND_UH_MYSQLND_OPTION_OPT_COMPRESS",
        MYSQLND_UH_MYSQLND_OPTION_OPT_NAMED_PIPE => "MYSQLND_UH_MYSQLND_OPTION_OPT_NAMED_PIPE",
        MYSQLND_UH_MYSQLND_OPTION_INIT_COMMAND => "MYSQLND_UH_MYSQLND_OPTION_INIT_COMMAND",
        MYSQLND_UH_MYSQLND_READ_DEFAULT_FILE => "MYSQLND_UH_MYSQLND_READ_DEFAULT_FILE",
        MYSQLND_UH_MYSQLND_READ_DEFAULT_GROUP => "MYSQLND_UH_MYSQLND_READ_DEFAULT_GROUP",
        MYSQLND_UH_MYSQLND_SET_CHARSET_DIR => "MYSQLND_UH_MYSQLND_SET_CHARSET_DIR",
        MYSQLND_UH_MYSQLND_SET_CHARSET_NAME => "MYSQLND_UH_MYSQLND_SET_CHARSET_NAME",
        MYSQLND_UH_MYSQLND_OPT_LOCAL_INFILE => "MYSQLND_UH_MYSQLND_OPT_LOCAL_INFILE",
        MYSQLND_UH_MYSQLND_OPT_PROTOCOL => "MYSQLND_UH_MYSQLND_OPT_PROTOCOL",
        MYSQLND_UH_MYSQLND_SHARED_MEMORY_BASE_NAME => "MYSQLND_UH_MYSQLND_SHARED_MEMORY_BASE_NAME",
        MYSQLND_UH_MYSQLND_OPT_READ_TIMEOUT => "MYSQLND_UH_MYSQLND_OPT_READ_TIMEOUT",
        MYSQLND_UH_MYSQLND_OPT_WRITE_TIMEOUT => "MYSQLND_UH_MYSQLND_OPT_WRITE_TIMEOUT",
        MYSQLND_UH_MYSQLND_OPT_USE_RESULT => "MYSQLND_UH_MYSQLND_OPT_USE_RESULT",
        MYSQLND_UH_MYSQLND_OPT_USE_REMOTE_CONNECTION => "MYSQLND_UH_MYSQLND_OPT_USE_REMOTE_CONNECTION",
        MYSQLND_UH_MYSQLND_OPT_USE_EMBEDDED_CONNECTION => "MYSQLND_UH_MYSQLND_OPT_USE_EMBEDDED_CONNECTION",
        MYSQLND_UH_MYSQLND_OPT_GUESS_CONNECTION => "MYSQLND_UH_MYSQLND_OPT_GUESS_CONNECTION",
        MYSQLND_UH_MYSQLND_SET_CLIENT_IP => "MYSQLND_UH_MYSQLND_SET_CLIENT_IP",
        MYSQLND_UH_MYSQLND_SECURE_AUTH => "MYSQLND_UH_MYSQLND_SECURE_AUTH",
        MYSQLND_UH_MYSQLND_REPORT_DATA_TRUNCATION => "MYSQLND_UH_MYSQLND_REPORT_DATA_TRUNCATION",
        MYSQLND_UH_MYSQLND_OPT_RECONNECT => "MYSQLND_UH_MYSQLND_OPT_RECONNECT",
        MYSQLND_UH_MYSQLND_OPT_SSL_VERIFY_SERVER_CERT => "MYSQLND_UH_MYSQLND_OPT_SSL_VERIFY_SERVER_CERT",
        MYSQLND_UH_MYSQLND_OPT_NET_CMD_BUFFER_SIZE => "MYSQLND_UH_MYSQLND_OPT_NET_CMD_BUFFER_SIZE",
        MYSQLND_UH_MYSQLND_OPT_NET_READ_BUFFER_SIZE => "MYSQLND_UH_MYSQLND_OPT_NET_READ_BUFFER_SIZE",
        MYSQLND_UH_MYSQLND_OPT_SSL_KEY => "MYSQLND_UH_MYSQLND_OPT_SSL_KEY",
        MYSQLND_UH_MYSQLND_OPT_SSL_CERT => "MYSQLND_UH_MYSQLND_OPT_SSL_CERT",
        MYSQLND_UH_MYSQLND_OPT_SSL_CA => "MYSQLND_UH_MYSQLND_OPT_SSL_CA",
        MYSQLND_UH_MYSQLND_OPT_SSL_CAPATH => "MYSQLND_UH_MYSQLND_OPT_SSL_CAPATH",
        MYSQLND_UH_MYSQLND_OPT_SSL_CIPHER => "MYSQLND_UH_MYSQLND_OPT_SSL_CIPHER",
        MYSQLND_UH_MYSQLND_OPT_SSL_PASSPHRASE => "MYSQLND_UH_MYSQLND_OPT_SSL_PASSPHRASE",
        MYSQLND_UH_SERVER_OPTION_PLUGIN_DIR => "MYSQLND_UH_SERVER_OPTION_PLUGIN_DIR",
        MYSQLND_UH_SERVER_OPTION_DEFAULT_AUTH => "MYSQLND_UH_SERVER_OPTION_DEFAULT_AUTH",
```

```

    MYSQLND_UH_SERVER_OPTION_SET_CLIENT_IP => "MYSQLND_UH_SERVER_OPTION_SET_CLIENT_IP"
);
if (version_compare(PHP_VERSION, '5.3.99-dev', '>')) {
    $mapping[MYSQLND_UH_MYSQLND_OPT_MAX_ALLOWED_PACKET] = "MYSQLND_UH_MYSQLND_OPT_MAX_ALLOWED_PACKET";
    $mapping[MYSQLND_UH_MYSQLND_OPT_AUTH_PROTOCOL] = "MYSQLND_UH_MYSQLND_OPT_AUTH_PROTOCOL";
}
if (defined("MYSQLND_UH_MYSQLND_OPT_INT_AND_FLOAT_NATIVE")) {
    /* special mysqlnd build */
    $mapping["MYSQLND_UH_MYSQLND_OPT_INT_AND_FLOAT_NATIVE"] = "MYSQLND_UH_MYSQLND_OPT_INT_AND_FLOAT_NATIVE";
}
return (isset($mapping[$option])) ? $mapping[$option] : 'unknown';
}

class proxy extends MysqlndUhConnection {
    public function setClientOption($res, $option, $value) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        printf("Option '%s' set to %s\n", client_option_to_string($option), var_export($value, true));
        $ret = parent::setClientOption($res, $option, $value);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
?>

```

The above example will output:

```

proxy::setClientOption(array (
    0 => NULL,
    1 => 210,
    2 => 3221225472,
))
Option 'MYSQLND_UH_MYSQLND_OPT_MAX_ALLOWED_PACKET' set to 3221225472
proxy::setClientOption returns true
proxy::setClientOption(array (
    0 => NULL,
    1 => 211,
    2 => 'mysql_native_password',
))
Option 'MYSQLND_UH_MYSQLND_OPT_AUTH_PROTOCOL' set to 'mysql_native_password'
proxy::setClientOption returns true
proxy::setClientOption(array (
    0 => NULL,
    1 => 8,
    2 => 1,
))
Option 'MYSQLND_UH_MYSQLND_OPT_LOCAL_INFILE' set to 1
proxy::setClientOption returns true

```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_real\\_connect](#)  
[mysqli\\_options](#)

### 21.9.8.7.42. MysqlndUhConnection::setServerOption

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::setServerOption](#)

Sets a server option

#### Description



```
public void MysqlndUhConnection::setServerOption(
    mysqlnd_connection connection,
    int option);
```

Sets a server option.

### Parameters

*connection* Mysqlnd connection handle. Do not modify!

*option* The option to be set.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.377. `MysqlndUhConnection::setServerOption` example

```
<?php
function server_option_to_string($option) {
    $ret = 'unknown';
    switch ($option) {
        case MYSQLND_UH_SERVER_OPTION_MULTI_STATEMENTS_ON:
            $ret = 'MYSQLND_UH_SERVER_OPTION_MULTI_STATEMENTS_ON';
            break;
        case MYSQLND_UH_SERVER_OPTION_MULTI_STATEMENTS_OFF:
            $ret = 'MYSQLND_UH_SERVER_OPTION_MULTI_STATEMENTS_ON';
            break;
    }
    return $ret;
}

class proxy extends MysqlndUhConnection {
    public function setServerOption($res, $option) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        printf("Option '%s' set\n", server_option_to_string($option));
        $ret = parent::setServerOption($res, $option);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}

mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->multi_query("SELECT 1; SELECT 2");
?>
```

The above example will output:

```
proxy::setServerOption(array (
    0 => NULL,
    1 => 0,
))
Option 'MYSQLND_UH_SERVER_OPTION_MULTI_STATEMENTS_ON' set
proxy::setServerOption returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_real\\_connect](#)  
[mysqli\\_options](#)

`mysqli_multi_query`

### 21.9.8.7.43. `MySQLndUhConnection::shutdownServer`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MySQLndUhConnection::shutdownServer`

The `shutdownServer` purpose

#### Description

```
public void MySQLndUhConnection::shutdownServer(
    string MYSQLND_UH_RES_MYSQLND_NAME,
    string "level");
```



#### Warning

This function is currently not documented; only its argument list is available.

#### Parameters

`MYSQLND_UH_RES_MYSQLND_NAME`

`"level"`

#### Return Values

### 21.9.8.7.44. `MySQLndUhConnection::simpleCommand`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MySQLndUhConnection::simpleCommand`

Sends a basic `COM_*` command

#### Description

```
public bool MySQLndUhConnection::simpleCommand(
    mysqlnd_connection connection,
    int command,
    string arg,
    int ok_packet,
    bool silent,
    bool ignore_upsert_status);
```

Sends a basic `COM_*` command to MySQL.

#### Parameters

<code>connection</code>	MySQLnd connection handle. Do not modify!
<code>command</code>	The COM command to be send.
<code>arg</code>	Optional COM command arguments.
<code>ok_packet</code>	The OK packet type.
<code>silent</code>	Whether mysqlnd may emit errors.
<code>ignore_upsert_status</code>	Whether to ignore <code>UPDATE/INSERT</code> status.

#### Return Values

Returns **TRUE** on success. Otherwise, returns **FALSE**

## Examples

### Example 21.378. MysqLndUhConnection::simpleCommand example

```
<?php
function server_cmd_2_string($command) {
    $mapping = array(
        MYSQLND_UH_MYSQLND_COM_SLEEP => "MYSQLND_UH_MYSQLND_COM_SLEEP",
        MYSQLND_UH_MYSQLND_COM_QUIT => "MYSQLND_UH_MYSQLND_COM_QUIT",
        MYSQLND_UH_MYSQLND_COM_INIT_DB => "MYSQLND_UH_MYSQLND_COM_INIT_DB",
        MYSQLND_UH_MYSQLND_COM_QUERY => "MYSQLND_UH_MYSQLND_COM_QUERY",
        MYSQLND_UH_MYSQLND_COM_FIELD_LIST => "MYSQLND_UH_MYSQLND_COM_FIELD_LIST",
        MYSQLND_UH_MYSQLND_COM_CREATE_DB => "MYSQLND_UH_MYSQLND_COM_CREATE_DB",
        MYSQLND_UH_MYSQLND_COM_DROP_DB => "MYSQLND_UH_MYSQLND_COM_DROP_DB",
        MYSQLND_UH_MYSQLND_COM_REFRESH => "MYSQLND_UH_MYSQLND_COM_REFRESH",
        MYSQLND_UH_MYSQLND_COM_SHUTDOWN => "MYSQLND_UH_MYSQLND_COM_SHUTDOWN",
        MYSQLND_UH_MYSQLND_COM_STATISTICS => "MYSQLND_UH_MYSQLND_COM_STATISTICS",
        MYSQLND_UH_MYSQLND_COM_PROCESS_INFO => "MYSQLND_UH_MYSQLND_COM_PROCESS_INFO",
        MYSQLND_UH_MYSQLND_COM_CONNECT => "MYSQLND_UH_MYSQLND_COM_CONNECT",
        MYSQLND_UH_MYSQLND_COM_PROCESS_KILL => "MYSQLND_UH_MYSQLND_COM_PROCESS_KILL",
        MYSQLND_UH_MYSQLND_COM_DEBUG => "MYSQLND_UH_MYSQLND_COM_DEBUG",
        MYSQLND_UH_MYSQLND_COM_PING => "MYSQLND_UH_MYSQLND_COM_PING",
        MYSQLND_UH_MYSQLND_COM_TIME => "MYSQLND_UH_MYSQLND_COM_TIME",
        MYSQLND_UH_MYSQLND_COM_DELAYED_INSERT => "MYSQLND_UH_MYSQLND_COM_DELAYED_INSERT",
        MYSQLND_UH_MYSQLND_COM_CHANGE_USER => "MYSQLND_UH_MYSQLND_COM_CHANGE_USER",
        MYSQLND_UH_MYSQLND_COM_BINLOG_DUMP => "MYSQLND_UH_MYSQLND_COM_BINLOG_DUMP",
        MYSQLND_UH_MYSQLND_COM_TABLE_DUMP => "MYSQLND_UH_MYSQLND_COM_TABLE_DUMP",
        MYSQLND_UH_MYSQLND_COM_CONNECT_OUT => "MYSQLND_UH_MYSQLND_COM_CONNECT_OUT",
        MYSQLND_UH_MYSQLND_COM_REGISTER_SLAVE => "MYSQLND_UH_MYSQLND_COM_REGISTER_SLAVE",
        MYSQLND_UH_MYSQLND_COM_STMT_PREPARE => "MYSQLND_UH_MYSQLND_COM_STMT_PREPARE",
        MYSQLND_UH_MYSQLND_COM_STMT_EXECUTE => "MYSQLND_UH_MYSQLND_COM_STMT_EXECUTE",
        MYSQLND_UH_MYSQLND_COM_STMT_SEND_LONG_DATA => "MYSQLND_UH_MYSQLND_COM_STMT_SEND_LONG_DATA",
        MYSQLND_UH_MYSQLND_COM_STMT_CLOSE => "MYSQLND_UH_MYSQLND_COM_STMT_CLOSE",
        MYSQLND_UH_MYSQLND_COM_STMT_RESET => "MYSQLND_UH_MYSQLND_COM_STMT_RESET",
        MYSQLND_UH_MYSQLND_COM_SET_OPTION => "MYSQLND_UH_MYSQLND_COM_SET_OPTION",
        MYSQLND_UH_MYSQLND_COM_STMT_FETCH => "MYSQLND_UH_MYSQLND_COM_STMT_FETCH",
        MYSQLND_UH_MYSQLND_COM_DAEMON => "MYSQLND_UH_MYSQLND_COM_DAEMON",
        MYSQLND_UH_MYSQLND_COM_END => "MYSQLND_UH_MYSQLND_COM_END",
    );
    return (isset($mapping[$command])) ? $mapping[$command] : 'unknown';
}

function ok_packet_2_string($ok_packet) {
    $mapping = array(
        MYSQLND_UH_MYSQLND_PROT_GREET_PACKET => "MYSQLND_UH_MYSQLND_PROT_GREET_PACKET",
        MYSQLND_UH_MYSQLND_PROT_AUTH_PACKET => "MYSQLND_UH_MYSQLND_PROT_AUTH_PACKET",
        MYSQLND_UH_MYSQLND_PROT_OK_PACKET => "MYSQLND_UH_MYSQLND_PROT_OK_PACKET",
        MYSQLND_UH_MYSQLND_PROT_EOF_PACKET => "MYSQLND_UH_MYSQLND_PROT_EOF_PACKET",
        MYSQLND_UH_MYSQLND_PROT_CMD_PACKET => "MYSQLND_UH_MYSQLND_PROT_CMD_PACKET",
        MYSQLND_UH_MYSQLND_PROT_RSET_HEADER_PACKET => "MYSQLND_UH_MYSQLND_PROT_RSET_HEADER_PACKET",
        MYSQLND_UH_MYSQLND_PROT_RSET_FLD_PACKET => "MYSQLND_UH_MYSQLND_PROT_RSET_FLD_PACKET",
        MYSQLND_UH_MYSQLND_PROT_ROW_PACKET => "MYSQLND_UH_MYSQLND_PROT_ROW_PACKET",
        MYSQLND_UH_MYSQLND_PROT_STATS_PACKET => "MYSQLND_UH_MYSQLND_PROT_STATS_PACKET",
        MYSQLND_UH_MYSQLND_PREPARE_RESP_PACKET => "MYSQLND_UH_MYSQLND_PREPARE_RESP_PACKET",
        MYSQLND_UH_MYSQLND_CHG_USER_RESP_PACKET => "MYSQLND_UH_MYSQLND_CHG_USER_RESP_PACKET",
        MYSQLND_UH_MYSQLND_PROT_LAST => "MYSQLND_UH_MYSQLND_PROT_LAST",
    );
    return (isset($mapping[$ok_packet])) ? $mapping[$ok_packet] : 'unknown';
}

class proxy extends MysqLndUhConnection {
    public function simpleCommand($conn, $command, $arg, $ok_packet, $silent, $ignore_upsert_status) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        printf("Command '%s'\n", server_cmd_2_string($command));
        printf("OK packet '%s'\n", ok_packet_2_string($ok_packet));
        $ret = parent::simpleCommand($conn, $command, $arg, $ok_packet, $silent, $ignore_upsert_status);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
```

```

}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("SELECT 1");
?>

```

The above example will output:

```

proxy::simpleCommand(array (
  0 => NULL,
  1 => 3,
  2 => 'SELECT 1',
  3 => 13,
  4 => false,
  5 => false,
))
Command 'MYSQLND_UH_MYSQLND_COM_QUERY'
OK packet 'MYSQLND_UH_MYSQLND_PROT_LAST'
proxy::simpleCommand returns true
:)proxy::simpleCommand(array (
  0 => NULL,
  1 => 1,
  2 => '',
  3 => 13,
  4 => true,
  5 => true,
))
Command 'MYSQLND_UH_MYSQLND_COM_QUIT'
OK packet 'MYSQLND_UH_MYSQLND_PROT_LAST'
proxy::simpleCommand returns true

```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

#### 21.9.8.7.45. [MysqlndUhConnection::simpleCommandHandleResponse](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::simpleCommandHandleResponse](#)

Process a response for a basic COM\_\* command send to the client

### Description

```

public bool MysqlndUhConnection::simpleCommandHandleResponse(
    mysqlnd_connection connection,
    int ok_packet,
    bool silent,
    int command,
    bool ignore_upsert_status);

```

Process a response for a basic COM\_\* command send to the client.

### Parameters

<i>connection</i>	Mysqlnd connection handle. Do not modify!
<i>ok_packet</i>	The OK packet type.
<i>silent</i>	Whether mysqlnd may emit errors.
<i>command</i>	The COM command to process results from.

*ignore\_upsert\_status* Whether to ignore UPDATE/INSERT status.

## Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

## Examples

### Example 21.379. `MysqldUhConnection::simpleCommandHandleResponse` example

```
<?php
function server_cmd_2_string($command) {
    $mapping = array(
        MYSQLND_UH_MYSQLND_COM_SLEEP => "MYSQLND_UH_MYSQLND_COM_SLEEP",
        MYSQLND_UH_MYSQLND_COM_QUIT => "MYSQLND_UH_MYSQLND_COM_QUIT",
        MYSQLND_UH_MYSQLND_COM_INIT_DB => "MYSQLND_UH_MYSQLND_COM_INIT_DB",
        MYSQLND_UH_MYSQLND_COM_QUERY => "MYSQLND_UH_MYSQLND_COM_QUERY",
        MYSQLND_UH_MYSQLND_COM_FIELD_LIST => "MYSQLND_UH_MYSQLND_COM_FIELD_LIST",
        MYSQLND_UH_MYSQLND_COM_CREATE_DB => "MYSQLND_UH_MYSQLND_COM_CREATE_DB",
        MYSQLND_UH_MYSQLND_COM_DROP_DB => "MYSQLND_UH_MYSQLND_COM_DROP_DB",
        MYSQLND_UH_MYSQLND_COM_REFRESH => "MYSQLND_UH_MYSQLND_COM_REFRESH",
        MYSQLND_UH_MYSQLND_COM_SHUTDOWN => "MYSQLND_UH_MYSQLND_COM_SHUTDOWN",
        MYSQLND_UH_MYSQLND_COM_STATISTICS => "MYSQLND_UH_MYSQLND_COM_STATISTICS",
        MYSQLND_UH_MYSQLND_COM_PROCESS_INFO => "MYSQLND_UH_MYSQLND_COM_PROCESS_INFO",
        MYSQLND_UH_MYSQLND_COM_CONNECT => "MYSQLND_UH_MYSQLND_COM_CONNECT",
        MYSQLND_UH_MYSQLND_COM_PROCESS_KILL => "MYSQLND_UH_MYSQLND_COM_PROCESS_KILL",
        MYSQLND_UH_MYSQLND_COM_DEBUG => "MYSQLND_UH_MYSQLND_COM_DEBUG",
        MYSQLND_UH_MYSQLND_COM_PING => "MYSQLND_UH_MYSQLND_COM_PING",
        MYSQLND_UH_MYSQLND_COM_TIME => "MYSQLND_UH_MYSQLND_COM_TIME",
        MYSQLND_UH_MYSQLND_COM_DELAYED_INSERT => "MYSQLND_UH_MYSQLND_COM_DELAYED_INSERT",
        MYSQLND_UH_MYSQLND_COM_CHANGE_USER => "MYSQLND_UH_MYSQLND_COM_CHANGE_USER",
        MYSQLND_UH_MYSQLND_COM_BINLOG_DUMP => "MYSQLND_UH_MYSQLND_COM_BINLOG_DUMP",
        MYSQLND_UH_MYSQLND_COM_TABLE_DUMP => "MYSQLND_UH_MYSQLND_COM_TABLE_DUMP",
        MYSQLND_UH_MYSQLND_COM_CONNECT_OUT => "MYSQLND_UH_MYSQLND_COM_CONNECT_OUT",
        MYSQLND_UH_MYSQLND_COM_REGISTER_SLAVED => "MYSQLND_UH_MYSQLND_COM_REGISTER_SLAVED",
        MYSQLND_UH_MYSQLND_COM_STMT_PREPARE => "MYSQLND_UH_MYSQLND_COM_STMT_PREPARE",
        MYSQLND_UH_MYSQLND_COM_STMT_EXECUTE => "MYSQLND_UH_MYSQLND_COM_STMT_EXECUTE",
        MYSQLND_UH_MYSQLND_COM_STMT_SEND_LONG_DATA => "MYSQLND_UH_MYSQLND_COM_STMT_SEND_LONG_DATA",
        MYSQLND_UH_MYSQLND_COM_STMT_CLOSE => "MYSQLND_UH_MYSQLND_COM_STMT_CLOSE",
        MYSQLND_UH_MYSQLND_COM_STMT_RESET => "MYSQLND_UH_MYSQLND_COM_STMT_RESET",
        MYSQLND_UH_MYSQLND_COM_SET_OPTION => "MYSQLND_UH_MYSQLND_COM_SET_OPTION",
        MYSQLND_UH_MYSQLND_COM_STMT_FETCH => "MYSQLND_UH_MYSQLND_COM_STMT_FETCH",
        MYSQLND_UH_MYSQLND_COM_DAEMON => "MYSQLND_UH_MYSQLND_COM_DAEMON",
        MYSQLND_UH_MYSQLND_COM_END => "MYSQLND_UH_MYSQLND_COM_END",
    );
    return (isset($mapping[$command])) ? $mapping[$command] : 'unknown';
}

function ok_packet_2_string($ok_packet) {
    $mapping = array(
        MYSQLND_UH_MYSQLND_PROT_GREET_PACKET => "MYSQLND_UH_MYSQLND_PROT_GREET_PACKET",
        MYSQLND_UH_MYSQLND_PROT_AUTH_PACKET => "MYSQLND_UH_MYSQLND_PROT_AUTH_PACKET",
        MYSQLND_UH_MYSQLND_PROT_OK_PACKET => "MYSQLND_UH_MYSQLND_PROT_OK_PACKET",
        MYSQLND_UH_MYSQLND_PROT_EOF_PACKET => "MYSQLND_UH_MYSQLND_PROT_EOF_PACKET",
        MYSQLND_UH_MYSQLND_PROT_CMD_PACKET => "MYSQLND_UH_MYSQLND_PROT_CMD_PACKET",
        MYSQLND_UH_MYSQLND_PROT_RSET_HEADER_PACKET => "MYSQLND_UH_MYSQLND_PROT_RSET_HEADER_PACKET",
        MYSQLND_UH_MYSQLND_PROT_RSET_FLD_PACKET => "MYSQLND_UH_MYSQLND_PROT_RSET_FLD_PACKET",
        MYSQLND_UH_MYSQLND_PROT_ROW_PACKET => "MYSQLND_UH_MYSQLND_PROT_ROW_PACKET",
        MYSQLND_UH_MYSQLND_PROT_STATS_PACKET => "MYSQLND_UH_MYSQLND_PROT_STATS_PACKET",
        MYSQLND_UH_MYSQLND_PREPARE_RESP_PACKET => "MYSQLND_UH_MYSQLND_PREPARE_RESP_PACKET",
        MYSQLND_UH_MYSQLND_CHG_USER_RESP_PACKET => "MYSQLND_UH_MYSQLND_CHG_USER_RESP_PACKET",
        MYSQLND_UH_MYSQLND_PROT_LAST => "MYSQLND_UH_MYSQLND_PROT_LAST",
    );
    return (isset($mapping[$ok_packet])) ? $mapping[$ok_packet] : 'unknown';
}

class proxy extends MysqldUhConnection {
    public function simpleCommandHandleResponse($conn, $ok_packet, $silent, $command, $ignore_upsert_status) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        printf("Command '%s'\n", server_cmd_2_string($command));
    }
}
```

```
printf("OK packet '%s'\n", ok_packet_2_string($ok_packet));
$ret = parent::simpleCommandHandleResponse($conn, $ok_packet, $silent, $command, $ignore_upsert_status);
printf("%s returns %s\n", __METHOD__, var_export($ret, true));
return $ret;
}
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysql = mysql_connect("localhost", "root", "");
mysql_query("SELECT 1 FROM DUAL", $mysql);
?>
```

The above example will output:

```
proxy::simpleCommandHandleResponse(array (
  0 => NULL,
  1 => 5,
  2 => false,
  3 => 27,
  4 => true,
))
Command 'MYSQLND_UH_MYSQLND_COM_SET_OPTION'
OK packet 'MYSQLND_UH_MYSQLND_PROT_EOF_PACKET'
proxy::simpleCommandHandleResponse returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

### 21.9.8.7.46. `MysqlndUhConnection::sslSet`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::sslSet`

Used for establishing secure connections using SSL

### Description

```
public bool MysqlndUhConnection::sslSet(
    mysqlnd_connection connection,
    string key,
    string cert,
    string ca,
    string capath,
    string cipher);
```

Used for establishing secure connections using SSL.

### Parameters

<i>connection</i>	Mysqlnd connection handle. Do not modify!
<i>key</i>	The path name to the key file.
<i>cert</i>	The path name to the certificate file.
<i>ca</i>	The path name to the certificate authority file.
<i>capath</i>	The pathname to a directory that contains trusted SSL CA certificates in PEM format.
<i>cipher</i>	A list of allowable ciphers to use for SSL encryption.

### Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.380. [MysqlndUhConnection::sslSet](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function sslSet($conn, $key, $cert, $ca, $capath, $cipher) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::sslSet($conn, $key, $cert, $ca, $capath, $cipher);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->ssl_set("key", "cert", "ca", "capath", "cipher");
?>
```

The above example will output:

```
proxy::sslSet(array (
  0 => NULL,
  1 => 'key',
  2 => 'cert',
  3 => 'ca',
  4 => 'capath',
  5 => 'cipher',
))
proxy::sslSet returns true
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_ssl\\_set](#)

### 21.9.8.7.47. [MysqlndUhConnection::stmtInit](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::stmtInit](#)

Initializes a statement and returns a resource for use with `mysqli_statement::prepare`

## Description

```
public resource MysqlndUhConnection::stmtInit(
    mysqlnd_connection connection);
```

Initializes a statement and returns a resource for use with `mysqli_statement::prepare`.

## Parameters

*connection* Mysqlnd connection handle. Do not modify!

## Return Values

Resource of type [Mysqlnd Prepared Statement \(internal only - you must not modify it!\)](#). The documentation may also refer to such resources using the alias name [mysqlnd\\_prepared\\_statement](#).

## Examples

### Example 21.381. `MysqlndUhConnection::stmtInit` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function stmtInit($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        var_dump($res);
        $ret = parent::stmtInit($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        var_dump($ret);
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());
$mysqli = new mysqli("localhost", "root", "", "test");
$stmt = $mysqli->prepare("SELECT 1 AS _one FROM DUAL");
$stmt->execute();
$one = NULL;
$stmt->bind_result($one);
$stmt->fetch();
var_dump($one);
?>
```

The above example will output:

```
proxy::stmtInit(array (
    0 => NULL,
))
resource(19) of type (Mysqlnd Connection)
proxy::stmtInit returns NULL
resource(246) of type (Mysqlnd Prepared Statement (internal only - you must not modify it!))
int(1)
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_stmt\\_init](#)

### 21.9.8.7.48. `MysqlndUhConnection::storeResult`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::storeResult`

Transfers a result set from the last query

## Description

```
public resource MysqlndUhConnection::storeResult(
    mysqlnd_connection connection);
```

Transfers a result set from the last query.

## Parameters

*connection* Mysqlnd connection handle. Do not modify!

## Return Values

Resource of type `Mysqlnd Resultset (internal only - you must not modify it!)`. The documentation may also refer to such resources using the alias name `mysqlnd_resultset`.



## Examples

### Example 21.382. [MysqlndUhConnection::storeResult](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function storeResult($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::storeResult($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        var_dump($ret);
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$res = $mysqli->query("SELECT 'Also called buffered result' AS _msg FROM DUAL");
var_dump($res->fetch_assoc());

$mysqli->real_query("SELECT 'Good morning!' AS _msg FROM DUAL");
$res = $mysqli->store_result();
var_dump($res->fetch_assoc());
?>
```

The above example will output:

```
proxy::storeResult(array (
    0 => NULL,
))
proxy::storeResult returns NULL
resource(475) of type (Mysqlnd Resultset (internal only - you must not modify it!))
array(1) {
    ["_msg"]=>
        string(27) "Also called buffered result"
}
proxy::storeResult(array (
    0 => NULL,
))
proxy::storeResult returns NULL
resource(730) of type (Mysqlnd Resultset (internal only - you must not modify it!))
array(1) {
    ["_msg"]=>
        string(13) "Good morning!"
}
```

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_store\\_result](#)  
[mysqli\\_real\\_query](#)

### 21.9.8.7.49. [MysqlndUhConnection::txCommit](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::txCommit](#)

Commits the current transaction

### Description

```
public bool MysqlndUhConnection::txCommit(
```

```
mysqlnd_connection connection);
```

Commits the current transaction.

### Parameters

`connection` Mysqlnd connection handle. Do not modify!

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.383. `MysqlndUhConnection::txCommit` example

```
<?php
class proxy extends MysqlndUhConnection {
    public function txCommit($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::txCommit($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->commit();
?>
```

The above example will output:

```
proxy::txCommit(array (
    0 => NULL,
))
proxy::txCommit returns true
```

### See Also

`mysqlnd_uh_set_connection_proxy`  
`mysqli_commit`

#### 21.9.8.7.50. `MysqlndUhConnection::txRollback`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhConnection::txRollback`

Rolls back current transaction

### Description

```
public bool MysqlndUhConnection::txRollback(
    mysqlnd_connection connection);
```

Rolls back current transaction.

### Parameters

`connection` Mysqlnd connection handle. Do not modify!

## Return Values

Returns [TRUE](#) on success. Otherwise, returns [FALSE](#)

## Examples

### Example 21.384. [MysqlndUhConnection::txRollback](#) example

```
<?php
class proxy extends MysqlndUhConnection {
    public function txRollback($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::txRollback($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->rollback();
?>
```

The above example will output:

```
proxy::txRollback(array (
    0 => NULL,
))
proxy::txRollback returns true
```

## See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_commit](#)

### 21.9.8.7.51. [MysqlndUhConnection::useResult](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [MysqlndUhConnection::useResult](#)

Initiate a result set retrieval

## Description

```
public resource MysqlndUhConnection::useResult(
    mysqlnd_connection connection);
```

Initiate a result set retrieval.

## Parameters

*connection* Mysqlnd connection handle. Do not modify!

## Return Values

Resource of type [Mysqlnd Resultset](#) (*internal only - you must not modify it!*). The documentation may also refer to such resources using the alias name [mysqlnd\\_resultset](#).

## Examples

**Example 21.385. [MysqlndUhConnection::useResult](#) example**

```
<?php
class proxy extends MysqlndUhConnection {
    public function useResult($res) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::useResult($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        var_dump($ret);
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->real_query("SELECT 'Good morning!' AS _msg FROM DUAL");
$res = $mysqli->use_result();
var_dump($res->fetch_assoc());
?>
```

The above example will output:

```
proxy::useResult(array (
    0 => NULL,
))
proxy::useResult returns NULL
resource(425) of type (Mysqlnd Resultset (internal only - you must not modify it!))
array(1) {
    ["_msg"]=>
        string(13) "Good morning!"
}
```

**See Also**

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqli\\_use\\_result](#)  
[mysqli\\_real\\_query](#)

**21.9.8.8. The [MysqlndUhPreparedStatement](#) class ([MysqlndUhPreparedStatement](#))**

Copyright 1997-2012 the PHP Documentation Group. [2906]

```
MysqlndUhPreparedStatement {
MysqlndUhPreparedStatement

    Methods

    public MysqlndUhPreparedStatement::__construct();

    public bool MysqlndUhPreparedStatement::execute(
        mysqlnd_prepared_statement statement);

    public bool MysqlndUhPreparedStatement::prepare(
        mysqlnd_prepared_statement statement,
        string query);
}
```

**21.9.8.8.1. [MysqlndUhPreparedStatement::\\_\\_construct](#)**

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhpStatement::__construct`

The `__construct` purpose

### Description

```
public MysqlndUhpStatement::__construct();
```



### Warning

This function is currently not documented; only its argument list is available.

### Parameters

This function has no parameters.

### Return Values

## 21.9.8.8.2. `MysqlndUhpStatement::execute`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhpStatement::execute`

Executes a prepared Query

### Description

```
public bool MysqlndUhpStatement::execute(
    mysqlnd_prepared_statement statement);
```

Executes a prepared Query.

### Parameters

*statement*

Mysqlnd prepared statement handle. Do not modify! Resource of type `Mysqlnd Prepared Statement (internal only - you must not modify it!)`.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.386. `MysqlndUhpStatement::execute` example

```
<?php
class stmt_proxy extends MysqlndUhpStatement {
    public function execute($res) {
        printf("%s(", __METHOD__);
        var_dump($res);
        printf(")\n");
        $ret = parent::execute($res);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        var_dump($ret);
        return $ret;
    }
}
mysqlnd_uh_set_statement_proxy(new stmt_proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
```

```
$stmt = $mysqli->prepare("SELECT 'Labskaus' AS _msg FROM DUAL");
$stmt->execute();
$msg = NULL;
$stmt->bind_result($msg);
$stmt->fetch();
var_dump($msg);
?>
```

The above example will output:

```
stmt_proxy::execute(resource(256) of type (Mysqlnd Prepared Statement (internal only - you must not modify
)
)
stmt_proxy::execute returns true
bool(true)
string(8) "Labskaus"
```

### See Also

[mysqlnd\\_uh\\_set\\_statement\\_proxy](#)  
[mysqli\\_stmt\\_execute](#)

### 21.9.8.8.3. `MysqlndUhPreparedStatement::prepare`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `MysqlndUhPreparedStatement::prepare`

Prepare an SQL statement for execution

#### Description

```
public bool MysqlndUhPreparedStatement::prepare(
    mysqlnd_prepared_statement statement,
    string query);
```

Prepare an SQL statement for execution.

#### Parameters

<i>statement</i>	Mysqlnd prepared statement handle. Do not modify! Resource of type <code>Mysqlnd Prepared Statement (internal only - you must not modify it!)</code> .
<i>query</i>	The query to be prepared.

#### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

#### Examples

##### Example 21.387. `MysqlndUhPreparedStatement::prepare` example

```
<?php
class stmt_proxy extends MysqlndUhPreparedStatement {
    public function prepare($res, $query) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $query = "SELECT 'No more you-know-what-I-mean for lunch, please' AS _msg FROM DUAL";
        $ret = parent::prepare($res, $query);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        var_dump($ret);
    }
}
```

```

    return $ret;
}
}
mysqlnd_uh_set_statement_proxy(new stmt_proxy());

$mysqli = new mysqli("localhost", "root", "", "test");
$stmt = $mysqli->prepare("SELECT 'Labskaus' AS _msg FROM DUAL");
$stmt->execute();
$msg = NULL;
$stmt->bind_result($msg);
$stmt->fetch();
var_dump($msg);
?>

```

The above example will output:

```

stmt_proxy::prepare(array (
  0 => NULL,
  1 => 'SELECT \'Labskaus\' AS _msg FROM DUAL',
))
stmt_proxy::prepare returns true
bool(true)
string(46) "No more you-know-what-I-mean for lunch, please"

```

### See Also

[mysqlnd\\_uh\\_set\\_statement\\_proxy](#)  
[mysqli\\_stmt\\_prepare](#)  
[mysqli\\_prepare](#)

## 21.9.8.9. Mysqlnd\_uh Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]

### 21.9.8.9.1. [mysqlnd\\_uh\\_convert\\_to\\_mysqlnd](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_uh\\_convert\\_to\\_mysqlnd](#)

Converts a MySQL connection handle into a mysqlnd connection handle

#### Description

```

resource mysqlnd_uh_convert_to_mysqlnd(
    mysqli mysql_connection);

```

Converts a MySQL connection handle into a mysqlnd connection handle. After conversion you can execute mysqlnd library calls on the connection handle. This can be used to access mysqlnd functionality not made available through user space API calls.

The function can be disabled with [mysqlnd\\_uh.enable](#). If [mysqlnd\\_uh.enable](#) is set to [FALSE](#) the function will not install the proxy and always return [TRUE](#). Additionally, an error of the type [E\\_WARNING](#) may be emitted. The error message may read like [PHP Warning: mysqlnd\\_uh\\_convert\\_to\\_mysqlnd\(\): \(Mysqlnd User Handler\) The plugin has been disabled by setting the configuration parameter mysqlnd\\_uh.enable = false. You are not allowed to call this function \[...\]](#).

#### Parameters

*MySQL connection handle* A MySQL connection handle of type [mysql](#), [mysqli](#) or [PDO\\_MySQL](#).

## Return Values

A mysqlnd connection handle.

## Changelog

Version	Description
5.4.0	The <code>mysql_connection</code> parameter can now be of type <code>mysql</code> , <code>PDO_MySQL</code> , or <code>mysqli</code> . Before, only the <code>mysqli</code> type was allowed.

## Examples

### Example 21.388. [mysqlnd\\_uh\\_convert\\_to\\_mysqlnd](#) example

```
<?php
/* PDO user API gives no access to connection thread id */
$mysql_connection = new PDO("mysql:host=localhost;dbname=test", "root", "");

/* Convert PDO MySQL handle to mysqlnd handle */
$mysqlnd = mysqlnd_uh_convert_to_mysqlnd($mysql_connection);

/* Create Proxy to call mysqlnd connection class methods */
$obj = new MySQLndUHConnection();
/* Call mysqlnd_conn::get_thread_id */
var_dump($obj->getThreadId($mysqlnd));

/* Use SQL to fetch connection thread id */
var_dump($mysql_connection->query("SELECT CONNECTION_ID()")->fetchAll());
?>
```

The above example will output:

```
int(27054)
array(1) {
  [0]=>
  array(2) {
    ["CONNECTION_ID()"]=>
    string(5) "27054"
    [0]=>
    string(5) "27054"
  }
}
```

## See Also

[mysqlnd\\_uh.enable](#)

### 21.9.8.9.2. [mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_uh\\_set\\_connection\\_proxy](#)

Installs a proxy for mysqlnd connections

## Description

```
bool mysqlnd_uh_set_connection_proxy(
    MySQLndUHConnection connection_proxy,
    mysqli mysqli_connection);
```



Installs a proxy object to hook mysqlnd's connection objects methods. Once installed, the proxy will be used for all MySQL connections opened with `mysqli`, `mysql` or `PDO_MYSQL`, assuming that the listed extensions are compiled to use the `mysqlnd` library.

The function can be disabled with `mysqlnd_uh.enable`. If `mysqlnd_uh.enable` is set to `FALSE` the function will not install the proxy and always return `TRUE`. Additionally, an error of the type `E_WARNING` may be emitted. The error message may read like `PHP Warning: mysqlnd_uh_set_connection_proxy(): (Mysqlnd User Handler) The plugin has been disabled by setting the configuration parameter mysqlnd_uh.enable = false. The proxy has not been installed [...]`.

### Parameters

<code>connection_proxy</code>	A proxy object of type <code>MysqlndUhConnection</code> .
<code>mysqli_connection</code>	Object of type <code>mysqli</code> . If given, the proxy will be set for this particular connection only.

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### Examples

#### Example 21.389. `mysqlnd_uh_set_connection_proxy` example

```
<?php
$mysqli = new mysqli("localhost", "root", "", "test");
$mysqli->query("SELECT 'No proxy installed, yet'");

class proxy extends MysqlndUhConnection {
    public function query($res, $query) {
        printf("%s(%s)\n", __METHOD__, var_export(func_get_args(), true));
        $ret = parent::query($res, $query);
        printf("%s returns %s\n", __METHOD__, var_export($ret, true));
        return $ret;
    }
}
mysqlnd_uh_set_connection_proxy(new proxy());

$mysqli->query("SELECT 'mysqlnd rocks!'");

$mysql = mysql_connect("localhost", "root", "", "test");
mysql_query("SELECT 'Ahoy Andrey!'", $mysql);

$pdo = new PDO("mysql:host=localhost;dbname=test", "root", "");
$pdo->query("SELECT 'Moin Johannes!'");
?>
```

The above example will output:

```
proxy::query(array (
  0 => NULL,
  1 => 'SELECT \'mysqlnd rocks!\'',
))
proxy::query returns true
proxy::query(array (
  0 => NULL,
  1 => 'SELECT \'Ahoy Andrey!\'',
))
proxy::query returns true
proxy::query(array (
  0 => NULL,
  1 => 'SELECT \'Moin Johannes!\'',
```

```
))  
proxy::query returns true
```

### See Also

[mysqlnd\\_uh\\_set\\_statement\\_proxy](#)  
[mysqlnd\\_uh.enable](#)

#### 21.9.8.9.3. `mysqlnd_uh_set_statement_proxy`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_uh\\_set\\_statement\\_proxy](#)

Installs a proxy for mysqlnd statements

### Description

```
bool mysqlnd_uh_set_statement_proxy(  
    MysqlndUhStatement statement_proxy);
```

Installs a proxy for mysqlnd statements. The proxy object will be used for all mysqlnd prepared statement objects, regardless which PHP MySQL extension ([mysqli](#), [mysql](#), [PDO\\_MYSQL](#)) has created them as long as the extension is compiled to use the [mysqlnd](#) library.

The function can be disabled with [mysqlnd\\_uh.enable](#). If [mysqlnd\\_uh.enable](#) is set to `FALSE` the function will not install the proxy and always return `TRUE`. Additionally, an error of the type `E_WARNING` may be emitted. The error message may read like `PHP Warning: mysqlnd_uh_set_statement_proxy(): (Mysqlnd User Handler) The plugin has been disabled by setting the configuration parameter mysqlnd_uh.enable = false. The proxy has not been installed [...]`.

### Parameters

*statement\_proxy*                      The mysqlnd statement proxy object of type `MysqlndUhStatement`

### Return Values

Returns `TRUE` on success. Otherwise, returns `FALSE`

### See Also

[mysqlnd\\_uh\\_set\\_connection\\_proxy](#)  
[mysqlnd\\_uh.enable](#)

#### 21.9.8.10. Change History

Copyright 1997-2012 the PHP Documentation Group. [2906]

The Change History lists major changes users need to be aware of if upgrading from one version to another. It is a high level summary of selected changes that may impact applications or might even break backwards compatibility. See also the [CHANGES](#) file contained in the source for additional changelog information. The commit history is also available.

##### 21.9.8.10.1. PECL/mysqlnd\_uh 1.0 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

1.0.1-alpha

- Release date: TBD
- Motto/theme: bug fix release

#### Feature changes

- Support of PHP 5.4.0 or later.
- BC break: `MysqlndUhConnection::changeUser` requires additional `passwd_len` parameter.
- BC break: `MYSQLND_UH_VERSION_STR` renamed to `MYSQLND_UH_VERSION`.  
`MYSQLND_UH_VERSION` renamed to `MYSQLND_UH_VERSION_ID`.
- BC break: `mysqlnd_uh.enabled` configuration setting renamed to `mysqlnd_uh.enable`.

#### 1.0.0-alpha

- Release date: 08/2010
- Motto/theme: Initial release

## 21.9.9. Mysqlnd connection multiplexing plugin ([mysqlnd\\_mux](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqlnd` multiplexing plugin ([mysqlnd\\_mux](#)) multiplexes MySQL connections established by all PHP MySQL extensions that use the MySQL native driver ([mysqlnd](#)) for PHP.

The MySQL native driver for PHP features an internal C API for plugins, such as the connection multiplexing plugin, which can extend the functionality of `mysqlnd`. See the [mysqlnd](#) for additional details about its benefits over the MySQL Client Library `libmysqlclient`.

`mysqlnd` plugins like [mysqlnd\\_mux](#) operate, for the most part, transparently from a user perspective. The connection multiplexing plugin supports all PHP applications, and all MySQL PHP extensions. It does not change existing APIs. Therefore, it can easily be used with existing PHP applications.



### Note

This is a proof-of-concept. All features are at an early stage. Not all kinds of queries are handled by the plugin yet. Thus, it cannot be used in a drop-in fashion at the moment.

Please, do not use this version in production environments.

### 21.9.9.1. Key Features

Copyright 1997-2012 the PHP Documentation Group. [2906]

The key features of `mysqlnd_mux` are as follows:

- Transparent and therefore easy to use:
  - Supports all of the PHP MySQL extensions.
  - Little to no application changes are required, dependent on the required usage scenario.
- Reduces server load and connection establishment latency:
  - Opens less connections to the MySQL server.
  - Less connections to MySQL mean less work for the MySQL server. In a client-server environment scaling the server is often more difficult than scaling the client. Multiplexing helps with horizontal scale-out (scale-by-client).

- Pooling saves connection time.
- Multiplexed connection: multiple user handles share the same network connection. Once opened, a network connection is cached and shared among multiple user handles. There is a 1:n relationship between internal network connection and user connection handles.
- Persistent connection: a network connection is kept open at the end of the web request, if the PHP deployment model allows. Thus, subsequently web requests can reuse a previously opened connection. Like other resources, network connections are bound to the scope of a process. Thus, they can be reused for all web requests served by a process.

### 21.9.9.2. Limitations

Copyright 1997-2012 the PHP Documentation Group. [2906]

The proof-of-concept does not support unbuffered queries, prepared statements, and asynchronous queries.

The connection pool is using a combination of the transport method and hostname as keys. As a consequence, two connections to the same host using the same transport method (TCP/IP, Unix socket, Windows named pipe) will be linked to the same pooled connection even if username and password differ. Be aware of the possible security implications.

The proof-of-concept is transaction agnostic. It does not about SQL transactions.



#### Note

Applications must be aware of the consequences of connection sharing connections.

### 21.9.9.3. About the name `mysqlnd_mux`

Copyright 1997-2012 the PHP Documentation Group. [2906]

The shortcut `mysqlnd_mux` stands for `mysqlnd connection multiplexing plugin`.

### 21.9.9.4. Concepts

Copyright 1997-2012 the PHP Documentation Group. [2906]

This explains the architecture and related concepts for this plugin. Reading and understanding these concepts is required to successfully use this plugin.

#### 21.9.9.4.1. Architecture

Copyright 1997-2012 the PHP Documentation Group. [2906]

The `mysqlnd` connection multiplexing plugin is implemented as a PHP extension. It is written in C and operates under the hood of PHP. During the startup of the PHP interpreter, in the module initialization phase of the PHP engine, it gets registered as a `mysqlnd` plugin to replace specific `mysqlnd` C methods.

The `mysqlnd` library uses PHP streams to communicate with the MySQL server. PHP streams are accessed by the `mysqlnd` library through its net module. The `mysqlnd` connection multiplexing plugin proxies methods of the `mysqlnd` library net module to control opening and closing of network streams.

Upon opening a user connection to MySQL using the appropriate connection functions of either `mysqli`, `PDO_MYSQL` or `ext/mysqli`, the plugin will search its connection pool for an open network connection. If the pool contains a network connection to the host specified by the connect function using the transport method requested (TCP/IP, Unix domain socket, Windows named pipe), the pooled connection is

linked to the user handle. Otherwise, a new network connection is opened, put into the pool and associated with the user connection handle. This way, multiple user handles can be linked to the same network connection.

#### 21.9.9.4.2. Connection pool

Copyright 1997-2012 the PHP Documentation Group. [2906]

The plugins connection pool is created when PHP initializes its modules (`MINIT`) and free'd when PHP shuts down the modules (`MSHUTDOWN`). This is the same as for persistent MySQL connections.

Depending on the deployment model, the pool is used for the duration of one or multiple web requests. Network connections are bound to the lifespan of an operating system level process. If the PHP process serves multiple web requests as it is the case for Fast-CGI or threaded web server deployments, then the pooled connections can be reused over multiple connections. Because multiplexing means sharing connections, it can even happen with a threaded deployment that two threads or two distinct web requests are linked to one pooled network connections.

A pooled connection is explicitly closed once the last reference to it is released. An implicit close happens when PHP shuts down its modules.

#### 21.9.9.4.3. Sharing connections

Copyright 1997-2012 the PHP Documentation Group. [2906]

The PHP mysqlnd connection multiplexing plugin changes the relationship between a users connection handle and the underlying MySQL connection. Without the plugin, every MySQL connection belongs to exactly one user connection at a time. The multiplexing plugin changes. A MySQL connection is shared among multiple user handles. There no one-to-one relation if using the plugin.

Sharing pooled connections has an impact on the connection state. State changing operations from multiple user handles pointing to one MySQL connection are not isolated from each other. If, for example, a session variable is set through one user connection handle, the session variable becomes visible to all other user handles that reference the same underlying MySQL connection.

This is similar in concept to connection state related phenomena described for the PHP mysqlnd replication and load balancing plugin. Please, check the [PECL/mysqlnd\\_ms documentation](#) for more details on the state of a connection.

The proof-of-concept takes no measures to isolate multiplexed connections from each other.

#### 21.9.9.5. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

##### 21.9.9.5.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

[PHP 5.5.0](#) or newer. Some advanced functionality requires [PHP 5.5.0](#) or newer.

The [mysqlnd\\_mux](#) replication and load balancing plugin supports all PHP applications and all available PHP MySQL extensions ([mysqli](#), [mysql](#), [PDO\\_MYSQL](#)). The PHP MySQL extension must be configured to use [mysqlnd](#) in order to be able to use the [mysqlnd\\_mux](#) plugin for [mysqlnd](#).

##### 21.9.9.5.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

Information for installing this PECL extension may be found in the manual chapter titled [Installation of PECL extensions](#). Additional information such as new releases, downloads, source files, maintainer information, and a CHANGELOG, can be located here: [http://pecl.php.net/package/mysqlnd\\_mux](http://pecl.php.net/package/mysqlnd_mux)

### 21.9.9.5.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.77. Mysqlnd\_mux Configure Options**

Name	Default	Changeable	Changelog
<code>mysqlnd_mux.enable</code>	0	PHP_INI_SYSTEM	

Here's a short explanation of the configuration directives.

`mysqlnd_mux.enable` integer Enables or disables the plugin. If disabled, the extension will not plug into `mysqlnd` to proxy internal `mysqlnd` C API calls.

### 21.9.9.6. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

*Other*

The plugins version number can be obtained using `MYSQLND_MUX_VERSION` or `MYSQLND_MUX_VERSION_ID`. `MYSQLND_MUX_VERSION` is the string representation of the numerical version number `MYSQLND_MUX_VERSION_ID`, which is an integer such as 10000. Developers can calculate the version number as follows.

Version (part)	Example
Major*10000	1*10000 = 10000
Minor*100	0*100 = 0
Patch	0 = 0
<code>MYSQLND_MUX_VERSION_ID</code>	10000

`MYSQLND_MUX_VERSION` (string) Plugin version string, for example, "1.0.0-prototype".

`MYSQLND_MUX_VERSION_ID` (integer) Plugin version number, for example, 10000.

### 21.9.9.7. Change History

Copyright 1997-2012 the PHP Documentation Group. [2906]

This change history is a high level summary of selected changes that may impact applications and/or break backwards compatibility.

See also the `CHANGES` file in the source distribution for a complete list of changes.

#### 21.9.9.7.1. PECL/mysqlnd\_mux 1.0 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

1.0.0-pre-alpha

- Release date: no package released, initial check-in 09/2012

- Motto/theme: Proof of concept

Initial check-in. Essentially a demo of the [mysqlnd](#) plugin API.



**Note**

This is the current development series. All features are at an early stage. Changes may happen at any time without prior notice. Please, do not use this version in production environments.

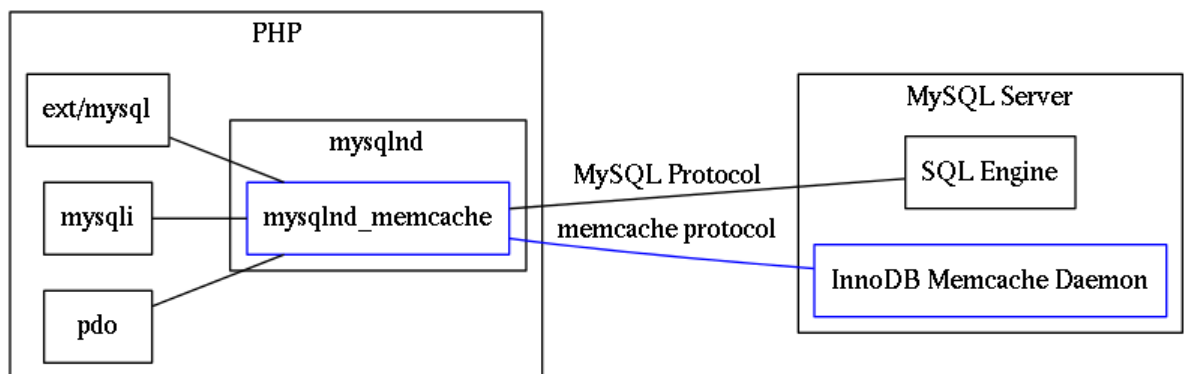
The documentation may not reflect all changes yet.

## 21.9.10. Mysqlnd Memcache plugin ([mysqlnd\\_memcache](#))

Copyright 1997-2012 the PHP Documentation Group. [2906]

The [mysqlnd memcache plugin](#) ([mysqlnd\\_memcache](#)) is an PHP extension for transparently translating SQL into requests for the MySQL InnoDB Memcached Daemon Plugin (server plugin). It includes experimental support for the MySQL Cluster Memcached Daemon. The server plugin provides access to data stored inside MySQL InnoDB (respectively MySQL Cluster NDB) tables using the Memcache protocol. This PHP extension, which supports all PHP MySQL extensions that use [mysqlnd](#), will identify tables exported in this way and will translate specific SELECT queries into Memcache requests.

**Figure 21.92. [mysqlnd\\_memcache](#) data flow**



**Note**

This plugin depends on the MySQL InnoDB Memcached Daemon Plugin. It is not provided to be used with a stand-alone Memcached. For a generic query cache using Memcached look at the [mysqlnd query cache plugin](#). For direct Memcache access look at the [memcache](#) and [memcached](#) extensions.

The MySQL native driver for PHP is a C library that ships together with PHP as of PHP 5.3.0. It serves as a drop-in replacement for the MySQL Client Library (libmysqlclient). Using [mysqlnd](#) has several advantages: no extra downloads are required because it's bundled with PHP, it's under the PHP license, there is lower memory consumption in certain cases, and it contains new functionality such as asynchronous queries.

The [mysqlnd\\_mmemcache](#) operates, for the most part, transparently from a user perspective. The [mysqlnd memcache plugin](#) supports all PHP applications, and all MySQL PHP extensions. It does not change existing APIs. Therefore, it can easily be used with existing PHP applications.

The MySQL Memcache plugins add key-value style access method for data stored in InnoDB resp. NDB (MySQL Cluster) SQL tables through the Memcache protocol. This type of key-value access is often faster than using SQL.

### 21.9.10.1. Key Features

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The key features of PECL/[mysqLnd\\_memcache](#) are as follows.

- Possible performance benefits
  - Client-side: light-weight protocol.
  - Server-side: no SQL parsing, direct access to the storage.
- Please, run your own benchmarks! Actual performance results are highly dependent on setup and hardware used.

### 21.9.10.2. Limitations

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The initial version is not binary safe. Due to the way the MySQL Memcache plugins works there are restrictions related to separators.

Prepared statements and asynchronous queries are not supported. Result set meta data support is limited.

The mapping information for tables accessible via Memcache is not cached in the plugin between requests but fetched from the MySQL server each time a MySQL connection is associated with a Memcache connection. See [mysqLnd\\_memcache\\_set](#) for details.

### 21.9.10.3. On the name

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The shortcut [mysqLnd\\_memcache](#) stands for [mysqLnd memcache plugin](#). Memcache refers to support of the MySQL Memcache plugins for InnoDB and NDB (MySQL Cluster). The plugin is not related to the Memcached cache server.

### 21.9.10.4. Quickstart and Examples

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The [mysqLnd memcache](#) plugin is easy to use. This quickstart will demo typical use-cases, and provide practical advice on getting started.

It is strongly recommended to read the reference sections in addition to the quickstart. The quickstart tries to avoid discussing theoretical concepts and limitations. Instead, it will link to the reference sections. It is safe to begin with the quickstart. However, before using the plugin in mission critical environments we urge you to read additionally the background information from the reference sections.

#### 21.9.10.4.1. Setup

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

The plugin is implemented as a PHP extension. See also the [installation instructions](#) to install this extension.

Compile or configure the PHP MySQL extension (API) ([mysqli](#), [PDO\\_MYSQL](#), [mysql](#)). That extension must use the [mysqLnd](#) library as because [mysqLnd\\_memcache](#) is a plugin for the [mysqLnd](#) library. For additional information, refer to the [mysqLnd\\_memcache installation instructions](#).

Then, load this extension into PHP and activate the plugin in the PHP configuration file using the PHP configuration directive named [mysqLnd\\_memcache.enable](#).



### Example 21.390. Enabling the plugin (php.ini)

```
; On Windows the filename is php_mysqlnd_memcache.dll
; Load the extension
extension=mysqlnd_memcache.so
; Enable it
mysqlnd_memcache.enable=1
```

Follow the instructions given in the [MySQL Reference Manual on installing the Memcache plugins](#) for the MySQL server. Activate the plugins and configure Memcache access for SQL tables.

The examples in this quickguide assume that the following table exists, and that Memcache is configured with access to it.

### Example 21.391. SQL table used for the Quickstart

```
CREATE TABLE test(
  id CHAR(16),
  f1 VARCHAR(255),
  f2 VARCHAR(255),
  f3 VARCHAR(255),
  flags INT NOT NULL,
  cas_column INT,
  expire_time_column INT,
  PRIMARY KEY(id)
) ENGINE=InnoDB;

INSERT INTO test (id, f1, f2, f3) VALUES (1, 'Hello', 'World', '');
INSERT INTO test (id, f1, f2, f3) VALUES (2, 'Lady', 'and', 'the tramp');

INSERT INTO innodb_memcache.containers(
  name, db_schema, db_table, key_columns, value_columns,
  flags, cas_column, expire_time_column, unique_idx_name_on_key)
VALUES (
  'plugin_test', 'test', 'test', 'id', 'f1,f2,f3',
  'flags', 'cas_column', 'expire_time_column', 'PRIMARY KEY');
```

## 21.9.10.4.2. Usage

[Copyright 1997-2012 the PHP Documentation Group. \[2906\]](#)

After associating a MySQL connection with a Memcache connection using [mysqlnd\\_memcache\\_set](#) the plugin attempts to transparently replace SQL [SELECT](#) statements by a memcache access. For that purpose the plugin monitors all SQL statements executed and tries to match the statement string against [MYSQLND\\_MEMCACHE\\_DEFAULT\\_REGEXP](#). In case of a match, the mysqlnd memcache plugin checks whether the [SELECT](#) is accessing only columns of a mapped table and the [WHERE](#) clause is limited to a single key lookup.

In case of the example SQL table, the plugin will use the Memcache interface of the MySQL server to fetch results for a SQL query like [SELECT f1, f2, f3 WHERE id = n](#).

### Example 21.392. Basic example.

```
<?php
$mysqli = new mysqli("host", "user", "passwd", "database");
$memc = new Memcached();
$memc->addServer("host", 11211);
mysqlnd_memcache_set($mysqli, $memc);

/*
  This is a query which queries table test using id as key in the WHERE part
```

```
and is accessing fields f1, f2 and f3. Therefore, mysqlnd_memcache
will intercept it and route it via memcache.
*/
$result = $mysqli->query("SELECT f1, f2, f3 FROM test WHERE id = 1");
while ($row = $result->fetch_row()) {
    print_r($row);
}

/*
This is a query which queries table test but using f1 in the WHERE clause.
Therefore, mysqlnd_memcache can't intercept it. This will be executed
using the MySQL protocol
*/
$mysqli->query("SELECT id FROM test WHERE f1 = 'Lady'");
while ($row = $result->fetch_row()) {
    print_r($row);
}
?>
```

The above example will output:

```
array(
  [f1] => Hello
  [f2] => World
  [f3] => !
)
array(
  [id] => 2
)
```

### 21.9.10.5. Installing/Configuring

Copyright 1997-2012 the PHP Documentation Group. [2906]

#### 21.9.10.5.1. Requirements

Copyright 1997-2012 the PHP Documentation Group. [2906]

PHP: this extension requires PHP 5.4+, version PHP 5.4.4 or newer. The required PHP extensions are [PCRE](#) (enabled by default), and the [memcached](#) extension version 2.0.x.

The `mysqlnd_memcache` Memcache plugin supports all PHP applications and all available PHP MySQL extensions ([mysqli](#), [mysql](#), [PDO\\_MYSQL](#)). The PHP MySQL extension must be configured with [mysqlnd](#) support.

For accessing [InnoDB](#) tables, this PHP extension requires [MySQL Server 5.6.6](#) or newer with the InnoDB Memcache Daemon Plugin enabled.

For accessing [MySQL Cluster NDB](#) tables, this PHP extension requires [MySQL Cluster 7.2](#) or newer with the NDB Memcache API nodes enabled.

#### 21.9.10.5.2. Installation

Copyright 1997-2012 the PHP Documentation Group. [2906]

This [PECL](#) extension is not bundled with PHP.

Information for installing this PECL extension may be found in the manual chapter titled [Installation of PECL extensions](#). Additional information such as new releases, downloads, source files, maintainer information, and a CHANGELOG, can be located here: [http://pecl.php.net/package/mysqlnd\\_memcache](http://pecl.php.net/package/mysqlnd_memcache)

A DLL for this PECL extension is currently unavailable. See also the [building on Windows](#) section.

### 21.9.10.5.3. Runtime Configuration

Copyright 1997-2012 the PHP Documentation Group. [2906]

The behaviour of these functions is affected by settings in `php.ini`.

**Table 21.78. MySQLnd\_memcache Configure Options**

Name	Default	Changeable	Changelog
<code>mysqlnd_memcache.enable</code>		PHP_INI_SYSTEM	Available since 1.0.0

Here's a short explanation of the configuration directives.

`mysqlnd_memcache.enable` integer Enables or disables the plugin. If disabled, the extension will not plug into `mysqlnd` to proxy internal `mysqlnd` C API calls.



#### Note

This option is mainly used by developers to build this extension statically into PHP. General users are encouraged to build this extension as a shared object, and to unload it completely when it is not needed.

### 21.9.10.6. Predefined Constants

Copyright 1997-2012 the PHP Documentation Group. [2906]

The constants below are defined by this extension, and will only be available when the extension has either been compiled into PHP or dynamically loaded at runtime.

#### MySQL Memcache Plugin related

`MYSQLND_MEMCACHE_DEFAULT_FILTER` Default regular expression (PCRE style) used for matching `SELECT` statements that will be mapped into a MySQL Memcache Plugin access point, if possible.

It is also possible to use `mysqlnd_memcache_set`, but the default approach is using this regular expression for pattern matching.

#### Assorted

The version number of this plugin can be obtained by using `MYSQLND_MEMCACHE_VERSION` or `MYSQLND_MEMCACHE_VERSION_ID`. `MYSQLND_MEMCACHE_VERSION` is the string representation of the numerical version number `MYSQLND_MEMCACHE_VERSION_ID`, which is an integer such as 10000. Developers can calculate the version number as follows.

Version (part)	Example
Major*10000	1*10000 = 10000
Minor*100	0*100 = 0
Patch	0 = 0
<code>MYSQLND_MEMCACHE_VERSION_ID</code>	10000

`MYSQLND_MEMCACHE_VERSION` Plugin version string, for example, "1.0.0-alpha". (string)

`MYSQLND_MEMCACHE_VERSION_ID` Plugin version number, for example, 10000. (integer)

## 21.9.10.7. Mysqlnd\_memcache Functions

Copyright 1997-2012 the PHP Documentation Group. [2906]

### 21.9.10.7.1. `mysqlnd_memcache_get_config`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- `mysqlnd_memcache_get_config`

Returns information about the plugin configuration

#### Description

```
array mysqlnd_memcache_get_config(
    mixed connection);
```

This function returns an array of all `mysqlnd_memcache` related configuration information that is attached to the MySQL connection. This includes MySQL, the Memcache object provided via `mysqlnd_memcache_set`, and the table mapping configuration that was automatically collected from the MySQL Server.

#### Parameters

*connection*

A handle to a MySQL Server using one of the MySQL API extensions for PHP, which are `PDO_MYSQL`, `mysqli` or `ext/mysql`.

#### Return Values

An array of `mysqlnd_memcache` configuration information on success, otherwise `FALSE`.


The returned array has these elements:

**Table 21.79.** `mysqlnd_memcache_get_config` array structure

Array Key	Description
<code>memcached</code>	Instance of Memcached associated to this MySQL connection by <code>mysqlnd_memcache_set</code> . You can use this to change settings of the memcache connection, or directly by querying the server on this connection.
<code>pattern</code>	The PCRE regular expression used to match the SQL query sent to the server. Queries matching this pattern will be further analyzed to decide whether the query can be intercepted and sent via the memcache interface or whether the query is sent using the general MySQL protocol to the server. The pattern is either the default pattern ( <code>MYSQLND_MEMCACHE_DEFAULT_REGEXP</code> ) or it is set via <code>mysqlnd_memcache_set</code> .
<code>mappings</code>	An associative array with a list of all configured containers as they were discovered by this plugin. The key for these elements is the name of the container in the MySQL configuration. The value is described below. The contents of this field is created by querying the MySQL Server during association to MySQL and a memcache connection using <code>mysqlnd_memcache_set</code> .
<code>mapping_query</code>	An SQL query used during <code>mysqlnd_memcache_set</code> to identify the

Array Key	Description
	available containers and mappings. The result of that query is provided in the mappings element.

**Table 21.80. Mapping entry structure**

Array Key	Description
prefix	A prefix used while accessing data via memcache. With the MySQL InnoDB Memcache Daemon plugin, this usually begins with @@ and ends with a configurable separator. This prefix is placed in front of the key value while using the memcache protocol.
schema_name	Name of the schema (database) which contains the table being accessed.
table_name	Name of the table which contains the data accessible via memcache protocol.
id_field_name	Name of the database field (column) with the id used as key when accessing the table via memcache. Often this is the database field having a primary key.
separator	<p>The separator used to split the different field values. This is needed as memcache only provides access to a single value while MySQL can map multiple columns to this value.</p> <div style="display: flex; align-items: center;"> <div style="margin-right: 10px;"></div> <div> <p><b>Note</b></p> <p>The separator, which can be set in the MySQL Server configuration, should not be part of any value retrieved via memcache because proper mapping can't be guaranteed.</p> </div> </div>
fields	An array with the name of all fields available for this mapping.

## Examples

### Example 21.393. `mysqld_memcache_get_config` example

```
<?php
$mysqli = new mysqli("host", "user", "passwd", "database");
$memc = new Memcached();
$memc->addServer("host", 11211);
mysqld_memcache_set($mysqli, $memc);

var_dump(mysqld_memcache_get_config($mysqli));
?>
```

The above example will output:

```
array(4) {
  ["memcached"]=>
```

```

object(Memcached)#2 (0) {
}
["pattern"]=>
string(125) "/^\s*SELECT\s*(.+?)\s*FROM\s*`?([a-z0-9_]+)`?\s*WHERE\s*`?([a-z0-9_]+)`?\s*=\s*(?(['"])[^
["mappings"]=>
array(1) {
  ["mymem_test"]=>
  array(6) {
    ["prefix"]=>
    string(13) "@@mymem_test."
    ["schema_name"]=>
    string(4) "test"
    ["table_name"]=>
    string(10) "mymem_test"
    ["id_field_name"]=>
    string(2) "id"
    ["separator"]=>
    string(1) "|"
    ["fields"]=>
    array(3) {
      [0]=>
      string(2) "f1"
      [1]=>
      string(2) "f2"
      [2]=>
      string(2) "f3"
    }
  }
}
["mapping_query"]=>
string(209) "      SELECT c.name,
              CONCAT('@@', c.name, (SELECT value FROM innodb_memcache.config_options WHERE name
              c.db_schema,
              c.db_table,
              c.key_columns,
              c.value_columns,
              (SELECT value FROM innodb_memcache.config_options WHERE name = 'separator') AS s
              FROM innodb_memcache.containers c"
}

```

## See Also

[mysqlnd\\_memcache\\_set](#)

### 21.9.10.7.2. `mysqlnd_memcache_set`

Copyright 1997-2012 the PHP Documentation Group. [2906]

- [mysqlnd\\_memcache\\_set](#)

Associate a MySQL connection with a Memcache connection

## Description

```

bool mysqlnd_memcache_set(
    mixed mysql_connection,
    Memcached memcache_connection,
    string pattern,
    callback callback);

```

Associate *mysql\_connection* with *memcache\_connection* using *pattern* as a PCRE regular expression, and *callback* as a notification callback or to unset the association of *mysql\_connection*.

While associating a MySQL connection with a Memcache connection, this function will query the MySQL Server for its configuration. It will automatically detect whether the server is configured to use the InnoDB Memcache Daemon Plugin or MySQL Cluster NDB Memcache support. It will also query

the server to automatically identify exported tables and other configuration options. The results of this automatic configuration can be retrieved using `mysqlnd_memcache_get_config`.

## Parameters

<code>mysql_connection</code>	A handle to a MySQL Server using one of the MySQL API extensions for PHP, which are <code>PDO_MYSQL</code> , <code>mysqli</code> or <code>ext/mysql</code> .
<code>memcache_connection</code>	A <code>Memcached</code> instance with a connection to the MySQL Memcache Daemon plugin. If this parameter is omitted, then <code>mysql_connection</code> will be unassociated from any memcache connection. And if a previous association exists, then it will be replaced.
<code>pattern</code>	A regular expression in <code>Perl Compatible Regular Expression</code> syntax used to identify potential Memcache-queries. The query should have three sub patterns. The first subpattern contains the requested field list, the second the name of the ID column from the query and the third the requested value. If this parameter is omitted or is set to <code>NULL</code> , then a default pattern will be used.
<code>callback</code>	A callback which will be used whenever a query is being sent to MySQL. The callback will receive a single boolean parameter telling if a query was sent via Memcache.

## Return Values

`TRUE` if the association or disassociation is successful, otherwise `FALSE` if there is an error.

## Examples

**Example 21.394.** `mysqlnd_memcache_set` example with `var_dump` as a simple debugging callback.

```
<?php
$mysqli = new mysqli("host", "user", "passwd", "database");
$memc = new Memcached();
$memc->addServer("host", 11211);
mysqlnd_memcache_set($mysqli, $memc, NULL, 'var_dump');

/* This query will be intercepted and executed via Memcache protocol */
echo "Sending query for id via Memcache: ";
$mysqli->query("SELECT f1, f2, f3 FROM test WHERE id = 1");

/* f1 is not configured as valid key field, this won't be sent via Memcache */
echo "Sending query for f1 via Memcache: ";
$mysqli->query("SELECT id FROM test WHERE f1 = 1");

mysqlnd_memcache_set($mysqli);

/* Now the regular MySQL protocol will be used */
echo "var_dump won't be invoked: ";
$mysqli->query("SELECT f1, f2, f3 WHERE id = 1");

?>
```

The above example will output:

```
Sending query for id via Memcache: bool(true)
Sending query for f1 via Memcache: bool(false)
var_dump won't be invoked:
```

---

**See Also**[mysqlnd\\_memcache\\_get\\_config](#)

### 21.9.10.8. Change History

Copyright 1997-2012 the PHP Documentation Group. [2906]

This change history is a high level summary of selected changes that may impact applications and/or break backwards compatibility.

See also the [CHANGES](#) file in the source distribution for a complete list of changes.

#### 21.9.10.8.1. PECL/mysqlnd\_memcache 1.0 series

Copyright 1997-2012 the PHP Documentation Group. [2906]

1.0.0-alpha

- Release date: TBD
- Motto/theme: Basic mapping of SQL SELECT to a MySQL Memcache plugin access.

The initial release does map basic SQL SELECT statements to a MySQL Memcache plugin access. This bares potential performance benefits as the direct key-value access to MySQL storage using the Memcache protocol is usually faster than using SQL access.

### 21.9.11. Connector/PHP

This documentation, and other publications, sometimes uses the term [Connector/PHP](#). This term refers to the full set of MySQL related functionality in PHP, which includes the three APIs that are described above, along with the [mysqlnd](#) core library and all of its plugins.

### 21.9.12. Common Problems with MySQL and PHP

- [Error: Maximum Execution Time Exceeded](#): This is a PHP limit; go into the [php.ini](#) file and set the maximum execution time up from 30 seconds to something higher, as needed. It is also not a bad idea to double the RAM allowed per script to 16MB instead of 8MB.
- [Fatal error: Call to unsupported or undefined function mysql\\_connect\(\) in ...](#): This means that your PHP version isn't compiled with MySQL support. You can either compile a dynamic MySQL module and load it into PHP or recompile PHP with built-in MySQL support. This process is described in detail in the PHP manual.
- [Error: Undefined reference to 'uncompress'](#): This means that the client library is compiled with support for a compressed client/server protocol. The fix is to add `-lz` last when linking with `-lmysqlclient`.
- [Error: Client does not support authentication protocol](#): This is most often encountered when trying to use the older [mysql](#) extension with MySQL 4.1.1 and later. Possible solutions are: downgrade to MySQL 4.0; switch to PHP 5 and the newer [mysqli](#) extension; or configure the MySQL server with the [old\\_passwords](#) [517] system variable set to 1. (See [Section C.5.2.4](#), "Client does not support authentication protocol", for more information.)

### 21.9.13. Enabling Both [mysql](#) and [mysqli](#) in PHP

If you're experiencing problems with enabling both the [mysql](#) and the [mysqli](#) extension when building PHP on Linux yourself, you should try the following procedure.



1. Configure PHP like this:

```
./configure --with-mysqli=/usr/bin/mysql_config --with-mysql=/usr
```

2. Edit the `Makefile` and search for a line that starts with `EXTRA_LIBS`. It might look like this (all on one line):

```
EXTRA_LIBS = -lcrypt -lcrypt -lmysqlclient -lz -lresolv -lm -ldl -lnsl
-lxml2 -lz -lm -lxml2 -lz -lm -lmysqlclient -lz -lcrypt -lnsl -lm
-lxml2 -lz -lm -lcrypt -lxml2 -lz -lm -lcrypt
```

Remove all duplicates, so that the line looks like this (all on one line):

```
EXTRA_LIBS = -lcrypt -lcrypt -lmysqlclient -lz -lresolv -lm -ldl -lnsl
-lxml2
```

3. Build and install PHP:

```
make
make install
```

## 21.10. MySQL Perl API

The Perl `DBI` module provides a generic interface for database access. You can write a DBI script that works with many different database engines without change. To use DBI with MySQL, install the following:

1. The `DBI` module.
2. The `DBD::mysql` module. This is the DataBase Driver (DBD) module for Perl.
3. Optionally, the DBD module for any other type of database server you want to access.

Perl DBI is the recommended Perl interface. It replaces an older interface called `mysqlperl`, which should be considered obsolete.

These sections contain information about using Perl with MySQL and writing MySQL applications in Perl:

- For installation instructions for Perl DBI support, see [Section 2.13, “Perl Installation Notes”](#).
- For an example of reading options from option files, see [Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”](#).
- For secure coding tips, see [Section 6.1.1, “Security Guidelines”](#).
- For debugging tips, see [Section 22.4.1.4, “Debugging mysqld under gdb”](#).
- For some Perl-specific environment variables, see [Section 2.12, “Environment Variables”](#).
- For considerations for running on Mac OS X, see [Section 2.4.5, “Using the Bundled MySQL on Mac OS X Server”](#).
- For ways to quote string literals, see [Section 9.1.1, “String Literals”](#).

DBI information is available at the command line, online, or in printed form:

- Once you have the `DBI` and `DBD::mysql` modules installed, you can get information about them at the command line with the `perldoc` command:

```
shell> perldoc DBI
shell> perldoc DBI:FAQ
shell> perldoc DBD:mysql
```

You can also use `pod2man`, `pod2html`, and so on to translate this information into other formats.

- For online information about Perl DBI, visit the DBI Web site, <http://dbi.perl.org/>. That site hosts a general DBI mailing list. Oracle Corporation hosts a list specifically about `DBD::mysql`; see [Section 1.6.1, “MySQL Mailing Lists”](#).
- For printed information, the official DBI book is *Programming the Perl DBI* (Alligator Descartes and Tim Bunce, O'Reilly & Associates, 2000). Information about the book is available at the DBI Web site, <http://dbi.perl.org/>.

For information that focuses specifically on using DBI with MySQL, see *MySQL and Perl for the Web* (Paul DuBois, New Riders, 2001). This book's Web site is <http://www.kitebird.com/mysql-perl/>.

## 21.11. MySQL Python API

`MySQLdb` is a third-party driver that provides MySQL support for Python, compliant with the Python DB API version 2.0. It can be found at <http://sourceforge.net/projects/mysql-python/>.

The new MySQL Connector/Python component provides an interface to the same Python API, and is built into the MySQL Server and supported by Oracle. See [Section 21.6, “MySQL Connector/Python”](#) for details on the Connector, as well as coding guidelines for Python applications and sample Python code.

## 21.12. MySQL Ruby APIs

Two APIs are available for Ruby programmers developing MySQL applications:

- The MySQL/Ruby API is based on the `libmysqlclient` API library. For information on installing and using the MySQL/Ruby API, see [Section 21.12.1, “The MySQL/Ruby API”](#).
- The Ruby/MySQL API is written to use the native MySQL network protocol (a native driver). For information on installing and using the Ruby/MySQL API, see [Section 21.12.2, “The Ruby/MySQL API”](#).

For background and syntax information about the Ruby language, see [Ruby Programming Language](#).

### 21.12.1. The MySQL/Ruby API

The MySQL/Ruby module provides access to MySQL databases using Ruby through `libmysqlclient`.

For information on installing the module, and the functions exposed, see [MySQL/Ruby](#).

### 21.12.2. The Ruby/MySQL API

The Ruby/MySQL module provides access to MySQL databases using Ruby through a native driver interface using the MySQL network protocol.

For information on installing the module, and the functions exposed, see [Ruby/MySQL](#).

## 21.13. MySQL Tcl API

`MySQLtcl` is a simple API for accessing a MySQL database server from the [Tcl programming language](#). It can be found at <http://www.xdobry.de/mysqltcl/>.

## 21.14. MySQL Eiffel Wrapper

Eiffel MySQL is an interface to the MySQL database server using the [Eiffel programming language](#), written by Michael Ravits. It can be found at <http://efsa.sourceforge.net/archive/ravits/mysql.htm>.

---

# Chapter 22. Extending MySQL

## Table of Contents

22.1. MySQL Internals .....	3483
22.1.1. MySQL Threads .....	3483
22.1.2. The MySQL Test Suite .....	3484
22.2. The MySQL Plugin API .....	3485
22.2.1. Plugin API Characteristics .....	3485
22.2.2. Plugin API Components .....	3486
22.2.3. Types of Plugins .....	3487
22.2.4. Writing Plugins .....	3490
22.2.5. MySQL Services for Plugins .....	3530
22.3. Adding New Functions to MySQL .....	3532
22.3.1. Features of the User-Defined Function Interface .....	3532
22.3.2. Adding a New User-Defined Function .....	3532
22.3.3. Adding a New Native Function .....	3542
22.4. Debugging and Porting MySQL .....	3543
22.4.1. Debugging a MySQL Server .....	3544
22.4.2. Debugging a MySQL Client .....	3550
22.4.3. The DEBUG Package .....	3550

## 22.1. MySQL Internals

This chapter describes a lot of things that you need to know when working on the MySQL code. To track or contribute to MySQL development, follow the instructions in [Section 2.9.3, “Installing MySQL from a Development Source Tree”](#). If you are interested in MySQL internals, you should also subscribe to our `internals` mailing list. This list has relatively low traffic. For details on how to subscribe, please see [Section 1.6.1, “MySQL Mailing Lists”](#). Many MySQL developers at Oracle Corporation are on the `internals` list and we help other people who are working on the MySQL code. Feel free to use this list both to ask questions about the code and to send patches that you would like to contribute to the MySQL project!

### 22.1.1. MySQL Threads

The MySQL server creates the following threads:

- Connection manager threads handle client connection requests on the network interfaces that the server listens to. On all platforms, one manager thread handles TCP/IP connection requests. On Unix, this manager thread also handles Unix socket file connection requests. On Windows, a manager thread handles shared-memory connection requests, and another handles named-pipe connection requests. The server does not create threads to handle interfaces that it does not listen to. For example, a Windows server that does not have support for named-pipe connections enabled does not create a thread to handle them.
- Connection manager threads associate each client connection with a thread dedicated to it that handles authentication and request processing for that connection. Manager threads create a new thread when necessary but try to avoid doing so by consulting the thread cache first to see whether it contains a thread that can be used for the connection. When a connection ends, its thread is returned to the thread cache if the cache is not full.

For information about tuning the parameters that control thread resources, see [Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”](#).

- On a master replication server, connections from slave servers are handled like client connections: There is one thread per connected slave.

- On a slave replication server, an I/O thread is started to connect to the master server and read updates from it. An SQL thread is started to apply updates read from the master. These two threads run independently and can be started and stopped independently.
- A signal thread handles all signals. This thread also normally handles alarms and calls `process_alarm()` to force timeouts on connections that have been idle too long.
- If InnoDB is used, there will be additional read and write threads by default. The number of these are controlled by the `innodb_read_io_threads` [1743] and `innodb_write_io_threads` [1754] parameters. See Section 14.2.6, “InnoDB Startup Options and System Variables”.
- If `mysqld` is compiled with `-DUSE_ALARM_THREAD`, a dedicated thread that handles alarms is created. This is only used on some systems where there are problems with `sigwait()` or if you want to use the `thr_alarm()` code in your application without a dedicated signal handling thread.
- If the server is started with the `--flush_time=val` [482] option, a dedicated thread is created to flush all tables every `val` seconds.
- Each table for which `INSERT DELAYED` statements are issued gets its own thread. See Section 13.2.5.2, “INSERT DELAYED Syntax”.
- If the event scheduler is active, there is one thread for the scheduler, and a thread for each event currently running. See Section 18.4.1, “Event Scheduler Overview”.

`mysqladmin processlist` only shows the connection, `INSERT DELAYED`, replication, and event threads.

## 22.1.2. The MySQL Test Suite

The test system that is included in Unix source and binary distributions makes it possible for users and developers to perform regression tests on the MySQL code. These tests can be run on Unix.

You can also write your own test cases. For information about the MySQL Test Framework, including system requirements, see the manual available at <http://dev.mysql.com/doc/mysqctest/2.0/en/>.

The current set of test cases doesn't test everything in MySQL, but it should catch most obvious bugs in the SQL processing code, operating system or library issues, and is quite thorough in testing replication. Our goal is to have the tests cover 100% of the code. We welcome contributions to our test suite. You may especially want to contribute tests that examine the functionality critical to your system because this ensures that all future MySQL releases work well with your applications.

The test system consists of a test language interpreter (`mysqctest`), a Perl script to run all tests (`mysql-test-run.pl`), the actual test cases written in a special test language, and their expected results. To run the test suite on your system after a build, type `make test` from the source root directory, or change location to the `mysql-test` directory and type `./mysql-test-run.pl`. If you have installed a binary distribution, change location to the `mysql-test` directory under the installation root directory (for example, `/usr/local/mysql/mysql-test`), and run `./mysql-test-run.pl`. All tests should succeed. If any do not, feel free to try to find out why and report the problem if it indicates a bug in MySQL. See Section 1.7, “How to Report Bugs or Problems”.

If one test fails, you should run `mysql-test-run.pl` with the `--force` option to check whether any other tests fail.

If you have a copy of `mysqld` running on the machine where you want to run the test suite, you do not have to stop it, as long as it is not using ports 9306 or 9307. If either of those ports is taken, you should set the `MTR_BUILD_THREAD` environment variable to an appropriate value, and the test suite will use a different set of ports for master, slave, and NDB). For example:

```
shell> export MTR_BUILD_THREAD=31
shell> ./mysql-test-run.pl [options] [test_name]
```

In the `mysql-test` directory, you can run an individual test case with `./mysql-test-run.pl test_name`.

If you have a question about the test suite, or have a test case to contribute, send an email message to the MySQL [internals](#) mailing list. See [Section 1.6.1, “MySQL Mailing Lists”](#).

## 22.2. The MySQL Plugin API

MySQL supports a plugin API that enables creation of server components. Plugins can be loaded at server startup, or loaded and unloaded at runtime without restarting the server. The API is generic and does not specify what plugins can do. The components supported by this interface include, but are not limited to, storage engines, full-text parser plugins, and server extensions.

For example, full-text parser plugins can be used to replace or augment the built-in full-text parser. A plugin can parse text into words using rules that differ from those used by the built-in parser. This can be useful if you need to parse text with characteristics different from those expected by the built-in parser.

The plugin interface is more general than the older user-defined function (UDF) interface.

The plugin interface uses the `plugin` table in the `mysql` database to record information about plugins that have been installed permanently with the `INSTALL PLUGIN` statement. This table is created as part of the MySQL installation process. Plugins can also be installed for a single server invocation with the `--plugin-load` option. Plugins installed this way are not recorded in the `plugin` table. See [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

MySQL 5.6 supports an API for client plugins in addition to that for server plugins. This is used, for example, by authentication plugins where a server-side plugin and a client-side plugin cooperate to enable clients to connect to the server through a variety of authentication methods.

### Additional Resources

The book *MySQL 5.1 Plugin Development* by Sergei Golubchik and Andrew Hutchings provides a wealth of detail about the plugin API. Despite the fact that the book's title refers to MySQL Server 5.1, most of the information in it applies to later versions as well.

#### 22.2.1. Plugin API Characteristics

The server plugin API has these characteristics:

- All plugins have several things in common.

Each plugin has a name that it can be referred to in SQL statements, as well as other metadata such as an author and a description that provide other information. This information can be examined in the `INFORMATION_SCHEMA.PLUGINS` table or using the `SHOW PLUGINS` statement.

- The plugin framework is extendable to accommodate different kinds of plugins.

Although some aspects of the plugin API are common to all types of plugins, the API also permits type-specific interface elements so that different types of plugins can be created. A plugin with one purpose can have an interface most appropriate to its own requirements and not the requirements of some other plugin type.

Interfaces for several types of plugins exist, such as storage engines, full-text parser, and `INFORMATION_SCHEMA` tables. Others can be added.

- Plugins can expose information to users.

A plugin can implement system and status variables that are available through the `SHOW VARIABLES` and `SHOW STATUS` statements.

- The plugin API includes versioning information.

The version information included in the plugin API enables a plugin library and each plugin that it contains to be self-identifying with respect to the API version that was used to build the library. If the

API changes over time, the version numbers will change, but a server can examine a given plugin library's version information to determine whether it supports the plugins in the library.

There are two types of version numbers. The first is the version for the general plugin framework itself. Each plugin library includes this kind of version number. The second type of version applies to individual plugins. Each specific type of plugin has a version for its interface, so each plugin in a library has a type-specific version number. For example, a library containing a full-text parser plugin has a general plugin API version number, and the plugin has a version number specific to the full-text plugin interface.

- The plugin API implements security restrictions.

A plugin library must be installed in a specific dedicated directory for which the location is controlled by the server and cannot be changed at runtime. Also, the library must contain specific symbols that identify it as a plugin library. The server will not load something as a plugin if it was not built as a plugin.

- Plugins have access to server services.

The services interface exposes server functionality that plugins can access using ordinary function calls. For details, see [Section 22.2.5, “MySQL Services for Plugins”](#).

In some respects, the server plugin API is similar to the older user-defined function (UDF) API that it supersedes, but the plugin API has several advantages over the older interface. For example, UDFs had no versioning information. Also, the newer plugin interface eliminates the security issues of the older UDF interface. The older interface for writing nonplugin UDFs permitted libraries to be loaded from any directory searched by the system's dynamic linker, and the symbols that identified the UDF library were relatively nonspecific.

The client plugin API has similar architectural characteristics, but client plugins have no direct access to the server the way server plugins do.

## 22.2.2. Plugin API Components

The server plugin implementation comprises several components.

SQL statements:

- `INSTALL PLUGIN` registers a plugin in the `mysql.plugin` table and loads the plugin code.
- `UNINSTALL PLUGIN` unregisters a plugin from the `mysql.plugin` table and unloads the plugin code.
- The `WITH PARSER` clause for full-text index creation associates a full-text parser plugin with a given `FULLTEXT` index.
- `SHOW PLUGINS` displays information about server plugins.

Command-line options and system variables:

- The `--plugin-load` [\[435\]](#) option enables plugins to be loaded at server startup time.
- The `plugin_dir` [\[525\]](#) system variable indicates the location of the directory where all plugins must be installed. The value of this variable can be specified at server startup with a `--plugin_dir=path` [\[525\]](#) option. `mysql_config --plugindir` displays the default plugin directory path name.

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

Plugin-related tables:

- The `INFORMATION_SCHEMA.PLUGINS` table contains plugin information.
- The `mysql.plugin` table lists each plugin that was installed with `INSTALL PLUGIN` and is required for plugin use. For new MySQL installations, this table is created during the installation process.

The client plugin implementation is simpler:

- For the `mysql_options()` C API function, the `MYSQL_DEFAULT_AUTH` and `MYSQL_PLUGIN_DIR` options enable client programs to load authentication plugins.
- There are C API functions that enable management of client plugins.

To examine how MySQL implements plugins, consult the following source files in a MySQL source distribution:

- In the `include/mysql` directory, `plugin.h` exposes the public plugin API. This file should be examined by anyone who wants to write a plugin library. `plugin_XXX.h` files provide additional information that pertains to specific types of plugins. `client_plugin.h` contains information specific to client plugins.
- In the `sql` directory, `sql_plugin.h` and `sql_plugin.cc` comprise the internal plugin implementation. `sql_acl.cc` is where the server uses authentication plugins. These files need not be consulted by plugin developers. They may be of interest for those who want to know more about how the server handles plugins.
- In the `sql-common` directory, `client_plugin.h` implements the C API client plugin functions, and `client.c` implements client authentication support. These files need not be consulted by plugin developers. They may be of interest for those who want to know more about how the server handles plugins.

### 22.2.3. Types of Plugins

The plugin API enables creation of plugins that implement several capabilities:

- Storage engines
- Full-text parsers
- Daemons
- `INFORMATION_SCHEMA` tables
- Semisynchronous replication
- Auditing
- Authentication

The following sections provide an overview of these plugin types.

#### 22.2.3.1. Storage Engine Plugins

The pluggable storage engine architecture used by MySQL Server enables storage engines to be written as plugins and loaded into and unloaded from a running server. For a description of this architecture, see [Section 14.12, “Overview of MySQL Storage Engine Architecture”](#).

For information on how to use the plugin API to write storage engines, see [MySQL Internals: Writing a Custom Storage Engine](#).

#### 22.2.3.2. Full-Text Parser Plugins

MySQL has a built-in parser that it uses by default for full-text operations (parsing text to be indexed, or parsing a query string to determine the terms to be used for a search). For full-text processing,



“parsing” means extracting words from text or a query string based on rules that define which character sequences make up a word and where word boundaries lie.

When parsing for indexing purposes, the parser passes each word to the server, which adds it to a full-text index. When parsing a query string, the parser passes each word to the server, which accumulates the words for use in a search.

The parsing properties of the built-in full-text parser are described in [Section 12.9, “Full-Text Search Functions”](#). These properties include rules for determining how to extract words from text. The parser is influenced by certain system variables such as `ft_min_word_len` [484] and `ft_max_word_len` [484] that cause words shorter or longer to be excluded, and by the stopword list that identifies common words to be ignored.

The plugin API enables you to provide a full-text parser of your own so that you have control over the basic duties of a parser. A parser plugin can operate in either of two roles:

- The plugin can replace the built-in parser. In this role, the plugin reads the input to be parsed, splits it up into words, and passes the words to the server (either for indexing or for word accumulation).

One reason to use a parser this way is that you need to use different rules from those of the built-in parser for determining how to split up input into words. For example, the built-in parser considers the text “case-sensitive” to consist of two words “case” and “sensitive,” whereas an application might need to treat the text as a single word.

- The plugin can act in conjunction with the built-in parser by serving as a front end for it. In this role, the plugin extracts text from the input and passes the text to the parser, which splits up the text into words using its normal parsing rules. In particular, this parsing will be affected by the `ft_XXX` system variables and the stopword list.

One reason to use a parser this way is that you need to index content such as PDF documents, XML documents, or `.doc` files. The built-in parser is not intended for those types of input but a plugin can pull out the text from these input sources and pass it to the built-in parser.

It is also possible for a parser plugin to operate in both roles. That is, it could extract text from nonplaintext input (the front end role), and also parse the text into words (thus replacing the built-in parser).

A full-text plugin is associated with full-text indexes on a per-index basis. That is, when you install a parser plugin initially, that does not cause it to be used for any full-text operations. It simply becomes available. For example, a full-text parser plugin becomes available to be named in a `WITH PARSE` clause when creating individual `FULLTEXT` indexes. To create such an index at table-creation time, do this:

```
CREATE TABLE t
(
  doc CHAR(255),
  FULLTEXT INDEX (doc) WITH PARSE my_parser
) ENGINE=MyISAM;
```

Or you can add the index after the table has been created:

```
ALTER TABLE t ADD FULLTEXT INDEX (doc) WITH PARSE my_parser;
```

The only SQL change for associating the parser with the index is the `WITH PARSE` clause. Searches are specified as before, with no changes needed for queries.

When you associate a parser plugin with a `FULLTEXT` index, the plugin is required for using the index. If the parser plugin is dropped, any index associated with it becomes unusable. Any attempt to use it a table for which a plugin is not available results in an error, although `DROP TABLE` is still possible.

For more information about full-text plugins, see [Section 22.2.4.4, “Writing Full-Text Parser Plugins”](#).

### 22.2.3.3. Daemon Plugins



A daemon plugin is a simple type of plugin used for code that should be run by the server but that does not communicate with it. MySQL distributions include an example daemon plugin that writes periodic heartbeat messages to a file.

For more information about daemon plugins, see [Section 22.2.4.5, “Writing Daemon Plugins”](#).

#### 22.2.3.4. **INFORMATION\_SCHEMA** Plugins

**INFORMATION\_SCHEMA** plugins enable the creation of tables containing server metadata that are exposed to users through the **INFORMATION\_SCHEMA** database. For example, **InnoDB** uses **INFORMATION\_SCHEMA** plugins to provide tables that contain information about current transactions and locks.

For more information about **INFORMATION\_SCHEMA** plugins, see [Section 22.2.4.6, “Writing INFORMATION\\_SCHEMA Plugins”](#).

#### 22.2.3.5. **Semisynchronous Replication** Plugins

MySQL replication is asynchronous by default. With semisynchronous replication, a commit performed on the master side blocks before returning to the session that performed the transaction until at least one slave acknowledges that it has received and logged the events for the transaction. Semisynchronous replication is implemented through complementary master and client plugins. See [Section 16.3.8, “Semisynchronous Replication”](#).

For more information about semisynchronous replication plugins, see [Section 22.2.4.7, “Writing Semisynchronous Replication Plugins”](#).

#### 22.2.3.6. **Audit** Plugins

In MySQL 5.6, the server provides a pluggable audit interface that enables information about server operations to be reported to interested parties. Currently, audit notification occurs for these operations (although the interface is general and the server could be modified to report others):

- Write a message to the general query log (if the log is enabled)
- Write a message to the error log
- Send a query result to a client

Audit plugins may register with the audit interface to receive notification about server operations. When an auditable event occurs within the server, the server determines whether notification is needed. For each registered audit plugin, the server checks the event against those event classes in which the plugin is interested and passes the event to the plugin if there is a match.

This interface enables audit plugins to receive notifications only about operations in event classes they consider significant and to ignore others. The interface provides for categorization of operations into event classes and further division into event subclasses within each class.

When an audit plugin is notified of an auditable event, it receives a pointer to the current THD structure and a pointer to a structure that contains information about the event. The plugin can examine the event and perform whatever auditing actions are appropriate. For example, the plugin can see what statement produced a result set or was logged, the number of rows in a result, who the current user was for an operation, or the error code for failed operations.

For more information about audit plugins, see [Section 22.2.4.8, “Writing Audit Plugins”](#).

#### 22.2.3.7. **Authentication** Plugins

MySQL 5.6 supports pluggable authentication. Authentication plugins exist on both the server and client sides. Plugins on the server side implement authentication methods for use by clients when they connect to the server. A plugin on the client side communicates with a server-side plugin to provide the authentication information that it requires. A client-side plugin may interact with the user, performing

tasks such as soliciting a password or other authentication credentials to be sent to the server. See [Section 6.3.7, “Pluggable Authentication”](#).

Pluggable authentication also enables proxy user capability, in which one user takes the identity of another user. A server-side authentication plugin can return to the server the name of the user whose identity the connecting user should have. See [Section 6.3.8, “Proxy Users”](#).

For more information about authentication plugins, see [Section 22.2.4.9, “Writing Authentication Plugins”](#).

### 22.2.3.8. Password-Validation Plugins

As of MySQL 5.6.6, the server provides an interface for writing plugins that test passwords. Such a plugin implements two capabilities:

- Rejection of too-weak passwords in statements that assign passwords (such as `CREATE USER`, `GRANT`, and `SET PASSWORD` statements), and passwords given as arguments to the `PASSWORD( )` [1263] and `OLD_PASSWORD( )` [1263] functions.
- Assessing the strength of potential passwords for the `VALIDATE_PASSWORD_STRENGTH( )` [1265] SQL function.

For information about writing this type of plugin, see [Section 22.2.4.10, “Writing Password-Validation Plugins”](#).

## 22.2.4. Writing Plugins

To create a plugin library, you must provide the required descriptor information that indicates what plugins the library file contains, and write the interface functions for each plugin.

Every server plugin must have a general descriptor that provides information to the plugin API, and a type-specific descriptor that provides information about the plugin interface for a given type of plugin. The structure of the general descriptor is the same for all plugin types. The structure of the type-specific descriptor varies among plugin types and is determined by the requirements of what the plugin needs to do. The server plugin interface also enables plugins to expose status and system variables. These variables become visible through the `SHOW STATUS` and `SHOW VARIABLES` statements and the corresponding `INFORMATION_SCHEMA` tables.

For client-side plugins, the architecture is a bit different. Each plugin must have a descriptor, but there is no division into separate general and type-specific descriptors. Instead, the descriptor begins with a fixed set of members common to all client plugin types, and the common members are followed by any additional members required to implement the specific plugin type.

You can write plugins in C or C++ (or another language that can use C calling conventions). Plugins are loaded and unloaded dynamically, so your operating system must support dynamic loading and you must have compiled the calling application dynamically (not statically). For server plugins, this means that `mysqld` must be compiled dynamically.

A server plugin contains code that becomes part of the running server, so when you write the plugin, you are bound by any and all constraints that otherwise apply to writing server code. For example, you may have problems if you attempt to use functions from the `libstdc++` library. These constraints may change in future versions of the server, so it is possible that server upgrades will require revisions to plugins originally written for older servers. For information about these constraints, see [Section 2.9.4, “MySQL Source-Configuration Options”](#), and [Section 2.9.5, “Dealing with Problems Compiling MySQL”](#).

Client plugin writers should avoid dependencies on what symbols the calling application has because you cannot be sure what applications will use the plugin.

### 22.2.4.1. Overview of Plugin Writing

The following procedure provides an overview of the steps needed to create a plugin library. The next sections provide additional details on setting plugin data structures and writing specific types of plugins.

1. In the plugin source file, include the header files that the plugin library needs. The `plugin.h` file is required, and the library might require other files as well. For example:

```
#include <stdlib.h>
#include <ctype.h>
#include <mysql/plugin.h>
```

2. Set up the descriptor information for the plugin library file. For server plugins, write the library descriptor, which must contain the general plugin descriptor for each server plugin in the file. For more information, see [Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”](#). In addition, set up the type-specific descriptor for each server plugin in the library. Each plugin's general descriptor points to its type-specific descriptor.

For client plugins, write the client descriptor. For more information, see [Section 22.2.4.2.3, “Client Plugin Descriptors”](#).

3. Write the plugin interface functions for each plugin. For example, each plugin's general plugin descriptor points to the initialization and deinitialization functions that the server should invoke when it loads and unloads the plugin. The plugin's type-specific description may also point to interface functions.
4. For server plugins, set up the status and system variables, if there are any.
5. Compile the plugin library as a shared library and install it in the plugin directory. For more information, see [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#).
6. For server plugins, register the plugin with the server. For more information, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).
7. Test the plugin to verify that it works properly.

### 22.2.4.2. Plugin Data Structures

A plugin library file includes descriptor information to indicate what plugins it contains.

If the plugin library contains any server plugins, it must include the following descriptor information:

- A library descriptor indicates the general server plugin API version number used by the library and contains a general plugin descriptor for each server plugin in the library. To provide the framework for this descriptor, invoke two macros from the `plugin.h` header file:

```
mysql_declare_plugin(name)
... one or more server plugin descriptors here ...
mysql_declare_plugin_end;
```

The macros expand to provide a declaration for the API version automatically. You must provide the plugin descriptors.

- Within the library descriptor, each general server plugin is described by a `st_mysql_plugin` structure. This plugin descriptor structure contains information that is common to every type of server plugin: A value that indicates the plugin type; the plugin name, author, description, and license type; pointers to the initialization and deinitialization functions that the server invokes when it loads and unloads the plugin, and pointers to any status or system variables the plugin implements.
- Each general server plugin descriptor within the library descriptor also contains a pointer to a type-specific plugin descriptor. The structure of the type-specific descriptors varies from one plugin type to another because each type of plugin can have its own API. A type-specific plugin descriptor contains a type-specific API version number and pointers to the functions that are needed to implement that plugin type. For example, a full-text parser plugin has initialization and deinitialization functions, and a main parsing function. The server invokes these functions when it uses the plugin to parse text.

The plugin library also contains the interface functions that are referenced by the general and type-specific descriptors for each plugin in the library.

If the plugin library contains a client plugin, it must include a descriptor for the plugin. The descriptor begins with a fixed set of members common to all client plugins, followed by any members specific to the plugin type. To provide the descriptor framework, invoke two macros from the `client_plugin.h` header file:

```
mysql_declare_client_plugin(plugin_type)
... members common to all client plugins ...
... type-specific extra members ...
mysql_end_client_plugin;
```

The plugin library also contains any interface functions referenced by the client descriptor.

The `mysql_declare_plugin()` and `mysql_declare_client_plugin()` macros differ somewhat in how they can be invoked, which has implications for the contents of plugin libraries. The following guidelines summarize the rules:

- `mysql_declare_plugin()` and `mysql_declare_client_plugin()` can both be used in the same source file, which means that a plugin library can contain both server and client plugins. However, each of `mysql_declare_plugin()` and `mysql_declare_client_plugin()` can be used at most once.
- `mysql_declare_plugin()` permits multiple server plugin declarations, so a plugin library can contain multiple server plugins.
- `mysql_declare_client_plugin()` permits only a single client plugin declaration. To create multiple client plugins, separate plugin libraries must be used.

When a client program looks for a client plugin that is in a plugin library and not built into `libmysqlclient`, it looks for a file with a basename that is the same as the plugin name. For example, if a program needs to use a client authentication plugin named `auth_xxx` on a system that uses `.so` as the library suffix, it looks in the file named `auth_xxx.so`. (On Mac OS X, the program looks first for `auth_xxx.dylib`, then for `auth_xxx.so`.) For this reason, if a plugin library contains a client plugin, the library must have the same basename as that plugin.

The same is not true for a library that contains server plugins. The `--plugin-load` [435] option and the `INSTALL PLUGIN` statement provide the library file name explicitly, so there need be no explicit relationship between the library name and the name of any server plugins it contains.

#### 22.2.4.2.1. Server Plugin Library and Plugin Descriptors

Every plugin library that contains server plugins must include a library descriptor that contains the general plugin descriptor for each server plugin in the file. This section discusses how to write the library and general descriptors for server plugins.

The library descriptor must define two symbols:

- `_mysql_plugin_interface_version` specifies the version number of the general plugin framework. This is given by the `MYSQL_PLUGIN_INTERFACE_VERSION` symbol, which is defined in the `plugin.h` file.
- `_mysql_plugin_declarations` defines an array of plugin declarations, terminated by a declaration with all members set to 0. Each declaration is an instance of the `st_mysql_plugin` structure (also defined in `plugin.h`). There must be one of these for each server plugin in the library.

If the server does not find those two symbols in a library, it does not accept it as a legal plugin library and rejects it with an error. This prevents use of a library for plugin purposes unless it was built specifically as a plugin library.

The conventional way to define the two required symbols is by using the `mysql_declare_plugin()` and `mysql_declare_plugin_end` macros from the `plugin.h` file:

```
mysql_declare_plugin(name)
... one or more server plugin descriptors here ...
```

```
mysql_declare_plugin_end;
```

Each server plugin must have a general descriptor that provides information to the server plugin API. The general descriptor has the same structure for all plugin types. The `st_mysql_plugin` structure in the `plugin.h` file defines this descriptor:

```
struct st_mysql_plugin
{
    int type; /* the plugin type (a MYSQL_XXX_PLUGIN value) */
    void *info; /* pointer to type-specific plugin descriptor */
    const char *name; /* plugin name */
    const char *author; /* plugin author (for I_S.PLUGINS) */
    const char *descr; /* general descriptive text (for I_S.PLUGINS) */
    int license; /* the plugin license (PLUGIN_LICENSE_XXX) */
    int (*init)(void *); /* the function to invoke when plugin is loaded */
    int (*deinit)(void *); /* the function to invoke when plugin is unloaded */
    unsigned int version; /* plugin version (for I_S.PLUGINS) */
    struct st_mysql_show_var *status_vars;
    struct st_mysql_sys_var **system_vars;
    void * __reserved1; /* reserved for dependency checking */
    unsigned long flags; /* flags for plugin */
};
```

The `st_mysql_plugin` descriptor structure members are used as follows. `char *` members should be specified as null-terminated strings.

- `type`: The plugin type. This must be one of the plugin-type values from `plugin.h`:

```
/*
 * The allowable types of plugins
 */
#define MYSQL_UDF_PLUGIN 0 /* User-defined function */
#define MYSQL_STORAGE_ENGINE_PLUGIN 1 /* Storage Engine */
#define MYSQL_FTPARSER_PLUGIN 2 /* Full-text parser plugin */
#define MYSQL_DAEMON_PLUGIN 3 /* The daemon/raw plugin type */
#define MYSQL_INFORMATION_SCHEMA_PLUGIN 4 /* The I_S plugin type */
#define MYSQL_AUDIT_PLUGIN 5 /* The Audit plugin type */
#define MYSQL_REPLICATION_PLUGIN 6 /* The replication plugin type */
#define MYSQL_AUTHENTICATION_PLUGIN 7 /* The authentication plugin type */
...
```

For example, for a full-text parser plugin, the `type` value is `MYSQL_FTPARSER_PLUGIN`.

- `info`: A pointer to the type-specific descriptor for the plugin. This descriptor's structure depends on the particular type of plugin, unlike that of the general plugin descriptor structure. For version-control purposes, the first member of the type-specific descriptor for every plugin type is expected to be the interface version for the type. This enables the server to check the type-specific version for every plugin no matter its type. Following the version number, the descriptor includes any other members needed, such as callback functions and other information needed by the server to invoke the plugin properly. Later sections on writing particular types of server plugins describe the structure of their type-specific descriptors.
- `name`: A string that gives the plugin name. This is the name that will be listed in the `mysql.plugin` table and by which you refer to the plugin in SQL statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`, or with the `--plugin-load` [435] option. The name is also visible in the `INFORMATION_SCHEMA.PLUGINS` table or the output from `SHOW PLUGINS`.

The plugin name should not begin with the name of any server option. If it does, the server will fail to initialize it. For example, the server has a `--socket` [442] option, so you should not use a plugin name such as `socket`, `socket_plugin`, and so forth.

- `author`: A string naming the plugin author. This can be whatever you like.
- `desc`: A string that provides a general description of the plugin. This can be whatever you like.
- `license`: The plugin license type. The value can be one of `PLUGIN_LICENSE_PROPRIETARY`, `PLUGIN_LICENSE_GPL`, or `PLUGIN_LICENSE_BSD`.

- `init`: A once-only initialization function, or `NULL` if there is no such function. The server executes this function when it loads the plugin, which happens for `INSTALL PLUGIN` or, for plugins listed in the `mysql.plugin` table, at server startup. The function takes one argument that points to the internal structure used to identify the plugin. It returns zero for success and nonzero for failure.
- `deinit`: A once-only deinitialization function, or `NULL` if there is no such function. The server executes this function when it unloads the plugin, which happens for `UNINSTALL PLUGIN` or, for plugins listed in the `mysql.plugin` table, at server shutdown. The function takes one argument that points to the internal structure used to identify the plugin. It returns zero for success and nonzero for failure.
- `version`: The plugin version number. When the plugin is installed, this value can be retrieved from the `INFORMATION_SCHEMA.PLUGINS` table. The value includes major and minor numbers. If you write the value as a hex constant, the format is `0xMMNN`, where `MM` and `NN` are the major and minor numbers, respectively. For example, `0x0302` represents version 3.2.

- `status_vars`: A pointer to a structure for status variables associated with the plugin, or `NULL` if there are no such variables. When the plugin is installed, these variables are displayed in the output of the `SHOW STATUS` statement.

The `status_vars` member, if not `NULL`, points to an array of `st_mysql_show_var` structures that describe status variables. See [Section 22.2.4.2.2, “Server Plugin Status and System Variables”](#).

- `system_vars`: A pointer to a structure for system variables associated with the plugin, or `NULL` if there are no such variables. These options and system variables can be used to help initialize variables within the plugin.

The `system_vars` member, if not `NULL`, points to an array of `st_mysql_sys_var` structures that describe system variables. See [Section 22.2.4.2.2, “Server Plugin Status and System Variables”](#).

- `__reserved1`: A placeholder for the future. Currently, it should be set to `NULL`.
- `flags`: Plugin flags. Individual bits correspond to different flags. The value should be set to the OR of the applicable flags. These flags are available:

```
#define PLUGIN_OPT_NO_INSTALL 1UL /* Not dynamically loadable */
#define PLUGIN_OPT_NO_UNINSTALL 2UL /* Not dynamically unloadable */
```

`PLUGIN_OPT_NO_INSTALL` indicates that the plugin cannot be loaded at runtime with the `INSTALL PLUGIN` statement. This is appropriate for plugins that must be loaded at server startup with the `--plugin-load [435]` option. `PLUGIN_OPT_NO_UNINSTALL` indicates that the plugin cannot be unloaded at runtime with the `UNINSTALL PLUGIN` statement.

This member was added in MySQL 5.6.3.

The server invokes the `init` and `deinit` functions in the general plugin descriptor only when loading and unloading the plugin. They have nothing to do with use of the plugin such as happens when an SQL statement causes the plugin to be invoked.

For example, the descriptor information for a library that contains a single full-text parser plugin named `simple_parser` looks like this:

```
mysql_declare_plugin(ftexample)
{
    MYSQL_FTPARSER_PLUGIN,          /* type */
    &simple_parser_descriptor,       /* descriptor */
    "simple_parser",                 /* name */
    "Oracle Corporation",          /* author */
    "Simple Full-Text Parser",     /* description */
    PLUGIN_LICENSE_GPL,            /* plugin license */
    simple_parser_plugin_init,     /* init function (when loaded) */
    simple_parser_plugin_deinit,   /* deinit function (when unloaded) */
    0x0001,                         /* version */
    simple_status,                  /* status variables */
}
```



```

simple_system_variables, /* system variables */
NULL,
0
}
mysql_declare_plugin_end;

```

For a full-text parser plugin, the type must be `MYSQL_FTPARSER_PLUGIN`. This is the value that identifies the plugin as being legal for use in a `WITH PARSER` clause when creating a `FULLTEXT` index. (No other plugin type is legal for this clause.)

`plugin.h` defines the `mysql_declare_plugin()` and `mysql_declare_plugin_end` macros like this:

```

#ifndef MYSQL_DYNAMIC_PLUGIN
#define __MYSQL_DECLARE_PLUGIN(NAME, VERSION, PSIZE, DECLS) \
MYSQL_PLUGIN_EXPORT int VERSION= MYSQL_PLUGIN_INTERFACE_VERSION; \
MYSQL_PLUGIN_EXPORT int PSIZE= sizeof(struct st_mysql_plugin); \
MYSQL_PLUGIN_EXPORT struct st_mysql_plugin DECLS[]= {
#else
#define __MYSQL_DECLARE_PLUGIN(NAME, VERSION, PSIZE, DECLS) \
MYSQL_PLUGIN_EXPORT int _mysql_plugin_interface_version= MYSQL_PLUGIN_INTERFACE_VERSION; \
MYSQL_PLUGIN_EXPORT int _mysql_sizeof_struct_st_plugin= sizeof(struct st_mysql_plugin); \
MYSQL_PLUGIN_EXPORT struct st_mysql_plugin _mysql_plugin_declarations[]= {
#endif

#define mysql_declare_plugin(NAME) \
__MYSQL_DECLARE_PLUGIN(NAME, \
    builtin_ ## NAME ## _plugin_interface_version, \
    builtin_ ## NAME ## _sizeof_struct_st_plugin, \
    builtin_ ## NAME ## _plugin)

#define mysql_declare_plugin_end ,{0,0,0,0,0,0,0,0,0,0,0,0}

```



#### Note

Those declarations define the `_mysql_plugin_interface_version` symbol only if the `MYSQL_DYNAMIC_PLUGIN` symbol is defined. This means that `-DMYSQL_DYNAMIC_PLUGIN` must be provided as part of the compilation command to build the plugin as a shared library.

When the macros are used as just shown, they expand to the following code, which defines both of the required symbols (`_mysql_plugin_interface_version` and `_mysql_plugin_declarations`):

```

int _mysql_plugin_interface_version= MYSQL_PLUGIN_INTERFACE_VERSION;
int _mysql_sizeof_struct_st_plugin= sizeof(struct st_mysql_plugin);
struct st_mysql_plugin _mysql_plugin_declarations[]= {
{
    MYSQL_FTPARSER_PLUGIN, /* type */
    &simple_parser_descriptor, /* descriptor */
    "simple_parser", /* name */
    "Oracle Corporation", /* author */
    "Simple Full-Text Parser", /* description */
    PLUGIN_LICENSE_GPL, /* plugin license */
    simple_parser_plugin_init, /* init function (when loaded) */
    simple_parser_plugin_deinit, /* deinit function (when unloaded) */
    0x0001, /* version */
    simple_status, /* status variables */
    simple_system_variables, /* system variables */
    NULL,
    0
}
, {0,0,0,0,0,0,0,0,0,0,0,0}
};

```

The preceding example declares a single plugin in the general descriptor, but it is possible to declare multiple plugins. List the declarations one after the other between `mysql_declare_plugin()` and `mysql_declare_plugin_end`, separated by commas.

MySQL server plugins can be written in C or C++ (or another language that can use C calling conventions). If you write a C++ plugin, one C++ feature that you should not use is nonconstant variables to initialize global structures. Members of structures such as the `st_mysql_plugin` structure should be initialized only with constant variables. The `simple_parser` descriptor shown earlier is permissible in a C++ plugin because it satisfies that requirement:

```
mysql_declare_plugin(ftexample)
{
    MYSQL_FTPARSER_PLUGIN,      /* type */
    &simple_parser_descriptor,   /* descriptor */
    "simple_parser",             /* name */
    "Oracle Corporation",       /* author */
    "Simple Full-Text Parser",  /* description */
    PLUGIN_LICENSE_GPL,        /* plugin license */
    simple_parser_plugin_init,  /* init function (when loaded) */
    simple_parser_plugin_deinit, /* deinit function (when unloaded) */
    0x0001,                     /* version */
    simple_status,              /* status variables */
    simple_system_variables,    /* system variables */
    NULL,
    0
}
mysql_declare_plugin_end;
```

Here is another valid way to write the general descriptor. It uses constant variables to indicate the plugin name, author, and description:

```
const char *simple_parser_name = "simple_parser";
const char *simple_parser_author = "Oracle Corporation";
const char *simple_parser_description = "Simple Full-Text Parser";

mysql_declare_plugin(ftexample)
{
    MYSQL_FTPARSER_PLUGIN,      /* type */
    &simple_parser_descriptor,   /* descriptor */
    simple_parser_name,         /* name */
    simple_parser_author,       /* author */
    simple_parser_description,  /* description */
    PLUGIN_LICENSE_GPL,        /* plugin license */
    simple_parser_plugin_init,  /* init function (when loaded) */
    simple_parser_plugin_deinit, /* deinit function (when unloaded) */
    0x0001,                     /* version */
    simple_status,              /* status variables */
    simple_system_variables,    /* system variables */
    NULL,
    0
}
mysql_declare_plugin_end;
```

However, the following general descriptor is invalid. It uses structure members to indicate the plugin name, author, and description, but structures are not considered constant initializers in C++:

```
typedef struct
{
    const char *name;
    const char *author;
    const char *description;
} plugin_info;

plugin_info parser_info = {
    "simple_parser",
    "Oracle Corporation",
    "Simple Full-Text Parser"
};

mysql_declare_plugin(ftexample)
{
    MYSQL_FTPARSER_PLUGIN,      /* type */
    &simple_parser_descriptor,   /* descriptor */
    parser_info.name,           /* name */
    parser_info.author,         /* author */

```



```

parser_info.description, /* description */
PLUGIN_LICENSE_GPL, /* plugin license */
simple_parser_plugin_init, /* init function (when loaded) */
simple_parser_plugin_deinit, /* deinit function (when unloaded) */
0x0001, /* version */
simple_status, /* status variables */
simple_system_variables, /* system variables */
NULL,
0
}
mysql_declare_plugin_end;

```

### 22.2.4.2.2. Server Plugin Status and System Variables

The server plugin interface enables plugins to expose status and system variables using the `status_vars` and `system_vars` members of the general plugin descriptor.

The `status_vars` member of the general plugin descriptor, if not 0, points to an array of `st_mysql_show_var` structures, each of which describes one status variable, followed by a structure with all members set to 0. The `st_mysql_show_var` structure has this definition:

```

struct st_mysql_show_var {
    const char *name;
    char *value;
    enum enum_mysql_show_type type;
};

```

When the plugin is installed, the plugin name and the `name` value are joined with an underscore to form the name displayed by `SHOW STATUS`.

The following table shows the permissible status variable `type` values and what the corresponding variable should be.

**Table 22.1. Server Plugin Status Variable Types**

Variable Type	Meaning
<code>SHOW_BOOL</code>	Pointer to a boolean variable
<code>SHOW_INT</code>	Pointer to an integer variable
<code>SHOW_LONG</code>	Pointer to a long integer variable
<code>SHOW_LONGLONG</code>	Pointer to a longlong integer variable
<code>SHOW_CHAR</code>	A string
<code>SHOW_CHAR_PTR</code>	Pointer to a string
<code>SHOW_ARRAY</code>	Pointer to another <code>st_mysql_show_var</code> array
<code>SHOW_FUNC</code>	Pointer to a function
<code>SHOW_DOUBLE</code>	Pointer to a double

For the `SHOW_FUNC` type, the function is called and fills in its `out` parameter, which then provides information about the variable to be displayed. The function has this signature:

```

#define SHOW_VAR_FUNC_BUFF_SIZE 1024

typedef int (*mysql_show_var_func) (void *thd,
                                   struct st_mysql_show_var *out,
                                   char *buf);

```

The `system_vars` member, if not 0, points to an array of `st_mysql_sys_var` structures, each of which describes one system variable (which can also be set from the command-line or configuration file), followed by a structure with all members set to 0. The `st_mysql_sys_var` structure is defined as follows:

```

struct st_mysql_sys_var {
    int flags;
    const char *name, *comment;
};

```

```
int (*check)(THD*, struct st_mysql_sys_var *, void*, st_mysql_value*);
void (*update)(THD*, struct st_mysql_sys_var *, void*, const void*);
};
```

Additional fields are append as required depending upon the flags.

For convenience, a number of macros are defined that make creating new system variables within a plugin much simpler.

Throughout the macros, the following fields are available:

- **name**: An unquoted identifier for the system variable.
- **varname**: The identifier for the static variable. Where not available, it is the same as the **name** field.
- **opt**: Additional use flags for the system variable. The following table shows the permissible flags.

**Table 22.2. Server Plugin System Variable Flags**

Flag Value	Description
<code>PLUGIN_VAR_READONLY</code>	The system variable is read only
<code>PLUGIN_VAR_NOSYSVAR</code>	The system variable is not user visible at runtime
<code>PLUGIN_VAR_NOCMDOPT</code>	The system variable is not configurable from the command line
<code>PLUGIN_VAR_NOCMDARG</code>	No argument is required at the command line (typically used for boolean variables)
<code>PLUGIN_VAR_RQCMDARG</code>	An argument is required at the command line (this is the default)
<code>PLUGIN_VAR_OPCMDARG</code>	An argument is optional at the command line
<code>PLUGIN_VAR_MEMALLOC</code>	Used for string variables; indicates that memory is to be allocated for storage of the string

- **comment**: A descriptive comment to be displayed in the server help message. `NULL` if this variable is to be hidden.
- **check**: The check function, `NULL` for default.
- **update**: The update function, `NULL` for default.
- **default**: The variable default value.
- **minimum**: The variable minimum value.
- **maximum**: The variable maximum value.
- **blocksize**: The variable block size. When the value is set, it is rounded to the nearest multiple of **blocksize**.

A system variable may be accessed either by using the static variable directly or by using the `SYSVAR()` accessor macro. The `SYSVAR()` macro is provided for completeness. Usually it should be used only when the code cannot directly access the underlying variable.

For example:

```
static int my_foo;
static MYSQL_SYSVAR_INT(foo_var, my_foo,
                        PLUGIN_VAR_RQCMDARG, "foo comment",
                        NULL, NULL, 0, 0, INT_MAX, 0);
...
SYSVAR(foo_var)= value;
value= SYSVAR(foo_var);
my_foo= value;
value= my_foo;
```

Session variables may be accessed only through the `THDVAR()` accessor macro. For example:

```
static MYSQL_THDVAR_BOOL(some_flag,
                        PLUGIN_VAR_NOCMDARG, "flag comment",
                        NULL, NULL, FALSE);
...
if (THDVAR(thd, some_flag))
{
    do_something();
    THDVAR(thd, some_flag)= FALSE;
}
```

All global and session system variables must be published to `mysqld` before use. This is done by constructing a `NULL`-terminated array of the variables and linking to it in the plugin public interface. For example:

```
static struct st_mysql_sys_var *my_plugin_vars[] = {
    MYSQL_SYSVAR(foo_var),
    MYSQL_SYSVAR(some_flag),
    NULL
};
mysql_declare_plugin(fooplug)
{
    MYSQL_..._PLUGIN,
    &plugin_data,
    "fooplug",
    "foo author",
    "This does foo!",
    PLUGIN_LICENSE_GPL,
    foo_init,
    foo_fini,
    0x0001,
    NULL,
    my_plugin_vars,
    NULL,
    0
}
mysql_declare_plugin_end;
```

The following convenience macros enable you to declare different types of system variables:

- Boolean system variables of type `my_bool`, which is a 1-byte boolean. (0 = FALSE, 1 = TRUE)

```
MYSQL_THDVAR_BOOL(name, opt, comment, check, update, default)
MYSQL_SYSVAR_BOOL(name, varname, opt, comment, check, update, default)
```

- String system variables of type `char*`, which is a pointer to a null-terminated string.

```
MYSQL_THDVAR_STR(name, opt, comment, check, update, default)
MYSQL_SYSVAR_STR(name, varname, opt, comment, check, update, default)
```

- Integer system variables, of which there are several varieties.

- An `int` system variable, which is typically a 4-byte signed word.

```
MYSQL_THDVAR_INT(name, opt, comment, check, update, default, min, max, blk)
MYSQL_SYSVAR_INT(name, varname, opt, comment, check, update, default,
                 minimum, maximum, blocksize)
```

- An `unsigned int` system variable, which is typically a 4-byte unsigned word.

```
MYSQL_THDVAR_UINT(name, opt, comment, check, update, default, min, max, blk)
MYSQL_SYSVAR_UINT(name, varname, opt, comment, check, update, default,
                  minimum, maximum, blocksize)
```

- A `long` system variable, which is typically either a 4- or 8-byte signed word.

```
MYSQL_THDVAR_LONG(name, opt, comment, check, update, default, min, max, blk)
MYSQL_SYSVAR_LONG(name, varname, opt, comment, check, update, default,
                  minimum, maximum, blocksize)
```

- An `unsigned long` system variable, which is typically either a 4- or 8-byte unsigned word.

```
MYSQL_THDVAR_ULONG(name, opt, comment, check, update, default, min, max, blk)
MYSQL_SYSVAR_ULONG(name, varname, opt, comment, check, update, default,
                    minimum, maximum, blocksize)
```

- A `long long` system variable, which is typically an 8-byte signed word.

```
MYSQL_THDVAR_LONGLONG(name, opt, comment, check, update,
                      default, minimum, maximum, blocksize)
MYSQL_SYSVAR_LONGLONG(name, varname, opt, comment, check, update,
                      default, minimum, maximum, blocksize)
```

- An `unsigned long long` system variable, which is typically an 8-byte unsigned word.

```
MYSQL_THDVAR_ULONGLONG(name, opt, comment, check, update,
                       default, minimum, maximum, blocksize)
MYSQL_SYSVAR_ULONGLONG(name, varname, opt, comment, check, update,
                       default, minimum, maximum, blocksize)
```

- An `unsigned long` system variable, which is typically either a 4- or 8-byte unsigned word. The range of possible values is an ordinal of the number of elements in the `typelib`, starting from 0.

```
MYSQL_THDVAR_ENUM(name, opt, comment, check, update, default, typelib)
MYSQL_SYSVAR_ENUM(name, varname, opt, comment, check, update,
                  default, typelib)
```

- An `unsigned long long` system variable, which is typically an 8-byte unsigned word. Each bit represents an element in the `typelib`.

```
MYSQL_THDVAR_SET(name, opt, comment, check, update, default, typelib)
MYSQL_SYSVAR_SET(name, varname, opt, comment, check, update,
                 default, typelib)
```

Internally, all mutable and plugin system variables are stored in a `HASH` structure.

Display of the server command-line help text is handled by compiling a `DYNAMIC_ARRAY` of all variables relevant to command-line options, sorting them, and then iterating through them to display each option.

When a command-line option has been handled, it is then removed from the `argv` by the `handle_option()` function (`my_getopt.c`); in effect, it is consumed.

The server processes command-line options during the plugin installation process, immediately after the plugin has been successfully loaded but before the plugin initialization function has been called

Plugins loaded at runtime do not benefit from any configuration options and must have usable defaults. Once they are installed, they are loaded at `mysqld` initialization time and configuration options can be set at the command line or within `my.cnf`.

Plugins should consider the `thd` parameter to be read only.

### 22.2.4.2.3. Client Plugin Descriptors

Each client plugin must have a descriptor that provides information to the client plugin API. The descriptor structure begins with a fixed set of members common to all client plugins, followed by any members specific to the plugin type.

The `st_mysql_client_plugin` structure in the `client_plugin.h` file defines a “generic” descriptor that contains the common members:

```
struct st_mysql_client_plugin
{
    int type;
    unsigned int interface_version;
    const char *name;
    const char *author;
    const char *desc;
```

```

unsigned int version[3];
const char *license;
void *mysql_api;
int (*init)(char *, size_t, int, va_list);
int (*deinit)();
int (*options)(const char *option, const void *);
};

```

The common `st_mysql_client_plugin` descriptor structure members are used as follows. `char *` members should be specified as null-terminated strings.

- `type`: The plugin type. This must be one of the plugin-type values from `client_plugin.h`, such as `MYSQL_CLIENT_AUTHENTICATION_PLUGIN`.
- `interface_version`: The plugin interface version. For example, this is `MYSQL_CLIENT_AUTHENTICATION_PLUGIN_INTERFACE_VERSION` for an authentication plugin.
- `name`: A string that gives the plugin name. This is the name by which you refer to the plugin when you call `mysql_options()` with the `MYSQL_DEFAULT_AUTH` option or specify the `--default-auth` option to a MySQL client program.
- `author`: A string naming the plugin author. This can be whatever you like.
- `desc`: A string that provides a general description of the plugin. This can be whatever you like.
- `version`: The plugin version as an array of three integers indicating the major, minor, and teeny versions. For example, `{1, 2, 3}` indicates version 1.2.3.
- `license`: A string that specifies the license type.
- `mysql_api`: For internal use. Specify it as `NULL` in the plugin descriptor.
- `init`: A once-only initialization function, or `NULL` if there is no such function. The client library executes this function when it loads the plugin. The function returns zero for success and nonzero for failure.

The `init` function uses its first two arguments to return an error message if an error occurs. The first argument is a pointer to a `char` buffer, and the second argument indicates the buffer length. Any message returned by the `init` function must be null-terminated, so the maximum message length is the buffer length minus one. The next arguments are passed to `mysql_load_plugin()`. The first indicates how many more arguments there are (0 if none), followed by any remaining arguments.

- `deinit`: A once-only deinitialization function, or `NULL` if there is no such function. The client library executes this function when it unloads the plugin. The function takes no arguments. It returns zero for success and nonzero for failure.
- `options`: A function for handling options passed to the plugin, or `NULL` if there is no such function. The function takes two arguments representing the option name and a pointer to its value.

For a given client plugin type, the common descriptor members may be followed by additional members necessary to implement plugins of that type. For example, the `st_mysql_client_plugin_AUTHENTICATION` structure for authentication plugins has a function at the end that the client library calls to perform authentication.

To declare a plugin, use the `mysql_declare_client_plugin()` and `mysql_end_client_plugin` macros:

```

mysql_declare_client_plugin(plugin_type)
... members common to all client plugins ...
... type-specific extra members ...
mysql_end_client_plugin;

```

Do not specify the `type` or `interface_version` member explicitly. The `mysql_declare_client_plugin()` macro uses the `plugin_type` argument to generate their values automatically. For example, declare an authentication client plugin like this:

```
mysql_declare_client_plugin(AUTHENTICATION)
    "my_auth_plugin",
    "Author Name",
    "My Client Authentication Plugin",
    {1,0,0},
    "GPL",
    NULL,
    my_auth_init,
    my_auth_deinit,
    my_auth_options,
    my_auth_main
mysql_end_client_plugin;
```

This declaration uses the `AUTHENTICATION` argument to set the `type` and `interface_version` members to `MYSQL_CLIENT_AUTHENTICATION_PLUGIN` and `MYSQL_CLIENT_AUTHENTICATION_PLUGIN_INTERFACE_VERSION`.

Depending on the plugin type, the descriptor may have other members following the common members. For example, for an authentication plugin, there is a function (`my_auth_main()` in the descriptor just shown) that handles communication with the server. See [Section 22.2.4.9, “Writing Authentication Plugins”](#).

Normally, a client program that supports the use of authentication plugins causes a plugin to be loaded by calling `mysql_options()` to set the `MYSQL_DEFAULT_AUTH` and `MYSQL_PLUGIN_DIR` options:

```
char *plugin_dir = "path_to_plugin_dir";
char *default_auth = "plugin_name";

/* ... process command-line options ... */

mysql_options(&mysql, MYSQL_PLUGIN_DIR, plugin_dir);
mysql_options(&mysql, MYSQL_DEFAULT_AUTH, default_auth);
```

Typically, the program will also accept `--plugin-dir` and `--default-auth` options that enable users to override the default values.

Should a client program require lower-level plugin management, the client library contains functions that take an `st_mysql_client_plugin` argument. See [Section 21.8.13, “C API Client Plugin Functions”](#).

### 22.2.4.3. Compiling and Installing Plugin Libraries

After your plugin is written, you must compile it and install it. The procedure for compiling shared objects varies from system to system. If you build your library using `CMake`, it should be able to generate the correct compilation commands for your system. If the library is named `somepluglib`, you should end up with a shared object file that has a name something like `somepluglib.so`. (The file name might have a different suffix on your system.)

To use `CMake`, you'll need to set up the configuration files to enable the plugin to be compiled and installed. Use the plugin examples under the `plugin` directory of a MySQL source distribution as a guide.

Create `CMakeLists.txt`, which should look something like this:

```
MYSQL_ADD_PLUGIN(somepluglib somepluglib.c
    MODULE_ONLY MODULE_OUTPUT_NAME "somepluglib")
```

When `CMake` generates the `Makefile`, it should take care of passing to the compilation command the `-DMYSQL_DYNAMIC_PLUGIN` flag, and passing to the linker the `-lmysqservices` flag, which is needed to link in any functions from services provided through the plugin services interface. See [Section 22.2.5, “MySQL Services for Plugins”](#).

Run `CMake`, then run `make`:

```
shell> cmake .
```

```
shell> make
```

If you need to specify configuration options to CMake, see [Section 2.9.4, “MySQL Source-Configuration Options”](#), for a list. For example, you might want to specify `CMAKE_INSTALL_PREFIX [132]` to indicate the MySQL base directory under which the plugin should be installed. You can see what value to use for this option with `SHOW VARIABLES`:

```
mysql> SHOW VARIABLES LIKE 'basedir';
+-----+-----+
| Variable_name | Value                |
+-----+-----+
| base          | /usr/local/mysql    |
+-----+-----+
```

The location of the plugin directory where you should install the library is given by the `plugin_dir [525]` system variable. For example:

```
mysql> SHOW VARIABLES LIKE 'plugin_dir';
+-----+-----+
| Variable_name | Value                |
+-----+-----+
| plugin_dir    | /usr/local/mysql/lib/mysql/plugin |
+-----+-----+
```

To install the plugin library, use `make`:

```
shell> make install
```

Verify that `make install` installed the plugin library in the proper directory. After installing it, make sure that the library permissions permit it to be executed by the server.

#### 22.2.4.4. Writing Full-Text Parser Plugins

A full-text parser server plugin can be used to replace or modify the built-in full-text parser. This section describes how to write a full-text parser plugin named `simple_parser`. This plugin performs parsing based on simpler rules than those used by the MySQL built-in full-text parser: Words are nonempty runs of whitespace characters.

The instructions use the source code in the `plugin/fulltext` directory of MySQL source distributions, so change location into that directory. The following procedure describes how the plugin library is created:

1. To write a full-text parser plugin, include the following header file in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <mysql/plugin.h>
```

`plugin.h` defines the `MYSQL_FTPARSER_PLUGIN` server plugin type and the data structures needed to declare the plugin.

2. Set up the library descriptor for the plugin library file.

This descriptor contains the general plugin descriptor for the server plugin. For a full-text parser plugin, the type must be `MYSQL_FTPARSER_PLUGIN`. This is the value that identifies the plugin as being legal for use in a `WITH PARSER` clause when creating a `FULLTEXT` index. (No other plugin type is legal for this clause.)

For example, the library descriptor for a library that contains a single full-text parser plugin named `simple_parser` looks like this:

```
mysql_declare_plugin(ftexample)
{
  MYSQL_FTPARSER_PLUGIN,      /* type                */
  &simple_parser_descriptor,   /* descriptor          */
  "simple_parser",             /* name                */
  "Oracle Corporation",      /* author              */
}
```



```

"Simple Full-Text Parser", /* description */
PLUGIN_LICENSE_GPL, /* plugin license */
simple_parser_plugin_init, /* init function (when loaded) */
simple_parser_plugin_deinit, /* deinit function (when unloaded) */
0x0001, /* version */
simple_status, /* status variables */
simple_system_variables, /* system variables */
NULL,
0
}
mysql_declare_plugin_end;

```

The `name` member (`simple_parser`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `SHOW PLUGINS` or `INFORMATION_SCHEMA.PLUGINS`.

For more information, see [Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”](#).

### 3. Set up the type-specific plugin descriptor.

Each general plugin descriptor in the library descriptor points to a type-specific descriptor. For a full-text parser plugin, the type-specific descriptor is an instance of the `st_mysql_ftparser` structure in the `plugin.h` file:

```

struct st_mysql_ftparser
{
    int interface_version;
    int (*parse)(MYSQL_FTPARSER_PARAM *param);
    int (*init)(MYSQL_FTPARSER_PARAM *param);
    int (*deinit)(MYSQL_FTPARSER_PARAM *param);
};

```

As shown by the structure definition, the descriptor has an interface version number and contains pointers to three functions. The version is specified using a symbol of the form `MYSQL_xxx_INTERFACE_VERSION` (such as `MYSQL_FTPARSER_INTERFACE_VERSION` for full-text parser plugins). The `init` and `deinit` members should point to a function or be set to 0 if the function is not needed. The `parse` member must point to the function that performs the parsing.

In the `simple_parser` declaration, that descriptor is indicated by `&simple_parser_descriptor`. The descriptor specifies the version number for the full-text plugin interface (as given by `MYSQL_FTPARSER_INTERFACE_VERSION`), and the plugin's parsing, initialization, and deinitialization functions:

```

static struct st_mysql_ftparser simple_parser_descriptor=
{
    MYSQL_FTPARSER_INTERFACE_VERSION, /* interface version */
    simple_parser_parse, /* parsing function */
    simple_parser_init, /* parser init function */
    simple_parser_deinit /* parser deinit function */
};

```

A full-text parser plugin is used in two different contexts, indexing and searching. In both contexts, the server calls the initialization and deinitialization functions at the beginning and end of processing each SQL statement that causes the plugin to be invoked. However, during statement processing, the server calls the main parsing function in context-specific fashion:

- For indexing, the server calls the parser for each column value to be indexed.
- For searching, the server calls the parser to parse the search string. The parser might also be called for rows processed by the statement. In natural language mode, there is no need for the server to call the parser. For boolean mode phrase searches or natural language searches with query expansion, the parser is used to parse column values for information that is not in the index. Also, if a boolean mode search is done for a column that has no `FULLTEXT` index, the built-in parser will be called. (Plugins are associated with specific indexes. If there is no index, no plugin is used.)



The plugin declaration in the general plugin descriptor has `init` and `deinit` members that point initialization and deinitialization functions, and so does the type-specific plugin descriptor to which it points. However, these pairs of functions have different purposes and are invoked for different reasons:

- For the plugin declaration in the general plugin descriptor, the initialization and deinitialization functions are invoked when the plugin is loaded and unloaded.
- For the type-specific plugin descriptor, the initialization and deinitialization functions are invoked per SQL statement for which the plugin is used.

Each interface function named in the plugin descriptor should return zero for success or nonzero for failure, and each of them receives an argument that points to a `MYSQL_FTPARSER_PARAM` structure containing the parsing context. The structure has this definition:

```
typedef struct st_mysql_ftparser_param
{
    int (*mysql_parse)(struct st_mysql_ftparser_param *,
                      char *doc, int doc_len);
    int (*mysql_add_word)(struct st_mysql_ftparser_param *,
                          char *word, int word_len,
                          MYSQL_FTPARSER_BOOLEAN_INFO *boolean_info);
    void *ftparser_state;
    void *mysql_ftparam;
    struct charset_info_st *cs;
    char *doc;
    int length;
    int flags;
    enum enum_ftparser_mode mode;
} MYSQL_FTPARSER_PARAM;
```

The structure members are used as follows:

- `mysql_parse`: A pointer to a callback function that invokes the server's built-in parser. Use this callback when the plugin acts as a front end to the built-in parser. That is, when the plugin parsing function is called, it should process the input to extract the text and pass the text to the `mysql_parse` callback.

The first parameter for this callback function should be the `param` value itself:

```
param->mysql_parse(param, ...);
```

A front end plugin can extract text and pass it all at once to the built-in parser, or it can extract and pass text to the built-in parser a piece at a time. However, in this case, the built-in parser treats the pieces of text as though there are implicit word breaks between them.

- `mysql_add_word`: A pointer to a callback function that adds a word to a full-text index or to the list of search terms. Use this callback when the parser plugin replaces the built-in parser. That is, when the plugin parsing function is called, it should parse the input into words and invoke the `mysql_add_word` callback for each word.

The first parameter for this callback function should be the `param` value itself:

```
param->mysql_add_word(param, ...);
```

- `ftparser_state`: This is a generic pointer. The plugin can set it to point to information to be used internally for its own purposes.
- `mysql_ftparam`: This is set by the server. It is passed as the first argument to the `mysql_parse` or `mysql_add_word` callback.
- `cs`: A pointer to information about the character set of the text, or 0 if no information is available.

- `doc`: A pointer to the text to be parsed.
- `length`: The length of the text to be parsed, in bytes.
- `flags`: Parser flags. This is zero if there are no special flags. Currently, the only nonzero flag is `MYSQL_FTFLAGS_NEED_COPY`, which means that `mysql_add_word()` must save a copy of the word (that is, it cannot use a pointer to the word because the word is in a buffer that will be overwritten.) This member was added in MySQL 5.1.12.

This flag might be set or reset by MySQL before calling the parser plugin, by the parser plugin itself, or by the `mysql_parse()` function.

- `mode`: The parsing mode. This value will be one of the following constants:
  - `MYSQL_FTPARSER_SIMPLE_MODE`: Parse in fast and simple mode, which is used for indexing and for natural language queries. The parser should pass to the server only those words that should be indexed. If the parser uses length limits or a stopword list to determine which words to ignore, it should not pass such words to the server.
  - `MYSQL_FTPARSER_WITH_STOPWORDS`: Parse in stopword mode. This is used in boolean searches for phrase matching. The parser should pass all words to the server, even stopwords or words that are outside any normal length limits.
  - `MYSQL_FTPARSER_FULL_BOOLEAN_INFO`: Parse in boolean mode. This is used for parsing boolean query strings. The parser should recognize not only words but also boolean-mode operators and pass them to the server as tokens using the `mysql_add_word` callback. To tell the server what kind of token is being passed, the plugin needs to fill in a `MYSQL_FTPARSER_BOOLEAN_INFO` structure and pass a pointer to it.

If the parser is called in boolean mode, the `param->mode` value will be `MYSQL_FTPARSER_FULL_BOOLEAN_INFO`. The `MYSQL_FTPARSER_BOOLEAN_INFO` structure that the parser uses for passing token information to the server looks like this:

```
typedef struct st_mysql_ftparser_boolean_info
{
    enum enum_ft_token_type type;
    int yesno;
    int weight_adjust;
    bool wasign;
    bool trunc;
    /* These are parser state and must be removed. */
    byte prev;
    byte *quot;
} MYSQL_FTPARSER_BOOLEAN_INFO;
```

The parser should fill in the structure members as follows:

- `type`: The token type. The following table shows the permissible types.

**Table 22.3. Full-Text Parser Token Types**

Token Value	Meaning
<code>FT_TOKEN_EOF</code>	End of data
<code>FT_TOKEN_WORD</code>	A regular word
<code>FT_TOKEN_LEFT_PAREN</code>	The beginning of a group or subexpression
<code>FT_TOKEN_RIGHT_PAREN</code>	The end of a group or subexpression
<code>FT_TOKEN_STOPWORD</code>	A stopword

- `yesno`: Whether the word must be present for a match to occur. 0 means that the word is optional but increases the match relevance if it is present. Values larger than 0 mean that the word must be present. Values smaller than 0 mean that the word must not be present.
- `weight_adjust`: A weighting factor that determines how much a match for the word counts. It can be used to increase or decrease the word's importance in relevance calculations. A value of zero indicates no weight adjustment. Values greater than or less than zero mean higher or lower weight, respectively. The examples at [Section 12.9.2, "Boolean Full-Text Searches"](#), that use the `<` and `>` operators illustrate how weighting works.
- `wasign`: The sign of the weighting factor. A negative value acts like the `~` boolean-search operator, which causes the word's contribution to the relevance to be negative.
- `trunc`: Whether matching should be done as if the boolean-mode `*` truncation operator had been given.

Plugins should not use the `prev` and `quot` members of the `MYSQL_FTPARSER_BOOLEAN_INFO` structure.

#### 4. Set up the plugin interface functions.

The general plugin descriptor in the library descriptor names the initialization and deinitialization functions that the server should invoke when it loads and unloads the plugin. For `simple_parser`, these functions do nothing but return zero to indicate that they succeeded:

```
static int simple_parser_plugin_init(void *arg __attribute__((unused)))
{
    return(0);
}

static int simple_parser_plugin_deinit(void *arg __attribute__((unused)))
{
    return(0);
}
```

Because those functions do not actually do anything, you could omit them and specify 0 for each of them in the plugin declaration.

The type-specific plugin descriptor for `simple_parser` names the initialization, deinitialization, and parsing functions that the server invokes when the plugin is used. For `simple_parser`, the initialization and deinitialization functions do nothing:

```
static int simple_parser_init(MYSQL_FTPARSER_PARAM *param
                             __attribute__((unused)))
{
    return(0);
}

static int simple_parser_deinit(MYSQL_FTPARSER_PARAM *param
                                __attribute__((unused)))
{
    return(0);
}
```

Here too, because those functions do nothing, you could omit them and specify 0 for each of them in the plugin descriptor.

The main parsing function, `simple_parser_parse()`, acts as a replacement for the built-in full-text parser, so it needs to split text into words and pass each word to the server. The parsing function's first argument is a pointer to a structure that contains the parsing context. This structure has a `doc` member that points to the text to be parsed, and a `length` member that indicates how long the text is. The simple parsing done by the plugin considers nonempty runs of whitespace characters to be words, so it identifies words like this:

```
static int simple_parser_parse(MYSQL_FTPARSER_PARAM *param)
{
    char *end, *start, *docend= param->doc + param->length;

    for (end= start= param->doc;; end++)
    {
        if (end == docend)
        {
            if (end > start)
                add_word(param, start, end - start);
            break;
        }
        else if (isspace(*end))
        {
            if (end > start)
                add_word(param, start, end - start);
            start= end + 1;
        }
    }
    return(0);
}
```

As the parser finds each word, it invokes a function `add_word()` to pass the word to the server. `add_word()` is a helper function only; it is not part of the plugin interface. The parser passes the parsing context pointer to `add_word()`, as well as a pointer to the word and a length value:

```
static void add_word(MYSQL_FTPARSER_PARAM *param, char *word, size_t len)
{
    MYSQL_FTPARSER_BOOLEAN_INFO bool_info=
        { FT_TOKEN_WORD, 0, 0, 0, 0, ' ', 0 };

    param->mysql_add_word(param, word, len, &bool_info);
}
```

For boolean-mode parsing, `add_word()` fills in the members of the `bool_info` structure as described earlier in the discussion of the `st_mysql_ftparser_boolean_info` structure.

5. Set up the status variables. For the `simple_parser` plugin, the following status variable array sets up one status variable with a value that is static text, and another with a value that is stored in a long integer variable:

```
long number_of_calls= 0;

struct st_mysql_show_var simple_status[]=
{
    {"static", (char *)"just a static text", SHOW_CHAR},
    {"called", (char *)&number_of_calls, SHOW_LONG},
    {0,0,0}
};
```

When the plugin is installed, the plugin name and the `name` value are joined with an underscore to form the name displayed by `SHOW STATUS`. For the array just shown, the resulting status variable names are `simple_parser_static` and `simple_parser_called`. This convention means that you can easily display the variables for a plugin using its name:

```
mysql> SHOW STATUS LIKE 'simple_parser%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| simple_parser_static | just a static text |
| simple_parser_called | 0 |
+-----+-----+
```

6. To compile and install a plugin library object file, use the instructions in [Section 22.2.4.3, "Compiling and Installing Plugin Libraries"](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable). For the `simple_parser` plugin, it is compiled and installed when you build MySQL from source. It is also included in binary

distributions. The build process produces a shared object library with a name of `mypluglib.so` (the suffix might differ depending on your platform).

- To use the plugin, register it with the server. For example, to register the plugin at runtime, use this statement (changing the suffix as necessary):

```
mysql> INSTALL PLUGIN simple_parser SONAME 'mypluglib.so';
```

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

- To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement.
- Test the plugin to verify that it works properly.

Create a table that contains a string column and associate the parser plugin with a `FULLTEXT` index on the column:

```
mysql> CREATE TABLE t (c VARCHAR(255),
-> FULLTEXT (c) WITH PARSER simple_parser
-> ) ENGINE=MyISAM;
Query OK, 0 rows affected (0.01 sec)
```

Insert some text into the table and try some searches. These should verify that the parser plugin treats all nonwhitespace characters as word characters:

```
mysql> INSERT INTO t VALUES
-> ('latin1_general_cs is a case-sensitive collation'),
-> ('I\'d like a case of oranges'),
-> ('this is sensitive information'),
-> ('another row'),
-> ('yet another row');
Query OK, 5 rows affected (0.02 sec)
Records: 5 Duplicates: 0 Warnings: 0
```

```
mysql> SELECT c FROM t;
+-----+
| c |
+-----+
| latin1_general_cs is a case-sensitive collation |
| I'd like a case of oranges |
| this is sensitive information |
| another row |
| yet another row |
+-----+
5 rows in set (0.00 sec)
```

```
mysql> SELECT MATCH(c) AGAINST('case') FROM t;
+-----+
| MATCH(c) AGAINST('case') |
+-----+
| 0 |
| 1.2968142032623 |
| 0 |
| 0 |
| 0 |
+-----+
5 rows in set (0.00 sec)
```

```
mysql> SELECT MATCH(c) AGAINST('sensitive') FROM t;
+-----+
| MATCH(c) AGAINST('sensitive') |
+-----+
| 0 |
| 0 |
| 1.3253291845322 |
| 0 |
| 0 |
+-----+
```

```

+-----+
5 rows in set (0.01 sec)

mysql> SELECT MATCH(c) AGAINST('case-sensitive') FROM t;
+-----+
| MATCH(c) AGAINST('case-sensitive') |
+-----+
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
+-----+
5 rows in set (0.01 sec)

mysql> SELECT MATCH(c) AGAINST('I\'d') FROM t;
+-----+
| MATCH(c) AGAINST('I\'d') |
+-----+
|                             |
|                             |
|                             |
|                             |
|                             |
+-----+
5 rows in set (0.01 sec)

```

Note how neither “case” nor “insensitive” match “case-insensitive” the way that they would for the built-in parser.

### 22.2.4.5. Writing Daemon Plugins

A daemon plugin is a simple type of plugin used for code that should be run by the server but that does not communicate with it. This section describes how to write a daemon server plugin, using the example plugin found in the `plugin/daemon_example` directory of MySQL source distributions. That directory contains the `daemon_example.cc` source file for a daemon plugin named `daemon_example` that writes a heartbeat string at regular intervals to a file named `mysql-heartbeat.log` in the data directory.

To write a daemon plugin, include the following header file in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <mysql/plugin.h>
```

`plugin.h` defines the `MYSQL_DAEMON_PLUGIN` server plugin type and the data structures needed to declare the plugin.

The `daemon_example.cc` file sets up the library descriptor as follows. The library descriptor includes a single general server plugin descriptor.

```

mysql_declare_plugin(daemon_example)
{
    MYSQL_DAEMON_PLUGIN,
    &daemon_example_plugin,
    "daemon_example",
    "Brian Aker",
    "Daemon example, creates a heartbeat beat file in mysql-heartbeat.log",
    PLUGIN_LICENSE_GPL,
    daemon_example_plugin_init, /* Plugin Init */
    daemon_example_plugin_deinit, /* Plugin Deinit */
    0x0100 /* 1.0 */,
    NULL, /* status variables */
    NULL, /* system variables */
    NULL, /* config options */
    0, /* flags */
}
mysql_declare_plugin_end;

```

The `name` member (`daemon_example`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `SHOW PLUGINS` or `INFORMATION_SCHEMA.PLUGINS`.

The second member of the plugin descriptor, `daemon_example_plugin`, points to the type-specific daemon plugin descriptor. This structure consists only of the type-specific API version number:

```
struct st_mysql_daemon daemon_example_plugin=
{ MYSQL_DAEMON_INTERFACE_VERSION  };
```

The type-specific structure has no interface functions. There is no communication between the server and the plugin, except that the server calls the initialization and deinitialization functions from the general plugin descriptor to start and stop the plugin:

- `daemon_example_plugin_init()` opens the heartbeat file and spawns a thread that wakes up periodically and writes the next message to the file.
- `daemon_example_plugin_deinit()` closes the file and performs other cleanup.

To compile and install a plugin library object file, use the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable). For the `daemon_example` plugin, it is compiled and installed when you build MySQL from source. It is also included in binary distributions. The build process produces a shared object library with a name of `libdaemon_example.so` (the suffix might differ depending on your platform).

To use the plugin, register it with the server. For example, to register the plugin at runtime, use this statement (change the suffix as necessary):

```
mysql> INSTALL PLUGIN daemon_example SONAME 'libdaemon_example.so';
```

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement.

While the plugin is loaded, it writes a heartbeat string at regular intervals to a file named `mysql-heartbeat.log` in the data directory. This file grows without limit, so after you have satisfied yourself that the plugin operates correctly, unload it:

```
mysql> UNINSTALL PLUGIN daemon_example;
```

## 22.2.4.6. Writing `INFORMATION_SCHEMA` Plugins

This section describes how to write an `INFORMATION_SCHEMA` table server plugin. For example code that implements such plugins, see the `sql/sql_show.cc` file of a MySQL source distribution. You can also look at the example plugins found in the InnoDB source. See the `handler/i_s.cc` and `handler/ha_innodb.cc` files within the InnoDB source tree (in the `storage/innobase` directory).

To write an `INFORMATION_SCHEMA` table plugin, include the following header files in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <sql_class.h>
#include <table.h>
```

These header files are located in the `sql` directory of MySQL source distributions. They contain C++ structures, so the source file for an `INFORMATION_SCHEMA` plugin must be compiled as C++ (not C) code.

The source file for the example plugin developed here is named `simple_i_s_table.cc`. It creates a simple `INFORMATION_SCHEMA` table named `SIMPLE_I_S_TABLE` that has two columns named `NAME` and `VALUE`. The general descriptor for a plugin library that implements the table looks like this:

```
mysql_declare_plugin(simple_i_s_library)
{
    MYSQL_INFORMATION_SCHEMA_PLUGIN,
    &simple_table_info,           /* type-specific descriptor */
    "SIMPLE_I_S_TABLE",         /* table name */
    "Author Name",             /* author */
    "Simple INFORMATION_SCHEMA table", /* description */
    PLUGIN_LICENSE_GPL,        /* license type */
    simple_table_init,         /* init function */
    NULL,
    0x0100,                    /* version = 1.0 */
    NULL,                      /* no status variables */
    NULL,                      /* no system variables */
    NULL,                      /* no reserved information */
    0                           /* no flags */
}
mysql_declare_plugin_end;
```

The `name` member (`SIMPLE_I_S_TABLE`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `SHOW PLUGINS` or `INFORMATION_SCHEMA.PLUGINS`.

The `simple_table_info` member of the general descriptor points to the type-specific descriptor, which consists only of the type-specific API version number:

```
static struct st_mysql_information_schema simple_table_info =
{ MYSQL_INFORMATION_SCHEMA_INTERFACE_VERSION };
```

The general descriptor points to the initialization and deinitialization functions:

- The initialization function provides information about the table structure and a function that populates the table.
- The deinitialization function performs any required cleanup. If no cleanup is needed, this descriptor member can be `NULL` (as in the example shown).

The initialization function should return 0 for success, 1 if an error occurs. The function receives a generic pointer, which it should interpret as a pointer to the table structure:

```
static int table_init(void *ptr)
{
    ST_SCHEMA_TABLE *schema_table= (ST_SCHEMA_TABLE*)ptr;

    schema_table->fields_info= simple_table_fields;
    schema_table->fill_table= simple_fill_table;
    return 0;
}
```

The function should set these two members of the table structure:

- `fields_info`: An array of `ST_FIELD_INFO` structures that contain information about each column.
- `fill_table`: A function that populates the table.

The array pointed to by `fields_info` should contain one element per column of the `INFORMATION_SCHEMA` plus a terminating element. The following `simple_table_fields` array for the example plugin indicates that `SIMPLE_I_S_TABLE` has two columns. `NAME` is string-valued with a length of 10 and `VALUE` is integer-valued with a display width of 20. The last structure marks the end of the array.

```
static ST_FIELD_INFO simple_table_fields[]=
{
    {"NAME", 10, MYSQL_TYPE_STRING, 0, 0 0, 0},
    {"VALUE", 6, MYSQL_TYPE_LONG, 0, MY_I_S_UNSIGNED, 0, 0},
    {0, 0, MYSQL_TYPE_NULL, 0, 0, 0, 0}
};
```



For more information about the column information structure, see the definition of `ST_FIELD_INFO` in the `table.h` header file. The permissible `MYSQL_TYPE_XXX` type values are those used in the C API; see [Section 21.8.4, “C API Data Structures”](#).

The `fill_table` member should be set to a function that populates the table and returns 0 for success, 1 if an error occurs. For the example plugin, the `simple_fill_table()` function looks like this:

```
static int simple_fill_table(THD *thd, TABLE_LIST *tables, Item *cond)
{
    TABLE *table= tables->table;

    table->field[0]->store("Name 1", 6, system_charset_info);
    table->field[1]->store(1);
    if (schema_table_store_record(thd, table))
        return 1;
    table->field[0]->store("Name 2", 6, system_charset_info);
    table->field[1]->store(2);
    if (schema_table_store_record(thd, table))
        return 1;
    return 0;
}
```

For each row of the `INFORMATION_SCHEMA` table, this function initializes each column, then calls `schema_table_store_record()` to install the row. The `store()` method arguments depend on the type of value to be stored. For column 0 (`NAME`, a string), `store()` takes a pointer to a string, its length, and information about the character set of the string:

```
store(const char *to, uint length, CHARSET_INFO *cs);
```

For column 1 (`VALUE`, an integer), `store()` takes the value and a flag indicating whether it is unsigned:

```
store(longlong nr, bool unsigned_value);
```

For other examples of how to populate `INFORMATION_SCHEMA` tables, search for instances of `schema_table_store_record()` in `sql_show.cc`.

To compile and install a plugin library object file, see the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable).

To test the plugin, install it:

```
mysql> INSTALL PLUGIN SIMPLE_I_S_TABLE SONAME 'simple_i_s_table.so';
```

Verify that the table is present:

```
mysql> SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES
-> WHERE TABLE_NAME = 'SIMPLE_I_S_TABLE';
+-----+
| TABLE_NAME |
+-----+
| SIMPLE_I_S_TABLE |
+-----+
```

Try to select from it:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.SIMPLE_I_S_TABLE;
+-----+-----+
| NAME | VALUE |
+-----+-----+
| Name 1 | 1 |
| Name 2 | 2 |
+-----+-----+
```

Uninstall it:

```
mysql> UNINSTALL PLUGIN SIMPLE_I_S_TABLE;
```

### 22.2.4.7. Writing Semisynchronous Replication Plugins

This section describes how to write semisynchronous replication server plugins, using the example plugins found in the `plugin/semisync` directory of MySQL source distributions. That directory contains the source files for master and slave plugins named `rpl_semi_sync_master` and `rpl_semi_sync_slave`. The information here covers only how to set up the plugin framework. For details about how the plugins implement replication functions, see the source.

To write a semisynchronous replication plugin, include the following header file in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <mysql/plugin.h>
```

`plugin.h` defines the `MYSQL_REPLICATION_PLUGIN` server plugin type and the data structures needed to declare the plugin.

For the master side, `semisync_master_plugin.cc` contains this general descriptor for a plugin named `rpl_semi_sync_master`:

```
mysql_declare_plugin(semi_sync_master)
{
    MYSQL_REPLICATION_PLUGIN,
    &semi_sync_master_plugin,
    "rpl_semi_sync_master",
    "He Zhenxing",
    "Semi-synchronous replication master",
    PLUGIN_LICENSE_GPL,
    semi_sync_master_plugin_init, /* Plugin Init */
    semi_sync_master_plugin_deinit, /* Plugin Deinit */
    0x0100 /* 1.0 */,
    semi_sync_master_status_vars, /* status variables */
    semi_sync_master_system_vars, /* system variables */
    NULL, /* config options */
    0, /* flags */
}
mysql_declare_plugin_end;
```

For the slave side, `semisync_slave_plugin.cc` contains this general descriptor for a plugin named `rpl_semi_sync_slave`:

```
mysql_declare_plugin(semi_sync_slave)
{
    MYSQL_REPLICATION_PLUGIN,
    &semi_sync_slave_plugin,
    "rpl_semi_sync_slave",
    "He Zhenxing",
    "Semi-synchronous replication slave",
    PLUGIN_LICENSE_GPL,
    semi_sync_slave_plugin_init, /* Plugin Init */
    semi_sync_slave_plugin_deinit, /* Plugin Deinit */
    0x0100 /* 1.0 */,
    semi_sync_slave_status_vars, /* status variables */
    semi_sync_slave_system_vars, /* system variables */
    NULL, /* config options */
    0, /* flags */
}
mysql_declare_plugin_end;
```

For both the master and slave plugins, the general descriptor has pointers to the type-specific descriptor, the initialization and deinitialization functions, and to the status and system variables implemented by the plugin. For information about variable setup, see [Section 22.2.4.2.2, “Server Plugin Status and System Variables”](#). The following remarks discuss the type-specific descriptor and the initialization and deinitialization functions for the master plugin but apply similarly to the slave plugin.

The `semi_sync_master_plugin` member of the master general descriptor points to the type-specific descriptor, which consists only of the type-specific API version number:

```
struct Mysql_replication semi_sync_master_plugin= {
    MYSQL_REPLICATION_INTERFACE_VERSION
};
```

The initialization and deinitialization function declarations look like this:

```
static int semi_sync_master_plugin_init(void *p);
static int semi_sync_master_plugin_deinit(void *p);
```

The initialization function uses the pointer to register transaction and binary logging “observers” with the server. After successful initialization, the server takes care of invoking the observers at the appropriate times. (For details on the observers, see the source files.) The deinitialization function cleans up by deregistering the observers. Each function returns 0 for success or 1 if an error occurs.

To compile and install a plugin library object file, use the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library files, they must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable). For the `rpl_semi_sync_master` and `rpl_semi_sync_slave` plugins, they are compiled and installed when you build MySQL from source. They are also included in binary distributions. The build process produces shared object libraries with names of `semisync_master.so` and `semisync_slave.so` (the suffix might differ depending on your platform).

### 22.2.4.8. Writing Audit Plugins

This section describes how to write an audit server plugin, using the example plugin found in the `plugin/audit_null` directory of MySQL source distributions. The `audit_null.c` source file in that directory implements a simple example audit plugin named `NULL_AUDIT`.

Within the server, the pluggable audit interface is implemented in the `sql_audit.h` and `sql_audit.cc` files in the `sql` directory of MySQL source distributions. Additionally, several places in the server are modified to call the audit interface when an auditable event occurs, so that registered audit plugins can be notified about the event if necessary. To see where such calls occur, look for invocations of functions with names of the form `mysql_audit_xxx()`. Audit notification occurs for server operations such as these:

- Writing a message to the general query log (if the log is enabled)
- Writing a message to the error log
- Sending a query result to a client
- Client connect and disconnect events

To write an audit plugin, include the following header file in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <mysql/plugin_audit.h>
```

`plugin_audit.h` includes `plugin.h`, so you need not include the latter file explicitly. `plugin.h` defines the `MYSQL_AUDIT_PLUGIN` server plugin type and the data structures needed to declare the plugin. `plugin_audit.h` defines data structures specific to audit plugins.

An audit plugin, like any MySQL server plugin, has a general plugin descriptor (see [Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”](#)). In `audit_null.c`, the general descriptor looks like this:

```
mysql_declare_plugin(audit_null)
{
    MYSQL_AUDIT_PLUGIN,          /* type */
    &audit_null_descriptor,      /* descriptor */
    "NULL_AUDIT",               /* name */
    "Oracle Corp",              /* author */
    "Simple NULL Audit",        /* description */
};
```

```

PLUGIN_LICENSE_GPL,
audit_null_plugin_init,      /* init function (when loaded) */
audit_null_plugin_deinit,   /* deinit function (when unloaded) */
0x0003,                      /* version */
simple_status,               /* status variables */
NULL,                       /* system variables */
NULL,
0,
}
mysql_declare_plugin_end;

```

The `name` member (`NULL_AUDIT`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `INFORMATION_SCHEMA.PLUGINS` or `SHOW PLUGINS`.

The general descriptor also refers to `simple_status`, a structure that exposes several status variables to the `SHOW STATUS` statement:

```

static struct st_mysql_show_var simple_status[]=
{
  { "Audit_null_called",
    (char *) &number_of_calls,
    SHOW_INT },
  { "Audit_null_general_log",
    (char *) &number_of_calls_general_log,
    SHOW_INT },
  { "Audit_null_general_error",
    (char *) &number_of_calls_general_error,
    SHOW_INT },
  { "Audit_null_general_result",
    (char *) &number_of_calls_general_result,
    SHOW_INT },
  { "Audit_null_general_status",
    (char *) &number_of_calls_general_status,
    SHOW_INT },
  { "Audit_null_connection_connect",
    (char *) &number_of_calls_connection_connect,
    SHOW_INT },
  { "Audit_null_connection_disconnect",
    (char *) &number_of_calls_connection_disconnect,
    SHOW_INT },
  { "Audit_null_connection_change_user",
    (char *) &number_of_calls_connection_change_user,
    SHOW_INT },
  { 0, 0, 0 }
};

```

The `audit_null_plugin_init` initialization function sets the status variables to zero when the plugin is loaded. The `audit_null_plugin_deinit` function performs cleanup with the plugin is unloaded. During operation, the plugin increments the first status variable for each notification it receives. It also increments the others according to the event class and subclass. In effect, the first variable is the aggregate of the counts for the event subclasses.

The `audit_null_descriptor` value in the general descriptor points to the type-specific descriptor. For audit plugins, this descriptor has the following structure:

```

struct st_mysql_audit
{
  int interface_version;
  void (*release_thd)(MYSQL_THD);
  void (*event_notify)(MYSQL_THD, unsigned int, const void *);
  unsigned long class_mask[MYSQL_AUDIT_CLASS_MASK_SIZE];
};

```

The type-specific descriptor has these members:

- `interface_version`: By convention, type-specific plugin descriptors begin with the interface version for the given plugin type. The server checks `interface_version` when it loads

the plugin to see whether the plugin is compatible with it. For audit plugins, the value of the `interface_version` member is `MYSQL_AUDIT_INTERFACE_VERSION` (defined in `plugin_audit.h`).

- `release_thd`: A function that the server calls to inform the plugin that it is being dissociated from its thread context. This should be `NULL` if there is no such function.
- `event_notify`: A function that the server calls to notify the plugin that an auditable event has occurred. This function should not be `NULL`; that would not make sense because no auditing would occur.
- `class_mask`: A bit mask that indicates the event classes for which the plugin wants to receive notification. If this value is 0, the server passes no events to the plugin.

The server uses the `event_notify` and `release_thd` functions together. They are called within the context of a specific thread, and a thread might perform an activity that produces several event notifications. The first time the server calls `event_notify` for a thread, it creates a binding of the plugin to the thread. The plugin cannot be uninstalled while this binding exists. When no more events for the thread will occur, the server informs the plugin of this by calling the `release_thd` function, and then destroys the binding. For example, when a client issues a statement, the thread processing the statement might notify audit plugins about the result set produced by the statement and about the statement being logged. After these notifications occur, the server releases the plugin before putting the thread to sleep until the client issues another statement.

This design enables the plugin to allocate resources needed for a given thread in the first call to the `event_notify` function and release them in the `release_thd` function:

```
event_notify function:
  if memory is needed to service the thread
    allocate memory
    ... rest of notification processing ...

release_thd function:
  if memory was allocated
    release memory
    ... rest of release processing ...
```

That is more efficient than allocating and releasing memory repeatedly in the notification function.

For the `NULL_AUDIT` example audit plugin, the type-specific descriptor looks like this:

```
static struct st_mysql_audit audit_null_descriptor=
{
  MYSQL_AUDIT_INTERFACE_VERSION,          /* interface version */
  NULL,                                    /* release_thd function */
  audit_null_notify,                       /* notify function */
  { (unsigned long) MYSQL_AUDIT_GENERAL_CLASSMASK |
    MYSQL_AUDIT_CONNECTION_CLASSMASK } /* class mask */
};
```

The server calls `audit_null_notify` to pass audit event information to the plugin. There is no `release_thd` function.

The event class mask indicates an interest in all events of the “general” and “connection” classes. `plugin_audit.h` defines symbols for these classes and their corresponding class masks:

```
#define MYSQL_AUDIT_GENERAL_CLASS 0
#define MYSQL_AUDIT_GENERAL_CLASSMASK (1 << MYSQL_AUDIT_GENERAL_CLASS)

#define MYSQL_AUDIT_CONNECTION_CLASS 1
#define MYSQL_AUDIT_CONNECTION_CLASSMASK (1 << MYSQL_AUDIT_CONNECTION_CLASS)
```

In the type-specific descriptor, the second and third parameters of the `event_notify` function prototype represent the event class and a generic pointer to an event structure:

```
void (*event_notify)(MYSQL_THD, unsigned int, const void *);
```

Events in different classes may have different structures, so the notification function should use the event class value to determine how to interpret the pointer to the event structure.

If the server calls the notification function with an event class of `MYSQL_AUDIT_GENERAL_CLASS`, it passes the event structure as a pointer to a `mysql_event_general` structure:

```
struct mysql_event_general
{
    unsigned int event_subclass;
    int general_error_code;
    unsigned long general_thread_id;
    const char *general_user;
    unsigned int general_user_length;
    const char *general_command;
    unsigned int general_command_length;
    const char *general_query;
    unsigned int general_query_length;
    struct charset_info_st *general_charset;
    unsigned long long general_time;
    unsigned long long general_rows;
};
```

Audit plugins can interpret `mysql_event_general` members as follows:

- `event_subclass`: The event subclass, one of the following values:

```
#define MYSQL_AUDIT_GENERAL_LOG 0
#define MYSQL_AUDIT_GENERAL_ERROR 1
#define MYSQL_AUDIT_GENERAL_RESULT 2
#define MYSQL_AUDIT_GENERAL_STATUS 3
```

- `general_error_code`: The error code. This is a value like that returned by the `mysql_errno()` C API function; 0 means “no error.”
- `general_thread_id`: The ID of the thread for which the event occurred.
- `general_user`: The current user for the event.
- `general_user_length`: The length of `general_user`, in bytes.
- `general_command`: For general query log events, the type of operation. Examples: `Connect`, `Query`, `Shutdown`. For error log events, the error message. This is a value like that returned by the `mysql_error()` C API function; an empty string means “no error.” For result events, this is empty.
- `general_command_length`: The length of `general_command`, in bytes.
- `general_query`: The SQL statement that was logged or produced a result.
- `general_query_length`: The length of `general_query`, in bytes.
- `general_charset`: Character set information for the event.
- `general_time`: A `TIMESTAMP` value indicating the time just before the notification function was called.
- `general_rows`: For general query log events, zero. For error log events, the row number at which an error occurred. For result events, the number of rows in the result plus one. For statements that produce no result set, the value is 0. This encoding enables statements that produce no result set to be distinguished from those that produce an empty result set. For example, for a `DELETE` statement, this value is 0. For a `SELECT`, the result is always 1 or more, where 1 represents an empty result set.

If the server calls the notification function with an event class of `MYSQL_AUDIT_CONNECTION_CLASS`, it passes the event structure as a pointer to a `mysql_event_connection` structure, which is similar to and interpreted much the same way as the `mysql_event_general` structure.

The `NULL_AUDIT` plugin notification function is quite simple. It increments a global event counter, determines the event class, then looks at the event subclass to determine which subclass counter to increment:

```
static void audit_null_notify(MYSQL_THD thd __attribute__((unused)),
                             unsigned int event_class,
                             const void *event)
{
    /* prone to races, oh well */
    number_of_calls++;
    if (event_class == MYSQL_AUDIT_GENERAL_CLASS)
    {
        const struct mysql_event_general *event_general=
            (const struct mysql_event_general *) event;
        switch (event_general->event_subclass)
        {
            case MYSQL_AUDIT_GENERAL_LOG:
                number_of_calls_general_log++;
                break;
            case MYSQL_AUDIT_GENERAL_ERROR:
                number_of_calls_general_error++;
                break;
            case MYSQL_AUDIT_GENERAL_RESULT:
                number_of_calls_general_result++;
                break;
            case MYSQL_AUDIT_GENERAL_STATUS:
                number_of_calls_general_status++;
                break;
            default:
                break;
        }
    }
    else if (event_class == MYSQL_AUDIT_CONNECTION_CLASS)
    {
        const struct mysql_event_connection *event_connection=
            (const struct mysql_event_connection *) event;
        switch (event_connection->event_subclass)
        {
            case MYSQL_AUDIT_CONNECTION_CONNECT:
                number_of_calls_connection_connect++;
                break;
            case MYSQL_AUDIT_CONNECTION_DISCONNECT:
                number_of_calls_connection_disconnect++;
                break;
            case MYSQL_AUDIT_CONNECTION_CHANGE_USER:
                number_of_calls_connection_change_user++;
                break;
            default:
                break;
        }
    }
}
```

To compile and install a plugin library object file, use the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable). For the `AUDIT_NULL` plugin, it is compiled and installed when you build MySQL from source. It is also included in binary distributions. The build process produces a shared object library with a name of `adt_null.so` (the suffix might differ depending on your platform).

To register the plugin at runtime, use this statement (change the suffix as necessary):

```
mysql> INSTALL PLUGIN NULL_AUDIT SONAME 'adt_null.so';
```

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement.

While the audit plugin is installed, it exposes status variables that indicate the events for which the plugin has been called:

```
mysql> SHOW STATUS LIKE 'Audit_null%';
+-----+
| Variable_name          | Value |
+-----+
| Audit_null_called      | 1388  |
| Audit_null_connection_change_user | 0     |
| Audit_null_connection_connect | 22    |
| Audit_null_connection_disconnect | 21    |
| Audit_null_general_error | 1     |
| Audit_null_general_log  | 513   |
| Audit_null_general_result | 415   |
| Audit_null_general_status | 416   |
+-----+

```

`Audit_null_called` counts all events, and the other variables count instances of event subclasses. For example, the preceding `SHOW STATUS` statement causes the server to send a result to the client and to write a message to the general query log if that log is enabled. Thus, a client that issues the statement repeatedly causes `Audit_null_called` and `Audit_null_general_result` to be incremented each time, and `Audit_null_general_log` to be incremented if the log is enabled.

To disable the plugin after testing it, use this statement to unload it:

```
mysql> UNINSTALL PLUGIN NULL_AUDIT;
```

### 22.2.4.9. Writing Authentication Plugins

MySQL supports pluggable authentication, in which plugins are invoked to authenticate client connections. Authentication plugins enable the use of authentication methods other than the built-in method of passwords stored in the `mysql.user` table. For example, plugins can be written to access external authentication methods. Also, authentication plugins can support the proxy user capability, such that the connecting user is a proxy for another user and is treated, for purposes of access control, as having the privileges of a different user. For more information, see [Section 6.3.7, “Pluggable Authentication”](#), and [Section 6.3.8, “Proxy Users”](#).

An authentication plugin can be written for the server side or the client side. Server-side plugins use the same plugin API that is used for the other server plugin types such as full-text parser or audit plugins (although with a different type-specific descriptor). Client-side plugins use the client plugin API.

Several header files contain information relevant to authentication plugins:

- `plugin.h`: Defines the `MYSQL_AUTHENTICATION_PLUGIN` server plugin type.
- `client_plugin.h`: Defines the API for client plugins. This includes the client plugin descriptor and function prototypes for client plugin C API calls (see [Section 21.8.13, “C API Client Plugin Functions”](#)).
- `plugin_auth.h`: Defines the part of the server plugin API specific to authentication plugins. This includes the type-specific descriptor for server-side authentication plugins and the `MYSQL_SERVER_AUTH_INFO` structure.
- `plugin_auth_common.h`: Contains common elements of client and server authentication plugins. This includes return value definitions and the `MYSQL_PLUGIN_VIO` structure.

To write an authentication plugin, include the following header files in the plugin source file. Other MySQL or general header files might also be needed.

- For a source file that implements a server authentication plugin, include this file:

```
#include <mysql/plugin_auth.h>
```

- For a source file that implements a client authentication plugin, or both client and server plugins, include these files:



```
#include <mysql/plugin_auth.h>
#include <mysql/client_plugin.h>
#include <mysql>
```

`plugin_auth.h` includes `plugin.h` and `plugin_auth_common.h`, so you need not include the latter files explicitly.

This section describes how to write a pair of simple server and client authentication plugins that work together.



### Warning

These plugins accept any non-empty password and the password is sent in clear text. This is insecure, so the plugins *should not be used in production environments*.

The server-side and client-side plugins developed here both are named `auth_simple`. As described in [Section 22.2.4.2, “Plugin Data Structures”](#), the plugin library file must have the same basename as the client plugin, so the source file name is `auth_simple.c` and produces a library named `auth_simple.so` (assuming that your system uses `.so` as the suffix for library files).

In MySQL source distributions, authentication plugin source is located in the `plugin/auth` directory and can be examined as a guide to writing other authentication plugins. Also, to see how the built-in authentication plugins are implemented, see `sql/sql_acl.cc` for plugins that are built in to the MySQL server and `sql-common/client.c` for plugins that are built in to the `libmysqlclient` client library. (For the built-in client plugins, note that the `auth_plugin_t` structures used there differ from the structures used with the usual client plugin declaration macros. In particular, the first two members are provided explicitly, not by declaration macros.)

#### 22.2.4.9.1. Writing the Server-Side Authentication Plugin

Declare the server-side plugin with the usual general descriptor format that is used for all server plugin types (see [Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”](#)). For the `auth_simple` plugin, the descriptor looks like this:

```
mysql_declare_plugin(auth_simple)
{
  MYSQL_AUTHENTICATION_PLUGIN,
  &auth_simple_handler,           /* type-specific descriptor */
  "auth_simple",                 /* plugin name */
  "Author Name",                 /* author */
  "Any-password authentication plugin", /* description */
  PLUGIN_LICENSE_GPL,           /* license type */
  NULL,                          /* no init function */
  NULL,                          /* no deinit function */
  0x0100,                        /* version = 1.0 */
  NULL,                          /* no status variables */
  NULL,                          /* no system variables */
  NULL,                          /* no reserved information */
  0                               /* no flags */
}
mysql_declare_plugin_end;
```

The `name` member (`auth_simple`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `SHOW PLUGINS` or `INFORMATION_SCHEMA.PLUGINS`.

The `auth_simple_handler` member of the general descriptor points to the type-specific descriptor. For an authentication plugin, the type-specific descriptor is an instance of the `st_mysql_auth` structure (defined in `plugin_auth.h`):

```
struct st_mysql_auth
{
  int interface_version;
```

```
const char *client_auth_plugin;
int (*authenticate_user)(MYSQL_PLUGIN_VIO *vio, MYSQL_SERVER_AUTH_INFO *info);
};
```

The `st_mysql_auth` structure has three members: The type-specific API version number, the client plugin name, and a pointer to the main plugin function that communicates with the client. The `client_auth_plugin` member should indicate the name of the client plugin if a specific plugin is required. A value of `NULL` means “any plugin.” In the latter case, whatever plugin the client uses will do. This is useful if the server plugin does not care about the client plugin or what user name or password it sends. For example, this might be true if the server plugin authenticates only local clients and uses some property of the operating system rather than the information sent by the client plugin.

For `auth_simple`, the type-specific descriptor looks like this:

```
static struct st_mysql_auth auth_simple_handler =
{
    MYSQL_AUTHENTICATION_INTERFACE_VERSION,
    "auth_simple",          /* required client-side plugin name */
    auth_simple_server      /* server-side plugin main function */
};
```

The main function takes two arguments representing an I/O structure and a `MYSQL_SERVER_AUTH_INFO` structure. The structure definition is found in `plugin_auth.h`. It looks like this:

```
typedef struct st_mysql_server_auth_info
{
    char *user_name;
    unsigned int user_name_length;
    const char *auth_string;
    unsigned long auth_string_length;
    char authenticated_as[MYSQL_USERNAME_LENGTH+1];
    char external_user[512];
    int password_used;
    const char *host_or_ip;
    unsigned int host_or_ip_length;
} MYSQL_SERVER_AUTH_INFO;
```

The character set for string members is UTF-8. If there is a `_length` member associated with a string, it indicates the string length in bytes. Strings are also null-terminated.

When an authentication plugin is invoked by the server, it should interpret the `MYSQL_SERVER_AUTH_INFO` structure members as follows. Some of these are used to set the value of SQL functions or system variables within the client session, as indicated.

- `user_name`: The user name sent by the client. The value becomes the `USER()` [1273] function value.
- `user_name_length`: The length of `user_name` in bytes.
- `auth_string`: The value of the `authentication_string` column of the `mysql.user` table row for the matching account name (that is, the row that matches the client user name and host name and that the server uses to determine how to authenticate the client).

Suppose that you create an account using the following statement:

```
CREATE USER 'my_user'@'localhost'
    IDENTIFIED WITH my_plugin AS 'my_auth_string';
```

When `my_user` connects from the local host, the server invokes `my_plugin` and passes `'my_auth_string'` to it as the `auth_string` value.

- `auth_string_length`: The length of `auth_string` in bytes.
- `authenticated_as`: The server sets this to the user name (the value of `user_name`). The plugin can alter it to indicate that the client should have the privileges of a different user. For example, if

the plugin supports proxy users, the initial value is the name of the connecting (proxy) user, and the plugin can change this member to the proxied user name. The server then treats the proxy user as having the privileges of the proxied user (assuming that the other conditions for proxy user support are satisfied; see [Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”](#)). The value is represented as a string at most `MYSQL_USER_NAME_LENGTH` bytes long, plus a terminating null. The value becomes the `CURRENT_USER() [1267]` function value.

- `external_user`: The server sets this to the empty string (null terminated). Its value becomes the `external_user [482]` system variable value. If the plugin wants that system variable to have a different value, it should set this member accordingly; for example, to the connecting user name. The value is represented as a string at most 511 bytes long, plus a terminating null.
- `password_used`: This member applies when authentication fails. The plugin can set it or ignore it. The value is used to construct the failure error message of `Authentication fails. Password used: %s`. The value of `password_used` determines how `%s` is handled, as shown in the following table.

<code>password_used</code>	<code>%s Handling</code>
0	NO
1	YES
2	There will be no <code>%s</code>

- `host_or_ip`: The name of the client host if it can be resolved, or the IP address otherwise.
- `host_or_ip_length`: The length of `host_or_ip` in bytes.

The `auth_simple` main function, `auth_simple_server()`, reads the password (a null-terminated string) from the client and succeeds if the password is nonempty (first byte not null):

```
static int auth_simple_server (MYSQL_PLUGIN_VIO *vio,
                              MYSQL_SERVER_AUTH_INFO *info)
{
    unsigned char *pkt;
    int pkt_len;

    /* read the password as null-terminated string, fail on error */
    if ((pkt_len= vio->read_packet(vio, &pkt)) < 0)
        return CR_ERROR;

    /* fail on empty password */
    if (!pkt_len || *pkt == '\0')
    {
        info->password_used= PASSWORD_USED_NO;
        return CR_ERROR;
    }

    /* accept any nonempty password */
    info->password_used= PASSWORD_USED_YES;

    return CR_OK;
}
```

The main function should return one of the error codes shown in the following table.

Error Code	Meaning
<code>CR_OK</code>	Success
<code>CR_OK_HANDSHAKE_COMPLETE</code>	Do not send a status packet back to client
<code>CR_ERROR</code>	Error
<code>CR_AUTH_USER_CREDENTIALS</code>	Authentication failure
<code>CR_AUTH_HANDSHAKE</code>	Authentication handshake failure

Error Code	Meaning
<a href="#">CR_AUTH_PLUGIN_ERROR</a>	Internal plugin error

For an example of how the handshake works, see the [plugin/auth/dialog.c](#) source file.

The error codes following [CR\\_ERROR](#) are available from MySQL 5.6.5. The server counts plugin errors in the Performance Schema [host\\_cache](#) table, also available as of MySQL 5.6.5.

[auth\\_simple\\_server\\_main\(\)](#) is so basic that it does not use the authentication information structure except to set the member that indicates whether a password was received.

A plugin that supports proxy users must return to the server the name of the proxied user (the MySQL user whose privileges the client user should get). To do this, the plugin must set the [info->authenticated\\_as](#) member to the proxied user name. For information about proxying, see [Section 6.3.8, “Proxy Users”](#), and [Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”](#).

### 22.2.4.9.2. Writing the Client-Side Authentication Plugin

Declare the client-side plugin descriptor with the [mysql\\_declare\\_client\\_plugin\(\)](#) and [mysql\\_end\\_client\\_plugin](#) macros (see [Section 22.2.4.2.3, “Client Plugin Descriptors”](#)). For the [auth\\_simple](#) plugin, the descriptor looks like this:

```
mysql_declare_client_plugin(AUTHENTICATION)
    "auth_simple",          /* plugin name */
    "Author Name",        /* author */
    "Any-password authentication plugin", /* description */
    {1,0,0},              /* version = 1.0.0 */
    "GPL",                 /* license type */
    NULL,                  /* for internal use */
    NULL,                  /* no init function */
    NULL,                  /* no deinit function */
    NULL,                  /* no option-handling function */
    auth_simple_client     /* main function */
mysql_end_client_plugin;
```

The descriptor members common to all client plugins are the plugin name through the option-handling function. Following the common members, an authentication plugin also has one extra member. This is the “main” function, which handles communication with the server. The function takes two arguments representing an I/O structure and a connection handler. For our simple any-password plugin, the main function does nothing but write to the server the password provided by the user:

```
static int auth_simple_client (MYSQL_PLUGIN_VIO *vio, MYSQL *mysql)
{
    int res;

    /* send password as null-terminated string in clear text */
    res= vio->write_packet(vio, (const unsigned char *) mysql->passwd,
        strlen(mysql->passwd) + 1);

    return res ? CR_ERROR : CR_OK;
}
```

The main function should return one of the error codes shown in the following table.

Error Code	Meaning
<a href="#">CR_OK</a>	Success
<a href="#">CR_OK_HANDSHAKE_COMPLETE</a>	Success, client done
<a href="#">CR_ERROR</a>	Error

[CR\\_OK\\_HANDSHAKE\\_COMPLETE](#) indicates that the client has done its part successfully and has read the last packet. A client plugin may return [CR\\_OK\\_HANDSHAKE\\_COMPLETE](#) if the number of round

trips in the authentication protocol is not known in advance and the plugin must read another packet to determine whether authentication is finished.

### 22.2.4.9.3. Using the Authentication Plugins

To compile and install a plugin library object file, see the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable).

Register the server-side plugin with the server. For example, to load the plugin at server startup, use a `--plugin-load=auth_simple.so` [435] option (change the library suffix as necessary for your system).

Create a user for whom the server will use the `auth_simple` plugin for authentication:

```
mysql> CREATE USER 'x'@'localhost'
-> IDENTIFIED WITH auth_simple;
```

Use a client program to connect to the server as user `x`. The server-side `auth_simple` plugin communicates with the client program that it should use the client-side `auth_simple` plugin, and the latter sends the password to the server. The server plugin should reject connections that send an empty password and accept connections that send a nonempty password. Invoke the client program each way to verify this:

```
shell> mysql --user=x --skip-password
ERROR 1045 (28000): Access denied for user 'x'@'localhost' (using password: NO)

shell> mysql --user=x --password=abc
mysql>
```

Because the server plugin accepts any nonempty password, it should be considered insecure. After testing the plugin to verify that it works, restart the server without the `--plugin-load` [435] option so as not to inadvertently leave the server running with an insecure authentication plugin loaded. Also, drop the user with `DROP USER 'x'@'localhost'`.

For additional information about loading and using authentication plugins, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#), and [Section 6.3.7, “Pluggable Authentication”](#).

If you are writing a client program that supports the use of authentication plugins, normally such a program causes a plugin to be loaded by calling `mysql_options()` to set the `MYSQL_DEFAULT_AUTH` and `MYSQL_PLUGIN_DIR` options:

```
char *plugin_dir = "path_to_plugin_dir";
char *default_auth = "plugin_name";

/* ... process command-line options ... */

mysql_options(&mysql, MYSQL_PLUGIN_DIR, plugin_dir);
mysql_options(&mysql, MYSQL_DEFAULT_AUTH, default_auth);
```

Typically, the program will also accept `--plugin-dir` and `--default-auth` options that enable users to override the default values.

Should a client program require lower-level plugin management, the client library contains functions that take an `st_mysql_client_plugin` argument. See [Section 21.8.13, “C API Client Plugin Functions”](#).

### 22.2.4.9.4. Implementing Proxy User Support in Authentication Plugins

One of the capabilities that pluggable authentication makes possible is proxy users (see [Section 6.3.8, “Proxy Users”](#)). For a server-side authentication plugin to participate in proxy user support, these conditions must be satisfied:

- When a connecting client should be treated as a proxy user, the plugin must return a different name in the `authenticated_as` member of the `MYSQL_SERVER_AUTH_INFO` structure, to indicate the proxied user name. It may also optionally set the `external_user` member, to set the value of the `external_user` [482] system variable.
- Proxy user accounts must be set up to be authenticated by the plugin. Use the `CREATE USER` or `GRANT` statement to associate accounts with plugins.
- Proxy user accounts must have the `PROXY` [755] privilege for the proxied accounts. Use the `GRANT` statement to grant this privilege.

In other words, the only aspect of proxy user support required of the plugin is that it set `authenticated_as` to the proxied user name. The rest is optional (setting `external_user`) or done by the DBA using SQL statements.

How does an authentication plugin determine which proxied user to return when the proxy user connects? That depends on the plugin. Typically, the plugin maps clients to proxied users based on the authentication string passed to it by the server. This string comes from the `AS` part of the `IDENTIFIED WITH` clause of the `CREATE USER` statement that specifies use of the plugin for authentication.

The plugin developer determines the syntax rules for the authentication string and implements the plugin according to those rules. Suppose that a plugin takes a comma-separated list of pairs that map external users to MySQL users. For example:

```
CREATE USER '@%.example.com'
  IDENTIFIED WITH my_plugin AS 'extuser1=mysqlusera, extuser2=mysqluserb'
CREATE USER '@%.example.org'
  IDENTIFIED WITH my_plugin AS 'extuser1=mysqluserc, extuser2=mysqluserd'
```

When the server invokes a plugin to authenticate a client, it passes the appropriate authentication string to the plugin. The plugin is responsible to:

1. Parse the string into its components to determine the mapping to use
2. Compare the client user name to the mapping
3. Return the proper MySQL user name

For example, if `extuser2` connects from an `example.com` host, the server passes `'extuser1=mysqlusera, extuser2=mysqluserb'` to the plugin, and the plugin should copy `mysqluserb` into `authenticated_as`, with a terminating null byte. If `extuser2` connects from an `example.org` host, the server passes `'extuser1=mysqluserc, extuser2=mysqluserd'`, and the plugin should copy `mysqluserd` instead.

If there is no match in the mapping, the action depends on the plugin. If a match is required, the plugin likely will return an error. Or the plugin might simply return the client name; in this case, it should not change `authenticated_as`, and the server will not treat the client as a proxy.

The following example demonstrates how to handle proxy users using a plugin named `auth_simple_proxy`. Like the `auth_simple` plugin described earlier, `auth_simple_proxy` accepts any nonempty password as valid (and thus should not be used in production environments). In addition, it examines the `auth_string` authentication string member and uses these very simple rules for interpreting it:

- If the string is empty, the plugin returns the user name as given and no proxying occurs. That is, the plugin leaves the value of `authenticated_as` unchanged.
- If the string is nonempty, the plugin treats it as the name of the proxied user and copies it to `authenticated_as` so that proxying occurs.

For testing, set up one account that is not proxied according to the preceding rules, and one that is. This means that one account has no `AS` clause, and one includes an `AS` clause that names the proxied user:

```
CREATE USER 'plugin_user1'@'localhost'
  IDENTIFIED WITH auth_simple_proxy;
CREATE USER 'plugin_user2'@'localhost'
  IDENTIFIED WITH auth_simple_proxy AS 'proxied_user';
```

In addition, create an account for the proxied user and grant `plugin_user2` the `PROXY` [755] privilege for it:

```
CREATE USER 'proxied_user'@'localhost'
  IDENTIFIED BY 'proxied_user_pass';
GRANT PROXY
  ON 'proxied_user'@'localhost'
  TO 'plugin_user2'@'localhost';
```

Before the server invokes an authentication plugin, it sets `authenticated_as` to the client user name. To indicate that the user is a proxy, the plugin should set `authenticated_as` to the proxied user name. For `auth_simple_proxy`, this means that it must examine the `auth_string` value, and, if the value is nonempty, copy it to the `authenticated_as` member to return it as the name of the proxied user. In addition, when proxying occurs, the plugin sets the `external_user` member to the client user name; this becomes the value of the `external_user` [482] system variable.

```
static int auth_simple_proxy_server (MYSQL_PLUGIN_VIO *vio,
                                     MYSQL_SERVER_AUTH_INFO *info)
{
  unsigned char *pkt;
  int pkt_len;

  /* read the password as null-terminated string, fail on error */
  if ((pkt_len= vio->read_packet(vio, &pkt)) < 0)
    return CR_ERROR;

  /* fail on empty password */
  if (!pkt_len || *pkt == '\0')
  {
    info->password_used= PASSWORD_USED_NO;
    return CR_ERROR;
  }

  /* accept any nonempty password */
  info->password_used= PASSWORD_USED_YES;

  /* if authentication string is nonempty, use as proxied user name */
  /* and use client name as external_user value */
  if (info->auth_string_length > 0)
  {
    strcpy (info->authenticated_as, info->auth_string);
    strcpy (info->external_user, info->user_name);
  }

  return CR_OK;
}
```

After a successful connection, the `USER()` [1273] function should indicate the connecting client user and host name, and `CURRENT_USER()` [1267] should indicate the account whose privileges apply during the session. The latter value should be the connecting user account if no proxying occurs or the proxied account if proxying does occur.

Compile and install the plugin, then test it. First, connect as `plugin_user1`:

```
shell> mysql --user=plugin_user1 --password=x
```

In this case, there should be no proxying:

```
mysql> SELECT USER(), CURRENT_USER(), @@proxy_user, @@external_user\G
***** 1. row *****
      USER(): plugin_user1@localhost
     CURRENT_USER(): plugin_user1@localhost
      @@proxy_user: NULL
     @@external_user: NULL
```



Then connect as `plugin_user2`:

```
shell> mysql --user=plugin_user2 --password=x
```

In this case, `plugin_user2` should be proxied to `proxied_user`:

```
mysql> SELECT USER(), CURRENT_USER(), @@proxy_user, @@external_user\G
***** 1. row *****
      USER(): plugin_user2@localhost
      CURRENT_USER(): proxied_user@localhost
      @@proxy_user: 'plugin_user2'@'localhost'
      @@external_user: 'plugin_user2'@'localhost'
```

## 22.2.4.10. Writing Password-Validation Plugins

This section describes how to write a password-validation server plugin. The instructions are based on the source code in the `plugin/password_validation` directory of MySQL source distributions. The `validate_password.cc` source file in that directory implements the the plugin named `validate_password`.

To write a password-validation plugin, include the following header file in the plugin source file. Other MySQL or general header files might also be needed.

```
#include <mysql/plugin_validate_password.h>
```

`plugin_validate_password.h` includes `plugin.h`, so you need not include the latter file explicitly. `plugin.h` defines the `MYSQL_VALIDATE_PASSWORD_PLUGIN` server plugin type and the data structures needed to declare the plugin. `plugin_validate_password.h` defines data structures specific to password-validation plugins.

A password-validation plugin, like any MySQL server plugin, has a general plugin descriptor (see [Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”](#)). In `validate_password.cc`, the general descriptor looks like this:

```
mysql_declare_plugin(validate_password)
{
    MYSQL_VALIDATE_PASSWORD_PLUGIN, /* type */
    &validate_password_descriptor, /* descriptor */
    "validate_password", /* name */
    "Oracle Corporation", /* author */
    "check password strength", /* description */
    PLUGIN_LICENSE_GPL,
    validate_password_init, /* init function (when loaded) */
    validate_password_deinit, /* deinit function (when unloaded) */
    0x0100, /* version */
    NULL,
    validate_password_system_variables, /* system variables */
    NULL,
    0,
}
mysql_declare_plugin_end;
```

The `name` member (`validate_password`) indicates the name to use for references to the plugin in statements such as `INSTALL PLUGIN` or `UNINSTALL PLUGIN`. This is also the name displayed by `INFORMATION_SCHEMA.PLUGINS` or `SHOW PLUGINS`.

The general descriptor also refers to `validate_password_system_variables`, a structure that exposes several system variables to the `SHOW VARIABLES` statement:

```
static struct st_mysql_sys_var* validate_password_system_variables[]= {
    MYSQL_SYSVAR(length),
    MYSQL_SYSVAR(number_count),
    MYSQL_SYSVAR(mixed_case_count),
    MYSQL_SYSVAR(special_char_count),
    MYSQL_SYSVAR(policy),
    MYSQL_SYSVAR(dictionary_file),
```



```
NULL
};
```

The `validate_password_init` initialization function reads the dictionary file if one was specified, and the `validate_password_deinit` function frees data structures associated with the file.

The `validate_password_descriptor` value in the general descriptor points to the type-specific descriptor. For password-validation plugins, this descriptor has the following structure:

```
struct st_mysql_validate_password
{
    int interface_version;
    /*
     * This function returns TRUE for passwords which satisfy the password
     * policy (as chosen by plugin variable) and FALSE for all other
     * password
     */
    int (*validate_password)(mysql_string_handle password);
    /*
     * This function returns the password strength (0-100) depending
     * upon the policies
     */
    int (*get_password_strength)(mysql_string_handle password);
};
```

The type-specific descriptor has these members:

- `interface_version`: By convention, type-specific plugin descriptors begin with the interface version for the given plugin type. The server checks `interface_version` when it loads the plugin to see whether the plugin is compatible with it. For password-validation plugins, the value of the `interface_version` member is `MYSQL_VALIDATE_PASSWORD_INTERFACE_VERSION` (defined in `plugin_validate_password.h`).
- `validate_password`: A function that the server calls to test whether a password satisfies the current password policy. It returns 1 if the password is okay and 0 otherwise. The argument is the password, passed as a `mysql_string_handle` value. This data type is implemented by the `mysql_string` server service. For details, see the `string_service.h` and `string_service.cc` source files in the `sql` directory.
- `get_password_strength`: A function that the server calls to assess the strength of a password. It returns a value from 0 (weak) to 100 (strong). The argument is the password, passed as a `mysql_string_handle` value.

For the `validate_password` plugin, the type-specific descriptor looks like this:

```
static struct st_mysql_validate_password validate_password_descriptor=
{
    MYSQL_VALIDATE_PASSWORD_INTERFACE_VERSION,
    validate_password, /* validate function */
    get_password_strength /* validate strength function */
};
```

To compile and install a plugin library object file, use the instructions in [Section 22.2.4.3, “Compiling and Installing Plugin Libraries”](#). To use the library file, it must be installed in the plugin directory (the directory named by the `plugin_dir` [525] system variable). For the `validate_password` plugin, it is compiled and installed when you build MySQL from source. It is also included in binary distributions. The build process produces a shared object library with a name of `validate_password.so` (the extension might be different depending on your platform).

To register the plugin at runtime, use this statement (change the extension as necessary):

```
mysql> INSTALL PLUGIN validate_password SONAME 'validate_password.so';
```

For additional information about plugin loading, see [Section 5.1.8.1, “Installing and Uninstalling Plugins”](#).

To verify plugin installation, examine the `INFORMATION_SCHEMA.PLUGINS` table or use the `SHOW PLUGINS` statement.

While the `validate_password` plugin is installed, it exposes system variables that indicate the password-checking parameters:

```
mysql> SHOW VARIABLES LIKE 'validate_password%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| validate_password_dictionary_file |      |
| validate_password_length | 8 |
| validate_password_mixed_case_count | 1 |
| validate_password_number_count | 1 |
| validate_password_policy | MEDIUM |
| validate_password_special_char_count | 1 |
+-----+-----+
```

For descriptions of these variables, see [Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”](#).

To disable the plugin after testing it, use this statement to unload it:

```
mysql> UNINSTALL PLUGIN validate_password;
```

## 22.2.5. MySQL Services for Plugins

MySQL server plugins have access to server “services.” The services interface exposes server functionality that plugins can call. It complements the plugin API and has these characteristics:

- Services enable plugins to access code inside the server using ordinary function calls.
- Services are portable and work on multiple platforms.
- The interface includes a versioning mechanism so that service versions supported by the server can be checked at load time against plugin versions. Versioning protects against incompatibilities between the version of a service that the server provides and the version of the service expected or required by a plugin.

Current services include the following, and others can be implemented:

- `my_plugin_log_service`: A service that enables plugins to report errors and specify error messages. The server writes the messages to the error log.
- `my_snprintf`: A string-formatting service that produces consistent results across platforms.
- `my_thd_scheduler`: A service for plugins to select a thread scheduler.
- `mysql_string`: A service for string manipulation.
- `thd_alloc`: A memory-allocation service.
- `thd_wait`: A service for plugins to report when they are going to sleep or stall.

The plugin services interface differs from the plugin API as follows:

- The plugin API enables plugins to be used by the server. The calling initiative lies with the server to invoke plugins. This enables plugins to extend server functionality or register to receive notifications about server processing.
- The plugin services interface enables plugins to call code inside the server. The calling initiative lies with plugins to invoke service functions. This enables functionality already implemented in the server to be used by many plugins; they need not individually implement it themselves.

For developers who wish to modify the server to add a new service, see [MySQL Services for Plugins](#).

The remainder of this section describes how a plugin uses server functionality that is available as a service. See also the source for the “daemon” example plugin, which uses the `my_snprintf` service. Within a MySQL source distribution, that plugin is located in the `plugin/daemon_example` directory.

To determine what services exist and what functions they provide, look in the `include/mysql` directory of a MySQL source distribution. The relevant files are:

- `plugin.h` includes `services.h`.
- `services.h` is the “umbrella” header that includes all available service-specific header files.
- Service-specific headers have names like `service_my_snprintf.h` or `service_thd_alloc.h`.

Each service-specific header should contain comments that provide full usage documentation for a given service, including what service functions are available, their calling sequences, and return values.

To use a service or services from within a plugin, the plugin source file must include the `plugin.h` header file to access service-related information:

```
#include <mysql/plugin.h>
```

This does not represent any additional setup cost. A plugin must include that file anyway because it contains definitions and structures that every plugin needs.

To access a service, a plugin calls service functions like any other function. For example, to format a string into a buffer for printing, call the `my_snprintf()` function provided by the service of the same name:

```
char buffer[BUFFER_SIZE];
my_snprintf(buffer, sizeof(buffer), format_string, argument_to_format, ...);
```

To report an error that the server will write to the error log, first choose an error level. `mysql/service_my_plugin_log.h` defines these levels:

```
enum plugin_log_level
{
    MY_ERROR_LEVEL,
    MY_WARNING_LEVEL,
    MY_INFORMATION_LEVEL
};
```

Then invoke `my_plugin_log_message()`:

```
int my_plugin_log_message(MYSQL_PLUGIN *plugin, enum plugin_log_level level,
    const char *format, ...);
```

For example:

```
my_plugin_log_message(plugin_ptr, MY_ERROR_LEVEL, "Cannot initialize plugin");
```

When you build your plugin, you must link in the `libmysqservices` library. Use the `-lmysqservices` flag at link time. For example, for `CMake`, put this in the top-level `CMakeLists.txt` file:

```
FIND_LIBRARY(MYSQLSERVICES_LIB mysqservices
    PATHS "${MYSQL_SRCDIR}/libservices" NO_DEFAULT_PATH)
```

Put this in the `CMakeLists.txt` file in the directory containing the plugin source:

```
# the plugin needs the mysql services library for error logging
```

```
TARGET_LINK_LIBRARIES (your_plugin_library_name ${MYSQLSERVICES_LIB})
```

## 22.3. Adding New Functions to MySQL

There are three ways to add new functions to MySQL:

- You can add functions through the user-defined function (UDF) interface. User-defined functions are compiled as object files and then added to and removed from the server dynamically using the `CREATE FUNCTION` and `DROP FUNCTION` statements. See [Section 13.7.3.1, “CREATE FUNCTION Syntax for User-Defined Functions”](#).
- You can add functions as native (built-in) MySQL functions. Native functions are compiled into the `mysqld` server and become available on a permanent basis.
- Another way to add functions is by creating stored functions. These are written using SQL statements rather than by compiling object code. The syntax for writing stored functions is not covered here. See [Section 18.2, “Using Stored Routines \(Procedures and Functions\)”](#).

Each method of creating compiled functions has advantages and disadvantages:

- If you write user-defined functions, you must install object files in addition to the server itself. If you compile your function into the server, you don't need to do that.
- Native functions require you to modify a source distribution. UDFs do not. You can add UDFs to a binary MySQL distribution. No access to MySQL source is necessary.
- If you upgrade your MySQL distribution, you can continue to use your previously installed UDFs, unless you upgrade to a newer version for which the UDF interface changes. For native functions, you must repeat your modifications each time you upgrade.

Whichever method you use to add new functions, they can be invoked in SQL statements just like native functions such as `ABS ( )` [1202] or `SOUNDEX ( )` [1184].

See [Section 9.2.4, “Function Name Parsing and Resolution”](#), for the rules describing how the server interprets references to different kinds of functions.

The following sections describe features of the UDF interface, provide instructions for writing UDFs, discuss security precautions that MySQL takes to prevent UDF misuse, and describe how to add native MySQL functions.

For example source code that illustrates how to write UDFs, take a look at the `sql/udf_example.c` file that is provided in MySQL source distributions.

### 22.3.1. Features of the User-Defined Function Interface

The MySQL interface for user-defined functions provides the following features and capabilities:

- Functions can return string, integer, or real values and can accept arguments of those same types.
- You can define simple functions that operate on a single row at a time, or aggregate functions that operate on groups of rows.
- Information is provided to functions that enables them to check the number, types, and names of the arguments passed to them.
- You can tell MySQL to coerce arguments to a given type before passing them to a function.
- You can indicate that a function returns `NULL` or that an error occurred.

### 22.3.2. Adding a New User-Defined Function

For the UDF mechanism to work, functions must be written in C or C++ and your operating system must support dynamic loading. MySQL source distributions include a file `sql/udf_example.c` that defines five UDF functions. Consult this file to see how UDF calling conventions work. The `include/mysql_com.h` header file defines UDF-related symbols and data structures, although you need not include this header file directly; it is included by `mysql.h`.

A UDF contains code that becomes part of the running server, so when you write a UDF, you are bound by any and all constraints that apply to writing server code. For example, you may have problems if you attempt to use functions from the `libstdc++` library. These constraints may change in future versions of the server, so it is possible that server upgrades will require revisions to UDFs that were originally written for older servers. For information about these constraints, see [Section 2.9.4, “MySQL Source-Configuration Options”](#), and [Section 2.9.5, “Dealing with Problems Compiling MySQL”](#).

To be able to use UDFs, you must link `mysqld` dynamically. If you want to use a UDF that needs to access symbols from `mysqld` (for example, the `metaphone` function in `sql/udf_example.c` uses `default_charset_info`), you must link the program with `-rdynamic` (see `man dlopen`).

For each function that you want to use in SQL statements, you should define corresponding C (or C++) functions. In the following discussion, the name “xxx” is used for an example function name. To distinguish between SQL and C/C++ usage, `XXX()` (uppercase) indicates an SQL function call, and `xxx()` (lowercase) indicates a C/C++ function call.



#### Note

When using C++ you can encapsulate your C functions within:

```
extern "C" { ... }
```

This ensures that your C++ function names remain readable in the completed UDF.

The following list describes the C/C++ functions that you write to implement the interface for a function named `xxx()`. The main function, `xxx()`, is required. In addition, a UDF requires at least one of the other functions described here, for reasons discussed in [Section 22.3.2.6, “User-Defined Function Security Precautions”](#).

- `xxx()`

The main function. This is where the function result is computed. The correspondence between the SQL function data type and the return type of your C/C++ function is shown here.

SQL Type	C/C++ Type
STRING	<code>char *</code>
INTEGER	<code>long long</code>
REAL	<code>double</code>

It is also possible to declare a `DECIMAL` function, but currently the value is returned as a string, so you should write the UDF as though it were a `STRING` function. `ROW` functions are not implemented.

- `xxx_init()`

The initialization function for `xxx()`. If present, it can be used for the following purposes:

- To check the number of arguments to `xxx()`.
- To verify that the arguments are of a required type or, alternatively, to tell MySQL to coerce arguments to the required types when the main function is called.
- To allocate any memory required by the main function.

- To specify the maximum length of the result.
- To specify (for `REAL` functions) the maximum number of decimal places in the result.
- To specify whether the result can be `NULL`.
- `xxx_deinit()`

The deinitialization function for `xxx()`. If present, it should deallocate any memory allocated by the initialization function.

When an SQL statement invokes `XXX()`, MySQL calls the initialization function `xxx_init()` to let it perform any required setup, such as argument checking or memory allocation. If `xxx_init()` returns an error, MySQL aborts the SQL statement with an error message and does not call the main or deinitialization functions. Otherwise, MySQL calls the main function `xxx()` once for each row. After all rows have been processed, MySQL calls the deinitialization function `xxx_deinit()` so that it can perform any required cleanup.

For aggregate functions that work like `SUM()` [1285], you must also provide the following functions:

- `xxx_clear()`

Reset the current aggregate value but do not insert the argument as the initial aggregate value for a new group.

- `xxx_add()`

Add the argument to the current aggregate value.

MySQL handles aggregate UDFs as follows:

1. Call `xxx_init()` to let the aggregate function allocate any memory it needs for storing results.
2. Sort the table according to the `GROUP BY` expression.
3. Call `xxx_clear()` for the first row in each new group.
4. Call `xxx_add()` for each row that belongs in the same group.
5. Call `xxx()` to get the result for the aggregate when the group changes or after the last row has been processed.
6. Repeat steps 3 to 5 until all rows has been processed
7. Call `xxx_deinit()` to let the UDF free any memory it has allocated.

All functions must be thread-safe. This includes not just the main function, but the initialization and deinitialization functions as well, and also the additional functions required by aggregate functions. A consequence of this requirement is that you are not permitted to allocate any global or static variables that change! If you need memory, you should allocate it in `xxx_init()` and free it in `xxx_deinit()`.

### 22.3.2.1. UDF Calling Sequences for Simple Functions

This section describes the different functions that you need to define when you create a simple UDF. [Section 22.3.2, “Adding a New User-Defined Function”](#), describes the order in which MySQL calls these functions.

The main `xxx()` function should be declared as shown in this section. Note that the return type and parameters differ, depending on whether you declare the SQL function `XXX()` to return `STRING`, `INTEGER`, or `REAL` in the `CREATE FUNCTION` statement:

For **STRING** functions:

```
char *xxx(UDF_INIT *initid, UDF_ARGS *args,
         char *result, unsigned long *length,
         char *is_null, char *error);
```

For **INTEGER** functions:

```
long long xxx(UDF_INIT *initid, UDF_ARGS *args,
             char *is_null, char *error);
```

For **REAL** functions:

```
double xxx(UDF_INIT *initid, UDF_ARGS *args,
          char *is_null, char *error);
```

**DECIMAL** functions return string values and should be declared the same way as **STRING** functions. **ROW** functions are not implemented.

The initialization and deinitialization functions are declared like this:

```
my_bool xxx_init(UDF_INIT *initid, UDF_ARGS *args, char *message);

void xxx_deinit(UDF_INIT *initid);
```

The `initid` parameter is passed to all three functions. It points to a `UDF_INIT` structure that is used to communicate information between functions. The `UDF_INIT` structure members follow. The initialization function should fill in any members that it wishes to change. (To use the default for a member, leave it unchanged.)

- `my_bool maybe_null`

`xxx_init()` should set `maybe_null` to 1 if `xxx()` can return `NULL`. The default value is 1 if any of the arguments are declared `maybe_null`.

- `unsigned int decimals`

The number of decimal digits to the right of the decimal point. The default value is the maximum number of decimal digits in the arguments passed to the main function. For example, if the function is passed 1.34, 1.345, and 1.3, the default would be 3, because 1.345 has 3 decimal digits.

For arguments that have no fixed number of decimals, the `decimals` value is set to 31, which is 1 more than the maximum number of decimals permitted for the **DECIMAL**, **FLOAT**, and **DOUBLE** data types. In MySQL 5.6, this value is available as the constant `NOT_FIXED_DEC` in the `mysql_com.h` header file.

A `decimals` value of 31 is used for arguments in cases such as a **FLOAT** or **DOUBLE** column declared without an explicit number of decimals (for example, **FLOAT** rather than **FLOAT(10, 3)**) and for floating-point constants such as `1345E-3`. It is also used for string and other nonnumber arguments that might be converted within the function to numeric form.

The value to which the `decimals` member is initialized is only a default. It can be changed within the function to reflect the actual calculation performed. The default is determined such that the largest number of decimals of the arguments is used. If the number of decimals is `NOT_FIXED_DEC` for even one of the arguments, that is the value used for `decimals`.

- `unsigned int max_length`

The maximum length of the result. The default `max_length` value differs depending on the result type of the function. For string functions, the default is the length of the longest argument. For integer functions, the default is 21 digits. For real functions, the default is 13 plus the number of decimal digits indicated by `initid->decimals`. (For numeric functions, the length includes any sign or decimal point characters.)



If you want to return a blob value, you can set `max_length` to 65KB or 16MB. This memory is not allocated, but the value is used to decide which data type to use if there is a need to temporarily store the data.

- `char *ptr`

A pointer that the function can use for its own purposes. For example, functions can use `initid->ptr` to communicate allocated memory among themselves. `xxx_init()` should allocate the memory and assign it to this pointer:

```
initid->ptr = allocated_memory;
```

In `xxx()` and `xxx_deinit()`, refer to `initid->ptr` to use or deallocate the memory.

- `my_bool const_item`

`xxx_init()` should set `const_item` to 1 if `xxx()` always returns the same value and to 0 otherwise.

### 22.3.2.2. UDF Calling Sequences for Aggregate Functions

This section describes the different functions that you need to define when you create an aggregate UDF. Section 22.3.2, “Adding a New User-Defined Function”, describes the order in which MySQL calls these functions.

- `xxx_reset()`

This function is called when MySQL finds the first row in a new group. It should reset any internal summary variables and then use the given `UDF_ARGS` argument as the first value in your internal summary value for the group. Declare `xxx_reset()` as follows:

```
void xxx_reset(UDF_INIT *initid, UDF_ARGS *args,
              char *is_null, char *error);
```

`xxx_reset()` is not needed or used in MySQL 5.6, in which the UDF interface uses `xxx_clear()` instead. However, you can define both `xxx_reset()` and `xxx_clear()` if you want to have your UDF work with older versions of the server. (If you do include both functions, the `xxx_reset()` function in many cases can be implemented internally by calling `xxx_clear()` to reset all variables, and then calling `xxx_add()` to add the `UDF_ARGS` argument as the first value in the group.)

- `xxx_clear()`

This function is called when MySQL needs to reset the summary results. It is called at the beginning for each new group but can also be called to reset the values for a query where there were no matching rows. Declare `xxx_clear()` as follows:

```
void xxx_clear(UDF_INIT *initid, char *is_null, char *error);
```

`is_null` is set to point to `CHAR(0)` before calling `xxx_clear()`.

If something went wrong, you can store a value in the variable to which the `error` argument points. `error` points to a single-byte variable, not to a string buffer.

`xxx_clear()` is required by MySQL 5.6.

- `xxx_add()`

This function is called for all rows that belong to the same group. You should use it to add the value in the `UDF_ARGS` argument to your internal summary variable.

```
void xxx_add(UDF_INIT *initid, UDF_ARGS *args,
            char *is_null, char *error);
```



The `xxx()` function for an aggregate UDF should be declared the same way as for a nonaggregate UDF. See [Section 22.3.2.1, “UDF Calling Sequences for Simple Functions”](#).

For an aggregate UDF, MySQL calls the `xxx()` function after all rows in the group have been processed. You should normally never access its `UDF_ARGS` argument here but instead return a value based on your internal summary variables.

Return value handling in `xxx()` should be done the same way as for a nonaggregate UDF. See [Section 22.3.2.4, “UDF Return Values and Error Handling”](#).

The `xxx_reset()` and `xxx_add()` functions handle their `UDF_ARGS` argument the same way as functions for nonaggregate UDFs. See [Section 22.3.2.3, “UDF Argument Processing”](#).

The pointer arguments to `is_null` and `error` are the same for all calls to `xxx_reset()`, `xxx_clear()`, `xxx_add()` and `xxx()`. You can use this to remember that you got an error or whether the `xxx()` function should return `NULL`. You should not store a string into `*error! error` points to a single-byte variable, not to a string buffer.

`*is_null` is reset for each group (before calling `xxx_clear()`). `*error` is never reset.

If `*is_null` or `*error` are set when `xxx()` returns, MySQL returns `NULL` as the result for the group function.

### 22.3.2.3. UDF Argument Processing

The `args` parameter points to a `UDF_ARGS` structure that has the members listed here:

- `unsigned int arg_count`

The number of arguments. Check this value in the initialization function if you require your function to be called with a particular number of arguments. For example:

```
if (args->arg_count != 2)
{
    strcpy(message, "XXX() requires two arguments");
    return 1;
}
```

For other `UDF_ARGS` member values that are arrays, array references are zero-based. That is, refer to array members using index values from 0 to `args->arg_count - 1`.

- `enum Item_result *arg_type`

A pointer to an array containing the types for each argument. The possible type values are `STRING_RESULT`, `INT_RESULT`, `REAL_RESULT`, and `DECIMAL_RESULT`.

To make sure that arguments are of a given type and return an error if they are not, check the `arg_type` array in the initialization function. For example:

```
if (args->arg_type[0] != STRING_RESULT ||
    args->arg_type[1] != INT_RESULT)
{
    strcpy(message, "XXX() requires a string and an integer");
    return 1;
}
```

Arguments of type `DECIMAL_RESULT` are passed as strings, so you should handle them the same way as `STRING_RESULT` values.

As an alternative to requiring your function's arguments to be of particular types, you can use the initialization function to set the `arg_type` elements to the types you want. This causes MySQL to coerce arguments to those types for each call to `xxx()`. For example, to specify that the first two arguments should be coerced to string and integer, respectively, do this in `xxx_init()`:

```
args->arg_type[0] = STRING_RESULT;
args->arg_type[1] = INT_RESULT;
```

Exact-value decimal arguments such as `1.3` or `DECIMAL` column values are passed with a type of `DECIMAL_RESULT`. However, the values are passed as strings. If you want to receive a number, use the initialization function to specify that the argument should be coerced to a `REAL_RESULT` value:

```
args->arg_type[2] = REAL_RESULT;
```

- `char **args`

`args->args` communicates information to the initialization function about the general nature of the arguments passed to your function. For a constant argument `i`, `args->args[i]` points to the argument value. (See later for instructions on how to access the value properly.) For a nonconstant argument, `args->args[i]` is `0`. A constant argument is an expression that uses only constants, such as `3` or `4*7-2` or `SIN(3.14)` [1208]. A nonconstant argument is an expression that refers to values that may change from row to row, such as column names or functions that are called with nonconstant arguments.

For each invocation of the main function, `args->args` contains the actual arguments that are passed for the row currently being processed.

If argument `i` represents `NULL`, `args->args[i]` is a null pointer (`0`). If the argument is not `NULL`, functions can refer to it as follows:

- An argument of type `STRING_RESULT` is given as a string pointer plus a length, to enable handling of binary data or data of arbitrary length. The string contents are available as `args->args[i]` and the string length is `args->lengths[i]`. Do not assume that the string is null-terminated.
- For an argument of type `INT_RESULT`, you must cast `args->args[i]` to a `long long` value:

```
long long int_val;
int_val = *((long long*) args->args[i]);
```

- For an argument of type `REAL_RESULT`, you must cast `args->args[i]` to a `double` value:

```
double real_val;
real_val = *((double*) args->args[i]);
```

- For an argument of type `DECIMAL_RESULT`, the value is passed as a string and should be handled like a `STRING_RESULT` value.
- `ROW_RESULT` arguments are not implemented.

- `unsigned long *lengths`

For the initialization function, the `lengths` array indicates the maximum string length for each argument. You should not change these. For each invocation of the main function, `lengths` contains the actual lengths of any string arguments that are passed for the row currently being processed. For arguments of types `INT_RESULT` or `REAL_RESULT`, `lengths` still contains the maximum length of the argument (as for the initialization function).

- `char *maybe_null`

For the initialization function, the `maybe_null` array indicates for each argument whether the argument value might be null (`0` if no, `1` if yes).

- `char **attributes`

`args->attributes` communicates information about the names of the UDF arguments. For argument `i`, the attribute name is available as a string in `args->attributes[i]` and the attribute length is `args->attribute_lengths[i]`. Do not assume that the string is null-terminated.

By default, the name of a UDF argument is the text of the expression used to specify the argument. For UDFs, an argument may also have an optional `[AS] alias_name` clause, in which case the argument name is `alias_name`. The `attributes` value for each argument thus depends on whether an alias was given.

Suppose that a UDF `my_udf()` is invoked as follows:

```
SELECT my_udf(expr1, expr2 AS alias1, expr3 alias2);
```

In this case, the `attributes` and `attribute_lengths` arrays will have these values:

```
args->attributes[0] = "expr1"
args->attribute_lengths[0] = 5

args->attributes[1] = "alias1"
args->attribute_lengths[1] = 6

args->attributes[2] = "alias2"
args->attribute_lengths[2] = 6
```

- `unsigned long *attribute_lengths`

The `attribute_lengths` array indicates the length of each argument name.

#### 22.3.2.4. UDF Return Values and Error Handling

The initialization function should return `0` if no error occurred and `1` otherwise. If an error occurs, `xxx_init()` should store a null-terminated error message in the `message` parameter. The message is returned to the client. The message buffer is `MYSQL_ERRMSG_SIZE` characters long, but you should try to keep the message to less than 80 characters so that it fits the width of a standard terminal screen.

The return value of the main function `xxx()` is the function value, for `long long` and `double` functions. A string function should return a pointer to the result and set `*length` to the length (in bytes) of the return value. For example:

```
memcpy(result, "result string", 13);
*length = 13;
```

MySQL passes a buffer to the `xxx()` function using the `result` parameter. This buffer is sufficiently long to hold 255 characters, which can be multi-byte characters. The `xxx()` function can store the result in this buffer if it fits, in which case the return value should be a pointer to the buffer. If the function stores the result in a different buffer, it should return a pointer to that buffer.

If your string function does not use the supplied buffer (for example, if it needs to return a string longer than 255 characters), you must allocate the space for your own buffer with `malloc()` in your `xxx_init()` function or your `xxx()` function and free it in your `xxx_deinit()` function. You can store the allocated memory in the `ptr` slot in the `UDF_INIT` structure for reuse by future `xxx()` calls. See [Section 22.3.2.1, "UDF Calling Sequences for Simple Functions"](#).

To indicate a return value of `NULL` in the main function, set `*is_null` to `1`:

```
*is_null = 1;
```

To indicate an error return in the main function, set `*error` to `1`:

```
*error = 1;
```

If `xxx()` sets `*error` to `1` for any row, the function value is `NULL` for the current row and for any subsequent rows processed by the statement in which `xxx()` was invoked. (`xxx()` is not even called for subsequent rows.)

#### 22.3.2.5. Compiling and Installing User-Defined Functions

Files implementing UDFs must be compiled and installed on the host where the server runs. This process is described below for the example UDF file `sql/udf_example.c` that is included in MySQL source distributions.

If a UDF will be referred to in statements that will be replicated to slave servers, you must ensure that every slave also has the function available. Otherwise, replication will fail on the slaves when they attempt to invoke the function.

The immediately following instructions are for Unix. Instructions for Windows are given later in this section.

The `udf_example.c` file contains the following functions:

- `metaphon()` returns a metaphon string of the string argument. This is something like a soundex string, but it is more tuned for English.
- `myfunc_double()` returns the sum of the ASCII values of the characters in its arguments, divided by the sum of the length of its arguments.
- `myfunc_int()` returns the sum of the length of its arguments.
- `sequence([const int])` returns a sequence starting from the given number or 1 if no number has been given.
- `lookup()` returns the IP address for a host name.
- `reverse_lookup()` returns the host name for an IP address. The function may be called either with a single string argument of the form `'xxx.xxx.xxx.xxx'` or with four numbers.
- `avgcost()` returns an average cost. This is an aggregate function.

A dynamically loadable file should be compiled as a sharable object file, using a command something like this:

```
shell> gcc -shared -o udf_example.so udf_example.c
```

If you are using `gcc` with `CMake` (which is how MySQL is configured), you should be able to create `udf_example.so` with a simpler command:

```
shell> make udf_example
```

After you compile a shared object containing UDFs, you must install it and tell MySQL about it. Compiling a shared object from `udf_example.c` using `gcc` directly produces a file named `udf_example.so`. Copy the shared object to the server's plugin directory and name it `udf_example.so`. This directory is given by the value of the `plugin_dir` [525] system variable.

On some systems, the `ldconfig` program that configures the dynamic linker does not recognize a shared object unless its name begins with `lib`. In this case you should rename a file such as `udf_example.so` to `libudf_example.so`.

On Windows, you can compile user-defined functions by using the following procedure:

1. Obtain a MySQL source distribution. See [Section 2.1.3, "How to Get MySQL"](#).
2. Obtain the `CMake` build utility, if necessary, from <http://www.cmake.org>. (Version 2.6 or later is required).
3. In the source tree, look in the `sql` directory. There are files named `udf_example.def` and `udf_example.c` there. Copy both files from this directory to your working directory.
4. Create a `CMake` makefile (`CMakeLists.txt`) with these contents:

```
PROJECT(udf_example)
```

```
# Path for MySQL include directory
INCLUDE_DIRECTORIES("c:/mysql/include")

ADD_DEFINITIONS("-DHAVE_DLOPEN")
ADD_LIBRARY(udf_example MODULE udf_example.c udf_example.def)
TARGET_LINK_LIBRARIES(udf_example wsock32)
```

5. Create the VC project and solution files:

```
cmake -G "<Generator>"
```

Invoking `cmake --help` shows you a list of valid Generators.

6. Create `udf_example.dll`:

```
devenv udf_example.sln /build Release
```

After the shared object file has been installed, notify `mysqld` about the new functions with the following statements. If object files have a suffix different from `.so` on your system, substitute the correct suffix throughout (for example, `.dll` on Windows).

```
mysql> CREATE FUNCTION metaphon RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE FUNCTION myfunc_double RETURNS REAL SONAME 'udf_example.so';
mysql> CREATE FUNCTION myfunc_int RETURNS INTEGER SONAME 'udf_example.so';
mysql> CREATE FUNCTION sequence RETURNS INTEGER SONAME 'udf_example.so';
mysql> CREATE FUNCTION lookup RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE FUNCTION reverse_lookup
-> RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE AGGREGATE FUNCTION avgcost
-> RETURNS REAL SONAME 'udf_example.so';
```

To delete functions, use `DROP FUNCTION`:

```
mysql> DROP FUNCTION metaphon;
mysql> DROP FUNCTION myfunc_double;
mysql> DROP FUNCTION myfunc_int;
mysql> DROP FUNCTION sequence;
mysql> DROP FUNCTION lookup;
mysql> DROP FUNCTION reverse_lookup;
mysql> DROP FUNCTION avgcost;
```

The `CREATE FUNCTION` and `DROP FUNCTION` statements update the `func` system table in the `mysql` database. The function's name, type and shared library name are saved in the table. You must have the `INSERT [755]` or `DELETE [754]` privilege for the `mysql` database to create or drop functions, respectively.

You should not use `CREATE FUNCTION` to add a function that has previously been created. If you need to reinstall a function, you should remove it with `DROP FUNCTION` and then reinstall it with `CREATE FUNCTION`. You would need to do this, for example, if you recompile a new version of your function, so that `mysqld` gets the new version. Otherwise, the server continues to use the old version.

An active function is one that has been loaded with `CREATE FUNCTION` and not removed with `DROP FUNCTION`. All active functions are reloaded each time the server starts, unless you start `mysqld` with the `--skip-grant-tables [439]` option. In this case, UDF initialization is skipped and UDFs are unavailable.

### 22.3.2.6. User-Defined Function Security Precautions

MySQL takes several measures to prevent misuse of user-defined functions.

UDF object files cannot be placed in arbitrary directories. They must be located in the server's plugin directory. This directory is given by the value of the `plugin_dir [525]` system variable.

To use `CREATE FUNCTION` or `DROP FUNCTION`, you must have the `INSERT [755]` or `DELETE [754]` privilege, respectively, for the `mysql` database. This is necessary because those statements add and delete rows from the `mysql.func` table.

UDFs should have at least one symbol defined in addition to the `xxx` symbol that corresponds to the main `xxx()` function. These auxiliary symbols correspond to the `xxx_init()`, `xxx_deinit()`, `xxx_reset()`, `xxx_clear()`, and `xxx_add()` functions. `mysqld` also supports an `--allow-suspicious-udfs` [413] option that controls whether UDFs that have only an `xxx` symbol can be loaded. By default, the option is off, to prevent attempts at loading functions from shared object files other than those containing legitimate UDFs. If you have older UDFs that contain only the `xxx` symbol and that cannot be recompiled to include an auxiliary symbol, it may be necessary to specify the `--allow-suspicious-udfs` [413] option. Otherwise, you should avoid enabling this capability.

### 22.3.3. Adding a New Native Function

To add a new native MySQL function, use the procedure described here, which requires that you use a source distribution. You cannot add native functions to a binary distribution because it is necessary to modify MySQL source code and compile MySQL from the modified source. If you migrate to another version of MySQL (for example, when a new version is released), you must repeat the procedure with the new version.

If the new native function will be referred to in statements that will be replicated to slave servers, you must ensure that every slave server also has the function available. Otherwise, replication will fail on the slaves when they attempt to invoke the function.

To add a new native function, follow these steps to modify source files in the `sql` directory:

1. Create a subclass for the function in `item_create.cc`:
  - If the function takes a fixed number of arguments, create a subclass of `Create_func_arg0`, `Create_func_arg1`, `Create_func_arg2`, or `Create_func_arg3`, respectively, depending on whether the function takes zero, one, two, or three arguments. For examples, see the `Create_func_uuid`, `Create_func_abs`, `Create_func_pow`, and `Create_func_lpad` classes.
  - If the function takes a variable number of arguments, create a subclass of `Create_native_func`. For an example, see `Create_func_concat`.
2. To provide a name by which the function can be referred to in SQL statements, register the name in `item_create.cc` by adding a line to this array:

```
static Native_func_registry func_array[]
```

You can register several names for the same function. For example, see the lines for `"LCASE"` and `"LOWER"`, which are aliases for `Create_func_lcase`.

3. In `item_func.h`, declare a class inheriting from `Item_num_func` or `Item_str_func`, depending on whether your function returns a number or a string.
4. In `item_func.cc`, add one of the following declarations, depending on whether you are defining a numeric or string function:

```
double   Item_func_newname::val()
longlong Item_func_newname::val_int()
String   *Item_func_newname::Str(String *str)
```

If you inherit your object from any of the standard items (like `Item_num_func`), you probably only have to define one of these functions and let the parent object take care of the other functions. For example, the `Item_str_func` class defines a `val()` function that executes `atof()` on the value returned by `::str()`.

5. If the function is nondeterministic, include the following statement in the item constructor to indicate that function results should not be cached:

```
current_thd->lex->safe_to_cache_query=0;
```

A function is nondeterministic if, given fixed values for its arguments, it can return different results for different invocations.

6. You should probably also define the following object function:

```
void Item_func_newname::fix_length_and_dec()
```

This function should at least calculate `max_length` based on the given arguments. `max_length` is the maximum number of characters the function may return. This function should also set `maybe_null = 0` if the main function can't return a `NULL` value. The function can check whether any of the function arguments can return `NULL` by checking the arguments' `maybe_null` variable. Look at `Item_func_mod::fix_length_and_dec` for a typical example of how to do this.

All functions must be thread-safe. In other words, do not use any global or static variables in the functions without protecting them with mutexes.

If you want to return `NULL` from `::val()`, `::val_int()`, or `::str()`, you should set `null_value` to 1 and return 0.

For `::str()` object functions, there are additional considerations to be aware of:

- The `String *str` argument provides a string buffer that may be used to hold the result. (For more information about the `String` type, take a look at the `sql_string.h` file.)
- The `::str()` function should return the string that holds the result, or `(char*) 0` if the result is `NULL`.
- All current string functions try to avoid allocating any memory unless absolutely necessary!

## 22.4. Debugging and Porting MySQL

This appendix helps you port MySQL to other operating systems. Do check the list of currently supported operating systems first. See [Section 2.1.1, “Operating Systems Supported by MySQL Community Server”](#). If you have created a new port of MySQL, please let us know so that we can list it here and on our Web site (<http://www.mysql.com/>), recommending it to other users.



### Note

If you create a new port of MySQL, you are free to copy and distribute it under the GPL license, but it does not make you a copyright holder of MySQL.

A working POSIX thread library is needed for the server.

To build MySQL from source, your system must satisfy the tool requirements listed at [Section 2.9, “Installing MySQL from Source”](#).



### Important

If you are trying to build MySQL 5.6 with `icc` on the IA64 platform, and need support for MySQL Cluster, you should first ensure that you are using `icc` version 9.1.043 or later. (For details, see Bug #21875.)

If you run into problems with a new port, you may have to do some debugging of MySQL! See [Section 22.4.1, “Debugging a MySQL Server”](#).



### Note

Before you start debugging `mysqld`, first get the test programs `mysys/thr_alarm` and `mysys/thr_lock` to work. This ensures that your thread installation has even a remote chance to work!



## 22.4.1. Debugging a MySQL Server

If you are using some functionality that is very new in MySQL, you can try to run `mysqld` with the `--skip-new` (which disables all new, potentially unsafe functionality). See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#).

If `mysqld` doesn't want to start, you should verify that you don't have any `my.cnf` files that interfere with your setup! You can check your `my.cnf` arguments with `mysqld --print-defaults` and avoid using them by starting with `mysqld --no-defaults . . . .`

If `mysqld` starts to eat up CPU or memory or if it “hangs,” you can use `mysqladmin processlist status` to find out if someone is executing a query that takes a long time. It may be a good idea to run `mysqladmin -i10 processlist status` in some window if you are experiencing performance problems or problems when new clients can't connect.

The command `mysqladmin debug` dumps some information about locks in use, used memory and query usage to the MySQL log file. This may help solve some problems. This command also provides some useful information even if you haven't compiled MySQL for debugging!

If the problem is that some tables are getting slower and slower you should try to optimize the table with `OPTIMIZE TABLE` or `myisamchk`. See [Chapter 5, \*MySQL Server Administration\*](#). You should also check the slow queries with `EXPLAIN`.

You should also read the OS-specific section in this manual for problems that may be unique to your environment. See [Section 2.1, “General Installation Guidance”](#).

### 22.4.1.1. Compiling MySQL for Debugging

If you have some very specific problem, you can always try to debug MySQL. To do this you must configure MySQL with the `-DWITH_DEBUG=1` [136] option. You can check whether MySQL was compiled with debugging by doing: `mysqld --help`. If the `--debug` [418] flag is listed with the options then you have debugging enabled. `mysqladmin ver` also lists the `mysqld` version as `mysql . . . --debug` in this case.

If `mysqld` stops crashing when you compile it with `-DWITH_DEBUG=1` [136], you probably have found a compiler bug or a timing bug within MySQL. In this case, you can try to add `-g` to the `CFLAGS` and `CXXFLAGS` environment variables and not use `-DWITH_DEBUG=1` [136]. If `mysqld` dies, you can at least attach to it with `gdb` or use `gdb` on the core file to find out what happened.

When you configure MySQL for debugging you automatically enable a lot of extra safety check functions that monitor the health of `mysqld`. If they find something “unexpected,” an entry is written to `stderr`, which `mysqld_safe` directs to the error log! This also means that if you are having some unexpected problems with MySQL and are using a source distribution, the first thing you should do is to configure MySQL for debugging! (The second thing is to send mail to a MySQL mailing list and ask for help. See [Section 1.6.1, “MySQL Mailing Lists”](#). If you believe that you have found a bug, please use the instructions at [Section 1.7, “How to Report Bugs or Problems”](#).)

In the Windows MySQL distribution, `mysqld.exe` is by default compiled with support for trace files.

### 22.4.1.2. Creating Trace Files

If the `mysqld` server doesn't start or if you can cause it to crash quickly, you can try to create a trace file to find the problem.

To do this, you must have a `mysqld` that has been compiled with debugging support. You can check this by executing `mysqld -V`. If the version number ends with `-debug`, it is compiled with support for trace files. (On Windows, the debugging server is named `mysqld-debug` rather than `mysqld` as of MySQL 4.1.)

Start the `mysqld` server with a trace log in `/tmp/mysqld.trace` on Unix or `C:\mysqld.trace` on Windows:



```
shell> mysqld --debug
```

On Windows, you should also use the `--standalone` [440] flag to not start `mysqld` as a service. In a console window, use this command:

```
C:\> mysqld-debug --debug --standalone
```

After this, you can use the `mysql.exe` command-line tool in a second console window to reproduce the problem. You can stop the `mysqld` server with `mysqladmin shutdown`.

The trace file can become **very large**! To generate a smaller trace file, you can use debugging options something like this:

```
mysqld --debug=d,info,error,query,general,where:0,/tmp/mysql.trace
```

This only prints information with the most interesting tags to the trace file.

If you make a bug report about this, please only send the lines from the trace file to the appropriate mailing list where something seems to go wrong! If you can't locate the wrong place, you can open a bug report and upload the trace file to the report, so that a MySQL developer can take a look at it. For instructions, see [Section 1.7, "How to Report Bugs or Problems"](#).

The trace file is made with the **DEBUG** package by Fred Fish. See [Section 22.4.3, "The DEBUG Package"](#).

### 22.4.1.3. Using `pdb` to create a Windows crashdump

Program Database files (extension `pdb`) are included in the Noinstall distribution of MySQL. These files provide information for debugging your MySQL installation in the event of a problem.

The PDB file contains more detailed information about `mysqld` and other tools that enables more detailed trace and dump files to be created. You can use these with Dr Watson, `WinDbg` and Visual Studio to debug `mysqld`.

For more information on PDB files, see [Microsoft Knowledge Base Article 121366](#). For more information on the debugging options available, see [Debugging Tools for Windows](#).

Dr Watson is installed with all Windows distributions, but if you have installed Windows development tools, Dr Watson may have been replaced with `WinDbg`, the debugger included with Visual Studio, or the debugging tools provided with Borland or Delphi.

To generate a crash file using Dr Watson, follow these steps:

1. Start Dr Watson by running `drwtsn32.exe` interactively using the `-i` option:

```
C:\> drwtsn32 -i
```

2. Set the **Log File Path** to the directory where you want to store trace files.
3. Make sure **Dump All Thread Contexts** and **Append To Existing Log File**.
4. Uncheck **Dump Symbol Table**, **Visual Notification**, **Sound Notification** and **Create Crash Dump File**.
5. Set the **Number of Instructions** to a suitable value to capture enough calls in the stacktrace. A value of at 25 should be enough.

Note that the file generated can become very large.

### 22.4.1.4. Debugging `mysqld` under `gdb`

On most systems you can also start `mysqld` from `gdb` to get more information if `mysqld` crashes.

With some older `gdb` versions on Linux you must use `run --one-thread` if you want to be able to debug `mysqld` threads. In this case, you can only have one thread active at a time. It is best to upgrade to `gdb` 5.1 because thread debugging works much better with this version!

NPTL threads (the new thread library on Linux) may cause problems while running `mysqld` under `gdb`. Some symptoms are:

- `mysqld` hangs during startup (before it writes `ready for connections`).
- `mysqld` crashes during a `pthread_mutex_lock()` or `pthread_mutex_unlock()` call.

In this case, you should set the following environment variable in the shell before starting `gdb`:

```
LD_ASSUME_KERNEL=2.4.1
export LD_ASSUME_KERNEL
```

When running `mysqld` under `gdb`, you should disable the stack trace with `--skip-stack-trace` [441] to be able to catch segfaults within `gdb`.

In MySQL 4.0.14 and above you should use the `--gdb` [423] option to `mysqld`. This installs an interrupt handler for `SIGINT` (needed to stop `mysqld` with `^C` to set breakpoints) and disable stack tracing and core file handling.

It is very hard to debug MySQL under `gdb` if you do a lot of new connections the whole time as `gdb` doesn't free the memory for old threads. You can avoid this problem by starting `mysqld` with `thread_cache_size` [552] set to a value equal to `max_connections` [503] + 1. In most cases just using `--thread_cache_size=5` [552] helps a lot!

If you want to get a core dump on Linux if `mysqld` dies with a `SIGSEGV` signal, you can start `mysqld` with the `--core-file` [418] option. This core file can be used to make a backtrace that may help you find out why `mysqld` died:

```
shell> gdb mysqld core
gdb> backtrace full
gdb> quit
```

See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#).

If you are using `gdb` 4.17.x or above on Linux, you should install a `.gdb` file, with the following information, in your current directory:

```
set print sevenbit off
handle SIGUSR1 nostop noprint
handle SIGUSR2 nostop noprint
handle SIGWAITING nostop noprint
handle SIGLWP nostop noprint
handle SIGPIPE nostop
handle SIGALRM nostop
handle SIGHUP nostop
handle SIGTERM nostop noprint
```

If you have problems debugging threads with `gdb`, you should download `gdb` 5.x and try this instead. The new `gdb` version has very improved thread handling!

Here is an example how to debug `mysqld`:

```
shell> gdb /usr/local/libexec/mysqld
gdb> run
...
backtrace full # Do this when mysqld crashes
```

Include the preceding output in a bug report, which you can file using the instructions in [Section 1.7, “How to Report Bugs or Problems”](#).

If `mysqld` hangs, you can try to use some system tools like `strace` or `/usr/proc/bin/pstack` to examine where `mysqld` has hung.

```
strace /tmp/log libexec/mysqld
```

If you are using the Perl [DBI](#) interface, you can turn on debugging information by using the `trace` method or by setting the `DBI_TRACE` environment variable.

### 22.4.1.5. Using a Stack Trace

On some operating systems, the error log contains a stack trace if `mysqld` dies unexpectedly. You can use this to find out where (and maybe why) `mysqld` died. See [Section 5.2.2, “The Error Log”](#). To get a stack trace, you must not compile `mysqld` with the `-fomit-frame-pointer` option to `gcc`. See [Section 22.4.1.1, “Compiling MySQL for Debugging”](#).

A stack trace in the error log looks something like this:

```
mysqld got signal 11;
Attempting backtrace. You can use the following information
to find out where mysqld died. If you see no messages after
this, something went terribly wrong...

stack_bottom = 0x41fd0110 thread_stack 0x40000
mysqld(my_print_stacktrace+0x32)[0x9da402]
mysqld(handle_segfault+0x28a)[0x6648e9]
/lib/libpthread.so.0[0x7f1a5af000f0]
/lib/libc.so.6(strcmp+0x2)[0x7f1a5a10f0f2]
mysqld(_Z21check_change_passwordP3THDPKcS2_Pcj+0x7c)[0x7412cb]
mysqld(_ZN16set_var_password5checkEP3THD+0xd0)[0x688354]
mysqld(_Z17sql_set_variablesP3THDP4ListI12set_var_baseE+0x68)[0x688494]
mysqld(_Z21mysql_execute_commandP3THD+0x41a0)[0x67a170]
mysqld(_Z11mysql_parseP3THDPKcjPS2_+0x282)[0x67f0ad]
mysqld(_Z16dispatch_command19enum_server_commandP3THDPcj+0xbb7)[0x67fdf8]
mysqld(_Z10do_commandP3THD+0x24d)[0x6811b6]
mysqld(handle_one_connection+0x11c)[0x66e05e]
```

If resolution of function names for the trace fails, the trace contains less information:

```
mysqld got signal 11;
Attempting backtrace. You can use the following information
to find out where mysqld died. If you see no messages after
this, something went terribly wrong...

stack_bottom = 0x41fd0110 thread_stack 0x40000
[0x9da402]
[0x6648e9]
[0x7f1a5af000f0]
[0x7f1a5a10f0f2]
[0x7412cb]
[0x688354]
[0x688494]
[0x67a170]
[0x67f0ad]
[0x67fdf8]
[0x6811b6]
[0x66e05e]
```

In the latter case, you can use the `resolve_stack_dump` utility to determine where `mysqld` died by using the following procedure:

1. Copy the numbers from the stack trace to a file, for example `mysqld.stack`. The numbers should not include the surrounding square brackets:

```
0x9da402
0x6648e9
0x7f1a5af000f0
0x7f1a5a10f0f2
0x7412cb
0x688354
0x688494
```

```
0x67a170
0x67f0ad
0x67fdf8
0x6811b6
0x66e05e
```

2. Make a symbol file for the `mysqld` server:

```
shell> nm -n libexec/mysqld > /tmp/mysqld.sym
```

If `mysqld` is not linked statically, use the following command instead:

```
shell> nm -D -n libexec/mysqld > /tmp/mysqld.sym
```

If you want to decode C++ symbols, use the `--demangle`, if available, to `nm`. If your version of `nm` does not have this option, you will need to use the `c++filt` command after the stack dump has been produced to demangle the C++ names.

3. Execute the following command:

```
shell> resolve_stack_dump -s /tmp/mysqld.sym -n mysqld.stack
```

If you were not able to include demangled C++ names in your symbol file, process the `resolve_stack_dump` output using `c++filt`:

```
shell> resolve_stack_dump -s /tmp/mysqld.sym -n mysqld.stack | c++filt
```

This prints out where `mysqld` died. If that does not help you find out why `mysqld` died, you should create a bug report and include the output from the preceding command with the bug report.

However, in most cases it does not help us to have just a stack trace to find the reason for the problem. To be able to locate the bug or provide a workaround, in most cases we need to know the statement that killed `mysqld` and preferably a test case so that we can repeat the problem! See [Section 1.7, “How to Report Bugs or Problems”](#).

Newer versions of `glibc` stack trace functions also print the address as relative to the object. On `glibc`-based systems (Linux), the trace for a crash within a plugin looks something like:

```
plugin/auth/auth_test_plugin.so(+0x9a6)[0x7ff4d11c29a6]
```

To translate the relative address (`+0x9a6`) into a file name and line number, use this command:

```
shell> addr2line -fie auth_test_plugin.so 0x9a6
auth_test_plugin
mysql-trunk/plugin/auth/test_plugin.c:65
```

The `addr2line` utility is part of the `binutils` package on Linux.

On Solaris, the procedure is similar. The Solaris `printstack()` already prints relative addresses:

```
plugin/auth/auth_test_plugin.so:0x1510
```

To translate, use this command:

```
shell> gaddr2line -fie auth_test_plugin.so 0x1510
mysql-trunk/plugin/auth/test_plugin.c:88
```

Windows already prints the address, function name and line:

```
000007FEF07E10A4 auth_test_plugin.dll!auth_test_plugin()[test_plugin.c:72]
```

### 22.4.1.6. Using Server Logs to Find Causes of Errors in `mysqld`

Note that before starting `mysqld` with the general query log enabled, you should check all your tables with `myisamchk`. See [Chapter 5, MySQL Server Administration](#).

If `mysqld` dies or hangs, you should start `mysqld` with the general query log enabled. See [Section 5.2.3, “The General Query Log”](#). When `mysqld` dies again, you can examine the end of the log file for the query that killed `mysqld`.

If you use the default general query log file, the log is stored in the database directory as `host_name.log`. In most cases it is the last query in the log file that killed `mysqld`, but if possible you should verify this by restarting `mysqld` and executing the found query from the `mysql` command-line tools. If this works, you should also test all complicated queries that didn't complete.

You can also try the command `EXPLAIN` on all `SELECT` statements that takes a long time to ensure that `mysqld` is using indexes properly. See [Section 13.8.2, “EXPLAIN Syntax”](#).

You can find the queries that take a long time to execute by starting `mysqld` with the slow query log enabled. See [Section 5.2.5, “The Slow Query Log”](#).

If you find the text `mysqld restarted` in the error log file (normally named `hostname.err`) you probably have found a query that causes `mysqld` to fail. If this happens, you should check all your tables with `myisamchk` (see [Chapter 5, MySQL Server Administration](#)), and test the queries in the MySQL log files to see whether one fails. If you find such a query, try first upgrading to the newest MySQL version. If this doesn't help and you can't find anything in the `mysql` mail archive, you should report the bug to a MySQL mailing list. The mailing lists are described at <http://lists.mysql.com/>, which also has links to online list archives.

If you have started `mysqld` with `--myisam-recover-options` [432], MySQL automatically checks and tries to repair `MyISAM` tables if they are marked as 'not closed properly' or 'crashed'. If this happens, MySQL writes an entry in the `hostname.err` file `'Warning: Checking table ...'` which is followed by `Warning: Repairing table` if the table needs to be repaired. If you get a lot of these errors, without `mysqld` having died unexpectedly just before, then something is wrong and needs to be investigated further. See [Section 5.1.3, “Server Command Options”](#).

In MySQL 5.6, when the server detects `MyISAM` table corruption, it writes additional information to the error log, such as the name and line number of the source file, and the list of threads accessing the table. Example: `Got an error from thread_id=1, mi_dynrec.c:368`. This is useful information to include in bug reports.

It is not a good sign if `mysqld` did die unexpectedly, but in this case, you should not investigate the `Checking table...` messages, but instead try to find out why `mysqld` died.

### 22.4.1.7. Making a Test Case If You Experience Table Corruption

If you get corrupted tables or if `mysqld` always fails after some update commands, you can test whether this bug is reproducible by doing the following:

- Take down the MySQL daemon (with `mysqladmin shutdown`).
- Make a backup of the tables (to guard against the very unlikely case that the repair does something bad).
- Check all tables with `myisamchk -s database/*.MYI`. Repair any wrong tables with `myisamchk -r database/table.MYI`.
- Make a second backup of the tables.
- Remove (or move away) any old log files from the MySQL data directory if you need more space.
- Start `mysqld` with the binary log enabled. If you want to find a query that crashes `mysqld`, you should start the server with both the general query log enabled as well. See [Section 5.2.3, “The General Query Log”](#), and [Section 5.2.4, “The Binary Log”](#).
- When you have gotten a crashed table, stop the `mysqld server`.

- Restore the backup.
- Restart the `mysqld` server **without** the binary log enabled.
- Re-execute the commands with `mysqlbinlog binary-log-file | mysql`. The binary log is saved in the MySQL database directory with the name `hostname-bin.NNNNNN`.
- If the tables are corrupted again or you can get `mysqld` to die with the above command, you have found a reproducible bug that should be easy to fix! FTP the tables and the binary log to our bugs database using the instructions given in [Section 1.7, "How to Report Bugs or Problems"](#). If you are a support customer, you can use the MySQL Customer Support Center <https://support.mysql.com/> to alert the MySQL team about the problem and have it fixed as soon as possible.

You can also use the script `mysql_find_rows` to just execute some of the update statements if you want to narrow down the problem.

## 22.4.2. Debugging a MySQL Client

To be able to debug a MySQL client with the integrated debug package, you should configure MySQL with `-DWITH_DEBUG=1` [136]. See [Section 2.9.4, "MySQL Source-Configuration Options"](#).

Before running a client, you should set the `MYSQL_DEBUG` environment variable:

```
shell> MYSQL_DEBUG=d:t:O,/tmp/client.trace
shell> export MYSQL_DEBUG
```

This causes clients to generate a trace file in `/tmp/client.trace`.

If you have problems with your own client code, you should attempt to connect to the server and run your query using a client that is known to work. Do this by running `mysql` in debugging mode (assuming that you have compiled MySQL with debugging on):

```
shell> mysql --debug=d:t:O,/tmp/client.trace
```

This provides useful information in case you mail a bug report. See [Section 1.7, "How to Report Bugs or Problems"](#).

If your client crashes at some 'legal' looking code, you should check that your `mysql.h` include file matches your MySQL library file. A very common mistake is to use an old `mysql.h` file from an old MySQL installation with new MySQL library.

## 22.4.3. The DBUG Package

The MySQL server and most MySQL clients are compiled with the DBUG package originally created by Fred Fish. When you have configured MySQL for debugging, this package makes it possible to get a trace file of what the program is debugging. See [Section 22.4.1.2, "Creating Trace Files"](#).

This section summarizes the argument values that you can specify in debug options on the command line for MySQL programs that have been built with debugging support. For more information about programming with the DBUG package, see the DBUG manual in the `debug` directory of MySQL source distributions. It is best to use a recent distribution to get the most updated DBUG manual.

You use the debug package by invoking a program with the `--debug="..."` or the `-#...` option.

Most MySQL programs have a default debug string that is used if you don't specify an option to `--debug`. The default trace file is usually `/tmp/program_name.trace` on Unix and `\program_name.trace` on Windows.

The debug control string is a sequence of colon-separated fields as follows:

```
<field_1>:<field_2>:...:<field_N>
```

Each field consists of a mandatory flag character followed by an optional “,” and comma-separated list of modifiers:

```
flag[,modifier,modifier,...,modifier]
```

The following table shows the currently recognized flag characters.

Flag	Description
d	Enable output from <code>DEBUG_&lt;N&gt;</code> macros for the current state. May be followed by a list of keywords which selects output only for the <code>DEBUG</code> macros with that keyword. An empty list of keywords implies output for all macros.
D	Delay after each debugger output line. The argument is the number of tenths of seconds to delay, subject to machine capabilities. For example, <code>-#D,20</code> specifies a delay of two seconds.
f	Limit debugging, tracing, and profiling to the list of named functions. Note that a null list disables all functions. The appropriate <code>d</code> or <code>t</code> flags must still be given; this flag only limits their actions if they are enabled.
F	Identify the source file name for each line of debug or trace output.
i	Identify the process with the PID or thread ID for each line of debug or trace output.
g	Enable profiling. Create a file called <code>ddebugmon.out</code> containing information that can be used to profile the program. May be followed by a list of keywords that select profiling only for the functions in that list. A null list implies that all functions are considered.
L	Identify the source file line number for each line of debug or trace output.
n	Print the current function nesting depth for each line of debug or trace output.
N	Number each line of debug output.
o	Redirect the debugger output stream to the specified file. The default output is <code>stderr</code> .
O	Like <code>o</code> , but the file is really flushed between each write. When needed, the file is closed and reopened between each write.
p	Limit debugger actions to specified processes. A process must be identified with the <code>DEBUG_PROCESS</code> macro and match one in the list for debugger actions to occur.
P	Print the current process name for each line of debug or trace output.
r	When pushing a new state, do not inherit the previous state's function nesting level. Useful when the output is to start at the left margin.
S	Do function <code>_sanity(_file_,_line_)</code> at each debugged function until <code>_sanity()</code> returns something that differs from 0.
t	Enable function call/exit trace lines. May be followed by a list (containing only one modifier) giving a numeric maximum trace level, beyond which no output occurs for either debugging or tracing macros. The default is a compile time option.

Some examples of debug control strings that might appear on a shell command line (the `-#` is typically used to introduce a control string to an application program) are:

```
-#d:t
-#d:f,main,subr1:F:L:t,20
-#d,input,output,files:n
-#d:t:i:0,\\mysqld.trace
```

In MySQL, common tags to print (with the `d` option) are `enter`, `exit`, `error`, `warning`, `info`, and `loop`.





---

## Chapter 23. MySQL Enterprise Edition

### Table of Contents

23.1. MySQL Enterprise Monitor .....	3553
23.2. MySQL Enterprise Backup .....	3554
23.3. MySQL Enterprise Security .....	3554
23.4. MySQL Enterprise Audit .....	3555
23.5. MySQL Thread Pool .....	3555

MySQL Enterprise Edition is a commercial product. Like MySQL Community Edition, MySQL Enterprise Edition includes MySQL Server, a fully integrated transaction-safe, ACID-compliant database with full commit, rollback, crash-recovery, and row-level locking capabilities. In addition, MySQL Enterprise Edition includes the following components designed to provide monitoring and online backup, as well as improved security and scalability:

The following sections briefly discuss each of these components and indicate where to find more detailed information. To learn more about commercial products, see <http://www.mysql.com/products/>.

- [MySQL Enterprise Monitor](#)
- [MySQL Enterprise Backup](#)
- [MySQL Enterprise Security](#)
- [MySQL Enterprise Audit](#)
- [MySQL Thread Pool](#)

### 23.1. MySQL Enterprise Monitor

MySQL Enterprise Monitor is an enterprise monitoring system for MySQL that keeps an eye on your MySQL servers, notifies you of potential issues and problems, and advises you how to fix the issues. MySQL Enterprise Monitor can monitor all kinds of configurations, from a single MySQL server that is important to your business, all the way up to a huge farm of MySQL servers powering a busy web site.

The following discussion briefly summarizes the basic components that make up the MySQL Enterprise Monitor product. For more information, see the MySQL Enterprise Monitor manual, available at <http://dev.mysql.com/doc/mysql-monitor/en/>.

MySQL Enterprise Monitor components can be installed in various configurations depending on your database and network topology, to give you the best combination of reliable and responsive monitoring data, with minimal overhead on the database server machines. A typical MySQL Enterprise Monitor installation consists of:

- One or more MySQL servers to monitor. MySQL Enterprise Monitor can monitor both Community and Enterprise MySQL server releases.
- A MySQL Enterprise Agent for each monitored MySQL server.
- A single MySQL Enterprise Service Manager, which collates information from the agents and provides the user interface to the collected data.

MySQL Enterprise Monitor is designed to monitor one or more MySQL servers. The monitoring information is collected by using an agent, *MySQL Enterprise Agent*. The agent communicates with the MySQL server that it monitors, collecting variables, status and health information, and sending this information to the MySQL Enterprise Service Manager. If you have multiple MySQL servers, then you have multiple MySQL Enterprise Agent processes monitoring each MySQL server.

The information collected by the agent about each MySQL server you are monitoring is sent to the *MySQL Enterprise Service Manager*. This server collates all of the information from the agents. As it collates the information sent by the agents, the MySQL Enterprise Service Manager continually tests the collected data, comparing the status of the server to reasonable values. When thresholds are reached, the server can trigger an event (including an alarm and notification) to highlight a potential issue, such as low memory, high CPU usage, or more complex conditions such as insufficient buffer sizes and status information. We call each test, with its associated threshold value, a *rule*.

These rules, and the alarms and notifications, are each known as a *MySQL Enterprise Advisor*. Advisors form a critical part of the MySQL Enterprise Service Manager, as they provide warning information and troubleshooting advice about potential problems.

The MySQL Enterprise Service Manager includes a web server, and you interact with it through any web browser. This interface, the MySQL Enterprise Dashboard, displays all of the information collected by the agents, and lets you view all of your servers and their current status as a group or individually. You control and configure all aspects of the service using the MySQL Enterprise Dashboard.

The information supplied by the MySQL Enterprise Agent processes also includes statistical and query information, which you can view in the form of graphs. For example, you can view aspects such as server load, query numbers, or index usage information as a graph over time. The graph lets you pinpoint problems or potential issues on your server, and can help diagnose the impact from database or external problems (such as external system or network failure) by examining the data from a specific time interval.

The MySQL Enterprise Agent can also be configured to collect detailed information about the queries executed on your server, including the row counts and performance times for executing each query. You can correlate the detailed query data with the graphical information to identify which queries were executing when you experienced a particularly high load, index or other issue. The query data is supported by a system called Query Analyzer, and the data can be presented in different ways depending on your needs.

## 23.2. MySQL Enterprise Backup

MySQL Enterprise Backup performs hot backup operations for MySQL databases. The product is architected for efficient and reliable backups of tables created by the InnoDB storage engine. For completeness, it can also back up tables from MyISAM and other storage engines.

The following discussion briefly summarizes MySQL Enterprise Backup. For more information, see the MySQL Enterprise Backup manual, available at <http://dev.mysql.com/doc/mysql-enterprise-backup/en/>.

Hot backups are performed while the database is running and applications are reading and writing to it. This type of backup does not block normal database operations, and it captures even changes that occur while the backup is happening. For these reasons, hot backups are desirable when your database “grows up” -- when the data is large enough that the backup takes significant time, and when your data is important enough to your business that you must capture every last change, without taking your application, web site, or web service offline.

MySQL Enterprise Backup does a hot backup of all tables that use the InnoDB storage engine. For tables using MyISAM or other non-InnoDB storage engines, it does a “warm” backup, where the database continues to run, but those tables cannot be modified while being backed up. For efficient backup operations, you can designate InnoDB as the default storage engine for new tables, or convert existing tables to use the InnoDB storage engine.

## 23.3. MySQL Enterprise Security

MySQL Enterprise Edition provides plugins that implement authentication using external services:

- MySQL Enterprise Edition includes an authentication plugin that enables MySQL Server to use PAM (Pluggable Authentication Modules) to authenticate MySQL users. PAM enables a system to use a

standard interface to access various kinds of authentication methods, such as Unix passwords or an LDAP directory.

- MySQL Enterprise Edition includes an authentication plugin that performs external authentication on Windows, enabling MySQL Server to use native Windows services to authenticate client connections. Users who have logged in to Windows can connect from MySQL client programs to the server based on the information in their environment without specifying an additional password.

For more information, see [Section 6.3.7.3, “The PAM Authentication Plugin”](#), and [Section 6.3.7.4, “The Windows Native Authentication Plugin”](#).

## 23.4. MySQL Enterprise Audit

MySQL Enterprise Edition includes MySQL Enterprise Audit, implemented using a server plugin. MySQL Enterprise Audit uses the open MySQL Audit API to enable standard, policy-based monitoring and logging of connection and query activity executed on specific MySQL servers. Designed to meet the Oracle audit specification, MySQL Enterprise Audit provides an out of box, easy to use auditing and compliance solution for applications that are governed by both internal and external regulatory guidelines.

When installed, the audit plugin enables MySQL Server to produce a log file containing an audit record of server activity. The log contents include when clients connect and disconnect, and what actions they perform while connected, such as which databases and tables they access.

For more information, see [Section 6.3.11, “MySQL Enterprise Audit Log Plugin”](#).

## 23.5. MySQL Thread Pool

MySQL Enterprise Edition includes the MySQL Thread Pool, implemented using a server plugin. The default thread-handling model in MySQL Server executes statements using one thread per client connection. As more clients connect to the server and execute statements, overall performance degrades. In MySQL Enterprise Edition, a thread pool plugin provides an alternative thread-handling model designed to reduce overhead and improve performance. The plugin implements a thread pool that increases server performance by efficiently managing statement execution threads for large numbers of client connections.

For more information, see [Section 8.11.6, “The Thread Pool Plugin”](#).



---

## Chapter 24. MySQL Workbench

MySQL Workbench provides a graphical tool for working with MySQL Servers and databases. MySQL Workbench fully supports MySQL Server versions 5.1 and above. It is also compatible with MySQL Server 5.0, but not every feature of 5.0 may be supported. It does not support MySQL Server versions 4.x.

The following discussion briefly describes MySQL Workbench capabilities. For more information, see the MySQL Workbench manual, available at <http://dev.mysql.com/doc/workbench/en/>.

MySQL Workbench provides three main areas of functionality:

- **SQL Development:** Enables you to create and manage connections to database servers. As well as enabling you to configure connection parameters, MySQL Workbench provides the capability to execute SQL queries on the database connections using the built-in SQL Editor. This functionality replaces that previously provided by the Query Browser standalone application.
- **Data Modeling:** Enables you to create models of your database schema graphically, reverse and forward engineer between a schema and a live database, and edit all aspects of your database using the comprehensive Table Editor. The Table Editor provides easy-to-use facilities for editing Tables, Columns, Indexes, Triggers, Partitioning, Options, Inserts and Privileges, Routines and Views.
- **Server Administration:** Enables you to create and administer server instances. This functionality replaces that previously provided by the MySQL Administrator standalone application.

MySQL Workbench is available in two editions, the Community Edition and the Standard Edition. The Community Edition is available free of charge. The Standard Edition provides additional Enterprise features, such as database documentation generation, at low cost.



# Chapter 25. MySQL for Excel

## Table of Contents

25.1. Installing and Configuring .....	3559
25.2. Edit MySQL Data .....	3560
25.3. Import MySQL Data .....	3561
25.4. Append MySQL Data .....	3562
25.5. Export MySQL Data .....	3565
25.6. MySQL for Excel FAQ .....	3568

MySQL for Excel enables you to work with a MySQL database from within Microsoft Excel. MySQL data can be imported into Excel, Excel data can be exported into MySQL as a new table or appended to a current table, and MySQL for Excel enables you to edit the MySQL data directly from within Excel.

Visit the [MySQL for Excel forum](#) for additional MySQL for Excel help and support.

For release notes detailing the changes in each release of MySQL for Excel, see [MySQL for Excel Release Notes](#).

## 25.1. Installing and Configuring

MySQL for Excel is a product for Microsoft Windows, and it is installed with MySQL Installer. And typically you will not be required to install or configure additional tools to use MySQL for Excel.



### Note

To install, download and execute the [MySQL Installer](#). Select the MySQL for Excel product and then proceed with the installation. See the [MySQL Installer manual](#) for additional details.


## MySQL for Excel Requirements

The MySQL Installer installation process will check if these requirements are met, or notify you if further action is required before proceeding with the installation.

- .NET Framework 4.0 (Client or Full Profile).
- Microsoft Office Excel 2007 or greater, for Microsoft Windows.
- Visual Studio Tools for Office 4.0, and MySQL Installer may install this for you.
- An available MySQL Server connection.

MySQL for Excel is loaded and executed by selecting the **Data** menu tab in Excel, and then choosing the "MySQL for Excel" Database icon. This opens a new Excel sidebar with the available MySQL for Excel options. The navigation bar with the MySQL for Excel icon is shown in the following screenshot:

Figure 25.1. The MySQL for Excel navigation bar



## 25.2. Edit MySQL Data

MySQL for Excel enables you to load and edit MySQL data directly from Microsoft Excel. Changes are immediately committed if the **Auto-Commit** option is enabled, or done manually by pressing **Commit Changes**.

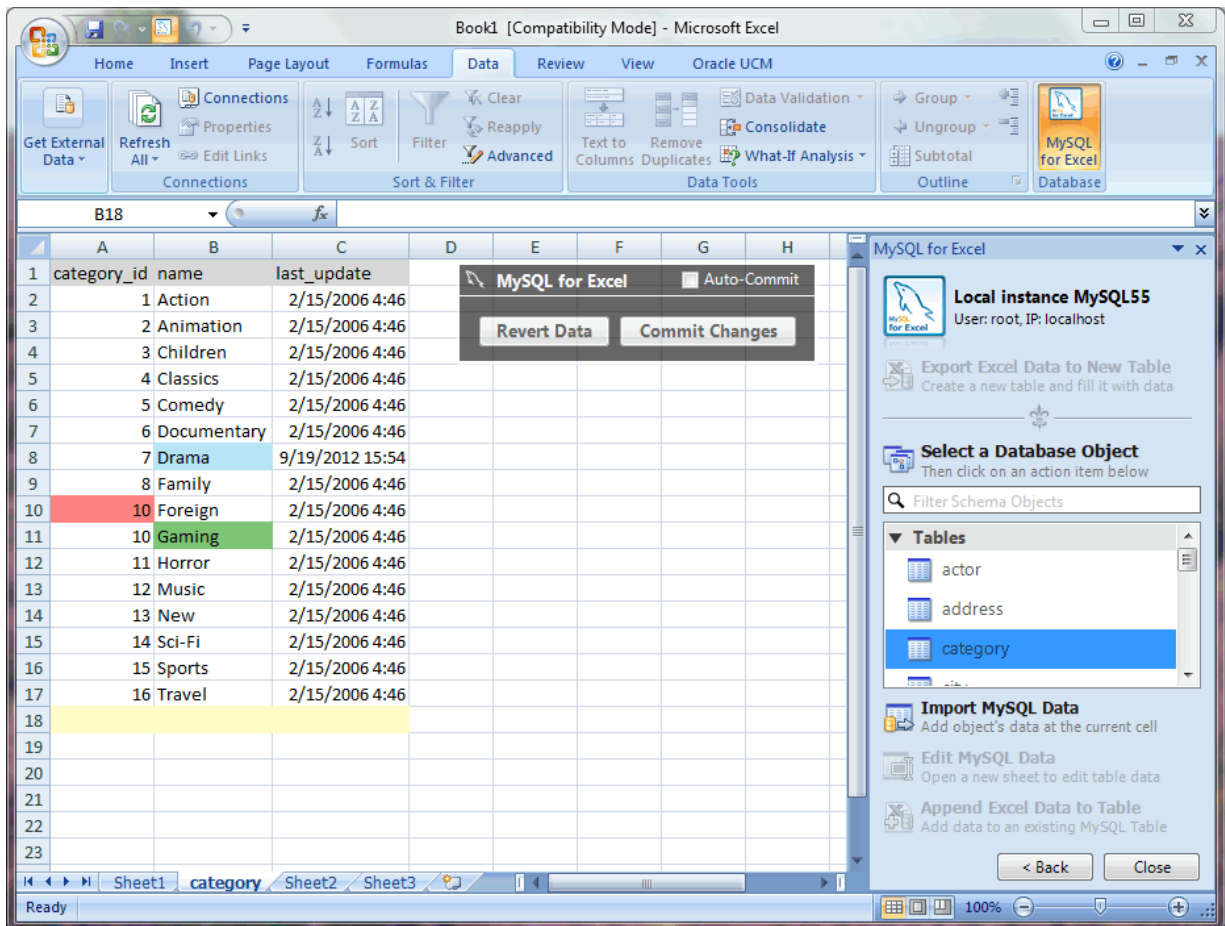


**Note**

The **Edit MySQL Data** feature exists as of MySQL for Excel 1.1.0.

The example below uses the `category` table of the example `sakila` database, but the screen will look the same for any table. Within MySQL for Excel, **Open a MySQL Connection**, click the `sakila` schema, **Next**, select the `category` table, click **Edit MySQL Data**, then choose **Import** to import the data into a new Microsoft Excel worksheet for editing.

**Figure 25.2. Editing table data with MySQL for Excel**



The background color represents the status of each cell, and there are four distinct colors that are used while editing table data:

**Table 25.1. Background cell colors**

Color	Description
White	Default color for all cells. This is either the original data, or the data after <b>Refresh from DB</b> is clicked.
Blue	Cells that were committed with success.



Color	Description
Green	Cells that were modified but have not yet been committed.
Red	Cells that generated an error when a commit was attempted. An error dialog is also displayed while the commit is attempted.
Yellow	Cells that accept new data. Data entered here is inserted into the MySQL table.

In our example, the blue "Drama" field was changed and then committed first, then the green "Gaming" field was changed but not committed, and then **Auto-Commit** was enabled before changing the "9" to a "10" in column 10, which generated an error because this commit would have added a duplicate value as primary key.

## 25.3. Import MySQL Data

Data can be imported from MySQL into a Microsoft Excel spreadsheet by using the **Import MySQL Data** option after selecting either a table, view, or procedure to import.

### Choosing columns to import

By default, all columns are selected and will be imported. Specific columns may be selected (or unselected) using the standard Microsoft Windows method of either **Control + Mouse click** to toggle the selection of individual columns, or **Shift + Mouse click** to select a range of columns.

The background color of a column shows the status of each column. The color white means that the column has been selected, and therefore it will be imported. Conversely, a gray background means that the column will not be imported.

Right-clicking anywhere in the preview grid opens a context-menu with either a **Select None** or **Select All** option, depending on the current status.

### Importing a table

The dialog while importing a table includes the following options:

- **Include Column Names as Headers**: Enabled by default, this inserts the column names at the top of the Microsoft Excel spreadsheet as a "headers" row.
- **Limit to \_\_\_ Rows and Start with Row \_\_\_**: Disabled by default, this limits the range of imported data. The **Limit to** option defaults to 1, and defines the number of rows to import. The **Start with Row** option defaults to 1 (the first row), and defines where the import begins. Each option has a maximum value of COUNT(rows) in the table.

Importing a table displays a dialog similar to the following:

**Figure 25.3. Importing table data with MySQL for Excel**

### Importing a view or procedure

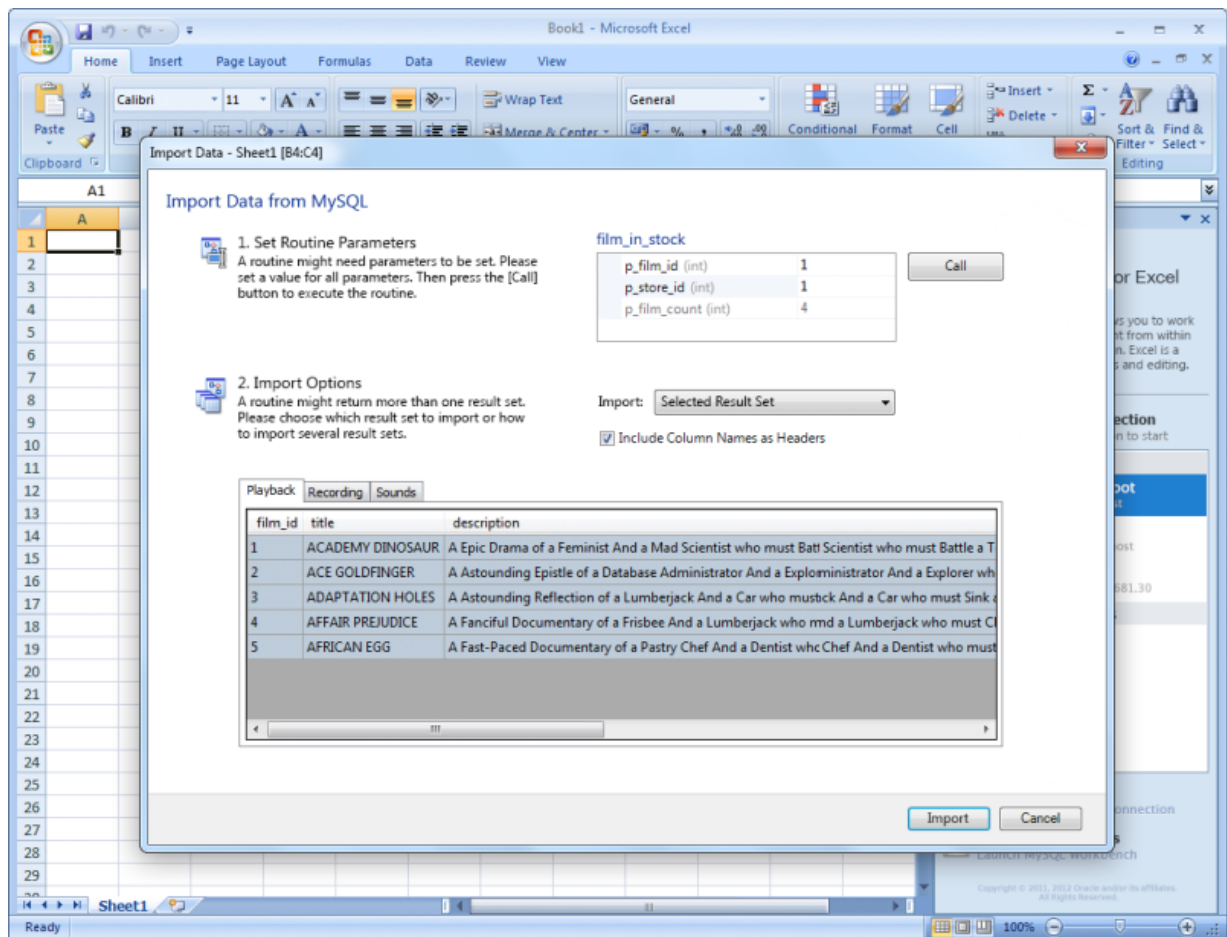
Importing a view or procedure displays a similar dialogue, but with the following options:

- **Include Column Names as Headers**: Enabled by default, this will insert the column names at the top of the Excel spreadsheet as a "headers" row.
- **Import**: Because a procedure might return multiple result sets, the import options include:
  - **Selected Result Set**: Imports the selected tab sheet. This is the default behavior.

- **All Result Sets - Arranged Horizontally:** Imports all result sets into the Excel Worksheet horizontally, and inserts one empty column between each result set.
- **All Result Sets - Arranged Vertically:** Imports all result sets into the Excel Worksheet vertically, and inserts one empty row between each result set.

For example, a dialogue like the following is displayed after importing a procedure and pressing the **Call** button to invoke the stored procedure:

**Figure 25.4. Importing called stored procedure data with MySQL for Excel**



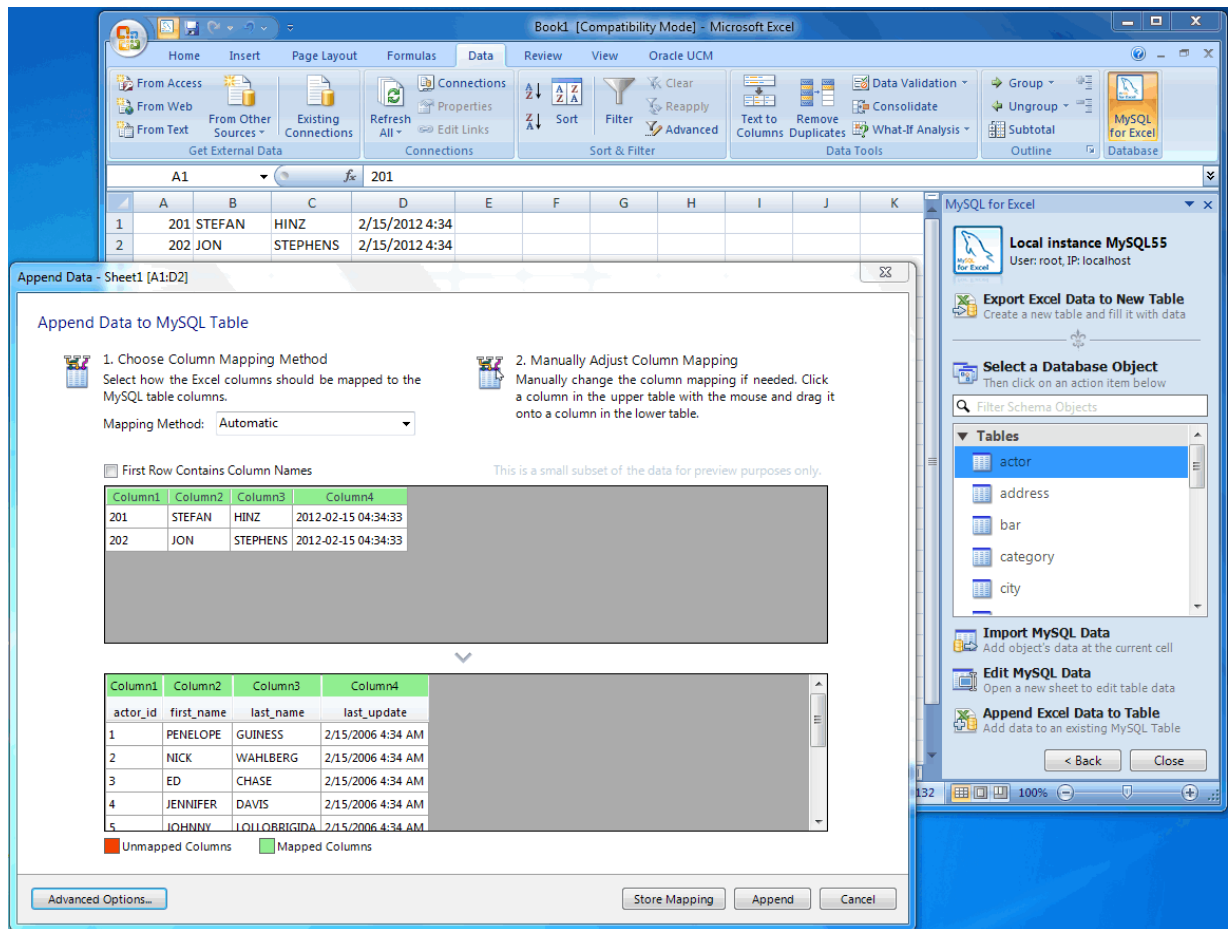
## 25.4. Append MySQL Data

Data from a Microsoft Excel spreadsheet can be appended to a MySQL database table by using the **Append Excel MySQL Data to Table** option.

### Column mappings

Mapping the Excel columns to the MySQL columns can be executed automatically (default), manually, or by using a stored mapping routine. An automatic mapping routine is the default, and can be can be tweaked if every column cannot be matched automatically. The following screenshot shows two columns of Excel data, and the preview dialog after choosing **Append Excel Data to Table**:

Figure 25.5. Appending Excel data to MySQL (Automatic mapping)



## General mapping information

It is common to tweak the column mappings. A few notes about the manual mapping process:

- Manual mapping is performed by dragging a column from the upper source grid (Excel spreadsheet) and dropping it into the lower target column MySQL table grid. Click anywhere within the column to initiate this dragging routine.
- The color of the header field for each column defines the current mapping status of the column. The colors include:
  - Green: A source column is mapped to a target column.
  - Red: A target column is not mapped.
  - Gray: A source column is not mapped.
- A source column may be mapped to multiple target columns, although this action generates a warning dialog.
- Right-clicking a target column shows a context menu with options to either **Remove Column Mapping** for a single column, or **Clear All Mappings** for all columns. Dragging a target column outside of the grid removes the mapping.

## Mapping methods

The three mapping methods are described below:

- **Automatic:** The automatic mapping method attempts to match the Excel source column names with the MySQL target table column names. It is then possible to manually tweak the mapping afterwards.

If the automatic process finds zero columns to match, then a simple 1 to 1 matching routine is attempted. Meaning, SourceColumn #1 to TargetColumn #1, SourceColumn #2 to TargetColumn #2, and so on.

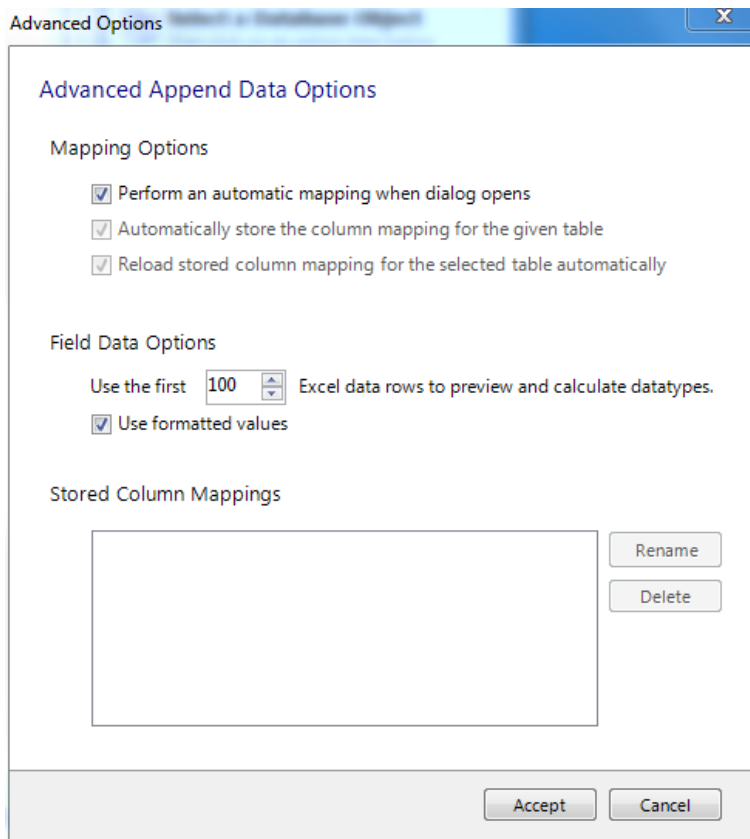
- **Manual:** The source column names are manually dragged (matched) with the target column names. Manual dragging can also be performed after the **Automatic** method is selected.
- **Stored:** Manual mapping styles may be saved using the **Store Mapping** button, which will also prompt for a name and then save it using a "*name* (dbname.tablename)" naming scheme. The saved mapping style will then be available alongside the **Automatic** and **Manual** options.

Stored mappings may be deleted or renamed within the **Advanced Options** dialog.

## Advanced Options

There are several advanced options that are configured and stored between sessions for each Excel user. The dialog looks similar to:

**Figure 25.6. Appending Excel data to MySQL (Advanced Options)**



The advanced **Mapping Options**:

- **Perform an automatic mapping when dialog opens:** Automatically attempt to map the target and source when the **Append Data** dialog is opened. This feature is enabled by default.
- **Automatically store the column mapping for the given table:** Stores each mapping routine after executing the **Append** operation. The mapping routine is saved using the "tablenameMapping (dbname.tablename)" format. This may also be performed manually using the **Store Mapping** button. It is enabled by default, and this feature was added in MySQL for Excel 1.1.0.

- **Reload stored column mapping for the selected table automatically:** If a stored mapping routine exists that matches all column names in the source grid with the target grid, then it is automatically be loaded. This is enabled by default, and this feature was added in MySQL for Excel 1.1.0.

The advanced **Field Data Options:**

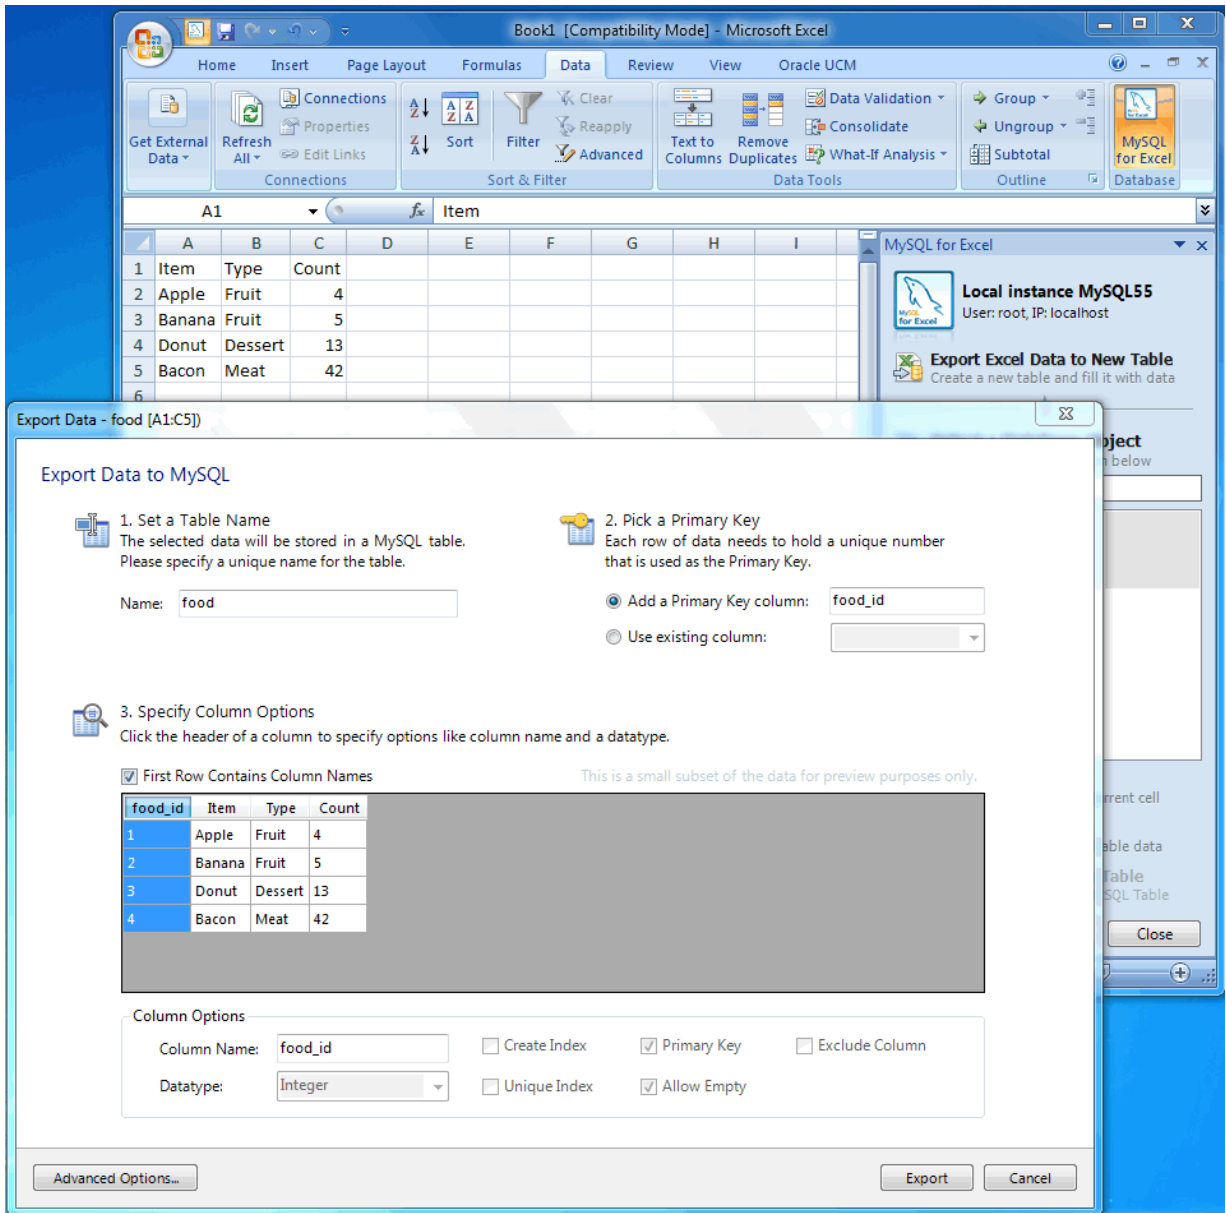
- **Use the first 100** (default) Excel data rows to preview and calculate data types. This determines the number of rows that the preview displays, and the values that affect the automatic mapping feature.
- **Use formatted values:** The data from Excel is treated as **Text**, **Double**, or **Date**. This is enabled by default. When disabled, data is never treated as a **Date** type, so for example, this means that a date would be represented as a number.

The **Stored Column Mappings** is a list of saved column mappings that were saved with with the "Automatically store the column mapping for the given table" feature, or manually with the **Store Mapping** option.

## 25.5. Export MySQL Data

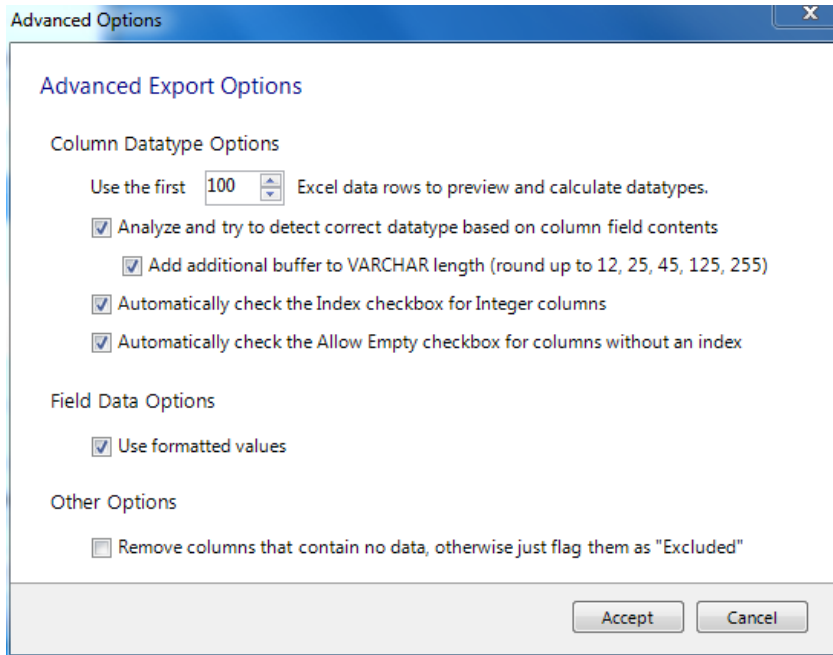
Data from a Microsoft Excel spreadsheet can be exported to a new MySQL database table by using the **Export Excel Data to New Table** option. Exporting data looks like so:

Figure 25.7. Exporting Excel data to MySQL



## Advanced Export options

Several advanced options enables you to tweak the exported data. The advanced options dialog looks like so:

**Figure 25.8. Exporting Excel data to MySQL (Advanced options)**

- **Column Datatype Options:**

- Use the first `100` (default) Excel data rows to preview and calculate data types: This determines the number of rows that the preview displays, and the values that affect the automatic mapping feature.
- Analyze and try to detect correct datatype based on column field contents: Attempts to analyze the data and determine the data type for the column. The column type is defined as `VARCHAR` if it contains multiple types.
- Add additional buffer to `VARCHAR` length (round up to 12, 25, 45, 125, 255): When the data type is automatically detected and is set to `VARCHAR`, then it calculates the maximum length for all rows within the column, and rounds up the maximum length to one of the defined lengths above.

If disabled, then the `VARCHAR` length is set to the length of the longest entry in the Excel spreadsheet.

- Automatically check the Index checkbox for Integer columns: If enabled (default), columns with an Integer data type will have the **Create Index** option enabled by default.
- Automatically check the Allow Empty checkbox for columns without an index: If enabled (default), columns without the **Create Index** checkbox checked will automatically enable the **Allow Empty** configuration option.

- **Field Data options:**

- Use formatted values: When enabled (default), the data from Excel is treated as `Text`, `Double`, or `Date`. When disabled, data is never treated as a `Date` type, so for example this means that a date would be represented as a number.

- **Other options:**

- Remove columns that contain no data, otherwise flag them as "Excluded": If enabled, columns without data in Excel are removed and not shown in the preview panel. If disabled (default), these columns will exist but have the **Exclude Column** option checked. This option was added in MySQL for Excel 1.1.0.

## 25.6. MySQL for Excel FAQ

Frequently Asked Questions, with Answers.

### Questions

- [26.6.1: \[3568\]](#) I installed the MySQL for Excel plugin, but can't find it in Microsoft Excel. How do I start it?

### Questions and Answers

**26.6.1: I installed the MySQL for Excel plugin, but can't find it in Microsoft Excel. How do I start it?**

The MySQL for Excel plugin is automatically added to Microsoft Excel's data menu when it is installed. Look for the MySQL for Excel icon, by default it will be listed on the right side of the main menu.

If it's not there, then you might have to reinstall the plugin. But before doing so, first check if it's listed under "Add/Remove Programs" in Microsoft Windows. If not, then it has not been installed. Next, check the Excel Add-Ins list. For Office 2007 this is found by clicking the Office logo in Excel (top left corner), click **Excel Options**, then select **Add-Ins**. Is MySQL for Excel listed as a COM Add-in? If so, then consider filing a bug report ([bugs.mysql.com](http://bugs.mysql.com)), or attempt to reinstall the plugin.



---

## Appendix A. Licenses for Third-Party Components

### Table of Contents

A.1. Ant-Contrib License .....	3572
A.2. ANTLR 3.3 License .....	3573
A.3. Boost Library License .....	3573
A.4. c3p0 JDBC Library License .....	3574
A.5. <a href="#">dtoa.c</a> License .....	3574
A.6. Editline Library ( <a href="#">libedit</a> ) License .....	3574
A.7. Facebook Fast Checksum Patch License .....	3577
A.8. Facebook Patches License .....	3578
A.9. <a href="#">FindGTest.cmake</a> License .....	3578
A.10. Fred Fish's Dbug Library License .....	3579
A.11. <a href="#">getarg</a> License .....	3580
A.12. GLib License (for MySQL Proxy) .....	3580
A.13. GNU General Public License Version 2.0, June 1991 .....	3581
A.14. GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009 .....	3586
A.15. GNU Lesser General Public License Version 2.1, February 1999 .....	3597
A.16. GNU Libtool License .....	3604
A.17. GNU Readline License .....	3605
A.18. GNU Standard C++ Library (libstdc++) License .....	3605
A.19. Google Controlling Master Thread I/O Rate Patch License .....	3606
A.20. Google Perftools (TCMalloc utility) License .....	3607
A.21. Google SMP Patch License .....	3607
A.22. jboss-common-jdbc-wrapper.jar License .....	3608
A.23. <a href="#">lib_sql.cc</a> License .....	3608
A.24. Libaio License .....	3608
A.25. <a href="#">libevent</a> License .....	3609
A.26. Libiconv License .....	3610
A.27. <a href="#">libintl</a> License .....	3611
A.28. Linux-PAM License .....	3611
A.29. <a href="#">LPeg</a> Library License .....	3612
A.30. Lua (liblua) License .....	3612
A.31. <a href="#">LuaFileSystem</a> Library License .....	3613
A.32. md5 (Message-Digest Algorithm 5) License .....	3613
A.33. <a href="#">memcached</a> License .....	3613
A.34. <a href="#">mkpasswd.pl</a> License .....	3614
A.35. nt_servc (Windows NT Service class library) License .....	3617
A.36. OpenPAM License .....	3618
A.37. OpenSSL v1.0 License .....	3618
A.38. PCRE License .....	3620
A.39. Percona Multiple I/O Threads Patch License .....	3621
A.40. Red HAT RPM Spec File License .....	3621
A.41. RegEX-Spencer Library License .....	3622
A.42. RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License .....	3622
A.43. Richard A. O'Keefe String Library License .....	3622
A.44. SHA-1 in C License .....	3623
A.45. Simple Logging Facade for Java (SLF4J) License .....	3623
A.46. Unicode Data Files .....	3623
A.47. <a href="#">zlib</a> License .....	3624
A.48. ZLIB.NET License .....	3624

The following is a list of the libraries we have included with the MySQL Server source and components used to test MySQL. We are thankful to all individuals that have created these. Some of the

components require that their licensing terms be included in the documentation of products that include them. Cross references to these licensing terms are given with the applicable items in the list.

- GroupLens Research Project

The MySQL Quality Assurance team would like to acknowledge the use of the MovieLens Data Sets (10 million ratings and 100,000 tags for 10681 movies by 71567 users) to help test MySQL products and to thank the GroupLens Research Project at the University of Minnesota for making the data sets available.

## MySQL 5.6

- [Section A.3, “Boost Library License”](#)
- [Section A.5, “`dtoa.c` License”](#)
- [Section A.6, “Editline Library \(`libedit`\) License”](#)
- [Section A.7, “Facebook Fast Checksum Patch License”](#)
- [Section A.8, “Facebook Patches License”](#)
- [Section A.9, “`FindGTest.cmake` License”](#)
- [Section A.10, “Fred Fish's Dbug Library License”](#)
- [Section A.11, “`getarg` License”](#)
- [Section A.13, “GNU General Public License Version 2.0, June 1991”](#)
- [Section A.14, “GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009”](#)
- [Section A.15, “GNU Lesser General Public License Version 2.1, February 1999”](#)
- [Section A.17, “GNU Readline License”](#)
- [Section A.18, “GNU Standard C++ Library \(`libstdc++`\) License”](#)
- [Section A.19, “Google Controlling Master Thread I/O Rate Patch License”](#)
- [Section A.20, “Google Perftools \(TCMalloc utility\) License”](#)
- [Section A.21, “Google SMP Patch License”](#)
- [Section A.23, “`lib\_sql.cc` License”](#)
- [Section A.24, “Libaio License”](#)
- [Section A.25, “`libevent` License”](#)
- [Section A.28, “Linux-PAM License”](#)
- [Section A.32, “md5 \(Message-Digest Algorithm 5\) License”](#)
- [Section A.33, “`memcached` License”](#)
- [Section A.34, “`mkpasswd.pl` License”](#)
- [Section A.35, “`nt\_servc` \(Windows NT Service class library\) License”](#)
- [Section A.36, “OpenPAM License”](#)
- [Section A.37, “OpenSSL v1.0 License”](#)
- [Section A.39, “Percona Multiple I/O Threads Patch License”](#)

- [Section A.40, “Red HAT RPM Spec File License”](#)
- [Section A.41, “RegEX-Spencer Library License”](#)
- [Section A.43, “Richard A. O’Keefe String Library License”](#)
- [Section A.44, “SHA-1 in C License”](#)
- [Section A.46, “Unicode Data Files”](#)
- [Section A.47, “zlib License”](#)

## MySQL Connector/C 6.0

- [Section A.10, “Fred Fish’s Dbug Library License”](#)
- [Section A.41, “RegEX-Spencer Library License”](#)
- [Section A.42, “RFC 3174 - US Secure Hash Algorithm 1 \(SHA1\) License”](#)
- [Section A.47, “zlib License”](#)

## MySQL Connector/C 6.1

- [Section A.5, “dtoa.c License”](#)
- [Section A.10, “Fred Fish’s Dbug Library License”](#)
- [Section A.13, “GNU General Public License Version 2.0, June 1991”](#)
- [Section A.32, “md5 \(Message-Digest Algorithm 5\) License”](#)
- [Section A.41, “RegEX-Spencer Library License”](#)
- [Section A.43, “Richard A. O’Keefe String Library License”](#)
- [Section A.44, “SHA-1 in C License”](#)
- [Section A.47, “zlib License”](#)

## MySQL Connector/C++

- [Section A.3, “Boost Library License”](#)
- [Section A.37, “OpenSSL v1.0 License”](#)

## MySQL Connector/J

- [Section A.1, “Ant-Contrib License”](#)
- [Section A.4, “c3p0 JDBC Library License”](#)
- [Section A.15, “GNU Lesser General Public License Version 2.1, February 1999”](#)
- [Section A.22, “jboss-common-jdbc-wrapper.jar License”](#)
- [Section A.45, “Simple Logging Facade for Java \(SLF4J\) License”](#)

## MySQL Connector/Net

- [Section A.2, “ANTLR 3.3 License”](#)
- [Section A.42, “RFC 3174 - US Secure Hash Algorithm 1 \(SHA1\) License”](#)
- [Section A.47, “zlib License”](#)
- [Section A.48, “ZLIB.NET License”](#)

## MySQL Connector/ODBC

- [Section A.16, “GNU Libtool License”](#)
- [Section A.37, “OpenSSL v1.0 License”](#)

## MySQL Proxy

- [Section A.12, “GLib License \(for MySQL Proxy\)”](#)
- [Section A.15, “GNU Lesser General Public License Version 2.1, February 1999”](#)
- [Section A.25, “libevent License”](#)
- [Section A.26, “Libiconv License”](#)
- [Section A.27, “libintl License”](#)
- [Section A.29, “LPeg Library License”](#)
- [Section A.30, “Lua \(liblua\) License”](#)
- [Section A.31, “LuaFileSystem Library License”](#)
- [Section A.38, “PCRE License”](#)

## A.1. Ant-Contrib License

The following software may be included in this product: Ant-Contrib

```
Ant-Contrib
Copyright (c) 2001-2003 Ant-Contrib project. All rights reserved.
Licensed under the Apache 1.1 License Agreement, a copy of which is reproduced below.

The Apache Software License, Version 1.1

Copyright (c) 2001-2003 Ant-Contrib project. All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in
the documentation and/or other materials provided with the
distribution.

3. The end-user documentation included with the redistribution, if
any, must include the following acknowledgement:
   "This product includes software developed by the
   Ant-Contrib project (http://sourceforge.net/projects/ant-contrib)."
   Alternately, this acknowledgement may appear in the software itself,
   if and wherever such third-party acknowledgements normally appear.

4. The name Ant-Contrib must not be used to endorse or promote
products derived from this software without prior written
permission. For written permission, please contact
ant-contrib-developers@lists.sourceforge.net.

5. Products derived from this software may not be called "Ant-Contrib"
nor may "Ant-Contrib" appear in their names without prior written
permission of the Ant-Contrib project.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED
WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
```

```
DISCLAIMED. IN NO EVENT SHALL THE ANT-CONTRIB PROJECT OR ITS
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF
USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT
OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

## A.2. ANTLR 3.3 License

The following software may be included in this product:

ANTLR 3.3

ANTLR 3.3 License

Copyright (c) 2010 Terence Parr

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.3. Boost Library License

The following software may be included in this product:

### Boost C++ Libraries

Use of any of this software is governed by the terms of the license below:

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND

```
NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE
DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER
LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT
OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS
IN THE SOFTWARE.
```

## A.4. c3p0 JDBC Library License

You are receiving a copy of `c3p0-0.9.1-pre6.jar` in both source and object code in the following `/src/lib/c3p0-0.9.1-pre6.jar`. The terms of the Oracle license do NOT apply to `c3p0-0.9.1-pre6.jar`; it is licensed under the following license, separately from the Oracle programs you receive. If you do not wish to install this library, you may remove the file `/src/lib/c3p0-0.9.1-pre6.jar`, but the Oracle program might not operate properly or at all without the library.

This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.5. dtoa.c License

The following software may be included in this product:

`dtoa.c`

```
The author of this software is David M. Gay.
```

```
Copyright (c) 1991, 2000, 2001 by Lucent Technologies.
```

```
Permission to use, copy, modify, and distribute this software for
any purpose without fee is hereby granted, provided that this entire
notice is included in all copies of any software which is or includes
a copy or modification of this software and in all copies of the
supporting documentation for such software.
```

```
THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR
IMPLIED WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR LUCENT
MAKES ANY REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE
MERCHANTABILITY OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR
PURPOSE.
```

## A.6. Editline Library (`libedit`) License

The following software may be included in this product:

Editline Library (`libedit`)

Some files are:

```
Copyright (c) 1992, 1993
The Regents of the University of California. All rights reserved.
```

```
This code is derived from software contributed to
Berkeley by Christos Zoulas of Cornell University.
```

```
Redistribution and use in source and binary forms,
with or without modification, are permitted provided
that the following conditions are met:
```

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING,
BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO
EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY
DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED
AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING
IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.
```

**Some files are:**

```
Copyright (c) 2001 The NetBSD Foundation, Inc.
All rights reserved.
```

```
This code is derived from software contributed to The NetBSD Foundation
```

```
by Anthony Mallet.
```

```
Redistribution and use in source and binary forms,
with or without modification, are permitted provided
that the following conditions are met:
```

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

```
THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC.
AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED
WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL
THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY
DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF
USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE
OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH
DAMAGE.
```

**Some files are:**

```
Copyright (c) 1997 The NetBSD Foundation, Inc.
All rights reserved.
```

```
This code is derived from software contributed to The NetBSD Foundation
```

```
by Jaromir Dolecek.
```

```
Redistribution and use in source and binary forms,
with or without modification, are permitted provided
that the following conditions are met:
```

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Some files are:**

Copyright (c) 1998 Todd C. Miller <Todd.Miller@courtesan.com>

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND TODD C. MILLER DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL TODD C. MILLER BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

**Some files are:**

Copyright (c) 1998 The NetBSD Foundation, Inc.  
All rights reserved.

This code is derived from software contributed to The NetBSD Foundation by Christos Zoulas.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Some files are:**

Copyright (c) 2009 The NetBSD Foundation, Inc.



All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
4. Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.7. Facebook Fast Checksum Patch License

The following software may be included in this product:

Facebook Fast Checksum Patch

Copyright (C) 2009-2010 Facebook, Inc. All Rights Reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY FACEBOOK, INC. "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL FACEBOOK, INC. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Also included:

crc32.c -- compute the CRC-32 of a buf stream  
Copyright (C) 1995-2005 Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it

freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org  
Mark Adler madler@alumni.caltech.edu

## A.8. Facebook Patches License

The following software may be included in this product:

```
Copyright (c) 2012, Facebook, Inc.
All rights reserved.
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:
* Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
* Redistributions in binary form must reproduce the above copyright notice,
this list of conditions and the following disclaimer in the documentation
and/or other materials provided with the distribution.
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.
```

## A.9. FindGTest.cmake License

The following software may be included in this product:

```
FindGTest.cmake helper script (part of CMake)

Copyright 2009 Kitware, Inc.
Copyright 2009 Philip Lowman
Copyright 2009 Daniel Blezek

Distributed under the OSI-approved BSD License (the "License");
see accompanying file Copyright.txt for details.

This software is distributed WITHOUT ANY WARRANTY; without even the
implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
See the License for more information.
=====
(To distributed this file outside of CMake, substitute the full
License text for the above reference.)

Thanks to Daniel Blezek for the GTEST_ADD_TESTS code

Text of Copyright.txt mentioned above:

CMake - Cross Platform Makefile Generator
Copyright 2000-2009 Kitware, Inc., Insight Software Consortium
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

* Redistributions of source code must retain the above copyright
```

notice, this list of conditions and the following disclaimer.

- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the names of Kitware, Inc., the Insight Software Consortium, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.10. Fred Fish's Dbug Library License

The following software may be included in this product:

Fred Fish's Dbug Library

### N O T I C E

Copyright Abandoned, 1987, Fred Fish

This previously copyrighted work has been placed into the public domain by the author and may be freely used for any purpose, private or commercial.

Because of the number of inquiries I was receiving about the use of this product in commercially developed works I have decided to simply make it public domain to further its unrestricted use. I specifically would be most happy to see this material become a part of the standard Unix distributions by AT&T and the Berkeley Computer Science Research Group, and a standard part of the GNU system from the Free Software Foundation.

I would appreciate it, as a courtesy, if this notice is left in all copies and derivative works. Thank you.

The author makes no warranty of any kind with respect to this product and explicitly disclaims any implied warranties of mer-

chantability or fitness for any particular purpose.

The `debug_analyze.c` file is subject to the following notice:

Copyright June 1987, Binayak Banerjee  
All rights reserved.

This program may be freely distributed under the same terms and conditions as Fred Fish's Dbug package.

## A.11. getarg License

The following software may be included in this product:

`getarg` Function (`getarg.h`, `getarg.c` files)

Copyright (c) 1997 - 2000 Kungliga Tekniska Högskolan  
(Royal Institute of Technology, Stockholm, Sweden).  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the Institute nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INSTITUTE AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE INSTITUTE OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.12. GLib License (for MySQL Proxy)

The following software may be included in this product:

GLib

You are receiving a copy of the GLib library in both source and object code in the following `[proxy install dir]/lib/` and `[proxy install dir]/licenses/lgpl` folders. The terms of the Oracle license do NOT apply to the GLib library; it is licensed under the following license, separately from the Oracle programs you receive. If you do not wish to install this library, you may create an "exclude" file and run tar with the X option, as in the following example, but the Oracle program might not operate properly or at all without the library:

```
tar -xvfX <package-tar-file> <exclude-file>
where the exclude-file contains, e.g.:
<package-name>/lib/libglib-2.0.so.0.1600.6
<package-name>/lib/libglib-2.0.so.0
...
```

```
Example:
tar -xvfX mysql-proxy-0.8.1-solaris10-x86-64bit.tar.gz Exclude

Exclude File:
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libglib-2.0.so
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libglib-2.0.so.0
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libglib-2.0.so.0.1600.6
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgmodule-2.0.so
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgmodule-2.0.so.0
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgmodule-2.0.so.0.1600.6
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgthread-2.0.so
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgthread-2.0.so.0
mysql-proxy-0.8.1-solaris10-x86-64bit/lib/libgthread-2.0.so.0.1600.6
mysql-proxy-0.8.1-solaris10-x86-64bit/licenses/lgpl/glib-2.16.6.tar.gz
```

This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.13. GNU General Public License Version 2.0, June 1991

The following applies to all products licensed under the GNU General Public License, Version 2.0: You may not use the identified files except in compliance with the GNU General Public License, Version 2.0 (the "License.") You may obtain a copy of the License at <http://www.gnu.org/licenses/gpl-2.0.txt>. A copy of the license is also reproduced below. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

GNU GENERAL PUBLIC LICENSE  
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.,  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA  
Everyone is permitted to copy and distribute verbatim  
copies of this license document, but changing it is not  
allowed.

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but

does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or

distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we



sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.>

Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details  
type 'show w'. This is free software, and you are welcome  
to redistribute it under certain conditions; type 'show c'  
for details.

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the

commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
program 'Gnomovision' (which makes passes at compilers) written
by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

## A.14. GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009

### GNU GENERAL PUBLIC LICENSE Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <<http://fsf.org/>>  
Everyone is permitted to copy and distribute verbatim copies of  
this license document, but changing it is not allowed.

#### Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

#### TERMS AND CONDITIONS

##### 0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

##### 1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

## 3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological

measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

#### 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

#### 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

#### 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product

(including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.

b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity

or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

#### 7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that

it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

#### 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

#### 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

#### 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.



You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

#### 11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that

contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING

WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or  
modify it under the terms of the GNU General Public License  
as published by the Free Software Foundation, either version 3  
of the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program. If not, see  
<http://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>  
This program comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type 'show c' for details.
```

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see [<http://www.gnu.org/licenses/>](http://www.gnu.org/licenses/).

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine

library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <<http://www.gnu.org/philosophy/why-not-lgpl.html>>.

==

==

GCC RUNTIME LIBRARY EXCEPTION

Version 3.1, 31 March 2009

Copyright © 2009 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This GCC Runtime Library Exception ("Exception") is an additional permission under section 7 of the GNU General Public License, version 3 ("GPLv3"). It applies to a given file (the "Runtime Library") that bears a notice placed by the copyright holder of the file stating that the file is governed by GPLv3 along with this Exception.

When you use GCC to compile a program, GCC may combine portions of certain GCC header files and runtime libraries with the compiled program. The purpose of this Exception is to allow compilation of non-GPL (including proprietary) programs to use, in this way, the header files and runtime libraries covered by this Exception.

#### 0. Definitions.

A file is an "Independent Module" if it either requires the Runtime Library for execution after a Compilation Process, or makes use of an interface provided by the Runtime Library, but is not otherwise based on the Runtime Library.

"GCC" means a version of the GNU Compiler Collection, with or without modifications, governed by version 3 (or a specified later version) of the GNU General Public License (GPL) with the option of using any subsequent versions published by the FSF.

"GPL-compatible Software" is software whose conditions of propagation, modification and use would permit combination with GCC in accord with the license of GCC.

"Target Code" refers to output from any compiler for a real or virtual target processor architecture, in executable form or suitable for input to an assembler, loader, linker and/or execution phase. Notwithstanding that, Target Code does not include data in any format that is used as a compiler intermediate representation, or used for producing a compiler intermediate representation.

The "Compilation Process" transforms code entirely represented in non-intermediate languages designed for human-written code, and/or in Java Virtual Machine byte code, into Target Code. Thus, for example, use of source code generators and preprocessors need not be considered part of the Compilation Process, since the Compilation Process can be understood as starting with the output of the generators or preprocessors.

A Compilation Process is "Eligible" if it is done using GCC, alone or with other GPL-compatible software, or if it is done without using any work based on GCC. For example, using non-GPL-compatible Software to optimize any GCC intermediate representations would not qualify as an Eligible Compilation Process.

#### 1. Grant of Additional Permission.

You have permission to propagate a work of Target Code formed by combining the Runtime Library with Independent Modules, even if such propagation would otherwise violate the terms of GPLv3, provided that all Target Code was generated by Eligible Compilation Processes. You may then convey such a combination under terms of your choice,

consistent with the licensing of the Independent Modules.

2. No Weakening of GCC Copyleft.

The availability of this Exception does not imply any general presumption that third-party software is unaffected by the copyleft requirements of the license of GCC.

==

## A.15. GNU Lesser General Public License Version 2.1, February 1999

The following applies to all products licensed under the GNU Lesser General Public License, Version 2.1: You may not use the identified files except in compliance with the GNU Lesser General Public License, Version 2.1 (the "License"). You may obtain a copy of the License at <http://www.gnu.org/licenses/lgpl-2.1.html>. A copy of the license is also reproduced below. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

GNU LESSER GENERAL PUBLIC LICENSE  
Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA  
Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling

it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object



file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the

Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or
modify it under the terms of the GNU Lesser General Public
License as published by the Free Software Foundation; either
version 2.1 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
Lesser General Public License for more details.
```

```
You should have received a copy of the GNU Lesser General Public
License along with this library; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA
02110-1301 USA
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
library `Frob' (a library for tweaking knobs) written by James
Random Hacker.
```

```
<signature of Ty Coon>, 1 April 1990
Ty Coon, President of Vice
```

That's all there is to it!

## A.16. GNU Libtool License

The following software may be included in this product:

### GNU Libtool (The GNU Portable Library Tool)

```
If you are receiving a copy of the Oracle software in
source code, you are also receiving a copy of two files
(ltmain.sh and ltdl.h) generated by the GNU Libtool in
source code. If you received the Oracle software under
a license other than a commercial (non-GPL) license,
then the terms of the Oracle license do NOT apply to
these files from GNU Libtool; they are licensed under
the following licenses, separately from the Oracle
programs you receive.
```

```
Oracle elects to use GNU General Public License version
2 (GPL) for any software where a choice of GPL or GNU
Lesser/Library General Public License (LGPL) license
versions are made available with the language indicating
that GPL/LGPL or any later version may be used, or where
a choice of which version of the GPL/LGPL is applied is
unspecified.
```

From GNU Libtool:

```
ltmain.sh - Provide generalized library-building support
services.
NOTE: Changing this file will not affect anything until
you rerun configure.
Copyright (C) 1996, 1997, 1998, 1999, 2000, 2001, 2003, 2004,
2005, 2006, 2007 Free Software Foundation, Inc.
Originally by Gordon Matzigkeit, 1996
```

```
This program is free software; you can redistribute it
and/or modify it under the terms of the GNU General
Public License as published by the Free Software Foundation;
either version 2 of the License, or (at your option) any
later version.
```

```
This program is distributed in the hope that it will be
```

```
useful, but WITHOUT ANY WARRANTY; without even the implied
warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR
PURPOSE. See the GNU General Public License for more
details. You should have received a copy of the GNU General
Public License along with this program; if not, write to
the Free Software Foundation, Inc., 51 Franklin Street,
Fifth Floor, Boston, MA 02110-1301, USA.
```

```
As a special exception to the GNU General Public License,
if you distribute this file as part of a program that
contains a configuration script generated by Autoconf,
you may include it under the same distribution terms that
you use for the rest of that program.
```

This component is licensed under [Section A.13, "GNU General Public License Version 2.0, June 1991"](#)

## A.17. GNU Readline License

The following software may be included in this product:

### GNU Readline Library

```
GNU Readline Library
With respect to MySQL Server/Cluster software licensed
under GNU General Public License, you are receiving a
copy of the GNU Readline Library in source code. The
terms of any Oracle license that might accompany the
Oracle programs do NOT apply to the GNU Readline Library;
it is licensed under the following license, separately
from the Oracle programs you receive. Oracle elects to
use GNU General Public License version 2 (GPL) for any
software where a choice of GPL license versions are
made available with the language indicating that GPLv2
or any later version may be used, or where a choice of
which version of the GPL is applied is unspecified.
```

This component is licensed under [Section A.13, "GNU General Public License Version 2.0, June 1991"](#)

## A.18. GNU Standard C++ Library (libstdc++) License

The following software may be included in this product: GNU Standard C++ Library (libstdc++)

This component is licensed under [Section A.14, "GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009"](#).

Additional notices:

```
==
Copyright (c) 1994
Hewlett-Packard Company

Permission to use, copy, modify, distribute and sell this software
and its documentation for any purpose is hereby granted without fee,
provided that the above copyright notice appear in all copies and
that both that copyright notice and this permission notice appear
in supporting documentation. Hewlett-Packard Company makes no
representations about the suitability of this software for any
purpose. It is provided "as is" without express or implied
warranty.
==

==
Copyright (c) 1996,1997
Silicon Graphics Computer Systems, Inc.

Permission to use, copy, modify, distribute and sell this software
and its documentation for any purpose is hereby granted without fee,
provided that the above copyright notice appear in all copies and
that both that copyright notice and this permission notice appear
```

```
in supporting documentation. Silicon Graphics makes no
representations about the suitability of this software for any
purpose. It is provided "as is" without express or implied
warranty.
==
==
shared_count.hpp
@ Copyright (c) 2001, 2002, 2003 Peter Dimov and Multi Media Ltd.

shared_ptr.hpp
Copyright (C) 1998, 1999 Greg Colvin and Beman Dawes.
Copyright (C) 2001, 2002, 2003 Peter Dimov

weak_ptr.hpp
Copyright (C) 2001, 2002, 2003 Peter Dimov

enable_shared_from_this.hpp
Copyright (C) 2002 Peter Dimov

Distributed under the Boost Software License, Version 1.0.

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or
organization obtaining a copy of the software and accompanying
documentation covered by this license (the "Software") to use,
reproduce, display, distribute, execute, and transmit the Software,
and to prepare derivative works of the Software, and to permit
third-parties to whom the Software is furnished to do so, all subject
to the following:

The copyright notices in the Software and this entire statement,
including the above license grant, this restriction and the following
disclaimer, must be included in all copies of the Software, in whole
or in part, and all derivative works of the Software, unless such
copies or derivative works are solely in the form of machine-executable
object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND
NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE
DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER
LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
DEALINGS IN THE SOFTWARE.
==
==
Copyright (C) 2004 Ami Tavory and Vladimir Dreizin, IBM-HRL.

Permission to use, copy, modify, sell, and distribute this software
is hereby granted without fee, provided that the above copyright
notice appears in all copies, and that both that copyright notice
and this permission notice appear in supporting documentation. None
of the above authors, nor IBM Haifa Research Laboratories, make any
representation about the suitability of this software for any
purpose. It is provided "as is" without express or implied warranty.
==
```

## A.19. Google Controlling Master Thread I/O Rate Patch License

The following software may be included in this product:

Google Controlling master thread I/O rate patch

```
Copyright (c) 2009, Google Inc.
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without
```

modification, are permitted provided that the following conditions are met:

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of the Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.20. Google Perftools (TCMalloc utility) License

The following software may be included in this product:

Google Perftools (TCMalloc utility)

Copyright (c) 1998-2006, Google Inc.  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.21. Google SMP Patch License

The following software may be included in this product:

Google SMP Patch

Google SMP patch

```
Copyright (c) 2008, Google Inc.  
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without  
modification, are permitted provided that the following conditions  
are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of the Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS  
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT  
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS  
FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE  
COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,  
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,  
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;  
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER  
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT  
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN  
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE  
POSSIBILITY OF SUCH DAMAGE.
```

## A.22. jboss-common-jdbc-wrapper.jar License

You are receiving a copy of `jboss-common-jdbc-wrapper.jar` in both source and object code in the following `/src/lib/jboss-common-jdbc-wrapper.jar`. The terms of the Oracle license do NOT apply to `jboss-common-jdbc-wrapper.jar`; it is licensed under the following license, separately from the Oracle programs you receive. If you do not wish to install this library, you may remove the file `/src/lib/jboss-common-jdbc-wrapper.jar`, but the Oracle program might not operate properly or at all without the library.

This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.23. lib\_sql.cc License

The following software may be included in this product:

`lib_sql.cc`

```
Copyright (c) 2000  
SWsoft company
```

```
This material is provided "as is", with absolutely no warranty  
expressed or implied. Any use is at your own risk.
```

```
Permission to use or copy this software for any purpose is hereby  
granted without fee, provided the above notices are retained on  
all copies. Permission to modify the code and to distribute modified  
code is granted, provided the above notices are retained, and a  
notice that the code was modified is included with the above copyright  
notice.
```

```
This code was modified by the MySQL team.
```

## A.24. Libaio License

The following software may be included in this product:

`libaio`



This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.25. libevent License

The following software may be included in this product:

```
libevent

Copyright (c) 2000-2007 Niels Provos <provos@citi.umich.edu>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE

==
Parts developed by Adam Langley
==

==
log.c
Based on err.c, which was adapted from OpenBSD libc *err*warncode.

Copyright (c) 2005 Nick Mathewson
Copyright (c) 2000 Dug Song
Copyright (c) 1993 The Regents of the University of California.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in
   the documentation and/or other materials provided with the
   distribution.
3. Neither the name of the University nor the names of its
   contributors may be used to endorse or promote products derived
   from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS
OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF
USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED
AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING
IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF
```

```
THE POSSIBILITY OF SUCH DAMAGE.
==
==
min_heap.h

Copyright (c) 2006 Maxim Yegorushkin
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in
   the documentation and/or other materials provided with the
   distribution.
3. The name of the author may not be used to endorse or promote
   products derived from this software without specific prior
   written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS
OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY
DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE
GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER
IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
==
==
win32.c

Copyright 2000-2002 Niels Provos
Copyright 2003 Michael A. Davis
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in
   the documentation and/or other materials provided with the
   distribution.
3. The name of the author may not be used to endorse or promote
   products derived from this software without specific prior
   written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS
OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY
DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE
GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER
IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
==
```

## A.26. Libiconv License

The following software may be included in this product:

## Libiconv

You are receiving a copy of the GNU LIBICONV Library. The terms of the Oracle license do NOT apply to the GNU LIBICONV Library; it is licensed under the following license, separately from the Oracle programs you receive. If you do not wish to install this program, you may delete [agent install dir]/lib/libiconv.\* and [agent install dir]/licenses/lgpl/iconv files.

This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.27. libintl License

The following software may be included in this product:

## libintl

Copyright (C) 1994 X Consortium

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE X CONSORTIUM BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of the X Consortium shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from the X Consortium.

FSF changes to this file are in the public domain.

Copyright 1996-2007 Free Software Foundation, Inc. Taken from GNU libtool, 2001

Originally by Gordon Matzigkeit <gord@gnu.ai.mit.edu>, 1996

This file is free software; the Free Software Foundation gives unlimited permission to copy and/or distribute it, with or without modifications, as long as this notice is preserved.

You are receiving a copy of the libintl library. The terms of the Oracle license do NOT apply to the libintl library; it is licensed under the following license, separately from the Oracle programs you receive. If you do not wish to install this program, you may create an "exclude" file and run tar with the X option.

This component is licensed under [Section A.15, "GNU Lesser General Public License Version 2.1, February 1999"](#).

## A.28. Linux-PAM License

The following software may be included in this product:

Linux-PAM (pam-devel, Pluggable authentication modules for Linux)

Copyright Theodore Ts'o, 1996. All rights reserved.

(For the avoidance of doubt, Oracle uses and distributes this component under the terms below and elects not to do so under the GPL even though the GPL is referenced as an option below.)

Redistribution and use in source and binary forms, with or

without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, and the entire permission notice in its entirety, including the disclaimer of warranties.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

ALTERNATIVELY, this product may be distributed under the terms of the GNU Public License, in which case the provisions of the GPL are required INSTEAD OF the above restrictions. (This clause is necessary due to a potential bad interaction between the GPL and the restrictions contained in a BSD-style copyright.)

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.29. LPeg Library License

The following software may be included in this product:

LPeg

Use of any of this software is governed by the terms of the license below:

Copyright © 2008 Lua.org, PUC-Rio.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## A.30. Lua (liblua) License

The following software may be included in this product:

Lua (liblua)

Copyright © 1994-2008 Lua.org, PUC-Rio.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish,

```
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject
to the following conditions:
```

```
The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.
```

```
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES
OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND
NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT
HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY,
WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR
OTHER DEALINGS IN THE SOFTWARE.
```

## A.31. LuaFileSystem Library License

The following software may be included in this product:

```
LuaFileSystem
```

```
Copyright © 2003 Kepler Project.
```

```
Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject
to the following conditions:
```

```
The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.
```

```
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.
IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR
ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION
WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

## A.32. md5 (Message-Digest Algorithm 5) License

The following software may be included in this product:

```
md5 (Message-Digest Algorithm 5)
```

```
This code implements the MD5 message-digest algorithm.
The algorithm is due to Ron Rivest. This code was
written by Colin Plumb in 1993, no copyright is claimed.
This code is in the public domain; do with it what you wish.
```

```
Equivalent code is available from RSA Data Security, Inc.
This code has been tested against that, and is equivalent,
except that you don't need to include two pages of legalese
with every copy.
```

```
The code has been modified by Mikael Ronstroem to handle
calculating a hash value of a key that is always a multiple
of 4 bytes long. Word 0 of the calculated 4-word hash value
is returned as the hash value.
```

## A.33. memcached License

The following software may be included in this product:

```
memcached
```

Copyright (c) 2003, Danga Interactive, Inc.  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

\* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

\* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

\* Neither the name of the Danga Interactive nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## A.34. mkpasswd.pl License

The following software may be included in this product:

mkpasswd.pl Perl module

Copyright (C) 2003-2004 by Chris Grau

This library is free software; you can redistribute it and/or modify it under the same terms as Perl itself, either Perl version 5.8.1 or, at your option, any later version of Perl 5 you may have available.

The Perl 5.8.1 license (from <http://www.cpan.org/src/5.0/perl-5.8.1.tar.gz> - main readme file):

Perl Kit, Version 5

Copyright (C) 1993, 1994, 1995, 1996, 1997, 1998  
1999, 2000, 2001, by Larry Wall and others

All rights reserved.

This program is free software; you can redistribute it and/or modify it under the terms of either:

a) the GNU General Public License as published by the Free Software Foundation; either version 1, or (at your option) any later version, or

b) the "Artistic License" which comes with this Kit.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See either the GNU General Public License or the Artistic License for more details.

You should have received a copy of the Artistic License with this Kit, in the file named "Artistic". If not, I'll be glad to provide one.

You should also have received a copy of the GNU General Public License

along with this program in the file named "Copying". If not, write to the  
  
Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  
02111-1307, USA or visit their web page on the internet at  
<http://www.gnu.org/copyleft/gpl.html>.

For those of you that choose to use the GNU General Public License,  
my interpretation of the GNU General Public License is that no Perl  
script falls under the terms of the GPL unless you explicitly put  
said script under the terms of the GPL yourself. Furthermore, any  
object code linked with perl does not automatically fall under the  
terms of the GPL, provided such object code only adds definitions  
of subroutines and variables, and does not otherwise impair the  
resulting interpreter from executing any standard Perl script. I  
consider linking in C subroutines in this manner to be the moral  
equivalent of defining subroutines in the Perl language itself. You  
may sell such an object file as proprietary provided that you provide  
or offer to provide the Perl source, as specified by the GNU General  
Public License. (This is merely an alternate way of specifying input  
to the program.) You may also sell a binary produced by the dumping of  
a running Perl script that belongs to you, provided that you provide or  
offer to provide the Perl source as specified by the GPL. (The  
fact that a Perl interpreter and your code are in the same binary file  
is, in this case, a form of mere aggregation.) This is my interpretation  
of the GPL. If you still have concerns or difficulties understanding  
my intent, feel free to contact me. Of course, the Artistic License  
spells all this out for your protection, so you may prefer to use that.

-----  
Perl is a language that combines some of the features of C, sed, awk  
and shell. See the manual page for more hype. There are also many Perl  
books available, covering a wide variety of topics, from various publishers.  
See `pod/perlbook.pod` for more information.

Please read all the directions below before you proceed any further, and  
then follow them carefully.

After you have unpacked your kit, you should have all the files listed  
in MANIFEST.

#### Installation

1) Detailed instructions are in the file "INSTALL", which you should  
read if you are either installing on a system resembling Unix  
or porting perl to another platform. For non-Unix platforms, see the  
corresponding README.

2) Read the manual entries before running perl.

3) IMPORTANT! Help save the world! Communicate any problems and suggested  
patches to [perlbug@perl.org](mailto:perlbug@perl.org) so we can keep the world in sync.  
If you have a problem, there's someone else out there who either has had  
or will have the same problem. It's usually helpful if you send the  
output of the "myconfig" script in the main perl directory.

If you've succeeded in compiling perl, the perlbug script in the "utils"  
subdirectory can be used to help mail in a bug report.

If possible, send in patches such that the patch program will apply them.  
Context diffs are the best, then normal diffs. Don't send ed scripts--  
I've probably changed my copy since the version you have.

The latest versions of perl are always available on the various CPAN  
(Comprehensive Perl Archive Network) sites around the world.  
See `<URL:http://www.cpan.org/src/>`.

Just a personal note: I want you to know that I create nice things like this  
because it pleases the Author of my story. If this bothers you, then your  
notion of Authorship needs some revision. But you can use perl anyway. :-)

The author.

=====

The "Artistic License"

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.

3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as unet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:



a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

## A.35. nt\_servc (Windows NT Service class library) License

The following software may be included in this product:

nt\_servc (Windows NT Service class library)

Windows NT Service class library  
Copyright Abandoned 1998 Irena Pancirov - Irnet Snc  
This file is public domain and comes with NO WARRANTY of any kind

## A.36. OpenPAM License

The following software may be included in this product:

### OpenPAM

```
Copyright (c) 2002-2003 Networks Associates Technology, Inc.
Copyright (c) 2004-2007 Dag-Erling Smørgrav
All rights reserved.
```

```
This software was developed for the FreeBSD Project by
ThinkSec AS and Network Associates Laboratories, the
Security Research Division of Network Associates, Inc.
under DARPA/SPAWAR contract N66001-01-C-8035 ("CBOSS"),
as part of the DARPA CHATS research program.
```

```
Redistribution and use in source and binary forms,
with or without modification, are permitted provided
that the following conditions are met:
```

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING,
BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN
NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR
ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED
AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

## A.37. OpenSSL v1.0 License

The following software may be included in this product:

```
NOTE: Does not apply to GPL licensed server (OpenSSL is not shipped with it)
```

```
OpenSSL v1.0
```

```
LICENSE ISSUES
```

```
=====
```

```
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of
the OpenSSL License and the original SSLeay license apply to the toolkit. See
below for the actual license texts. Actually both licenses are BSD-style Open
Source licenses. In case of any license issues related to OpenSSL please
contact openssl-core@openssl.org.
```

```
OpenSSL License
```

```
-----
```

```
/ =====
```

```
Copyright (c) 1998-2008 The OpenSSL Project.
```

```
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:
```

1. Redistributions of source code must retain the above copyright notice,

this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (Link1 /)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (Link2 /)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====  
This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

-----  
/ Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)  
All rights reserved.  
This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL. This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com). Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-). 4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)" THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT

```
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. The
license and distribution terms for any publically available version or
derivative of this code cannot be changed. i.e. this code cannot simply be
copied and put under another distribution license [including the GNU Public
License.]
```

## A.38. PCRE License

The following software may be included in this product:

### PCRE (Perl Compatible Regular Expressions) Library

#### PCRE LICENCE

PCRE is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language.

Release 7 of PCRE is distributed under the terms of the "BSD" licence, as specified below. The documentation for PCRE, supplied in the "doc" directory, is distributed under the same terms as the software itself.

The basic library functions are written in C and are freestanding. Also included in the distribution is a set of C++ wrapper functions.

#### THE BASIC LIBRARY FUNCTIONS

```
-----
Written by:      Philip Hazel
Email local part: ph10
Email domain:    cam.ac.uk
```

University of Cambridge Computing Service,  
Cambridge, England. Phone: +44 1223 334714.

Copyright (c) 1997-2006 University of Cambridge  
All rights reserved.

#### THE C++ WRAPPER FUNCTIONS

```
-----
Contributed by:  Google Inc.
```

Copyright (c) 2006, Google Inc.  
All rights reserved.

#### THE "BSD" LICENCE

```
-----
Redistribution and use in source and binary forms,
with or without modification, are permitted provided
that the following conditions are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of the University of Cambridge nor the name of Google Inc. nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS

```
BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY,  
OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT  
OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;  
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF  
LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING  
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF  
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

End

## A.39. Percona Multiple I/O Threads Patch License

The following software may be included in this product:

Percona Multiple I/O threads patch

```
Copyright (c) 2008, 2009 Percona Inc  
All rights reserved.
```

```
Redistribution and use of this software in source and binary forms,  
with or without modification, are permitted provided that the  
following conditions are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of Percona Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission of Percona Inc.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS  
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT  
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS  
FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE  
COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,  
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,  
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;  
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER  
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT  
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN  
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE  
POSSIBILITY OF SUCH DAMAGE.
```

## A.40. Red HAT RPM Spec File License

The following software may be included in this product:

Red Hat RPM Spec File

```
You are receiving a copy of the Red Hat spec file. The terms of the Oracle  
license do NOT apply to the Red Hat spec file; it is licensed under the  
following license, separately from the Oracle programs you receive.
```

```
GNU GENERAL PUBLIC LICENSE  
Version 2, June 1991
```

```
Copyright (C) 1989, 1991 Free Software Foundation, Inc.,  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.  
[for rest of text, see following link]
```

This component is licensed under [Section A.13, "GNU General Public License Version 2.0, June 1991"](#)

## A.41. RegEX-Spencer Library License

The following software may be included in this product: Henry Spencer's Regular-Expression Library (RegEX-Spencer)

```
Copyright 1992, 1993, 1994 Henry Spencer. All rights reserved.
This software is not subject to any license of the American Telephone
and Telegraph Company or of the Regents of the University of California.
```

```
Permission is granted to anyone to use this software for any purpose on
any computer system, and to alter it and redistribute it, subject
to the following restrictions:
```

1. The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
2. The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
4. This notice may not be removed or altered.

## A.42. RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License

The following software may be included in this product:

RFC 3174 - US Secure Hash Algorithm 1 (SHA1)

```
RFC 3174 - US Secure Hash Algorithm 1 (SHA1)
```

```
Copyright (C) The Internet Society (2001). All Rights Reserved.
```

```
This document and translations of it may be copied and furnished to
others, and derivative works that comment on or otherwise explain it
or assist in its implementation may be prepared, copied, published
and distributed, in whole or in part, without restriction of any
kind, provided that the above copyright notice and this paragraph are
included on all such copies and derivative works. However, this
document itself may not be modified in any way, such as by removing
the copyright notice or references to the Internet Society or other
Internet organizations, except as needed for the purpose of
developing Internet standards in which case the procedures for
copyrights defined in the Internet Standards process must be
followed, or as required to translate it into languages other than
English.
```

```
The limited permissions granted above are perpetual and will not be
revoked by the Internet Society or its successors or assigns.
```

```
This document and the information contained herein is provided on an
"AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING
TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING
BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION
HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF
MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.
```

```
Acknowledgement
```

```
Funding for the RFC Editor function is currently provided by the
Internet Society.
```

## A.43. Richard A. O'Keefe String Library License

The following software may be included in this product:

Richard A. O'Keefe String Library

The Richard O'Keefe String Library is subject to the following notice:

These files are in the public domain. This includes getopt.c, which is the work of Henry Spencer, University of Toronto Zoology, who says of it "None of this software is derived from Bell software. I had no access to the source for Bell's versions at the time I wrote it. This software is hereby explicitly placed in the public domain. It may be used for any purpose on any machine by anyone." I would greatly prefer it if \*my\* material received no military use.

The t\_ctype.h file is subject to the following notice:

Copyright (C) 1998, 1999 by Pruet Boonma, all rights reserved.  
Copyright (C) 1998 by Theppitak Karoonboonyanan, all rights reserved.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies.

Smaphan Raruenrom and Pruet Boonma makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

## A.44. SHA-1 in C License

The following software may be included in this product:

SHA-1 in C

```
SHA-1 in C
By Steve Reid <steve@edmweb.com>
100% Public Domain
```

## A.45. Simple Logging Facade for Java (SLF4J) License

The following software may be included in this product:

Simple Logging Facade for Java (SLF4J)

Copyright (c) 2004-2008 QOS.ch  
All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## A.46. Unicode Data Files

The following software may be included in this product:

## Unicode Data Files

Copyright 2001-2009 Unicode, Inc.

### Disclaimer

This source code is provided as is by Unicode, Inc. No claims are made as to fitness for any particular purpose. No warranties of any kind are expressed or implied. The recipient agrees to determine applicability of information provided. If this file has been purchased on magnetic or optical media from Unicode, Inc., the sole remedy for any claim will be exchange of defective media within 90 days of receipt.

### Limitations on Rights to Redistribute This Code

Unicode, Inc. hereby grants the right to freely use the information supplied in this file in the creation of products supporting the Unicode Standard, and to make copies of this file in any form for internal or external distribution as long as this notice remains attached.

## A.47. zlib License

The following software may be included in this product:

zlib

Oracle gratefully acknowledges the contributions of Jean-loup Gailly and Mark Adler in creating the zlib general purpose compression library which is used in this product.

```
zlib.h -- interface of the 'zlib' general purpose compression library
Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler
```

```
zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.3, July 18th, 2005
Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler
```

```
zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.5, April 19th, 2010
Copyright (C) 1995-2010 Jean-loup Gailly and Mark Adler
```

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly [jloup@gzip.org](mailto:jloup@gzip.org)  
Mark Adler [madler@alumni.caltech.edu](mailto:madler@alumni.caltech.edu)

## A.48. ZLIB.NET License

The following software may be included in this product:

ZLIB.NET

Copyright (c) 2006-2007, ComponentAce  
<http://www.componentace.com>  
All rights reserved.



Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of ComponentAce nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



---

# Appendix B. MySQL 5.6 Frequently Asked Questions

## Table of Contents

B.1. MySQL 5.6 FAQ: General .....	3627
B.2. MySQL 5.6 FAQ: Storage Engines .....	3629
B.3. MySQL 5.6 FAQ: Server SQL Mode .....	3630
B.4. MySQL 5.6 FAQ: Stored Procedures and Functions .....	3631
B.5. MySQL 5.6 FAQ: Triggers .....	3634
B.6. MySQL 5.6 FAQ: Views .....	3637
B.7. MySQL 5.6 FAQ: <a href="#">INFORMATION_SCHEMA</a> .....	3637
B.8. MySQL 5.6 FAQ: Migration .....	3638
B.9. MySQL 5.6 FAQ: Security .....	3639
B.10. MySQL 5.6 FAQ: MySQL Cluster .....	3640
B.11. MySQL 5.6 FAQ: MySQL Chinese, Japanese, and Korean Character Sets .....	3640
B.12. MySQL 5.6 FAQ: Connectors & APIs .....	3652
B.13. MySQL 5.6 FAQ: Replication .....	3652
B.14. MySQL 5.6 FAQ: MySQL Enterprise Scalability Thread Pool .....	3656

## B.1. MySQL 5.6 FAQ: General

### Questions

- [B.1.1: \[3627\]](#) Which version of MySQL is production-ready (GA)?
- [B.1.2: \[3628\]](#) What is the state of development (non-GA) versions?
- [B.1.3: \[3628\]](#) Can MySQL 5.6 do subqueries?
- [B.1.4: \[3628\]](#) Can MySQL 5.6 perform multiple-table inserts, updates, and deletes?
- [B.1.5: \[3628\]](#) Does MySQL 5.6 have a Query Cache? Does it work on Server, Instance or Database?
- [B.1.6: \[3628\]](#) Does MySQL 5.6 have Sequences?
- [B.1.7: \[3628\]](#) Does MySQL 5.6 have a `NOW( ) [1220]` function with fractions of seconds?
- [B.1.8: \[3628\]](#) Does MySQL 5.6 work with multi-core processors?
- [B.1.9: \[3628\]](#) Why do I see multiple processes for `mysqld`?
- [B.1.10: \[3629\]](#) Have there been there any improvements in error reporting when foreign keys fail? Does MySQL now report which column and reference failed?
- [B.1.11: \[3629\]](#) Can MySQL 5.6 perform ACID transactions?

### Questions and Answers

#### B.1.1: Which version of MySQL is production-ready (GA)?

MySQL 5.6, MySQL 5.5, MySQL 5.1, and MySQL 5.0 are supported for production use.

MySQL 5.6 achieved General Availability (GA) status with MySQL 5.6.10, which was released for production use on 5 February 2013.

MySQL 5.5 achieved General Availability (GA) status with MySQL 5.5.8, which was released for production use on 3 December 2010.

MySQL 5.1 achieved General Availability (GA) status with MySQL 5.1.30, which was released for production use on 14 November 2008.

MySQL 5.0 achieved General Availability (GA) status with MySQL 5.0.15, which was released for production use on 19 October 2005. Note that active development for MySQL 5.0 has ended.

#### **B.1.2: What is the state of development (non-GA) versions?**

MySQL follows a milestone release model that introduces pre-production-quality features and stabilizes them to release quality (see <http://dev.mysql.com/doc/mysql-development-cycle/en/index.html>). This process then repeats, so releases cycle between pre-production and release quality status. Please check the change logs to identify the status of a given release.

MySQL 5.4 was a development series. Work on this series has ceased.

MySQL 5.7 is being actively developed using the milestone release methodology described above.

MySQL 6.0 was a development series. Work on this series has ceased.

#### **B.1.3: Can MySQL 5.6 do subqueries?**

Yes. See [Section 13.2.10, “Subquery Syntax”](#).

#### **B.1.4: Can MySQL 5.6 perform multiple-table inserts, updates, and deletes?**

Yes. For the syntax required to perform multiple-table updates, see [Section 13.2.11, “UPDATE Syntax”](#); for that required to perform multiple-table deletes, see [Section 13.2.2, “DELETE Syntax”](#).

A multiple-table insert can be accomplished using a trigger whose `FOR EACH ROW` clause contains multiple `INSERT` statements within a `BEGIN . . . END` block. See [Section 18.3, “Using Triggers”](#).

#### **B.1.5: Does MySQL 5.6 have a Query Cache? Does it work on Server, Instance or Database?**

Yes. The query cache operates on the server level, caching complete result sets matched with the original query string. If an exactly identical query is made (which often happens, particularly in web applications), no parsing or execution is necessary; the result is sent directly from the cache. Various tuning options are available. See [Section 8.9.3, “The MySQL Query Cache”](#).

#### **B.1.6: Does MySQL 5.6 have Sequences?**

No. However, MySQL has an `AUTO_INCREMENT` system, which in MySQL 5.6 can also handle inserts in a multi-master replication setup. With the `auto_increment_increment` [1965] and `auto_increment_offset` [1968] system variables, you can set each server to generate auto-increment values that don't conflict with other servers. The `auto_increment_increment` [1965] value should be greater than the number of servers, and each server should have a unique offset.

#### **B.1.7: Does MySQL 5.6 have a `NOW()` [1220] function with fractions of seconds?**

No. This is on the MySQL roadmap as a “rolling feature”. This means that it is not a flagship feature, but will be implemented, development time permitting. Specific customer demand may change this scheduling.

However, MySQL does parse time strings with a fractional component. See [Section 11.3.2, “The `TIME` Type”](#).

#### **B.1.8: Does MySQL 5.6 work with multi-core processors?**

Yes. MySQL is fully multi-threaded, and will make use of multiple CPUs, provided that the operating system supports them.

#### **B.1.9: Why do I see multiple processes for `mysqld`?**

When using LinuxThreads, you should see a minimum of three `mysqld` processes running. These are in fact threads. There is one thread for the LinuxThreads manager, one thread to handle connections, and one thread to handle alarms and signals.

**B.1.10: Have there been any improvements in error reporting when foreign keys fail? Does MySQL now report which column and reference failed?**

The foreign key support in `InnoDB` has seen improvements in each major version of MySQL. Foreign key support generic to all storage engines is scheduled for MySQL 6.x; this should resolve any inadequacies in the current storage engine specific implementation.

**B.1.11: Can MySQL 5.6 perform ACID transactions?**

Yes. All current MySQL versions support transactions. The `InnoDB` storage engine offers full ACID transactions with row-level locking, multi-versioning, nonlocking repeatable reads, and all four SQL standard isolation levels.

The `NDB` storage engine supports the `READ COMMITTED` [1469] transaction isolation level only.

## B.2. MySQL 5.6 FAQ: Storage Engines

### Questions

- **B.2.1:** [3629] Where can I obtain complete documentation for MySQL storage engines?
- **B.2.2:** [3629] Are there any new storage engines in MySQL 5.6?
- **B.2.3:** [3629] Have any storage engines been removed in MySQL 5.6?
- **B.2.4:** [3629] What are the unique benefits of the `ARCHIVE` storage engine?
- **B.2.5:** [3629] Do the new features in MySQL 5.6 apply to all storage engines?

### Questions and Answers

**B.2.1: Where can I obtain complete documentation for MySQL storage engines?**

See [Chapter 14, Storage Engines](#). That chapter contains information about all MySQL storage engines except for the `NDB` storage engine used for MySQL Cluster; `NDB` is covered in [MySQL Cluster NDB 7.2](#).

**B.2.2: Are there any new storage engines in MySQL 5.6?**

The features from the optional `InnoDB Plugin` from MySQL 5.1 are folded into the built-in `InnoDB` storage engine, so you can take advantage of features such as the Barracuda file format, `InnoDB` table compression, and the new configuration options for performance. See [Section 14.2, “The InnoDB Storage Engine”](#) for details. `InnoDB` also becomes the default storage engine for new tables. See [Section 14.2.1.1, “InnoDB as the Default MySQL Storage Engine”](#) for details.

**B.2.3: Have any storage engines been removed in MySQL 5.6?**

No.

**B.2.4: What are the unique benefits of the `ARCHIVE` storage engine?**

The `ARCHIVE` storage engine is ideally suited for storing large amounts of data without indexes; it has a very small footprint, and performs selects using table scans. See [Section 14.6, “The ARCHIVE Storage Engine”](#), for details.

**B.2.5: Do the new features in MySQL 5.6 apply to all storage engines?**

The general new features such as views, stored procedures, triggers, `INFORMATION_SCHEMA`, precision math (`DECIMAL` column type), and the `BIT` column type, apply to all storage engines. There are also additions and changes for specific storage engines.

## B.3. MySQL 5.6 FAQ: Server SQL Mode

### Questions

- [B.3.1: \[3630\]](#) What are server SQL modes?
- [B.3.2: \[3630\]](#) How many server SQL modes are there?
- [B.3.3: \[3630\]](#) How do you determine the server SQL mode?
- [B.3.4: \[3630\]](#) Is the mode dependent on the database or connection?
- [B.3.5: \[3630\]](#) Can the rules for strict mode be extended?
- [B.3.6: \[3630\]](#) Does strict mode impact performance?
- [B.3.7: \[3630\]](#) What is the default server SQL mode when My SQL 5.6 is installed?

### Questions and Answers

#### B.3.1: What are server SQL modes?

Server SQL modes define what SQL syntax MySQL should support and what kind of data validation checks it should perform. This makes it easier to use MySQL in different environments and to use MySQL together with other database servers. The MySQL Server apply these modes individually to different clients. For more information, see [Section 5.1.7, “Server SQL Modes”](#).

#### B.3.2: How many server SQL modes are there?

Each mode can be independently switched on and off. See [Section 5.1.7, “Server SQL Modes”](#), for a complete list of available modes.

#### B.3.3: How do you determine the server SQL mode?

You can set the default SQL mode (for `mysqld` startup) with the `--sql-mode [442]` option. Using the statement `SET [GLOBAL|SESSION] sql_mode='modes'`, you can change the settings from within a connection, either locally to the connection, or to take effect globally. You can retrieve the current mode by issuing a `SELECT @@sql_mode` statement.

#### B.3.4: Is the mode dependent on the database or connection?

A mode is not linked to a particular database. Modes can be set locally to the session (connection), or globally for the server. you can change these settings using `SET [GLOBAL|SESSION] sql_mode='modes'`.

#### B.3.5: Can the rules for strict mode be extended?

When we refer to *strict mode*, we mean a mode where at least one of the modes `TRADITIONAL [609]`, `STRICT_TRANS_TABLES [607]`, or `STRICT_ALL_TABLES [607]` is enabled. Options can be combined, so you can add restrictions to a mode. See [Section 5.1.7, “Server SQL Modes”](#), for more information.

#### B.3.6: Does strict mode impact performance?

The intensive validation of input data that some settings requires more time than if the validation is not done. While the performance impact is not that great, if you do not require such validation (perhaps your application already handles all of this), then MySQL gives you the option of leaving strict mode disabled. However—if you do require it—strict mode can provide such validation.

#### B.3.7: What is the default server SQL mode when My SQL 5.6 is installed?

By default, no special modes are enabled. See [Section 5.1.7, “Server SQL Modes”](#), for information about all available modes and MySQL's default behavior.

## B.4. MySQL 5.6 FAQ: Stored Procedures and Functions

### Questions

- [B.4.1: \[3632\]](#) Does MySQL 5.6 support stored procedures and functions?
- [B.4.2: \[3632\]](#) Where can I find documentation for MySQL stored procedures and stored functions?
- [B.4.3: \[3632\]](#) Is there a discussion forum for MySQL stored procedures?
- [B.4.4: \[3632\]](#) Where can I find the ANSI SQL 2003 specification for stored procedures?
- [B.4.5: \[3632\]](#) How do you manage stored routines?
- [B.4.6: \[3632\]](#) Is there a way to view all stored procedures and stored functions in a given database?
- [B.4.7: \[3632\]](#) Where are stored procedures stored?
- [B.4.8: \[3632\]](#) Is it possible to group stored procedures or stored functions into packages?
- [B.4.9: \[3632\]](#) Can a stored procedure call another stored procedure?
- [B.4.10: \[3632\]](#) Can a stored procedure call a trigger?
- [B.4.11: \[3632\]](#) Can a stored procedure access tables?
- [B.4.12: \[3633\]](#) Do stored procedures have a statement for raising application errors?
- [B.4.13: \[3633\]](#) Do stored procedures provide exception handling?
- [B.4.14: \[3633\]](#) Can MySQL 5.6 stored routines return result sets?
- [B.4.15: \[3633\]](#) Is `WITH RECOMPILE` supported for stored procedures?
- [B.4.16: \[3633\]](#) Is there a MySQL equivalent to using `mod_plsql` as a gateway on Apache to talk directly to a stored procedure in the database?
- [B.4.17: \[3633\]](#) Can I pass an array as input to a stored procedure?
- [B.4.18: \[3633\]](#) Can I pass a cursor as an `IN` parameter to a stored procedure?
- [B.4.19: \[3633\]](#) Can I return a cursor as an `OUT` parameter from a stored procedure?
- [B.4.20: \[3633\]](#) Can I print out a variable's value within a stored routine for debugging purposes?
- [B.4.21: \[3633\]](#) Can I commit or roll back transactions inside a stored procedure?
- [B.4.22: \[3633\]](#) Do MySQL 5.6 stored procedures and functions work with replication?
- [B.4.23: \[3633\]](#) Are stored procedures and functions created on a master server replicated to a slave?
- [B.4.24: \[3634\]](#) How are actions that take place inside stored procedures and functions replicated?
- [B.4.25: \[3634\]](#) Are there special security requirements for using stored procedures and functions together with replication?
- [B.4.26: \[3634\]](#) What limitations exist for replicating stored procedure and function actions?
- [B.4.27: \[3634\]](#) Do the preceding limitations affect MySQL's ability to do point-in-time recovery?
- [B.4.28: \[3634\]](#) What is being done to correct the aforementioned limitations?

### Questions and Answers

**B.4.1: Does MySQL 5.6 support stored procedures and functions?**

Yes. MySQL 5.6 supports two types of stored routines—stored procedures and stored functions.

**B.4.2: Where can I find documentation for MySQL stored procedures and stored functions?**

See [Section 18.2, “Using Stored Routines \(Procedures and Functions\)”](#).

**B.4.3: Is there a discussion forum for MySQL stored procedures?**

Yes. See <http://forums.mysql.com/list.php?98>.

**B.4.4: Where can I find the ANSI SQL 2003 specification for stored procedures?**

Unfortunately, the official specifications are not freely available (ANSI makes them available for purchase). However, there are books—such as *SQL-99 Complete, Really* by Peter Gulutzan and Trudy Pelzer—which give a comprehensive overview of the standard, including coverage of stored procedures.

**B.4.5: How do you manage stored routines?**

It is always good practice to use a clear naming scheme for your stored routines. You can manage stored procedures with `CREATE [FUNCTION|PROCEDURE]`, `ALTER [FUNCTION|PROCEDURE]`, `DROP [FUNCTION|PROCEDURE]`, and `SHOW CREATE [FUNCTION|PROCEDURE]`. You can obtain information about existing stored procedures using the `ROUTINES` table in the `INFORMATION_SCHEMA` database (see [Section 19.19, “The `INFORMATION\_SCHEMA` `ROUTINES` Table”](#)).

**B.4.6: Is there a way to view all stored procedures and stored functions in a given database?**

Yes. For a database named `dbname`, use this query on the `INFORMATION_SCHEMA.ROUTINES` table:

```
SELECT ROUTINE_TYPE, ROUTINE_NAME
FROM INFORMATION_SCHEMA.ROUTINES
WHERE ROUTINE_SCHEMA='dbname' ;
```

For more information, see [Section 19.19, “The `INFORMATION\_SCHEMA` `ROUTINES` Table”](#).

The body of a stored routine can be viewed using `SHOW CREATE FUNCTION` (for a stored function) or `SHOW CREATE PROCEDURE` (for a stored procedure). See [Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”](#), for more information.

**B.4.7: Where are stored procedures stored?**

In the `proc` table of the `mysql` system database. However, you should not access the tables in the system database directly. Instead, use `SHOW CREATE FUNCTION` to obtain information about stored functions, and `SHOW CREATE PROCEDURE` to obtain information about stored procedures. See [Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”](#), for more information about these statements.

You can also query the `ROUTINES` table in the `INFORMATION_SCHEMA` database—see [Section 19.19, “The `INFORMATION\_SCHEMA` `ROUTINES` Table”](#), for information about this table.

**B.4.8: Is it possible to group stored procedures or stored functions into packages?**

No. This is not supported in MySQL 5.6.

**B.4.9: Can a stored procedure call another stored procedure?**

Yes.

**B.4.10: Can a stored procedure call a trigger?**

A stored procedure can execute an SQL statement, such as an `UPDATE`, that causes a trigger to activate.

**B.4.11: Can a stored procedure access tables?**



Yes. A stored procedure can access one or more tables as required.

**B.4.12: Do stored procedures have a statement for raising application errors?**

Yes. MySQL 5.6 implements the SQL standard [SIGNAL](#) and [RESIGNAL](#) statements. See [Section 13.6.7, “Condition Handling”](#).

**B.4.13: Do stored procedures provide exception handling?**

MySQL implements [HANDLER](#) definitions according to the SQL standard. See [Section 13.6.7.2, “DECLARE . . . HANDLER Syntax”](#), for details.

**B.4.14: Can MySQL 5.6 stored routines return result sets?**

*Stored procedures* can, but stored functions cannot. If you perform an ordinary [SELECT](#) inside a stored procedure, the result set is returned directly to the client. You need to use the MySQL 4.1 (or above) client/server protocol for this to work. This means that—for instance—in PHP, you need to use the [mysqli](#) extension rather than the old [mysql](#) extension.

**B.4.15: Is [WITH RECOMPILE](#) supported for stored procedures?**

Not in MySQL 5.6.

**B.4.16: Is there a MySQL equivalent to using [mod\\_plsql](#) as a gateway on Apache to talk directly to a stored procedure in the database?**

There is no equivalent in MySQL 5.6.

**B.4.17: Can I pass an array as input to a stored procedure?**

Not in MySQL 5.6.

**B.4.18: Can I pass a cursor as an [IN](#) parameter to a stored procedure?**

In MySQL 5.6, cursors are available inside stored procedures only.

**B.4.19: Can I return a cursor as an [OUT](#) parameter from a stored procedure?**

In MySQL 5.6, cursors are available inside stored procedures only. However, if you do not open a cursor on a [SELECT](#), the result will be sent directly to the client. You can also [SELECT INTO](#) variables. See [Section 13.2.9, “SELECT Syntax”](#).

**B.4.20: Can I print out a variable's value within a stored routine for debugging purposes?**

Yes, you can do this in a *stored procedure*, but not in a stored function. If you perform an ordinary [SELECT](#) inside a stored procedure, the result set is returned directly to the client. You will need to use the MySQL 4.1 (or above) client/server protocol for this to work. This means that—for instance—in PHP, you need to use the [mysqli](#) extension rather than the old [mysql](#) extension.

**B.4.21: Can I commit or roll back transactions inside a stored procedure?**

Yes. However, you cannot perform transactional operations within a stored function.

**B.4.22: Do MySQL 5.6 stored procedures and functions work with replication?**

Yes, standard actions carried out in stored procedures and functions are replicated from a master MySQL server to a slave server. There are a few limitations that are described in detail in [Section 18.7, “Binary Logging of Stored Programs”](#).

**B.4.23: Are stored procedures and functions created on a master server replicated to a slave?**

Yes, creation of stored procedures and functions carried out through normal DDL statements on a master server are replicated to a slave, so the objects will exist on both servers. [ALTER](#) and [DROP](#) statements for stored procedures and functions are also replicated.

**B.4.24: How are actions that take place inside stored procedures and functions replicated?**

MySQL records each DML event that occurs in a stored procedure and replicates those individual actions to a slave server. The actual calls made to execute stored procedures are not replicated.

Stored functions that change data are logged as function invocations, not as the DML events that occur inside each function.

**B.4.25: Are there special security requirements for using stored procedures and functions together with replication?**

Yes. Because a slave server has authority to execute any statement read from a master's binary log, special security constraints exist for using stored functions with replication. If replication or binary logging in general (for the purpose of point-in-time recovery) is active, then MySQL DBAs have two security options open to them:

1. Any user wishing to create stored functions must be granted the `SUPER` [756] privilege.
2. Alternatively, a DBA can set the `log_bin_trust_function_creators` [497] system variable to 1, which enables anyone with the standard `CREATE ROUTINE` [754] privilege to create stored functions.

**B.4.26: What limitations exist for replicating stored procedure and function actions?**

Nondeterministic (random) or time-based actions embedded in stored procedures may not replicate properly. By their very nature, randomly produced results are not predictable and cannot be exactly reproduced, and therefore, random actions replicated to a slave will not mirror those performed on a master. Note that declaring stored functions to be `DETERMINISTIC` or setting the `log_bin_trust_function_creators` [497] system variable to 0 will not allow random-valued operations to be invoked.

In addition, time-based actions cannot be reproduced on a slave because the timing of such actions in a stored procedure is not reproducible through the binary log used for replication. It records only DML events and does not factor in timing constraints.

Finally, nontransactional tables for which errors occur during large DML actions (such as bulk inserts) may experience replication issues in that a master may be partially updated from DML activity, but no updates are done to the slave because of the errors that occurred. A workaround is for a function's DML actions to be carried out with the `IGNORE` keyword so that updates on the master that cause errors are ignored and updates that do not cause errors are replicated to the slave.

**B.4.27: Do the preceding limitations affect MySQL's ability to do point-in-time recovery?**

The same limitations that affect replication do affect point-in-time recovery.

**B.4.28: What is being done to correct the aforementioned limitations?**

You can choose either statement-based replication or row-based replication. The original replication implementation is based on statement-based binary logging. Row-based binary logging resolves the limitations mentioned earlier.

*Mixed* replication is also available (by starting the server with `--binlog-format=mixed` [415]). This hybrid, "smart" form of replication "knows" whether statement-level replication can safely be used, or row-level replication is required.

For additional information, see [Section 16.1.2, "Replication Formats"](#).

## B.5. MySQL 5.6 FAQ: Triggers

### Questions

- [B.5.1: \[3635\]](#) Where can I find the documentation for MySQL 5.6 triggers?

- [B.5.2: \[3635\]](#) Is there a discussion forum for MySQL Triggers?
- [B.5.3: \[3635\]](#) Does MySQL 5.6 have statement-level or row-level triggers?
- [B.5.4: \[3635\]](#) Are there any default triggers?
- [B.5.5: \[3635\]](#) How are triggers managed in MySQL?
- [B.5.6: \[3635\]](#) Is there a way to view all triggers in a given database?
- [B.5.7: \[3636\]](#) Where are triggers stored?
- [B.5.8: \[3636\]](#) Can a trigger call a stored procedure?
- [B.5.9: \[3636\]](#) Can triggers access tables?
- [B.5.10: \[3636\]](#) Can triggers call an external application through a UDF?
- [B.5.11: \[3636\]](#) Is it possible for a trigger to update tables on a remote server?
- [B.5.12: \[3636\]](#) Do triggers work with replication?
- [B.5.13: \[3636\]](#) How are actions carried out through triggers on a master replicated to a slave?

### Questions and Answers

#### **B.5.1: Where can I find the documentation for MySQL 5.6 triggers?**

See [Section 18.3, “Using Triggers”](#).

#### **B.5.2: Is there a discussion forum for MySQL Triggers?**

Yes. It is available at <http://forums.mysql.com/list.php?99>.

#### **B.5.3: Does MySQL 5.6 have statement-level or row-level triggers?**

In MySQL 5.6, all triggers are `FOR EACH ROW`—that is, the trigger is activated for each row that is inserted, updated, or deleted. MySQL 5.6 does not support triggers using `FOR EACH STATEMENT`.

#### **B.5.4: Are there any default triggers?**

Not explicitly. MySQL does have specific special behavior for some `TIMESTAMP` columns, as well as for columns which are defined using `AUTO_INCREMENT`.

#### **B.5.5: How are triggers managed in MySQL?**

In MySQL 5.6, triggers can be created using the `CREATE TRIGGER` statement, and dropped using `DROP TRIGGER`. See [Section 13.1.15, “CREATE TRIGGER Syntax”](#), and [Section 13.1.24, “DROP TRIGGER Syntax”](#), for more about these statements.

Information about triggers can be obtained by querying the `INFORMATION_SCHEMA.TRIGGERS` table. See [Section 19.27, “The INFORMATION\\_SCHEMA TRIGGERS Table”](#).

#### **B.5.6: Is there a way to view all triggers in a given database?**

Yes. You can obtain a listing of all triggers defined on database `dbname` using a query on the `INFORMATION_SCHEMA.TRIGGERS` table such as the one shown here:

```
SELECT TRIGGER_NAME, EVENT_MANIPULATION, EVENT_OBJECT_TABLE, ACTION_STATEMENT
FROM INFORMATION_SCHEMA.TRIGGERS
WHERE TRIGGER_SCHEMA='dbname' ;
```

For more information about this table, see [Section 19.27, “The INFORMATION\\_SCHEMA TRIGGERS Table”](#).

You can also use the `SHOW TRIGGERS` statement, which is specific to MySQL. See [Section 13.7.5.39](#), “`SHOW TRIGGERS Syntax`”.

**B.5.7: Where are triggers stored?**

Triggers for a table are currently stored in `.TRG` files, with one such file one per table.

**B.5.8: Can a trigger call a stored procedure?**

Yes.

**B.5.9: Can triggers access tables?**

A trigger can access both old and new data in its own table. A trigger can also affect other tables, but it is not permitted to modify a table that is already being used (for reading or writing) by the statement that invoked the function or trigger.

**B.5.10: Can triggers call an external application through a UDF?**

Yes. For example, a trigger could invoke the `sys_exec()` UDF.

**B.5.11: Is it possible for a trigger to update tables on a remote server?**

Yes. A table on a remote server could be updated using the `FEDERATED` storage engine. (See [Section 14.9](#), “`The FEDERATED Storage Engine`”).

**B.5.12: Do triggers work with replication?**

Yes. However, the way in which they work depends whether you are using MySQL's “classic” statement-based replication available in all versions of MySQL, or the row-based replication format introduced in MySQL 5.1.

When using statement-based replication, triggers on the slave are executed by statements that are executed on the master (and replicated to the slave).

When using row-based replication, triggers are not executed on the slave due to statements that were run on the master and then replicated to the slave. Instead, when using row-based replication, the changes caused by executing the trigger on the master are applied on the slave.

For more information, see [Section 16.4.1.30](#), “`Replication and Triggers`”.

**B.5.13: How are actions carried out through triggers on a master replicated to a slave?**

Again, this depends on whether you are using statement-based or row-based replication.

**Statement-based replication.** First, the triggers that exist on a master must be re-created on the slave server. Once this is done, the replication flow works as any other standard DML statement that participates in replication. For example, consider a table `EMP` that has an `AFTER` insert trigger, which exists on a master MySQL server. The same `EMP` table and `AFTER` insert trigger exist on the slave server as well. The replication flow would be:

1. An `INSERT` statement is made to `EMP`.
2. The `AFTER` trigger on `EMP` activates.
3. The `INSERT` statement is written to the binary log.
4. The replication slave picks up the `INSERT` statement to `EMP` and executes it.
5. The `AFTER` trigger on `EMP` that exists on the slave activates.

**Row-based replication.** When you use row-based replication, the changes caused by executing the trigger on the master are applied on the slave. However, the triggers themselves are not actually executed on the slave under row-based replication. This is because, if both the master and the slave

applied the changes from the master and—in addition—the trigger causing these changes were applied on the slave, the changes would in effect be applied twice on the slave, leading to different data on the master and the slave.

In most cases, the outcome is the same for both row-based and statement-based replication. However, if you use different triggers on the master and slave, you cannot use row-based replication. (This is because the row-based format replicates the changes made by triggers executing on the master to the slaves, rather than the statements that caused the triggers to execute, and the corresponding triggers on the slave are not executed.) Instead, any statements causing such triggers to be executed must be replicated using statement-based replication.

For more information, see [Section 16.4.1.30, “Replication and Triggers”](#).

## B.6. MySQL 5.6 FAQ: Views

### Questions

- [B.6.1: \[3637\]](#) Where can I find documentation covering MySQL Views?
- [B.6.2: \[3637\]](#) Is there a discussion forum for MySQL Views?
- [B.6.3: \[3637\]](#) What happens to a view if an underlying table is dropped or renamed?
- [B.6.4: \[3637\]](#) Does MySQL 5.6 have table snapshots?
- [B.6.5: \[3637\]](#) Does MySQL 5.6 have materialized views?
- [B.6.6: \[3637\]](#) Can you insert into views that are based on joins?

### Questions and Answers

#### **B.6.1: Where can I find documentation covering MySQL Views?**

See [Section 18.5, “Using Views”](#).

#### **B.6.2: Is there a discussion forum for MySQL Views?**

Yes. See <http://forums.mysql.com/list.php?100>

#### **B.6.3: What happens to a view if an underlying table is dropped or renamed?**

After a view has been created, it is possible to drop or alter a table or view to which the definition refers. To check a view definition for problems of this kind, use the `CHECK TABLE` statement. (See [Section 13.7.2.2, “CHECK TABLE Syntax”](#).)

#### **B.6.4: Does MySQL 5.6 have table snapshots?**

No.

#### **B.6.5: Does MySQL 5.6 have materialized views?**

No.

#### **B.6.6: Can you insert into views that are based on joins?**

It is possible, provided that your `INSERT` statement has a column list that makes it clear there is only one table involved.

You *cannot* insert into multiple tables with a single insert on a view.

## B.7. MySQL 5.6 FAQ: [INFORMATION\\_SCHEMA](#)

### Questions

- **B.7.1:** [3638] Where can I find documentation for the MySQL `INFORMATION_SCHEMA` database?
- **B.7.2:** [3638] Is there a discussion forum for `INFORMATION_SCHEMA`?
- **B.7.3:** [3638] Where can I find the ANSI SQL 2003 specification for `INFORMATION_SCHEMA`?
- **B.7.4:** [3638] What is the difference between the Oracle Data Dictionary and MySQL's `INFORMATION_SCHEMA`?
- **B.7.5:** [3638] Can I add to or otherwise modify the tables found in the `INFORMATION_SCHEMA` database?

### Questions and Answers

#### **B.7.1: Where can I find documentation for the MySQL `INFORMATION_SCHEMA` database?**

See [Chapter 19, `INFORMATION\_SCHEMA` Tables](#)

#### **B.7.2: Is there a discussion forum for `INFORMATION_SCHEMA`?**

See <http://forums.mysql.com/list.php?101>.

#### **B.7.3: Where can I find the ANSI SQL 2003 specification for `INFORMATION_SCHEMA`?**

Unfortunately, the official specifications are not freely available. (ANSI makes them available for purchase.) However, there are books available—such as *SQL-99 Complete, Really* by Peter Gulutzan and Trudy Pelzer—which give a comprehensive overview of the standard, including `INFORMATION_SCHEMA`.

#### **B.7.4: What is the difference between the Oracle Data Dictionary and MySQL's `INFORMATION_SCHEMA`?**

Both Oracle and MySQL provide metadata in tables. However, Oracle and MySQL use different table names and column names. MySQL's implementation is more similar to those found in DB2 and SQL Server, which also support `INFORMATION_SCHEMA` as defined in the SQL standard.

#### **B.7.5: Can I add to or otherwise modify the tables found in the `INFORMATION_SCHEMA` database?**

No. Since applications may rely on a certain standard structure, this should not be modified. For this reason, *we cannot support bugs or other issues which result from modifying `INFORMATION_SCHEMA` tables or data.*

## B.8. MySQL 5.6 FAQ: Migration

### Questions

- **B.8.1:** [3638] Where can I find information on how to migrate from MySQL 5.5 to MySQL 5.6?
- **B.8.2:** [3639] How has storage engine (table type) support changed in MySQL 5.6 from previous versions?

### Questions and Answers

#### **B.8.1: Where can I find information on how to migrate from MySQL 5.5 to MySQL 5.6?**

For detailed upgrade information, see [Section 2.11.1, “Upgrading MySQL”](#). Do not skip a major version when upgrading, but rather complete the process in steps, upgrading from one major version to the next in each step. This may seem more complicated, but it will you save time and trouble—if you encounter problems during the upgrade, their origin will be easier to identify, either by you or—if you have a MySQL Enterprise subscription—by MySQL support.

**B.8.2: How has storage engine (table type) support changed in MySQL 5.6 from previous versions?**

Storage engine support has changed as follows:

- Support for [ISAM](#) tables was removed in MySQL 5.0 and you should now use the [MyISAM](#) storage engine in place of [ISAM](#). To convert a table *tblname* from [ISAM](#) to [MyISAM](#), simply issue a statement such as this one:

```
ALTER TABLE tblname ENGINE=MYISAM;
```

- Internal [RAID](#) for [MyISAM](#) tables was also removed in MySQL 5.0. This was formerly used to allow large tables in file systems that did not support file sizes greater than 2GB. All modern file systems allow for larger tables; in addition, there are now other solutions such as [MERGE](#) tables and views.
- The [VARCHAR](#) column type now retains trailing spaces in all storage engines.
- [MEMORY](#) tables (formerly known as [HEAP](#) tables) can also contain [VARCHAR](#) columns.

## B.9. MySQL 5.6 FAQ: Security

### Questions

- [B.9.1: \[3639\]](#) Where can I find documentation that addresses security issues for MySQL?
- [B.9.2: \[3639\]](#) Does MySQL 5.6 have native support for SSL?
- [B.9.3: \[3640\]](#) Is SSL support be built into MySQL binaries, or must I recompile the binary myself to enable it?
- [B.9.4: \[3640\]](#) Does MySQL 5.6 have built-in authentication against LDAP directories?
- [B.9.5: \[3640\]](#) Does MySQL 5.6 include support for Roles Based Access Control (RBAC)?

### Questions and Answers

**B.9.1: Where can I find documentation that addresses security issues for MySQL?**

The best place to start is [Chapter 6, Security](#).

Other portions of the MySQL Documentation which you may find useful with regard to specific security concerns include the following:

- [Section 6.1.1, "Security Guidelines"](#).
- [Section 6.1.3, "Making MySQL Secure Against Attackers"](#).
- [Section C.5.4.1, "How to Reset the Root Password"](#).
- [Section 6.1.5, "How to Run MySQL as a Normal User"](#).
- [Section 22.3.2.6, "User-Defined Function Security Precautions"](#).
- [Section 6.1.4, "Security-Related `mysqld` Options and Variables"](#).
- [Section 6.1.6, "Security Issues with `LOAD DATA LOCAL`"](#).
- [Section 2.10, "Postinstallation Setup and Testing"](#).
- [Section 6.3.9.1, "Basic SSL Concepts"](#).

**B.9.2: Does MySQL 5.6 have native support for SSL?**

Most 5.6 binaries have support for SSL connections between the client and server. See [Section 6.3.9, "Using SSL for Secure Connections"](#).



You can also tunnel a connection using SSH, if (for example) the client application does not support SSL connections. For an example, see [Section 6.3.10, “Connecting to MySQL Remotely from Windows with SSH”](#).

**B.9.3: Is SSL support be built into MySQL binaries, or must I recompile the binary myself to enable it?**

Most 5.6 binaries have SSL enabled for client-server connections that are secured, authenticated, or both. See [Section 6.3.9, “Using SSL for Secure Connections”](#).

**B.9.4: Does MySQL 5.6 have built-in authentication against LDAP directories?**

Not at this time.

**B.9.5: Does MySQL 5.6 include support for Roles Based Access Control (RBAC)?**

Not at this time.

## B.10. MySQL 5.6 FAQ: MySQL Cluster

In the following section, we answer questions that are frequently asked about MySQL Cluster and the [NDB](#) storage engine.

### Questions

- [B.10.1: \[3640\]](#) Which versions of the MySQL software support Cluster? Do I have to compile from source?

### Questions and Answers

**B.10.1: Which versions of the MySQL software support Cluster? Do I have to compile from source?**

MySQL Cluster is not supported in MySQL Server 5.6 releases. Instead, MySQL Cluster is released as a separate product, available as MySQL Cluster NDB 7.1 and MySQL Cluster NDB 7.2. You should use MySQL Cluster NDB 7.2 for new deployments, and plan to upgrade if you are using a previous version of MySQL Cluster. For an overview of improvements made in MySQL Cluster NDB 7.1, see [MySQL Cluster Development in MySQL Cluster NDB 7.1](#); for information about improvements made in MySQL Cluster NDB 7.2, see [MySQL Cluster Development in MySQL Cluster NDB 7.2](#).

For detailed information about deploying and using MySQL Cluster, see [MySQL Cluster NDB 6.1 - 7.1](#) and [MySQL Cluster NDB 7.2](#).

## B.11. MySQL 5.6 FAQ: MySQL Chinese, Japanese, and Korean Character Sets

This set of Frequently Asked Questions derives from the experience of MySQL's Support and Development groups in handling many inquiries about CJK (Chinese-Japanese-Korean) issues.

### Questions

- [B.11.1: \[3641\]](#) What CJK character sets are available in MySQL?
- [B.11.2: \[3642\]](#) I have inserted CJK characters into my table. Why does `SELECT` display them as “?” characters?
- [B.11.3: \[3643\]](#) What problems should I be aware of when working with the Big5 Chinese character set?
- [B.11.4: \[3644\]](#) Why do Japanese character set conversions fail?



- [B.11.5: \[3645\]](#) What should I do if I want to convert SJIS 81CA to cp932?
- [B.11.6: \[3645\]](#) How does MySQL represent the Yen (¥) sign?
- [B.11.7: \[3645\]](#) Does MySQL plan to make a separate character set where 5C is the Yen sign, as at least one other major DBMS does?
- [B.11.8: \[3645\]](#) Of what issues should I be aware when working with Korean character sets in MySQL?
- [B.11.9: \[3645\]](#) Why do I get `Incorrect string value` error messages?
- [B.11.10: \[3646\]](#) Why does my GUI front end or browser not display CJK characters correctly in my application using Access, PHP, or another API?
- [B.11.11: \[3647\]](#) I've upgraded to MySQL 5.6. How can I revert to behavior like that in MySQL 4.0 with regard to character sets?
- [B.11.12: \[3648\]](#) Why do some `LIKE [1190]` and `FULLTEXT` searches with CJK characters fail?
- [B.11.13: \[3649\]](#) How do I know whether character `x` is available in all character sets?
- [B.11.14: \[3650\]](#) Why do CJK strings sort incorrectly in Unicode? (I)
- [B.11.15: \[3651\]](#) Why do CJK strings sort incorrectly in Unicode? (II)
- [B.11.16: \[3651\]](#) Why are my supplementary characters rejected by MySQL?
- [B.11.17: \[3651\]](#) Shouldn't it be "CJKV"?
- [B.11.18: \[3652\]](#) Does MySQL allow CJK characters to be used in database and table names?
- [B.11.19: \[3652\]](#) Where can I find translations of the MySQL Manual into Chinese, Japanese, and Korean?
- [B.11.20: \[3652\]](#) Where can I get help with CJK and related issues in MySQL?

## Questions and Answers

### B.11.1: What CJK character sets are available in MySQL?

The list of CJK character sets may vary depending on your MySQL version. For example, the `ucjps` character set was not supported prior to MySQL 5.0.3. However, since the name of the applicable language appears in the `DESCRIPTION` column for every entry in the `INFORMATION_SCHEMA.CHARACTER_SETS` table, you can obtain a current list of all the non-Unicode CJK character sets using this query:

```
mysql> SELECT CHARACTER_SET_NAME, DESCRIPTION
-> FROM INFORMATION_SCHEMA.CHARACTER_SETS
-> WHERE DESCRIPTION LIKE '%Chinese%'
-> OR DESCRIPTION LIKE '%Japanese%'
-> OR DESCRIPTION LIKE '%Korean%'
-> ORDER BY CHARACTER_SET_NAME;
```

CHARACTER_SET_NAME	DESCRIPTION
big5	Big5 Traditional Chinese
cp932	SJIS for Windows Japanese
ucjps	UJIS for Windows Japanese
euckr	EUC-KR Korean
gb2312	GB2312 Simplified Chinese
gbk	GBK Simplified Chinese
sjis	Shift-JIS Japanese
ujis	EUC-JP Japanese

```
8 rows in set (0.01 sec)
```

(See [Section 19.1](#), “[The INFORMATION\\_SCHEMA CHARACTER\\_SETS Table](#)”, for more information.)

MySQL supports the two common variants of the *GB* (*Guojia Biaozhun*, or *National Standard*, or *Simplified Chinese*) character sets which are official in the People's Republic of China: [gb2312](#) and [gbk](#). Sometimes people try to insert [gbk](#) characters into [gb2312](#), and it works most of the time because [gbk](#) is a superset of [gb2312](#)—but eventually they try to insert a rarer Chinese character and it doesn't work. (See [Bug #16072](#) for an example).

Here, we try to clarify exactly what characters are legitimate in [gb2312](#) or [gbk](#), with reference to the official documents. Please check these references before reporting [gb2312](#) or [gbk](#) bugs.

- For a complete listing of the [gb2312](#) characters, ordered according to the [gb2312\\_chinese\\_ci](#) collation: [gb2312](#)
- MySQL's [gbk](#) is in reality “Microsoft code page 936”. This differs from the official [gbk](#) for characters [A1A4](#) (middle dot), [A1AA](#) (em dash), [A6E0–A6F5](#), and [A8BB–A8C0](#).
- For a listing of [gbk](#)/Unicode mappings, see <http://www.unicode.org/Public/MAPPINGS/VENDORS/MICSFT/WINDOWS/CP936.TXT>.
- For MySQL's listing of [gbk](#) characters, see [gbk](#).

### B.11.2: I have inserted CJK characters into my table. Why does `SELECT` display them as “?” characters?

This problem is usually due to a setting in MySQL that doesn't match the settings for the application program or the operating system. Here are some common steps for correcting these types of issues:

- *Be certain of what MySQL version you are using.*

Use the statement `SELECT VERSION();` to determine this.

- *Make sure that the database is actually using the desired character set.*

People often think that the client character set is always the same as either the server character set or the character set used for display purposes. However, both of these are false assumptions. You can make sure by checking the result of `SHOW CREATE TABLE tablename` or—better yet—by using this statement:

```
SELECT character_set_name, collation_name
FROM information_schema.columns
WHERE table_schema = your_database_name
AND table_name = your_table_name
AND column_name = your_column_name;
```

- *Determine the hexadecimal value of the character or characters that are not being displayed correctly.*

You can obtain this information for a column `column_name` in the table `table_name` using the following query:

```
SELECT HEX(column_name)
FROM table_name;
```

`3F` is the encoding for the `?` character; this means that `?` is the character actually stored in the column. This most often happens because of a problem converting a particular character from your client character set to the target character set.

- *Make sure that a round trip possible—that is, when you select `literal` (or `_introducer hexadecimal-value`), you obtain `literal` as a result.*

For example, the Japanese *Katakana* character *Pe* (ペ) exists in all CJK character sets, and has the code point value (hexadecimal coding) `0x30da`. To test a round trip for this character, use this query:

```
SELECT 'ㄹ' AS `ㄹ`;          /* or SELECT _ucs2 0x30da; */
```

If the result is not also ㄹ, then the round trip has failed.

For bug reports regarding such failures, we might ask you to follow up with `SELECT HEX('ㄹ');`. Then we can determine whether the client encoding is correct.

- *Make sure that the problem is not with the browser or other application, rather than with MySQL.*

Use the `mysql` client program (on Windows: `mysql.exe`) to accomplish this task. If `mysql` displays correctly but your application doesn't, then your problem is probably due to system settings.

To find out what your settings are, use the `SHOW VARIABLES` statement, whose output should resemble what is shown here:

```
mysql> SHOW VARIABLES LIKE 'char%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_client | utf8 |
| character_set_connection | utf8 |
| character_set_database | latin1 |
| character_set_filesystem | binary |
| character_set_results | utf8 |
| character_set_server | latin1 |
| character_set_system | utf8 |
| character_sets_dir | /usr/local/mysql/share/mysqlCharsets/ |
+-----+-----+
8 rows in set (0.03 sec)
```

These are typical character-set settings for an international-oriented client (notice the use of `utf8` Unicode) connected to a server in the West (`latin1` is a West Europe character set and a default for MySQL).

Although Unicode (usually the `utf8` variant on Unix, and the `ucs2` variant on Windows) is preferable to Latin, it is often not what your operating system utilities support best. Many Windows users find that a Microsoft character set, such as `cp932` for Japanese Windows, is suitable.

If you cannot control the server settings, and you have no idea what your underlying computer is, then try changing to a common character set for the country that you're in (`euckr` = Korea; `gb2312` or `gbk` = People's Republic of China; `big5` = Taiwan; `sjis`, `ujis`, `cp932`, or `eucljms` = Japan; `ucs2` or `utf8` = anywhere). Usually it is necessary to change only the client and connection and results settings. There is a simple statement which changes all three at once: `SET NAMES`. For example:

```
SET NAMES 'big5';
```

Once the setting is correct, you can make it permanent by editing `my.cnf` or `my.ini`. For example you might add lines looking like these:

```
[mysqld]
character-set-server=big5
[client]
default-character-set=big5
```

It is also possible that there are issues with the API configuration setting being used in your application; see *Why does my GUI front end or browser not display CJK characters correctly...?* for more information.

### B.11.3: What problems should I be aware of when working with the Big5 Chinese character set?

MySQL supports the Big5 character set which is common in Hong Kong and Taiwan (Republic of China). MySQL's `big5` is in reality Microsoft code page 950, which is very similar to the original `big5` character set. We changed to this character set starting with MySQL version 4.1.16 / 5.0.16 (as a result

of Bug #12476). For example, the following statements work in current versions of MySQL, but not in old versions:

```
mysql> CREATE TABLE big5 (BIG5 CHAR(1) CHARACTER SET BIG5);
Query OK, 0 rows affected (0.13 sec)

mysql> INSERT INTO big5 VALUES (0xf9dc);
Query OK, 1 row affected (0.00 sec)

mysql> SELECT * FROM big5;
+-----+
| big5 |
+-----+
| 嫺 |
+-----+
1 row in set (0.02 sec)
```

A feature request for adding [HKSCS](#) extensions has been filed. People who need this extension may find the suggested patch for Bug #13577 to be of interest.

#### B.11.4: Why do Japanese character set conversions fail?

MySQL supports the [sjis](#), [ujis](#), [cp932](#), and [eucjpms](#) character sets, as well as Unicode. A common need is to convert between character sets. For example, there might be a Unix server (typically with [sjis](#) or [ujis](#)) and a Windows client (typically with [cp932](#)).

In the following conversion table, the [ucs2](#) column represents the source, and the [sjis](#), [cp932](#), [ujis](#), and [eucjpms](#) columns represent the destinations—that is, the last 4 columns provide the hexadecimal result when we use [CONVERT\(ucs2\)](#) [1245] or we assign a [ucs2](#) column containing the value to an [sjis](#), [cp932](#), [ujis](#), or [eucjpms](#) column.

Character Name	ucs2	sjis	cp932	ujis	eucjpms
BROKEN BAR	00A6	3F	3F	8FA2C3	3F
FULLWIDTH BROKEN BAR	FFE4	3F	FA55	3F	8FA2
YEN SIGN	00A5	3F	3F	20	3F
FULLWIDTH YEN SIGN	FFE5	818F	818F	A1EF	3F
TILDE	007E	7E	7E	7E	7E
OVERLINE	203E	3F	3F	20	3F
HORIZONTAL BAR	2015	815C	815C	A1BD	A1BD
EM DASH	2014	3F	3F	3F	3F
REVERSE SOLIDUS	005C	815F	5C	5C	5C
FULLWIDTH ""	FF3C	3F	815F	3F	A1C0
WAVE DASH	301C	8160	3F	A1C1	3F
FULLWIDTH TILDE	FF5E	3F	8160	3F	A1C1
DOUBLE VERTICAL LINE	2016	8161	3F	A1C2	3F
PARALLEL TO	2225	3F	8161	3F	A1C2
MINUS SIGN	2212	817C	3F	A1DD	3F
FULLWIDTH HYPHEN-MINUS	FF0D	3F	817C	3F	A1DD
CENT SIGN	00A2	8191	3F	A1F1	3F
FULLWIDTH CENT SIGN	FFE0	3F	8191	3F	A1F1
POUND SIGN	00A3	8192	3F	A1F2	3F
FULLWIDTH POUND SIGN	FFE1	3F	8192	3F	A1F2
NOT SIGN	00AC	81CA	3F	A2CC	3F

Character Name	ucs2	sjis	cp932	ujis	eucjpms
FULLWIDTH NOT SIGN	FFE2	3F	81CA	3F	A2CC

Now consider the following portion of the table.

	ucs2	sjis	cp932
NOT SIGN	00AC	81CA	3F
FULLWIDTH NOT SIGN	FFE2	3F	81CA

This means that MySQL converts the `NOT SIGN` (Unicode `U+00AC`) to `sjis` code point `0x81CA` and to `cp932` code point `3F`. (`3F` is the question mark (“?”)—this is what is always used when the conversion cannot be performed.

#### B.11.5: What should I do if I want to convert SJIS `81CA` to `cp932`?

Our answer is: “?”. There are serious complaints about this: many people would prefer a “loose” conversion, so that `81CA` (`NOT SIGN`) in `sjis` becomes `81CA` (`FULLWIDTH NOT SIGN`) in `cp932`. We are considering a change to this behavior.

#### B.11.6: How does MySQL represent the Yen (¥) sign?

A problem arises because some versions of Japanese character sets (both `sjis` and `euc`) treat `5C` as a *reverse solidus* (`\`—also known as a backslash), and others treat it as a yen sign (¥).

MySQL follows only one version of the JIS (Japanese Industrial Standards) standard description. In MySQL, `5C` is always the *reverse solidus* (`\`).

#### B.11.7: Does MySQL plan to make a separate character set where `5C` is the Yen sign, as at least one other major DBMS does?

This is one possible solution to the Yen sign issue; however, this will not happen in MySQL 5.1 or 6.0.

#### B.11.8: Of what issues should I be aware when working with Korean character sets in MySQL?

In theory, while there have been several versions of the `euckr` (*Extended Unix Code Korea*) character set, only one problem has been noted.

We use the “ASCII” variant of EUC-KR, in which the code point `0x5c` is REVERSE SOLIDUS, that is `\`, instead of the “KS-Roman” variant of EUC-KR, in which the code point `0x5c` is `WON SIGN`(₩). This means that you cannot convert Unicode `U+20A9` to `euckr`:

```
mysql> SELECT
  ->   CONVERT('₩' USING euckr) AS euckr,
  ->   HEX(CONVERT('₩' USING euckr)) AS hexeuckr;
+-----+-----+
| euckr | hexeuckr |
+-----+-----+
| ?    | 3F      |
+-----+-----+
1 row in set (0.00 sec)
```

MySQL's graphic Korean chart is here: [euckr](#).

#### B.11.9: Why do I get **Incorrect string value** error messages?

For illustration, we'll create a table with one Unicode (`ucs2`) column and one Chinese (`gb2312`) column.

```
mysql> CREATE TABLE ch
  ->   (ucs2 CHAR(3) CHARACTER SET ucs2,
  ->   gb2312 CHAR(3) CHARACTER SET gb2312);
```

```
Query OK, 0 rows affected (0.05 sec)
```

We'll try to place the rare character ㄸ in both columns.

```
mysql> INSERT INTO ch VALUES ('AㄸB','AㄸB');
Query OK, 1 row affected, 1 warning (0.00 sec)
```

Ah, there is a warning. Use the following statement to see what it is:

```
mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1366
Message: Incorrect string value: '\xE6\xB1\x8CB' for column 'gb2312' at row 1
1 row in set (0.00 sec)
```

So it is a warning about the `gb2312` column only.

```
mysql> SELECT ucs2,HEX(ucs2),gb2312,HEX(gb2312) FROM ch;
+-----+-----+-----+-----+
| ucs2 | HEX(ucs2) | gb2312 | HEX(gb2312) |
+-----+-----+-----+-----+
| AㄸB | 00416C4C0042 | A?B | 413F42 |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

Several things need explanation here:

1. The fact that it is a “warning” rather than an “error” is characteristic of MySQL. We like to try to do what we can, to get the best fit, rather than give up.
2. The ㄸ character is not in the `gb2312` character set. We described that problem earlier.
3. If you are using an old version of MySQL, you will probably see a different message.
4. With `sql_mode=TRADITIONAL` [545], there would be an error message, rather than a warning.

#### B.11.10: Why does my GUI front end or browser not display CJK characters correctly in my application using Access, PHP, or another API?

Obtain a direct connection to the server using the `mysql` client (Windows: `mysql.exe`), and try the same query there. If `mysql` responds correctly, then the trouble may be that your application interface requires initialization. Use `mysql` to tell you what character set or sets it uses with the statement `SHOW VARIABLES LIKE 'char%'`. If you are using Access, then you are most likely connecting with Connector/ODBC. In this case, you should check [Section 21.1.4, “Configuring Connector/ODBC”](#). If, for instance, you use `big5`, you would enter `SET NAMES 'big5'`. (Note that no `;` is required in this case). If you are using ASP, you might need to add `SET NAMES` in the code. Here is an example that has worked in the past:

```
<%
Session.CodePage=0
Dim strConnection
Dim Conn
strConnection="driver={MySQL ODBC 3.51 Driver};server=server;uid=username;" \
& "pwd=password;database=database;stmt=SET NAMES 'big5';"
Set Conn = Server.CreateObject("ADODB.Connection")
Conn.Open strConnection
%>
```

In much the same way, if you are using any character set other than `latin1` with Connector/Net, then you must specify the character set in the connection string. See [Section 21.2.5.1, “Connecting to MySQL Using Connector/Net”](#), for more information.

If you are using PHP, try this:

```
<?php
```

```

$link = mysql_connect($host, $usr, $pwd);

mysql_select_db($db);

if( mysql_error() ) { print "Database ERROR: " . mysql_error(); }
mysql_query("SET NAMES 'utf8'", $link);
?>

```

In this case, we used `SET NAMES` to change `character_set_client` [467] and `character_set_connection` [467] and `character_set_results` [468].

We encourage the use of the newer `mysqli` extension, rather than `mysql`. Using `mysqli`, the previous example could be rewritten as shown here:

```

<?php
$link = new mysqli($host, $usr, $pwd, $db);

if( mysqli_connect_errno() )
{
    printf("Connect failed: %s\n", mysqli_connect_error());
    exit();
}

$link->query("SET NAMES 'utf8'");
?>

```

Another issue often encountered in PHP applications has to do with assumptions made by the browser. Sometimes adding or changing a `<meta>` tag suffices to correct the problem: for example, to insure that the user agent interprets page content as UTF-8, you should include `<meta http-equiv="Content-Type" content="text/html; charset=utf-8">` in the `<head>` of the HTML page.

If you are using Connector/J, see [Section 21.3.5.4, “Using Character Sets and Unicode”](#).

#### **B.11.11: I've upgraded to MySQL 5.6. How can I revert to behavior like that in MySQL 4.0 with regard to character sets?**

In MySQL Version 4.0, there was a single “global” character set for both server and client, and the decision as to which character to use was made by the server administrator. This changed starting with MySQL Version 4.1. What happens now is a “handshake”, as described in [Section 10.1.4, “Connection Character Sets and Collations”](#):

When a client connects, it sends to the server the name of the character set that it wants to use. The server uses the name to set the `character_set_client` [467], `character_set_results` [468], and `character_set_connection` [467] system variables. In effect, the server performs a `SET NAMES` operation using the character set name.

The effect of this is that you cannot control the client character set by starting `mysqld` with `--character-set-server=utf8` [417]. However, some of our Asian customers have said that they prefer the MySQL 4.0 behavior. To make it possible to retain this behavior, we added a `mysqld` switch, `--character-set-client-handshake` [416], which can be turned off with `--skip-character-set-client-handshake` [416]. If you start `mysqld` with `--skip-character-set-client-handshake` [416], then, when a client connects, it sends to the server the name of the character set that it wants to use—however, *the server ignores this request from the client*.

By way of example, suppose that your favorite server character set is `latin1` (unlikely in a CJK area, but this is the default value). Suppose further that the client uses `utf8` because this is what the client's operating system supports. Now, start the server with `latin1` as its default character set:

```
mysqld --character-set-server=latin1
```

And then start the client with the default character set `utf8`:



```
mysql --default-character-set=utf8
```

The current settings can be seen by viewing the output of `SHOW VARIABLES`:

```
mysql> SHOW VARIABLES LIKE 'char%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_client | utf8 |
| character_set_connection | utf8 |
| character_set_database | latin1 |
| character_set_filesystem | binary |
| character_set_results | utf8 |
| character_set_server | latin1 |
| character_set_system | utf8 |
| character_sets_dir | /usr/local/mysql/share/mysqlCharsets/ |
+-----+-----+
8 rows in set (0.01 sec)
```

Now stop the client, and then stop the server using `mysqldadmin`. Then start the server again, but this time tell it to skip the handshake like so:

```
mysqld --character-set-server=utf8 --skip-character-set-client-handshake
```

Start the client with `utf8` once again as the default character set, then display the current settings:

```
mysql> SHOW VARIABLES LIKE 'char%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_client | latin1 |
| character_set_connection | latin1 |
| character_set_database | latin1 |
| character_set_filesystem | binary |
| character_set_results | latin1 |
| character_set_server | latin1 |
| character_set_system | utf8 |
| character_sets_dir | /usr/local/mysql/share/mysqlCharsets/ |
+-----+-----+
8 rows in set (0.01 sec)
```

As you can see by comparing the differing results from `SHOW VARIABLES`, the server ignores the client's initial settings if the `--skip-character-set-client-handshake` [416] is used.

#### B.11.12: Why do some `LIKE` [1190] and `FULLTEXT` searches with CJK characters fail?

There is a very simple problem with `LIKE` [1190] searches on `BINARY` and `BLOB` columns: we need to know the end of a character. With multi-byte character sets, different characters might have different octet lengths. For example, in `utf8`, `A` requires one byte but `^` requires three bytes, as shown here:

```
+-----+-----+
| OCTET_LENGTH(utf8 'A') | OCTET_LENGTH(utf8 '^') |
+-----+-----+
| 1 | 3 |
+-----+-----+
1 row in set (0.00 sec)
```

If we don't know where the first character ends, then we don't know where the second character begins, in which case even very simple searches such as `LIKE '_A%'` [1190] fail. The solution is to use a regular CJK character set in the first place, or to convert to a CJK character set before comparing.

This is one reason why MySQL cannot allow encodings of nonexistent characters. If it is not strict about rejecting bad input, then it has no way of knowing where characters end.

For `FULLTEXT` searches, we need to know where words begin and end. With Western languages, this is rarely a problem because most (if not all) of these use an easy-to-identify word boundary—



the space character. However, this is not usually the case with Asian writing. We could use arbitrary halfway measures, like assuming that all Han characters represent words, or (for Japanese) depending on changes from Katakana to Hiragana due to grammatical endings. However, the only sure solution requires a comprehensive word list, which means that we would have to include a dictionary in the server for each Asian language supported. This is simply not feasible.

### B.11.13: How do I know whether character *x* is available in all character sets?

The majority of simplified Chinese and basic nonhalfwidth Japanese *Kana* characters appear in all CJK character sets. This stored procedure accepts a [UCS-2](#) Unicode character, converts it to all other character sets, and displays the results in hexadecimal.

```
DELIMITER //

CREATE PROCEDURE p_convert(ucs2_char CHAR(1) CHARACTER SET ucs2)
BEGIN

CREATE TABLE tj
    (ucs2 CHAR(1) character set ucs2,
    utf8 CHAR(1) character set utf8,
    big5 CHAR(1) character set big5,
    cp932 CHAR(1) character set cp932,
    eucjpms CHAR(1) character set eucjpms,
    euckr CHAR(1) character set euckr,
    gb2312 CHAR(1) character set gb2312,
    gbk CHAR(1) character set gbk,
    sjis CHAR(1) character set sjis,
    ujis CHAR(1) character set ujis);

INSERT INTO tj (ucs2) VALUES (ucs2_char);

UPDATE tj SET utf8=ucs2,
    big5=ucs2,
    cp932=ucs2,
    eucjpms=ucs2,
    euckr=ucs2,
    gb2312=ucs2,
    gbk=ucs2,
    sjis=ucs2,
    ujis=ucs2;

/* If there is a conversion problem, UPDATE will produce a warning. */

SELECT hex(ucs2) AS ucs2,
    hex(utf8) AS utf8,
    hex(big5) AS big5,
    hex(cp932) AS cp932,
    hex(eucjpms) AS eucjpms,
    hex(euckr) AS euckr,
    hex(gb2312) AS gb2312,
    hex(gbk) AS gbk,
    hex(sjis) AS sjis,
    hex(ujis) AS ujis
FROM tj;

DROP TABLE tj;

END//
```

The input can be any single `ucs2` character, or it can be the code point value (hexadecimal representation) of that character. For example, from Unicode's list of `ucs2` encodings and names (<http://www.unicode.org/Public/UNIDATA/UnicodeData.txt>), we know that the *Katakana* character *Pe* appears in all CJK character sets, and that its code point value is `0x30da`. If we use this value as the argument to `p_convert()`, the result is as shown here:

```
mysql> CALL p_convert(0x30da)//
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ucs2 | utf8   | big5   | cp932 | eucjpms | euckr | gb2312 | gbk   | sjis   | ujis   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0x30 | 0x30da | 0x30da | 0x30da | 0x30da  | 0x30da | 0x30da  | 0x30da | 0x30da | 0x30da |
```

```
| 30DA | E3839A | C772 | 8379 | A5DA | ABDA | A5DA | A5DA | 8379 | A5DA |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.04 sec)
```

Since none of the column values is 3F—that is, the question mark character (?)—we know that every conversion worked.

#### B.11.14: Why do CJK strings sort incorrectly in Unicode? (I)

Sometimes people observe that the result of a `utf8_unicode_ci` or `ucs2_unicode_ci` search, or of an `ORDER BY` sort is not what they think a native would expect. Although we never rule out the possibility that there is a bug, we have found in the past that many people do not read correctly the standard table of weights for the Unicode Collation Algorithm. MySQL uses the table found at <http://www.unicode.org/Public/UCA/4.0.0/allkeys-4.0.0.txt>. This is not the first table you will find by navigating from the [unicode.org](http://www.unicode.org) home page, because MySQL uses the older 4.0.0 “allkeys” table, rather than the more recent 4.1.0 table. (The newer ‘520’ collations in MySQL 5.6 use the 5.2 “allkeys” table.) This is because we are very wary about changing ordering which affects indexes, lest we bring about situations such as that reported in Bug #16526, illustrated as follows:

```
mysql> CREATE TABLE tj (s1 CHAR(1) CHARACTER SET utf8 COLLATE utf8_unicode_ci);
Query OK, 0 rows affected (0.05 sec)

mysql> INSERT INTO tj VALUES ('が'),('か');
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM tj WHERE s1 = 'か';
+-----+
| s1    |
+-----+
| が    |
| か    |
+-----+
2 rows in set (0.00 sec)
```

The character in the first result row is not the one that we searched for. Why did MySQL retrieve it? First we look for the Unicode code point value, which is possible by reading the hexadecimal number for the `ucs2` version of the characters:

```
mysql> SELECT s1, HEX(CONVERT(s1 USING ucs2)) FROM tj;
+-----+-----+
| s1    | HEX(CONVERT(s1 USING ucs2)) |
+-----+-----+
| が    | 304C                          |
| か    | 304B                          |
+-----+-----+
2 rows in set (0.03 sec)
```

Now we search for 304B and 304C in the 4.0.0 `allkeys` table, and find these lines:

```
304B ; [.1E57.0020.000E.304B] # HIRAGANA LETTER KA
304C ; [.1E57.0020.000E.304B][.0000.0140.0002.3099] # HIRAGANA LETTER GA; QQCM
```

The official Unicode names (following the “#” mark) tell us the Japanese syllabary (Hiragana), the informal classification (letter, digit, or punctuation mark), and the Western identifier (`KA` or `GA`, which happen to be voiced and unvoiced components of the same letter pair). More importantly, the *primary weight* (the first hexadecimal number inside the square brackets) is 1E57 on both lines. For comparisons in both searching and sorting, MySQL pays attention to the primary weight only, ignoring all the other numbers. This means that we are sorting が and か correctly according to the Unicode specification. If we wanted to distinguish them, we’d have to use a non-UCA (Unicode Collation Algorithm) collation (`utf8_bin` or `utf8_general_ci`), or to compare the `HEX()` [1180] values, or use `ORDER BY CONVERT(s1 USING sjis)`. Being correct “according to Unicode” isn’t enough, of course: the person who submitted the bug was equally correct. We plan to add another collation for Japanese according to the JIS X 4061 standard, in which voiced/unvoiced letter pairs like `KA/GA` are distinguishable for ordering purposes.

**B.11.15: Why do CJK strings sort incorrectly in Unicode? (II)**

If you are using Unicode (`ucs2` or `utf8`), and you know what the Unicode sort order is (see [Section B.11, “MySQL 5.6 FAQ: MySQL Chinese, Japanese, and Korean Character Sets”](#)), but MySQL still seems to sort your table incorrectly, then you should first verify the table character set:

```
mysql> SHOW CREATE TABLE t\G
***** 1. row *****
Table: t
Create Table: CREATE TABLE `t` (
`s1` char(1) CHARACTER SET ucs2 DEFAULT NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
1 row in set (0.00 sec)
```

Since the character set appears to be correct, let's see what information the `INFORMATION_SCHEMA.COLUMNS` table can provide about this column:

```
mysql> SELECT COLUMN_NAME, CHARACTER_SET_NAME, COLLATION_NAME
-> FROM INFORMATION_SCHEMA.COLUMNS
-> WHERE COLUMN_NAME = 's1'
-> AND TABLE_NAME = 't';
+-----+-----+-----+
| COLUMN_NAME | CHARACTER_SET_NAME | COLLATION_NAME |
+-----+-----+-----+
| s1          | ucs2               | ucs2_general_ci |
+-----+-----+-----+
1 row in set (0.01 sec)
```

(See [Section 19.4, “The INFORMATION\\_SCHEMA COLUMNS Table”](#), for more information.)

You can see that the collation is `ucs2_general_ci` instead of `ucs2_unicode_ci`. The reason why this is so can be found using `SHOW CHARSET`, as shown here:

```
mysql> SHOW CHARSET LIKE 'ucs2%';
+-----+-----+-----+-----+
| Charset | Description      | Default collation | Maxlen |
+-----+-----+-----+-----+
| ucs2    | UCS-2 Unicode   | ucs2_general_ci  | 2      |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

For `ucs2` and `utf8`, the default collation is “general”. To specify a Unicode collation, use `COLLATE ucs2_unicode_ci`.

**B.11.16: Why are my supplementary characters rejected by MySQL?**

Before MySQL 5.5.3, MySQL does not support supplementary characters—that is, characters which need more than 3 bytes—for `UTF-8`. We support only what Unicode calls the *Basic Multilingual Plane / Plane 0*. Only a few very rare Han characters are supplementary; support for them is uncommon. This has led to reports such as that found in Bug #12600, which we rejected as “not a bug”. With `utf8`, we must truncate an input string when we encounter bytes that we don't understand. Otherwise, we wouldn't know how long the bad multi-byte character is.

One possible workaround is to use `ucs2` instead of `utf8`, in which case the “bad” characters are changed to question marks; however, no truncation takes place. You can also change the data type to `BLOB` or `BINARY`, which perform no validity checking.

As of MySQL 5.5.3, Unicode support is extended to include supplementary characters by means of additional Unicode character sets: `utf16`, `utf32`, and 4-byte `utf8mb4`. These character sets support supplementary Unicode characters outside the Basic Multilingual Plane (BMP).

**B.11.17: Shouldn't it be “CJKV”?**

No. The term “CJKV” (*Chinese Japanese Korean Vietnamese*) refers to Vietnamese character sets which contain Han (originally Chinese) characters. MySQL has no plan to support the old Vietnamese

script using Han characters. MySQL does of course support the modern Vietnamese script with Western characters.

As of MySQL 5.6, there are Vietnamese collations for Unicode character sets, as described in [Section 10.1.14.1, “Unicode Character Sets”](#).

#### **B.11.18: Does MySQL allow CJK characters to be used in database and table names?**

This issue is fixed in MySQL 5.1, by automatically rewriting the names of the corresponding directories and files.

For example, if you create a database named 椿 on a server whose operating system does not support CJK in directory names, MySQL creates a directory named @0w@00a5@00ae. which is just a fancy way of encoding E6A5AE—that is, the Unicode hexadecimal representation for the 椿 character. However, if you run a `SHOW DATABASES` statement, you can see that the database is listed as 椿.

#### **B.11.19: Where can I find translations of the MySQL Manual into Chinese, Japanese, and Korean?**

A Simplified Chinese version of the Manual, current for MySQL 5.1.12, can be found at <http://dev.mysql.com/doc/>. The Japanese translation of the MySQL 4.1 manual can be downloaded from <http://dev.mysql.com/doc/>.

#### **B.11.20: Where can I get help with CJK and related issues in MySQL?**

The following resources are available:

- A listing of MySQL user groups can be found at <http://dev.mysql.com/user-groups/>.
- View feature requests relating to character set issues at <http://tinyurl.com/y6xcuf>.
- Visit the MySQL [Character Sets, Collation, Unicode Forum](#). We are also in the process of adding foreign-language forums at <http://forums.mysql.com/>.

## **B.12. MySQL 5.6 FAQ: Connectors & APIs**

For common questions, issues, and answers relating to the MySQL Connectors and other APIs, see the following areas of the Manual:

- [Section 21.8.14, “Common Questions and Problems When Using the C API”](#)
- [Section 21.9.12, “Common Problems with MySQL and PHP”](#)
- [Section 21.1.7, “Connector/ODBC Notes and Tips”](#)
- [Section 21.2.5, “Connector/Net Programming”](#)
- [Section 21.3, “MySQL Connector/J”](#)

## **B.13. MySQL 5.6 FAQ: Replication**

In the following section, we provide answers to questions that are most frequently asked about MySQL Replication.

### **Questions**

- **B.13.1:** [\[3653\]](#) Must the slave be connected to the master all the time?
- **B.13.2:** [\[3653\]](#) Must I enable networking on my master and slave to enable replication?
- **B.13.3:** [\[3653\]](#) How do I know how late a slave is compared to the master? In other words, how do I know the date of the last statement replicated by the slave?

- [B.13.4: \[3654\]](#) How do I force the master to block updates until the slave catches up?
- [B.13.5: \[3654\]](#) What issues should I be aware of when setting up two-way replication?
- [B.13.6: \[3654\]](#) How can I use replication to improve performance of my system?
- [B.13.7: \[3654\]](#) What should I do to prepare client code in my own applications to use performance-enhancing replication?
- [B.13.8: \[3654\]](#) When and how much can MySQL replication improve the performance of my system?
- [B.13.9: \[3655\]](#) How can I use replication to provide redundancy or high availability?
- [B.13.10: \[3656\]](#) How do I tell whether a master server is using statement-based or row-based binary logging format?
- [B.13.11: \[3656\]](#) How do I tell a slave to use row-based replication?
- [B.13.12: \[3656\]](#) How do I prevent `GRANT` and `REVOKE` statements from replicating to slave machines?
- [B.13.13: \[3656\]](#) Does replication work on mixed operating systems (for example, the master runs on Linux while slaves run on Mac OS X and Windows)?
- [B.13.14: \[3656\]](#) Does replication work on mixed hardware architectures (for example, the master runs on a 64-bit machine while slaves run on 32-bit machines)?

## Questions and Answers

### **B.13.1: Must the slave be connected to the master all the time?**

No, it does not. The slave can go down or stay disconnected for hours or even days, and then reconnect and catch up on updates. For example, you can set up a master/slave relationship over a dial-up link where the link is up only sporadically and for short periods of time. The implication of this is that, at any given time, the slave is not guaranteed to be in synchrony with the master unless you take some special measures.

To ensure that catchup can occur for a slave that has been disconnected, you must not remove binary log files from the master that contain information that has not yet been replicated to the slaves. Asynchronous replication can work only if the slave is able to continue reading the binary log from the point where it last read events.

### **B.13.2: Must I enable networking on my master and slave to enable replication?**

Yes, networking must be enabled on the master and slave. If networking is not enabled, the slave cannot connect to the master and transfer the binary log. Check that the `skip-networking` [\[440\]](#) option has not been enabled in the configuration file for either server.

### **B.13.3: How do I know how late a slave is compared to the master? In other words, how do I know the date of the last statement replicated by the slave?**

Check the `Seconds_Behind_Master` column in the output from `SHOW SLAVE STATUS`. See [Section 16.1.5.1, "Checking Replication Status"](#).

When the slave SQL thread executes an event read from the master, it modifies its own time to the event timestamp. (This is why `TIMESTAMP` is well replicated.) In the `Time` column in the output of `SHOW PROCESSLIST`, the number of seconds displayed for the slave SQL thread is the number of seconds between the timestamp of the last replicated event and the real time of the slave machine. You can use this to determine the date of the last replicated event. Note that if your slave has been disconnected from the master for one hour, and then reconnects, you may immediately see large `Time` values such

as 3600 for the slave SQL thread in `SHOW PROCESSLIST`. This is because the slave is executing statements that are one hour old. See [Section 16.2.1, “Replication Implementation Details”](#).

#### B.13.4: How do I force the master to block updates until the slave catches up?

Use the following procedure:

1. On the master, execute these statements:

```
mysql> FLUSH TABLES WITH READ LOCK;  
mysql> SHOW MASTER STATUS;
```

Record the replication coordinates (the current binary log file name and position) from the output of the `SHOW` statement.

2. On the slave, issue the following statement, where the arguments to the `MASTER_POS_WAIT()` [1279] function are the replication coordinate values obtained in the previous step:

```
mysql> SELECT MASTER_POS_WAIT('log_name', log_pos);
```

The `SELECT` statement blocks until the slave reaches the specified log file and position. At that point, the slave is in synchrony with the master and the statement returns.

3. On the master, issue the following statement to enable the master to begin processing updates again:

```
mysql> UNLOCK TABLES;
```

#### B.13.5: What issues should I be aware of when setting up two-way replication?

MySQL replication currently does not support any locking protocol between master and slave to guarantee the atomicity of a distributed (cross-server) update. In other words, it is possible for client A to make an update to co-master 1, and in the meantime, before it propagates to co-master 2, client B could make an update to co-master 2 that makes the update of client A work differently than it did on co-master 1. Thus, when the update of client A makes it to co-master 2, it produces tables that are different from what you have on co-master 1, even after all the updates from co-master 2 have also propagated. This means that you should not chain two servers together in a two-way replication relationship unless you are sure that your updates can safely happen in any order, or unless you take care of mis-ordered updates somehow in the client code.

You should also realize that two-way replication actually does not improve performance very much (if at all) as far as updates are concerned. Each server must do the same number of updates, just as you would have a single server do. The only difference is that there is a little less lock contention because the updates originating on another server are serialized in one slave thread. Even this benefit might be offset by network delays.

#### B.13.6: How can I use replication to improve performance of my system?

Set up one server as the master and direct all writes to it. Then configure as many slaves as you have the budget and rackspace for, and distribute the reads among the master and the slaves. You can also start the slaves with the `--skip-innodb` [1698], `--low-priority-updates` [430], and `--delay-key-write=ALL` [421] options to get speed improvements on the slave end. In this case, the slave uses nontransactional `MyISAM` tables instead of `InnoDB` tables to get more speed by eliminating transactional overhead.

#### B.13.7: What should I do to prepare client code in my own applications to use performance-enhancing replication?

See the guide to using replication as a scale-out solution, [Section 16.3.3, “Using Replication for Scale-Out”](#).

#### B.13.8: When and how much can MySQL replication improve the performance of my system?

MySQL replication is most beneficial for a system that processes frequent reads and infrequent writes. In theory, by using a single-master/multiple-slave setup, you can scale the system by adding more slaves until you either run out of network bandwidth, or your update load grows to the point that the master cannot handle it.

To determine how many slaves you can use before the added benefits begin to level out, and how much you can improve performance of your site, you must know your query patterns, and determine empirically by benchmarking the relationship between the throughput for reads and writes on a typical master and a typical slave. The example here shows a rather simplified calculation of what you can get with replication for a hypothetical system. Let `reads` and `writes` denote the number of reads and writes per second, respectively.

Let's say that system load consists of 10% writes and 90% reads, and we have determined by benchmarking that `reads` is  $1200 - 2 * \text{writes}$ . In other words, the system can do 1,200 reads per second with no writes, the average write is twice as slow as the average read, and the relationship is linear. Suppose that the master and each slave have the same capacity, and that we have one master and  $N$  slaves. Then we have for each server (master or slave):

$$\text{reads} = 1200 - 2 * \text{writes}$$

$$\text{reads} = 9 * \text{writes} / (N + 1) \text{ (reads are split, but writes replicated to all slaves)}$$

$$9 * \text{writes} / (N + 1) + 2 * \text{writes} = 1200$$

$$\text{writes} = 1200 / (2 + 9/(N + 1))$$

The last equation indicates the maximum number of writes for  $N$  slaves, given a maximum possible read rate of 1,200 per second and a ratio of nine reads per write.

This analysis yields the following conclusions:

- If  $N = 0$  (which means we have no replication), our system can handle about  $1200/11 = 109$  writes per second.
- If  $N = 1$ , we get up to 184 writes per second.
- If  $N = 8$ , we get up to 400 writes per second.
- If  $N = 17$ , we get up to 480 writes per second.
- Eventually, as  $N$  approaches infinity (and our budget negative infinity), we can get very close to 600 writes per second, increasing system throughput about 5.5 times. However, with only eight servers, we increase it nearly four times.

Note that these computations assume infinite network bandwidth and neglect several other factors that could be significant on your system. In many cases, you may not be able to perform a computation similar to the one just shown that accurately predicts what will happen on your system if you add  $N$  replication slaves. However, answering the following questions should help you decide whether and by how much replication will improve the performance of your system:

- What is the read/write ratio on your system?
- How much more write load can one server handle if you reduce the reads?
- For how many slaves do you have bandwidth available on your network?

#### **B.13.9: How can I use replication to provide redundancy or high availability?**

How you implement redundancy is entirely dependent on your application and circumstances. High-availability solutions (with automatic failover) require active monitoring and either custom scripts or third party tools to provide the failover support from the original MySQL server to the slave.



To handle the process manually, you should be able to switch from a failed master to a pre-configured slave by altering your application to talk to the new server or by adjusting the DNS for the MySQL server from the failed server to the new server.

For more information and some example solutions, see [Section 16.3.6, “Switching Masters During Failover”](#).

#### **B.13.10: How do I tell whether a master server is using statement-based or row-based binary logging format?**

Check the value of the `binlog_format` [2010] system variable:

```
mysql> SHOW VARIABLES LIKE 'binlog_format';
```

The value shown will be one of `STATEMENT`, `ROW`, or `MIXED`. For `MIXED` mode, row-based logging is preferred but replication switches automatically to statement-based logging under certain conditions; for information about when this may occur, see [Section 5.2.4.3, “Mixed Binary Logging Format”](#).

#### **B.13.11: How do I tell a slave to use row-based replication?**

Slaves automatically know which format to use.

#### **B.13.12: How do I prevent `GRANT` and `REVOKE` statements from replicating to slave machines?**

Start the server with the `--replicate-wild-ignore-table=mysql.%` [1980] option to ignore replication for tables in the `mysql` database.

#### **B.13.13: Does replication work on mixed operating systems (for example, the master runs on Linux while slaves run on Mac OS X and Windows)?**

Yes.

#### **B.13.14: Does replication work on mixed hardware architectures (for example, the master runs on a 64-bit machine while slaves run on 32-bit machines)?**

Yes.

## **B.14. MySQL 5.6 FAQ: MySQL Enterprise Scalability Thread Pool**

### **Questions**

- [B.14.1: \[3656\]](#) What is the Thread Pool and what problem does it solve?
- [B.14.2: \[3657\]](#) How does the Thread Pool limit and manage concurrent sessions and transactions for optimal performance and throughput?
- [B.14.3: \[3657\]](#) How is the Thread Pool different from the client side Connection Pool?
- [B.14.4: \[3657\]](#) When should I use the Thread Pool?
- [B.14.5: \[3657\]](#) Are there recommended Thread Pool configurations?

### **Questions and Answers**

#### **B.14.1: What is the Thread Pool and what problem does it solve?**

The MySQL Thread Pool is a MySQL server plugin that extends the default connection-handling capabilities of the MySQL server to limit the number of concurrently executing statements/queries and transactions to ensure that each has sufficient CPU and memory resources to fulfill its task. Commercial distributions of MySQL 5.5 and 5.6 include the Thread Pool plugin.

The default thread-handling model in MySQL Server executes statements using one thread per client connection. As more clients connect to the server and execute statements, overall performance



degrades. The Thread Pool plugin provides an alternative thread-handling model designed to reduce overhead and improve performance. The Thread Pool plugin increases server performance by efficiently managing statement execution threads for large numbers of client connections, especially on modern multi-CPU/Core systems.

For more information, see [Section 8.11.6, “The Thread Pool Plugin”](#).

**B.14.2: How does the Thread Pool limit and manage concurrent sessions and transactions for optimal performance and throughput?**

The Thread Pool uses a “divide and conquer” approach to limiting and balancing concurrency. Unlike the default connection handling of the MySQL Server, the Thread Pool separates connections and threads, so there is no fixed relationship between connections and the threads that execute statements received from those connections. The Thread Pool then manages client connections within configurable thread groups, where they are prioritized and queued based on the nature of the work they were submitted to accomplish.

For more information, see [Section 8.11.6.2, “Thread Pool Operation”](#).

**B.14.3: How is the Thread Pool different from the client side Connection Pool?**

The MySQL Connection Pool operates on the client side to ensure that a MySQL client does not constantly connect to and disconnect from the MySQL server. It is designed to cache idle connections in the MySQL client for use by other users as they are needed. This minimizes the overhead and expense of establishing and tearing down connections as queries are submitted to the MySQL server. The MySQL Connection Pool has no visibility as to the query handling capabilities or load of the backend MySQL server. By contrast, the Thread Pool operates on the MySQL server side and is designed to manage the execution of inbound concurrent connections and queries as they are received from the client connections accessing the backend MySQL database. Because of the separation of duties, the MySQL Connection Pool and Thread Pool are orthogonal and can be used independent of each other.

MySQL Connection Pooling via the MySQL Connectors is covered in [Chapter 21, Connectors and APIs](#).

**B.14.4: When should I use the Thread Pool?**

There are a few rules of thumb to consider for optimal Thread Pool use cases:

The MySQL `Threads_running` [601] variable keeps track of the number of concurrent statements currently executing in the MySQL Server. If this variable consistently exceeds a region where the server won't operate optimally (usually going beyond 40 for InnoDB workloads), the Thread Pool will be beneficial, especially in extreme parallel overload situations.

If you are using the `innodb_thread_concurrency` [1751] to limit the number of concurrently executing statements, you will find the Thread Pool solves the same problem, only better, by assigning connections to thread groups, then queuing executions based on transactional content, user defined designations, and so forth.

Lastly, if your workload comprises mainly short queries, the Thread Pool will be beneficial.

To learn more, see [Section 8.11.6.3, “Thread Pool Tuning”](#).

**B.14.5: Are there recommended Thread Pool configurations?**

The Thread Pool has a number of user case driven configuration parameters that affect its performance. To learn about these and tips on tuning, see [Section 8.11.6.3, “Thread Pool Tuning”](#).



---

# Appendix C. Errors, Error Codes, and Common Problems

## Table of Contents

C.1. Sources of Error Information .....	3659
C.2. Types of Error Values .....	3659
C.3. Server Error Codes and Messages .....	3660
C.4. Client Error Codes and Messages .....	3718
C.5. Problems and Common Errors .....	3722
C.5.1. How to Determine What Is Causing a Problem .....	3722
C.5.2. Common Errors When Using MySQL Programs .....	3723
C.5.3. Installation-Related Issues .....	3735
C.5.4. Administration-Related Issues .....	3736
C.5.5. Query-Related Issues .....	3743
C.5.6. Optimizer-Related Issues .....	3749
C.5.7. Table Definition-Related Issues .....	3750
C.5.8. Known Issues in MySQL .....	3751

This appendix lists common problems and errors that may occur and potential resolutions, in addition to listing the errors that may appear when you call MySQL from any host language. The first section covers problems and resolutions. Detailed information on errors is provided: One list displays server error messages. Another list displays client program messages.

## C.1. Sources of Error Information

There are several sources of error information in MySQL:

- Each SQL statement executed results in an error code, an SQLSTATE value, and an error message, as described in [Section C.2, “Types of Error Values”](#). These errors are returned from the server side; see [Section C.3, “Server Error Codes and Messages”](#).
- Errors can occur on the client side, usually involving problems communicating with the server; see [Section C.4, “Client Error Codes and Messages”](#).
- SQL statement warning and error information is available through the `SHOW WARNINGS` and `SHOW ERRORS` statements. The `warning_count` [564] system variable indicates the number of errors, warnings, and notes. The `error_count` [480] system variable indicates the number of errors. Its value excludes warnings and notes.
- `SHOW SLAVE STATUS` statement output includes information about replication errors occurring on the slave side.
- `SHOW ENGINE INNODB STATUS` statement output includes information about the most recent foreign key error if a `CREATE TABLE` statement for an `InnoDB` table fails.
- The `percona` program provides information from the command line about error numbers. See [Section 4.8.1, “percona — Explain Error Codes”](#).

Descriptions of server and client errors are provided later in this Appendix. For information about errors related to `InnoDB`, see [Section 14.2.3.14, “InnoDB Error Handling”](#).

## C.2. Types of Error Values

When an error occurs in MySQL, the server returns two types of error values:

- A MySQL-specific error code. This value is numeric. It is not portable to other database systems.
- An SQLSTATE value. The value is a five-character string (for example, `'42S02'`). The values are taken from ANSI SQL and ODBC and are more standardized.

A message string that provides a textual description of the error is also available.

When an error occurs, the MySQL error code, SQLSTATE value, and message string are available using C API functions:

- MySQL error code: Call `mysql_errno()`
- SQLSTATE value: Call `mysql_sqlstate()`
- Error message: Call `mysql_error()`

For prepared statements, the corresponding error functions are `mysql_stmt_errno()`, `mysql_stmt_sqlstate()`, and `mysql_stmt_error()`. All error functions are described in [Section 21.8, “MySQL C API”](#).

The number of errors, warnings, and notes for the previous statement can be obtained by calling `mysql_warning_count()`. See [Section 21.8.6.73, “mysql\\_warning\\_count\(\)”](#).

The first two characters of an SQLSTATE value indicate the error class:

- Class = '00' indicates success.
- Class = '01' indicates a warning.
- Class = '02' indicates “not found.” This is relevant within the context of cursors and is used to control what happens when a cursor reaches the end of a data set. This condition also occurs for `SELECT ... INTO var_list` statements that retrieve no rows.
- Class > '02' indicates an exception.

## C.3. Server Error Codes and Messages

MySQL programs have access to several types of error information when the server returns an error. For example, the `mysql` client program displays errors using the following format:

```
shell> SELECT * FROM no_such_table;
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```

The message displayed contains three types of information:

- A numeric error code ([1146](#)). This number is MySQL-specific and is not portable to other database systems.
- A five-character SQLSTATE value ('[42S02](#)'). The values are taken from ANSI SQL and ODBC and are more standardized. Not all MySQL error numbers have corresponding SQLSTATE values. In these cases, '[HY000](#)' (general error) is used.
- A message string that provides a textual description of the error.

For error checking, use error codes, not error messages. Error messages do not change often, but it is possible. Also if the database administrator changes the language setting, that affects the language of error messages.

Error codes are stable across GA releases of a given MySQL series. Before a series reaches GA status, new codes may still be under development and subject to change.

Server error information comes from the following source files. For details about the way that error information is defined, see the [MySQL Internals Manual](#).

- Error message information is listed in the `share/errmsg.txt` file. `%d` and `%s` represent numbers and strings, respectively, that are substituted into the Message values when they are displayed.
- The Error values listed in `share/errmsg.txt` are used to generate the definitions in the `include/mysql_d_error.h` and `include/mysql_d_ename.h` MySQL source files.

- The SQLSTATE values listed in [share/errmsg.txt](#) are used to generate the definitions in the [include/sql\\_state.h](#) MySQL source file.

Because updates are frequent, it is possible that those files will contain additional error information not listed here.

- Error: 1000 SQLSTATE: HY000 ([ER\\_HASHCHK \[3661\]](#))  
Message: hashchk
- Error: 1001 SQLSTATE: HY000 ([ER\\_NISAMCHK \[3661\]](#))  
Message: isamchk
- Error: 1002 SQLSTATE: HY000 ([ER\\_NO \[3661\]](#))  
Message: NO
- Error: 1003 SQLSTATE: HY000 ([ER\\_YES \[3661\]](#))  
Message: YES
- Error: 1004 SQLSTATE: HY000 ([ER\\_CANT\\_CREATE\\_FILE \[3661\]](#))  
Message: Can't create file '%s' (errno: %d - %s)
- Error: 1005 SQLSTATE: HY000 ([ER\\_CANT\\_CREATE\\_TABLE \[3661\]](#))  
Message: Can't create table '%s' (errno: %d)
- Error: 1006 SQLSTATE: HY000 ([ER\\_CANT\\_CREATE\\_DB \[3661\]](#))  
Message: Can't create database '%s' (errno: %d)
- Error: 1007 SQLSTATE: HY000 ([ER\\_DB\\_CREATE\\_EXISTS \[3661\]](#))  
Message: Can't create database '%s'; database exists
- Error: 1008 SQLSTATE: HY000 ([ER\\_DB\\_DROP\\_EXISTS \[3661\]](#))  
Message: Can't drop database '%s'; database doesn't exist
- Error: 1009 SQLSTATE: HY000 ([ER\\_DB\\_DROP\\_DELETE \[3661\]](#))  
Message: Error dropping database (can't delete '%s', errno: %d)
- Error: 1010 SQLSTATE: HY000 ([ER\\_DB\\_DROP\\_RMDIR \[3661\]](#))  
Message: Error dropping database (can't rmdir '%s', errno: %d)
- Error: 1011 SQLSTATE: HY000 ([ER\\_CANT\\_DELETE\\_FILE \[3661\]](#))  
Message: Error on delete of '%s' (errno: %d - %s)
- Error: 1012 SQLSTATE: HY000 ([ER\\_CANT\\_FIND\\_SYSTEM\\_REC \[3661\]](#))  
Message: Can't read record in system table
- Error: 1013 SQLSTATE: HY000 ([ER\\_CANT\\_GET\\_STAT \[3661\]](#))  
Message: Can't get status of '%s' (errno: %d - %s)
- Error: 1014 SQLSTATE: HY000 ([ER\\_CANT\\_GET\\_WD \[3661\]](#))  
Message: Can't get working directory (errno: %d - %s)

- Error: 1015 SQLSTATE: HY000 (ER\_CANT\_LOCK [3662])  
Message: Can't lock file (errno: %d - %s)
- Error: 1016 SQLSTATE: HY000 (ER\_CANT\_OPEN\_FILE [3662])  
Message: Can't open file: '%s' (errno: %d - %s)
- Error: 1017 SQLSTATE: HY000 (ER\_FILE\_NOT\_FOUND [3662])  
Message: Can't find file: '%s' (errno: %d - %s)
- Error: 1018 SQLSTATE: HY000 (ER\_CANT\_READ\_DIR [3662])  
Message: Can't read dir of '%s' (errno: %d - %s)
- Error: 1019 SQLSTATE: HY000 (ER\_CANT\_SET\_WD [3662])  
Message: Can't change dir to '%s' (errno: %d - %s)
- Error: 1020 SQLSTATE: HY000 (ER\_CHECKREAD [3662])  
Message: Record has changed since last read in table '%s'
- Error: 1021 SQLSTATE: HY000 (ER\_DISK\_FULL [3662])  
Message: Disk full (%s); waiting for someone to free some space... (errno: %d - %s)
- Error: 1022 SQLSTATE: 23000 (ER\_DUP\_KEY [3662])  
Message: Can't write; duplicate key in table '%s'
- Error: 1023 SQLSTATE: HY000 (ER\_ERROR\_ON\_CLOSE [3662])  
Message: Error on close of '%s' (errno: %d - %s)
- Error: 1024 SQLSTATE: HY000 (ER\_ERROR\_ON\_READ [3662])  
Message: Error reading file '%s' (errno: %d - %s)
- Error: 1025 SQLSTATE: HY000 (ER\_ERROR\_ON\_RENAME [3662])  
Message: Error on rename of '%s' to '%s' (errno: %d - %s)
- Error: 1026 SQLSTATE: HY000 (ER\_ERROR\_ON\_WRITE [3662])  
Message: Error writing file '%s' (errno: %d - %s)
- Error: 1027 SQLSTATE: HY000 (ER\_FILE\_USED [3662])  
Message: '%s' is locked against change
- Error: 1028 SQLSTATE: HY000 (ER\_FILSORT\_ABORT [3662])  
Message: Sort aborted
- Error: 1029 SQLSTATE: HY000 (ER\_FORM\_NOT\_FOUND [3662])  
Message: View '%s' doesn't exist for '%s'
- Error: 1030 SQLSTATE: HY000 (ER\_GET\_ERRNO [3662])  
Message: Got error %d from storage engine
- Error: 1031 SQLSTATE: HY000 (ER\_ILLEGAL HA [3662])

Message: Table storage engine for '%s' doesn't have this option

- Error: [1032](#) SQLSTATE: [HY000](#) ([ER\\_KEY\\_NOT\\_FOUND](#) [\[3663\]](#))

Message: Can't find record in '%s'

- Error: [1033](#) SQLSTATE: [HY000](#) ([ER\\_NOT\\_FORM\\_FILE](#) [\[3663\]](#))

Message: Incorrect information in file: '%s'

- Error: [1034](#) SQLSTATE: [HY000](#) ([ER\\_NOT\\_KEYFILE](#) [\[3663\]](#))

Message: Incorrect key file for table '%s'; try to repair it

- Error: [1035](#) SQLSTATE: [HY000](#) ([ER\\_OLD\\_KEYFILE](#) [\[3663\]](#))

Message: Old key file for table '%s'; repair it!

- Error: [1036](#) SQLSTATE: [HY000](#) ([ER\\_OPEN\\_AS\\_READONLY](#) [\[3663\]](#))

Message: Table '%s' is read only

- Error: [1037](#) SQLSTATE: [HY001](#) ([ER\\_OUTOFMEMORY](#) [\[3663\]](#))

Message: Out of memory; restart server and try again (needed %d bytes)

- Error: [1038](#) SQLSTATE: [HY001](#) ([ER\\_OUT\\_OF\\_SORTMEMORY](#) [\[3663\]](#))

Message: Out of sort memory, consider increasing server sort buffer size

- Error: [1039](#) SQLSTATE: [HY000](#) ([ER\\_UNEXPECTED\\_EOF](#) [\[3663\]](#))

Message: Unexpected EOF found when reading file '%s' (errno: %d - %s)

- Error: [1040](#) SQLSTATE: [08004](#) ([ER\\_CON\\_COUNT\\_ERROR](#) [\[3663\]](#))

Message: Too many connections

- Error: [1041](#) SQLSTATE: [HY000](#) ([ER\\_OUT\\_OF\\_RESOURCES](#) [\[3663\]](#))

Message: Out of memory; check if mysqld or some other process uses all available memory; if not, you may have to use 'ulimit' to allow mysqld to use more memory or you can add more swap space

- Error: [1042](#) SQLSTATE: [08S01](#) ([ER\\_BAD\\_HOST\\_ERROR](#) [\[3663\]](#))

Message: Can't get hostname for your address

- Error: [1043](#) SQLSTATE: [08S01](#) ([ER\\_HANDSHAKE\\_ERROR](#) [\[3663\]](#))

Message: Bad handshake

- Error: [1044](#) SQLSTATE: [42000](#) ([ER\\_DBACCESS\\_DENIED\\_ERROR](#) [\[3663\]](#))

Message: Access denied for user '%s'@'%s' to database '%s'

- Error: [1045](#) SQLSTATE: [28000](#) ([ER\\_ACCESS\\_DENIED\\_ERROR](#) [\[3663\]](#))

Message: Access denied for user '%s'@'%s' (using password: %s)

- Error: [1046](#) SQLSTATE: [3D000](#) ([ER\\_NO\\_DB\\_ERROR](#) [\[3663\]](#))

Message: No database selected

- Error: [1047](#) SQLSTATE: [08S01](#) ([ER\\_UNKNOWN\\_COM\\_ERROR](#) [\[3663\]](#))

Message: Unknown command

- Error: 1048 SQLSTATE: 23000 (ER\_BAD\_NULL\_ERROR [3664])

Message: Column '%s' cannot be null

- Error: 1049 SQLSTATE: 42000 (ER\_BAD\_DB\_ERROR [3664])

Message: Unknown database '%s'

- Error: 1050 SQLSTATE: 42S01 (ER\_TABLE\_EXISTS\_ERROR [3664])

Message: Table '%s' already exists

- Error: 1051 SQLSTATE: 42S02 (ER\_BAD\_TABLE\_ERROR [3664])

Message: Unknown table '%s'

- Error: 1052 SQLSTATE: 23000 (ER\_NON\_UNIQ\_ERROR [3664])

Message: Column '%s' in %s is ambiguous

- Error: 1053 SQLSTATE: 08S01 (ER\_SERVER\_SHUTDOWN [3664])

Message: Server shutdown in progress

- Error: 1054 SQLSTATE: 42S22 (ER\_BAD\_FIELD\_ERROR [3664])

Message: Unknown column '%s' in '%s'

- Error: 1055 SQLSTATE: 42000 (ER\_WRONG\_FIELD\_WITH\_GROUP [3664])

Message: '%s' isn't in GROUP BY

- Error: 1056 SQLSTATE: 42000 (ER\_WRONG\_GROUP\_FIELD [3664])

Message: Can't group on '%s'

- Error: 1057 SQLSTATE: 42000 (ER\_WRONG\_SUM\_SELECT [3664])

Message: Statement has sum functions and columns in same statement

- Error: 1058 SQLSTATE: 21S01 (ER\_WRONG\_VALUE\_COUNT [3664])

Message: Column count doesn't match value count

- Error: 1059 SQLSTATE: 42000 (ER\_TOO\_LONG\_IDENT [3664])

Message: Identifier name '%s' is too long

- Error: 1060 SQLSTATE: 42S21 (ER\_DUP\_FIELDNAME [3664])

Message: Duplicate column name '%s'

- Error: 1061 SQLSTATE: 42000 (ER\_DUP\_KEYNAME [3664])

Message: Duplicate key name '%s'

- Error: 1062 SQLSTATE: 23000 (ER\_DUP\_ENTRY [3664])

Message: Duplicate entry '%s' for key %d

As of MySQL 5.1.20, the message returned with this error uses the format string for [ER\\_DUP\\_ENTRY\\_WITH\\_KEY\\_NAME \[3698\]](#).



- Error: 1063 SQLSTATE: 42000 ([ER\\_WRONG\\_FIELD\\_SPEC \[3665\]](#))  
Message: Incorrect column specifier for column '%s'
- Error: 1064 SQLSTATE: 42000 ([ER\\_PARSE\\_ERROR \[3665\]](#))  
Message: %s near '%s' at line %d
- Error: 1065 SQLSTATE: 42000 ([ER\\_EMPTY\\_QUERY \[3665\]](#))  
Message: Query was empty
- Error: 1066 SQLSTATE: 42000 ([ER\\_NONUNIQ\\_TABLE \[3665\]](#))  
Message: Not unique table/alias: '%s'
- Error: 1067 SQLSTATE: 42000 ([ER\\_INVALID\\_DEFAULT \[3665\]](#))  
Message: Invalid default value for '%s'
- Error: 1068 SQLSTATE: 42000 ([ER\\_MULTIPLE\\_PRI\\_KEY \[3665\]](#))  
Message: Multiple primary key defined
- Error: 1069 SQLSTATE: 42000 ([ER\\_TOO\\_MANY\\_KEYS \[3665\]](#))  
Message: Too many keys specified; max %d keys allowed
- Error: 1070 SQLSTATE: 42000 ([ER\\_TOO\\_MANY\\_KEY\\_PARTS \[3665\]](#))  
Message: Too many key parts specified; max %d parts allowed
- Error: 1071 SQLSTATE: 42000 ([ER\\_TOO\\_LONG\\_KEY \[3665\]](#))  
Message: Specified key was too long; max key length is %d bytes
- Error: 1072 SQLSTATE: 42000 ([ER\\_KEY\\_COLUMN\\_DOES\\_NOT\\_EXIST \[3665\]](#))  
Message: Key column '%s' doesn't exist in table
- Error: 1073 SQLSTATE: 42000 ([ER\\_BLOB\\_USED\\_AS\\_KEY \[3665\]](#))  
Message: BLOB column '%s' can't be used in key specification with the used table type
- Error: 1074 SQLSTATE: 42000 ([ER\\_TOO\\_BIG\\_FIELDLENGTH \[3665\]](#))  
Message: Column length too big for column '%s' (max = %lu); use BLOB or TEXT instead
- Error: 1075 SQLSTATE: 42000 ([ER\\_WRONG\\_AUTO\\_KEY \[3665\]](#))  
Message: Incorrect table definition; there can be only one auto column and it must be defined as a key
- Error: 1076 SQLSTATE: HY000 ([ER\\_READY \[3665\]](#))  
Message: %s: ready for connections. Version: '%s' socket: '%s' port: %d
- Error: 1077 SQLSTATE: HY000 ([ER\\_NORMAL\\_SHUTDOWN \[3665\]](#))  
Message: %s: Normal shutdown
- Error: 1078 SQLSTATE: HY000 ([ER\\_GOT\\_SIGNAL \[3665\]](#))  
Message: %s: Got signal %d. Aborting!

- Error: [1079](#) SQLSTATE: [HY000](#) ([ER\\_SHUTDOWN\\_COMPLETE](#) [\[3666\]](#))  
Message: %s: Shutdown complete
- Error: [1080](#) SQLSTATE: [08S01](#) ([ER\\_FORCING\\_CLOSE](#) [\[3666\]](#))  
Message: %s: Forcing close of thread %ld user: '%s'
- Error: [1081](#) SQLSTATE: [08S01](#) ([ER\\_IPSOCK\\_ERROR](#) [\[3666\]](#))  
Message: Can't create IP socket
- Error: [1082](#) SQLSTATE: [42S12](#) ([ER\\_NO\\_SUCH\\_INDEX](#) [\[3666\]](#))  
Message: Table '%s' has no index like the one used in CREATE INDEX; recreate the table
- Error: [1083](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_FIELD\\_TERMINATORS](#) [\[3666\]](#))  
Message: Field separator argument is not what is expected; check the manual
- Error: [1084](#) SQLSTATE: [42000](#) ([ER\\_BLOBS\\_AND\\_NO\\_TERMINATED](#) [\[3666\]](#))  
Message: You can't use fixed rowlength with BLOBs; please use 'fields terminated by'
- Error: [1085](#) SQLSTATE: [HY000](#) ([ER\\_TEXTFILE\\_NOT\\_READABLE](#) [\[3666\]](#))  
Message: The file '%s' must be in the database directory or be readable by all
- Error: [1086](#) SQLSTATE: [HY000](#) ([ER\\_FILE\\_EXISTS\\_ERROR](#) [\[3666\]](#))  
Message: File '%s' already exists
- Error: [1087](#) SQLSTATE: [HY000](#) ([ER\\_LOAD\\_INFO](#) [\[3666\]](#))  
Message: Records: %ld Deleted: %ld Skipped: %ld Warnings: %ld
- Error: [1088](#) SQLSTATE: [HY000](#) ([ER\\_ALTER\\_INFO](#) [\[3666\]](#))  
Message: Records: %ld Duplicates: %ld
- Error: [1089](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_SUB\\_KEY](#) [\[3666\]](#))  
Message: Incorrect prefix key; the used key part isn't a string, the used length is longer than the key part, or the storage engine doesn't support unique prefix keys
- Error: [1090](#) SQLSTATE: [42000](#) ([ER\\_CANT\\_REMOVE\\_ALL\\_FIELDS](#) [\[3666\]](#))  
Message: You can't delete all columns with ALTER TABLE; use DROP TABLE instead
- Error: [1091](#) SQLSTATE: [42000](#) ([ER\\_CANT\\_DROP\\_FIELD\\_OR\\_KEY](#) [\[3666\]](#))  
Message: Can't DROP '%s'; check that column/key exists
- Error: [1092](#) SQLSTATE: [HY000](#) ([ER\\_INSERT\\_INFO](#) [\[3666\]](#))  
Message: Records: %ld Duplicates: %ld Warnings: %ld
- Error: [1093](#) SQLSTATE: [HY000](#) ([ER\\_UPDATE\\_TABLE\\_USED](#) [\[3666\]](#))  
Message: You can't specify target table '%s' for update in FROM clause
- Error: [1094](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_SUCH\\_THREAD](#) [\[3666\]](#))  
Message: Unknown thread id: %lu

- Error: [1095](#) SQLSTATE: [HY000](#) ([ER\\_KILL\\_DENIED\\_ERROR](#) [\[3667\]](#))  
Message: You are not owner of thread %lu
- Error: [1096](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_TABLES\\_USED](#) [\[3667\]](#))  
Message: No tables used
- Error: [1097](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_BIG\\_SET](#) [\[3667\]](#))  
Message: Too many strings for column %s and SET
- Error: [1098](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_UNIQUE\\_LOGFILE](#) [\[3667\]](#))  
Message: Can't generate a unique log-filename %s.(1-999)
- Error: [1099](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_NOT\\_LOCKED\\_FOR\\_WRITE](#) [\[3667\]](#))  
Message: Table '%s' was locked with a READ lock and can't be updated
- Error: [1100](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_NOT\\_LOCKED](#) [\[3667\]](#))  
Message: Table '%s' was not locked with LOCK TABLES
- Error: [1101](#) SQLSTATE: [42000](#) ([ER\\_BLOB\\_CANT\\_HAVE\\_DEFAULT](#) [\[3667\]](#))  
Message: BLOB/TEXT column '%s' can't have a default value
- Error: [1102](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_DB\\_NAME](#) [\[3667\]](#))  
Message: Incorrect database name '%s'
- Error: [1103](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_TABLE\\_NAME](#) [\[3667\]](#))  
Message: Incorrect table name '%s'
- Error: [1104](#) SQLSTATE: [42000](#) ([ER\\_TOO\\_BIG\\_SELECT](#) [\[3667\]](#))  
Message: The SELECT would examine more than MAX\_JOIN\_SIZE rows; check your WHERE and use SET SQL\_BIG\_SELECTS=1 or SET MAX\_JOIN\_SIZE=# if the SELECT is okay
- Error: [1105](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_ERROR](#) [\[3667\]](#))  
Message: Unknown error
- Error: [1106](#) SQLSTATE: [42000](#) ([ER\\_UNKNOWN\\_PROCEDURE](#) [\[3667\]](#))  
Message: Unknown procedure '%s'
- Error: [1107](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_PARAMCOUNT\\_TO\\_PROCEDURE](#) [\[3667\]](#))  
Message: Incorrect parameter count to procedure '%s'
- Error: [1108](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_PARAMETERS\\_TO\\_PROCEDURE](#) [\[3667\]](#))  
Message: Incorrect parameters to procedure '%s'
- Error: [1109](#) SQLSTATE: [42S02](#) ([ER\\_UNKNOWN\\_TABLE](#) [\[3667\]](#))  
Message: Unknown table '%s' in %s
- Error: [1110](#) SQLSTATE: [42000](#) ([ER\\_FIELD\\_SPECIFIED\\_TWICE](#) [\[3667\]](#))  
Message: Column '%s' specified twice

- Error: [1111](#) SQLSTATE: [HY000](#) ([ER\\_INVALID\\_GROUP\\_FUNC\\_USE](#) [\[3668\]](#))  
Message: Invalid use of group function
- Error: [1112](#) SQLSTATE: [42000](#) ([ER\\_UNSUPPORTED\\_EXTENSION](#) [\[3668\]](#))  
Message: Table '%s' uses an extension that doesn't exist in this MySQL version
- Error: [1113](#) SQLSTATE: [42000](#) ([ER\\_TABLE\\_MUST\\_HAVE\\_COLUMNS](#) [\[3668\]](#))  
Message: A table must have at least 1 column
- Error: [1114](#) SQLSTATE: [HY000](#) ([ER\\_RECORD\\_FILE\\_FULL](#) [\[3668\]](#))  
Message: The table '%s' is full
- Error: [1115](#) SQLSTATE: [42000](#) ([ER\\_UNKNOWN\\_CHARACTER\\_SET](#) [\[3668\]](#))  
Message: Unknown character set: '%s'
- Error: [1116](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_MANY\\_TABLES](#) [\[3668\]](#))  
Message: Too many tables; MySQL can only use %d tables in a join
- Error: [1117](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_MANY\\_FIELDS](#) [\[3668\]](#))  
Message: Too many columns
- Error: [1118](#) SQLSTATE: [42000](#) ([ER\\_TOO\\_BIG\\_ROWSIZE](#) [\[3668\]](#))  
Message: Row size too large. The maximum row size for the used table type, not counting BLOBs, is %ld. This includes storage overhead, check the manual. You have to change some columns to TEXT or BLOBs
- Error: [1119](#) SQLSTATE: [HY000](#) ([ER\\_STACK\\_OVERRUN](#) [\[3668\]](#))  
Message: Thread stack overrun: Used: %ld of a %ld stack. Use 'mysqld --thread\_stack=#' to specify a bigger stack if needed
- Error: [1120](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_OUTER\\_JOIN](#) [\[3668\]](#))  
Message: Cross dependency found in OUTER JOIN; examine your ON conditions
- Error: [1121](#) SQLSTATE: [42000](#) ([ER\\_NULL\\_COLUMN\\_IN\\_INDEX](#) [\[3668\]](#))  
Message: Table handler doesn't support NULL in given index. Please change column '%s' to be NOT NULL or use another handler
- Error: [1122](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_FIND\\_UDF](#) [\[3668\]](#))  
Message: Can't load function '%s'
- Error: [1123](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_INITIALIZE\\_UDF](#) [\[3668\]](#))  
Message: Can't initialize function '%s'; %s
- Error: [1124](#) SQLSTATE: [HY000](#) ([ER\\_UDF\\_NO\\_PATHS](#) [\[3668\]](#))  
Message: No paths allowed for shared library
- Error: [1125](#) SQLSTATE: [HY000](#) ([ER\\_UDF\\_EXISTS](#) [\[3668\]](#))  
Message: Function '%s' already exists

- Error: [1126](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_OPEN\\_LIBRARY](#) [\[3669\]](#))  
Message: Can't open shared library '%s' (errno: %d %s)
- Error: [1127](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_FIND\\_DL\\_ENTRY](#) [\[3669\]](#))  
Message: Can't find symbol '%s' in library
- Error: [1128](#) SQLSTATE: [HY000](#) ([ER\\_FUNCTION\\_NOT\\_DEFINED](#) [\[3669\]](#))  
Message: Function '%s' is not defined
- Error: [1129](#) SQLSTATE: [HY000](#) ([ER\\_HOST\\_IS\\_BLOCKED](#) [\[3669\]](#))  
Message: Host '%s' is blocked because of many connection errors; unblock with 'mysqladmin flush-hosts'
- Error: [1130](#) SQLSTATE: [HY000](#) ([ER\\_HOST\\_NOT\\_PRIVILEGED](#) [\[3669\]](#))  
Message: Host '%s' is not allowed to connect to this MySQL server
- Error: [1131](#) SQLSTATE: [42000](#) ([ER\\_PASSWORD\\_ANONYMOUS\\_USER](#) [\[3669\]](#))  
Message: You are using MySQL as an anonymous user and anonymous users are not allowed to change passwords
- Error: [1132](#) SQLSTATE: [42000](#) ([ER\\_PASSWORD\\_NOT\\_ALLOWED](#) [\[3669\]](#))  
Message: You must have privileges to update tables in the mysql database to be able to change passwords for others
- Error: [1133](#) SQLSTATE: [42000](#) ([ER\\_PASSWORD\\_NO\\_MATCH](#) [\[3669\]](#))  
Message: Can't find any matching row in the user table
- Error: [1134](#) SQLSTATE: [HY000](#) ([ER\\_UPDATE\\_INFO](#) [\[3669\]](#))  
Message: Rows matched: %ld Changed: %ld Warnings: %ld
- Error: [1135](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_CREATE\\_THREAD](#) [\[3669\]](#))  
Message: Can't create a new thread (errno %d); if you are not out of available memory, you can consult the manual for a possible OS-dependent bug
- Error: [1136](#) SQLSTATE: [21S01](#) ([ER\\_WRONG\\_VALUE\\_COUNT\\_ON\\_ROW](#) [\[3669\]](#))  
Message: Column count doesn't match value count at row %ld
- Error: [1137](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_REOPEN\\_TABLE](#) [\[3669\]](#))  
Message: Can't reopen table: '%s'
- Error: [1138](#) SQLSTATE: [22004](#) ([ER\\_INVALID\\_USE\\_OF\\_NULL](#) [\[3669\]](#))  
Message: Invalid use of NULL value
- Error: [1139](#) SQLSTATE: [42000](#) ([ER\\_REGEXP\\_ERROR](#) [\[3669\]](#))  
Message: Got error '%s' from regexp
- Error: [1140](#) SQLSTATE: [42000](#) ([ER\\_MIX\\_OF\\_GROUP\\_FUNC\\_AND\\_FIELDS](#) [\[3669\]](#))  
Message: Mixing of GROUP columns (MIN(),MAX(),COUNT(),...) with no GROUP columns is illegal if there is no GROUP BY clause

- Error: 1141 SQLSTATE: 42000 (ER\_NONEXISTING\_GRANT [3670])  
Message: There is no such grant defined for user '%s' on host '%s'
- Error: 1142 SQLSTATE: 42000 (ER\_TABLEACCESS\_DENIED\_ERROR [3670])  
Message: %s command denied to user '%s'@'%s' for table '%s'
- Error: 1143 SQLSTATE: 42000 (ER\_COLUMNACCESS\_DENIED\_ERROR [3670])  
Message: %s command denied to user '%s'@'%s' for column '%s' in table '%s'
- Error: 1144 SQLSTATE: 42000 (ER\_ILLEGAL\_GRANT\_FOR\_TABLE [3670])  
Message: Illegal GRANT/REVOKE command; please consult the manual to see which privileges can be used
- Error: 1145 SQLSTATE: 42000 (ER\_GRANT\_WRONG\_HOST\_OR\_USER [3670])  
Message: The host or user argument to GRANT is too long
- Error: 1146 SQLSTATE: 42S02 (ER\_NO\_SUCH\_TABLE [3670])  
Message: Table '%s.%s' doesn't exist
- Error: 1147 SQLSTATE: 42000 (ER\_NONEXISTING\_TABLE\_GRANT [3670])  
Message: There is no such grant defined for user '%s' on host '%s' on table '%s'
- Error: 1148 SQLSTATE: 42000 (ER\_NOT\_ALLOWED\_COMMAND [3670])  
Message: The used command is not allowed with this MySQL version
- Error: 1149 SQLSTATE: 42000 (ER\_SYNTAX\_ERROR [3670])  
Message: You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use
- Error: 1150 SQLSTATE: HY000 (ER\_DELAYED\_CANT\_CHANGE\_LOCK [3670])  
Message: Delayed insert thread couldn't get requested lock for table %s
- Error: 1151 SQLSTATE: HY000 (ER\_TOO\_MANY\_DELAYED\_THREADS [3670])  
Message: Too many delayed threads in use
- Error: 1152 SQLSTATE: 08S01 (ER\_ABORTING\_CONNECTION [3670])  
Message: Aborted connection %ld to db: '%s' user: '%s' (%s)
- Error: 1153 SQLSTATE: 08S01 (ER\_NET\_PACKET\_TOO\_LARGE [3670])  
Message: Got a packet bigger than 'max\_allowed\_packet' bytes
- Error: 1154 SQLSTATE: 08S01 (ER\_NET\_READ\_ERROR\_FROM\_PIPE [3670])  
Message: Got a read error from the connection pipe
- Error: 1155 SQLSTATE: 08S01 (ER\_NET\_FCNTL\_ERROR [3670])  
Message: Got an error from fcntl()
- Error: 1156 SQLSTATE: 08S01 (ER\_NET\_PACKETS\_OUT\_OF\_ORDER [3670])  
Message: Got packets out of order

- Error: 1157 SQLSTATE: 08S01 ([ER\\_NET\\_UNCOMPRESS\\_ERROR \[3671\]](#))  
Message: Couldn't uncompress communication packet
- Error: 1158 SQLSTATE: 08S01 ([ER\\_NET\\_READ\\_ERROR \[3671\]](#))  
Message: Got an error reading communication packets
- Error: 1159 SQLSTATE: 08S01 ([ER\\_NET\\_READ\\_INTERRUPTED \[3671\]](#))  
Message: Got timeout reading communication packets
- Error: 1160 SQLSTATE: 08S01 ([ER\\_NET\\_ERROR\\_ON\\_WRITE \[3671\]](#))  
Message: Got an error writing communication packets
- Error: 1161 SQLSTATE: 08S01 ([ER\\_NET\\_WRITE\\_INTERRUPTED \[3671\]](#))  
Message: Got timeout writing communication packets
- Error: 1162 SQLSTATE: 42000 ([ER\\_TOO\\_LONG\\_STRING \[3671\]](#))  
Message: Result string is longer than 'max\_allowed\_packet' bytes
- Error: 1163 SQLSTATE: 42000 ([ER\\_TABLE\\_CANT\\_HANDLE\\_BLOB \[3671\]](#))  
Message: The used table type doesn't support BLOB/TEXT columns
- Error: 1164 SQLSTATE: 42000 ([ER\\_TABLE\\_CANT\\_HANDLE\\_AUTO\\_INCREMENT \[3671\]](#))  
Message: The used table type doesn't support AUTO\_INCREMENT columns
- Error: 1165 SQLSTATE: HY000 ([ER\\_DELAYED\\_INSERT\\_TABLE\\_LOCKED \[3671\]](#))  
Message: INSERT DELAYED can't be used with table '%s' because it is locked with LOCK TABLES
- Error: 1166 SQLSTATE: 42000 ([ER\\_WRONG\\_COLUMN\\_NAME \[3671\]](#))  
Message: Incorrect column name '%s'
- Error: 1167 SQLSTATE: 42000 ([ER\\_WRONG\\_KEY\\_COLUMN \[3671\]](#))  
Message: The used storage engine can't index column '%s'
- Error: 1168 SQLSTATE: HY000 ([ER\\_WRONG\\_MRG\\_TABLE \[3671\]](#))  
Message: Unable to open underlying table which is differently defined or of non-MyISAM type or doesn't exist
- Error: 1169 SQLSTATE: 23000 ([ER\\_DUP\\_UNIQUE \[3671\]](#))  
Message: Can't write, because of unique constraint, to table '%s'
- Error: 1170 SQLSTATE: 42000 ([ER\\_BLOB\\_KEY\\_WITHOUT\\_LENGTH \[3671\]](#))  
Message: BLOB/TEXT column '%s' used in key specification without a key length
- Error: 1171 SQLSTATE: 42000 ([ER\\_PRIMARY\\_CANT\\_HAVE\\_NULL \[3671\]](#))  
Message: All parts of a PRIMARY KEY must be NOT NULL; if you need NULL in a key, use UNIQUE instead
- Error: 1172 SQLSTATE: 42000 ([ER\\_TOO\\_MANY\\_ROWS \[3671\]](#))  
Message: Result consisted of more than one row

- Error: 1173 SQLSTATE: 42000 (ER\_REQUIRES\_PRIMARY\_KEY [3672])  
Message: This table type requires a primary key
- Error: 1174 SQLSTATE: HY000 (ER\_NO\_RAID\_COMPILED [3672])  
Message: This version of MySQL is not compiled with RAID support
- Error: 1175 SQLSTATE: HY000 (ER\_UPDATE\_WITHOUT\_KEY\_IN\_SAFE\_MODE [3672])  
Message: You are using safe update mode and you tried to update a table without a WHERE that uses a KEY column
- Error: 1176 SQLSTATE: 42000 (ER\_KEY\_DOES\_NOT\_EXISTS [3672])  
Message: Key '%s' doesn't exist in table '%s'
- Error: 1177 SQLSTATE: 42000 (ER\_CHECK\_NO\_SUCH\_TABLE [3672])  
Message: Can't open table
- Error: 1178 SQLSTATE: 42000 (ER\_CHECK\_NOT\_IMPLEMENTED [3672])  
Message: The storage engine for the table doesn't support %s
- Error: 1179 SQLSTATE: 25000 (ER\_CANT\_DO\_THIS\_DURING\_AN\_TRANSACTION [3672])  
Message: You are not allowed to execute this command in a transaction
- Error: 1180 SQLSTATE: HY000 (ER\_ERROR\_DURING\_COMMIT [3672])  
Message: Got error %d during COMMIT
- Error: 1181 SQLSTATE: HY000 (ER\_ERROR\_DURING\_ROLLBACK [3672])  
Message: Got error %d during ROLLBACK
- Error: 1182 SQLSTATE: HY000 (ER\_ERROR\_DURING\_FLUSH\_LOGS [3672])  
Message: Got error %d during FLUSH\_LOGS
- Error: 1183 SQLSTATE: HY000 (ER\_ERROR\_DURING\_CHECKPOINT [3672])  
Message: Got error %d during CHECKPOINT
- Error: 1184 SQLSTATE: 08S01 (ER\_NEW\_ABORTING\_CONNECTION [3672])  
Message: Aborted connection %ld to db: '%s' user: '%s' host: '%s' (%s)
- Error: 1185 SQLSTATE: HY000 (ER\_DUMP\_NOT\_IMPLEMENTED [3672])  
Message: The storage engine for the table does not support binary table dump
- Error: 1186 SQLSTATE: HY000 (ER\_FLUSH\_MASTER\_BINLOG\_CLOSED [3672])  
Message: Binlog closed, cannot RESET MASTER
- Error: 1187 SQLSTATE: HY000 (ER\_INDEX\_REBUILD [3672])  
Message: Failed rebuilding the index of dumped table '%s'
- Error: 1188 SQLSTATE: HY000 (ER\_MASTER [3672])  
Message: Error from master: '%s'



- Error: [1189](#) SQLSTATE: [08S01](#) ([ER\\_MASTER\\_NET\\_READ](#) [\[3673\]](#))  
Message: Net error reading from master
- Error: [1190](#) SQLSTATE: [08S01](#) ([ER\\_MASTER\\_NET\\_WRITE](#) [\[3673\]](#))  
Message: Net error writing to master
- Error: [1191](#) SQLSTATE: [HY000](#) ([ER\\_FT\\_MATCHING\\_KEY\\_NOT\\_FOUND](#) [\[3673\]](#))  
Message: Can't find FULLTEXT index matching the column list
- Error: [1192](#) SQLSTATE: [HY000](#) ([ER\\_LOCK\\_OR\\_ACTIVE\\_TRANSACTION](#) [\[3673\]](#))  
Message: Can't execute the given command because you have active locked tables or an active transaction
- Error: [1193](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_SYSTEM\\_VARIABLE](#) [\[3673\]](#))  
Message: Unknown system variable '%s'
- Error: [1194](#) SQLSTATE: [HY000](#) ([ER\\_CRASHED\\_ON\\_USAGE](#) [\[3673\]](#))  
Message: Table '%s' is marked as crashed and should be repaired
- Error: [1195](#) SQLSTATE: [HY000](#) ([ER\\_CRASHED\\_ON\\_REPAIR](#) [\[3673\]](#))  
Message: Table '%s' is marked as crashed and last (automatic?) repair failed
- Error: [1196](#) SQLSTATE: [HY000](#) ([ER\\_WARNING\\_NOT\\_COMPLETE\\_ROLLBACK](#) [\[3673\]](#))  
Message: Some non-transactional changed tables couldn't be rolled back
- Error: [1197](#) SQLSTATE: [HY000](#) ([ER\\_TRANS\\_CACHE\\_FULL](#) [\[3673\]](#))  
Message: Multi-statement transaction required more than 'max\_binlog\_cache\_size' bytes of storage; increase this mysqld variable and try again
- Error: [1198](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_MUST\\_STOP](#) [\[3673\]](#))  
Message: This operation cannot be performed with a running slave; run STOP SLAVE first
- Error: [1199](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_NOT\\_RUNNING](#) [\[3673\]](#))  
Message: This operation requires a running slave; configure slave and do START SLAVE
- Error: [1200](#) SQLSTATE: [HY000](#) ([ER\\_BAD\\_SLAVE](#) [\[3673\]](#))  
Message: The server is not configured as slave; fix in config file or with CHANGE MASTER TO
- Error: [1201](#) SQLSTATE: [HY000](#) ([ER\\_MASTER\\_INFO](#) [\[3673\]](#))  
Message: Could not initialize master info structure; more error messages can be found in the MySQL error log
- Error: [1202](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_THREAD](#) [\[3673\]](#))  
Message: Could not create slave thread; check system resources
- Error: [1203](#) SQLSTATE: [42000](#) ([ER\\_TOO\\_MANY\\_USER\\_CONNECTIONS](#) [\[3673\]](#))  
Message: User %s already has more than 'max\_user\_connections' active connections
- Error: [1204](#) SQLSTATE: [HY000](#) ([ER\\_SET\\_CONSTANTS\\_ONLY](#) [\[3673\]](#))

Message: You may only use constant expressions with SET

- Error: 1205 SQLSTATE: HY000 (ER\_LOCK\_WAIT\_TIMEOUT [3674])

Message: Lock wait timeout exceeded; try restarting transaction

- Error: 1206 SQLSTATE: HY000 (ER\_LOCK\_TABLE\_FULL [3674])

Message: The total number of locks exceeds the lock table size

- Error: 1207 SQLSTATE: 25000 (ER\_READ\_ONLY\_TRANSACTION [3674])

Message: Update locks cannot be acquired during a READ UNCOMMITTED transaction

- Error: 1208 SQLSTATE: HY000 (ER\_DROP\_DB\_WITH\_READ\_LOCK [3674])

Message: DROP DATABASE not allowed while thread is holding global read lock

- Error: 1209 SQLSTATE: HY000 (ER\_CREATE\_DB\_WITH\_READ\_LOCK [3674])

Message: CREATE DATABASE not allowed while thread is holding global read lock

- Error: 1210 SQLSTATE: HY000 (ER\_WRONG\_ARGUMENTS [3674])

Message: Incorrect arguments to %s

- Error: 1211 SQLSTATE: 42000 (ER\_NO\_PERMISSION\_TO\_CREATE\_USER [3674])

Message: '%s'@'%s' is not allowed to create new users

- Error: 1212 SQLSTATE: HY000 (ER\_UNION\_TABLES\_IN\_DIFFERENT\_DIR [3674])

Message: Incorrect table definition; all MERGE tables must be in the same database

- Error: 1213 SQLSTATE: 40001 (ER\_LOCK\_DEADLOCK [3674])

Message: Deadlock found when trying to get lock; try restarting transaction

- Error: 1214 SQLSTATE: HY000 (ER\_TABLE\_CANT\_HANDLE\_FT [3674])

Message: The used table type doesn't support FULLTEXT indexes

- Error: 1215 SQLSTATE: HY000 (ER\_CANNOT\_ADD\_FOREIGN [3674])

Message: Cannot add foreign key constraint

- Error: 1216 SQLSTATE: 23000 (ER\_NO\_REFERENCED\_ROW [3674])

Message: Cannot add or update a child row: a foreign key constraint fails

- Error: 1217 SQLSTATE: 23000 (ER\_ROW\_IS\_REFERENCED [3674])

Message: Cannot delete or update a parent row: a foreign key constraint fails

- Error: 1218 SQLSTATE: 08S01 (ER\_CONNECT\_TO\_MASTER [3674])

Message: Error connecting to master: %s

- Error: 1219 SQLSTATE: HY000 (ER\_QUERY\_ON\_MASTER [3674])

Message: Error running query on master: %s

- Error: 1220 SQLSTATE: HY000 (ER\_ERROR\_WHEN\_EXECUTING\_COMMAND [3674])

Message: Error when executing command %s: %s

- Error: [1221 SQLSTATE: HY000 \(ER\\_WRONG\\_USAGE \[3675\]\)](#)

Message: Incorrect usage of %s and %s

- Error: [1222 SQLSTATE: 21000 \(ER\\_WRONG\\_NUMBER\\_OF\\_COLUMNS\\_IN\\_SELECT \[3675\]\)](#)

Message: The used SELECT statements have a different number of columns

- Error: [1223 SQLSTATE: HY000 \(ER\\_CANT\\_UPDATE\\_WITH\\_READLOCK \[3675\]\)](#)

Message: Can't execute the query because you have a conflicting read lock

- Error: [1224 SQLSTATE: HY000 \(ER\\_MIXING\\_NOT\\_ALLOWED \[3675\]\)](#)

Message: Mixing of transactional and non-transactional tables is disabled

- Error: [1225 SQLSTATE: HY000 \(ER\\_DUP\\_ARGUMENT \[3675\]\)](#)

Message: Option '%s' used twice in statement

- Error: [1226 SQLSTATE: 42000 \(ER\\_USER\\_LIMIT\\_REACHED \[3675\]\)](#)

Message: User '%s' has exceeded the '%s' resource (current value: %ld)

- Error: [1227 SQLSTATE: 42000 \(ER\\_SPECIFIC\\_ACCESS\\_DENIED\\_ERROR \[3675\]\)](#)

Message: Access denied; you need (at least one of) the %s privilege(s) for this operation

- Error: [1228 SQLSTATE: HY000 \(ER\\_LOCAL\\_VARIABLE \[3675\]\)](#)

Message: Variable '%s' is a SESSION variable and can't be used with SET GLOBAL

- Error: [1229 SQLSTATE: HY000 \(ER\\_GLOBAL\\_VARIABLE \[3675\]\)](#)

Message: Variable '%s' is a GLOBAL variable and should be set with SET GLOBAL

- Error: [1230 SQLSTATE: 42000 \(ER\\_NO\\_DEFAULT \[3675\]\)](#)

Message: Variable '%s' doesn't have a default value

- Error: [1231 SQLSTATE: 42000 \(ER\\_WRONG\\_VALUE\\_FOR\\_VAR \[3675\]\)](#)

Message: Variable '%s' can't be set to the value of '%s'

- Error: [1232 SQLSTATE: 42000 \(ER\\_WRONG\\_TYPE\\_FOR\\_VAR \[3675\]\)](#)

Message: Incorrect argument type to variable '%s'

- Error: [1233 SQLSTATE: HY000 \(ER\\_VAR\\_CANT\\_BE\\_READ \[3675\]\)](#)

Message: Variable '%s' can only be set, not read

- Error: [1234 SQLSTATE: 42000 \(ER\\_CANT\\_USE\\_OPTION\\_HERE \[3675\]\)](#)

Message: Incorrect usage/placement of '%s'

- Error: [1235 SQLSTATE: 42000 \(ER\\_NOT\\_SUPPORTED\\_YET \[3675\]\)](#)

Message: This version of MySQL doesn't yet support '%s'

- Error: [1236 SQLSTATE: HY000 \(ER\\_MASTER\\_FATAL\\_ERROR\\_READING\\_BINLOG \[3675\]\)](#)

Message: Got fatal error %d from master when reading data from binary log: '%s'

- Error: 1237 SQLSTATE: HY000 (ER\_SLAVE\_IGNORED\_TABLE [3676])

Message: Slave SQL thread ignored the query because of replicate-\*-\*table rules

- Error: 1238 SQLSTATE: HY000 (ER\_INCORRECT\_GLOBAL\_LOCAL\_VAR [3676])

Message: Variable '%s' is a %s variable

- Error: 1239 SQLSTATE: 42000 (ER\_WRONG\_FK\_DEF [3676])

Message: Incorrect foreign key definition for '%s': %s

- Error: 1240 SQLSTATE: HY000 (ER\_KEY\_REF\_DO\_NOT\_MATCH\_TABLE\_REF [3676])

Message: Key reference and table reference don't match

- Error: 1241 SQLSTATE: 21000 (ER\_OPERAND\_COLUMNS [3676])

Message: Operand should contain %d column(s)

- Error: 1242 SQLSTATE: 21000 (ER\_SUBQUERY\_NO\_1\_ROW [3676])

Message: Subquery returns more than 1 row

- Error: 1243 SQLSTATE: HY000 (ER\_UNKNOWN\_STMT\_HANDLER [3676])

Message: Unknown prepared statement handler (%.\*s) given to %s

- Error: 1244 SQLSTATE: HY000 (ER\_CORRUPT\_HELP\_DB [3676])

Message: Help database is corrupt or does not exist

- Error: 1245 SQLSTATE: HY000 (ER\_CYCLIC\_REFERENCE [3676])

Message: Cyclic reference on subqueries

- Error: 1246 SQLSTATE: HY000 (ER\_AUTO\_CONVERT [3676])

Message: Converting column '%s' from %s to %s

- Error: 1247 SQLSTATE: 42S22 (ER\_ILLEGAL\_REFERENCE [3676])

Message: Reference '%s' not supported (%s)

- Error: 1248 SQLSTATE: 42000 (ER\_DERIVED\_MUST\_HAVE\_ALIAS [3676])

Message: Every derived table must have its own alias

- Error: 1249 SQLSTATE: 01000 (ER\_SELECT\_REDUCED [3676])

Message: Select %u was reduced during optimization

- Error: 1250 SQLSTATE: 42000 (ER\_TABLENAME\_NOT\_ALLOWED\_HERE [3676])

Message: Table '%s' from one of the SELECTs cannot be used in %s

- Error: 1251 SQLSTATE: 08004 (ER\_NOT\_SUPPORTED\_AUTH\_MODE [3676])

Message: Client does not support authentication protocol requested by server; consider upgrading MySQL client

- Error: 1252 SQLSTATE: 42000 (ER\_SPATIAL\_CANT\_HAVE\_NULL [3676])

Message: All parts of a SPATIAL index must be NOT NULL

- Error: [1253](#) SQLSTATE: [42000](#) ([ER\\_COLLATION\\_CHARSET\\_MISMATCH](#) [\[3677\]](#))

Message: COLLATION '%s' is not valid for CHARACTER SET '%s'

- Error: [1254](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_WAS\\_RUNNING](#) [\[3677\]](#))

Message: Slave is already running

- Error: [1255](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_WAS\\_NOT\\_RUNNING](#) [\[3677\]](#))

Message: Slave already has been stopped

- Error: [1256](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_BIG\\_FOR\\_UNCOMPRESS](#) [\[3677\]](#))

Message: Uncompressed data size too large; the maximum size is %d (probably, length of uncompressed data was corrupted)

- Error: [1257](#) SQLSTATE: [HY000](#) ([ER\\_ZLIB\\_Z\\_MEM\\_ERROR](#) [\[3677\]](#))

Message: ZLIB: Not enough memory

- Error: [1258](#) SQLSTATE: [HY000](#) ([ER\\_ZLIB\\_Z\\_BUF\\_ERROR](#) [\[3677\]](#))

Message: ZLIB: Not enough room in the output buffer (probably, length of uncompressed data was corrupted)

- Error: [1259](#) SQLSTATE: [HY000](#) ([ER\\_ZLIB\\_Z\\_DATA\\_ERROR](#) [\[3677\]](#))

Message: ZLIB: Input data corrupted

- Error: [1260](#) SQLSTATE: [HY000](#) ([ER\\_CUT\\_VALUE\\_GROUP\\_CONCAT](#) [\[3677\]](#))

Message: Row %u was cut by GROUP\_CONCAT()

- Error: [1261](#) SQLSTATE: [01000](#) ([ER\\_WARN\\_TOO\\_FEW\\_RECORDS](#) [\[3677\]](#))

Message: Row %ld doesn't contain data for all columns

- Error: [1262](#) SQLSTATE: [01000](#) ([ER\\_WARN\\_TOO\\_MANY\\_RECORDS](#) [\[3677\]](#))

Message: Row %ld was truncated; it contained more data than there were input columns

- Error: [1263](#) SQLSTATE: [22004](#) ([ER\\_WARN\\_NULL\\_TO\\_NOTNULL](#) [\[3677\]](#))

Message: Column set to default value; NULL supplied to NOT NULL column '%s' at row %ld

- Error: [1264](#) SQLSTATE: [22003](#) ([ER\\_WARN\\_DATA\\_OUT\\_OF\\_RANGE](#) [\[3677\]](#))

Message: Out of range value for column '%s' at row %ld

- Error: [1265](#) SQLSTATE: [01000](#) ([WARN\\_DATA\\_TRUNCATED](#) [\[3677\]](#))

Message: Data truncated for column '%s' at row %ld

- Error: [1266](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_USING\\_OTHER\\_HANDLER](#) [\[3677\]](#))

Message: Using storage engine %s for table '%s'

- Error: [1267](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_AGGREGATE\\_2COLLATIONS](#) [\[3677\]](#))

Message: Illegal mix of collations (%s,%s) and (%s,%s) for operation '%s'

- Error: [1268](#) SQLSTATE: [HY000](#) ([ER\\_DROP\\_USER](#) [\[3678\]](#))  
Message: Cannot drop one or more of the requested users
- Error: [1269](#) SQLSTATE: [HY000](#) ([ER\\_REVOKE\\_GRANTS](#) [\[3678\]](#))  
Message: Can't revoke all privileges for one or more of the requested users
- Error: [1270](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_AGGREGATE\\_3COLLATIONS](#) [\[3678\]](#))  
Message: Illegal mix of collations (%s,%s), (%s,%s), (%s,%s) for operation '%s'
- Error: [1271](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_AGGREGATE\\_NCOLLATIONS](#) [\[3678\]](#))  
Message: Illegal mix of collations for operation '%s'
- Error: [1272](#) SQLSTATE: [HY000](#) ([ER\\_VARIABLE\\_IS\\_NOT\\_STRUCT](#) [\[3678\]](#))  
Message: Variable '%s' is not a variable component (can't be used as XXXX.variable\_name)
- Error: [1273](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_COLLATION](#) [\[3678\]](#))  
Message: Unknown collation: '%s'
- Error: [1274](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_IGNORED\\_SSL\\_PARAMS](#) [\[3678\]](#))  
Message: SSL parameters in CHANGE MASTER are ignored because this MySQL slave was compiled without SSL support; they can be used later if MySQL slave with SSL is started
- Error: [1275](#) SQLSTATE: [HY000](#) ([ER\\_SERVER\\_IS\\_IN\\_SECURE\\_AUTH\\_MODE](#) [\[3678\]](#))  
Message: Server is running in --secure-auth mode, but '%s'@'%s' has a password in the old format; please change the password to the new format
- Error: [1276](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_FIELD\\_RESOLVED](#) [\[3678\]](#))  
Message: Field or reference '%s%s%s%s%s' of SELECT #d was resolved in SELECT #d
- Error: [1277](#) SQLSTATE: [HY000](#) ([ER\\_BAD\\_SLAVE\\_UNTIL\\_COND](#) [\[3678\]](#))  
Message: Incorrect parameter or combination of parameters for START SLAVE UNTIL
- Error: [1278](#) SQLSTATE: [HY000](#) ([ER\\_MISSING\\_SKIP\\_SLAVE](#) [\[3678\]](#))  
Message: It is recommended to use --skip-slave-start when doing step-by-step replication with START SLAVE UNTIL; otherwise, you will get problems if you get an unexpected slave's mysqld restart
- Error: [1279](#) SQLSTATE: [HY000](#) ([ER\\_UNTIL\\_COND\\_IGNORED](#) [\[3678\]](#))  
Message: SQL thread is not to be started so UNTIL options are ignored
- Error: [1280](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_NAME\\_FOR\\_INDEX](#) [\[3678\]](#))  
Message: Incorrect index name '%s'
- Error: [1281](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_NAME\\_FOR\\_CATALOG](#) [\[3678\]](#))  
Message: Incorrect catalog name '%s'
- Error: [1282](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_QC\\_RESIZE](#) [\[3678\]](#))  
Message: Query cache failed to set size %lu; new query cache size is %lu

- Error: [1283](#) SQLSTATE: [HY000](#) ([ER\\_BAD\\_FT\\_COLUMN](#) [\[3679\]](#))  
Message: Column '%s' cannot be part of FULLTEXT index
- Error: [1284](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_KEY\\_CACHE](#) [\[3679\]](#))  
Message: Unknown key cache '%s'
- Error: [1285](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_HOSTNAME\\_WONT\\_WORK](#) [\[3679\]](#))  
Message: MySQL is started in --skip-name-resolve mode; you must restart it without this switch for this grant to work
- Error: [1286](#) SQLSTATE: [42000](#) ([ER\\_UNKNOWN\\_STORAGE\\_ENGINE](#) [\[3679\]](#))  
Message: Unknown storage engine '%s'
- Error: [1287](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_DEPRECATED\\_SYNTAX](#) [\[3679\]](#))  
Message: '%s' is deprecated and will be removed in a future release. Please use %s instead
- Error: [1288](#) SQLSTATE: [HY000](#) ([ER\\_NON\\_UPDATABLE\\_TABLE](#) [\[3679\]](#))  
Message: The target table %s of the %s is not updatable
- Error: [1289](#) SQLSTATE: [HY000](#) ([ER\\_FEATURE\\_DISABLED](#) [\[3679\]](#))  
Message: The '%s' feature is disabled; you need MySQL built with '%s' to have it working
- Error: [1290](#) SQLSTATE: [HY000](#) ([ER\\_OPTION\\_PREVENTS\\_STATEMENT](#) [\[3679\]](#))  
Message: The MySQL server is running with the %s option so it cannot execute this statement
- Error: [1291](#) SQLSTATE: [HY000](#) ([ER\\_DUPLICATED\\_VALUE\\_IN\\_TYPE](#) [\[3679\]](#))  
Message: Column '%s' has duplicated value '%s' in %s
- Error: [1292](#) SQLSTATE: [22007](#) ([ER\\_TRUNCATED\\_WRONG\\_VALUE](#) [\[3679\]](#))  
Message: Truncated incorrect %s value: '%s'
- Error: [1293](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_MUCH\\_AUTO\\_TIMESTAMP\\_COLS](#) [\[3679\]](#))  
Message: Incorrect table definition; there can be only one TIMESTAMP column with CURRENT\_TIMESTAMP in DEFAULT or ON UPDATE clause
- Error: [1294](#) SQLSTATE: [HY000](#) ([ER\\_INVALID\\_ON\\_UPDATE](#) [\[3679\]](#))  
Message: Invalid ON UPDATE clause for '%s' column
- Error: [1295](#) SQLSTATE: [HY000](#) ([ER\\_UNSUPPORTED\\_PS](#) [\[3679\]](#))  
Message: This command is not supported in the prepared statement protocol yet
- Error: [1296](#) SQLSTATE: [HY000](#) ([ER\\_GET\\_ERRMSG](#) [\[3679\]](#))  
Message: Got error %d '%s' from %s
- Error: [1297](#) SQLSTATE: [HY000](#) ([ER\\_GET\\_TEMPORARY\\_ERRMSG](#) [\[3679\]](#))  
Message: Got temporary error %d '%s' from %s
- Error: [1298](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_TIME\\_ZONE](#) [\[3679\]](#))  
Message: Unknown or incorrect time zone: '%s'

- Error: [1299](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_INVALID\\_TIMESTAMP](#) [\[3680\]](#))  
Message: Invalid `TIMESTAMP` value in column '%s' at row %ld
- Error: [1300](#) SQLSTATE: [HY000](#) ([ER\\_INVALID\\_CHARACTER\\_STRING](#) [\[3680\]](#))  
Message: Invalid %s character string: '%s'
- Error: [1301](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_ALLOWED\\_PACKET\\_OVERFLOWED](#) [\[3680\]](#))  
Message: Result of %s() was larger than `max_allowed_packet` (%ld) - truncated
- Error: [1302](#) SQLSTATE: [HY000](#) ([ER\\_CONFLICTING\\_DECLARATIONS](#) [\[3680\]](#))  
Message: Conflicting declarations: '%s%s' and '%s%s'
- Error: [1303](#) SQLSTATE: [2F003](#) ([ER\\_SP\\_NO\\_RECURSIVE\\_CREATE](#) [\[3680\]](#))  
Message: Can't create a %s from within another stored routine
- Error: [1304](#) SQLSTATE: [42000](#) ([ER\\_SP\\_ALREADY\\_EXISTS](#) [\[3680\]](#))  
Message: %s %s already exists
- Error: [1305](#) SQLSTATE: [42000](#) ([ER\\_SP\\_DOES\\_NOT\\_EXIST](#) [\[3680\]](#))  
Message: %s %s does not exist
- Error: [1306](#) SQLSTATE: [HY000](#) ([ER\\_SP\\_DROP\\_FAILED](#) [\[3680\]](#))  
Message: Failed to DROP %s %s
- Error: [1307](#) SQLSTATE: [HY000](#) ([ER\\_SP\\_STORE\\_FAILED](#) [\[3680\]](#))  
Message: Failed to CREATE %s %s
- Error: [1308](#) SQLSTATE: [42000](#) ([ER\\_SP\\_LILABEL\\_MISMATCH](#) [\[3680\]](#))  
Message: %s with no matching label: %s
- Error: [1309](#) SQLSTATE: [42000](#) ([ER\\_SP\\_LABEL\\_REDEFINE](#) [\[3680\]](#))  
Message: Redefining label %s
- Error: [1310](#) SQLSTATE: [42000](#) ([ER\\_SP\\_LABEL\\_MISMATCH](#) [\[3680\]](#))  
Message: End-label %s without match
- Error: [1311](#) SQLSTATE: [01000](#) ([ER\\_SP\\_UNINIT\\_VAR](#) [\[3680\]](#))  
Message: Referring to uninitialized variable %s
- Error: [1312](#) SQLSTATE: [0A000](#) ([ER\\_SP\\_BADSELECT](#) [\[3680\]](#))  
Message: PROCEDURE %s can't return a result set in the given context
- Error: [1313](#) SQLSTATE: [42000](#) ([ER\\_SP\\_BADRETURN](#) [\[3680\]](#))  
Message: RETURN is only allowed in a FUNCTION
- Error: [1314](#) SQLSTATE: [0A000](#) ([ER\\_SP\\_BADSTATEMENT](#) [\[3680\]](#))  
Message: %s is not allowed in stored procedures
- Error: [1315](#) SQLSTATE: [42000](#) ([ER\\_UPDATE\\_LOG\\_DEPRECATED\\_IGNORED](#) [\[3680\]](#))



Message: The update log is deprecated and replaced by the binary log; SET SQL\_LOG\_UPDATE has been ignored.

- Error: 1316 SQLSTATE: 42000 (ER\_UPDATE\_LOG\_DEPRECATED\_TRANSLATED [3681])

Message: The update log is deprecated and replaced by the binary log; SET SQL\_LOG\_UPDATE has been translated to SET SQL\_LOG\_BIN.

- Error: 1317 SQLSTATE: 70100 (ER\_QUERY\_INTERRUPTED [3681])

Message: Query execution was interrupted

- Error: 1318 SQLSTATE: 42000 (ER\_SP\_WRONG\_NO\_OF\_ARGS [3681])

Message: Incorrect number of arguments for %s %s; expected %u, got %u

- Error: 1319 SQLSTATE: 42000 (ER\_SP\_COND\_MISMATCH [3681])

Message: Undefined CONDITION: %s

- Error: 1320 SQLSTATE: 42000 (ER\_SP\_NORETURN [3681])

Message: No RETURN found in FUNCTION %s

- Error: 1321 SQLSTATE: 2F005 (ER\_SP\_NORETURNEND [3681])

Message: FUNCTION %s ended without RETURN

- Error: 1322 SQLSTATE: 42000 (ER\_SP\_BAD\_CURSOR\_QUERY [3681])

Message: Cursor statement must be a SELECT

- Error: 1323 SQLSTATE: 42000 (ER\_SP\_BAD\_CURSOR\_SELECT [3681])

Message: Cursor SELECT must not have INTO

- Error: 1324 SQLSTATE: 42000 (ER\_SP\_CURSOR\_MISMATCH [3681])

Message: Undefined CURSOR: %s

- Error: 1325 SQLSTATE: 24000 (ER\_SP\_CURSOR\_ALREADY\_OPEN [3681])

Message: Cursor is already open

- Error: 1326 SQLSTATE: 24000 (ER\_SP\_CURSOR\_NOT\_OPEN [3681])

Message: Cursor is not open

- Error: 1327 SQLSTATE: 42000 (ER\_SP\_UNDECLARED\_VAR [3681])

Message: Undeclared variable: %s

- Error: 1328 SQLSTATE: HY000 (ER\_SP\_WRONG\_NO\_OF\_FETCH\_ARGS [3681])

Message: Incorrect number of FETCH variables

- Error: 1329 SQLSTATE: 02000 (ER\_SP\_FETCH\_NO\_DATA [3681])

Message: No data - zero rows fetched, selected, or processed

- Error: 1330 SQLSTATE: 42000 (ER\_SP\_DUP\_PARAM [3681])

Message: Duplicate parameter: %s

- Error: 1331 SQLSTATE: 42000 (ER\_SP\_DUP\_VAR [3682])  
Message: Duplicate variable: %s
- Error: 1332 SQLSTATE: 42000 (ER\_SP\_DUP\_COND [3682])  
Message: Duplicate condition: %s
- Error: 1333 SQLSTATE: 42000 (ER\_SP\_DUP\_CURS [3682])  
Message: Duplicate cursor: %s
- Error: 1334 SQLSTATE: HY000 (ER\_SP\_CANT\_ALTER [3682])  
Message: Failed to ALTER %s %s
- Error: 1335 SQLSTATE: 0A000 (ER\_SP\_SUBSELECT\_NYI [3682])  
Message: Subquery value not supported
- Error: 1336 SQLSTATE: 0A000 (ER\_STMT\_NOT\_ALLOWED\_IN\_SF\_OR\_TRG [3682])  
Message: %s is not allowed in stored function or trigger
- Error: 1337 SQLSTATE: 42000 (ER\_SP\_VARCOND\_AFTER\_CURSHNDLR [3682])  
Message: Variable or condition declaration after cursor or handler declaration
- Error: 1338 SQLSTATE: 42000 (ER\_SP\_CURSOR\_AFTER\_HANDLER [3682])  
Message: Cursor declaration after handler declaration
- Error: 1339 SQLSTATE: 20000 (ER\_SP\_CASE\_NOT\_FOUND [3682])  
Message: Case not found for CASE statement
- Error: 1340 SQLSTATE: HY000 (ER\_FPARSER\_TOO\_BIG\_FILE [3682])  
Message: Configuration file '%s' is too big
- Error: 1341 SQLSTATE: HY000 (ER\_FPARSER\_BAD\_HEADER [3682])  
Message: Malformed file type header in file '%s'
- Error: 1342 SQLSTATE: HY000 (ER\_FPARSER\_EOF\_IN\_COMMENT [3682])  
Message: Unexpected end of file while parsing comment '%s'
- Error: 1343 SQLSTATE: HY000 (ER\_FPARSER\_ERROR\_IN\_PARAMETER [3682])  
Message: Error while parsing parameter '%s' (line: '%s')
- Error: 1344 SQLSTATE: HY000 (ER\_FPARSER\_EOF\_IN\_UNKNOWN\_PARAMETER [3682])  
Message: Unexpected end of file while skipping unknown parameter '%s'
- Error: 1345 SQLSTATE: HY000 (ER\_VIEW\_NO\_EXPLAIN [3682])  
Message: EXPLAIN/SHOW can not be issued; lacking privileges for underlying table
- Error: 1346 SQLSTATE: HY000 (ER\_FRM\_UNKNOWN\_TYPE [3682])  
Message: File '%s' has unknown type '%s' in its header
- Error: 1347 SQLSTATE: HY000 (ER\_WRONG\_OBJECT [3682])

Message: '%s.%s' is not %s

- Error: [1348](#) SQLSTATE: HY000 ([ER\\_NONUPDATEABLE\\_COLUMN](#) [\[3683\]](#))

Message: Column '%s' is not updatable

- Error: [1349](#) SQLSTATE: HY000 ([ER\\_VIEW\\_SELECT\\_DERIVED](#) [\[3683\]](#))

Message: View's SELECT contains a subquery in the FROM clause

- Error: [1350](#) SQLSTATE: HY000 ([ER\\_VIEW\\_SELECT\\_CLAUSE](#) [\[3683\]](#))

Message: View's SELECT contains a '%s' clause

- Error: [1351](#) SQLSTATE: HY000 ([ER\\_VIEW\\_SELECT\\_VARIABLE](#) [\[3683\]](#))

Message: View's SELECT contains a variable or parameter

- Error: [1352](#) SQLSTATE: HY000 ([ER\\_VIEW\\_SELECT\\_TMPTABLE](#) [\[3683\]](#))

Message: View's SELECT refers to a temporary table '%s'

- Error: [1353](#) SQLSTATE: HY000 ([ER\\_VIEW\\_WRONG\\_LIST](#) [\[3683\]](#))

Message: View's SELECT and view's field list have different column counts

- Error: [1354](#) SQLSTATE: HY000 ([ER\\_WARN\\_VIEW\\_MERGE](#) [\[3683\]](#))

Message: View merge algorithm can't be used here for now (assumed undefined algorithm)

- Error: [1355](#) SQLSTATE: HY000 ([ER\\_WARN\\_VIEW\\_WITHOUT\\_KEY](#) [\[3683\]](#))

Message: View being updated does not have complete key of underlying table in it

- Error: [1356](#) SQLSTATE: HY000 ([ER\\_VIEW\\_INVALID](#) [\[3683\]](#))

Message: View '%s.%s' references invalid table(s) or column(s) or function(s) or definer/invoker of view lack rights to use them

- Error: [1357](#) SQLSTATE: HY000 ([ER\\_SP\\_NO\\_DROP\\_SP](#) [\[3683\]](#))

Message: Can't drop or alter a %s from within another stored routine

- Error: [1358](#) SQLSTATE: HY000 ([ER\\_SP\\_GOTO\\_IN\\_HNDLR](#) [\[3683\]](#))

Message: GOTO is not allowed in a stored procedure handler

- Error: [1359](#) SQLSTATE: HY000 ([ER\\_TRG\\_ALREADY\\_EXISTS](#) [\[3683\]](#))

Message: Trigger already exists

- Error: [1360](#) SQLSTATE: HY000 ([ER\\_TRG\\_DOES\\_NOT\\_EXIST](#) [\[3683\]](#))

Message: Trigger does not exist

- Error: [1361](#) SQLSTATE: HY000 ([ER\\_TRG\\_ON\\_VIEW\\_OR\\_TEMP\\_TABLE](#) [\[3683\]](#))

Message: Trigger's '%s' is view or temporary table

- Error: [1362](#) SQLSTATE: HY000 ([ER\\_TRG\\_CANT\\_CHANGE\\_ROW](#) [\[3683\]](#))

Message: Updating of %s row is not allowed in %strigger

- Error: [1363](#) SQLSTATE: HY000 ([ER\\_TRG\\_NO\\_SUCH\\_ROW\\_IN\\_TRG](#) [\[3683\]](#))

Message: There is no %s row in %s trigger

- Error: 1364 SQLSTATE: HY000 (ER\_NO\_DEFAULT\_FOR\_FIELD [3684])

Message: Field '%s' doesn't have a default value

- Error: 1365 SQLSTATE: 22012 (ER\_DIVISION\_BY\_ZERO [3684])

Message: Division by 0

- Error: 1366 SQLSTATE: HY000 (ER\_TRUNCATED\_WRONG\_VALUE\_FOR\_FIELD [3684])

Message: Incorrect %s value: '%s' for column '%s' at row %ld

- Error: 1367 SQLSTATE: 22007 (ER\_ILLEGAL\_VALUE\_FOR\_TYPE [3684])

Message: Illegal %s '%s' value found during parsing

- Error: 1368 SQLSTATE: HY000 (ER\_VIEW\_NONUPD\_CHECK [3684])

Message: CHECK OPTION on non-updatable view '%s.%s'

- Error: 1369 SQLSTATE: HY000 (ER\_VIEW\_CHECK\_FAILED [3684])

Message: CHECK OPTION failed '%s.%s'

- Error: 1370 SQLSTATE: 42000 (ER\_PROCACCESS\_DENIED\_ERROR [3684])

Message: %s command denied to user '%s'@'%s' for routine '%s'

- Error: 1371 SQLSTATE: HY000 (ER\_RELAY\_LOG\_FAIL [3684])

Message: Failed purging old relay logs: %s

- Error: 1372 SQLSTATE: HY000 (ER\_PASSWD\_LENGTH [3684])

Message: Password hash should be a %d-digit hexadecimal number

- Error: 1373 SQLSTATE: HY000 (ER\_UNKNOWN\_TARGET\_BINLOG [3684])

Message: Target log not found in binlog index

- Error: 1374 SQLSTATE: HY000 (ER\_IO\_ERR\_LOG\_INDEX\_READ [3684])

Message: I/O error reading log index file

- Error: 1375 SQLSTATE: HY000 (ER\_BINLOG\_PURGE\_PROHIBITED [3684])

Message: Server configuration does not permit binlog purge

- Error: 1376 SQLSTATE: HY000 (ER\_FSEEK\_FAIL [3684])

Message: Failed on fseek()

- Error: 1377 SQLSTATE: HY000 (ER\_BINLOG\_PURGE\_FATAL\_ERR [3684])

Message: Fatal error during log purge

- Error: 1378 SQLSTATE: HY000 (ER\_LOG\_IN\_USE [3684])

Message: A purgeable log is in use, will not purge

- Error: 1379 SQLSTATE: HY000 (ER\_LOG\_PURGE\_UNKNOWN\_ERR [3684])

Message: Unknown error during log purge

- Error: [1380](#) SQLSTATE: HY000 ([ER\\_RELAY\\_LOG\\_INIT](#) [\[3685\]](#))

Message: Failed initializing relay log position: %s

- Error: [1381](#) SQLSTATE: HY000 ([ER\\_NO\\_BINARY\\_LOGGING](#) [\[3685\]](#))

Message: You are not using binary logging

- Error: [1382](#) SQLSTATE: HY000 ([ER\\_RESERVED\\_SYNTAX](#) [\[3685\]](#))

Message: The '%s' syntax is reserved for purposes internal to the MySQL server

- Error: [1383](#) SQLSTATE: HY000 ([ER\\_WSAS\\_FAILED](#) [\[3685\]](#))

Message: WSASStartup Failed

- Error: [1384](#) SQLSTATE: HY000 ([ER\\_DIFF\\_GROUPS\\_PROC](#) [\[3685\]](#))

Message: Can't handle procedures with different groups yet

- Error: [1385](#) SQLSTATE: HY000 ([ER\\_NO\\_GROUP\\_FOR\\_PROC](#) [\[3685\]](#))

Message: Select must have a group with this procedure

- Error: [1386](#) SQLSTATE: HY000 ([ER\\_ORDER\\_WITH\\_PROC](#) [\[3685\]](#))

Message: Can't use ORDER clause with this procedure

- Error: [1387](#) SQLSTATE: HY000 ([ER\\_LOGGING\\_PROHIBIT\\_CHANGING\\_OF](#) [\[3685\]](#))

Message: Binary logging and replication forbid changing the global server %s

- Error: [1388](#) SQLSTATE: HY000 ([ER\\_NO\\_FILE\\_MAPPING](#) [\[3685\]](#))

Message: Can't map file: %s, errno: %d

- Error: [1389](#) SQLSTATE: HY000 ([ER\\_WRONG\\_MAGIC](#) [\[3685\]](#))

Message: Wrong magic in %s

- Error: [1390](#) SQLSTATE: HY000 ([ER\\_PS\\_MANY\\_PARAM](#) [\[3685\]](#))

Message: Prepared statement contains too many placeholders

- Error: [1391](#) SQLSTATE: HY000 ([ER\\_KEY\\_PART\\_0](#) [\[3685\]](#))

Message: Key part '%s' length cannot be 0

- Error: [1392](#) SQLSTATE: HY000 ([ER\\_VIEW\\_CHECKSUM](#) [\[3685\]](#))

Message: View text checksum failed

- Error: [1393](#) SQLSTATE: HY000 ([ER\\_VIEW\\_MULTIUPDATE](#) [\[3685\]](#))

Message: Can not modify more than one base table through a join view '%s.%s'

- Error: [1394](#) SQLSTATE: HY000 ([ER\\_VIEW\\_NO\\_INSERT\\_FIELD\\_LIST](#) [\[3685\]](#))

Message: Can not insert into join view '%s.%s' without fields list

- Error: [1395](#) SQLSTATE: HY000 ([ER\\_VIEW\\_DELETE\\_MERGE\\_VIEW](#) [\[3685\]](#))

Message: Can not delete from join view '%s.%s'

- Error: 1396 SQLSTATE: HY000 (ER\_CANNOT\_USER [3686])

Message: Operation %s failed for %s

- Error: 1397 SQLSTATE: XAE04 (ER\_XAER\_NOTA [3686])

Message: XAER\_NOTA: Unknown XID

- Error: 1398 SQLSTATE: XAE05 (ER\_XAER\_INVAL [3686])

Message: XAER\_INVAL: Invalid arguments (or unsupported command)

- Error: 1399 SQLSTATE: XAE07 (ER\_XAER\_RMFAIL [3686])

Message: XAER\_RMFAIL: The command cannot be executed when global transaction is in the %s state

- Error: 1400 SQLSTATE: XAE09 (ER\_XAER\_OUTSIDE [3686])

Message: XAER\_OUTSIDE: Some work is done outside global transaction

- Error: 1401 SQLSTATE: XAE03 (ER\_XAER\_RMERR [3686])

Message: XAER\_RMERR: Fatal error occurred in the transaction branch - check your data for consistency

- Error: 1402 SQLSTATE: XA100 (ER\_XA\_RBROLLBACK [3686])

Message: XA\_RBROLLBACK: Transaction branch was rolled back

- Error: 1403 SQLSTATE: 42000 (ER\_NONEXISTING\_PROC\_GRANT [3686])

Message: There is no such grant defined for user '%s' on host '%s' on routine '%s'

- Error: 1404 SQLSTATE: HY000 (ER\_PROC\_AUTO\_GRANT\_FAIL [3686])

Message: Failed to grant EXECUTE and ALTER ROUTINE privileges

- Error: 1405 SQLSTATE: HY000 (ER\_PROC\_AUTO\_REVOKE\_FAIL [3686])

Message: Failed to revoke all privileges to dropped routine

- Error: 1406 SQLSTATE: 22001 (ER\_DATA\_TOO\_LONG [3686])

Message: Data too long for column '%s' at row %ld

- Error: 1407 SQLSTATE: 42000 (ER\_SP\_BAD\_SQLSTATE [3686])

Message: Bad SQLSTATE: '%s'

- Error: 1408 SQLSTATE: HY000 (ER\_STARTUP [3686])

Message: %s: ready for connections. Version: '%s' socket: '%s' port: %d %s

- Error: 1409 SQLSTATE: HY000 (ER\_LOAD\_FROM\_FIXED\_SIZE\_ROWS\_TO\_VAR [3686])

Message: Can't load value from file with fixed size rows to variable

- Error: 1410 SQLSTATE: 42000 (ER\_CANT\_CREATE\_USER\_WITH\_GRANT [3686])

Message: You are not allowed to create a user with GRANT

- Error: [1411](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_VALUE\\_FOR\\_TYPE](#) [\[3687\]](#))  
Message: Incorrect %s value: '%s' for function %s
- Error: [1412](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_DEF\\_CHANGED](#) [\[3687\]](#))  
Message: Table definition has changed, please retry transaction
- Error: [1413](#) SQLSTATE: [42000](#) ([ER\\_SP\\_DUP\\_HANDLER](#) [\[3687\]](#))  
Message: Duplicate handler declared in the same block
- Error: [1414](#) SQLSTATE: [42000](#) ([ER\\_SP\\_NOT\\_VAR\\_ARG](#) [\[3687\]](#))  
Message: OUT or INOUT argument %d for routine %s is not a variable or NEW pseudo-variable in BEFORE trigger
- Error: [1415](#) SQLSTATE: [0A000](#) ([ER\\_SP\\_NO\\_RESET](#) [\[3687\]](#))  
Message: Not allowed to return a result set from a %s
- Error: [1416](#) SQLSTATE: [22003](#) ([ER\\_CANT\\_CREATE\\_GEOMETRY\\_OBJECT](#) [\[3687\]](#))  
Message: Cannot get geometry object from data you send to the GEOMETRY field
- Error: [1417](#) SQLSTATE: [HY000](#) ([ER\\_FAILED\\_ROUTINE\\_BREAK\\_BINLOG](#) [\[3687\]](#))  
Message: A routine failed and has neither NO SQL nor READS SQL DATA in its declaration and binary logging is enabled; if non-transactional tables were updated, the binary log will miss their changes
- Error: [1418](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_ROUTINE](#) [\[3687\]](#))  
Message: This function has none of DETERMINISTIC, NO SQL, or READS SQL DATA in its declaration and binary logging is enabled (you \*might\* want to use the less safe log\_bin\_trust\_function\_creators variable)
- Error: [1419](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_CREATE\\_ROUTINE\\_NEED\\_SUPER](#) [\[3687\]](#))  
Message: You do not have the SUPER privilege and binary logging is enabled (you \*might\* want to use the less safe log\_bin\_trust\_function\_creators variable)
- Error: [1420](#) SQLSTATE: [HY000](#) ([ER\\_EXEC\\_STMT\\_WITH\\_OPEN\\_CURSOR](#) [\[3687\]](#))  
Message: You can't execute a prepared statement which has an open cursor associated with it. Reset the statement to re-execute it.
- Error: [1421](#) SQLSTATE: [HY000](#) ([ER\\_STMT\\_HAS\\_NO\\_OPEN\\_CURSOR](#) [\[3687\]](#))  
Message: The statement (%lu) has no open cursor.
- Error: [1422](#) SQLSTATE: [HY000](#) ([ER\\_COMMIT\\_NOT\\_ALLOWED\\_IN\\_SF\\_OR\\_TRG](#) [\[3687\]](#))  
Message: Explicit or implicit commit is not allowed in stored function or trigger.
- Error: [1423](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_DEFAULT\\_FOR\\_VIEW\\_FIELD](#) [\[3687\]](#))  
Message: Field of view '%s.%s' underlying table doesn't have a default value
- Error: [1424](#) SQLSTATE: [HY000](#) ([ER\\_SP\\_NO\\_RECURSION](#) [\[3687\]](#))  
Message: Recursive stored functions and triggers are not allowed.
- Error: [1425](#) SQLSTATE: [42000](#) ([ER\\_TOO\\_BIG\\_SCALE](#) [\[3687\]](#))

Message: Too big scale %d specified for column '%s'. Maximum is %lu.

- Error: 1426 SQLSTATE: 42000 (ER\_TOO\_BIG\_PRECISION [3688])

Message: Too big precision %d specified for column '%s'. Maximum is %lu.

- Error: 1427 SQLSTATE: 42000 (ER\_M\_BIGGER\_THAN\_D [3688])

Message: For float(M,D), double(M,D) or decimal(M,D), M must be >= D (column '%s').

- Error: 1428 SQLSTATE: HY000 (ER\_WRONG\_LOCK\_OF\_SYSTEM\_TABLE [3688])

Message: You can't combine write-locking of system tables with other tables or lock types

- Error: 1429 SQLSTATE: HY000 (ER\_CONNECT\_TO\_FOREIGN\_DATA\_SOURCE [3688])

Message: Unable to connect to foreign data source: %s

- Error: 1430 SQLSTATE: HY000 (ER\_QUERY\_ON\_FOREIGN\_DATA\_SOURCE [3688])

Message: There was a problem processing the query on the foreign data source. Data source error: %s

- Error: 1431 SQLSTATE: HY000 (ER\_FOREIGN\_DATA\_SOURCE\_DOESNT\_EXIST [3688])

Message: The foreign data source you are trying to reference does not exist. Data source error: %s

- Error: 1432 SQLSTATE: HY000 (ER\_FOREIGN\_DATA\_STRING\_INVALID\_CANT\_CREATE [3688])

Message: Can't create federated table. The data source connection string '%s' is not in the correct format

- Error: 1433 SQLSTATE: HY000 (ER\_FOREIGN\_DATA\_STRING\_INVALID [3688])

Message: The data source connection string '%s' is not in the correct format

- Error: 1434 SQLSTATE: HY000 (ER\_CANT\_CREATE\_FEDERATED\_TABLE [3688])

Message: Can't create federated table. Foreign data src error: %s

- Error: 1435 SQLSTATE: HY000 (ER\_TRG\_IN\_WRONG\_SCHEMA [3688])

Message: Trigger in wrong schema

- Error: 1436 SQLSTATE: HY000 (ER\_STACK\_OVERRUN\_NEED\_MORE [3688])

Message: Thread stack overrun: %ld bytes used of a %ld byte stack, and %ld bytes needed. Use 'mysqld --thread\_stack=#' to specify a bigger stack.

- Error: 1437 SQLSTATE: 42000 (ER\_TOO\_LONG\_BODY [3688])

Message: Routine body for '%s' is too long

- Error: 1438 SQLSTATE: HY000 (ER\_WARN\_CANT\_DROP\_DEFAULT\_KEYCACHE [3688])

Message: Cannot drop default keycache

- Error: 1439 SQLSTATE: 42000 (ER\_TOO\_BIG\_DISPLAYWIDTH [3688])

Message: Display width out of range for column '%s' (max = %lu)

- Error: 1440 SQLSTATE: XAE08 (ER\_XAER\_DUPID [3688])

Message: XAER\_DUPID: The XID already exists



- Error: [1441](#) SQLSTATE: [22008](#) ([ER\\_DATETIME\\_FUNCTION\\_OVERFLOW](#) [\[3689\]](#))  
Message: Datetime function: %s field overflow
- Error: [1442](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_UPDATE\\_USED\\_TABLE\\_IN\\_SF\\_OR\\_TRG](#) [\[3689\]](#))  
Message: Can't update table '%s' in stored function/trigger because it is already used by statement which invoked this stored function/trigger.
- Error: [1443](#) SQLSTATE: [HY000](#) ([ER\\_VIEW\\_PREVENT\\_UPDATE](#) [\[3689\]](#))  
Message: The definition of table '%s' prevents operation %s on table '%s'.
- Error: [1444](#) SQLSTATE: [HY000](#) ([ER\\_PS\\_NO\\_RECURSION](#) [\[3689\]](#))  
Message: The prepared statement contains a stored routine call that refers to that same statement. It's not allowed to execute a prepared statement in such a recursive manner
- Error: [1445](#) SQLSTATE: [HY000](#) ([ER\\_SP\\_CANT\\_SET\\_AUTOCOMMIT](#) [\[3689\]](#))  
Message: Not allowed to set autocommit from a stored function or trigger
- Error: [1446](#) SQLSTATE: [HY000](#) ([ER\\_MALFORMED\\_DEFINER](#) [\[3689\]](#))  
Message: Definer is not fully qualified
- Error: [1447](#) SQLSTATE: [HY000](#) ([ER\\_VIEW\\_FRM\\_NO\\_USER](#) [\[3689\]](#))  
Message: View '%s'.'%s' has no definer information (old table format). Current user is used as definer. Please recreate the view!
- Error: [1448](#) SQLSTATE: [HY000](#) ([ER\\_VIEW\\_OTHER\\_USER](#) [\[3689\]](#))  
Message: You need the SUPER privilege for creation view with '%s'@'%s' definer
- Error: [1449](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_SUCH\\_USER](#) [\[3689\]](#))  
Message: The user specified as a definer ('%s'@'%s') does not exist
- Error: [1450](#) SQLSTATE: [HY000](#) ([ER\\_FORBID\\_SCHEMA\\_CHANGE](#) [\[3689\]](#))  
Message: Changing schema from '%s' to '%s' is not allowed.
- Error: [1451](#) SQLSTATE: [23000](#) ([ER\\_ROW\\_IS\\_REFERENCED\\_2](#) [\[3689\]](#))  
Message: Cannot delete or update a parent row: a foreign key constraint fails (%s)
- Error: [1452](#) SQLSTATE: [23000](#) ([ER\\_NO\\_REFERENCED\\_ROW\\_2](#) [\[3689\]](#))  
Message: Cannot add or update a child row: a foreign key constraint fails (%s)
- Error: [1453](#) SQLSTATE: [42000](#) ([ER\\_SP\\_BAD\\_VAR\\_SHADOW](#) [\[3689\]](#))  
Message: Variable '%s' must be quoted with `...`, or renamed
- Error: [1454](#) SQLSTATE: [HY000](#) ([ER\\_TRG\\_NO\\_DEFINER](#) [\[3689\]](#))  
Message: No definer attribute for trigger '%s'.'%s'. The trigger will be activated under the authorization of the invoker, which may have insufficient privileges. Please recreate the trigger.
- Error: [1455](#) SQLSTATE: [HY000](#) ([ER\\_OLD\\_FILE\\_FORMAT](#) [\[3689\]](#))  
Message: '%s' has an old format, you should re-create the '%s' object(s)

- Error: 1456 SQLSTATE: HY000 (ER\_SP\_RECURSION\_LIMIT [3690])  
Message: Recursive limit %d (as set by the max\_sp\_recursion\_depth variable) was exceeded for routine %s
- Error: 1457 SQLSTATE: HY000 (ER\_SP\_PROC\_TABLE\_CORRUPT [3690])  
Message: Failed to load routine %s. The table mysql.proc is missing, corrupt, or contains bad data (internal code %d)
- Error: 1458 SQLSTATE: 42000 (ER\_SP\_WRONG\_NAME [3690])  
Message: Incorrect routine name '%s'
- Error: 1459 SQLSTATE: HY000 (ER\_TABLE\_NEEDS\_UPGRADE [3690])  
Message: Table upgrade required. Please do "REPAIR TABLE `%s`" or dump/reload to fix it!
- Error: 1460 SQLSTATE: 42000 (ER\_SP\_NO\_AGGREGATE [3690])  
Message: AGGREGATE is not supported for stored functions
- Error: 1461 SQLSTATE: 42000 (ER\_MAX\_PREPARED\_STMT\_COUNT\_REACHED [3690])  
Message: Can't create more than max\_prepared\_stmt\_count statements (current value: %lu)
- Error: 1462 SQLSTATE: HY000 (ER\_VIEW\_RECURSIVE [3690])  
Message: `%s`.`%s` contains view recursion
- Error: 1463 SQLSTATE: 42000 (ER\_NON\_GROUPING\_FIELD\_USED [3690])  
Message: Non-grouping field '%s' is used in %s clause
- Error: 1464 SQLSTATE: HY000 (ER\_TABLE\_CANT\_HANDLE\_SPKEYS [3690])  
Message: The used table type doesn't support SPATIAL indexes
- Error: 1465 SQLSTATE: HY000 (ER\_NO\_TRIGGERS\_ON\_SYSTEM\_SCHEMA [3690])  
Message: Triggers can not be created on system tables
- Error: 1466 SQLSTATE: HY000 (ER\_REMOVED\_SPACES [3690])  
Message: Leading spaces are removed from name '%s'
- Error: 1467 SQLSTATE: HY000 (ER\_AUTOINC\_READ\_FAILED [3690])  
Message: Failed to read auto-increment value from storage engine
- Error: 1468 SQLSTATE: HY000 (ER\_USERNAME [3690])  
Message: user name
- Error: 1469 SQLSTATE: HY000 (ER\_HOSTNAME [3690])  
Message: host name
- Error: 1470 SQLSTATE: HY000 (ER\_WRONG\_STRING\_LENGTH [3690])  
Message: String '%s' is too long for %s (should be no longer than %d)
- Error: 1471 SQLSTATE: HY000 (ER\_NON\_INSERTABLE\_TABLE [3690])  
Message: The target table %s of the %s is not insertable-into

- Error: [1472](#) SQLSTATE: HY000 ([ER\\_ADMIN\\_WRONG\\_MRG\\_TABLE](#) [\[3691\]](#))  
Message: Table '%s' is differently defined or of non-MyISAM type or doesn't exist
- Error: [1473](#) SQLSTATE: HY000 ([ER\\_TOO\\_HIGH\\_LEVEL\\_OF\\_NESTING\\_FOR\\_SELECT](#) [\[3691\]](#))  
Message: Too high level of nesting for select
- Error: [1474](#) SQLSTATE: HY000 ([ER\\_NAME\\_BECOMES\\_EMPTY](#) [\[3691\]](#))  
Message: Name '%s' has become "
- Error: [1475](#) SQLSTATE: HY000 ([ER\\_AMBIGUOUS\\_FIELD\\_TERM](#) [\[3691\]](#))  
Message: First character of the FIELDS TERMINATED string is ambiguous; please use non-optional and non-empty FIELDS ENCLOSED BY
- Error: [1476](#) SQLSTATE: HY000 ([ER\\_FOREIGN\\_SERVER\\_EXISTS](#) [\[3691\]](#))  
Message: The foreign server, %s, you are trying to create already exists.
- Error: [1477](#) SQLSTATE: HY000 ([ER\\_FOREIGN\\_SERVER\\_DOESNT\\_EXIST](#) [\[3691\]](#))  
Message: The foreign server name you are trying to reference does not exist. Data source error: %s
- Error: [1478](#) SQLSTATE: HY000 ([ER\\_ILLEGAL\\_HA\\_CREATE\\_OPTION](#) [\[3691\]](#))  
Message: Table storage engine '%s' does not support the create option '%s'
- Error: [1479](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_REQUIRES\\_VALUES\\_ERROR](#) [\[3691\]](#))  
Message: Syntax error: %s PARTITIONING requires definition of VALUES %s for each partition
- Error: [1480](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_WRONG\\_VALUES\\_ERROR](#) [\[3691\]](#))  
Message: Only %s PARTITIONING can use VALUES %s in partition definition
- Error: [1481](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_MAXVALUE\\_ERROR](#) [\[3691\]](#))  
Message: MAXVALUE can only be used in last partition definition
- Error: [1482](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_SUBPARTITION\\_ERROR](#) [\[3691\]](#))  
Message: Subpartitions can only be hash partitions and by key
- Error: [1483](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_SUBPART\\_MIX\\_ERROR](#) [\[3691\]](#))  
Message: Must define subpartitions on all partitions if on one partition
- Error: [1484](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_WRONG\\_NO\\_PART\\_ERROR](#) [\[3691\]](#))  
Message: Wrong number of partitions defined, mismatch with previous setting
- Error: [1485](#) SQLSTATE: HY000 ([ER\\_PARTITION\\_WRONG\\_NO\\_SUBPART\\_ERROR](#) [\[3691\]](#))  
Message: Wrong number of subpartitions defined, mismatch with previous setting
- Error: [1486](#) SQLSTATE: HY000 ([ER\\_WRONG\\_EXPR\\_IN\\_PARTITION\\_FUNC\\_ERROR](#) [\[3691\]](#))  
Message: Constant, random or timezone-dependent expressions in (sub)partitioning function are not allowed
- Error: [1487](#) SQLSTATE: HY000 ([ER\\_NO\\_CONST\\_EXPR\\_IN\\_RANGE\\_OR\\_LIST\\_ERROR](#) [\[3691\]](#))  
Message: Expression in RANGE/LIST VALUES must be constant

- Error: 1488 SQLSTATE: HY000 (ER\_FIELD\_NOT\_FOUND\_PART\_ERROR [3692])  
Message: Field in list of fields for partition function not found in table
- Error: 1489 SQLSTATE: HY000 (ER\_LIST\_OF\_FIELDS\_ONLY\_IN\_HASH\_ERROR [3692])  
Message: List of fields is only allowed in KEY partitions
- Error: 1490 SQLSTATE: HY000 (ER\_INCONSISTENT\_PARTITION\_INFO\_ERROR [3692])  
Message: The partition info in the frm file is not consistent with what can be written into the frm file
- Error: 1491 SQLSTATE: HY000 (ER\_PARTITION\_FUNC\_NOT\_ALLOWED\_ERROR [3692])  
Message: The %s function returns the wrong type
- Error: 1492 SQLSTATE: HY000 (ER\_PARTITIONS\_MUST\_BE\_DEFINED\_ERROR [3692])  
Message: For %s partitions each partition must be defined
- Error: 1493 SQLSTATE: HY000 (ER\_RANGE\_NOT\_INCREASING\_ERROR [3692])  
Message: VALUES LESS THAN value must be strictly increasing for each partition
- Error: 1494 SQLSTATE: HY000 (ER\_INCONSISTENT\_TYPE\_OF\_FUNCTIONS\_ERROR [3692])  
Message: VALUES value must be of same type as partition function
- Error: 1495 SQLSTATE: HY000 (ER\_MULTIPLE\_DEF\_CONST\_IN\_LIST\_PART\_ERROR [3692])  
Message: Multiple definition of same constant in list partitioning
- Error: 1496 SQLSTATE: HY000 (ER\_PARTITION\_ENTRY\_ERROR [3692])  
Message: Partitioning can not be used stand-alone in query
- Error: 1497 SQLSTATE: HY000 (ER\_MIX\_HANDLER\_ERROR [3692])  
Message: The mix of handlers in the partitions is not allowed in this version of MySQL
- Error: 1498 SQLSTATE: HY000 (ER\_PARTITION\_NOT\_DEFINED\_ERROR [3692])  
Message: For the partitioned engine it is necessary to define all %s
- Error: 1499 SQLSTATE: HY000 (ER\_TOO\_MANY\_PARTITIONS\_ERROR [3692])  
Message: Too many partitions (including subpartitions) were defined
- Error: 1500 SQLSTATE: HY000 (ER\_SUBPARTITION\_ERROR [3692])  
Message: It is only possible to mix RANGE/LIST partitioning with HASH/KEY partitioning for subpartitioning
- Error: 1501 SQLSTATE: HY000 (ER\_CANT\_CREATE\_HANDLER\_FILE [3692])  
Message: Failed to create specific handler file
- Error: 1502 SQLSTATE: HY000 (ER\_BLOB\_FIELD\_IN\_PART\_FUNC\_ERROR [3692])  
Message: A BLOB field is not allowed in partition function
- Error: 1503 SQLSTATE: HY000 (ER\_UNIQUE\_KEY\_NEED\_ALL\_FIELDS\_IN\_PF [3692])  
Message: A %s must include all columns in the table's partitioning function

- Error: [1504](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_PARTS\\_ERROR](#) [\[3693\]](#))  
Message: Number of %s = 0 is not an allowed value
- Error: [1505](#) SQLSTATE: [HY000](#) ([ER\\_PARTITION\\_MGMT\\_ON\\_NONPARTITIONED](#) [\[3693\]](#))  
Message: Partition management on a not partitioned table is not possible
- Error: [1506](#) SQLSTATE: [HY000](#) ([ER\\_FOREIGN\\_KEY\\_ON\\_PARTITIONED](#) [\[3693\]](#))  
Message: Foreign key clause is not yet supported in conjunction with partitioning
- Error: [1507](#) SQLSTATE: [HY000](#) ([ER\\_DROP\\_PARTITION\\_NON\\_EXISTENT](#) [\[3693\]](#))  
Message: Error in list of partitions to %s
- Error: [1508](#) SQLSTATE: [HY000](#) ([ER\\_DROP\\_LAST\\_PARTITION](#) [\[3693\]](#))  
Message: Cannot remove all partitions, use DROP TABLE instead
- Error: [1509](#) SQLSTATE: [HY000](#) ([ER\\_COALESCE\\_ONLY\\_ON\\_HASH\\_PARTITION](#) [\[3693\]](#))  
Message: COALESCE PARTITION can only be used on HASH/KEY partitions
- Error: [1510](#) SQLSTATE: [HY000](#) ([ER\\_REORG\\_HASH\\_ONLY\\_ON\\_SAME\\_NO](#) [\[3693\]](#))  
Message: REORGANIZE PARTITION can only be used to reorganize partitions not to change their numbers
- Error: [1511](#) SQLSTATE: [HY000](#) ([ER\\_REORG\\_NO\\_PARAM\\_ERROR](#) [\[3693\]](#))  
Message: REORGANIZE PARTITION without parameters can only be used on auto-partitioned tables using HASH PARTITIONS
- Error: [1512](#) SQLSTATE: [HY000](#) ([ER\\_ONLY\\_ON\\_RANGE\\_LIST\\_PARTITION](#) [\[3693\]](#))  
Message: %s PARTITION can only be used on RANGE/LIST partitions
- Error: [1513](#) SQLSTATE: [HY000](#) ([ER\\_ADD\\_PARTITION\\_SUBPART\\_ERROR](#) [\[3693\]](#))  
Message: Trying to Add partition(s) with wrong number of subpartitions
- Error: [1514](#) SQLSTATE: [HY000](#) ([ER\\_ADD\\_PARTITION\\_NO\\_NEW\\_PARTITION](#) [\[3693\]](#))  
Message: At least one partition must be added
- Error: [1515](#) SQLSTATE: [HY000](#) ([ER\\_COALESCE\\_PARTITION\\_NO\\_PARTITION](#) [\[3693\]](#))  
Message: At least one partition must be coalesced
- Error: [1516](#) SQLSTATE: [HY000](#) ([ER\\_REORG\\_PARTITION\\_NOT\\_EXIST](#) [\[3693\]](#))  
Message: More partitions to reorganize than there are partitions
- Error: [1517](#) SQLSTATE: [HY000](#) ([ER\\_SAME\\_NAME\\_PARTITION](#) [\[3693\]](#))  
Message: Duplicate partition name %s
- Error: [1518](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_BINLOG\\_ERROR](#) [\[3693\]](#))  
Message: It is not allowed to shut off binlog on this command
- Error: [1519](#) SQLSTATE: [HY000](#) ([ER\\_CONSECUTIVE\\_REORG\\_PARTITIONS](#) [\[3693\]](#))  
Message: When reorganizing a set of partitions they must be in consecutive order

- Error: [1520](#) SQLSTATE: [HY000](#) ([ER\\_REORG\\_OUTSIDE\\_RANGE](#) [\[3694\]](#))  
Message: Reorganize of range partitions cannot change total ranges except for last partition where it can extend the range
- Error: [1521](#) SQLSTATE: [HY000](#) ([ER\\_PARTITION\\_FUNCTION\\_FAILURE](#) [\[3694\]](#))  
Message: Partition function not supported in this version for this handler
- Error: [1522](#) SQLSTATE: [HY000](#) ([ER\\_PART\\_STATE\\_ERROR](#) [\[3694\]](#))  
Message: Partition state cannot be defined from CREATE/ALTER TABLE
- Error: [1523](#) SQLSTATE: [HY000](#) ([ER\\_LIMITED\\_PART\\_RANGE](#) [\[3694\]](#))  
Message: The %s handler only supports 32 bit integers in VALUES
- Error: [1524](#) SQLSTATE: [HY000](#) ([ER\\_PLUGIN\\_IS\\_NOT\\_LOADED](#) [\[3694\]](#))  
Message: Plugin '%s' is not loaded
- Error: [1525](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_VALUE](#) [\[3694\]](#))  
Message: Incorrect %s value: '%s'
- Error: [1526](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_PARTITION\\_FOR\\_GIVEN\\_VALUE](#) [\[3694\]](#))  
Message: Table has no partition for value %s
- Error: [1527](#) SQLSTATE: [HY000](#) ([ER\\_FILEGROUP\\_OPTION\\_ONLY\\_ONCE](#) [\[3694\]](#))  
Message: It is not allowed to specify %s more than once
- Error: [1528](#) SQLSTATE: [HY000](#) ([ER\\_CREATE\\_FILEGROUP\\_FAILED](#) [\[3694\]](#))  
Message: Failed to create %s
- Error: [1529](#) SQLSTATE: [HY000](#) ([ER\\_DROP\\_FILEGROUP\\_FAILED](#) [\[3694\]](#))  
Message: Failed to drop %s
- Error: [1530](#) SQLSTATE: [HY000](#) ([ER\\_TABLESPACE\\_AUTO\\_EXTEND\\_ERROR](#) [\[3694\]](#))  
Message: The handler doesn't support autoextend of tablespaces
- Error: [1531](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_SIZE\\_NUMBER](#) [\[3694\]](#))  
Message: A size parameter was incorrectly specified, either number or on the form 10M
- Error: [1532](#) SQLSTATE: [HY000](#) ([ER\\_SIZE\\_OVERFLOW\\_ERROR](#) [\[3694\]](#))  
Message: The size number was correct but we don't allow the digit part to be more than 2 billion
- Error: [1533](#) SQLSTATE: [HY000](#) ([ER\\_ALTER\\_FILEGROUP\\_FAILED](#) [\[3694\]](#))  
Message: Failed to alter: %s
- Error: [1534](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_ROW\\_LOGGING\\_FAILED](#) [\[3694\]](#))  
Message: Writing one row to the row-based binary log failed
- Error: [1535](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_ROW\\_WRONG\\_TABLE\\_DEF](#) [\[3694\]](#))  
Message: Table definition on master and slave does not match: %s

- Error: [1536](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_ROW\\_RBR\\_TO\\_SBR](#) [\[3695\]](#))  
Message: Slave running with --log-slave-updates must use row-based binary logging to be able to replicate row-based binary log events
- Error: [1537](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_ALREADY\\_EXISTS](#) [\[3695\]](#))  
Message: Event '%s' already exists
- Error: [1538](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_STORE\\_FAILED](#) [\[3695\]](#))  
Message: Failed to store event %s. Error code %d from storage engine.
- Error: [1539](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_DOES\\_NOT\\_EXIST](#) [\[3695\]](#))  
Message: Unknown event '%s'
- Error: [1540](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_CANT\\_ALTER](#) [\[3695\]](#))  
Message: Failed to alter event '%s'
- Error: [1541](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_DROP\\_FAILED](#) [\[3695\]](#))  
Message: Failed to drop %s
- Error: [1542](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_INTERVAL\\_NOT\\_POSITIVE\\_OR\\_TOO\\_BIG](#) [\[3695\]](#))  
Message: INTERVAL is either not positive or too big
- Error: [1543](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_ENDS\\_BEFORE\\_STARTS](#) [\[3695\]](#))  
Message: ENDS is either invalid or before STARTS
- Error: [1544](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_EXEC\\_TIME\\_IN\\_THE\\_PAST](#) [\[3695\]](#))  
Message: Event execution time is in the past. Event has been disabled
- Error: [1545](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_OPEN\\_TABLE\\_FAILED](#) [\[3695\]](#))  
Message: Failed to open mysql.event
- Error: [1546](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_NEITHER\\_M\\_EXPR\\_NOR\\_M\\_AT](#) [\[3695\]](#))  
Message: No datetime expression provided
- Error: [1547](#) SQLSTATE: [HY000](#) ([ER\\_OBSOLETE\\_COL\\_COUNT\\_DOESNT\\_MATCH\\_CORRUPTED](#) [\[3695\]](#))  
Message: Column count of mysql.%s is wrong. Expected %d, found %d. The table is probably corrupted
- Error: [1548](#) SQLSTATE: [HY000](#) ([ER\\_OBSOLETE\\_CANNOT\\_LOAD\\_FROM\\_TABLE](#) [\[3695\]](#))  
Message: Cannot load from mysql.%s. The table is probably corrupted
- Error: [1549](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_CANNOT\\_DELETE](#) [\[3695\]](#))  
Message: Failed to delete the event from mysql.event
- Error: [1550](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_COMPILE\\_ERROR](#) [\[3695\]](#))  
Message: Error during compilation of event's body
- Error: [1551](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_SAME\\_NAME](#) [\[3695\]](#))

Message: Same old and new event name

- Error: 1552 SQLSTATE: HY000 (ER\_EVENT\_DATA\_TOO\_LONG [3696])

Message: Data for column '%s' too long

- Error: 1553 SQLSTATE: HY000 (ER\_DROP\_INDEX\_FK [3696])

Message: Cannot drop index '%s': needed in a foreign key constraint

- Error: 1554 SQLSTATE: HY000 (ER\_WARN\_DEPRECATED\_SYNTAX\_WITH\_VER [3696])

Message: The syntax '%s' is deprecated and will be removed in MySQL %s. Please use %s instead

- Error: 1555 SQLSTATE: HY000 (ER\_CANT\_WRITE\_LOCK\_LOG\_TABLE [3696])

Message: You can't write-lock a log table. Only read access is possible

- Error: 1556 SQLSTATE: HY000 (ER\_CANT\_LOCK\_LOG\_TABLE [3696])

Message: You can't use locks with log tables.

- Error: 1557 SQLSTATE: 23000 (ER\_FOREIGN\_DUPLICATE\_KEY\_OLD\_UNUSED [3696])

Message: Upholding foreign key constraints for table '%s', entry '%s', key %d would lead to a duplicate entry

- Error: 1558 SQLSTATE: HY000 (ER\_COL\_COUNT\_DOESNT\_MATCH\_PLEASE\_UPDATE [3696])

Message: Column count of mysql.%s is wrong. Expected %d, found %d. Created with MySQL %d, now running %d. Please use mysql\_upgrade to fix this error.

- Error: 1559 SQLSTATE: HY000 (ER\_TEMP\_TABLE\_PREVENTS\_SWITCH\_OUT\_OF\_RBR [3696])

Message: Cannot switch out of the row-based binary log format when the session has open temporary tables

- Error: 1560 SQLSTATE: HY000 (ER\_STORED\_FUNCTION\_PREVENTS\_SWITCH\_BINLOG\_FORMAT [3696])

Message: Cannot change the binary logging format inside a stored function or trigger

- Error: 1561 SQLSTATE: HY000 (ER\_NDB\_CANT\_SWITCH\_BINLOG\_FORMAT [3696])

Message: The NDB cluster engine does not support changing the binlog format on the fly yet

- Error: 1562 SQLSTATE: HY000 (ER\_PARTITION\_NO\_TEMPORARY [3696])

Message: Cannot create temporary table with partitions

- Error: 1563 SQLSTATE: HY000 (ER\_PARTITION\_CONST\_DOMAIN\_ERROR [3696])

Message: Partition constant is out of partition function domain

- Error: 1564 SQLSTATE: HY000 (ER\_PARTITION\_FUNCTION\_IS\_NOT\_ALLOWED [3696])

Message: This partition function is not allowed

- Error: 1565 SQLSTATE: HY000 (ER\_DDL\_LOG\_ERROR [3696])

Message: Error in DDL log

- Error: 1566 SQLSTATE: HY000 (ER\_NULL\_IN\_VALUES\_LESS\_THAN [3696])



Message: Not allowed to use NULL value in VALUES LESS THAN

- Error: [1567](#) SQLSTATE: [HY000](#) ([ER\\_WRONG\\_PARTITION\\_NAME](#) [\[3697\]](#))

Message: Incorrect partition name

- Error: [1568](#) SQLSTATE: [25001](#) ([ER\\_CANT\\_CHANGE\\_TX\\_CHARACTERISTICS](#) [\[3697\]](#))

Message: Transaction characteristics can't be changed while a transaction is in progress

- Error: [1569](#) SQLSTATE: [HY000](#) ([ER\\_DUP\\_ENTRY\\_AUTOINCREMENT\\_CASE](#) [\[3697\]](#))

Message: ALTER TABLE causes auto\_increment resequencing, resulting in duplicate entry '%s' for key '%s'

- Error: [1570](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_MODIFY\\_QUEUE\\_ERROR](#) [\[3697\]](#))

Message: Internal scheduler error %d

- Error: [1571](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_SET\\_VAR\\_ERROR](#) [\[3697\]](#))

Message: Error during starting/stopping of the scheduler. Error code %u

- Error: [1572](#) SQLSTATE: [HY000](#) ([ER\\_PARTITION\\_MERGE\\_ERROR](#) [\[3697\]](#))

Message: Engine cannot be used in partitioned tables

- Error: [1573](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_ACTIVATE\\_LOG](#) [\[3697\]](#))

Message: Cannot activate '%s' log

- Error: [1574](#) SQLSTATE: [HY000](#) ([ER\\_RBR\\_NOT\\_AVAILABLE](#) [\[3697\]](#))

Message: The server was not built with row-based replication

- Error: [1575](#) SQLSTATE: [HY000](#) ([ER\\_BASE64\\_DECODE\\_ERROR](#) [\[3697\]](#))

Message: Decoding of base64 string failed

- Error: [1576](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_RECURSION\\_FORBIDDEN](#) [\[3697\]](#))

Message: Recursion of EVENT DDL statements is forbidden when body is present

- Error: [1577](#) SQLSTATE: [HY000](#) ([ER\\_EVENTS\\_DB\\_ERROR](#) [\[3697\]](#))

Message: Cannot proceed because system tables used by Event Scheduler were found damaged at server start

- Error: [1578](#) SQLSTATE: [HY000](#) ([ER\\_ONLY\\_INTEGERS\\_ALLOWED](#) [\[3697\]](#))

Message: Only integers allowed as number here

- Error: [1579](#) SQLSTATE: [HY000](#) ([ER\\_UNSUPPORTED\\_LOG\\_ENGINE](#) [\[3697\]](#))

Message: This storage engine cannot be used for log tables"

- Error: [1580](#) SQLSTATE: [HY000](#) ([ER\\_BAD\\_LOG\\_STATEMENT](#) [\[3697\]](#))

Message: You cannot '%s' a log table if logging is enabled

- Error: [1581](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_RENAME\\_LOG\\_TABLE](#) [\[3697\]](#))

Message: Cannot rename '%s'. When logging enabled, rename to/from log table must rename two tables: the log table to an archive table and another table back to '%s'

- Error: [1582](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_PARAMCOUNT\\_TO\\_NATIVE\\_FCT](#) [\[3698\]](#))  
Message: Incorrect parameter count in the call to native function '%s'
- Error: [1583](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_PARAMETERS\\_TO\\_NATIVE\\_FCT](#) [\[3698\]](#))  
Message: Incorrect parameters in the call to native function '%s'
- Error: [1584](#) SQLSTATE: [42000](#) ([ER\\_WRONG\\_PARAMETERS\\_TO\\_STORED\\_FCT](#) [\[3698\]](#))  
Message: Incorrect parameters in the call to stored function '%s'
- Error: [1585](#) SQLSTATE: [HY000](#) ([ER\\_NATIVE\\_FCT\\_NAME\\_COLLISION](#) [\[3698\]](#))  
Message: This function '%s' has the same name as a native function
- Error: [1586](#) SQLSTATE: [23000](#) ([ER\\_DUP\\_ENTRY\\_WITH\\_KEY\\_NAME](#) [\[3698\]](#))  
Message: Duplicate entry '%s' for key '%s'  
  
As of MySQL 5.1.20, the format string for this error is also used with [ER\\_DUP\\_ENTRY](#) [\[3664\]](#).
- Error: [1587](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_PURGE\\_EMFILE](#) [\[3698\]](#))  
Message: Too many files opened, please execute the command again
- Error: [1588](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_CANNOT\\_CREATE\\_IN\\_THE\\_PAST](#) [\[3698\]](#))  
Message: Event execution time is in the past and ON COMPLETION NOT PRESERVE is set. The event was dropped immediately after creation.
- Error: [1589](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_CANNOT ALTER\\_IN\\_THE\\_PAST](#) [\[3698\]](#))  
Message: Event execution time is in the past and ON COMPLETION NOT PRESERVE is set. The event was not changed. Specify a time in the future.
- Error: [1590](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_INCIDENT](#) [\[3698\]](#))  
Message: The incident %s occurred on the master. Message: %s
- Error: [1591](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_PARTITION\\_FOR\\_GIVEN\\_VALUE\\_SILENT](#) [\[3698\]](#))  
Message: Table has no partition for some existing values
- Error: [1592](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_STATEMENT](#) [\[3698\]](#))  
Message: Unsafe statement written to the binary log using statement format since BINLOG\_FORMAT = STATEMENT. %s
- Error: [1593](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_FATAL\\_ERROR](#) [\[3698\]](#))  
Message: Fatal error: %s
- Error: [1594](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_RELAY\\_LOG\\_READ\\_FAILURE](#) [\[3698\]](#))  
Message: Relay log read failure: %s
- Error: [1595](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_RELAY\\_LOG\\_WRITE\\_FAILURE](#) [\[3698\]](#))  
Message: Relay log write failure: %s
- Error: [1596](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_CREATE\\_EVENT\\_FAILURE](#) [\[3698\]](#))  
Message: Failed to create %s

- Error: [1597](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_MASTER\\_COM\\_FAILURE](#) [\[3699\]](#))  
Message: Master command %s failed: %s
- Error: [1598](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_LOGGING\\_IMPOSSIBLE](#) [\[3699\]](#))  
Message: Binary logging not possible. Message: %s
- Error: [1599](#) SQLSTATE: [HY000](#) ([ER\\_VIEW\\_NO\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: View `%s`.`%s` has no creation context
- Error: [1600](#) SQLSTATE: [HY000](#) ([ER\\_VIEW\\_INVALID\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: Creation context of view `%s`.`%s` is invalid
- Error: [1601](#) SQLSTATE: [HY000](#) ([ER\\_SR\\_INVALID\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: Creation context of stored routine `%s`.`%s` is invalid
- Error: [1602](#) SQLSTATE: [HY000](#) ([ER\\_TRG\\_CORRUPTED\\_FILE](#) [\[3699\]](#))  
Message: Corrupted TRG file for table `%s`.`%s`
- Error: [1603](#) SQLSTATE: [HY000](#) ([ER\\_TRG\\_NO\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: Triggers for table `%s`.`%s` have no creation context
- Error: [1604](#) SQLSTATE: [HY000](#) ([ER\\_TRG\\_INVALID\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: Trigger creation context of table `%s`.`%s` is invalid
- Error: [1605](#) SQLSTATE: [HY000](#) ([ER\\_EVENT\\_INVALID\\_CREATION\\_CTX](#) [\[3699\]](#))  
Message: Creation context of event `%s`.`%s` is invalid
- Error: [1606](#) SQLSTATE: [HY000](#) ([ER\\_TRG\\_CANT\\_OPEN\\_TABLE](#) [\[3699\]](#))  
Message: Cannot open table for trigger `%s`.`%s`
- Error: [1607](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_CREATE\\_SROUTINE](#) [\[3699\]](#))  
Message: Cannot create stored routine `%s`. Check warnings
- Error: [1608](#) SQLSTATE: [HY000](#) ([ER\\_NEVER\\_USED](#) [\[3699\]](#))  
Message: Ambiguous slave modes combination. %s
- Error: [1609](#) SQLSTATE: [HY000](#)  
([ER\\_NO\\_FORMAT\\_DESCRIPTION\\_EVENT\\_BEFORE\\_BINLOG\\_STATEMENT](#) [\[3699\]](#))  
Message: The BINLOG statement of type `%s` was not preceded by a format description BINLOG statement.
- Error: [1610](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_CORRUPT\\_EVENT](#) [\[3699\]](#))  
Message: Corrupted replication event was detected
- Error: [1611](#) SQLSTATE: [HY000](#) ([ER\\_LOAD\\_DATA\\_INVALID\\_COLUMN](#) [\[3699\]](#))  
Message: Invalid column reference (%s) in LOAD DATA
- Error: [1612](#) SQLSTATE: [HY000](#) ([ER\\_LOG\\_PURGE\\_NO\\_FILE](#) [\[3699\]](#))  
Message: Being purged log %s was not found

- Error: 1613 SQLSTATE: XA106 (ER\_XA\_RBTIMEOUT [3700])  
Message: XA\_RBTIMEOUT: Transaction branch was rolled back: took too long
- Error: 1614 SQLSTATE: XA102 (ER\_XA\_RBDEADLOCK [3700])  
Message: XA\_RBDEADLOCK: Transaction branch was rolled back: deadlock was detected
- Error: 1615 SQLSTATE: HY000 (ER\_NEED\_REPREPARE [3700])  
Message: Prepared statement needs to be re-prepared
- Error: 1616 SQLSTATE: HY000 (ER\_DELAYED\_NOT\_SUPPORTED [3700])  
Message: DELAYED option not supported for table '%s'
- Error: 1617 SQLSTATE: HY000 (WARN\_NO\_MASTER\_INFO [3700])  
Message: The master info structure does not exist
- Error: 1618 SQLSTATE: HY000 (WARN\_OPTION\_IGNORED [3700])  
Message: <%s> option ignored
- Error: 1619 SQLSTATE: HY000 (WARN\_PLUGIN\_DELETE\_BUILTIN [3700])  
Message: Built-in plugins cannot be deleted
- Error: 1620 SQLSTATE: HY000 (WARN\_PLUGIN\_BUSY [3700])  
Message: Plugin is busy and will be uninstalled on shutdown
- Error: 1621 SQLSTATE: HY000 (ER\_VARIABLE\_IS\_READONLY [3700])  
Message: %s variable '%s' is read-only. Use SET %s to assign the value
- Error: 1622 SQLSTATE: HY000 (ER\_WARN\_ENGINE\_TRANSACTION\_ROLLBACK [3700])  
Message: Storage engine %s does not support rollback for this statement. Transaction rolled back and must be restarted
- Error: 1623 SQLSTATE: HY000 (ER\_SLAVE\_HEARTBEAT\_FAILURE [3700])  
Message: Unexpected master's heartbeat data: %s
- Error: 1624 SQLSTATE: HY000 (ER\_SLAVE\_HEARTBEAT\_VALUE\_OUT\_OF\_RANGE [3700])  
Message: The requested value for the heartbeat period is either negative or exceeds the maximum allowed (%s seconds).
- Error: 1625 SQLSTATE: HY000 (ER\_NDB\_REPLICATION\_SCHEMA\_ERROR [3700])  
Message: Bad schema for mysql.ndb\_replication table. Message: %s
- Error: 1626 SQLSTATE: HY000 (ER\_CONFLICT\_FN\_PARSE\_ERROR [3700])  
Message: Error in parsing conflict function. Message: %s
- Error: 1627 SQLSTATE: HY000 (ER\_EXCEPTIONS\_WRITE\_ERROR [3700])  
Message: Write to exceptions table failed. Message: %s"
- Error: 1628 SQLSTATE: HY000 (ER\_TOO\_LONG\_TABLE\_COMMENT [3700])  
Message: Comment for table '%s' is too long (max = %lu)

- Error: [1629](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_LONG\\_FIELD\\_COMMENT](#) [\[3701\]](#))  
Message: Comment for field '%s' is too long (max = %lu)
- Error: [1630](#) SQLSTATE: [42000](#) ([ER\\_FUNC\\_INEXISTENT\\_NAME\\_COLLISION](#) [\[3701\]](#))  
Message: FUNCTION %s does not exist. Check the 'Function Name Parsing and Resolution' section in the Reference Manual
- Error: [1631](#) SQLSTATE: [HY000](#) ([ER\\_DATABASE\\_NAME](#) [\[3701\]](#))  
Message: Database
- Error: [1632](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_NAME](#) [\[3701\]](#))  
Message: Table
- Error: [1633](#) SQLSTATE: [HY000](#) ([ER\\_PARTITION\\_NAME](#) [\[3701\]](#))  
Message: Partition
- Error: [1634](#) SQLSTATE: [HY000](#) ([ER\\_SUBPARTITION\\_NAME](#) [\[3701\]](#))  
Message: Subpartition
- Error: [1635](#) SQLSTATE: [HY000](#) ([ER\\_TEMPORARY\\_NAME](#) [\[3701\]](#))  
Message: Temporary
- Error: [1636](#) SQLSTATE: [HY000](#) ([ER\\_RENAMED\\_NAME](#) [\[3701\]](#))  
Message: Renamed
- Error: [1637](#) SQLSTATE: [HY000](#) ([ER\\_TOO\\_MANY\\_CONCURRENT\\_TRXS](#) [\[3701\]](#))  
Message: Too many active concurrent transactions
- Error: [1638](#) SQLSTATE: [HY000](#) ([WARN\\_NON\\_ASCII\\_SEPARATOR\\_NOT\\_IMPLEMENTED](#) [\[3701\]](#))  
Message: Non-ASCII separator arguments are not fully supported
- Error: [1639](#) SQLSTATE: [HY000](#) ([ER\\_DEBUG\\_SYNC\\_TIMEOUT](#) [\[3701\]](#))  
Message: debug sync point wait timed out
- Error: [1640](#) SQLSTATE: [HY000](#) ([ER\\_DEBUG\\_SYNC\\_HIT\\_LIMIT](#) [\[3701\]](#))  
Message: debug sync point hit limit reached
- Error: [1641](#) SQLSTATE: [42000](#) ([ER\\_DUP\\_SIGNAL\\_SET](#) [\[3701\]](#))  
Message: Duplicate condition information item '%s'
- Error: [1642](#) SQLSTATE: [01000](#) ([ER\\_SIGNAL\\_WARN](#) [\[3701\]](#))  
Message: Unhandled user-defined warning condition
- Error: [1643](#) SQLSTATE: [02000](#) ([ER\\_SIGNAL\\_NOT\\_FOUND](#) [\[3701\]](#))  
Message: Unhandled user-defined not found condition
- Error: [1644](#) SQLSTATE: [HY000](#) ([ER\\_SIGNAL\\_EXCEPTION](#) [\[3701\]](#))  
Message: Unhandled user-defined exception condition
- Error: [1645](#) SQLSTATE: [0K000](#) ([ER\\_RESIGNAL\\_WITHOUT\\_ACTIVE\\_HANDLER](#) [\[3701\]](#))

Message: RESIGNAL when handler not active

- Error: 1646 SQLSTATE: HY000 (ER\_SIGNAL\_BAD\_CONDITION\_TYPE [3702])

Message: SIGNAL/RESIGNAL can only use a CONDITION defined with SQLSTATE

- Error: 1647 SQLSTATE: HY000 (WARN\_COND\_ITEM\_TRUNCATED [3702])

Message: Data truncated for condition item '%s'

- Error: 1648 SQLSTATE: HY000 (ER\_COND\_ITEM\_TOO\_LONG [3702])

Message: Data too long for condition item '%s'

- Error: 1649 SQLSTATE: HY000 (ER\_UNKNOWN\_LOCALE [3702])

Message: Unknown locale: '%s'

- Error: 1650 SQLSTATE: HY000 (ER\_SLAVE\_IGNORE\_SERVER\_IDS [3702])

Message: The requested server id %d clashes with the slave startup option --replicate-same-server-id

- Error: 1651 SQLSTATE: HY000 (ER\_QUERY\_CACHE\_DISABLED [3702])

Message: Query cache is disabled; restart the server with query\_cache\_type=1 to enable it

- Error: 1652 SQLSTATE: HY000 (ER\_SAME\_NAME\_PARTITION\_FIELD [3702])

Message: Duplicate partition field name '%s'

- Error: 1653 SQLSTATE: HY000 (ER\_PARTITION\_COLUMN\_LIST\_ERROR [3702])

Message: Inconsistency in usage of column lists for partitioning

- Error: 1654 SQLSTATE: HY000 (ER\_WRONG\_TYPE\_COLUMN\_VALUE\_ERROR [3702])

Message: Partition column values of incorrect type

- Error: 1655 SQLSTATE: HY000 (ER\_TOO\_MANY\_PARTITION\_FUNC\_FIELDS\_ERROR [3702])

Message: Too many fields in '%s'

- Error: 1656 SQLSTATE: HY000 (ER\_MAXVALUE\_IN\_VALUES\_IN [3702])

Message: Cannot use MAXVALUE as value in VALUES IN

- Error: 1657 SQLSTATE: HY000 (ER\_TOO\_MANY\_VALUES\_ERROR [3702])

Message: Cannot have more than one value for this type of %s partitioning

- Error: 1658 SQLSTATE: HY000 (ER\_ROW\_SINGLE\_PARTITION\_FIELD\_ERROR [3702])

Message: Row expressions in VALUES IN only allowed for multi-field column partitioning

- Error: 1659 SQLSTATE: HY000 (ER\_FIELD\_TYPE\_NOT\_ALLOWED\_AS\_PARTITION\_FIELD [3702])

Message: Field '%s' is of a not allowed type for this type of partitioning

- Error: 1660 SQLSTATE: HY000 (ER\_PARTITION\_FIELDS\_TOO\_LONG [3702])

Message: The total length of the partitioning fields is too large

- Error: 1661 SQLSTATE: HY000 (ER\_BINLOG\_ROW\_ENGINE\_AND\_STMT\_ENGINE [3702])

Message: Cannot execute statement: impossible to write to binary log since both row-incapable engines and statement-incapable engines are involved.

- Error: 1662 SQLSTATE: HY000 ([ER\\_BINLOG\\_ROW\\_MODE\\_AND\\_STMT\\_ENGINE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since BINLOG\_FORMAT = ROW and at least one table uses a storage engine limited to statement-based logging.

- Error: 1663 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_AND\\_STMT\\_ENGINE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since statement is unsafe, storage engine is limited to statement-based logging, and BINLOG\_FORMAT = MIXED. %s

- Error: 1664 SQLSTATE: HY000 ([ER\\_BINLOG\\_ROW\\_INJECTION\\_AND\\_STMT\\_ENGINE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since statement is in row format and at least one table uses a storage engine limited to statement-based logging.

- Error: 1665 SQLSTATE: HY000 ([ER\\_BINLOG\\_STMT\\_MODE\\_AND\\_ROW\\_ENGINE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since BINLOG\_FORMAT = STATEMENT and at least one table uses a storage engine limited to row-based logging. %s

- Error: 1666 SQLSTATE: HY000 ([ER\\_BINLOG\\_ROW\\_INJECTION\\_AND\\_STMT\\_MODE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since statement is in row format and BINLOG\_FORMAT = STATEMENT.

- Error: 1667 SQLSTATE: HY000 ([ER\\_BINLOG\\_MULTIPLE\\_ENGINES\\_AND\\_SELF\\_LOGGING\\_ENGINE \[3703\]](#))

Message: Cannot execute statement: impossible to write to binary log since more than one engine is involved and at least one engine is self-logging.

- Error: 1668 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_LIMIT \[3703\]](#))

Message: The statement is unsafe because it uses a LIMIT clause. This is unsafe because the set of rows included cannot be predicted.

- Error: 1669 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_INSERT\\_DELAYED \[3703\]](#))

Message: The statement is unsafe because it uses INSERT DELAYED. This is unsafe because the times when rows are inserted cannot be predicted.

- Error: 1670 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_SYSTEM\\_TABLE \[3703\]](#))

Message: The statement is unsafe because it uses the general log, slow query log, or performance\_schema table(s). This is unsafe because system tables may differ on slaves.

- Error: 1671 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_AUTOINC\\_COLUMNS \[3703\]](#))

Message: Statement is unsafe because it invokes a trigger or a stored function that inserts into an AUTO\_INCREMENT column. Inserted values cannot be logged correctly.

- Error: 1672 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_UDF \[3703\]](#))

Message: Statement is unsafe because it uses a UDF which may not return the same value on the slave.

- Error: 1673 SQLSTATE: HY000 ([ER\\_BINLOG\\_UNSAFE\\_SYSTEM\\_VARIABLE \[3703\]](#))

Message: Statement is unsafe because it uses a system variable that may have a different value on the slave.

- Error: 1674 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_SYSTEM\_FUNCTION [3704])  
Message: Statement is unsafe because it uses a system function that may return a different value on the slave.
- Error: 1675 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_NONTRANS\_AFTER\_TRANS [3704])  
Message: Statement is unsafe because it accesses a non-transactional table after accessing a transactional table within the same transaction.
- Error: 1676 SQLSTATE: HY000 (ER\_MESSAGE\_AND\_STATEMENT [3704])  
Message: %s Statement: %s
- Error: 1677 SQLSTATE: HY000 (ER\_SLAVE\_CONVERSION\_FAILED [3704])  
Message: Column %d of table '%s.%s' cannot be converted from type '%s' to type '%s'
- Error: 1678 SQLSTATE: HY000 (ER\_SLAVE\_CANT\_CREATE\_CONVERSION [3704])  
Message: Can't create conversion table for table '%s.%s'
- Error: 1679 SQLSTATE: HY000  
(ER\_INSIDE\_TRANSACTION\_PREVENTS\_SWITCH\_BINLOG\_FORMAT [3704])  
Message: Cannot modify @@session.binlog\_format inside a transaction
- Error: 1680 SQLSTATE: HY000 (ER\_PATH\_LENGTH [3704])  
Message: The path specified for %s is too long.
- Error: 1681 SQLSTATE: HY000 (ER\_WARN\_DEPRECATED\_SYNTAX\_NO\_REPLACEMENT [3704])  
Message: '%s' is deprecated and will be removed in a future release.
- Error: 1682 SQLSTATE: HY000 (ER\_WRONG\_NATIVE\_TABLE\_STRUCTURE [3704])  
Message: Native table '%s'.'%s' has the wrong structure
- Error: 1683 SQLSTATE: HY000 (ER\_WRONG\_PERFSCHEMA\_USAGE [3704])  
Message: Invalid performance\_schema usage.
- Error: 1684 SQLSTATE: HY000 (ER\_WARN\_I\_S\_SKIPPED\_TABLE [3704])  
Message: Table '%s'.'%s' was skipped since its definition is being modified by concurrent DDL statement
- Error: 1685 SQLSTATE: HY000  
(ER\_INSIDE\_TRANSACTION\_PREVENTS\_SWITCH\_BINLOG\_DIRECT [3704])  
Message: Cannot modify @@session.binlog\_direct\_non\_transactional\_updates inside a transaction
- Error: 1686 SQLSTATE: HY000  
(ER\_STORED\_FUNCTION\_PREVENTS\_SWITCH\_BINLOG\_DIRECT [3704])  
Message: Cannot change the binlog direct flag inside a stored function or trigger
- Error: 1687 SQLSTATE: 42000 (ER\_SPATIAL\_MUST\_HAVE\_GEOM\_COL [3704])  
Message: A SPATIAL index may only contain a geometrical type column
- Error: 1688 SQLSTATE: HY000 (ER\_TOO\_LONG\_INDEX\_COMMENT [3704])



Message: Comment for index '%s' is too long (max = %lu)

- Error: 1689 SQLSTATE: HY000 (ER\_LOCK\_ABORTED [3705])

Message: Wait on a lock was aborted due to a pending exclusive lock

- Error: 1690 SQLSTATE: 22003 (ER\_DATA\_OUT\_OF\_RANGE [3705])

Message: %s value is out of range in '%s'

- Error: 1691 SQLSTATE: HY000 (ER\_WRONG\_SPVAR\_TYPE\_IN\_LIMIT [3705])

Message: A variable of a non-integer based type in LIMIT clause

- Error: 1692 SQLSTATE: HY000  
(ER\_BINLOG\_UNSAFE\_MULTIPLE\_ENGINES\_AND\_SELF\_LOGGING\_ENGINE [3705])

Message: Mixing self-logging and non-self-logging engines in a statement is unsafe.

- Error: 1693 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_MIXED\_STATEMENT [3705])

Message: Statement accesses nontransactional table as well as transactional or temporary table, and writes to any of them.

- Error: 1694 SQLSTATE: HY000  
(ER\_INSIDE\_TRANSACTION\_PREVENTS\_SWITCH\_SQL\_LOG\_BIN [3705])

Message: Cannot modify @@session.sql\_log\_bin inside a transaction

- Error: 1695 SQLSTATE: HY000  
(ER\_STORED\_FUNCTION\_PREVENTS\_SWITCH\_SQL\_LOG\_BIN [3705])

Message: Cannot change the sql\_log\_bin inside a stored function or trigger

- Error: 1696 SQLSTATE: HY000 (ER\_FAILED\_READ\_FROM\_PAR\_FILE [3705])

Message: Failed to read from the .par file

- Error: 1697 SQLSTATE: HY000 (ER\_VALUES\_IS\_NOT\_INT\_TYPE\_ERROR [3705])

Message: VALUES value for partition '%s' must have type INT

- Error: 1698 SQLSTATE: 28000 (ER\_ACCESS\_DENIED\_NO\_PASSWORD\_ERROR [3705])

Message: Access denied for user '%s'@'%s'

- Error: 1699 SQLSTATE: HY000 (ER\_SET\_PASSWORD\_AUTH\_PLUGIN [3705])

Message: SET PASSWORD has no significance for users authenticating via plugins

- Error: 1700 SQLSTATE: HY000 (ER\_GRANT\_PLUGIN\_USER\_EXISTS [3705])

Message: GRANT with IDENTIFIED WITH is illegal because the user %-\*s already exists

- Error: 1701 SQLSTATE: 42000 (ER\_TRUNCATE\_ILLEGAL\_FK [3705])

Message: Cannot truncate a table referenced in a foreign key constraint (%s)

- Error: 1702 SQLSTATE: HY000 (ER\_PLUGIN\_IS\_PERMANENT [3705])

Message: Plugin '%s' is force\_plus\_permanent and can not be unloaded

- Error: 1703 SQLSTATE: HY000 (ER\_SLAVE\_HEARTBEAT\_VALUE\_OUT\_OF\_RANGE\_MIN [3705])

Message: The requested value for the heartbeat period is less than 1 millisecond. The value is reset to 0, meaning that heartbeating will effectively be disabled.

- Error: 1704 SQLSTATE: HY000 (ER\_SLAVE\_HEARTBEAT\_VALUE\_OUT\_OF\_RANGE\_MAX [3706])

Message: The requested value for the heartbeat period exceeds the value of `slave\_net\_timeout` seconds. A sensible value for the period should be less than the timeout.

- Error: 1705 SQLSTATE: HY000 (ER\_STMT\_CACHE\_FULL [3706])

Message: Multi-row statements required more than 'max\_binlog\_stmt\_cache\_size' bytes of storage; increase this mysqld variable and try again

- Error: 1706 SQLSTATE: HY000 (ER\_MULTI\_UPDATE\_KEY\_CONFLICT [3706])

Message: Primary key/partition key update is not allowed since the table is updated both as '%s' and '%s'.

- Error: 1707 SQLSTATE: HY000 (ER\_TABLE\_NEEDS\_REBUILD [3706])

Message: Table rebuild required. Please do "ALTER TABLE `'%s` FORCE" or dump/reload to fix it!

- Error: 1708 SQLSTATE: HY000 (WARN\_OPTION\_BELOW\_LIMIT [3706])

Message: The value of '%s' should be no less than the value of '%s'

- Error: 1709 SQLSTATE: HY000 (ER\_INDEX\_COLUMN\_TOO\_LONG [3706])

Message: Index column size too large. The maximum column size is %lu bytes.

- Error: 1710 SQLSTATE: HY000 (ER\_ERROR\_IN\_TRIGGER\_BODY [3706])

Message: Trigger '%s' has an error in its body: '%s'

- Error: 1711 SQLSTATE: HY000 (ER\_ERROR\_IN\_UNKNOWN\_TRIGGER\_BODY [3706])

Message: Unknown trigger has an error in its body: '%s'

- Error: 1712 SQLSTATE: HY000 (ER\_INDEX\_CORRUPT [3706])

Message: Index %s is corrupted

- Error: 1713 SQLSTATE: HY000 (ER\_UNDO\_RECORD\_TOO\_BIG [3706])

Message: Undo log record is too big.

- Error: 1714 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_INSERT\_IGNORE\_SELECT [3706])

Message: INSERT IGNORE... SELECT is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are ignored. This order cannot be predicted and may differ on master and the slave.

- Error: 1715 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_INSERT\_SELECT\_UPDATE [3706])

Message: INSERT... SELECT... ON DUPLICATE KEY UPDATE is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are updated. This order cannot be predicted and may differ on master and the slave.

- Error: 1716 SQLSTATE: HY000 (ER\_BINLOG\_UNSAFE\_REPLACE\_SELECT [3706])

Message: REPLACE... SELECT is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are replaced. This order cannot be predicted and may differ on master and the slave.

- Error: [1717](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_CREATE\\_IGNORE\\_SELECT](#) [\[3707\]](#))  
Message: CREATE... IGNORE SELECT is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are ignored. This order cannot be predicted and may differ on master and the slave.
- Error: [1718](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_CREATE\\_REPLACE\\_SELECT](#) [\[3707\]](#))  
Message: CREATE... REPLACE SELECT is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are replaced. This order cannot be predicted and may differ on master and the slave.
- Error: [1719](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_UPDATE\\_IGNORE](#) [\[3707\]](#))  
Message: UPDATE IGNORE is unsafe because the order in which rows are updated determines which (if any) rows are ignored. This order cannot be predicted and may differ on master and the slave.
- Error: [1720](#) SQLSTATE: [HY000](#) ([ER\\_PLUGIN\\_NO\\_UNINSTALL](#) [\[3707\]](#))  
Message: Plugin '%s' is marked as not dynamically uninstallable. You have to stop the server to uninstall it.
- Error: [1721](#) SQLSTATE: [HY000](#) ([ER\\_PLUGIN\\_NO\\_INSTALL](#) [\[3707\]](#))  
Message: Plugin '%s' is marked as not dynamically installable. You have to stop the server to install it.
- Error: [1722](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_WRITE\\_AUTOINC\\_SELECT](#) [\[3707\]](#))  
Message: Statements writing to a table with an auto-increment column after selecting from another table are unsafe because the order in which rows are retrieved determines what (if any) rows will be written. This order cannot be predicted and may differ on master and the slave.
- Error: [1723](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_CREATE\\_SELECT\\_AUTOINC](#) [\[3707\]](#))  
Message: CREATE TABLE... SELECT... on a table with an auto-increment column is unsafe because the order in which rows are retrieved by the SELECT determines which (if any) rows are inserted. This order cannot be predicted and may differ on master and the slave.
- Error: [1724](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_INSERT\\_TWO\\_KEYS](#) [\[3707\]](#))  
Message: INSERT... ON DUPLICATE KEY UPDATE on a table with more than one UNIQUE KEY is unsafe
- Error: [1725](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_IN\\_FK\\_CHECK](#) [\[3707\]](#))  
Message: Table is being used in foreign key check.
- Error: [1726](#) SQLSTATE: [HY000](#) ([ER\\_UNSUPPORTED\\_ENGINE](#) [\[3707\]](#))  
Message: Storage engine '%s' does not support system tables. [%s.%s]
- Error: [1727](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_UNSAFE\\_AUTOINC\\_NOT\\_FIRST](#) [\[3707\]](#))  
Message: INSERT into autoincrement field which is not the first part in the composed primary key is unsafe.
- Error: [1728](#) SQLSTATE: [HY000](#) ([ER\\_CANNOT\\_LOAD\\_FROM\\_TABLE\\_V2](#) [\[3707\]](#))  
Message: Cannot load from %s.%s. The table is probably corrupted
- Error: [1729](#) SQLSTATE: [HY000](#) ([ER\\_MASTER\\_DELAY\\_VALUE\\_OUT\\_OF\\_RANGE](#) [\[3707\]](#))

Message: The requested value %u for the master delay exceeds the maximum %u

- Error: 1730 SQLSTATE: HY000 ([ER\\_ONLY\\_FD\\_AND\\_RBR\\_EVENTS\\_ALLOWED\\_IN\\_BINLOG\\_STATEMENT \[3708\]](#))

Message: Only Format\_description\_log\_event and row events are allowed in BINLOG statements (but %s was provided)

- Error: 1731 SQLSTATE: HY000 ([ER\\_PARTITION\\_EXCHANGE\\_DIFFERENT\\_OPTION \[3708\]](#))

Message: Non matching attribute '%s' between partition and table

- Error: 1732 SQLSTATE: HY000 ([ER\\_PARTITION\\_EXCHANGE\\_PART\\_TABLE \[3708\]](#))

Message: Table to exchange with partition is partitioned: '%s'

- Error: 1733 SQLSTATE: HY000 ([ER\\_PARTITION\\_EXCHANGE\\_TEMP\\_TABLE \[3708\]](#))

Message: Table to exchange with partition is temporary: '%s'

- Error: 1734 SQLSTATE: HY000 ([ER\\_PARTITION\\_INSTEAD\\_OF\\_SUBPARTITION \[3708\]](#))

Message: Subpartitioned table, use subpartition instead of partition

- Error: 1735 SQLSTATE: HY000 ([ER\\_UNKNOWN\\_PARTITION \[3708\]](#))

Message: Unknown partition '%s' in table '%s'

- Error: 1736 SQLSTATE: HY000 ([ER\\_TABLES\\_DIFFERENT\\_METADATA \[3708\]](#))

Message: Tables have different definitions

- Error: 1737 SQLSTATE: HY000 ([ER\\_ROW\\_DOES\\_NOT\\_MATCH\\_PARTITION \[3708\]](#))

Message: Found a row that does not match the partition

- Error: 1738 SQLSTATE: HY000 ([ER\\_BINLOG\\_CACHE\\_SIZE\\_GREATER\\_THAN\\_MAX \[3708\]](#))

Message: Option binlog\_cache\_size (%lu) is greater than max\_binlog\_cache\_size (%lu); setting binlog\_cache\_size equal to max\_binlog\_cache\_size.

- Error: 1739 SQLSTATE: HY000 ([ER\\_WARN\\_INDEX\\_NOT\\_APPLICABLE \[3708\]](#))

Message: Cannot use %s access on index '%s' due to type or collation conversion on field '%s'

- Error: 1740 SQLSTATE: HY000 ([ER\\_PARTITION\\_EXCHANGE\\_FOREIGN\\_KEY \[3708\]](#))

Message: Table to exchange with partition has foreign key references: '%s'

- Error: 1741 SQLSTATE: HY000 ([ER\\_NO\\_SUCH\\_KEY\\_VALUE \[3708\]](#))

Message: Key value '%s' was not found in table '%s.%s'

- Error: 1742 SQLSTATE: HY000 ([ER\\_RPL\\_INFO\\_DATA\\_TOO\\_LONG \[3708\]](#))

Message: Data for column '%s' too long

- Error: 1743 SQLSTATE: HY000 ([ER\\_NETWORK\\_READ\\_EVENT\\_CHECKSUM\\_FAILURE \[3708\]](#))

Message: Replication event checksum verification failed while reading from network.

- Error: 1744 SQLSTATE: HY000 ([ER\\_BINLOG\\_READ\\_EVENT\\_CHECKSUM\\_FAILURE \[3708\]](#))

Message: Replication event checksum verification failed while reading from a log file.

- Error: [1745](#) SQLSTATE: [HY000](#) ([ER\\_BINLOG\\_STMT\\_CACHE\\_SIZE\\_GREATER\\_THAN\\_MAX](#) [\[3709\]](#))  
Message: Option binlog\_stmt\_cache\_size (%lu) is greater than max\_binlog\_stmt\_cache\_size (%lu); setting binlog\_stmt\_cache\_size equal to max\_binlog\_stmt\_cache\_size.
- Error: [1746](#) SQLSTATE: [HY000](#) ([ER\\_CANT\\_UPDATE\\_TABLE\\_IN\\_CREATE\\_TABLE\\_SELECT](#) [\[3709\]](#))  
Message: Can't update table '%s' while '%s' is being created.
- Error: [1747](#) SQLSTATE: [HY000](#) ([ER\\_PARTITION\\_CLAUSE\\_ON\\_NONPARTITIONED](#) [\[3709\]](#))  
Message: PARTITION () clause on non partitioned table
- Error: [1748](#) SQLSTATE: [HY000](#) ([ER\\_ROW\\_DOES\\_NOT\\_MATCH\\_GIVEN\\_PARTITION\\_SET](#) [\[3709\]](#))  
Message: Found a row not matching the given partition set
- Error: [1749](#) SQLSTATE: [HY000](#) ([ER\\_NO\\_SUCH\\_PARTITION\\_\\_UNUSED](#) [\[3709\]](#))  
Message: partition '%s' doesn't exist
- Error: [1750](#) SQLSTATE: [HY000](#) ([ER\\_CHANGE\\_RPL\\_INFO\\_REPOSITORY\\_FAILURE](#) [\[3709\]](#))  
Message: Failure while changing the type of replication repository: %s.
- Error: [1751](#) SQLSTATE: [HY000](#)  
([ER\\_WARNING\\_NOT\\_COMPLETE\\_ROLLBACK\\_WITH\\_CREATED\\_TEMP\\_TABLE](#) [\[3709\]](#))  
Message: The creation of some temporary tables could not be rolled back.
- Error: [1752](#) SQLSTATE: [HY000](#)  
([ER\\_WARNING\\_NOT\\_COMPLETE\\_ROLLBACK\\_WITH\\_DROPPED\\_TEMP\\_TABLE](#) [\[3709\]](#))  
Message: Some temporary tables were dropped, but these operations could not be rolled back.
- Error: [1753](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_FEATURE\\_IS\\_NOT\\_SUPPORTED](#) [\[3709\]](#))  
Message: %s is not supported in multi-threaded slave mode. %s
- Error: [1754](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_UPDATED\\_DBS\\_GREATER\\_MAX](#) [\[3709\]](#))  
Message: The number of modified databases exceeds the maximum %d; the database names will not be included in the replication event metadata.
- Error: [1755](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_CANT\\_PARALLEL](#) [\[3709\]](#))  
Message: Cannot execute the current event group in the parallel mode. Encountered event %s, relay-log name %s, position %s which prevents execution of this event group in parallel mode. Reason: %s.
- Error: [1756](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_INCONSISTENT\\_DATA](#) [\[3709\]](#))  
Message: %s
- Error: [1757](#) SQLSTATE: [HY000](#) ([ER\\_FULLTEXT\\_NOT\\_SUPPORTED\\_WITH\\_PARTITIONING](#) [\[3709\]](#))  
Message: FULLTEXT index is not supported for partitioned tables.
- Error: [1758](#) SQLSTATE: [35000](#) ([ER\\_DA\\_INVALID\\_CONDITION\\_NUMBER](#) [\[3709\]](#))  
Message: Invalid condition number
- Error: [1759](#) SQLSTATE: [HY000](#) ([ER\\_INSECURE\\_PLAIN\\_TEXT](#) [\[3709\]](#))

Message: Sending passwords in plain text without SSL/TLS is extremely insecure.

- Error: 1760 SQLSTATE: HY000 (ER\_INSECURE\_CHANGE\_MASTER [3710])

Message: Storing MySQL user name or password information in the master.info repository is not secure and is therefore not recommended. Please see the MySQL Manual for more about this issue and possible alternatives.

- Error: 1761 SQLSTATE: 23000 (ER\_FOREIGN\_DUPLICATE\_KEY\_WITH\_CHILD\_INFO [3710])

Message: Foreign key constraint for table '%s', record '%s' would lead to a duplicate entry in table '%s', key '%s'

- Error: 1762 SQLSTATE: 23000 (ER\_FOREIGN\_DUPLICATE\_KEY\_WITHOUT\_CHILD\_INFO [3710])

Message: Foreign key constraint for table '%s', record '%s' would lead to a duplicate entry in a child table

- Error: 1763 SQLSTATE: HY000 (ER\_SQLTHREAD\_WITH\_SECURE\_SLAVE [3710])

Message: Setting authentication options is not possible when only the Slave SQL Thread is being started.

- Error: 1764 SQLSTATE: HY000 (ER\_TABLE\_HAS\_NO\_FT [3710])

Message: The table does not have FULLTEXT index to support this query

- Error: 1765 SQLSTATE: HY000 (ER\_VARIABLE\_NOT\_SETTABLE\_IN\_SF\_OR\_TRIGGER [3710])

Message: The system variable %s cannot be set in stored functions or triggers.

- Error: 1766 SQLSTATE: HY000 (ER\_VARIABLE\_NOT\_SETTABLE\_IN\_TRANSACTION [3710])

Message: The system variable %s cannot be set when there is an ongoing transaction.

- Error: 1767 SQLSTATE: HY000 (ER\_GTID\_NEXT\_IS\_NOT\_IN\_GTID\_NEXT\_LIST [3710])

Message: The system variable @@SESSION.GTID\_NEXT has the value %s, which is not listed in @@SESSION.GTID\_NEXT\_LIST.

- Error: 1768 SQLSTATE: HY000 (ER\_CANT\_CHANGE\_GTID\_NEXT\_IN\_TRANSACTION\_WHEN\_GTID\_NEXT\_LIST\_IS\_NULL [3710])

Message: When @@SESSION.GTID\_NEXT\_LIST == NULL, the system variable @@SESSION.GTID\_NEXT cannot change inside a transaction.

- Error: 1769 SQLSTATE: HY000 (ER\_SET\_STATEMENT\_CANNOT\_INVOKE\_FUNCTION [3710])

Message: The statement 'SET %s' cannot invoke a stored function.

- Error: 1770 SQLSTATE: HY000 (ER\_GTID\_NEXT\_CANT\_BE\_AUTOMATIC\_IF\_GTID\_NEXT\_LIST\_IS\_NON\_NULL [3710])

Message: The system variable @@SESSION.GTID\_NEXT cannot be 'AUTOMATIC' when @@SESSION.GTID\_NEXT\_LIST is non-NULL.

- Error: 1771 SQLSTATE: HY000 (ER\_SKIPPING\_LOGGED\_TRANSACTION [3710])

Message: Skipping transaction %s because it has already been executed and logged.

- Error: 1772 SQLSTATE: HY000 (ER\_MALFORMED\_GTID\_SET\_SPECIFICATION [3710])

Message: Malformed GTID set specification '%s'.

- Error: 1773 SQLSTATE: HY000 (ER\_MALFORMED\_GTID\_SET\_ENCODING [3711])  
Message: Malformed GTID set encoding.
- Error: 1774 SQLSTATE: HY000 (ER\_MALFORMED\_GTID\_SPECIFICATION [3711])  
Message: Malformed GTID specification '%s'.
- Error: 1775 SQLSTATE: HY000 (ER\_GNO\_EXHAUSTED [3711])  
Message: Impossible to generate Global Transaction Identifier: the integer component reached the maximal value. Restart the server with a new server\_uuid.
- Error: 1776 SQLSTATE: HY000 (ER\_BAD\_SLAVE\_AUTO\_POSITION [3711])  
Message: Parameters MASTER\_LOG\_FILE, MASTER\_LOG\_POS, RELAY\_LOG\_FILE and RELAY\_LOG\_POS cannot be set when MASTER\_AUTO\_POSITION is active.
- Error: 1777 SQLSTATE: HY000 (ER\_AUTO\_POSITION\_REQUIRES\_GTID\_MODE\_ON [3711])  
Message: CHANGE MASTER TO MASTER\_AUTO\_POSITION = 1 can only be executed when @@GLOBAL.GTID\_MODE = ON.
- Error: 1778 SQLSTATE: HY000 (ER\_CANT\_DO\_IMPLICIT\_COMMIT\_IN\_TRX\_WHEN\_GTID\_NEXT\_IS\_SET [3711])  
Message: Cannot execute statements with implicit commit inside a transaction when @@SESSION.GTID\_NEXT != AUTOMATIC or @@SESSION.GTID\_NEXT\_LIST != NULL.
- Error: 1779 SQLSTATE: HY000 (ER\_GTID\_MODE\_2\_OR\_3\_REQUIRES\_ENFORCE\_GTID\_CONSISTENCY\_ON [3711])  
Message: @@GLOBAL.GTID\_MODE = ON or UPGRADE\_STEP\_2 requires @@GLOBAL.ENFORCE\_GTID\_CONSISTENCY = 1.
- Error: 1780 SQLSTATE: HY000 (ER\_GTID\_MODE\_REQUIRES\_BINLOG [3711])  
Message: @@GLOBAL.GTID\_MODE = ON or UPGRADE\_STEP\_1 or UPGRADE\_STEP\_2 requires --log-bin and --log-slave-updates.
- Error: 1781 SQLSTATE: HY000 (ER\_CANT\_SET\_GTID\_NEXT\_TO\_GTID\_WHEN\_GTID\_MODE\_IS\_OFF [3711])  
Message: @@SESSION.GTID\_NEXT cannot be set to UUID:NUMBER when @@GLOBAL.GTID\_MODE = OFF.
- Error: 1782 SQLSTATE: HY000 (ER\_CANT\_SET\_GTID\_NEXT\_TO\_ANONYMOUS\_WHEN\_GTID\_MODE\_IS\_ON [3711])  
Message: @@SESSION.GTID\_NEXT cannot be set to ANONYMOUS when @@GLOBAL.GTID\_MODE = ON.
- Error: 1783 SQLSTATE: HY000 (ER\_CANT\_SET\_GTID\_NEXT\_LIST\_TO\_NON\_NULL\_WHEN\_GTID\_MODE\_IS\_OFF [3711])  
Message: @@SESSION.GTID\_NEXT\_LIST cannot be set to a non-NULL value when @@GLOBAL.GTID\_MODE = OFF.
- Error: 1784 SQLSTATE: HY000 (ER\_FOUND\_GTID\_EVENT\_WHEN\_GTID\_MODE\_IS\_OFF [3711])  
Message: Found a Gtid\_log\_event or Previous\_gtids\_log\_event when @@GLOBAL.GTID\_MODE = OFF.



- Error: 1785 SQLSTATE: HY000 ([ER\\_GTID\\_UNSAFE\\_NON\\_TRANSACTIONAL\\_TABLE \[3712\]](#))

Message: When @@GLOBAL.ENFORCE\_GTID\_CONSISTENCY = 1, updates to non-transactional tables can only be done in either autocommitted statements or single-statement transactions, and never in the same statement as updates to transactional tables.

- Error: 1786 SQLSTATE: HY000 ([ER\\_GTID\\_UNSAFE\\_CREATE\\_SELECT \[3712\]](#))

Message: CREATE TABLE ... SELECT is forbidden when @@GLOBAL.ENFORCE\_GTID\_CONSISTENCY = 1.

- Error: 1787 SQLSTATE: HY000 ([ER\\_GTID\\_UNSAFE\\_CREATE\\_DROP\\_TEMPORARY\\_TABLE\\_IN\\_TRANSACTION \[3712\]](#))

Message: When @@GLOBAL.ENFORCE\_GTID\_CONSISTENCY = 1, the statements CREATE TEMPORARY TABLE and DROP TEMPORARY TABLE can be executed in a non-transactional context only, and require that AUTOCOMMIT = 1.

- Error: 1788 SQLSTATE: HY000 ([ER\\_GTID\\_MODE\\_CAN\\_ONLY\\_CHANGE\\_ONE\\_STEP\\_AT\\_A\\_TIME \[3712\]](#))

Message: The value of @@GLOBAL.GTID\_MODE can only change one step at a time: OFF <-> UPGRADE\_STEP\_1 <-> UPGRADE\_STEP\_2 <-> ON. Also note that this value must be stepped up or down simultaneously on all servers; see the Manual for instructions.

- Error: 1789 SQLSTATE: HY000 ([ER\\_MASTER\\_HAS\\_PURGED\\_REQUIRED\\_GTIDS \[3712\]](#))

Message: The slave is connecting using CHANGE MASTER TO MASTER\_AUTO\_POSITION = 1, but the master has purged binary logs containing GTIDs that the slave requires.

- Error: 1790 SQLSTATE: HY000 ([ER\\_CANT\\_SET\\_GTID\\_NEXT\\_WHEN\\_OWNING\\_GTID \[3712\]](#))

Message: @@SESSION.GTID\_NEXT cannot be changed by a client that owns a GTID. The client owns %s. Ownership is released on COMMIT or ROLLBACK.

- Error: 1791 SQLSTATE: HY000 ([ER\\_UNKNOWN\\_EXPLAIN\\_FORMAT \[3712\]](#))

Message: Unknown EXPLAIN format name: '%s'

- Error: 1792 SQLSTATE: 25006 ([ER\\_CANT\\_EXECUTE\\_IN\\_READ\\_ONLY\\_TRANSACTION \[3712\]](#))

Message: Cannot execute statement in a READ ONLY transaction.

- Error: 1793 SQLSTATE: HY000 ([ER\\_TOO\\_LONG\\_TABLE\\_PARTITION\\_COMMENT \[3712\]](#))

Message: Comment for table partition '%s' is too long (max = %lu)

- Error: 1794 SQLSTATE: HY000 ([ER\\_SLAVE\\_CONFIGURATION \[3712\]](#))

Message: Slave is not configured or failed to initialize properly. You must at least set --server-id to enable either a master or a slave. Additional error messages can be found in the MySQL error log.

- Error: 1795 SQLSTATE: HY000 ([ER\\_INNODB\\_FT\\_LIMIT \[3712\]](#))

Message: InnoDB presently supports one FULLTEXT index creation at a time

- Error: 1796 SQLSTATE: HY000 ([ER\\_INNODB\\_NO\\_FT\\_TEMP\\_TABLE \[3712\]](#))

Message: Cannot create FULLTEXT index on temporary InnoDB table

- Error: 1797 SQLSTATE: HY000 ([ER\\_INNODB\\_FT\\_WRONG\\_DOCID\\_COLUMN \[3712\]](#))

Message: Column '%s' is of wrong type for an InnoDB FULLTEXT index



- Error: [1798](#) SQLSTATE: [HY000](#) ([ER\\_INNOODB\\_FT\\_WRONG\\_DOCID\\_INDEX](#) [\[3713\]](#))  
Message: Index '%s' is of wrong type for an InnoDB FULLTEXT index
- Error: [1799](#) SQLSTATE: [HY000](#) ([ER\\_INNOODB\\_ONLINE\\_LOG\\_TOO\\_BIG](#) [\[3713\]](#))  
Message: Creating index '%s' required more than 'innodb\_online\_alter\_log\_max\_size' bytes of modification log. Please try again.
- Error: [1800](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_ALTER\\_ALGORITHM](#) [\[3713\]](#))  
Message: Unknown ALGORITHM '%s'
- Error: [1801](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_ALTER\\_LOCK](#) [\[3713\]](#))  
Message: Unknown LOCK type '%s'
- Error: [1802](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_CHANGE\\_MASTER\\_CANT\\_RUN\\_WITH\\_GAPS](#) [\[3713\]](#))  
Message: CHANGE MASTER cannot be executed when the slave was stopped with an error or killed in MTS mode. Consider using RESET SLAVE or START SLAVE UNTIL.
- Error: [1803](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_RECOVERY\\_FAILURE](#) [\[3713\]](#))  
Message: Cannot recover after SLAVE errored out in parallel execution mode. Additional error messages can be found in the MySQL error log.
- Error: [1804](#) SQLSTATE: [HY000](#) ([ER\\_MTS\\_RESET\\_WORKERS](#) [\[3713\]](#))  
Message: Cannot clean up worker info tables. Additional error messages can be found in the MySQL error log.
- Error: [1805](#) SQLSTATE: [HY000](#) ([ER\\_COL\\_COUNT\\_DOESNT\\_MATCH\\_CORRUPTED\\_V2](#) [\[3713\]](#))  
Message: Column count of %s.%s is wrong. Expected %d, found %d. The table is probably corrupted
- Error: [1806](#) SQLSTATE: [HY000](#) ([ER\\_SLAVE\\_SILENT\\_RETRY\\_TRANSACTION](#) [\[3713\]](#))  
Message: Slave must silently retry current transaction
- Error: [1807](#) SQLSTATE: [HY000](#) ([ER\\_DISCARD\\_FK\\_CHECKS\\_RUNNING](#) [\[3713\]](#))  
Message: There is a foreign key check running on table '%s'. Cannot discard the table.
- Error: [1808](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_SCHEMA\\_MISMATCH](#) [\[3713\]](#))  
Message: Schema mismatch (%s)
- Error: [1809](#) SQLSTATE: [HY000](#) ([ER\\_TABLE\\_IN\\_SYSTEM\\_TABLESPACE](#) [\[3713\]](#))  
Message: Table '%s' in system tablespace
- Error: [1810](#) SQLSTATE: [HY000](#) ([ER\\_IO\\_READ\\_ERROR](#) [\[3713\]](#))  
Message: IO Read error: (%lu, %s) %s
- Error: [1811](#) SQLSTATE: [HY000](#) ([ER\\_IO\\_WRITE\\_ERROR](#) [\[3713\]](#))  
Message: IO Write error: (%lu, %s) %s
- Error: [1812](#) SQLSTATE: [HY000](#) ([ER\\_TABLESPACE\\_MISSING](#) [\[3713\]](#))  
Message: Tablespace is missing for table '%s'

- Error: [1813](#) SQLSTATE: [HY000](#) ([ER\\_TABLESPACE\\_EXISTS](#) [\[3714\]](#))  
Message: Tablespace for table '%s' exists. Please DISCARD the tablespace before IMPORT.
- Error: [1814](#) SQLSTATE: [HY000](#) ([ER\\_TABLESPACE\\_DISCARDED](#) [\[3714\]](#))  
Message: Tablespace has been discarded for table '%s'
- Error: [1815](#) SQLSTATE: [HY000](#) ([ER\\_INTERNAL\\_ERROR](#) [\[3714\]](#))  
Message: Internal error: %s
- Error: [1816](#) SQLSTATE: [HY000](#) ([ER\\_INNODB\\_IMPORT\\_ERROR](#) [\[3714\]](#))  
Message: ALTER TABLE '%s' IMPORT TABLESPACE failed with error %lu : '%s'
- Error: [1817](#) SQLSTATE: [HY000](#) ([ER\\_INNODB\\_INDEX\\_CORRUPT](#) [\[3714\]](#))  
Message: Index corrupt: %s
- Error: [1818](#) SQLSTATE: [HY000](#) ([ER\\_INVALID\\_YEAR\\_COLUMN\\_LENGTH](#) [\[3714\]](#))  
Message: YEAR(%lu) column type is deprecated. Creating YEAR(4) column instead.
- Error: [1819](#) SQLSTATE: [HY000](#) ([ER\\_NOT\\_VALID\\_PASSWORD](#) [\[3714\]](#))  
Message: Your password does not satisfy the current policy requirements
- Error: [1820](#) SQLSTATE: [HY000](#) ([ER\\_MUST\\_CHANGE\\_PASSWORD](#) [\[3714\]](#))  
Message: You must SET PASSWORD before executing this statement
- Error: [1821](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_NO\\_INDEX\\_CHILD](#) [\[3714\]](#))  
Message: Failed to add the foreign key constraint. Missing index for constraint '%s' in the foreign table '%s'
- Error: [1822](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_NO\\_INDEX\\_PARENT](#) [\[3714\]](#))  
Message: Failed to add the foreign key constraint. Missing index for constraint '%s' in the referenced table '%s'
- Error: [1823](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_FAIL\\_ADD\\_SYSTEM](#) [\[3714\]](#))  
Message: Failed to add the foreign key constraint '%s' to system tables
- Error: [1824](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_CANNOT\\_OPEN\\_PARENT](#) [\[3714\]](#))  
Message: Failed to open the referenced table '%s'
- Error: [1825](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_INCORRECT\\_OPTION](#) [\[3714\]](#))  
Message: Failed to add the foreign key constraint on table '%s'. Incorrect options in FOREIGN KEY constraint '%s'
- Error: [1826](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_DUP\\_NAME](#) [\[3714\]](#))  
Message: Duplicate foreign key constraint name '%s'
- Error: [1827](#) SQLSTATE: [HY000](#) ([ER\\_PASSWORD\\_FORMAT](#) [\[3714\]](#))  
Message: The password hash doesn't have the expected format. Check if the correct password algorithm is being used with the PASSWORD() function.

- Error: [1828](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_COLUMN\\_CANNOT\\_DROP](#) [\[3715\]](#))  
Message: Cannot drop column '%s': needed in a foreign key constraint '%s'
- Error: [1829](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_COLUMN\\_CANNOT\\_DROP\\_CHILD](#) [\[3715\]](#))  
Message: Cannot drop column '%s': needed in a foreign key constraint '%s' of table '%s'
- Error: [1830](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_COLUMN\\_NOT\\_NULL](#) [\[3715\]](#))  
Message: Column '%s' cannot be NOT NULL: needed in a foreign key constraint '%s' SET NULL
- Error: [1831](#) SQLSTATE: [HY000](#) ([ER\\_DUP\\_INDEX](#) [\[3715\]](#))  
Message: Duplicate index '%s' defined on the table '%s.%s'. This is deprecated and will be disallowed in a future release.
- Error: [1832](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_COLUMN\\_CANNOT\\_CHANGE](#) [\[3715\]](#))  
Message: Cannot change column '%s': used in a foreign key constraint '%s'
- Error: [1833](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_COLUMN\\_CANNOT\\_CHANGE\\_CHILD](#) [\[3715\]](#))  
Message: Cannot change column '%s': used in a foreign key constraint '%s' of table '%s'
- Error: [1834](#) SQLSTATE: [HY000](#) ([ER\\_FK\\_CANNOT\\_DELETE\\_PARENT](#) [\[3715\]](#))  
Message: Cannot delete rows from table which is parent in a foreign key constraint '%s' of table '%s'
- Error: [1835](#) SQLSTATE: [HY000](#) ([ER\\_MALFORMED\\_PACKET](#) [\[3715\]](#))  
Message: Malformed communication packet.
- Error: [1836](#) SQLSTATE: [HY000](#) ([ER\\_READ\\_ONLY\\_MODE](#) [\[3715\]](#))  
Message: Running in read-only mode
- Error: [1837](#) SQLSTATE: [HY000](#) ([ER\\_GTID\\_NEXT\\_TYPE\\_UNDEFINED\\_GROUP](#) [\[3715\]](#))  
Message: When @@SESSION.GTID\_NEXT is set to a GTID, you must explicitly set it again after a COMMIT or ROLLBACK. If you see this error message in the slave SQL thread, it means that a table in the current transaction is transactional on the master and non-transactional on the slave. In a client connection, it means that you executed SET @@SESSION.GTID\_NEXT before a transaction and forgot to set @@SESSION.GTID\_NEXT to a different identifier or to 'AUTOMATIC' after COMMIT or ROLLBACK. Current @@SESSION.GTID\_NEXT is '%s'.
- Error: [1838](#) SQLSTATE: [HY000](#) ([ER\\_VARIABLE\\_NOT\\_SETTABLE\\_IN\\_SP](#) [\[3715\]](#))  
Message: The system variable %s cannot be set in stored procedures.
- Error: [1839](#) SQLSTATE: [HY000](#)  
([ER\\_CANT\\_SET\\_GTID\\_PURGED\\_WHEN\\_GTID\\_MODE\\_IS\\_OFF](#) [\[3715\]](#))  
Message: @@GLOBAL.GTID\_PURGED can only be set when @@GLOBAL.GTID\_MODE = ON.
- Error: [1840](#) SQLSTATE: [HY000](#)  
([ER\\_CANT\\_SET\\_GTID\\_PURGED\\_WHEN\\_GTID\\_EXECUTED\\_IS\\_NOT\\_EMPTY](#) [\[3715\]](#))  
Message: @@GLOBAL.GTID\_PURGED can only be set when @@GLOBAL.GTID\_EXECUTED is empty.
- Error: [1841](#) SQLSTATE: [HY000](#)  
([ER\\_CANT\\_SET\\_GTID\\_PURGED\\_WHEN\\_OWNED\\_GTIDS\\_IS\\_NOT\\_EMPTY](#) [\[3715\]](#))

Message: @@GLOBAL.GTID\_PURGED can only be set when there are no ongoing transactions (not even in other clients).

- Error: 1842 SQLSTATE: HY000 (ER\_GTID\_PURGED\_WAS\_CHANGED [3716])

Message: @@GLOBAL.GTID\_PURGED was changed from '%s' to '%s'.

- Error: 1843 SQLSTATE: HY000 (ER\_GTID\_EXECUTED\_WAS\_CHANGED [3716])

Message: @@GLOBAL.GTID\_EXECUTED was changed from '%s' to '%s'.

- Error: 1844 SQLSTATE: HY000 (ER\_BINLOG\_STMT\_MODE\_AND\_NO\_REPL\_TABLES [3716])

Message: Cannot execute statement: impossible to write to binary log since BINLOG\_FORMAT = STATEMENT, and both replicated and non replicated tables are written to.

- Error: 1845 SQLSTATE: 0A000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED [3716])

Message: %s is not supported for this operation. Try %s.

- Error: 1846 SQLSTATE: 0A000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON [3716])

Message: %s is not supported. Reason: %s. Try %s.

- Error: 1847 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_COPY [3716])

Message: COPY algorithm requires a lock

- Error: 1848 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_PARTITION [3716])

Message: Partition specific operations do not yet support LOCK/ALGORITHM

- Error: 1849 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_FK\_RENAME [3716])

Message: Columns participating in a foreign key are renamed

- Error: 1850 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_COLUMN\_TYPE [3716])

Message: Cannot change column type INPLACE

- Error: 1851 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_FK\_CHECK [3716])

Message: Adding foreign keys needs foreign\_key\_checks=OFF

- Error: 1852 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_IGNORE [3716])

Message: Creating unique indexes with IGNORE requires COPY algorithm to remove duplicate rows

- Error: 1853 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_NOPK [3716])

Message: Dropping a primary key is not allowed without also adding a new primary key

- Error: 1854 SQLSTATE: HY000 (ER\_ALTER\_OPERATION\_NOT\_SUPPORTED\_REASON\_AUTOINC [3716])

Message: Adding an auto-increment column requires a lock

- Error: 1855 SQLSTATE: HY000  
([ER\\_ALTER\\_OPERATION\\_NOT\\_SUPPORTED\\_REASON\\_HIDDEN\\_FTS \[3717\]](#))  
Message: Cannot replace hidden FTS\_DOC\_ID with a user-visible one
- Error: 1856 SQLSTATE: HY000  
([ER\\_ALTER\\_OPERATION\\_NOT\\_SUPPORTED\\_REASON\\_CHANGE\\_FTS \[3717\]](#))  
Message: Cannot drop or rename FTS\_DOC\_ID
- Error: 1857 SQLSTATE: HY000 ([ER\\_ALTER\\_OPERATION\\_NOT\\_SUPPORTED\\_REASON\\_FTS \[3717\]](#))  
Message: Fulltext index creation requires a lock
- Error: 1858 SQLSTATE: HY000  
([ER\\_SQL\\_SLAVE\\_SKIP\\_COUNTER\\_NOT\\_SETTABLE\\_IN\\_GTID\\_MODE \[3717\]](#))  
Message: sql\_slave\_skip\_counter can not be set when the server is running with @@GLOBAL.GTID\_MODE = ON. Instead, for each transaction that you want to skip, generate an empty transaction with the same GTID as the transaction
- Error: 1859 SQLSTATE: 23000 ([ER\\_DUP\\_UNKNOWN\\_IN\\_INDEX \[3717\]](#))  
Message: Duplicate entry for key '%s'
- Error: 1860 SQLSTATE: HY000 ([ER\\_IDENT\\_CAUSES\\_TOO\\_LONG\\_PATH \[3717\]](#))  
Message: Long database name and identifier for object resulted in path length exceeding %d characters. Path: '%s'.
- Error: 1861 SQLSTATE: HY000  
([ER\\_ALTER\\_OPERATION\\_NOT\\_SUPPORTED\\_REASON\\_NOT\\_NULL \[3717\]](#))  
Message: cannot silently convert NULL values, as required in this SQL\_MODE
- Error: 1862 SQLSTATE: HY000 ([ER\\_MUST\\_CHANGE\\_PASSWORD\\_LOGIN \[3717\]](#))  
Message: Your password has expired. To log in you must change it using a client that supports expired passwords.
- Error: 1863 SQLSTATE: HY000 ([ER\\_ROW\\_IN\\_WRONG\\_PARTITION \[3717\]](#))  
Message: Found a row in wrong partition %s
- Error: 1864 SQLSTATE: HY000 ([ER\\_MTS\\_EVENT\\_BIGGER\\_PENDING\\_JOBS\\_SIZE\\_MAX \[3717\]](#))  
Message: Cannot schedule event %s, relay-log name %s, position %s to Worker thread because its size %lu exceeds %lu of slave\_pending\_jobs\_size\_max.
- Error: 1865 SQLSTATE: HY000 ([ER\\_INNODB\\_NO\\_FT\\_USES\\_PARSER \[3717\]](#))  
Message: Cannot CREATE FULLTEXT INDEX WITH PARSER on InnoDB table
- Error: 1866 SQLSTATE: HY000 ([ER\\_BINLOG\\_LOGICAL\\_CORRUPTION \[3717\]](#))  
Message: The binary log file '%s' is logically corrupted: %s
- Error: 1867 SQLSTATE: HY000 ([ER\\_WARN\\_PURGE\\_LOG\\_IN\\_USE \[3717\]](#))  
Message: file %s was not purged because it was being read by %d thread(s), purged only %d out of %d files.
- Error: 1868 SQLSTATE: HY000 ([ER\\_WARN\\_PURGE\\_LOG\\_IS\\_ACTIVE \[3717\]](#))

Message: file %s was not purged because it is the active log file.

## C.4. Client Error Codes and Messages

Client error information comes from the following source files:

- The Error values and the symbols in parentheses correspond to definitions in the `include/errmsg.h` MySQL source file.
- The Message values correspond to the error messages that are listed in the `libmysql/errmsg.c` file. `%d` and `%s` represent numbers and strings, respectively, that are substituted into the messages when they are displayed.

Because updates are frequent, it is possible that those files will contain additional error information not listed here.

- Error: [2000 \(CR\\_UNKNOWN\\_ERROR \[3718\]\)](#)  
Message: Unknown MySQL error
- Error: [2001 \(CR\\_SOCKET\\_CREATE\\_ERROR \[3718\]\)](#)  
Message: Can't create UNIX socket (%d)
- Error: [2002 \(CR\\_CONNECTION\\_ERROR \[3718\]\)](#)  
Message: Can't connect to local MySQL server through socket '%s' (%d)
- Error: [2003 \(CR\\_CONN\\_HOST\\_ERROR \[3718\]\)](#)  
Message: Can't connect to MySQL server on '%s' (%d)
- Error: [2004 \(CR\\_IPSOCK\\_ERROR \[3718\]\)](#)  
Message: Can't create TCP/IP socket (%d)
- Error: [2005 \(CR\\_UNKNOWN\\_HOST \[3718\]\)](#)  
Message: Unknown MySQL server host '%s' (%d)
- Error: [2006 \(CR\\_SERVER\\_GONE\\_ERROR \[3718\]\)](#)  
Message: MySQL server has gone away
- Error: [2007 \(CR\\_VERSION\\_ERROR \[3718\]\)](#)  
Message: Protocol mismatch; server version = %d, client version = %d
- Error: [2008 \(CR\\_OUT\\_OF\\_MEMORY \[3718\]\)](#)  
Message: MySQL client ran out of memory
- Error: [2009 \(CR\\_WRONG\\_HOST\\_INFO \[3718\]\)](#)  
Message: Wrong host info
- Error: [2010 \(CR\\_LOCALHOST\\_CONNECTION \[3718\]\)](#)  
Message: Localhost via UNIX socket
- Error: [2011 \(CR\\_TCP\\_CONNECTION \[3718\]\)](#)  
Message: %s via TCP/IP

- Error: [2012 \(CR\\_SERVER\\_HANDSHAKE\\_ERR \[3719\]\)](#)  
Message: Error in server handshake
- Error: [2013 \(CR\\_SERVER\\_LOST \[3719\]\)](#)  
Message: Lost connection to MySQL server during query
- Error: [2014 \(CR\\_COMMANDS\\_OUT\\_OF\\_SYNC \[3719\]\)](#)  
Message: Commands out of sync; you can't run this command now
- Error: [2015 \(CR\\_NAMEDPIPE\\_CONNECTION \[3719\]\)](#)  
Message: Named pipe: %s
- Error: [2016 \(CR\\_NAMEDPIPEWAIT\\_ERROR \[3719\]\)](#)  
Message: Can't wait for named pipe to host: %s pipe: %s (%lu)
- Error: [2017 \(CR\\_NAMEDPIPEOPEN\\_ERROR \[3719\]\)](#)  
Message: Can't open named pipe to host: %s pipe: %s (%lu)
- Error: [2018 \(CR\\_NAMEDPIPESETSTATE\\_ERROR \[3719\]\)](#)  
Message: Can't set state of named pipe to host: %s pipe: %s (%lu)
- Error: [2019 \(CR\\_CANT\\_READ\\_CHARSET \[3719\]\)](#)  
Message: Can't initialize character set %s (path: %s)
- Error: [2020 \(CR\\_NET\\_PACKET\\_TOO\\_LARGE \[3719\]\)](#)  
Message: Got packet bigger than 'max\_allowed\_packet' bytes
- Error: [2021 \(CR\\_EMBEDDED\\_CONNECTION \[3719\]\)](#)  
Message: Embedded server
- Error: [2022 \(CR\\_PROBE\\_SLAVE\\_STATUS \[3719\]\)](#)  
Message: Error on SHOW SLAVE STATUS:
- Error: [2023 \(CR\\_PROBE\\_SLAVE\\_HOSTS \[3719\]\)](#)  
Message: Error on SHOW SLAVE HOSTS:
- Error: [2024 \(CR\\_PROBE\\_SLAVE\\_CONNECT \[3719\]\)](#)  
Message: Error connecting to slave:
- Error: [2025 \(CR\\_PROBE\\_MASTER\\_CONNECT \[3719\]\)](#)  
Message: Error connecting to master:
- Error: [2026 \(CR\\_SSL\\_CONNECTION\\_ERROR \[3719\]\)](#)  
Message: SSL connection error: %s
- Error: [2027 \(CR\\_MALFORMED\\_PACKET \[3719\]\)](#)  
Message: Malformed packet
- Error: [2028 \(CR\\_WRONG\\_LICENSE \[3719\]\)](#)

Message: This client library is licensed only for use with MySQL servers having '%s' license

- Error: 2029 (CR\_NULL\_POINTER [3720])

Message: Invalid use of null pointer

- Error: 2030 (CR\_NO\_PREPARE\_STMT [3720])

Message: Statement not prepared

- Error: 2031 (CR\_PARAMS\_NOT\_BOUND [3720])

Message: No data supplied for parameters in prepared statement

- Error: 2032 (CR\_DATA\_TRUNCATED [3720])

Message: Data truncated

- Error: 2033 (CR\_NO\_PARAMETERS\_EXISTS [3720])

Message: No parameters exist in the statement

- Error: 2034 (CR\_INVALID\_PARAMETER\_NO [3720])

Message: Invalid parameter number

- Error: 2035 (CR\_INVALID\_BUFFER\_USE [3720])

Message: Can't send long data for non-string/non-binary data types (parameter: %d)

- Error: 2036 (CR\_UNSUPPORTED\_PARAM\_TYPE [3720])

Message: Using unsupported buffer type: %d (parameter: %d)

- Error: 2037 (CR\_SHARED\_MEMORY\_CONNECTION [3720])

Message: Shared memory: %s

- Error: 2038 (CR\_SHARED\_MEMORY\_CONNECT\_REQUEST\_ERROR [3720])

Message: Can't open shared memory; client could not create request event (%lu)

- Error: 2039 (CR\_SHARED\_MEMORY\_CONNECT\_ANSWER\_ERROR [3720])

Message: Can't open shared memory; no answer event received from server (%lu)

- Error: 2040 (CR\_SHARED\_MEMORY\_CONNECT\_FILE\_MAP\_ERROR [3720])

Message: Can't open shared memory; server could not allocate file mapping (%lu)

- Error: 2041 (CR\_SHARED\_MEMORY\_CONNECT\_MAP\_ERROR [3720])

Message: Can't open shared memory; server could not get pointer to file mapping (%lu)

- Error: 2042 (CR\_SHARED\_MEMORY\_FILE\_MAP\_ERROR [3720])

Message: Can't open shared memory; client could not allocate file mapping (%lu)

- Error: 2043 (CR\_SHARED\_MEMORY\_MAP\_ERROR [3720])

Message: Can't open shared memory; client could not get pointer to file mapping (%lu)

- Error: 2044 (CR\_SHARED\_MEMORY\_EVENT\_ERROR [3720])



Message: Can't open shared memory; client could not create %s event (%lu)

- Error: 2045 (CR\_SHARED\_MEMORY\_CONNECT\_ABANDONED\_ERROR [3721])

Message: Can't open shared memory; no answer from server (%lu)

- Error: 2046 (CR\_SHARED\_MEMORY\_CONNECT\_SET\_ERROR [3721])

Message: Can't open shared memory; cannot send request event to server (%lu)

- Error: 2047 (CR\_CONN\_UNKNOW\_PROTOCOL [3721])

Message: Wrong or unknown protocol

- Error: 2048 (CR\_INVALID\_CONN\_HANDLE [3721])

Message: Invalid connection handle

- Error: 2049 (CR\_SECURE\_AUTH [3721])

Message: Connection using old (pre-4.1.1) authentication protocol refused (client option 'secure\_auth' enabled)

- Error: 2050 (CR\_FETCH\_CANCELED [3721])

Message: Row retrieval was canceled by mysql\_stmt\_close() call

- Error: 2051 (CR\_NO\_DATA [3721])

Message: Attempt to read column without prior row fetch

- Error: 2052 (CR\_NO\_STMT\_METADATA [3721])

Message: Prepared statement contains no metadata

- Error: 2053 (CR\_NO\_RESULT\_SET [3721])

Message: Attempt to read a row while there is no result set associated with the statement

- Error: 2054 (CR\_NOT\_IMPLEMENTED [3721])

Message: This feature is not implemented yet

- Error: 2055 (CR\_SERVER\_LOST\_EXTENDED [3721])

Message: Lost connection to MySQL server at '%s', system error: %d

- Error: 2056 (CR\_STMT\_CLOSED [3721])

Message: Statement closed indirectly because of a preceding %s() call

- Error: 2057 (CR\_NEW\_STMT\_METADATA [3721])

Message: The number of columns in the result set differs from the number of bound buffers. You must reset the statement, rebind the result set columns, and execute the statement again

- Error: 2058 (CR\_ALREADY\_CONNECTED [3721])

Message: This handle is already connected. Use a separate handle for each connection.

- Error: 2059 (CR\_AUTH\_PLUGIN\_CANNOT\_LOAD [3721])

Message: Authentication plugin '%s' cannot be loaded: %s

- Error: 2060 ([CR\\_DUPLICATE\\_CONNECTION\\_ATTR \[3722\]](#))

Message: There is an attribute with the same name already

- Error: 2061 ([CR\\_AUTH\\_PLUGIN\\_ERR \[3722\]](#))

Message: Authentication plugin '%s' reported error: %s

## C.5. Problems and Common Errors

This section lists some common problems and error messages that you may encounter. It describes how to determine the causes of the problems and what to do to solve them.

### C.5.1. How to Determine What Is Causing a Problem

When you run into a problem, the first thing you should do is to find out which program or piece of equipment is causing it:

- If you have one of the following symptoms, then it is probably a hardware problems (such as memory, motherboard, CPU, or hard disk) or kernel problem:
  - The keyboard does not work. This can normally be checked by pressing the Caps Lock key. If the Caps Lock light does not change, you have to replace your keyboard. (Before doing this, you should try to restart your computer and check all cables to the keyboard.)
  - The mouse pointer does not move.
  - The machine does not answer to a remote machine's pings.
  - Other programs that are not related to MySQL do not behave correctly.
  - Your system restarted unexpectedly. (A faulty user-level program should never be able to take down your system.)

In this case, you should start by checking all your cables and run some diagnostic tool to check your hardware! You should also check whether there are any patches, updates, or service packs for your operating system that could likely solve your problem. Check also that all your libraries (such as [glibc](#)) are up to date.

It is always good to use a machine with ECC memory to discover memory problems early.

- If your keyboard is locked up, you may be able to recover by logging in to your machine from another machine and executing `kbd_mode -a`.
- Please examine your system log file (`/var/log/messages` or similar) for reasons for your problem. If you think the problem is in MySQL, you should also examine MySQL's log files. See [Section 5.2](#), "MySQL Server Logs".
- If you do not think you have hardware problems, you should try to find out which program is causing problems. Try using `top`, `ps`, Task Manager, or some similar program, to check which program is taking all CPU or is locking the machine.
- Use `top`, `df`, or a similar program to check whether you are out of memory, disk space, file descriptors, or some other critical resource.
- If the problem is some runaway process, you can always try to kill it. If it does not want to die, there is probably a bug in the operating system.

If after you have examined all other possibilities and you have concluded that the MySQL server or a MySQL client is causing the problem, it is time to create a bug report for our mailing list or our support team. In the bug report, try to give a very detailed description of how the system is behaving and what

you think is happening. You should also state why you think that MySQL is causing the problem. Take into consideration all the situations in this chapter. State any problems exactly how they appear when you examine your system. Use the “copy and paste” method for any output and error messages from programs and log files.

Try to describe in detail which program is not working and all symptoms you see. We have in the past received many bug reports that state only “the system does not work.” This provides us with no information about what could be the problem.

If a program fails, it is always useful to know the following information:

- Has the program in question made a segmentation fault (did it dump core)?
- Is the program taking up all available CPU time? Check with `top`. Let the program run for a while, it may simply be evaluating something computationally intensive.
- If the `mysqld` server is causing problems, can you get any response from it with `mysqladmin -u root ping` or `mysqladmin -u root processlist`?
- What does a client program say when you try to connect to the MySQL server? (Try with `mysql`, for example.) Does the client jam? Do you get any output from the program?

When sending a bug report, you should follow the outline described in [Section 1.7, “How to Report Bugs or Problems”](#).

## C.5.2. Common Errors When Using MySQL Programs

This section lists some errors that users frequently encounter when running MySQL programs. Although the problems show up when you try to run client programs, the solutions to many of the problems involves changing the configuration of the MySQL server.

### C.5.2.1. Access denied

An `Access denied` error can have many causes. Often the problem is related to the MySQL accounts that the server permits client programs to use when connecting. See [Section 6.2, “The MySQL Access Privilege System”](#), and [Section 6.2.7, “Causes of Access-Denied Errors”](#).

### C.5.2.2. Can't connect to [local] MySQL server

A MySQL client on Unix can connect to the `mysqld` server in two different ways: By using a Unix socket file to connect through a file in the file system (default `/tmp/mysql.sock`), or by using TCP/IP, which connects through a port number. A Unix socket file connection is faster than TCP/IP, but can be used only when connecting to a server on the same computer. A Unix socket file is used if you do not specify a host name or if you specify the special host name `localhost`.

If the MySQL server is running on Windows, you can connect using TCP/IP. If the server is started with the `--enable-named-pipe [421]` option, you can also connect with named pipes if you run the client on the host where the server is running. The name of the named pipe is `MySQL` by default. If you do not give a host name when connecting to `mysqld`, a MySQL client first tries to connect to the named pipe. If that does not work, it connects to the TCP/IP port. You can force the use of named pipes on Windows by using `.` as the host name.

The error (2002) `Can't connect to ...` normally means that there is no MySQL server running on the system or that you are using an incorrect Unix socket file name or TCP/IP port number when trying to connect to the server. You should also check that the TCP/IP port you are using has not been blocked by a firewall or port blocking service.

The error (2003) `Can't connect to MySQL server on 'server' (10061)` indicates that the network connection has been refused. You should check that there is a MySQL server running, that it has network connections enabled, and that the network port you specified is the one configured on the server.

Start by checking whether there is a process named `mysqld` running on your server host. (Use `ps xa | grep mysqld` on Unix or the Task Manager on Windows.) If there is no such process, you should start the server. See [Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”](#).

If a `mysqld` process is running, you can check it by trying the following commands. The port number or Unix socket file name might be different in your setup. `host_ip` represents the IP address of the machine where the server is running.

```
shell> mysqladmin version
shell> mysqladmin variables
shell> mysqladmin -h `hostname` version variables
shell> mysqladmin -h `hostname` --port=3306 version
shell> mysqladmin -h host_ip version
shell> mysqladmin --protocol=SOCKET --socket=/tmp/mysql.sock version
```

Note the use of backticks rather than forward quotation marks with the `hostname` command; these cause the output of `hostname` (that is, the current host name) to be substituted into the `mysqladmin` command. If you have no `hostname` command or are running on Windows, you can manually type the host name of your machine (without backticks) following the `-h` option. You can also try `-h 127.0.0.1` to connect with TCP/IP to the local host.

Make sure that the server has not been configured to ignore network connections or (if you are attempting to connect remotely) that it has not been configured to listen only locally on its network interfaces. If the server was started with `--skip-networking` [440], it will not accept TCP/IP connections at all. If the server was started with `--bind-address=127.0.0.1` [414], it will listen for TCP/IP connections only locally on the loopback interface and will not accept remote connections.

Check to make sure that there is no firewall blocking access to MySQL. Your firewall may be configured on the basis of the application being executed, or the port number used by MySQL for communication (3306 by default). Under Linux or Unix, check your IP tables (or similar) configuration to ensure that the port has not been blocked. Under Windows, applications such as ZoneAlarm or the Windows XP personal firewall may need to be configured not to block the MySQL port.

Here are some reasons the [Can't connect to local MySQL server](#) error might occur:

- `mysqld` is not running on the local host. Check your operating system's process list to ensure the `mysqld` process is present.
- You're running a MySQL server on Windows with many TCP/IP connections to it. If you're experiencing that quite often your clients get that error, you can find a workaround here: [Section C.5.2.2.1, “Connection to MySQL Server Failing on Windows”](#).
- Someone has removed the Unix socket file that `mysqld` uses (`/tmp/mysql.sock` by default). For example, you might have a `cron` job that removes old files from the `/tmp` directory. You can always run `mysqladmin version` to check whether the Unix socket file that `mysqladmin` is trying to use really exists. The fix in this case is to change the `cron` job to not remove `mysql.sock` or to place the socket file somewhere else. See [Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”](#).
- You have started the `mysqld` server with the `--socket=/path/to/socket` [442] option, but forgotten to tell client programs the new name of the socket file. If you change the socket path name for the server, you must also notify the MySQL clients. You can do this by providing the same `--socket` [211] option when you run client programs. You also need to ensure that clients have permission to access the `mysql.sock` file. To find out where the socket file is, you can do:

```
shell> netstat -ln | grep mysql
```

See [Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”](#).

- You are using Linux and one server thread has died (dumped core). In this case, you must kill the other `mysqld` threads (for example, with `kill` or with the `mysql_zap` script) before you can restart the MySQL server. See [Section C.5.4.2, “What to Do If MySQL Keeps Crashing”](#).

- The server or client program might not have the proper access privileges for the directory that holds the Unix socket file or the socket file itself. In this case, you must either change the access privileges for the directory or socket file so that the server and clients can access them, or restart `mysqld` with a `--socket [442]` option that specifies a socket file name in a directory where the server can create it and where client programs can access it.

If you get the error message `Can't connect to MySQL server on some_host`, you can try the following things to find out what the problem is:

- Check whether the server is running on that host by executing `telnet some_host 3306` and pressing the Enter key a couple of times. (3306 is the default MySQL port number. Change the value if your server is listening to a different port.) If there is a MySQL server running and listening to the port, you should get a response that includes the server's version number. If you get an error such as `telnet: Unable to connect to remote host: Connection refused`, then there is no server running on the given port.
- If the server is running on the local host, try using `mysqladmin -h localhost variables` to connect using the Unix socket file. Verify the TCP/IP port number that the server is configured to listen to (it is the value of the `port [525]` variable.)
- If you are running under Linux and Security-Enhanced Linux (SELinux) is enabled, make sure you have disabled SELinux protection for the `mysqld` process.

#### C.5.2.2.1. Connection to MySQL Server Failing on Windows

When you're running a MySQL server on Windows with many TCP/IP connections to it, and you're experiencing that quite often your clients get a `Can't connect to MySQL server` error, the reason might be that Windows does not allow for enough ephemeral (short-lived) ports to serve those connections.

The purpose of `TIME_WAIT` is to keep a connection accepting packets even after the connection has been closed. This is because Internet routing can cause a packet to take a slow route to its destination and it may arrive after both sides have agreed to close. If the port is in use for a new connection, that packet from the old connection could break the protocol or compromise personal information from the original connection. The `TIME_WAIT` delay prevents this by ensuring that the port cannot be reused until after some time has been permitted for those delayed packets to arrive.

It is safe to reduce `TIME_WAIT` greatly on LAN connections because there is little chance of packets arriving at very long delays, as they could through the Internet with its comparatively large distances and latencies.

Windows permits ephemeral (short-lived) TCP ports to the user. After any port is closed it will remain in a `TIME_WAIT` status for 120 seconds. The port will not be available again until this time expires. The default range of port numbers depends on the version of Windows, with a more limited number of ports in older versions:

- Windows through Server 2003: Ports in range 1025–5000
- Windows Vista, Server 2008, and newer: Ports in range 49152–65535

With a small stack of available TCP ports (5000) and a high number of TCP ports being open and closed over a short period of time along with the `TIME_WAIT` status you have a good chance for running out of ports. There are two ways to address this problem:

- Reduce the number of TCP ports consumed quickly by investigating connection pooling or persistent connections where possible
- Tune some settings in the Windows registry (see below)

**IMPORTANT: The following procedure involves modifying the Windows registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore,**

and edit the registry, view the following article in the Microsoft Knowledge Base: <http://support.microsoft.com/kb/256986/EN-US/>.

1. Start Registry Editor (`Regedt32.exe`).
2. Locate the following key in the registry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value Name: MaxUserPort
Data Type: REG_DWORD
Value: 65534
```

This sets the number of ephemeral ports available to any user. The valid range is between 5000 and 65534 (decimal). The default value is 0x1388 (5000 decimal).

4. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value Name: TcpTimedWaitDelay
Data Type: REG_DWORD
Value: 30
```

This sets the number of seconds to hold a TCP port connection in `TIME_WAIT` state before closing. The valid range is between 0 (zero) and 300 (decimal). The default value is 0x78 (120 decimal).

5. Quit Registry Editor.
6. Reboot the machine.

Note: Undoing the above should be as simple as deleting the registry entries you've created.

### C.5.2.3. Lost connection to MySQL server

There are three likely causes for this error message.

Usually it indicates network connectivity trouble and you should check the condition of your network if this error occurs frequently. If the error message includes “during query,” this is probably the case you are experiencing.

Sometimes the “during query” form happens when millions of rows are being sent as part of one or more queries. If you know that this is happening, you should try increasing `net_read_timeout` [515] from its default of 30 seconds to 60 seconds or longer, sufficient for the data transfer to complete.

More rarely, it can happen when the client is attempting the initial connection to the server. In this case, if your `connect_timeout` [472] value is set to only a few seconds, you may be able to resolve the problem by increasing it to ten seconds, perhaps more if you have a very long distance or slow connection. You can determine whether you are experiencing this more uncommon cause by using `SHOW GLOBAL STATUS LIKE 'Aborted_connects'`. It will increase by one for each initial connection attempt that the server aborts. You may see “reading authorization packet” as part of the error message; if so, that also suggests that this is the solution that you need.

If the cause is none of those just described, you may be experiencing a problem with `BLOB` values that are larger than `max_allowed_packet` [502], which can cause this error with some clients. Sometime you may see an `ER_NET_PACKET_TOO_LARGE` [3670] error, and that confirms that you need to increase `max_allowed_packet` [502].

### C.5.2.4. Client does not support authentication protocol

The current implementation of the authentication protocol uses a password hashing algorithm that is incompatible with that used by older (pre-4.1) clients. Attempts to connect to a 4.1 or newer server with an older client may fail with the following message:



```
shell> mysql
Client does not support authentication protocol requested
by server; consider upgrading MySQL client
```

To deal with this problem, the preferred solution is to upgrade all client programs to use a 4.1.1 or newer client library. If that is not possible, use one of the following approaches:

- To connect to the server with a pre-4.1 client program, use an account that still has a pre-4.1-style password.
- Reset the password to pre-4.1 style for each user that needs to use a pre-4.1 client program. This can be done using the `SET PASSWORD` statement and the `OLD_PASSWORD()` [1263] function. As of MySQL 5.6.6, it is also necessary to first ensure that the authentication plugin for the account is `mysql_old_password`:

```
mysql> UPDATE mysql.user SET plugin = 'mysql_old_password'
mysql> WHERE User = 'some_user' AND Host = 'some_host';
mysql> FLUSH PRIVILEGES;
mysql> SET PASSWORD FOR
-> 'some_user'@'some_host' = OLD_PASSWORD('newpwd');
```

Substitute the password you want to use for “*newpwd*” in the preceding example. MySQL cannot tell you what the original password was, so you'll need to pick a new one.

- Tell the server to use the older password hashing algorithm by default:
  1. Start `mysqld` with the `old_passwords` [517] system variable set to 1.
  2. Assign an old-format password to each account that has had its password updated to the longer 4.1 format. You can identify these accounts with the following query:

```
mysql> SELECT Host, User, Password FROM mysql.user
-> WHERE LENGTH>Password) > 16;
```

For each account record displayed by the query, use the `Host` and `User` values and assign a password using one of the methods described previously.

The `Client does not support authentication protocol` error also can occur if multiple versions of MySQL are installed but client programs are dynamically linked and link to an older library. Make sure that clients use the most recent library version with which they are compatible. The procedure to do this will depend on your system.



#### Note

The `mysql` extension does not support the authentication protocol in MySQL 4.1.1 and higher. This is true regardless of the PHP version being used. If you wish to use the `mysql` extension with MySQL 4.1 or newer, you may need to follow one of the options discussed above for configuring MySQL to work with old clients. The `mysqli` extension (stands for “MySQL, Improved”; added in PHP 5) is compatible with the improved password hashing employed in MySQL 4.1 and higher, and no special configuration of MySQL need be done to use this MySQL client library. For more information about the `mysqli` extension, see <http://php.net/mysqli>.

For additional background on password hashing and authentication, see [Section 6.1.2.4, “Password Hashing in MySQL”](#).

### C.5.2.5. Password Fails When Entered Interactively

MySQL client programs prompt for a password when invoked with a `--password` [211] or `-p` option that has no following password value:

```
shell> mysql -u user_name -p
```

```
Enter password:
```

On some systems, you may find that your password works when specified in an option file or on the command line, but not when you enter it interactively at the `Enter password:` prompt. This occurs when the library provided by the system to read passwords limits password values to a small number of characters (typically eight). That is a problem with the system library, not with MySQL. To work around it, change your MySQL password to a value that is eight or fewer characters long, or put your password in an option file.

### C.5.2.6. Host '*host\_name*' is blocked

If the following error occurs, it means that `mysqld` has received many connection requests from the given host that were interrupted in the middle:

```
Host 'host_name' is blocked because of many connection errors.
Unblock with 'mysqladmin flush-hosts'
```

The value of the `max_connect_errors` [502] system variable determines how many successive interrupted connection requests are permitted. (See [Section 5.1.4, "Server System Variables"](#).) After `max_connect_errors` [502] failed requests without a successful connection, `mysqld` assumes that something is wrong (for example, that someone is trying to break in), and blocks the host from further connections until you issue a `FLUSH HOSTS` statement or execute a `mysqladmin flush-hosts` command.

By default, `mysqld` blocks a host after 10 connection errors. You can adjust the value by setting `max_connect_errors` [502] at server startup:

```
shell> mysqld_safe --max_connect_errors=10000 &
```

The value can also be set at runtime:

```
mysql> SET GLOBAL max_connect_errors=10000;
```

If you get the `Host 'host_name' is blocked` error message for a given host, you should first verify that there is nothing wrong with TCP/IP connections from that host. If you are having network problems, it does you no good to increase the value of the `max_connect_errors` [502] variable.

### C.5.2.7. Too many connections

If you get a `Too many connections` error when you try to connect to the `mysqld` server, this means that all available connections are in use by other clients.

The number of connections permitted is controlled by the `max_connections` [503] system variable. The default value is 151 to improve performance when MySQL is used with the Apache Web server. (Previously, the default was 100.) If you need to support more connections, you should set a larger value for this variable.

`mysqld` actually permits `max_connections+1` [503] clients to connect. The extra connection is reserved for use by accounts that have the `SUPER` [756] privilege. By granting the `SUPER` [756] privilege to administrators and not to normal users (who should not need it), an administrator can connect to the server and use `SHOW PROCESSLIST` to diagnose problems even if the maximum number of unprivileged clients are connected. See [Section 13.7.5.30, "SHOW PROCESSLIST Syntax"](#).

The maximum number of connections MySQL can support depends on the quality of the thread library on a given platform, the amount of RAM available, how much RAM is used for each connection, the workload from each connection, and the desired response time. Linux or Solaris should be able to support at 500 to 1000 simultaneous connections routinely and as many as 10,000 connections if you have many gigabytes of RAM available and the workload from each is low or the response time target undemanding. Windows is limited to (open tables × 2 + open connections) < 2048 due to the Posix compatibility layer used on that platform.



Increasing `open-files-limit` [434] may be necessary. Also see [Section 2.5, “Installing MySQL on Linux”](#), for how to raise the operating system limit on how many handles can be used by MySQL.

### C.5.2.8. Out of memory

If you issue a query using the `mysql` client program and receive an error like the following one, it means that `mysql` does not have enough memory to store the entire query result:

```
mysql: Out of memory at line 42, 'malloc.c'
mysql: needed 8136 byte (8k), memory in use: 12481367 bytes (12189k)
ERROR 2008: MySQL client ran out of memory
```

To remedy the problem, first check whether your query is correct. Is it reasonable that it should return so many rows? If not, correct the query and try again. Otherwise, you can invoke `mysql` with the `--quick` [252] option. This causes it to use the `mysql_use_result()` C API function to retrieve the result set, which places less of a load on the client (but more on the server).

### C.5.2.9. MySQL server has gone away

This section also covers the related `Lost connection to server during query` error.

The most common reason for the `MySQL server has gone away` error is that the server timed out and closed the connection. In this case, you normally get one of the following error codes (which one you get is operating system-dependent).

Error Code	Description
<code>CR_SERVER_GONE_ERROR</code> [3718]	The client couldn't send a question to the server.
<code>CR_SERVER_LOST</code> [3719]	The client didn't get an error when writing to the server, but it didn't get a full answer (or any answer) to the question.

By default, the server closes the connection after eight hours if nothing has happened. You can change the time limit by setting the `wait_timeout` [563] variable when you start `mysqld`. See [Section 5.1.4, “Server System Variables”](#).

If you have a script, you just have to issue the query again for the client to do an automatic reconnection. This assumes that you have automatic reconnection in the client enabled (which is the default for the `mysql` command-line client).

Some other common reasons for the `MySQL server has gone away` error are:

- You (or the db administrator) has killed the running thread with a `KILL` statement or a `mysqladmin kill` command.
- You tried to run a query after closing the connection to the server. This indicates a logic error in the application that should be corrected.
- A client application running on a different host does not have the necessary privileges to connect to the MySQL server from that host.
- You got a timeout from the TCP/IP connection on the client side. This may happen if you have been using the commands: `mysql_options(..., MYSQL_OPT_READ_TIMEOUT,...)` or `mysql_options(..., MYSQL_OPT_WRITE_TIMEOUT,...)`. In this case increasing the timeout may help solve the problem.
- You have encountered a timeout on the server side and the automatic reconnection in the client is disabled (the `reconnect` flag in the `MYSQL` structure is equal to 0).
- You are using a Windows client and the server had dropped the connection (probably because `wait_timeout` [563] expired) before the command was issued.

The problem on Windows is that in some cases MySQL does not get an error from the OS when writing to the TCP/IP connection to the server, but instead gets the error when trying to read the answer from the connection.

The solution to this is to either do a `mysql_ping()` on the connection if there has been a long time since the last query (this is what Connector/ODBC does) or set `wait_timeout` [563] on the `mysqld` server so high that it in practice never times out.

- You can also get these errors if you send a query to the server that is incorrect or too large. If `mysqld` receives a packet that is too large or out of order, it assumes that something has gone wrong with the client and closes the connection. If you need big queries (for example, if you are working with big `BLOB` columns), you can increase the query limit by setting the server's `max_allowed_packet` [502] variable, which has a default value of 1MB. You may also need to increase the maximum packet size on the client end. More information on setting the packet size is given in [Section C.5.2.10, "Packet Too Large"](#).

An `INSERT` or `REPLACE` statement that inserts a great many rows can also cause these sorts of errors. Either one of these statements sends a single request to the server irrespective of the number of rows to be inserted; thus, you can often avoid the error by reducing the number of rows sent per `INSERT` or `REPLACE`.

- You also get a lost connection if you are sending a packet 16MB or larger if your client is older than 4.0.8 and your server is 4.0.8 and above, or the other way around.
- It is also possible to see this error if host name lookups fail (for example, if the DNS server on which your server or network relies goes down). This is because MySQL is dependent on the host system for name resolution, but has no way of knowing whether it is working—from MySQL's point of view the problem is indistinguishable from any other network timeout.

You may also see the `MySQL server has gone away` error if MySQL is started with the `--skip-networking` [440] option.

Another networking issue that can cause this error occurs if the MySQL port (default 3306) is blocked by your firewall, thus preventing any connections at all to the MySQL server.

- You can also encounter this error with applications that fork child processes, all of which try to use the same connection to the MySQL server. This can be avoided by using a separate connection for each child process.
- You have encountered a bug where the server died while executing the query.

You can check whether the MySQL server died and restarted by executing `mysqladmin version` and examining the server's uptime. If the client connection was broken because `mysqld` crashed and restarted, you should concentrate on finding the reason for the crash. Start by checking whether issuing the query again kills the server again. See [Section C.5.4.2, "What to Do If MySQL Keeps Crashing"](#).

You can get more information about the lost connections by starting `mysqld` with the `--log-warnings=2` [430] option. This logs some of the disconnected errors in the `hostname.err` file. See [Section 5.2.2, "The Error Log"](#).

If you want to create a bug report regarding this problem, be sure that you include the following information:

- Indicate whether the MySQL server died. You can find information about this in the server error log. See [Section C.5.4.2, "What to Do If MySQL Keeps Crashing"](#).
- If a specific query kills `mysqld` and the tables involved were checked with `CHECK TABLE` before you ran the query, can you provide a reproducible test case? See [MySQL Internals: Porting to Other Systems](#).

- What is the value of the `wait_timeout` [563] system variable in the MySQL server? (`mysqladmin variables` gives you the value of this variable.)
- Have you tried to run `mysqld` with the general query log enabled to determine whether the problem query appears in the log? (See [Section 5.2.3, “The General Query Log”](#).)

See also [Section C.5.2.11, “Communication Errors and Aborted Connections”](#), and [Section 1.7, “How to Report Bugs or Problems”](#).

### C.5.2.10. Packet Too Large

A communication packet is a single SQL statement sent to the MySQL server, a single row that is sent to the client, or a binary log event sent from a master replication server to a slave.

The largest possible packet that can be transmitted to or from a MySQL 5.6 server or client is 1GB.

When a MySQL client or the `mysqld` server receives a packet bigger than `max_allowed_packet` [502] bytes, it issues an `ER_NET_PACKET_TOO_LARGE` [3670] error and closes the connection. With some clients, you may also get a `Lost connection to MySQL server during query` error if the communication packet is too large.

Both the client and the server have their own `max_allowed_packet` [502] variable, so if you want to handle big packets, you must increase this variable both in the client and in the server.

If you are using the `mysql` client program, its default `max_allowed_packet` [502] variable is 16MB. To set a larger value, start `mysql` like this:

```
shell> mysql --max_allowed_packet=32M
```

That sets the packet size to 32MB.

The server's default `max_allowed_packet` [502] value is 1MB. You can increase this if the server needs to handle big queries (for example, if you are working with big `BLOB` columns). For example, to set the variable to 16MB, start the server like this:

```
shell> mysqld --max_allowed_packet=16M
```

You can also use an option file to set `max_allowed_packet` [502]. For example, to set the size for the server to 16MB, add the following lines in an option file:

```
[mysqld]
max_allowed_packet=16M
```

It is safe to increase the value of this variable because the extra memory is allocated only when needed. For example, `mysqld` allocates more memory only when you issue a long query or when `mysqld` must return a large result row. The small default value of the variable is a precaution to catch incorrect packets between the client and server and also to ensure that you do not run out of memory by using large packets accidentally.

You can also get strange problems with large packets if you are using large `BLOB` values but have not given `mysqld` access to enough memory to handle the query. If you suspect this is the case, try adding `ulimit -d 256000` to the beginning of the `mysqld_safe` script and restarting `mysqld`.

### C.5.2.11. Communication Errors and Aborted Connections

If connection problems occur such as communication errors or aborted connections, use these sources of information to diagnose problems:

- The error log. See [Section 5.2.2, “The Error Log”](#).
- The general query log. See [Section 5.2.3, “The General Query Log”](#).
- The `Aborted_xxx` and `Connection_errors_xxx` status variables. See [Section 5.1.6, “Server Status Variables”](#).

- The host cache, which is accessible using the `host_cache` Performance Schema table. See [Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#), and [Section 20.9.9.1, “The host\\_cache Table”](#).

If you start the server with the `--log-warnings [430]` option, you might find messages like this in your error log:

```
010301 14:38:23 Aborted connection 854 to db:
'users' user: 'josh'
```

If a client successfully connects but later disconnects improperly or is terminated, the server increments the `Aborted_clients [585]` status variable, and logs an `Aborted connection` message to the error log. The cause can be any of the following:

- The client program did not call `mysql_close()` before exiting.
- The client had been sleeping more than `wait_timeout [563]` or `interactive_timeout [490]` seconds without issuing any requests to the server. See [Section 5.1.4, “Server System Variables”](#).
- The client program ended abruptly in the middle of a data transfer.

If a client is unable even to connect, the server increments the `Aborted_connects [585]` status variable. Unsuccessful connection attempts can occur for the following reasons:

- A client does not have privileges to connect to a database.
- A client uses an incorrect password.
- A connection packet does not contain the right information.
- It takes more than `connect_timeout [472]` seconds to get a connect packet. See [Section 5.1.4, “Server System Variables”](#).

If these kinds of things happen, it might indicate that someone is trying to break into your server! Messages for these types of problems are logged to the general query log if it is enabled.

Other reasons for problems with aborted clients or aborted connections:

- The `max_allowed_packet [502]` variable value is too small or queries require more memory than you have allocated for `mysqld`. See [Section C.5.2.10, “Packet Too Large”](#).
- Use of Ethernet protocol with Linux, both half and full duplex. Many Linux Ethernet drivers have this bug. You should test for this bug by transferring a huge file using FTP between the client and server machines. If a transfer goes in burst-pause-burst-pause mode, you are experiencing a Linux duplex syndrome. Switch the duplex mode for both your network card and hub/switch to either full duplex or to half duplex and test the results to determine the best setting.
- A problem with the thread library that causes interrupts on reads.
- Badly configured TCP/IP.
- Faulty Ethernets, hubs, switches, cables, and so forth. This can be diagnosed properly only by replacing hardware.

See also [Section C.5.2.9, “MySQL server has gone away”](#).

### C.5.2.12. The table is full

If a table-full error occurs, it may be that the disk is full or that the table has reached its maximum size. The effective maximum table size for MySQL databases is usually determined by operating system constraints on file sizes, not by MySQL internal limits. See [Section E.10.3, “Limits on Table Size”](#).

### C.5.2.13. Can't create/write to file

If you get an error of the following type for some queries, it means that MySQL cannot create a temporary file for the result set in the temporary directory:

```
Can't create/write to file '\\sqla3fe_0.ism'.
```

The preceding error is a typical message for Windows; the Unix message is similar.

One fix is to start `mysqld` with the `--tmpdir` [446] option or to add the option to the `[mysqld]` section of your option file. For example, to specify a directory of `C:\temp`, use these lines:

```
[mysqld]
tmpdir=C:/temp
```

The `C:\temp` directory must exist and have sufficient space for the MySQL server to write to. See [Section 4.2.3.3, "Using Option Files"](#).

Another cause of this error can be permissions issues. Make sure that the MySQL server can write to the `tmpdir` [559] directory.

Check also the error code that you get with `pererror`. One reason the server cannot write to a table is that the file system is full:

```
shell> pererror 28
OS error code 28: No space left on device
```

If you get an error of the following type during startup, it indicates that the file system or directory used for storing data files is write protected. Providing the write error is to a test file, This error is not serious and can be safely ignored.

```
Can't create test file /usr/local/mysql/data/master.lower-test
```

#### C.5.2.14. Commands out of sync

If you get `Commands out of sync; you can't run this command now` in your client code, you are calling client functions in the wrong order.

This can happen, for example, if you are using `mysql_use_result()` and try to execute a new query before you have called `mysql_free_result()`. It can also happen if you try to execute two queries that return data without calling `mysql_use_result()` or `mysql_store_result()` in between.

#### C.5.2.15. Ignoring user

If you get the following error, it means that when `mysqld` was started or when it reloaded the grant tables, it found an account in the `user` table that had an invalid password.

```
Found wrong password for user 'some_user'@'some_host'; ignoring user
```

As a result, the account is simply ignored by the permission system.

The following list indicates possible causes of and fixes for this problem:

- You may be running a new version of `mysqld` with an old `user` table. You can check this by executing `mysqlshow mysql user` to see whether the `Password` column is shorter than 16 characters. If so, you can correct this condition by running the `scripts/add_long_password` script.
- The account has an old password (eight characters long). Update the account in the `user` table to have a new password.
- You have specified a password in the `user` table without using the `PASSWORD()` [1263] function. Use `mysql` to update the account in the `user` table with a new password, making sure to use the `PASSWORD()` [1263] function:

```
mysql> UPDATE user SET Password=PASSWORD('newpwd')
```

```
-> WHERE User='some_user' AND Host='some_host';
```

### C.5.2.16. Table '*tbl\_name*' doesn't exist

If you get either of the following errors, it usually means that no table exists in the default database with the given name:

```
Table 'tbl_name' doesn't exist
Can't find file: 'tbl_name' (errno: 2)
```

In some cases, it may be that the table does exist but that you are referring to it incorrectly:

- Because MySQL uses directories and files to store databases and tables, database and table names are case sensitive if they are located on a file system that has case-sensitive file names.
- Even for file systems that are not case sensitive, such as on Windows, all references to a given table within a query must use the same lettercase.

You can check which tables are in the default database with `SHOW TABLES`. See [Section 13.7.5, “SHOW Syntax”](#).

### C.5.2.17. Can't initialize character set

You might see an error like this if you have character set problems:

```
MySQL Connection Failed: Can't initialize character set charset_name
```

This error can have any of the following causes:

- The character set is a multi-byte character set and you have no support for the character set in the client. In this case, you need to recompile the client by running `CMake` with the `-DDEFAULT_CHARSET=charset_name` [134] or `-DWITH_EXTRA_CHARSETS=charset_name` [136] option. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

All standard MySQL binaries are compiled with `-DWITH_EXTRA_CHARSETS=complex` [136], which enables support for all multi-byte character sets. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

- The character set is a simple character set that is not compiled into `mysqld`, and the character set definition files are not in the place where the client expects to find them.

In this case, you need to use one of the following methods to solve the problem:

- Recompile the client with support for the character set. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).
- Specify to the client the directory where the character set definition files are located. For many clients, you can do this with the `--character-sets-dir` option.
- Copy the character definition files to the path where the client expects them to be.

### C.5.2.18. '*File*' Not Found and Similar Errors

If you get `ERROR '...' not found (errno: 23), Can't open file: ... (errno: 24)`, or any other error with `errno 23` or `errno 24` from MySQL, it means that you haven't allocated enough file descriptors for the MySQL server. You can use the `pererror` utility to get a description of what the error number means:

```
shell> pererror 23
OS error code 23: File table overflow
shell> pererror 24
OS error code 24: Too many open files
shell> pererror 11
OS error code 11: Resource temporarily unavailable
```



The problem here is that `mysqld` is trying to keep open too many files simultaneously. You can either tell `mysqld` not to open so many files at once or increase the number of file descriptors available to `mysqld`.

To tell `mysqld` to keep open fewer files at a time, you can make the table cache smaller by reducing the value of the `table_open_cache` [551] system variable (the default value is 64). This may not entirely prevent running out of file descriptors because in some circumstances the server may attempt to extend the cache size temporarily, as described in [Section 8.4.3.1, “How MySQL Opens and Closes Tables”](#). Reducing the value of `max_connections` [503] also reduces the number of open files (the default value is 100).

To change the number of file descriptors available to `mysqld`, you can use the `--open-files-limit` [228] option to `mysqld_safe` or set the `open_files_limit` [518] system variable. See [Section 5.1.4, “Server System Variables”](#). The easiest way to set these values is to add an option to your option file. See [Section 4.2.3.3, “Using Option Files”](#). If you have an old version of `mysqld` that does not support setting the open files limit, you can edit the `mysqld_safe` script. There is a commented-out line `ulimit -n 256` in the script. You can remove the “#” character to uncomment this line, and change the number 256 to set the number of file descriptors to be made available to `mysqld`.

`--open-files-limit` [228] and `ulimit` can increase the number of file descriptors, but only up to the limit imposed by the operating system. There is also a “hard” limit that can be overridden only if you start `mysqld_safe` or `mysqld` as `root` (just remember that you also need to start the server with the `--user` [446] option in this case so that it does not continue to run as `root` after it starts up). If you need to increase the operating system limit on the number of file descriptors available to each process, consult the documentation for your system.



#### Note

If you run the `tcsh` shell, `ulimit` does not work! `tcsh` also reports incorrect values when you ask for the current limits. In this case, you should start `mysqld_safe` using `sh`.

### C.5.2.19. Table-Corruption Issues

If you have started `mysqld` with `--myisam-recover-options` [432], MySQL automatically checks and tries to repair `MyISAM` tables if they are marked as 'not closed properly' or 'crashed'. If this happens, MySQL writes an entry in the `hostname.err` file `'Warning: Checking table ...'` which is followed by `Warning: Repairing table` if the table needs to be repaired. If you get a lot of these errors, without `mysqld` having died unexpectedly just before, then something is wrong and needs to be investigated further.

In MySQL 5.6, when the server detects `MyISAM` table corruption, it writes additional information to the error log, such as the name and line number of the source file, and the list of threads accessing the table. Example: `Got an error from thread_id=1, mi_dynrec.c:368`. This is useful information to include in bug reports.

See also [Section 5.1.3, “Server Command Options”](#), and [Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”](#).

## C.5.3. Installation-Related Issues

### C.5.3.1. Problems with File Permissions

If you have problems with file permissions, the `UMASK` environment variable might be set incorrectly when `mysqld` starts. For example, MySQL might issue the following error message when you create a table:

```
ERROR: Can't find file: 'path/with/filename.frm' (Errcode: 13)
```

The default `UMASK` value is 0660. You can change this behavior by starting `mysqld_safe` as follows:

```
shell> UMASK=384 # = 600 in octal
shell> export UMASK
shell> mysqld_safe &
```

By default, MySQL creates database directories with an access permission value of `0700`. You can modify this behavior by setting the `UMASK_DIR` variable. If you set its value, new directories are created with the combined `UMASK` and `UMASK_DIR` values. For example, if you want to give group access to all new directories, you can do this:

```
shell> UMASK_DIR=504 # = 770 in octal
shell> export UMASK_DIR
shell> mysqld_safe &
```

MySQL assumes that the value for `UMASK` or `UMASK_DIR` is in octal if it starts with a zero.

See [Section 2.12, “Environment Variables”](#).

## C.5.4. Administration-Related Issues

### C.5.4.1. How to Reset the Root Password

If you have never set a `root` password for MySQL, the server does not require a password at all for connecting as `root`. However, this is insecure. For instructions on assigning passwords, see [Section 2.10.2, “Securing the Initial MySQL Accounts”](#).

If you know the `root` password, but want to change it, see [Section 13.7.1.7, “SET PASSWORD Syntax”](#).

If you set a `root` password previously, but have forgotten it, you can set a new password. The following sections provide instructions for Windows and Unix systems, as well as generic instructions that apply to any system.

#### C.5.4.1.1. Resetting the Root Password: Windows Systems

On Windows, use the following procedure to reset the password for all MySQL `root` accounts:

1. Log on to your system as Administrator.
2. Stop the MySQL server if it is running. For a server that is running as a Windows service, go to the Services manager: From the Start menu, select Control Panel, then Administrative Tools, then Services. Find the MySQL service in the list and stop it.

If your server is not running as a service, you may need to use the Task Manager to force it to stop.

3. Create a text file containing the following statements. Replace the password with the password that you want to use.

```
UPDATE mysql.user SET Password=PASSWORD('MyNewPass') WHERE User='root';
FLUSH PRIVILEGES;
```

Write the `UPDATE` and `FLUSH` statements each on a single line. The `UPDATE` statement resets the password for all `root` accounts, and the `FLUSH` statement tells the server to reload the grant tables into memory so that it notices the password change.

4. Save the file. For this example, the file will be named `C:\mysql-init.txt`.
5. Open a console window to get to the command prompt: From the Start menu, select Run, then enter `cmd` as the command to be run.
6. Start the MySQL server with the special `--init-file [424]` option (notice that the backslash in the option value is doubled):

```
C:\> C:\mysql\bin\mysqld --init-file=C:\mysql-init.txt
```



If you installed MySQL to a location other than `C:\mysql`, adjust the command accordingly.

The server executes the contents of the file named by the `--init-file` [424] option at startup, changing each `root` account password.

You can also add the `--console` [417] option to the command if you want server output to appear in the console window rather than in a log file.

If you installed MySQL using the MySQL Installation Wizard, you may need to specify a `--defaults-file` [219] option:

```
C:\> "C:\Program Files\MySQL\MySQL Server 5.6\bin\mysqld.exe"
      --defaults-file="C:\Program Files\MySQL\MySQL Server 5.6\my.ini"
      --init-file=C:\mysql-init.txt
```

The appropriate `--defaults-file` [219] setting can be found using the Services Manager: From the **Start** menu, select Control Panel, then Administrative Tools, then Services. Find the MySQL service in the list, right-click it, and choose the **Properties** option. The **Path to executable** field contains the `--defaults-file` [219] setting.

7. After the server has started successfully, delete `C:\mysql-init.txt`.

You should now be able to connect to the MySQL server as `root` using the new password. Stop the MySQL server, then restart it in normal mode again. If you run the server as a service, start it from the Windows Services window. If you start the server manually, use whatever command you normally use.

#### C.5.4.1.2. Resetting the Root Password: Unix Systems

On Unix, use the following procedure to reset the password for all MySQL `root` accounts. The instructions assume that you will start the server so that it runs using the Unix login account that you normally use for running the server. For example, if you run the server using the `mysql` login account, you should log in as `mysql` before using the instructions. Alternatively, you can log in as `root`, but in this case you *must* start `mysqld` with the `--user=mysql` [446] option. If you start the server as `root` without using `--user=mysql` [446], the server may create `root`-owned files in the data directory, such as log files, and these may cause permission-related problems for future server startups. If that happens, you will need to either change the ownership of the files to `mysql` or remove them.

1. Log on to your system as the Unix user that the `mysqld` server runs as (for example, `mysql`).
2. Locate the `.pid` file that contains the server's process ID. The exact location and name of this file depend on your distribution, host name, and configuration. Common locations are `/var/lib/mysql/`, `/var/run/mysqld/`, and `/usr/local/mysql/data/`. Generally, the file name has an extension of `.pid` and begins with either `mysqld` or your system's host name.

You can stop the MySQL server by sending a normal `kill` (not `kill -9`) to the `mysqld` process, using the path name of the `.pid` file in the following command:

```
shell> kill `cat /mysql-data-directory/host_name.pid`
```

Use backticks (not forward quotation marks) with the `cat` command. These cause the output of `cat` to be substituted into the `kill` command.

3. Create a text file containing the following statements. Replace the password with the password that you want to use.

```
UPDATE mysql.user SET Password=PASSWORD('MyNewPass') WHERE User='root';
FLUSH PRIVILEGES;
```

Write the `UPDATE` and `FLUSH` statements each on a single line. The `UPDATE` statement resets the password for all `root` accounts, and the `FLUSH` statement tells the server to reload the grant tables into memory so that it notices the password change.

4. Save the file. For this example, the file will be named `/home/me/mysql-init`. The file contains the password, so it should not be saved where it can be read by other users. If you are not logged in as `mysql` (the user the server runs as), make sure that the file has permissions that permit `mysql` to read it.
5. Start the MySQL server with the special `--init-file` [424] option:

```
shell> mysqld_safe --init-file=/home/me/mysql-init &
```

The server executes the contents of the file named by the `--init-file` [424] option at startup, changing each `root` account password.

6. After the server has started successfully, delete `/home/me/mysql-init`.

You should now be able to connect to the MySQL server as `root` using the new password. Stop the server and restart it normally.

### C.5.4.1.3. Resetting the Root Password: Generic Instructions

The preceding sections provide password-resetting instructions for Windows and Unix systems. Alternatively, on any platform, you can set the new password using the `mysql` client (but this approach is less secure):

1. Stop `mysqld` and restart it with the `--skip-grant-tables` [439] option. This enables anyone to connect without a password and with all privileges. Because this is insecure, you might want to use `--skip-grant-tables` [439] in conjunction with `--skip-networking` [440] to prevent remote clients from connecting.
2. Connect to the `mysqld` server with this command:

```
shell> mysql
```

3. Issue the following statements in the `mysql` client. Replace the password with the password that you want to use.

```
mysql> UPDATE mysql.user SET Password=PASSWORD('MyNewPass')
-> WHERE User='root';
mysql> FLUSH PRIVILEGES;
```

The `FLUSH` statement tells the server to reload the grant tables into memory so that it notices the password change.

You should now be able to connect to the MySQL server as `root` using the new password. Stop the server, then restart it normally (without the `--skip-grant-tables` [439] and `--skip-networking` [440] options).

### C.5.4.2. What to Do If MySQL Keeps Crashing

Each MySQL version is tested on many platforms before it is released. This does not mean that there are no bugs in MySQL, but if there are bugs, they should be very few and can be hard to find. If you have a problem, it always helps if you try to find out exactly what crashes your system, because you have a much better chance of getting the problem fixed quickly.

First, you should try to find out whether the problem is that the `mysqld` server dies or whether your problem has to do with your client. You can check how long your `mysqld` server has been up by executing `mysqladmin version`. If `mysqld` has died and restarted, you may find the reason by looking in the server's error log. See [Section 5.2.2, "The Error Log"](#).

On some systems, you can find in the error log a stack trace of where `mysqld` died that you can resolve with the `resolve_stack_dump` program. See [MySQL Internals: Porting to Other Systems](#). Note that the variable values written in the error log may not always be 100% correct.

Many server crashes are caused by corrupted data files or index files. MySQL updates the files on disk with the `write()` system call after every SQL statement and before the client is notified about the result. (This is not true if you are running with `--delay-key-write` [421], in which case data files are written but not index files.) This means that data file contents are safe even if `mysqld` crashes, because the operating system ensures that the unflushed data is written to disk. You can force MySQL to flush everything to disk after every SQL statement by starting `mysqld` with the `--flush` [423] option.

The preceding means that normally you should not get corrupted tables unless one of the following happens:

- The MySQL server or the server host was killed in the middle of an update.
- You have found a bug in `mysqld` that caused it to die in the middle of an update.
- Some external program is manipulating data files or index files at the same time as `mysqld` without locking the table properly.
- You are running many `mysqld` servers using the same data directory on a system that does not support good file system locks (normally handled by the `lockd` lock manager), or you are running multiple servers with external locking disabled.
- You have a crashed data file or index file that contains very corrupt data that confused `mysqld`.
- You have found a bug in the data storage code. This isn't likely, but it is at least possible. In this case, you can try to change the storage engine to another engine by using `ALTER TABLE` on a repaired copy of the table.

Because it is very difficult to know why something is crashing, first try to check whether things that work for others crash for you. Try the following things:

- Stop the `mysqld` server with `mysqladmin shutdown`, run `myisamchk --silent --force */*.MYI` from the data directory to check all MyISAM tables, and restart `mysqld`. This ensures that you are running from a clean state. See [Chapter 5, MySQL Server Administration](#).
- Start `mysqld` with the general query log enabled (see [Section 5.2.3, "The General Query Log"](#)). Then try to determine from the information written to the log whether some specific query kills the server. About 95% of all bugs are related to a particular query. Normally, this is one of the last queries in the log file just before the server restarts. See [Section 5.2.3, "The General Query Log"](#). If you can repeatedly kill MySQL with a specific query, even when you have checked all tables just before issuing it, then you have isolated the bug and should submit a bug report for it. See [Section 1.7, "How to Report Bugs or Problems"](#).
- Try to make a test case that we can use to repeat the problem. See [MySQL Internals: Porting to Other Systems](#).
- Try running the tests in the `mysql-test` directory and the MySQL benchmarks. See [Section 22.1.2, "The MySQL Test Suite"](#). They test MySQL rather well. You can also add code to the benchmarks that simulates your application. The benchmarks are in the `sql-bench` directory in a source distribution or, for a binary distribution, in the `sql-bench` directory under your MySQL installation directory.
- Try the `fork_big.pl` script. (It is located in the `tests` directory of source distributions.)
- Configuring MySQL for debugging makes it much easier to gather information about possible errors if something goes wrong. Reconfigure MySQL with the `-DWITH_DEBUG=1` [136] option to `CMake` and then recompile. See [MySQL Internals: Porting to Other Systems](#).
- Make sure that you have applied the latest patches for your operating system.
- Use the `--skip-external-locking` [439] option to `mysqld`. On some systems, the `lockd` lock manager does not work properly; the `--skip-external-locking` [439] option tells `mysqld` not

to use external locking. (This means that you cannot run two `mysqld` servers on the same data directory and that you must be careful if you use `myisamchk`. Nevertheless, it may be instructive to try the option as a test.)

- If `mysqld` appears to be running but not responding, try `mysqladmin -u root processlist`. Sometimes `mysqld` is not hung even though it seems unresponsive. The problem may be that all connections are in use, or there may be some internal lock problem. `mysqladmin -u root processlist` usually is able to make a connection even in these cases, and can provide useful information about the current number of connections and their status.
- Run the command `mysqladmin -i 5 status` or `mysqladmin -i 5 -r status` in a separate window to produce statistics while running other queries.
- Try the following:
  1. Start `mysqld` from `gdb` (or another debugger). See [MySQL Internals: Porting to Other Systems](#).
  2. Run your test scripts.
  3. Print the backtrace and the local variables at the three lowest levels. In `gdb`, you can do this with the following commands when `mysqld` has crashed inside `gdb`:

```
backtrace
info local
up
info local
up
info local
```

With `gdb`, you can also examine which threads exist with `info threads` and switch to a specific thread with `thread N`, where `N` is the thread ID.

- Try to simulate your application with a Perl script to force MySQL to crash or misbehave.
- Send a normal bug report. See [Section 1.7, “How to Report Bugs or Problems”](#). Be even more detailed than usual. Because MySQL works for many people, the crash might result from something that exists only on your computer (for example, an error that is related to your particular system libraries).
- If you have a problem with tables containing dynamic-length rows and you are using only `VARCHAR` columns (not `BLOB` or `TEXT` columns), you can try to change all `VARCHAR` to `CHAR` with `ALTER TABLE`. This forces MySQL to use fixed-size rows. Fixed-size rows take a little extra space, but are much more tolerant to corruption.

The current dynamic row code has been in use for several years with very few problems, but dynamic-length rows are by nature more prone to errors, so it may be a good idea to try this strategy to see whether it helps.

- Consider the possibility of hardware faults when diagnosing problems. Defective hardware can be the cause of data corruption. Pay particular attention to your memory and disk subsystems when troubleshooting hardware.

### C.5.4.3. How MySQL Handles a Full Disk

This section describes how MySQL responds to disk-full errors (such as “no space left on device”), and to quota-exceeded errors (such as “write failed” or “user block limit reached”).

This section is relevant for writes to `MyISAM` tables. It also applies for writes to binary log files and binary log index file, except that references to “row” and “record” should be understood to mean “event.”

When a disk-full condition occurs, MySQL does the following:

- It checks once every minute to see whether there is enough space to write the current row. If there is enough space, it continues as if nothing had happened.
- Every 10 minutes it writes an entry to the log file, warning about the disk-full condition.

To alleviate the problem, take the following actions:

- To continue, you only have to free enough disk space to insert all records.
- Alternatively, to abort the thread, use `mysqladmin kill`. The thread is aborted the next time it checks the disk (in one minute).
- Other threads might be waiting for the table that caused the disk-full condition. If you have several “locked” threads, killing the one thread that is waiting on the disk-full condition enables the other threads to continue.

Exceptions to the preceding behavior are when you use `REPAIR TABLE` or `OPTIMIZE TABLE` or when the indexes are created in a batch after `LOAD DATA INFILE` or after an `ALTER TABLE` statement. All of these statements may create large temporary files that, if left to themselves, would cause big problems for the rest of the system. If the disk becomes full while MySQL is doing any of these operations, it removes the big temporary files and mark the table as crashed. The exception is that for `ALTER TABLE`, the old table is left unchanged.

#### C.5.4.4. Where MySQL Stores Temporary Files

On Unix, MySQL uses the value of the `TMPDIR` environment variable as the path name of the directory in which to store temporary files. If `TMPDIR` is not set, MySQL uses the system default, which is usually `/tmp`, `/var/tmp`, or `/usr/tmp`.

On Windows, MySQL checks in order the values of the `TMPDIR`, `TEMP`, and `TMP` environment variables. For the first one found to be set, MySQL uses it and does not check those remaining. If none of `TMPDIR`, `TEMP`, or `TMP` are set, MySQL uses the Windows system default, which is usually `C:\windows\temp\`.

If the file system containing your temporary file directory is too small, you can use the `--tmpdir [446]` option to `mysqld` to specify a directory in a file system where you have enough space.

In MySQL 5.6, the `--tmpdir [446]` option can be set to a list of several paths that are used in round-robin fashion. Paths should be separated by colon characters (“:”) on Unix and semicolon characters (“;”) on Windows.



#### Note

To spread the load effectively, these paths should be located on different *physical* disks, not different partitions of the same disk.

If the MySQL server is acting as a replication slave, you should not set `--tmpdir [446]` to point to a directory on a memory-based file system or to a directory that is cleared when the server host restarts. A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails.

MySQL arranges that temporary files are removed if `mysqld` is terminated. On platforms that support it (such as Unix), this is done by unlinking the file after opening it. The disadvantage of this is that the name does not appear in directory listings and you do not see a big temporary file that fills up the file system in which the temporary file directory is located. (In such cases, `ls +Ll` may be helpful in identifying large files associated with `mysqld`.)

When sorting (`ORDER BY` or `GROUP BY`), MySQL normally uses one or two temporary files. The maximum disk space required is determined by the following expression:

```
(length of what is sorted + sizeof(row pointer))
```

```
* number of matched rows
* 2
```

The row pointer size is usually four bytes, but may grow in the future for really big tables.

For some `SELECT` queries, MySQL also creates temporary SQL tables. These are not hidden and have names of the form `SQL_*`.

`ALTER TABLE` creates a temporary table in the same directory as the original table.

#### C.5.4.5. How to Protect or Change the MySQL Unix Socket File

The default location for the Unix socket file that the server uses for communication with local clients is `/tmp/mysql.sock`. (For some distribution formats, the directory might be different, such as `/var/lib/mysql` for RPMs.)

On some versions of Unix, anyone can delete files in the `/tmp` directory or other similar directories used for temporary files. If the socket file is located in such a directory on your system, this might cause problems.

On most versions of Unix, you can protect your `/tmp` directory so that files can be deleted only by their owners or the superuser (`root`). To do this, set the `sticky` bit on the `/tmp` directory by logging in as `root` and using the following command:

```
shell> chmod +t /tmp
```

You can check whether the `sticky` bit is set by executing `ls -ld /tmp`. If the last permission character is `t`, the bit is set.

Another approach is to change the place where the server creates the Unix socket file. If you do this, you should also let client programs know the new location of the file. You can specify the file location in several ways:

- Specify the path in a global or local option file. For example, put the following lines in `/etc/my.cnf`:

```
[mysqld]
socket=/path/to/socket

[client]
socket=/path/to/socket
```

See [Section 4.2.3.3, “Using Option Files”](#).

- Specify a `--socket [211]` option on the command line to `mysqld_safe` and when you run client programs.
- Set the `MYSQL_UNIX_PORT` environment variable to the path of the Unix socket file.
- Recompile MySQL from source to use a different default Unix socket file location. Define the path to the file with the `MYSQL_UNIX_ADDR [135]` option when you run `CMake`. See [Section 2.9.4, “MySQL Source-Configuration Options”](#).

You can test whether the new socket location works by attempting to connect to the server with this command:

```
shell> mysqladmin --socket=/path/to/socket version
```

#### C.5.4.6. Time Zone Problems

If you have a problem with `SELECT NOW()` returning values in UTC and not your local time, you have to tell the server your current time zone. The same applies if `UNIX_TIMESTAMP()` [\[1226\]](#) returns the wrong value. This should be done for the environment in which the server runs; for example, in `mysqld_safe` or `mysql.server`. See [Section 2.12, “Environment Variables”](#).



You can set the time zone for the server with the `--timezone=timezone_name` [229] option to `mysqld_safe`. You can also set it by setting the `TZ` environment variable before you start `mysqld`.

The permissible values for `--timezone` [229] or `TZ` are system dependent. Consult your operating system documentation to see what values are acceptable.

## C.5.5. Query-Related Issues

### C.5.5.1. Case Sensitivity in String Searches

For nonbinary strings (`CHAR`, `VARCHAR`, `TEXT`), string searches use the collation of the comparison operands. For binary strings (`BINARY`, `VARBINARY`, `BLOB`), comparisons use the numeric values of the bytes in the operands; this means that for alphabetic characters, comparisons will be case sensitive.

A comparison between a nonbinary string and binary string is treated as a comparison of binary strings.

Simple comparison operations (`>=`, `>`, `=`, `<`, `<=`, sorting, and grouping) are based on each character's "sort value." Characters with the same sort value are treated as the same character. For example, if "e" and "é" have the same sort value in a given collation, they compare as equal.

The default character set and collation are `latin1` and `latin1_swedish_ci`, so nonbinary string comparisons are case insensitive by default. This means that if you search with `col_name LIKE 'a %'`, you get all column values that start with `A` or `a`. To make this search case sensitive, make sure that one of the operands has a case sensitive or binary collation. For example, if you are comparing a column and a string that both have the `latin1` character set, you can use the `COLLATE` operator to cause either operand to have the `latin1_general_cs` or `latin1_bin` collation:

```
col_name COLLATE latin1_general_cs LIKE 'a%'
col_name LIKE 'a%' COLLATE latin1_general_cs
col_name COLLATE latin1_bin LIKE 'a%'
col_name LIKE 'a%' COLLATE latin1_bin
```

If you want a column always to be treated in case-sensitive fashion, declare it with a case sensitive or binary collation. See [Section 13.1.14, "CREATE TABLE Syntax"](#).

To cause a case-sensitive comparison of nonbinary strings to be case insensitive, use `COLLATE` to name a case-insensitive collation. The strings in the following example normally are case sensitive, but `COLLATE` changes the comparison to be case insensitive:

```
mysql> SET @s1 = 'MySQL' COLLATE latin1_bin,
->      @s2 = 'mysql' COLLATE latin1_bin;
mysql> SELECT @s1 = @s2;
+-----+
| @s1 = @s2 |
+-----+
|          0 |
+-----+
mysql> SELECT @s1 COLLATE latin1_swedish_ci = @s2;
+-----+
| @s1 COLLATE latin1_swedish_ci = @s2 |
+-----+
|                                     1 |
+-----+
```

A binary string is case sensitive in comparisons. To compare the string as case insensitive, convert it to a nonbinary string and use `COLLATE` to name a case-insensitive collation:

```
mysql> SET @s = BINARY 'MySQL';
mysql> SELECT @s = 'mysql';
+-----+
| @s = 'mysql' |
+-----+
|          0 |
+-----+
mysql> SELECT CONVERT(@s USING latin1) COLLATE latin1_swedish_ci = 'mysql';
+-----+
```

```
| CONVERT(@s USING latin1) COLLATE latin1_swedish_ci = 'mysql' |
+-----+
|  |
|  |
|  |
+-----+
|  |
|  |
|  |
+-----+
```

To determine whether a value will compare as a nonbinary or binary string, use the `COLLATION()` [1267] function. This example shows that `VERSION()` [1273] returns a string that has a case-insensitive collation, so comparisons are case insensitive:

```
mysql> SELECT COLLATION(VERSION());
+-----+
| COLLATION(VERSION()) |
+-----+
| utf8_general_ci     |
+-----+
```

For binary strings, the collation value is `binary`, so comparisons will be case sensitive. One context in which you will see `binary` is for compression and encryption functions, which return binary strings as a general rule: string:

```
mysql> SELECT COLLATION(ENCRYPT('x')), COLLATION(SHA1('x'));
+-----+
| COLLATION(ENCRYPT('x')) | COLLATION(SHA1('x')) |
+-----+
| binary                 | binary                |
+-----+
```

To check the sort value of a string, the `WEIGHT_STRING()` [1187] may be helpful. See [Section 12.5, “String Functions”](#).

### C.5.5.2. Problems Using `DATE` Columns

The format of a `DATE` value is `'YYYY-MM-DD'`. According to standard SQL, no other format is permitted. You should use this format in `UPDATE` expressions and in the `WHERE` clause of `SELECT` statements. For example:

```
SELECT * FROM t1 WHERE date >= '2003-05-05';
```

As a convenience, MySQL automatically converts a date to a number if the date is used in a numeric context and vice versa. MySQL also permits a “relaxed” string format when updating and in a `WHERE` clause that compares a date to a `DATE`, `DATETIME`, or `TIMESTAMP` column. “Relaxed” format means that any punctuation character may be used as the separator between parts. For example, `'2004-08-15'` and `'2004#08#15'` are equivalent. MySQL can also convert a string containing no separators (such as `'20040815'`), provided it makes sense as a date.

When you compare a `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` to a constant string with the `<`, `<=`, `=`, `>=`, `>`, or `BETWEEN` operators, MySQL normally converts the string to an internal long integer for faster comparison (and also for a bit more “relaxed” string checking). However, this conversion is subject to the following exceptions:

- When you compare two columns
- When you compare a `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` column to an expression
- When you use any comparison method other than those just listed, such as `IN` or `STRCMP()` [1192].

For those exceptions, the comparison is done by converting the objects to strings and performing a string comparison.

To be on the safe side, assume that strings are compared as strings and use the appropriate string functions if you want to compare a temporal value to a string.

The special “zero” date `'0000-00-00'` can be stored and retrieved as `'0000-00-00'`. When a `'0000-00-00'` date is used through Connector/ODBC, it is automatically converted to `NULL` because ODBC cannot handle that kind of date.



Because MySQL performs the conversions just described, the following statements work (assume that `idate` is a `DATE` column):

```
INSERT INTO t1 (idate) VALUES (19970505);
INSERT INTO t1 (idate) VALUES ('19970505');
INSERT INTO t1 (idate) VALUES ('97-05-05');
INSERT INTO t1 (idate) VALUES ('1997.05.05');
INSERT INTO t1 (idate) VALUES ('1997 05 05');
INSERT INTO t1 (idate) VALUES ('0000-00-00');

SELECT idate FROM t1 WHERE idate >= '1997-05-05';
SELECT idate FROM t1 WHERE idate >= 19970505;
SELECT MOD(idate,100) FROM t1 WHERE idate >= 19970505;
SELECT idate FROM t1 WHERE idate >= '19970505';
```

However, the following statement does not work:

```
SELECT idate FROM t1 WHERE STRCMP(idate,'20030505')=0;
```

`STRCMP()` [1192] is a string function, so it converts `idate` to a string in `'YYYY-MM-DD'` format and performs a string comparison. It does not convert `'20030505'` to the date `'2003-05-05'` and perform a date comparison.

If you enable the `ALLOW_INVALID_DATES` [603] SQL mode, MySQL permits you to store dates that are given only limited checking: MySQL requires only that the day is in the range from 1 to 31 and the month is in the range from 1 to 12. This makes MySQL very convenient for Web applications where you obtain year, month, and day in three different fields and you want to store exactly what the user inserted (without date validation).

MySQL permits you to store dates where the day or month and day are zero. This is convenient if you want to store a birthdate in a `DATE` column and you know only part of the date. To disallow zero month or day parts in dates, enable the `NO_ZERO_IN_DATE` [606] SQL mode.

MySQL permits you to store a “zero” value of `'0000-00-00'` as a “dummy date.” This is in some cases more convenient than using `NULL` values. If a date to be stored in a `DATE` column cannot be converted to any reasonable value, MySQL stores `'0000-00-00'`. To disallow `'0000-00-00'`, enable the `NO_ZERO_DATE` [606] SQL mode.

To have MySQL check all dates and accept only legal dates (unless overridden by `IGNORE`), set the `sql_mode` [545] system variable to `"NO_ZERO_IN_DATE,NO_ZERO_DATE"`.

### C.5.5.3. Problems with `NULL` Values

The concept of the `NULL` value is a common source of confusion for newcomers to SQL, who often think that `NULL` is the same thing as an empty string `' '`. This is not the case. For example, the following statements are completely different:

```
mysql> INSERT INTO my_table (phone) VALUES (NULL);
mysql> INSERT INTO my_table (phone) VALUES ('');
```

Both statements insert a value into the `phone` column, but the first inserts a `NULL` value and the second inserts an empty string. The meaning of the first can be regarded as “phone number is not known” and the meaning of the second can be regarded as “the person is known to have no phone, and thus no phone number.”

To help with `NULL` handling, you can use the `IS NULL` [1168] and `IS NOT NULL` [1168] operators and the `IFNULL()` [1175] function.

In SQL, the `NULL` value is never true in comparison to any other value, even `NULL`. An expression that contains `NULL` always produces a `NULL` value unless otherwise indicated in the documentation for the operators and functions involved in the expression. All columns in the following example return `NULL`:

```
mysql> SELECT NULL, 1+NULL, CONCAT('Invisible',NULL);
```

To search for column values that are `NULL`, you cannot use an `expr = NULL` test. The following statement returns no rows, because `expr = NULL` is never true for any expression:

```
mysql> SELECT * FROM my_table WHERE phone = NULL;
```

To look for `NULL` values, you must use the `IS NULL` [1168] test. The following statements show how to find the `NULL` phone number and the empty phone number:

```
mysql> SELECT * FROM my_table WHERE phone IS NULL;
mysql> SELECT * FROM my_table WHERE phone = '';
```

See [Section 3.3.4.6, “Working with `NULL` Values”](#), for additional information and examples.

You can add an index on a column that can have `NULL` values if you are using the `MyISAM`, `InnoDB`, or `MEMORY` storage engine. Otherwise, you must declare an indexed column `NOT NULL`, and you cannot insert `NULL` into the column.

When reading data with `LOAD DATA INFILE`, empty or missing columns are updated with `' '`. To load a `NULL` value into a column, use `\N` in the data file. The literal word “`NULL`” may also be used under some circumstances. See [Section 13.2.6, “`LOAD DATA INFILE` Syntax”](#).

When using `DISTINCT`, `GROUP BY`, or `ORDER BY`, all `NULL` values are regarded as equal.

When using `ORDER BY`, `NULL` values are presented first, or last if you specify `DESC` to sort in descending order.

Aggregate (summary) functions such as `COUNT()` [1283], `MIN()` [1284], and `SUM()` [1285] ignore `NULL` values. The exception to this is `COUNT(*)` [1283], which counts rows and not individual column values. For example, the following statement produces two counts. The first is a count of the number of rows in the table, and the second is a count of the number of non-`NULL` values in the `age` column:

```
mysql> SELECT COUNT(*), COUNT(age) FROM person;
```

For some data types, MySQL handles `NULL` values specially. If you insert `NULL` into a `TIMESTAMP` column, the current date and time is inserted. If you insert `NULL` into an integer or floating-point column that has the `AUTO_INCREMENT` attribute, the next number in the sequence is inserted.

#### C.5.5.4. Problems with Column Aliases

An alias can be used in a query select list to give a column a different name. You can use the alias in `GROUP BY`, `ORDER BY`, or `HAVING` clauses to refer to the column:

```
SELECT SQRT(a*b) AS root FROM tbl_name
GROUP BY root HAVING root > 0;
SELECT id, COUNT(*) AS cnt FROM tbl_name
GROUP BY id HAVING cnt > 0;
SELECT id AS 'Customer identity' FROM tbl_name;
```

Standard SQL disallows references to column aliases in a `WHERE` clause. This restriction is imposed because when the `WHERE` clause is evaluated, the column value may not yet have been determined. For example, the following query is illegal:

```
SELECT id, COUNT(*) AS cnt FROM tbl_name
WHERE cnt > 0 GROUP BY id;
```

The `WHERE` clause determines which rows should be included in the `GROUP BY` clause, but it refers to the alias of a column value that is not known until after the rows have been selected, and grouped by the `GROUP BY`.

In the select list of a query, a quoted column alias can be specified using identifier or string quoting characters:

```
SELECT 1 AS `one`, 2 AS 'two';
```

Elsewhere in the statement, quoted references to the alias must use identifier quoting or the reference is treated as a string literal. For example, this statement groups by the values in column `id`, referenced using the alias ``a``:

```
SELECT id AS 'a', COUNT(*) AS cnt FROM tbl_name
GROUP BY `a`;
```

But this statement groups by the literal string `'a'` and will not work as expected:

```
SELECT id AS 'a', COUNT(*) AS cnt FROM tbl_name
GROUP BY 'a';
```

### C.5.5.5. Rollback Failure for Nontransactional Tables

If you receive the following message when trying to perform a `ROLLBACK`, it means that one or more of the tables you used in the transaction do not support transactions:

```
Warning: Some non-transactional changed tables couldn't be rolled back
```

These nontransactional tables are not affected by the `ROLLBACK` statement.

If you were not deliberately mixing transactional and nontransactional tables within the transaction, the most likely cause for this message is that a table you thought was transactional actually is not. This can happen if you try to create a table using a transactional storage engine that is not supported by your `mysqld` server (or that was disabled with a startup option). If `mysqld` does not support a storage engine, it instead creates the table as a `MyISAM` table, which is nontransactional.

You can check the storage engine for a table by using either of these statements:

```
SHOW TABLE STATUS LIKE 'tbl_name';
SHOW CREATE TABLE tbl_name;
```

See [Section 13.7.5.37](#), “`SHOW TABLE STATUS Syntax`”, and [Section 13.7.5.12](#), “`SHOW CREATE TABLE Syntax`”.

To check which storage engines your `mysqld` server supports, use this statement:

```
SHOW ENGINES;
```

See [Section 13.7.5.17](#), “`SHOW ENGINES Syntax`” for full details.

### C.5.5.6. Deleting Rows from Related Tables

If the total length of the `DELETE` statement for `related_table` is more than 1MB (the default value of the `max_allowed_packet` [502] system variable), you should split it into smaller parts and execute multiple `DELETE` statements. You probably get the fastest `DELETE` by specifying only 100 to 1,000 `related_column` values per statement if the `related_column` is indexed. If the `related_column` isn't indexed, the speed is independent of the number of arguments in the `IN` clause.

### C.5.5.7. Solving Problems with No Matching Rows

If you have a complicated query that uses many tables but that returns no rows, you should use the following procedure to find out what is wrong:

1. Test the query with `EXPLAIN` to check whether you can find something that is obviously wrong. See [Section 13.8.2](#), “`EXPLAIN Syntax`”.
2. Select only those columns that are used in the `WHERE` clause.
3. Remove one table at a time from the query until it returns some rows. If the tables are large, it is a good idea to use `LIMIT 10` with the query.

4. Issue a `SELECT` for the column that should have matched a row against the table that was last removed from the query.
5. If you are comparing `FLOAT` or `DOUBLE` columns with numbers that have decimals, you cannot use equality (`=`) comparisons. This problem is common in most computer languages because not all floating-point values can be stored with exact precision. In some cases, changing the `FLOAT` to a `DOUBLE` fixes this. See [Section C.5.5.8, “Problems with Floating-Point Values”](#).
6. If you still cannot figure out what is wrong, create a minimal test that can be run with `mysql test < query.sql` that shows your problems. You can create a test file by dumping the tables with `mysqldump --quick db_name tbl_name_1 ... tbl_name_n > query.sql`. Open the file in an editor, remove some insert lines (if there are more than needed to demonstrate the problem), and add your `SELECT` statement at the end of the file.

Verify that the test file demonstrates the problem by executing these commands:

```
shell> mysqladmin create test2
shell> mysql test2 < query.sql
```

Attach the test file to a bug report, which you can file using the instructions in [Section 1.7, “How to Report Bugs or Problems”](#).

### C.5.5.8. Problems with Floating-Point Values

Floating-point numbers sometimes cause confusion because they are approximate and not stored as exact values. A floating-point value as written in an SQL statement may not be the same as the value represented internally. Attempts to treat floating-point values as exact in comparisons may lead to problems. They are also subject to platform or implementation dependencies. The `FLOAT` and `DOUBLE` data types are subject to these issues. For `DECIMAL` columns, MySQL performs operations with a precision of 65 decimal digits, which should solve most common inaccuracy problems.

The following example uses `DOUBLE` to demonstrate how calculations that are done using floating-point operations are subject to floating-point error.

```
mysql> CREATE TABLE t1 (i INT, d1 DOUBLE, d2 DOUBLE);
mysql> INSERT INTO t1 VALUES (1, 101.40, 21.40), (1, -80.00, 0.00),
-> (2, 0.00, 0.00), (2, -13.20, 0.00), (2, 59.60, 46.40),
-> (2, 30.40, 30.40), (3, 37.00, 7.40), (3, -29.60, 0.00),
-> (4, 60.00, 15.40), (4, -10.60, 0.00), (4, -34.00, 0.00),
-> (5, 33.00, 0.00), (5, -25.80, 0.00), (5, 0.00, 7.20),
-> (6, 0.00, 0.00), (6, -51.40, 0.00);

mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b
-> FROM t1 GROUP BY i HAVING a <> b;
```

i	a	b
1	21.4	21.4
2	76.8	76.8
3	7.4	7.4
4	15.4	15.4
5	7.2	7.2
6	-51.4	0

The result is correct. Although the first five records look like they should not satisfy the comparison (the values of `a` and `b` do not appear to be different), they may do so because the difference between the numbers shows up around the tenth decimal or so, depending on factors such as computer architecture or the compiler version or optimization level. For example, different CPUs may evaluate floating-point numbers differently.

If columns `d1` and `d2` had been defined as `DECIMAL` rather than `DOUBLE`, the result of the `SELECT` query would have contained only one row—the last one shown above.

The correct way to do floating-point number comparison is to first decide on an acceptable tolerance for differences between the numbers and then do the comparison against the tolerance value. For example, if we agree that floating-point numbers should be regarded the same if they are same within a precision of one in ten thousand (0.0001), the comparison should be written to find differences larger than the tolerance value:

```
mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b FROM t1
-> GROUP BY i HAVING ABS(a - b) > 0.0001;
+-----+-----+-----+
| i     | a     | b     |
+-----+-----+-----+
| 6     | -51.4 | 0     |
+-----+-----+-----+
1 row in set (0.00 sec)
```

Conversely, to get rows where the numbers are the same, the test should find differences within the tolerance value:

```
mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b FROM t1
-> GROUP BY i HAVING ABS(a - b) <= 0.0001;
+-----+-----+-----+
| i     | a     | b     |
+-----+-----+-----+
| 1     | 21.4  | 21.4  |
| 2     | 76.8  | 76.8  |
| 3     | 7.4   | 7.4   |
| 4     | 15.4  | 15.4  |
| 5     | 7.2   | 7.2   |
+-----+-----+-----+
5 rows in set (0.03 sec)
```

Floating-point values are subject to platform or implementation dependencies. Suppose that you execute the following statements:

```
CREATE TABLE t1(c1 FLOAT(53,0), c2 FLOAT(53,0));
INSERT INTO t1 VALUES('1e+52', '-1e+52');
SELECT * FROM t1;
```

On some platforms, the `SELECT` statement returns `inf` and `-inf`. On others, it returns `0` and `-0`.

An implication of the preceding issues is that if you attempt to create a replication slave by dumping table contents with `mysqldump` on the master and reloading the dump file into the slave, tables containing floating-point columns might differ between the two hosts.

## C.5.6. Optimizer-Related Issues

MySQL uses a cost-based optimizer to determine the best way to resolve a query. In many cases, MySQL can calculate the best possible query plan, but sometimes MySQL does not have enough information about the data at hand and has to make “educated” guesses about the data.

For the cases when MySQL does not do the “right” thing, tools that you have available to help MySQL are:

- Use the `EXPLAIN` statement to get information about how MySQL processes a query. To use it, just add the keyword `EXPLAIN` to the front of your `SELECT` statement:

```
mysql> EXPLAIN SELECT * FROM t1, t2 WHERE t1.i = t2.i;
```

`EXPLAIN` is discussed in more detail in [Section 13.8.2, “EXPLAIN Syntax”](#).

- Use `ANALYZE TABLE tbl_name` to update the key distributions for the scanned table. See [Section 13.7.2.1, “ANALYZE TABLE Syntax”](#).
- Use `FORCE INDEX` for the scanned table to tell MySQL that table scans are very expensive compared to using the given index:

```
SELECT * FROM t1, t2 FORCE INDEX (index_for_column)
WHERE t1.col_name=t2.col_name;
```

`USE INDEX` and `IGNORE INDEX` may also be useful. See [Section 13.2.9.3, “Index Hint Syntax”](#).

- Global and table-level `STRAIGHT_JOIN`. See [Section 13.2.9, “SELECT Syntax”](#).
- You can tune global or thread-specific system variables. For example, start `mysqld` with the `--max-seeks-for-key=1000 [507]` option or use `SET max_seeks_for_key=1000` to tell the optimizer to assume that no key scan causes more than 1,000 key seeks. See [Section 5.1.4, “Server System Variables”](#).

## C.5.7. Table Definition-Related Issues

### C.5.7.1. Problems with `ALTER TABLE`

If you get a duplicate-key error when using `ALTER TABLE` to change the character set or collation of a character column, the cause is either that the new column collation maps two keys to the same value or that the table is corrupted. In the latter case, you should run `REPAIR TABLE` on the table.

If `ALTER TABLE` dies with the following error, the problem may be that MySQL crashed during an earlier `ALTER TABLE` operation and there is an old table named `A-xxx` or `B-xxx` lying around:

```
Error on rename of './database/name.frm'
to './database/B-xxx.frm' (Errcode: 17)
```

In this case, go to the MySQL data directory and delete all files that have names starting with `A-` or `B-`. (You may want to move them elsewhere instead of deleting them.)

`ALTER TABLE` works in the following way:

- Create a new table named `A-xxx` with the requested structural changes.
- Copy all rows from the original table to `A-xxx`.
- Rename the original table to `B-xxx`.
- Rename `A-xxx` to your original table name.
- Delete `B-xxx`.

If something goes wrong with the renaming operation, MySQL tries to undo the changes. If something goes seriously wrong (although this shouldn't happen), MySQL may leave the old table as `B-xxx`. A simple rename of the table files at the system level should get your data back.

If you use `ALTER TABLE` on a transactional table or if you are using Windows, `ALTER TABLE` unlocks the table if you had done a `LOCK TABLE` on it. This is done because `InnoDB` and these operating systems cannot drop a table that is in use.

### C.5.7.2. `TEMPORARY` Table Problems

The following list indicates limitations on the use of `TEMPORARY` tables:

- A `TEMPORARY` table can only be of type `MEMORY`, `MyISAM`, `MERGE`, or `InnoDB`.

Temporary tables are not supported for MySQL Cluster.

- You cannot refer to a `TEMPORARY` table more than once in the same query. For example, the following does not work:

```
mysql> SELECT * FROM temp_table, temp_table AS t2;
ERROR 1137: Can't reopen table: 'temp_table'
```

This error also occurs if you refer to a temporary table multiple times in a stored function under different aliases, even if the references occur in different statements within the function.

- The `SHOW TABLES` statement does not list `TEMPORARY` tables.
- You cannot use `RENAME` to rename a `TEMPORARY` table. However, you can use `ALTER TABLE` instead:

```
mysql> ALTER TABLE orig_name RENAME new_name;
```

- There are known issues in using temporary tables with replication. See [Section 16.4.1, “Replication Features and Issues”](#), for more information.

## C.5.8. Known Issues in MySQL

This section lists known issues in recent versions of MySQL.

For information about platform-specific issues, see the installation and porting instructions in [Section 2.1, “General Installation Guidance”](#), and [MySQL Internals: Porting to Other Systems](#).

The following problems are known:

- Subquery optimization for `IN` is not as effective as for `=`.
- Even if you use `lower_case_table_names=2` (which enables MySQL to remember the case used for databases and table names), MySQL does not remember the case used for database names for the function `DATABASE()` [1268] or within the various logs (on case-insensitive systems).
- Dropping a `FOREIGN KEY` constraint does not work in replication because the constraint may have another name on the slave.
- `REPLACE` (and `LOAD DATA` with the `REPLACE` option) does not trigger `ON DELETE CASCADE`.
- `DISTINCT` with `ORDER BY` does not work inside `GROUP_CONCAT()` [1283] if you do not use all and only those columns that are in the `DISTINCT` list.
- When inserting a big integer value (between  $2^{63}$  and  $2^{64}-1$ ) into a decimal or string column, it is inserted as a negative value because the number is evaluated in a signed integer context.
- `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` may cause problems on nontransactional tables for which you are using `INSERT DELAYED`.
- With statement-based binary logging, the master writes the executed queries to the binary log. This is a very fast, compact, and efficient logging method that works perfectly in most cases. However, it is possible for the data on the master and slave to become different if a query is designed in such a way that the data modification is nondeterministic (generally not a recommended practice, even outside of replication).

For example:

- `CREATE TABLE ... SELECT` or `INSERT ... SELECT` statements that insert zero or `NULL` values into an `AUTO_INCREMENT` column.
- `DELETE` if you are deleting rows from a table that has foreign keys with `ON DELETE CASCADE` properties.
- `REPLACE ... SELECT`, `INSERT IGNORE ... SELECT` if you have duplicate key values in the inserted data.

**If and only if the preceding queries have no `ORDER BY` clause guaranteeing a deterministic order.**



For example, for `INSERT ... SELECT` with no `ORDER BY`, the `SELECT` may return rows in a different order (which results in a row having different ranks, hence getting a different number in the `AUTO_INCREMENT` column), depending on the choices made by the optimizers on the master and slave.

A query is optimized differently on the master and slave only if:

- The table is stored using a different storage engine on the master than on the slave. (It is possible to use different storage engines on the master and slave. For example, you can use `InnoDB` on the master, but `MyISAM` on the slave if the slave has less available disk space.)
- MySQL buffer sizes (`key_buffer_size` [492], and so on) are different on the master and slave.
- The master and slave run different MySQL versions, and the optimizer code differs between these versions.

This problem may also affect database restoration using `mysqlbinlog|mysql`.

The easiest way to avoid this problem is to add an `ORDER BY` clause to the aforementioned nondeterministic queries to ensure that the rows are always stored or modified in the same order. Using row-based or mixed logging format also avoids the problem.

- Log file names are based on the server host name if you do not specify a file name with the startup option. To retain the same log file names if you change your host name to something else, you must explicitly use options such as `--log-bin=old_host_name-bin` [2002]. See Section 5.1.3, “Server Command Options”. Alternatively, rename the old files to reflect your host name change. If these are binary logs, you must edit the binary log index file and fix the binary log file names there as well. (The same is true for the relay logs on a slave server.)
- `mysqlbinlog` does not delete temporary files left after a `LOAD DATA INFILE` statement. See Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”.
- `RENAME` does not work with `TEMPORARY` tables or tables used in a `MERGE` table.
- When using `SET CHARACTER SET`, you cannot use translated characters in database, table, and column names.
- You cannot use “\_” or “%” with `ESCAPE` in `LIKE ... ESCAPE` [1190].
- Only the first `max_sort_length` [508] bytes are used when comparing data values. This means that values cannot reliably be used in `GROUP BY`, `ORDER BY` or `DISTINCT` if they are not distinct in the first `max_sort_length` [508] bytes. To work around this, increase the variable value. The default value of `max_sort_length` [508] is 1024 and can be changed at server startup time or at runtime.
- Numeric calculations are done with `BIGINT` or `DOUBLE` (both are normally 64 bits long). Which precision you get depends on the function. The general rule is that bit functions are performed with `BIGINT` precision, `IF()` [1174] and `ELT()` [1179] with `BIGINT` or `DOUBLE` precision, and the rest with `DOUBLE` precision. You should try to avoid using unsigned long long values if they resolve to be larger than 63 bits (9223372036854775807) for anything other than bit fields.
- You can have up to 255 `ENUM` and `SET` columns in one table.
- In `MIN()` [1284], `MAX()` [1284], and other aggregate functions, MySQL currently compares `ENUM` and `SET` columns by their string value rather than by the string's relative position in the set.
- In an `UPDATE` statement, columns are updated from left to right. If you refer to an updated column, you get the updated value instead of the original value. For example, the following statement increments `KEY` by 2, **not** 1:

```
mysql> UPDATE tbl_name SET KEY=KEY+1,KEY=KEY+1;
```



- You can refer to multiple temporary tables in the same query, but you cannot refer to any given temporary table more than once. For example, the following does not work:

```
mysql> SELECT * FROM temp_table, temp_table AS t2;
ERROR 1137: Can't reopen table: 'temp_table'
```

- The optimizer may handle `DISTINCT` differently when you are using “hidden” columns in a join than when you are not. In a join, hidden columns are counted as part of the result (even if they are not shown), whereas in normal queries, hidden columns do not participate in the `DISTINCT` comparison.

An example of this is:

```
SELECT DISTINCT mp3id FROM band_downloads
WHERE userid = 9 ORDER BY id DESC;
```

and

```
SELECT DISTINCT band_downloads.mp3id
FROM band_downloads,band_mp3
WHERE band_downloads.userid = 9
AND band_mp3.id = band_downloads.mp3id
ORDER BY band_downloads.id DESC;
```

In the second case, using MySQL Server 3.23.x, you may get two identical rows in the result set (because the values in the hidden `id` column may differ).

Note that this happens only for queries that do not have the `ORDER BY` columns in the result.

- If you execute a `PROCEDURE` on a query that returns an empty set, in some cases the `PROCEDURE` does not transform the columns.
- Creation of a table of type `MERGE` does not check whether the underlying tables are compatible types.
- If you use `ALTER TABLE` to add a `UNIQUE` index to a table used in a `MERGE` table and then add a normal index on the `MERGE` table, the key order is different for the tables if there was an old, non-`UNIQUE` key in the table. This is because `ALTER TABLE` puts `UNIQUE` indexes before normal indexes to be able to detect duplicate keys as early as possible.



---

## Appendix D. MySQL Release Notes

Release notes for MySQL Server and associated products are no longer published in the MySQL Reference Manual. Instead, see these release notes documents:

- [MySQL 5.6 Release Notes](#)
- [MySQL Connector/ODBC Release Notes](#)
- [MySQL Connector/Net Release Notes](#) (includes MySQL Visual Studio Plugin release notes)
- [MySQL Connector/J Release Notes](#)
- [MySQL Connector/C++ Release Notes](#)
- [MySQL Proxy Release Notes](#)
- [MySQL for Excel Release Notes](#)
- [MySQL Installer Release Notes](#)



---

## Appendix E. Restrictions and Limits

### Table of Contents

E.1. Restrictions on Stored Programs .....	3757
E.2. Restrictions on Condition Handling .....	3760
E.3. Restrictions on Server-Side Cursors .....	3760
E.4. Restrictions on Subqueries .....	3761
E.5. Restrictions on Views .....	3763
E.6. Restrictions on XA Transactions .....	3764
E.7. Restrictions on Character Sets .....	3765
E.8. Restrictions on Performance Schema .....	3765
E.9. Restrictions on Pluggable Authentication .....	3766
E.10. Limits in MySQL .....	3767
E.10.1. Limits of Joins .....	3767
E.10.2. Limits on Number of Databases and Tables .....	3768
E.10.3. Limits on Table Size .....	3768
E.10.4. Limits on Table Column Count and Row Size .....	3769
E.10.5. Limits Imposed by <code>.frm</code> File Structure .....	3770
E.10.6. Windows Platform Limitations .....	3771

The discussion here describes restrictions that apply to the use of MySQL features such as subqueries or views.

### E.1. Restrictions on Stored Programs

These restrictions apply to the features described in [Chapter 18, \*Stored Programs and Views\*](#).

Some of the restrictions noted here apply to all stored routines; that is, both to stored procedures and stored functions. There are also some [restrictions specific to stored functions](#) but not to stored procedures.

The restrictions for stored functions also apply to triggers. There are also some [restrictions specific to triggers](#).

The restrictions for stored procedures also apply to the `DO` clause of Event Scheduler event definitions. There are also some [restrictions specific to events](#).

### SQL Statements Not Permitted in Stored Routines

Stored routines cannot contain arbitrary SQL statements. The following statements are not permitted:

- The locking statements `LOCK TABLES` and `UNLOCK TABLES`.
- `ALTER VIEW`.
- `LOAD DATA` and `LOAD TABLE`.
- SQL prepared statements (`PREPARE`, `EXECUTE`, `DEALLOCATE PREPARE`) can be used in stored procedures, but not stored functions or triggers. Thus, stored functions and triggers cannot use dynamic SQL (where you construct statements as strings and then execute them).
- Generally, statements not permitted in SQL prepared statements are also not permitted in stored programs. For a list of statements supported as prepared statements, see [Section 13.5, “SQL Syntax for Prepared Statements”](#). Exceptions are `SIGNAL`, `RESIGNAL`, and `GET DIAGNOSTICS`, which are not permissible as prepared statements but are permitted in stored programs.

- Because local variables are in scope only during stored program execution, references to them are not permitted in prepared statements created within a stored program. Prepared statement scope is the current session, not the stored program, so the statement could be executed after the program ends, at which point the variables would no longer be in scope. For example, `SELECT ... INTO local_var` cannot be used as a prepared statement. This restriction also applies to stored procedure and function parameters. See [Section 13.5.1, “PREPARE Syntax”](#).
- Inserts cannot be delayed. `INSERT DELAYED` syntax is accepted, but the statement is handled as a normal `INSERT`.
- Within all stored programs (stored procedures and functions, triggers, and events), the parser treats `BEGIN [WORK]` as the beginning of a `BEGIN ... END` block. To begin a transaction in this context, use `START TRANSACTION` instead.

## Restrictions for Stored Functions

The following additional statements or operations are not permitted within stored functions. They are permitted within stored procedures, except stored procedures that are invoked from within a stored function or trigger. For example, if you use `FLUSH` in a stored procedure, that stored procedure cannot be called from a stored function or trigger.

- Statements that perform explicit or implicit commit or rollback. Support for these statements is not required by the SQL standard, which states that each DBMS vendor may decide whether to permit them.
- Statements that return a result set. This includes `SELECT` statements that do not have an `INTO var_list` clause and other statements such as `SHOW`, `EXPLAIN`, and `CHECK TABLE`. A function can process a result set either with `SELECT ... INTO var_list` or by using a cursor and `FETCH` statements. See [Section 13.2.9.1, “SELECT ... INTO Syntax”](#), and [Section 13.6.6, “Cursors”](#).
- `FLUSH` statements.
- Stored functions cannot be used recursively.
- A stored function or trigger cannot modify a table that is already being used (for reading or writing) by the statement that invoked the function or trigger.
- If you refer to a temporary table multiple times in a stored function under different aliases, a `Can't reopen table: 'tbl_name'` error occurs, even if the references occur in different statements within the function.
- `HANDLER ... READ` statements that invoke stored functions can cause replication errors and are disallowed.

## Restrictions for Triggers

For triggers, the following additional restrictions apply:

- Triggers currently are not activated by foreign key actions.
- When using row-based replication, triggers on the slave are not activated by statements originating on the master. The triggers on the slave are activated when using statement-based replication. For more information, see [Section 16.4.1.30, “Replication and Triggers”](#).
- The `RETURN` statement is not permitted in triggers, which cannot return a value. To exit a trigger immediately, use the `LEAVE` statement.
- Triggers are not permitted on tables in the `mysql` database.
- The trigger cache does not detect when metadata of the underlying objects has changed. If a trigger uses a table and the table has changed since the trigger was loaded into the cache, the trigger operates using the outdated metadata.

## Name Conflicts within Stored Routines

The same identifier might be used for a routine parameter, a local variable, and a table column. Also, the same local variable name can be used in nested blocks. For example:

```
CREATE PROCEDURE p (i INT)
BEGIN
  DECLARE i INT DEFAULT 0;
  SELECT i FROM t;
  BEGIN
    DECLARE i INT DEFAULT 1;
    SELECT i FROM t;
  END;
END;
```

In such cases, the identifier is ambiguous and the following precedence rules apply:

- A local variable takes precedence over a routine parameter or table column.
- A routine parameter takes precedence over a table column.
- A local variable in an inner block takes precedence over a local variable in an outer block.

The behavior that variables take precedence over table columns is nonstandard.

## Replication Considerations

Use of stored routines can cause replication problems. This issue is discussed further in [Section 18.7, “Binary Logging of Stored Programs”](#).

The `--replicate-wild-do-table=db_name.tbl_name [1979]` option applies to tables, views, and triggers. It does not apply to stored procedures and functions, or events. To filter statements operating on the latter objects, use one or more of the `--replicate-*-db` options.

## Debugging Considerations

There are no stored routine debugging facilities.

## Unsupported Syntax from the SQL:2003 Standard

The MySQL stored routine syntax is based on the SQL:2003 standard. The following items from that standard are not currently supported:

- [UNDO](#) handlers
- [FOR](#) loops

## Concurrency Considerations

To prevent problems of interaction between sessions, when a client issues a statement, the server uses a snapshot of routines and triggers available for execution of the statement. That is, the server calculates a list of procedures, functions, and triggers that may be used during execution of the statement, loads them, and then proceeds to execute the statement. While the statement executes, it does not see changes to routines performed by other sessions.

For maximum concurrency, stored functions should minimize their side-effects; in particular, updating a table within a stored function can reduce concurrent operations on that table. A stored function acquires table locks before executing, to avoid inconsistency in the binary log due to mismatch of the order in which statements execute and when they appear in the log. When statement-based binary logging is used, statements that invoke a function are recorded rather than the statements executed within the function. Consequently, stored functions that update the same underlying tables do not execute in parallel. In contrast, stored procedures do not acquire table-level locks. All statements executed within stored procedures are written to the binary log, even for statement-based binary logging. See [Section 18.7, “Binary Logging of Stored Programs”](#).

## Event Scheduler Restrictions

The following limitations are specific to the Event Scheduler:

- Event names are handled in case-insensitive fashion. For example, you cannot have two events in the same database with the names `anEvent` and `AnEvent`.
- An event may not be created, altered, or dropped by a stored routine, trigger, or another event. An event also may not create, alter, or drop stored routines or triggers. (Bug #16409, Bug #18896)
- DDL statements on events are prohibited while a `LOCK TABLES` statement is in effect.
- Event timings using the intervals `YEAR`, `QUARTER`, `MONTH`, and `YEAR_MONTH` are resolved in months; those using any other interval are resolved in seconds. There is no way to cause events scheduled to occur at the same second to execute in a given order. In addition—due to rounding, the nature of threaded applications, and the fact that a nonzero length of time is required to create events and to signal their execution—events may be delayed by as much as 1 or 2 seconds. However, the time shown in the `INFORMATION_SCHEMA.EVENTS` table's `LAST_EXECUTED` column or the `mysql.event` table's `last_executed` column is always accurate to within one second of the actual event execution time. (See also Bug #16522.)
- Each execution of the statements contained in the body of an event takes place in a new connection; thus, these statements has no effect in a given user session on the server's statement counts such as `Com_select` and `Com_insert` that are displayed by `SHOW STATUS`. However, such counts are updated in the global scope. (Bug #16422)
- Events do not support times later than the end of the Unix Epoch; this is approximately the beginning of the year 2038. Such dates are specifically not permitted by the Event Scheduler. (Bug #16396)
- References to stored functions, user-defined functions, and tables in the `ON SCHEDULE` clauses of `CREATE EVENT` and `ALTER EVENT` statements are not supported. These sorts of references are not permitted. (See Bug #22830 for more information.)

## E.2. Restrictions on Condition Handling

`SIGNAL`, `RESIGNAL`, and `GET DIAGNOSTICS` are not permissible as prepared statements. For example, this statement is invalid:

```
PREPARE stmt1 FROM 'SIGNAL SQLSTATE "02000"';
```

`SQLSTATE` values in class `'04'` are not treated specially. They are handled the same as other exceptions.

Standard SQL has a diagnostics area stack, containing a diagnostics area for each nested execution context. Standard SQL syntax includes `GET STACKED DIAGNOSTICS` for referring to stacked areas. MySQL does not support the `STACKED` keyword because there is a single diagnostics area containing information from the most recent statement that wrote to it. See also [Section 13.6.7.7, “The MySQL Diagnostics Area”](#).

In standard SQL, the first condition relates to the `SQLSTATE` value returned for the previous SQL statement. In MySQL, this is not guaranteed, so to get the main error, you cannot do this:

```
GET DIAGNOSTICS CONDITION 1 @errno = MYSQL_ERRNO;
```

Instead, do this:

```
GET DIAGNOSTICS @cno = NUMBER;
GET DIAGNOSTICS CONDITION @cno @errno = MYSQL_ERRNO;
```

## E.3. Restrictions on Server-Side Cursors

Server-side cursors are implemented in the C API using the `mysql_stmt_attr_set()` function. The same implementation is used for cursors in stored routines. A server-side cursor enables a result set



to be generated on the server side, but not transferred to the client except for those rows that the client requests. For example, if a client executes a query but is only interested in the first row, the remaining rows are not transferred.

In MySQL, a server-side cursor is materialized into an internal temporary table. Initially, this is a `MEMORY` table, but is converted to a `MyISAM` table when its size exceeds the minimum value of the `max_heap_table_size` [504] and `tmp_table_size` [558] system variables. Note that the same restrictions apply to internal temporary tables created to hold the result set for a cursor as for other uses of internal temporary tables. See [Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#). One limitation of the implementation is that for a large result set, retrieving its rows through a cursor might be slow.

Cursors are read only; you cannot use a cursor to update rows.

`UPDATE WHERE CURRENT OF` and `DELETE WHERE CURRENT OF` are not implemented, because updatable cursors are not supported.

Cursors are nonholdable (not held open after a commit).

Cursors are asensitive.

Cursors are nonscrollable.

Cursors are not named. The statement handler acts as the cursor ID.

You can have open only a single cursor per prepared statement. If you need several cursors, you must prepare several statements.

You cannot use a cursor for a statement that generates a result set if the statement is not supported in prepared mode. This includes statements such as `CHECK TABLE`, `HANDLER READ`, and `SHOW BINLOG EVENTS`.

## E.4. Restrictions on Subqueries

- Subquery optimization for `IN` is not as effective as for the `=` operator or for the `IN(value_list)` [1169] operator.

A typical case for poor `IN` subquery performance is when the subquery returns a small number of rows but the outer query returns a large number of rows to be compared to the subquery result.

The problem is that, for a statement that uses an `IN` subquery, the optimizer rewrites it as a correlated subquery. Consider the following statement that uses an uncorrelated subquery:

```
SELECT ... FROM t1 WHERE t1.a IN (SELECT b FROM t2);
```

The optimizer rewrites the statement to a correlated subquery:

```
SELECT ... FROM t1 WHERE EXISTS (SELECT 1 FROM t2 WHERE t2.b = t1.a);
```

If the inner and outer queries return  $M$  and  $N$  rows, respectively, the execution time becomes on the order of  $O(M \times N)$ , rather than  $O(M+N)$  as it would be for an uncorrelated subquery.

An implication is that an `IN` subquery can be much slower than a query written using an `IN(value_list)` [1169] operator that lists the same values that the subquery would return.

- In general, you cannot modify a table and select from the same table in a subquery. For example, this limitation applies to statements of the following forms:

```
DELETE FROM t WHERE ... (SELECT ... FROM t ...);
UPDATE t ... WHERE col = (SELECT ... FROM t ...);
{INSERT|REPLACE} INTO t (SELECT ... FROM t ...);
```

Exception: The preceding prohibition does not apply if you are using a subquery for the modified table in the `FROM` clause. Example:

```
UPDATE t ... WHERE col = (SELECT * FROM (SELECT ... FROM t...) AS _t ...);
```

Here the result from the subquery in the `FROM` clause is stored as a temporary table, so the relevant rows in `t` have already been selected by the time the update to `t` takes place.

- Row comparison operations are only partially supported:
  - For `expr [NOT] IN subquery`, `expr` can be an  $n$ -tuple (specified using row constructor syntax) and the subquery can return rows of  $n$ -tuples. The permitted syntax is therefore more specifically expressed as `row_constructor [NOT] IN table_subquery`
  - For `expr op {ALL|ANY|SOME} subquery`, `expr` must be a scalar value and the subquery must be a column subquery; it cannot return multiple-column rows.

In other words, for a subquery that returns rows of  $n$ -tuples, this is supported:

```
(expr_1, ..., expr_n) [NOT] IN table_subquery
```

But this is not supported:

```
(expr_1, ..., expr_n) op {ALL|ANY|SOME} subquery
```

The reason for supporting row comparisons for `IN` but not for the others is that `IN` is implemented by rewriting it as a sequence of `=` [1166] comparisons and `AND` [1171] operations. This approach cannot be used for `ALL`, `ANY`, or `SOME`.

- Subqueries in the `FROM` clause cannot be correlated subqueries. They are materialized in whole (evaluated to produce a result set) during query execution, so they cannot be evaluated per row of the outer query. Before MySQL 5.6.3, materialization takes place before evaluation of the outer query. As of 5.6.3, the optimizer delays materialization until the result is needed, which may permit materialization to be avoided. See [Section 8.13.16.3](#), “Optimizing Subqueries in the `FROM` Clause (Derived Tables)”.
- MySQL does not support `LIMIT` in subqueries for certain subquery operators:

```
mysql> SELECT * FROM t1
-> WHERE s1 IN (SELECT s2 FROM t2 ORDER BY s1 LIMIT 1);
ERROR 1235 (42000): This version of MySQL doesn't yet support
'LIMIT & IN/ALL/ANY/SOME subquery'
```

- The optimizer is more mature for joins than for subqueries, so in many cases a statement that uses a subquery can be executed more efficiently if you rewrite it as a join.

An exception occurs for the case where an `IN` subquery can be rewritten as a `SELECT DISTINCT` join. Example:

```
SELECT col FROM t1 WHERE id_col IN (SELECT id_col2 FROM t2 WHERE condition);
```

That statement can be rewritten as follows:

```
SELECT DISTINCT col FROM t1, t2 WHERE t1.id_col = t2.id_col AND condition;
```

- MySQL permits a subquery to refer to a stored function that has data-modifying side effects such as inserting rows into a table. For example, if `f()` inserts rows, the following query can modify data:

```
SELECT ... WHERE x IN (SELECT f() ...);
```

This behavior is an extension to the SQL standard. In MySQL, it can produce indeterminate results because `f()` might be executed a different number of times for different executions of a given query depending on how the optimizer chooses to handle it.

For statement-based or mixed-format replication, one implication of this indeterminism is that such a query can produce different results on the master and its slaves.

- Before MySQL 5.6.3, a subquery in the `FROM` clause is evaluated by materializing the result into a temporary table, and this table does not use indexes. As of 5.6.3, the optimizer creates an index on the materialized table if this will result in faster query execution. See [Section 8.13.16.3, “Optimizing Subqueries in the `FROM` Clause \(Derived Tables\)”](#).

## E.5. Restrictions on Views

View processing is not optimized:

- It is not possible to create an index on a view.
- Indexes can be used for views processed using the merge algorithm. However, a view that is processed with the temptable algorithm is unable to take advantage of indexes on its underlying tables (although indexes can be used during generation of the temporary tables).

Subqueries cannot be used in the `FROM` clause of a view.

There is a general principle that you cannot modify a table and select from the same table in a subquery. See [Section E.4, “Restrictions on Subqueries”](#).

The same principle also applies if you select from a view that selects from the table, if the view selects from the table in a subquery and the view is evaluated using the merge algorithm. Example:

```
CREATE VIEW v1 AS
SELECT * FROM t2 WHERE EXISTS (SELECT 1 FROM t1 WHERE t1.a = t2.a);
UPDATE t1, v2 SET t1.a = 1 WHERE t1.b = v2.b;
```

If the view is evaluated using a temporary table, you *can* select from the table in the view subquery and still modify that table in the outer query. In this case the view will be stored in a temporary table and thus you are not really selecting from the table in a subquery and modifying it “at the same time.” (This is another reason you might wish to force MySQL to use the temptable algorithm by specifying `ALGORITHM = TEMPTABLE` in the view definition.)

You can use `DROP TABLE` or `ALTER TABLE` to drop or alter a table that is used in a view definition. No warning results from the `DROP` or `ALTER` operation, even though this invalidates the view. Instead, an error occurs later, when the view is used. `CHECK TABLE` can be used to check for views that have been invalidated by `DROP` or `ALTER` operations.

A view definition is “frozen” by certain statements:

- If a statement prepared by `PREPARE` refers to a view, the view definition seen each time the statement is executed later will be the definition of the view at the time it was prepared. This is true even if the view definition is changed after the statement is prepared and before it is executed.

Example:

```
CREATE VIEW v AS SELECT RAND();
PREPARE s FROM 'SELECT * FROM v';
ALTER VIEW v AS SELECT NOW();
EXECUTE s;
```

The result returned by the `EXECUTE` statement is a random number, not the current date and time.

With regard to view updatability, the overall goal for views is that if any view is theoretically updatable, it should be updatable in practice. This includes views that have `UNION` in their definition. Currently, not all views that are theoretically updatable can be updated. The initial view implementation was deliberately written this way to get usable, updatable views into MySQL as quickly as possible. Many theoretically updatable views can be updated now, but limitations still exist:

- Updatable views with subqueries anywhere other than in the `WHERE` clause. Some views that have subqueries in the `SELECT` list may be updatable.
- You cannot use `UPDATE` to update more than one underlying table of a view that is defined as a join.

- You cannot use `DELETE` to update a view that is defined as a join.

There exists a shortcoming with the current implementation of views. If a user is granted the basic privileges necessary to create a view (the `CREATE VIEW` [754] and `SELECT` [755] privileges), that user will be unable to call `SHOW CREATE VIEW` on that object unless the user is also granted the `SHOW VIEW` [756] privilege.

That shortcoming can lead to problems backing up a database with `mysqldump`, which may fail due to insufficient privileges. This problem is described in Bug #22062.

The workaround to the problem is for the administrator to manually grant the `SHOW VIEW` [756] privilege to users who are granted `CREATE VIEW` [754], since MySQL doesn't grant it implicitly when views are created.

Views do not have indexes, so index hints do not apply. Use of index hints when selecting from a view is not permitted.

`SHOW CREATE VIEW` displays view definitions using an `AS alias_name` clause for each column. If a column is created from an expression, the default alias is the expression text, which can be quite long. Aliases for column names in `CREATE VIEW` statements are checked against the maximum column length of 64 characters (not the maximum alias length of 256 characters). As a result, views created from the output of `SHOW CREATE VIEW` fail if any column alias exceeds 64 characters. This can cause problems in the following circumstances for views with too-long aliases:

- View definitions fail to replicate to newer slaves that enforce the column-length restriction.
- Dump files created with `mysqldump` cannot be loaded into servers that enforce the column-length restriction.

A workaround for either problem is to modify each problematic view definition to use aliases that provide shorter column names. Then the view will replicate properly, and can be dumped and reloaded without causing an error. To modify the definition, drop and create the view again with `DROP VIEW` and `CREATE VIEW`, or replace the definition with `CREATE OR REPLACE VIEW`.

For problems that occur when reloading view definitions in dump files, another workaround is to edit the dump file to modify its `CREATE VIEW` statements. However, this does not change the original view definitions, which may cause problems for subsequent dump operations.

## E.6. Restrictions on XA Transactions

XA transaction support is limited to the `InnoDB` storage engine.

For “external XA,” a MySQL server acts as a Resource Manager and client programs act as Transaction Managers. For “Internal XA”, storage engines within a MySQL server act as RMs, and the server itself acts as a TM. Internal XA support is limited by the capabilities of individual storage engines. Internal XA is required for handling XA transactions that involve more than one storage engine. The implementation of internal XA requires that a storage engine support two-phase commit at the table handler level, and currently this is true only for `InnoDB`.

For `XA START`, the `JOIN` and `RESUME` clauses are not supported.

For `XA END`, the `SUSPEND [FOR MIGRATE]` clause is not supported.

The requirement that the `bqual` part of the `xid` value be different for each XA transaction within a global transaction is a limitation of the current MySQL XA implementation. It is not part of the XA specification.

If an XA transaction has reached the `PREPARED` state and the MySQL server is killed (for example, with `kill -9` on Unix) or shuts down abnormally, the transaction can be continued after the server restarts. However, if the client reconnects and commits the transaction, the transaction will be absent

from the binary log even though it has been committed. This means the data and the binary log have gone out of synchrony. An implication is that XA cannot be used safely together with replication.

It is possible that the server will roll back a pending XA transaction, even one that has reached the `PREPARED` state. This happens if a client connection terminates and the server continues to run, or if clients are connected and the server shuts down gracefully. (In the latter case, the server marks each connection to be terminated, and then rolls back the `PREPARED` XA transaction associated with it.) It should be possible to commit or roll back a `PREPARED` XA transaction, but this cannot be done without changes to the binary logging mechanism.

## E.7. Restrictions on Character Sets

- Identifiers are stored in `mysql` database tables (`user`, `db`, and so forth) using `utf8`, but identifiers can contain only characters in the Basic Multilingual Plane (BMP). Supplementary characters are not permitted in identifiers.
- The `ucs2`, `utf16`, `utf16le`, and `utf32` character sets have the following restrictions:
  - They cannot be used as a client character set, which means that they do not work for `SET NAMES` or `SET CHARACTER SET`. (See [Section 10.1.4, “Connection Character Sets and Collations”](#).)
  - It is currently not possible to use `LOAD DATA INFILE` to load data files that use these character sets.
  - `FULLTEXT` indexes cannot be created on a column that uses any of these character sets. However, you can perform `IN BOOLEAN MODE` searches on the column without an index.
  - The use of `ENCRYPT()` [1262] with these character sets is not recommended because the underlying system call expects a string terminated by a zero byte.
- The `REGEXP` [1193] and `RLIKE` [1193] operators work in byte-wise fashion, so they are not multi-byte safe and may produce unexpected results with multi-byte character sets. In addition, these operators compare characters by their byte values and accented characters may not compare as equal even if a given collation treats them as equal.

## E.8. Restrictions on Performance Schema

The Performance Schema avoids using mutexes to collect or produce data, so there are no guarantees of consistency and results can sometimes be incorrect. Event values in `performance_schema` tables are nondeterministic and nonrepeatable.

If you save event information in another table, you should not assume that the original events will still be available later. For example, if you select events from a `performance_schema` table into a temporary table, intending to join that table with the original table later, there might be no matches.

`mysqldump` and `BACKUP DATABASE` ignore tables in the `performance_schema` database.

Tables in the `performance_schema` database cannot be locked with `LOCK TABLES`, except the `setup_XXX` tables.

Tables in the `performance_schema` database cannot be indexed.

Results for queries that refer to tables in the `performance_schema` database are not saved in the query cache.

Tables in the `performance_schema` database are not replicated.

The Performance Schema is not available in `libmysqld`, the embedded server.

The types of timers might vary per platform. The `performance_timers` table shows which event timers are available. If the values in this table for a given timer name are `NULL`, that timer is not supported on your platform.

Instruments that apply to storage engines might not be implemented for all storage engines. Instrumentation of each third-party engine is the responsibility of the engine maintainer.

## E.9. Restrictions on Pluggable Authentication

The first part of this section describes general restrictions on the applicability of the pluggable authentication framework described at [Section 6.3.7, “Pluggable Authentication”](#). The second part describes how third-party connector developers can determine the extent to which a connector can take advantage of pluggable authentication capabilities and what steps to take to become more compliant.

The term “native authentication” used here refers to authentication against passwords stored in the `Password` column of the `mysql.user` table. This is the same authentication method provided by older MySQL servers, before pluggable authentication was implemented. It remains the default method, although now it is implemented using plugins. “Windows native authentication” refers to authentication using the credentials of a user who has already logged in to Windows, as implemented by the Windows Native Authentication plugin (“Windows plugin” for short).

### General Pluggable Authentication Restrictions

- **Connector/C, Connector/C++:** Clients that use these connectors can connect to the server only through accounts that use native authentication.

Exception: A connector supports pluggable authentication if it was built to link to `libmysqlclient` dynamically (rather than statically) and it loads the current version of `libmysqlclient` if that version is installed, or if the connector is recompiled from source to link against the current `libmysqlclient`.

- **Connector/J:** Clients that use this connector can connect to the server only through accounts that use native authentication.
- **Connector/Net:** Before Connector/Net 6.4.4, clients that use this connector can connect to the server only through accounts that use native authentication. As of 6.4.4, clients can also connect to the server through accounts that use the Windows plugin.
- **Connector/ODBC:** Before Connector/ODBC 3.51.29 and 5.1.9, clients that use this connector can connect to the server only through accounts that use native authentication. As of 3.51.29 and 5.1.9, clients that use binary releases of this connector for Windows can also connect to the server through accounts that use the PAM or Windows plugins. (These capabilities result from linking the Connector/ODBC binaries against the MySQL 5.5.16 `libmysqlclient` rather than the MySQL 5.1 `libmysqlclient` used previously. The newer `libmysqlclient` includes the client-side support needed for the server-side PAM and Windows authentication plugins.)
- **Connector/PHP:** Clients that use this connector can connect to the server only through accounts that use native authentication, when compiled using the MySQL native driver for PHP (`mysqlnd`).
- **MySQL Proxy:** Before MySQL Proxy 0.8.2, clients can connect to the server only through accounts that use native authentication. As of 0.8.2, clients can also connect to the server through accounts that use the PAM plugin. As of 0.8.3, clients can also connect to the server through accounts that use the Windows plugin.
- **MySQL Enterprise Backup:** MySQL Enterprise Backup before version 3.6.1 supports connections to the server only through accounts that use native authentication. As of 3.6.1, MySQL Enterprise Backup can connect to the server through accounts that use nonnative authentication.
- **Windows native authentication:** Connecting through an account that uses the Windows plugin requires Windows Domain setup. Without it, NTLM authentication is used and then only local connections are possible; that is, the client and server must run on the same computer.
- **Proxy users:** Proxy user support is available to the extent that clients can connect through accounts authenticated with plugins that implement proxy user capability (that is, plugins that can return a user



name different from that of the connecting user). For example, the native authentication plugins do not support proxy users, whereas the PAM and Windows plugins do.

- **Replication:** Before MySQL 5.6.4, replication slaves can connect to the master server only through master accounts that use native authentication. As of 5.6.4, replication slaves can also connect through master accounts that use nonnative authentication if the required client-side plugin is available. If the plugin is built into `libmysqlclient`, it is available by default. Otherwise, the plugin must be installed on the slave side in the directory named by the slave `plugin_dir` [525] system variable.
- **FEDERATED tables:** A `FEDERATED` table can access the remote table only through accounts on the remote server that use native authentication.

## Pluggable Authentication and Third-Party Connectors

Third-party connector developers can use the following guidelines to determine readiness of a connector to take advantage of pluggable authentication capabilities and what steps to take to become more compliant:

- An existing connector to which no changes have been made uses native authentication and clients that use the connector can connect to the server only through accounts that use native authentication. *However, you should test the connector against a recent version of the server to verify that such connections still work without problem.*

Exception: A connector might work with pluggable authentication without any changes if it links to `libmysqlclient` dynamically (rather than statically) and it loads the current version of `libmysqlclient` if that version is installed.

- To take advantage of pluggable authentication capabilities, a connector that is `libmysqlclient`-based should be relinked against the current version of `libmysqlclient`. This enables the connector to support connections through accounts that require client-side plugins now built into `libmysqlclient` (such as the cleartext plugin needed for PAM authentication and the Windows plugin needed for Windows native authentication). Linking with a current `libmysqlclient` also enables the connector to access client-side plugins installed in the default MySQL plugin directory (typically the directory named by the default value of the local server's `plugin_dir` [525] system variable).

If a connector links to `libmysqlclient` dynamically, it must be ensured that the newer version of `libmysqlclient` is installed on the client host and that the connector loads it at runtime.

- Another way for a connector to support a given authentication method is to implement it directly in the client/server protocol. Connector/Net uses this approach to provide support for Windows native authentication.
- If a connector should be able to load client-side plugins from a directory different from the default plugin directory, it must implement some means for client users to specify the directory. Possibilities for this include a command-line option or environment variable from which the connector can obtain the directory name. Standard MySQL client programs such as `mysql` and `mysqladmin` implement a `--plugin-dir` option. See also [Section 21.8.13, “C API Client Plugin Functions”](#).
- Proxy user support by a connector depends, as described earlier in this section, on whether the authentication methods that it supports permit proxy users.

## E.10. Limits in MySQL

This section lists current limits in MySQL 5.6.

### E.10.1. Limits of Joins

The maximum number of tables that can be referenced in a single join is 61. This also applies to the number of tables that can be referenced in the definition of a view.

## E.10.2. Limits on Number of Databases and Tables

MySQL has no limit on the number of databases. The underlying file system may have a limit on the number of directories.

MySQL has no limit on the number of tables. The underlying file system may have a limit on the number of files that represent tables. Individual storage engines may impose engine-specific constraints. [InnoDB](#) permits up to 4 billion tables.

## E.10.3. Limits on Table Size

The effective maximum table size for MySQL databases is usually determined by operating system constraints on file sizes, not by MySQL internal limits. The following table lists some examples of operating system file-size limits. This is only a rough guide and is not intended to be definitive. For the most up-to-date information, be sure to check the documentation specific to your operating system.

Operating System	File-size Limit
Win32 w/ FAT/FAT32	2GB/4GB
Win32 w/ NTFS	2TB (possibly larger)
Linux 2.2-Intel 32-bit	2GB (LFS: 4GB)
Linux 2.4+	(using ext3 file system) 4TB
Solaris 9/10	16TB
Mac OS X w/ HFS+	2TB

Windows users, please note that FAT and VFAT (FAT32) are *not* considered suitable for production use with MySQL. Use NTFS instead.

On Linux 2.2, you can get [MyISAM](#) tables larger than 2GB in size by using the Large File Support (LFS) patch for the ext2 file system. Most current Linux distributions are based on kernel 2.4 or higher and include all the required LFS patches. On Linux 2.4, patches also exist for ReiserFS to get support for big files (up to 2TB). With JFS and XFS, petabyte and larger files are possible on Linux.

For a detailed overview about LFS in Linux, have a look at Andreas Jaeger's *Large File Support in Linux* page at [http://www.suse.de/~aj/linux\\_lfs.html](http://www.suse.de/~aj/linux_lfs.html).

If you do encounter a full-table error, there are several reasons why it might have occurred:

- The disk might be full.
- The [InnoDB](#) storage engine maintains [InnoDB](#) tables within a tablespace that can be created from several files. This enables a table to exceed the maximum individual file size. The tablespace can include raw disk partitions, which permits extremely large tables. The maximum tablespace size is 64TB.

If you are using [InnoDB](#) tables and run out of room in the [InnoDB](#) tablespace. In this case, the solution is to extend the [InnoDB](#) tablespace. See [Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”](#).

- You are using [MyISAM](#) tables on an operating system that supports files only up to 2GB in size and you have hit this limit for the data file or index file.
- You are using a [MyISAM](#) table and the space required for the table exceeds what is permitted by the internal pointer size. [MyISAM](#) permits data and index files to grow up to 256TB by default, but this limit can be changed up to the maximum permissible size of 65,536TB ( $256^7 - 1$  bytes).

If you need a [MyISAM](#) table that is larger than the default limit and your operating system supports large files, the `CREATE TABLE` statement supports `AVG_ROW_LENGTH` and `MAX_ROWS` options. See [Section 13.1.14, “CREATE TABLE Syntax”](#). The server uses these options to determine how large a table to permit.



If the pointer size is too small for an existing table, you can change the options with `ALTER TABLE` to increase a table's maximum permissible size. See [Section 13.1.6, “ALTER TABLE Syntax”](#).

```
ALTER TABLE tbl_name MAX_ROWS=1000000000 AVG_ROW_LENGTH=nnn;
```

You have to specify `AVG_ROW_LENGTH` only for tables with `BLOB` or `TEXT` columns; in this case, MySQL can't optimize the space required based only on the number of rows.

To change the default size limit for `MyISAM` tables, set the `myisam_data_pointer_size` [511], which sets the number of bytes used for internal row pointers. The value is used to set the pointer size for new tables if you do not specify the `MAX_ROWS` option. The value of `myisam_data_pointer_size` [511] can be from 2 to 7. A value of 4 permits tables up to 4GB; a value of 6 permits tables up to 256TB.

You can check the maximum data and index sizes by using this statement:

```
SHOW TABLE STATUS FROM db_name LIKE 'tbl_name' ;
```

You also can use `myisamchk -dv /path/to/table-index-file`. See [Section 13.7.5, “SHOW Syntax”](#), or [Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

Other ways to work around file-size limits for `MyISAM` tables are as follows:

- If your large table is read only, you can use `myisampack` to compress it. `myisampack` usually compresses a table by at least 50%, so you can have, in effect, much bigger tables. `myisampack` also can merge multiple tables into a single table. See [Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).
- MySQL includes a `MERGE` library that enables you to handle a collection of `MyISAM` tables that have identical structure as a single `MERGE` table. See [Section 14.8, “The MERGE Storage Engine”](#).
- You are using the `MEMORY (HEAP)` storage engine; in this case you need to increase the value of the `max_heap_table_size` [504] system variable. See [Section 5.1.4, “Server System Variables”](#).

## E.10.4. Limits on Table Column Count and Row Size

There is a hard limit of 4096 columns per table, but the effective maximum may be less for a given table. The exact limit depends on several interacting factors.

- Every table (regardless of storage engine) has a maximum row size of 65,535 bytes. Storage engines may place additional constraints on this limit, reducing the effective maximum row size.

The maximum row size constrains the number (and possibly size) of columns because the total length of all columns cannot exceed this size. For example, `utf8` characters require up to three bytes per character, so for a `CHAR(255) CHARACTER SET utf8` column, the server must allocate  $255 \times 3 = 765$  bytes per value. Consequently, a table cannot contain more than  $65,535 / 765 = 85$  such columns.

Storage for variable-length columns includes length bytes, which are assessed against the row size. For example, a `VARCHAR(255) CHARACTER SET utf8` column takes two bytes to store the length of the value, so each value can take up to 767 bytes.

`BLOB` and `TEXT` columns count from one to four plus eight bytes each toward the row-size limit because their contents are stored separately from the rest of the row.

Declaring columns `NULL` can reduce the maximum number of columns permitted. For `MyISAM` tables, `NULL` columns require additional space in the row to record whether their values are `NULL`. Each `NULL` column takes one bit extra, rounded up to the nearest byte. The maximum row length in bytes can be calculated as follows:

```
row length = 1
```

```
+ (sum of column lengths)
+ (number of NULL columns + delete_flag + 7)/8
+ (number of variable-length columns)
```

`delete_flag` is 1 for tables with static row format. Static tables use a bit in the row record for a flag that indicates whether the row has been deleted. `delete_flag` is 0 for dynamic tables because the flag is stored in the dynamic row header. For information about `MyISAM` table formats, see [Section 14.3.3, “MyISAM Table Storage Formats”](#).

For `InnoDB` tables, storage size is the same for `NULL` and `NOT NULL` columns, so the preceding calculations do not apply.

The following statement to create table `t1` succeeds because the columns require 32,765 + 2 bytes and 32,766 + 2 bytes, which falls within the maximum row size of 65,535 bytes:

```
mysql> CREATE TABLE t1
-> (c1 VARCHAR(32765) NOT NULL, c2 VARCHAR(32766) NOT NULL)
-> ENGINE = MyISAM CHARACTER SET latin1;
Query OK, 0 rows affected (0.02 sec)
```

The following statement to create table `t2` fails because the columns are `NULL` and `MyISAM` requires additional space that causes the row size to exceed 65,535 bytes:

```
mysql> CREATE TABLE t2
-> (c1 VARCHAR(32765) NULL, c2 VARCHAR(32766) NULL)
-> ENGINE = MyISAM CHARACTER SET latin1;
ERROR 1118 (42000): Row size too large. The maximum row size for the
used table type, not counting BLOBs, is 65535. You have to change some
columns to TEXT or BLOBs
```

The following statement to create table `t3` fails because, although the column length is within the maximum length of 65,535 bytes, two additional bytes are required to record the length, which causes the row size to exceed 65,535 bytes:

```
mysql> CREATE TABLE t3
-> (c1 VARCHAR(65535) NOT NULL)
-> ENGINE = MyISAM CHARACTER SET latin1;
ERROR 1118 (42000): Row size too large. The maximum row size for the
used table type, not counting BLOBs, is 65535. You have to change some
columns to TEXT or BLOBs
```

Reducing the column length to 65,533 or less permits the statement to succeed.

- Individual storage engines might impose additional restrictions that limit table column count. Examples:
  - `InnoDB` permits up to 1000 columns.
  - `InnoDB` restricts row size to something less than half a database page (approximately 8000 bytes), not including `VARBINARY`, `VARCHAR`, `BLOB`, or `TEXT` columns.
  - Different `InnoDB` storage formats (`COMPRESSED`, `REDUNDANT`) use different amounts of page header and trailer data, which affects the amount of storage available for rows.
- Each table has an `.frm` file that contains the table definition. The definition affects the content of this file in ways that may affect the number of columns permitted in the table. For more information, see [Section E.10.5, “Limits Imposed by `.frm` File Structure”](#).

## E.10.5. Limits Imposed by `.frm` File Structure

Each table has an `.frm` file that contains the table definition. The server uses the following expression to check some of the table information stored in the file against an upper limit of 64KB:

```
if (info_length+(ulong) create_fields.elements*FCOMP+288+
    n_length+int_length+com_length > 65535L || int_count > 255)
```

The portion of the information stored in the `.frm` file that is checked against the expression cannot grow beyond the 64KB limit, so if the table definition reaches this size, no more columns can be added.

The relevant factors in the expression are:

- `info_length` is space needed for “screens.” This is related to MySQL's Unireg heritage.
- `create_fields.elements` is the number of columns.
- `FCOMP` is 17.
- `n_length` is the total length of all column names, including one byte per name as a separator.
- `int_length` is related to the list of values for `ENUM` and `SET` columns. In this context, “int” does not mean “integer.” It means “interval,” a term that refers collectively to `ENUM` and `SET` columns.
- `int_count` is the number of unique `ENUM` and `SET` definitions.
- `com_length` is the total length of column comments.

The expression just described has several implications for permitted table definitions:

- Using long column names can reduce the maximum number of columns, as can the inclusion of `ENUM` or `SET` columns, or use of column comments.
- A table can have no more than 255 unique `ENUM` and `SET` definitions. Columns with identical element lists are considered the same against this limit. For example, if a table contains these two columns, they count as one (not two) toward this limit because the definitions are identical:

```
e1 ENUM('a','b','c')
e2 ENUM('a','b','c')
```

- The sum of the length of element names in the unique `ENUM` and `SET` definitions counts toward the 64KB limit, so although the theoretical limit on number of elements in a given `ENUM` column is 65,535, the practical limit is less than 3000.

## E.10.6. Windows Platform Limitations

The following limitations apply to use of MySQL on the Windows platform:

- **Process memory**

On Windows 32-bit platforms, it is not possible by default to use more than 2GB of RAM within a single process, including MySQL. This is because the physical address limit on Windows 32-bit is 4GB and the default setting within Windows is to split the virtual address space between kernel (2GB) and user/applications (2GB).

Some versions of Windows have a boot time setting to enable larger applications by reducing the kernel application. Alternatively, to use more than 2GB, use a 64-bit version of Windows.

- **File system aliases**

When using `MyISAM` tables, you cannot use aliases within Windows link to the data files on another volume and then link back to the main MySQL `datadir` [418] location.

This facility is often used to move the data and index files to a RAID or other fast solution, while retaining the main `.frm` files in the default data directory configured with the `datadir` [418] option.

- **Limited number of ports**

Windows systems have about 4,000 ports available for client connections, and after a connection on a port closes, it takes two to four minutes before the port can be reused. In situations where clients

connect to and disconnect from the server at a high rate, it is possible for all available ports to be used up before closed ports become available again. If this happens, the MySQL server appears to be unresponsive even though it is running. Note that ports may be used by other applications running on the machine as well, in which case the number of ports available to MySQL is lower.

For more information about this problem, see <http://support.microsoft.com/default.aspx?scid=kb;en-us;196271>.

- **DATA DIRECTORY and INDEX DIRECTORY**

The `DATA DIRECTORY` option for `CREATE TABLE` is supported on Windows only for `InnoDB` tables, as described in [Section 5.4.1.2, “Specifying the Location of a Tablespace”](#). For `MyISAM` and other storage engines, the `DATA DIRECTORY` and `INDEX DIRECTORY` options for `CREATE TABLE` are ignored on Windows and any other platforms with a nonfunctional `realpath()` call.

- **DROP DATABASE**

You cannot drop a database that is in use by another session.

- **Case-insensitive names**

File names are not case sensitive on Windows, so MySQL database and table names are also not case sensitive on Windows. The only restriction is that database and table names must be specified using the same case throughout a given statement. See [Section 9.2.2, “Identifier Case Sensitivity”](#).

- **Directory and file names**

On Windows, MySQL Server supports only directory and file names that are compatible with the current ANSI code pages. For example, the following Japanese directory name will not work in the Western locale (code page 1252):

```
datadir="C:/私たちのプロジェクトのデータ"
```

The same limitation applies to directory and file names referred to in SQL statements, such as the data file path name in `LOAD DATA INFILE`.

- **The “\” path name separator character**

Path name components in Windows are separated by the “\” character, which is also the escape character in MySQL. If you are using `LOAD DATA INFILE` or `SELECT ... INTO OUTFILE`, use Unix-style file names with “/” characters:

```
mysql> LOAD DATA INFILE 'C:/tmp/skr.txt' INTO TABLE skr;
mysql> SELECT * INTO OUTFILE 'C:/tmp/skr.txt' FROM skr;
```

Alternatively, you must double the “\” character:

```
mysql> LOAD DATA INFILE 'C:\\tmp\\skr.txt' INTO TABLE skr;
mysql> SELECT * INTO OUTFILE 'C:\\tmp\\skr.txt' FROM skr;
```

- **Problems with pipes**

Pipes do not work reliably from the Windows command-line prompt. If the pipe includes the character `^Z / CHAR(24)`, Windows thinks that it has encountered end-of-file and aborts the program.

This is mainly a problem when you try to apply a binary log as follows:

```
C:\> mysqlbinlog binary_log_file | mysql --user=root
```

If you have a problem applying the log and suspect that it is because of a `^Z / CHAR(24)` character, you can use the following workaround:

```
C:\> mysqlbinlog binary_log_file --result-file=/tmp/bin.sql
C:\> mysql --user=root --execute "source /tmp/bin.sql"
```

The latter command also can be used to reliably read in any SQL file that may contain binary data.



---

# MySQL Glossary

These terms are commonly used in information about the MySQL database server. This glossary originated as a reference for terminology about the InnoDB storage engine, and the majority of definitions are InnoDB-related.

## A

### .ARM file

Metadata for ARCHIVE tables. Contrast with **.ARZ file**. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product. See Also [.ARZ file](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

### .ARZ file

Data for ARCHIVE tables. Contrast with **.ARM file**. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product. See Also [.ARM file](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

### ACID

An acronym standing for atomicity, consistency, isolation, and durability. These properties are all desirable in a database system, and are all closely tied to the notion of a **transaction**. The transactional features of InnoDB adhere to the ACID principles.

Transactions are **atomic** units of work that can be **committed** or **rolled back**. When a transaction makes multiple changes to the database, either all the changes succeed when the transaction is committed, or all the changes are undone when the transaction is rolled back.

The database remains in a consistent state at all times -- after each commit or rollback, and while transactions are in progress. If related data is being updated across multiple tables, queries see either all old values or all new values, not a mix of old and new values.

Transactions are protected (isolated) from each other while they are in progress; they cannot interfere with each other or see each other's uncommitted data. This isolation is achieved through the **locking** mechanism. Experienced users can adjust the **isolation level**, trading off less protection in favor of increased performance and **concurrency**, when they can be sure that the transactions really do not interfere with each other.

The results of transactions are durable: once a commit operation succeeds, the changes made by that transaction are safe from power failures, system crashes, race conditions, or other potential dangers that many non-database applications are vulnerable to. Durability typically involves writing to disk storage, with a certain amount of redundancy to protect against power failures or software crashes during write operations. (In InnoDB, the **doublewrite buffer** assists with durability.)

See Also [atomic](#), [commit](#), [concurrency](#), [doublewrite buffer](#), [isolation level](#), [locking](#), [rollback](#), [transaction](#).

### adaptive flushing

An algorithm for **InnoDB** tables that smooths out the I/O overhead introduced by **checkpoints**. Instead of **flushing** all modified **pages** from the **buffer pool** to the **data files** at once, MySQL periodically flushes small sets of modified pages. The adaptive flushing algorithm extends this process by estimating the optimal rate to perform these periodic flushes, based on the rate of flushing and how fast **redo** information is generated. First introduced in MySQL 5.1, in the InnoDB Plugin.

See Also [buffer pool](#), [checkpoint](#), [data files](#), [flush](#), [InnoDB](#), [page](#), [redo log](#).

### adaptive hash index

An optimization for InnoDB tables that can speed up lookups using `=` and `IN` operators, by constructing a **hash index** in memory. MySQL monitors index searches for InnoDB tables, and if queries could benefit from a hash index, it builds one automatically for index **pages** that are frequently accessed. In a sense, the adaptive hash index configures MySQL at runtime to take advantage of ample main memory, coming closer to the architecture of main-memory databases. This feature is controlled by the `innodb_adaptive_hash_index` [1701] configuration option. Because this feature benefits some workloads and not others, and the memory used for the hash index is reserved in the **buffer pool**, typically you should benchmark with this feature both enabled and disabled.

---

The hash index is always built based on an existing InnoDB **secondary index**, which is organized as a **B-tree** structure. MySQL can build a hash index on a prefix of any length of the key defined for the B-tree, depending on the pattern of searches against the index. A hash index can be partial; the whole B-tree index does not need to be cached in the buffer pool.

In MySQL 5.6 and higher, another way to take advantage of fast single-value lookups with InnoDB tables is to use the **memcached** interface to InnoDB. See [Section 14.2.9, “InnoDB Integration with memcached”](#) for details.

See Also [B-tree](#), [buffer pool](#), [hash index](#), [memcached](#), [page](#), [secondary index](#).

#### AHI

Acronym for **adaptive hash index**.

See Also [adaptive hash index](#).

#### AIO

Acronym for **asynchronous I/O**. You might see this acronym in InnoDB messages or keywords.

See Also [asynchronous I/O](#).

#### Antelope

The code name for the original InnoDB **file format**. It supports the **redundant** and **compact** row formats, but not the newer **dynamic** and **compressed** row formats available in the **Barracuda** file format.

If your application could benefit from InnoDB table **compression**, or uses BLOBs or large text columns that could benefit from the dynamic row format, you might switch some tables to Barracuda format. You select the file format to use by setting the `innodb_file_format` [1716] option before creating the table.

See Also [Barracuda](#), [compact row format](#), [compressed row format](#), [dynamic row format](#), [file format](#), [innodb\\_file\\_format](#), [redundant row format](#).

#### application programming interface (API)

A set of functions or procedures. An API provides a stable set of names and types for functions, procedures, parameters, and return values.

#### apply

When a backup produced by the **MySQL Enterprise Backup** product does not include the most recent changes that occurred while the backup was underway, the process of updating the backup files to include those changes is known as the **apply** step. It is specified by the `apply-log` option of the `mysqlbackup` command.

Before the changes are applied, we refer to the files as a **raw backup**. After the changes are applied, we refer to the files as a **prepared backup**. The changes are recorded in the `ibbackup_logfile` file; once the apply step is finished, this file is no longer necessary.

See Also [hot backup](#), [ibbackup\\_logfile](#), [MySQL Enterprise Backup](#), [prepared backup](#), [raw backup](#).

#### asynchronous I/O

A type of I/O operation that allows other processing to proceed before the I/O is completed. Also known as **non-blocking I/O** and abbreviated as **AIO**. InnoDB uses this type of I/O for certain operations that can run in parallel without affecting the reliability of the database, such as reading pages into the **buffer pool** that have not actually been requested, but might be needed soon.

Historically, InnoDB has used asynchronous I/O on Windows systems only. Starting with the InnoDB Plugin 1.1, InnoDB uses asynchronous I/O on Linux systems. This change introduces a dependency on `libaio`. On other Unix-like systems, InnoDB uses synchronous I/O only.

See Also [buffer pool](#), [non-blocking I/O](#).

#### atomic

In the SQL context, **transactions** are units of work that either succeed entirely (when **committed**) or have no effect at all (when **rolled back**). The indivisible ("atomic") property of transactions is the "A" in the acronym **ACID**.

See Also [ACID](#), [commit](#), [rollback](#), [transaction](#).



---

## atomic instruction

Special instructions provided by the CPU, to ensure that critical low-level operations cannot be interrupted.

## auto-increment

A property of a table column (specified by the [AUTO\\_INCREMENT](#) keyword) that automatically adds an ascending sequence of values in the column. InnoDB supports auto-increment only for **primary key** columns.

It saves work for the developer, not to have to produce new unique values when inserting new rows. It provides useful information for the query optimizer, because the column is known to be not null and with unique values. The values from such a column can be used as lookup keys in various contexts, and because they are auto-generated there is no reason to ever change them; for this reason, primary key columns are often specified as auto-incrementing.

Auto-increment columns can be problematic with statement-based replication, because replaying the statements on a slave might not produce the same set of column values as on the master, due to timing issues. When you have an auto-incrementing primary key, you can use statement-based replication only with the setting `innodb_autoinc_lock_mode=1` [1705]. If you have `innodb_autoinc_lock_mode=2`, which allows higher concurrency for insert operations, use **row-based replication** rather than **statement-based replication**. The setting `innodb_autoinc_lock_mode=0` is the previous (traditional) default setting and should not be used except for compatibility purposes.

See Also [auto-increment locking](#), `innodb_autoinc_lock_mode`, [primary key](#), [row-based replication](#), [statement-based replication](#).

## auto-increment locking

The convenience of an **auto-increment** primary key involves some tradeoff with concurrency. In the simplest case, if one transaction is inserting values into the table, any other transactions must wait to do their own inserts into that table, so that rows inserted by the first transaction receive consecutive primary key values. InnoDB includes optimizations, and the `innodb_autoinc_lock_mode` [1705] option, so that you can choose how to trade off between predictable sequences of auto-increment values and maximum **concurrency** for insert operations.

See Also [auto-increment](#), [concurrency](#), `innodb_autoinc_lock_mode`.

## autocommit

A setting that causes a **commit** operation after each **SQL** statement. This mode is not recommended for working with InnoDB tables with **transactions** that span several statements. It can help performance for **read-only transactions** on InnoDB tables, where it minimizes overhead from **locking** and generation of **undo** data, especially in MySQL 5.6.4 and up. It is also appropriate for working with MyISAM tables, where transactions are not applicable.

See Also [commit](#), [locking](#), [read-only transaction](#), [SQL](#), [transaction](#), [undo](#).

## availability

The ability to cope with, and if necessary recover from, failures on the host, including failures of MySQL, the operating system, or the hardware and maintenance activity that may otherwise cause downtime. Often paired with **scalability** as critical aspects of a large-scale deployment.

See Also [scalability](#).

# B

## B-tree

A tree data structure that is popular for use in database indexes. The structure is kept sorted at all times, enabling fast lookup for exact matches (equals operator) and ranges (for example, greater than, less than, and [BETWEEN](#) operators). This type of index is available for most storage engines, such as InnoDB and MyISAM.

Because B-tree nodes can have many children, a B-tree is not the same as a binary tree, which is limited to 2 children per node.

Contrast with **hash index**, which is only available in the MEMORY storage engine. The MEMORY storage engine can also use B-tree indexes, and you should choose B-tree indexes for MEMORY tables if some queries use range operators.

---

See Also [hash index](#).

#### backticks

Identifiers within MySQL SQL statements must be quoted using the backtick character ( ``` ) if they contain special characters or reserved words. For example, to refer to a table named `FOO#BAR` or a column named `SELECT`, you would specify the identifiers as ``FOO#BAR`` and ``SELECT``. Since the backticks provide an extra level of safety, they are used extensively in program-generated SQL statements, where the identifier names might not be known in advance.

Many other database systems use double quotation marks ( `"` ) around such special names. For portability, you can enable `ANSI_QUOTES` mode in MySQL and use double quotation marks instead of backticks to qualify identifier names.

See Also [SQL](#).

#### backup

The process of copying some or all table data and metadata from a MySQL instance, for safekeeping. Can also refer to the set of copied files. This is a crucial task for DBAs. The reverse of this process is the **restore** operation.

With MySQL, **physical backups** are performed by the **MySQL Enterprise Backup** product, and **logical backups** are performed by the `mysqldump` command. These techniques have different characteristics in terms of size and representation of the backup data, and speed (especially speed of the restore operation).

Backups are further classified as **hot**, **warm**, or **cold** depending on how much they interfere with normal database operation. (Hot backups have the least interference, cold backups the most.)

See Also [cold backup](#), [hot backup](#), [logical backup](#), [MySQL Enterprise Backup](#), [mysqldump](#), [physical backup](#), [warm backup](#).

#### Barracuda

The code name for an InnoDB **file format** that supports compression for table data. This file format was first introduced in the InnoDB Plugin. It supports the **compressed** row format that enables InnoDB table compression, and the **dynamic** row format that improves the storage layout for BLOB and large text columns. You can select it through the `innodb_file_format [1716]` option.

Because the InnoDB **system tablespace** is stored in the original **Antelope** file format, to use the Barracuda file format you must also enable the **file-per-table** setting, which puts newly created tables in their own tablespaces separate from the system tablespace.

The **MySQL Enterprise Backup** product version 3.5 and above supports backing up tablespaces that use the Barracuda file format.

See Also [Antelope](#), [compact row format](#), [compressed row format](#), [dynamic row format](#), [file format](#), [file-per-table](#), [innodb\\_file\\_format](#), [MySQL Enterprise Backup](#), [row format](#), [system tablespace](#).

#### beta

An early stage in the life of a software product, when it is available only for evaluation, typically without a definite release number or a number less than 1. InnoDB does not use the beta designation, preferring an **early adopter** phase that can extend over several point releases, leading to a **GA** release.

See Also [early adopter](#), [GA](#).

#### binary log

A file containing a record of all statements that attempt to change table data. These statements can be replayed to bring slave servers up to date in a **replication** scenario, or to bring a database up to date after restoring table data from a backup. The binary logging feature can be turned on and off, although Oracle recommends always enabling it if you use replication or perform backups.

You can examine the contents of the binary log, or replay those statements during replication or recovery, by using the `mysqlbinlog` command. For full information about the binary log, see [Section 5.2.4, “The Binary Log”](#). For MySQL configuration options related to the binary log, see [Section 16.1.4.4, “Binary Log Options and Variables”](#).

---

For the **MySQL Enterprise Backup** product, the file name of the binary log and the current position within the file are important details. To record this information for the master server when taking a backup in a replication context, you can specify the `--slave-info` option.

Prior to MySQL 5.0, a similar capability was available, known as the update log. In MySQL 5.0 and higher, the binary log replaces the update log.

See Also [binlog](#), [MySQL Enterprise Backup](#), [replication](#).

#### binlog

An informal name for the **binary log** file. For example, you might see this abbreviation used in e-mail messages or forum discussions.

See Also [binary log](#).

#### blind query expansion

A special mode of **full-text search** enabled by the `WITH QUERY EXPANSION` clause. It performs the search twice, where the search phrase for the second search is the original search phrase concatenated with the few most highly relevant documents from the first search. This technique is mainly applicable for short search phrases, perhaps only a single word. It can uncover relevant matches where the precise search term does not occur in the document.

See Also [full-text search](#).

#### bottleneck

A portion of a system that is constrained in size or capacity, that has the effect of limiting overall throughput. For example, a memory area might be smaller than necessary; access to a single required resource might prevent multiple CPU cores from running simultaneously; or waiting for disk I/O to complete might prevent the CPU from running at full capacity. Removing bottlenecks tends to improve **concurrency**. For example, the ability to have multiple InnoDB **buffer pool** instances reduces contention when multiple sessions read from and write to the buffer pool simultaneously.

See Also [buffer pool](#), [concurrency](#).

#### bounce

A **shutdown** operation immediately followed by a restart. Ideally with a relatively short **warmup** period so that performance and throughput quickly return to a high level.

See Also [shutdown](#).

#### buddy allocator

A mechanism for managing different-sized **pages** in the InnoDB **buffer pool**.

See Also [buffer pool](#), [page](#), [page size](#).

#### buffer

A memory or disk area used for temporary storage. Data is buffered in memory so that it can be written to disk efficiently, with a few large I/O operations rather than many small ones. Data is buffered on disk for greater reliability, so that it can be recovered even when a **crash** or other failure occurs at the worst possible time. The main types of buffers used by InnoDB are the **buffer pool**, the **doublewrite buffer**, and the **insert buffer**.

See Also [buffer pool](#), [crash](#), [doublewrite buffer](#), [insert buffer](#).

#### buffer pool

The memory area that holds cached InnoDB data for both tables and indexes. For efficiency of high-volume read operations, the buffer pool is divided into **pages** that can potentially hold multiple rows. For efficiency of cache management, the buffer pool is implemented as a linked list of pages; data that is rarely used is aged out of the cache, using a variation of the **LRU** algorithm. On systems with large memory, you can improve concurrency by dividing the buffer pool into multiple **buffer pool instances**.

Several InnoDB status variables, `information_schema` tables, and `performance_schema` tables help to monitor the internal workings of the buffer pool. Starting in MySQL 5.6, you can also dump and restore the contents of the buffer pool, either automatically during shutdown and restart, or manually at any time, through a set of InnoDB configuration variables such as `innodb_buffer_pool_dump_at_shutdown` [1705] and `innodb_buffer_pool_load_at_startup` [1707].

See Also [buffer pool instance](#), [LRU](#), [page](#), [warm up](#).

---

## buffer pool instance

Any of the multiple regions into which the **buffer pool** can be divided, controlled by the `innodb_buffer_pool_instances` [1706] configuration option. The total memory size specified by the `innodb_buffer_pool_size` [1708] is divided among all the instances. Typically, multiple buffer pool instances are appropriate for systems devoting multiple gigabytes to the InnoDB buffer pool, with each instance 1 gigabyte or larger. On systems loading or looking up large amounts of data in the buffer pool from many concurrent sessions, having multiple instances reduces the contention for exclusive access to the data structures that manage the buffer pool.

See Also [buffer pool](#).

## built-in

The built-in InnoDB storage engine within MySQL is the original form of distribution for the storage engine. Contrast with the **InnoDB Plugin**. Starting with MySQL 5.5, the InnoDB Plugin is merged back into the MySQL code base as the built-in InnoDB storage engine (known as InnoDB 1.1).

This distinction is important mainly in MySQL 5.1, where a feature or bug fix might apply to the InnoDB Plugin but not the built-in InnoDB, or vice versa.

See Also [InnoDB](#), [plugin](#).

## business rules

The relationships and sequences of actions that form the basis of business software, used to run a commercial company. Sometimes these rules are dictated by law, other times by company policy. Careful planning ensures that the relationships encoded and enforced by the database, and the actions performed through application logic, accurately reflect the real policies of the company and can handle real-life situations.

For example, an employee leaving a company might trigger a sequence of actions from the human resources department. The human resources database might also need the flexibility to represent data about a person who has been hired, but not yet started work. Closing an account at an online service might result in data being removed from a database, or the data might be moved or flagged so that it could be recovered if the account is re-opened. A company might establish policies regarding salary maximums, minimums, and adjustments, in addition to basic sanity checks such as the salary not being a negative number. A retail database might not allow a purchase with the same serial number to be returned more than once, or might not allow credit card purchases above a certain value, while a database used to detect fraud might allow these kinds of things.

See Also [relational](#).

# C

## .cfg file

A metadata file used with the **InnoDB transportable tablespace** feature. It is produced by the command `FLUSH TABLES ... FOR EXPORT`, puts one or more tables in a consistent state that can be copied to another server. The `.cfg` file is copied along with the corresponding **.ibd file**, and used to adjust the internal values of the `.ibd` file, such as the **space ID**, during the `ALTER TABLE ... IMPORT TABLESPACE` step.

See Also [.ibd file](#), [space ID](#), [transportable tablespace](#).

## cache

The general term for any memory area that stores copies of data for frequent or high-speed retrieval. In InnoDB, the primary kind of cache structure is the **buffer pool**.

See Also [buffer](#), [buffer pool](#).

## cardinality

The number of different values in a table **column**. When queries refer to columns that have an associated **index**, the cardinality of each column influences which access method is most efficient. For example, for a column with a **unique constraint**, the number of different values is equal to the number of rows in the table. If a table has a million rows but only 10 different values for a particular column, each value occurs (on average) 100,000 times. A query such as `SELECT c1 FROM t1 WHERE c1 = 50;` thus might return 1 row or a huge number of rows, and the database server might process the query differently depending on the cardinality of `c1`.

---

If the values in a column have a very uneven distribution, the cardinality might not be a good way to determine the best query plan. For example, `SELECT c1 FROM t1 WHERE c1 = x;` might return 1 row when `x=50` and a million rows when `x=30`. In such a case, you might need to use **index hints** to pass along advice about which lookup method is more efficient for a particular query.

Cardinality can also apply to the number of distinct values present in multiple columns, as in a **composite index**.

For InnoDB, the process of estimating cardinality for indexes is influenced by the [innodb\\_stats\\_sample\\_pages \[1748\]](#) and the [innodb\\_stats\\_on\\_metadata \[1747\]](#) configuration options. The estimated values are more stable when the **persistent statistics** feature is enabled (in MySQL 5.6 and higher).

See Also [column](#), [composite index](#), [index](#), [index hint](#), [persistent statistics](#), [random dive](#), [selectivity](#), [unique constraint](#).

#### change buffer

A special data structure that records changes to **pages** in **secondary indexes**. These values could result from SQL `INSERT`, `UPDATE`, or `DELETE` statements (**DML**). The set of features involving the change buffer is known collectively as **change buffering**, consisting of **insert buffering**, **delete buffering**, and **purge buffering**.

Changes are only recorded in the change buffer when the relevant page from the secondary index is not in the **buffer pool**. When the relevant index page is brought into the buffer pool while associated changes are still in the change buffer, the changes for that page are applied in the buffer pool (**merged**) using the data from the change buffer. Periodically, the **purge** operation that runs during times when the system is mostly idle, or during a slow shutdown, writes the new index pages to disk. The purge operation can write the disk blocks for a series of index values more efficiently than if each value were written to disk immediately.

Physically, the change buffer is part of the **system tablespace**, so that the index changes remain buffered across database restarts. The changes are only applied (**merged**) when the pages are brought into the buffer pool due to some other read operation.

The kinds and amount of data stored in the change buffer are governed by the [innodb\\_change\\_buffering \[1709\]](#) and [innodb\\_change\\_buffer\\_max\\_size \[1709\]](#) configuration options. To see information about the current data in the change buffer, issue the `SHOW ENGINE INNODB STATUS` command.

Formerly known as the **insert buffer**.

See Also [buffer pool](#), [change buffering](#), [delete buffering](#), [DML](#), [insert buffer](#), [insert buffering](#), [merge](#), [page](#), [purge](#), [purge buffering](#), [secondary index](#), [system tablespace](#).

#### change buffering

The general term for the features involving the **change buffer**, consisting of **insert buffering**, **delete buffering**, and **purge buffering**. Index changes resulting from SQL statements, which could normally involve random I/O operations, are held back and performed periodically by a background **thread**. This sequence of operations can write the disk blocks for a series of index values more efficiently than if each value were written to disk immediately. Controlled by the [innodb\\_change\\_buffering \[1709\]](#) and [innodb\\_change\\_buffer\\_max\\_size \[1709\]](#) configuration options.

See Also [change buffer](#), [delete buffering](#), [insert buffering](#), [purge buffering](#).

#### checkpoint

As changes are made to data pages that are cached in the **buffer pool**, those changes are written to the **data files** sometime later, a process known as **flushing**. The checkpoint is a record of the latest changes (represented by an **LSN** value) that have been successfully written to the data files.

See Also [buffer pool](#), [data files](#), [flush](#), [fuzzy checkpointing](#), [LSN](#).

#### checksum

In **InnoDB**, a validation mechanism to detect corruption when a **page** in a **tablespace** is read from disk into the InnoDB **buffer pool**. This feature is turned on and off by the [innodb\\_checksums \[1711\]](#) configuration option. In MySQL 5.6, you can enable a faster checksum algorithm by also specifying the configuration option [innodb\\_checksum\\_algorithm=crc32](#).

---

The [innochecksum](#) command helps to diagnose corruption problems by testing the checksum values for a specified **tablespace** file while the MySQL server is shut down.

MySQL also uses checksums for replication purposes. For details, see the configuration options [binlog\\_checksum](#) [2009], [master\\_verify\\_checksum](#) [2015], and [slave\\_sql\\_verify\\_checksum](#) [1995].  
See Also [buffer pool](#), [page](#), [tablespace](#).

#### child table

In a **foreign key** relationship, a child table is one whose rows refer (or point) to rows in another table with an identical value for a specific column. This is the table that contains the `FOREIGN KEY . . . REFERENCES` clause and optionally `ON UPDATE` and `ON DELETE` clauses. The corresponding row in the **parent table** must exist before the row can be created in the child table. The values in the child table can prevent delete or update operations on the parent table, or can cause automatic deletion or updates in the child table, based on the `ON CASCADE` option used when creating the foreign key.  
See Also [foreign key](#), [parent table](#).

#### clean page

A **page** in the InnoDB **buffer pool** where all changes made in memory have also been written (**flushed**) to the [data files](#). The opposite of a **dirty page**.  
See Also [buffer pool](#), [data files](#), [dirty page](#), [flush](#), [page](#).

#### clean shutdown

A **shutdown** that completes without errors and applies all changes to InnoDB tables before finishing, as opposed to a **crash** or a **fast shutdown**. Synonym for **slow shutdown**.  
See Also [crash](#), [fast shutdown](#), [shutdown](#), [slow shutdown](#).

#### client

A type of program that sends requests to a server, and interprets or processes the results. The client software might run only some of the time (such as a mail or chat program), and might run interactively (such as the `mysql` command processor).  
See Also [mysql](#), [server](#).

#### clustered index

The InnoDB term for a **primary key** index. InnoDB table storage is organized based on the values of the primary key columns, to speed up queries and sorts involving the primary key columns. For best performance, choose the primary key columns carefully based on the most performance-critical queries. Because modifying the columns of the clustered index is an expensive operation, choose primary columns that are rarely or never updated.

In the Oracle Database product, this type of table is known as an **index-organized table**.  
See Also [index](#), [primary key](#), [secondary index](#).

#### cold backup

A **backup** taken while the database is shut down. For busy applications and web sites, this might not be practical, and you might prefer a **warm backup** or a **hot backup**.  
See Also [backup](#), [hot backup](#), [warm backup](#).

#### column

A data item within a **row**, whose storage and semantics are defined by a data type. Each **table** and **index** is largely defined by the set of columns it contains.

Each column has a **cardinality** value. A column can be the **primary key** for its table, or part of the primary key. A column can be subject to a **unique constraint**, a **NOT NULL constraint**, or both. Values in different columns, even across different tables, can be linked by a **foreign key** relationship.

In discussions of MySQL internal operations, sometimes **field** is used as a synonym, while row is the preferred term when discussing **SQL**.  
See Also [cardinality](#), [foreign key](#), [index](#), [primary key](#), [row](#), [SQL](#), [table](#), [unique constraint](#).



---

#### column index

An **index** on a single column.

See Also [composite index](#), [index](#).

#### column prefix

When an index is created with a length specification, such as `CREATE INDEX idx ON t1 (c1(N))`, only the first N characters of the column value are stored in the index. Keeping the index prefix small makes the index compact, and the memory and disk I/O savings help performance. (Although making the index prefix too small can hinder query optimization by making rows with different values appear to the query optimizer to be duplicates.)

For columns containing binary values or long text strings, where sorting is not a major consideration and storing the entire value in the index would waste space, the index automatically uses the first N (typically 768) characters of the value to do lookups and sorts.

See Also [index](#).

#### commit

A **SQL** statement that ends a **transaction**, making permanent any changes made by the transaction. It is the opposite of **rollback**, which undoes any changes made in the transaction.

InnoDB uses an **optimistic** mechanism for commits, so that changes can be written to the data files before the commit actually occurs. This technique makes the commit itself faster, with the tradeoff that more work is required in case of a rollback.

By default, MySQL uses the **autocommit** setting, which automatically issues a commit following each SQL statement.

See Also [autocommit](#), [optimistic](#), [rollback](#), [SQL](#), [transaction](#).

#### compact row format

The default InnoDB **row format** since MySQL 5.0.3. Available for tables that use the **Antelope file format**. It has a more compact representation for nulls and variable-length fields than the prior default (**redundant row format**).

Because of the **B-tree** indexes that make row lookups so fast in InnoDB, there is little if any performance benefit to keeping all rows the same size.

See Also [Antelope](#), [file format](#), [redundant row format](#), [row format](#).

#### composite index

An **index** that includes multiple columns.

See Also [index](#), [index prefix](#).

#### compressed backup

The compression feature of the **MySQL Enterprise Backup** product makes a compressed copy of each tablespace, changing the extension from `.ibd` to `.ibz`. Compressing the backup data allows you to keep more backups on hand, and reduces the time to transfer backups to a different server. The data is uncompressed during the restore operation. When a compressed backup operation processes a table that is already compressed, it skips the compression step for that table, because compressing again would result in little or no space savings.

A set of files produced by the **MySQL Enterprise Backup** product, where each **tablespace** is compressed. The compressed files are renamed with a `.ibz` file extension.

Applying **compression** right at the start of the backup process helps to avoid storage overhead during the compression process, and to avoid network overhead when transferring the backup files to another server.

The process of **applying** the **binary log** takes longer, and requires uncompressing the backup files.

See Also [apply](#), [binary log](#), [compression](#), [hot backup](#), [MySQL Enterprise Backup](#), [tablespace](#).

#### compressed row format

A **row format** that enables data and index **compression** for InnoDB tables. It was introduced in the InnoDB Plugin, available as part of the **Barracuda** file format. Large fields are stored away from the page that holds the rest of the row data, as in **dynamic row format**. Both index pages and the large fields are compressed,

---

yielding memory and disk savings. Depending on the structure of the data, the decrease in memory and disk usage might or might not outweigh the performance overhead of uncompressing the data as it is used. See [Section 5.4.6, “Working with InnoDB Compressed Tables”](#) for usage details.  
See Also [Barracuda](#), [compression](#), [dynamic row format](#), [row format](#).

#### compression

A feature with wide-ranging benefits from using less disk space, performing less I/O, and using less memory for caching. InnoDB table and index data can be kept in a compressed format during database operation.

The data is uncompressed when needed for queries, and re-compressed when changes are made by **DML** operations. After you enable compression for a table, this processing is transparent to users and application developers. DBAs can consult **information\_schema** tables to monitor how efficiently the compression parameters work for the MySQL instance and for particular compressed tables.

When InnoDB table data is compressed, the compression applies to the **table** itself, any associated **index** data, and the pages loaded into the **buffer pool**. Compression does not apply to pages in the **undo buffer**.

The table compression feature requires using MySQL 5.5 or higher, or the InnoDB Plugin in MySQL 5.1 or earlier, and creating the table using the **Barracuda** file format and **compressed row format**, with the **innodb\_file\_per\_table** setting turned on. The compression for each table is influenced by the **KEY\_BLOCK\_SIZE** clause of the **CREATE TABLE** and **ALTER TABLE** statements. In MySQL 5.6 and higher, compression is also affected by the server-wide configuration options [innodb\\_compression\\_failure\\_threshold\\_pct \[1713\]](#), [innodb\\_compression\\_level \[1713\]](#), and [innodb\\_compression\\_pad\\_pct\\_max \[1713\]](#). See [Section 5.4.6, “Working with InnoDB Compressed Tables”](#) for usage details.

Another type of compression is the **compressed backup** feature of the **MySQL Enterprise Backup** product. See Also [Barracuda](#), [buffer pool](#), [compressed row format](#), [DML](#), [hot backup](#), [index](#), [INFORMATION\\_SCHEMA](#), [innodb\\_file\\_per\\_table](#), [plugin](#), [table](#), [undo buffer](#).

#### compression failure

Not actually an error, rather an expensive operation that can occur when using **compression** in combination with **DML** operations. It occurs when: updates to a compressed **page** overflow the area on the page reserved for recording modifications; the page is compressed again, with all changes applied to the table data; the re-compressed data does not fit on the original page, requiring MySQL to split the data into two new pages and compress each one separately. To check the frequency of this condition, query the table **INFORMATION\_SCHEMA.INNODB\_CMP** and check how much the value of the **COMPRESS\_OPS** column exceeds the value of the **COMPRESS\_OPS\_OK** column. Ideally, compression failures do not occur often; when they do, you can adjust the configuration options [innodb\\_compression\\_level \[1713\]](#), [innodb\\_compression\\_failure\\_threshold\\_pct \[1713\]](#), and [innodb\\_compression\\_pad\\_pct\\_max \[1713\]](#).

See Also [compression](#), [DML](#), [page](#).

#### concatenated index

See [composite index](#).

#### concurrency

The ability of multiple operations (in database terminology, **transactions**) to run simultaneously, without interfering with each other. Concurrency is also involved with performance, because ideally the protection for multiple simultaneous transactions works with a minimum of performance overhead, using efficient mechanisms for **locking**.

See Also [ACID](#), [locking](#), [transaction](#).

#### configuration file

The file that holds the **option** values used by MySQL at startup. Traditionally, on Linux and UNIX this file is named **my.cnf**, and on Windows it is named **my.ini**. You can set a number of options related to InnoDB under the **[mysqld]** section of the file.

Typically, this file is searched for in the locations **/etc/my.cnf** **/etc/mysql/my.cnf** **/usr/local/mysql/etc/my.cnf** and **~/my.cnf**. See [Section 4.2.3.3, “Using Option Files”](#) for details about the search path for this file.



---

When you use the **MySQL Enterprise Backup** product, you typically use two configuration files: one that specifies where the data comes from and how it is structured (which could be the original configuration file for your real server), and a stripped-down one containing only a small set of options that specify where the backup data goes and how it is structured. The configuration files used with the **MySQL Enterprise Backup** product must contain certain options that are typically left out of regular configuration files, so you might need to add some options to your existing configuration file for use with **MySQL Enterprise Backup**. See Also [my.cnf](#), [option file](#).

#### consistent read

A read operation that uses snapshot information to present query results based on a point in time, regardless of changes performed by other transactions running at the same time. If queried data has been changed by another transaction, the original data is reconstructed based on the contents of the **undo log**. This technique avoids some of the **locking** issues that can reduce **concurrency** by forcing transactions to wait for other transactions to finish.

With the **repeatable read** isolation level, the snapshot is based on the time when the first read operation is performed. With the **read committed** isolation level, the snapshot is reset to the time of each consistent read operation.

Consistent read is the default mode in which InnoDB processes **SELECT** statements in **READ COMMITTED** and **REPEATABLE READ** isolation levels. Because a consistent read does not set any locks on the tables it accesses, other sessions are free to modify those tables while a consistent read is being performed on the table.

For technical details about the applicable isolation levels, see [Section 14.2.3.3, “Consistent Nonlocking Reads”](#).

See Also [ACID](#), [concurrency](#), [isolation level](#), [locking](#), [MVCC](#), [READ COMMITTED](#), [READ UNCOMMITTED](#), [REPEATABLE READ](#), [SERIALIZABLE](#), [transaction](#), [undo log](#).

#### constraint

An automatic test that can block database changes to prevent data from becoming inconsistent. (In computer science terms, a kind of assertion related to an invariant condition.) Constraints are a crucial component of the **ACID** philosophy, to maintain data consistency. Constraints supported by MySQL include **FOREIGN KEY constraints** and **unique constraints**.

See Also [ACID](#), [foreign key](#), [relational](#), [unique constraint](#).

#### counter

A value that is incremented by a particular kind of **InnoDB** operation. Useful for measuring how busy a server is, troubleshooting the sources of performance issues, and testing whether changes (for example, to configuration settings or indexes used by queries) have the desired low-level effects. Different kinds of counters are available through **performance\_schema** tables and **information\_schema** tables, particularly [information\\_schema.innodb\\_metrics](#).

See Also [INFORMATION\\_SCHEMA](#), [metrics counter](#), [Performance Schema](#).

#### covering index

An **index** that includes all the columns retrieved by a query. Instead of using the index values as pointers to find the full table rows, the query returns values from the index structure, saving disk I/O. InnoDB can apply this optimization technique to more indexes than MyISAM can, because InnoDB **secondary indexes** also include the primary key columns. InnoDB cannot apply this technique for queries against tables modified by a transactions, until that transaction ends.

Any **column index** or **composite index** could act as a covering index, given the right query. Design your indexes and queries to take advantage of this optimization technique wherever possible.

See Also [column index](#), [composite index](#), [index](#), [secondary index](#).

#### crash

MySQL uses the term "crash" to refer generally to any unexpected [shutdown](#) operation where the server cannot do its normal cleanup. For example, a crash could happen due to a hardware fault on the database server machine or storage device; a power failure; a potential data mismatch that causes the MySQL server to

---

halt; a **fast shutdown** initiated by the DBA; or many other reasons. The robust, automatic **crash recovery** for **InnoDB** tables ensures that data is made consistent when the server is restarted, without any extra work for the DBA.

See Also [crash recovery](#), [fast shutdown](#), [InnoDB](#), [redo log](#), [shutdown](#).

#### crash recovery

The cleanup activities that occur when MySQL is started again after a **crash**. For **InnoDB** tables, changes from incomplete transactions are replayed using data from the **redo log**. Changes that were **committed** before the crash, but not yet written into the [data files](#), are reconstructed from the **doublewrite buffer**. When the database is shut down normally, this type of activity is performed during shutdown by the **purge** operation.

During normal operation, committed data can be stored in the **change buffer** for a period of time before being written to the data files. There is always a tradeoff between keeping the data files up-to-date, which introduces performance overhead during normal operation, and buffering the data, which can make shutdown and crash recovery take longer.

See Also [change buffer](#), [commit](#), [crash](#), [data files](#), [doublewrite buffer](#), [InnoDB](#), [purge](#), [redo log](#).

#### CRUD

Acronym for "create, read, update, delete", a common sequence of operations in database applications. Often denotes a class of applications with relatively simple database usage (basic **DDL**, **DML** and **query** statements in **SQL**) that can be implemented quickly in any language.

See Also [DDL](#), [DML](#), [query](#), [SQL](#).

#### cursor

An internal data structure that is used to represent the result set of a **query**, or other operation that performs a search using an SQL [WHERE](#) clause. It works like an iterator in other high-level languages, producing each value from the result set as requested.

Although usually SQL handles the processing of cursors for you, you might delve into the inner workings when dealing with performance-critical code.

See Also [query](#).

## D

#### data definition language

See [DDL](#).

#### data dictionary

Metadata that keeps track of InnoDB-related objects such as **tables**, **indexes**, and table **columns**. This metadata is physically located in the InnoDB **system tablespace**. For historical reasons, it overlaps to some degree with information stored in the **.frm files**.

Because the **MySQL Enterprise Backup** product always backs up the system tablespace, all backups include the contents of the data dictionary.

See Also [column](#), [.frm file](#), [hot backup](#), [index](#), [MySQL Enterprise Backup](#), [system tablespace](#), [table](#).

#### data directory

The directory under which each MySQL **instance** keeps the **data files** for InnoDB and the directories representing individual databases. Controlled by the [datadir \[472\]](#) configuration option.

See Also [data files](#), [instance](#).

#### data files

The files that physically contain the InnoDB **table** and **index** data. There can be a one-to-many relationship between data files and tables, as in the case of the **system tablespace**, which can hold multiple InnoDB tables as well as the **data dictionary**. There can also be a one-to-one relationship between data files and tables, as when the **file-per-table** setting is enabled, causing each newly created table to be stored in a separate **tablespace**.

See Also [data dictionary](#), [file-per-table](#), [index](#), [system tablespace](#), [table](#), [tablespace](#).

---

data manipulation language  
See [DML](#).

data warehouse

A database system or application that primarily runs large **queries**. The read-only or read-mostly data might be organized in **denormalized** form for query efficiency. Can benefit from the optimizations for **read-only transactions** in MySQL 5.6 and higher. Contrast with **OLTP**.

See Also [denormalized](#), [OLTP](#), [query](#), [read-only transaction](#).

database

Within the MySQL **data directory**, each database is represented by a separate directory. The InnoDB **system tablespace**, which can hold table data from multiple databases within a MySQL **instance**, is kept in its **data files** that reside outside the individual database directories. When **file-per-table** mode is enabled, the **.ibd files** representing individual InnoDB tables are stored inside the database directories.

For long-time MySQL users, a database is a familiar notion. Users coming from an Oracle Database background will find that the MySQL meaning of a database is closer to what Oracle Database calls a **schema**.

See Also [data files](#), [file-per-table](#), [.ibd file](#), [instance](#), [schema](#), [system tablespace](#).

DCL

Data control language, a set of **SQL** statements for managing privileges. In MySQL, consists of the [GRANT](#) and [REVOKE](#) statements. Contrast with **DDL** and **DML**.

See Also [DDL](#), [DML](#), [SQL](#).

DDL

Data definition language, a set of **SQL** statements for manipulating the database itself rather than individual table rows. Includes all forms of the [CREATE](#), [ALTER](#), and [DROP](#) statements. Also includes the [TRUNCATE](#) statement, because it works differently than a [DELETE FROM table\\_name](#) statement, even though the ultimate effect is similar.

DDL statements automatically **commit** the current **transaction**; they cannot be **rolled back**.

The InnoDB-related aspects of DDL include speed improvements for [CREATE INDEX](#) and [DROP INDEX](#) statements, and the way the **file-per-table** setting affects the behavior of the [TRUNCATE TABLE](#) statement.

Contrast with **DML** and **DCL**.

See Also [commit](#), [DCL](#), [DML](#), [file-per-table](#), [rollback](#), [SQL](#), [transaction](#).

deadlock

A situation where different **transactions** are unable to proceed, because each holds a **lock** that the other needs. Because both transactions are waiting for a resource to become available, neither will ever release the locks it holds.

A deadlock can occur when the transactions lock rows in multiple tables (through statements such as [UPDATE](#) or [SELECT ... FOR UPDATE](#)), but in the opposite order. A deadlock can also occur when such statements lock ranges of index records and **gaps**, with each transaction acquiring some locks but not others due to a timing issue.

To reduce the possibility of deadlocks, use transactions rather than [LOCK TABLE](#) statements; keep transactions that insert or update data small enough that they do not stay open for long periods of time; when different transactions update multiple tables or large ranges of rows, use the same order of operations (such as [SELECT ... FOR UPDATE](#)) in each transaction; create indexes on the columns used in [SELECT ... FOR UPDATE](#) and [UPDATE ... WHERE](#) statements. The possibility of deadlocks is not affected by the **isolation level**, because the isolation level changes the behavior of read operations, while deadlocks occur because of write operations.

If a deadlock does occur, InnoDB detects the condition and **rolls back** one of the transactions (the **victim**). Thus, even if your application logic is perfectly correct, you must still handle the case where a transaction must be retried. To see the last deadlock in an InnoDB user transaction, use the command [SHOW ENGINE INNODB STATUS](#). If frequent deadlocks highlight a problem with transaction structure or application error

---

handling, run with the `innodb_print_all_deadlocks` [1741] setting enabled to print information about all deadlocks to the `mysqld` error log.

For background information on how deadlocks are automatically detected and handled, see [Section 14.2.3.9, “Deadlock Detection and Rollback”](#). For tips on avoiding and recovering from deadlock conditions, see [Section 14.2.3.10, “How to Cope with Deadlocks”](#).

See Also [concurrency](#), [gap](#), [isolation level](#), [lock](#), [locking](#), [rollback](#), [transaction](#), [victim](#).

#### deadlock detection

A mechanism that automatically detects when a **deadlock** occurs, and automatically **rolls back** one of the **transactions** involved (the **victim**).

See Also [deadlock](#), [rollback](#), [transaction](#), [victim](#).

#### delete

When InnoDB processes a `DELETE` statement, the rows are immediately marked for deletion and no longer are returned by queries. The storage is reclaimed sometime later, during the periodic garbage collection known as the **purge** operation, performed by a separate thread. For removing large quantities of data, related operations with their own performance characteristics are **truncate** and **drop**.

See Also [drop](#), [purge](#), [truncate](#).

#### delete buffering

The technique of storing index changes due to `DELETE` operations in the **insert buffer** rather than writing them immediately, so that the physical writes can be performed to minimize random I/O. (Because delete operations are a two-step process, this operation buffers the write that normally marks an index record for deletion.) It is one of the types of **change buffering**; the others are **insert buffering** and **purge buffering**.

See Also [change buffer](#), [change buffering](#), [insert buffer](#), [insert buffering](#), [purge buffering](#).

#### denormalized

A data storage strategy that duplicates data across different tables, rather than linking the tables with **foreign keys** and **join** queries. Typically used in **data warehouse** applications, where the data is not updated after loading. In such applications, query performance is more important than making it simple to maintain consistent data during updates. Contrast with **normalized**.

See Also [normalized](#).

#### descending index

A type of index available with some database systems, where index storage is optimized to process `ORDER BY column DESC` clauses. Currently, although MySQL allows the `DESC` keyword in the `CREATE TABLE` statement, it does not use any special storage layout for the resulting index.

See Also [index](#).

#### dirty page

A **page** in the InnoDB **buffer pool** that has been updated in memory, where the changes are not yet written (**flushed**) to the [data files](#). The opposite of a **clean page**.

See Also [buffer pool](#), [clean page](#), [data files](#), [flush](#), [page](#).

#### dirty read

An operation that retrieves unreliable data, data that was updated by another transaction but not yet **committed**. It is only possible with the **isolation level** known as **read uncommitted**.

This kind of operation does not adhere to the **ACID** principle of database design. It is considered very risky, because the data could be **rolled back**, or updated further before being committed; then, the transaction doing the dirty read would be using data that was never confirmed as accurate.

Its polar opposite is **consistent read**, where InnoDB goes to great lengths to ensure that a transaction does not read information updated by another transaction, even if the other transaction commits in the meantime.

See Also [ACID](#), [commit](#), [consistent read](#), [isolation level](#), [READ COMMITTED](#), [READ UNCOMMITTED](#), [rollback](#).

#### disk-based

A kind of database that primarily organizes data on disk storage (hard drives or equivalent). Data is brought back and forth between disk and memory to be operated upon. It is the opposite of an **in-memory database**.

---

Although InnoDB is disk-based, it also contains features such as **the buffer pool**, multiple buffer pool instances, and the **adaptive hash index** that allow certain kinds of workloads to work primarily from memory. See Also [adaptive hash index](#), [buffer pool](#), [in-memory database](#).

#### disk-bound

A type of **workload** where the primary **bottleneck** is CPU operations in memory. Typically involves read-intensive operations where the results can all be cached in the **buffer pool**.

See Also [bottleneck](#), [buffer pool](#), [disk-bound](#), [workload](#).

#### disk-bound

A type of **workload** where the primary **bottleneck** is disk I/O. (Also known as **I/O-bound**.) Typically involves frequent writes to disk, or random reads of more data than can fit into the **buffer pool**.

See Also [bottleneck](#), [buffer pool](#), [disk-bound](#), [workload](#).

#### DML

Data manipulation language, a set of **SQL** statements for performing insert, update, and delete operations. The **SELECT** statement is sometimes considered as a DML statement, because the **SELECT ... FOR UPDATE** form is subject to the same considerations for **locking** as **INSERT**, **UPDATE**, and **DELETE**.

DML statements for an InnoDB table operate in the context of a **transaction**, so their effects can be **committed** or **rolled back** as a single unit.

Contrast with **DDL** and **DCL**.

See Also [commit](#), [DCL](#), [DDL](#), [locking](#), [rollback](#), [SQL](#), [transaction](#).

#### document id

In the InnoDB **full-text search** feature, a special column in the table containing the **FULLTEXT index**, to uniquely identify the document associated with each **ilist** value. Its name is **FTS\_DOC\_ID** (uppercase required). The column itself must be of **BIGINT UNSIGNED NOT NULL** type, with a unique index named **FTS\_DOC\_ID\_INDEX**. Preferably, you define this column when creating the table. If InnoDB must add the column to the table while creating a **FULLTEXT** index, the indexing operation is considerably more expensive. See Also [full-text search](#), [FULLTEXT index](#), [ilist](#).

#### doublewrite buffer

InnoDB uses a novel file flush technique called doublewrite. Before writing **pages** to the **data files**, InnoDB first writes them to a contiguous area called the doublewrite buffer. Only after the write and the flush to the doublewrite buffer have completed, does InnoDB write the pages to their proper positions in the data file. If the operating system crashes in the middle of a page write, InnoDB can later find a good copy of the page from the doublewrite buffer during **crash recovery**.

Although data is always written twice, the doublewrite buffer does not require twice as much I/O overhead or twice as many I/O operations. Data is written to the buffer itself as a large sequential chunk, with a single **fsync()** call to the operating system.

To turn off the doublewrite buffer, specify the option **innodb\_doublewrite=0** [1715].

See Also [crash recovery](#), [data files](#), [page](#), [purge](#).

#### drop

A kind of **DDL** operation that removes a schema object, through a statement such as **DROP TABLE** or **DROP INDEX**. It maps internally to an **ALTER TABLE** statement. From an InnoDB perspective, the performance considerations of such operations involve the time that the **data dictionary** is locked to ensure that interrelated objects are all updated, and the time to update memory structures such as the **buffer pool**. For a **table**, the drop operation has somewhat different characteristics than a **truncate** operation (**TRUNCATE TABLE** statement).

See Also [buffer pool](#), [data dictionary](#), [DDL](#), [table](#), [truncate](#).

#### dynamic row format

A row format introduced in the InnoDB Plugin, available as part of the **Barracuda file format**. Because **TEXT** and **BLOB** fields are stored outside of the rest of the page that holds the row data, it is very efficient for rows

---

that include large objects. Since the large fields are typically not accessed to evaluate query conditions, they are not brought into the **buffer pool** as often, resulting in fewer I/O operations and better utilization of cache memory.

See Also [Barracuda](#), [buffer pool](#), [file format](#), [row format](#).

## E

early adopter

A stage similar to beta, when a software product is typically evaluated for performance, functionality, and compatibility in a non-mission-critical setting. InnoDB uses the **early adopter** designation rather than **beta**, through a succession of point releases leading up to a **GA** release.

See Also [beta](#), [GA](#).

error log

A type of **log** showing information about MySQL startup and critical runtime errors and **crash** information. For details, see [Section 5.2.2, “The Error Log”](#).

See Also [crash](#), [log](#).

eviction

The process of removing an item from a cache or other temporary storage area, such as the InnoDB **buffer pool**. Often, but not always, uses the **LRU** algorithm to determine which item to remove. When a **dirty page** is evicted, its contents are **flushed** to disk, and any **dirty neighbor** pages might be flushed also.

See Also [buffer pool](#), [dirty page](#), [flush](#), [LRU](#).

exclusive lock

A kind of **lock** that prevents any other **transaction** from locking the same row. Depending on the transaction **isolation level**, this kind of lock might block other transactions from writing to the same row, or might also block other transactions from reading the same row. The default InnoDB isolation level, **REPEATABLE READ**, enables higher **concurrency** by allowing transactions to read rows that have exclusive locks, a technique known as **consistent read**.

See Also [concurrency](#), [consistent read](#), [isolation level](#), [lock](#), [REPEATABLE READ](#), [shared lock](#), [transaction](#).

extent

A group of **pages** within a **tablespace** totaling 1 megabyte. With the default **page size** of 16KB, an extent contains 64 pages. In MySQL 5.6, the page size can also be 4KB or 8KB, in which case an extent contains more pages, still adding up to 1MB.

InnoDB features such as **segments**, **read-ahead** requests and the **doublewrite buffer** use I/O operations that read, write, allocate, or free data one extent at a time.

See Also [doublewrite buffer](#), [neighbor page](#), [page](#), [page size](#), [read-ahead](#), [segment](#), [tablespace](#).

## F

.frm file

A file containing the metadata, such as the table definition, of a MySQL table.

Although each InnoDB table has a `.frm` file, InnoDB maintains its own table metadata (the **data dictionary**) in the **system tablespace**; the `.frm` files are not needed for InnoDB to operate on InnoDB tables. Problems can occur if a **crash** happens while the `.frm` file is being written, leading to inconsistencies between the `.frm` file and the data dictionary,

For backups, you must always keep the full set of `.frm` files along with the backup data, to be able to restore tables that are altered or dropped after the backup. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product. If you use the `ibbackup` command instead, you must copy the `.frm` files yourself.

These files are backed up by the **MySQL Enterprise Backup** product. These files must not be modified by an `ALTER TABLE` operation while the backup is taking place, which is why backups that include non-InnoDB



---

tables perform a `FLUSH TABLES WITH READ LOCK` operation to freeze such activity while backing up the `.frm` files. Restoring a backup can result in `.frm` files being created, changed, or removed to match the state of the database at the time of the backup.

See Also [crash](#), [data dictionary](#), [ibbackup command](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#), [system tablespace](#).

#### Fast Index Creation

A capability first introduced in the InnoDB Plugin, now part of the MySQL server in 5.5 and higher, that speeds up creation of InnoDB **secondary indexes** by avoiding the need to completely rewrite the associated table. The speedup applies to dropping secondary indexes also.

Because index maintenance can add performance overhead to many data transfer operations, consider doing operations such as `ALTER TABLE ... ENGINE=INNODB` or `INSERT INTO ... SELECT * FROM ...` without any secondary indexes in place, and creating the indexes afterward.

In MySQL 5.6, this feature becomes more general: you can read and write to tables while an index is being created, and many more kinds of `ALTER TABLE` operations can be performed without copying the table, without blocking **DML** operations, or both. Thus in MySQL 5.6 and higher, we typically refer to this set of features as **online DDL** rather than Fast Index Creation.

See Also [DML](#), [index](#), [online DDL](#), [secondary index](#).

#### fast shutdown

The default **shutdown** procedure for InnoDB, based on the configuration setting `innodb_fast_shutdown=1` [1716]. To save time, certain **flush** operations are skipped. This type of shutdown is safe during normal usage, because the flush operations are performed during the next startup, using the same mechanism as in **crash recovery**. In cases where the database is being shut down for an upgrade or downgrade, do a **slow shutdown** instead to ensure that all relevant changes are applied to the **data files** during the shutdown.

See Also [crash recovery](#), [data files](#), [flush](#), [shutdown](#), [slow shutdown](#).

#### file format

The format used by InnoDB for each table, typically with the **file-per-table** setting enabled so that each table is stored in a separate `.ibd` file. Currently, the file formats available in InnoDB are known as **Antelope** and **Barracuda**. Each file format supports one or more **row formats**. The row formats available for Barracuda tables, **COMPRESSED** and **DYNAMIC**, enable important new storage features for InnoDB tables.

See Also [Antelope](#), [Barracuda](#), [file-per-table](#), [.ibd file](#), [ibdata file](#), [row format](#).

#### file-per-table

A general name for the setting controlled by the `innodb_file_per_table` [1718] option. That is a very important configuration option that affects many aspects of InnoDB file storage, availability of features, and I/O characteristics. In MySQL 5.6.7 and higher, it is enabled by default. Prior to MySQL 5.6.7, it is disabled by default.

For each table created while this setting is in effect, the data is stored in a separate `.ibd` file rather than in the `ibdata` files of the **system tablespace**. When table data is stored in individual files, you have more flexibility to choose nondefault **file formats** and **row formats**, which are required for features such as data **compression**. The `TRUNCATE TABLE` operation is also much faster, and the reclaimed space can be used by the operating system rather than remaining reserved for InnoDB.

The **MySQL Enterprise Backup** product is more flexible for tables that are in their own files. For example, tables can be excluded from a backup, but only if they are in separate files. Thus, this setting is suitable for tables that are backed up less frequently or on a different schedule.

See Also [compressed row format](#), [compression](#), [file format](#), [.ibd file](#), [ibdata file](#), [innodb\\_file\\_per\\_table](#), [row format](#), [system tablespace](#).

#### fill factor

In an InnoDB **index**, the proportion of a **page** that is taken up by index data before the page is split. The unused space when index data is first divided between pages allows for rows to be updated with longer string values without requiring expensive index maintenance operations. If the fill factor is too low, the index

---

consumes more space than needed, causing extra I/O overhead when reading the index. If the fill factor is too high, any update that increases the length of column values can cause extra I/O overhead for index maintenance. See [Section 14.2.3.12.4, “Physical Structure of an InnoDB Index”](#) for more information. See Also [index](#), [page](#).

#### fixed row format

This row format is used by the MyISAM storage engine, not by InnoDB. If you create an InnoDB table with the option `row_format=fixed`, InnoDB translates this option to use the **compact row format** instead, although the `fixed` value might still show up in output such as `SHOW TABLE STATUS` reports. See Also [compact row format](#), [row format](#).

#### flush

To write changes to the database files, that had been buffered in a memory area or a temporary disk storage area. The InnoDB storage structures that are periodically flushed include the **redo log**, the **undo log**, and the **buffer pool**.

Flushing can happen because a memory area becomes full and the system needs to free some space, because a **commit** operation means the changes from a transaction can be finalized, or because a **slow shutdown** operation means that all outstanding work should be finalized. When it is not critical to flush all the buffered data at once, InnoDB can use a technique called **fuzzy checkpointing** to flush small batches of pages to spread out the I/O overhead.

See Also [buffer pool](#), [commit](#), [fuzzy checkpointing](#), [neighbor page](#), [redo log](#), [slow shutdown](#), [undo log](#).

#### flush list

An internal InnoDB data structure that tracks **dirty pages** in the **buffer pool**: that is, **pages** that have been changed and need to be written back out to disk. This data structure is updated frequently by InnoDB's internal **mini-transactions**, and so is protected by its own **mutex** to allow concurrent access to the buffer pool.

See Also [buffer pool](#), [dirty page](#), [LRU](#), [mini-transaction](#), [mutex](#), [page](#), [page cleaner](#).

#### foreign key

A type of pointer relationship, between rows in separate InnoDB tables. The foreign key relationship is defined on one column in both the **parent table** and the **child table**.

In addition to enabling fast lookup of related information, foreign keys help to enforce **referential integrity**, by preventing any of these pointers from becoming invalid as data is inserted, updated, and deleted. This enforcement mechanism is a type of **constraint**. A row that points to another table cannot be inserted if the associated foreign key value does not exist in the other table. If a row is deleted or its foreign key value changed, and rows in another table point to that foreign key value, the foreign key can be set up to prevent the deletion, cause the corresponding column values in the other table to become **null**, or automatically delete the corresponding rows in the other table.

One of the stages in designing a **normalized** database is to identify data that is duplicated, separate that data into a new table, and set up a foreign key relationship so that the multiple tables can be queried like a single table, using a **join** operation.

See Also [child table](#), [FOREIGN KEY constraint](#), [join](#), [normalized](#), [NULL](#), [parent table](#), [referential integrity](#), [relational](#).

#### FOREIGN KEY constraint

The type of **constraint** that maintains database consistency through a **foreign key** relationship. Like other kinds of constraints, it can prevent data from being inserted or updated if data would become inconsistent; in this case, the inconsistency being prevented is between data in multiple tables. Alternatively, when a **DML** operation is performed, **FOREIGN KEY** constraints can cause data in **child rows** to be deleted, changed to different values, or set to **null**, based on the `ON CASCADE` option specified when creating the foreign key.

See Also [child table](#), [constraint](#), [DML](#), [foreign key](#), [NULL](#).

#### FTS

In most contexts, an acronym for **full-text search**. Sometimes in performance discussions, an acronym for **full table scan**.



---

See Also [full table scan](#), [full-text search](#).

#### full backup

A **backup** that includes all the **tables** in each MySQL **database**, and all the databases in a MySQL **instance**. Contrast with **partial backup**.

See Also [backup](#), [database](#), [instance](#), [partial backup](#), [table](#).

#### full table scan

An operation that requires reading the entire contents of a table, rather than just selected portions using an index. Typically performed either with small lookup tables, or in data warehousing situations with large tables where all available data is aggregated and analyzed. How frequently these operations occur, and the sizes of the tables relative to available memory, have implications for the algorithms used in query optimization and managing the buffer pool.

The purpose of **indexes** is to allow lookups for specific values or ranges of values within a large table, thus avoiding full table scans when practical.

See Also [buffer pool](#), [index](#), [LRU](#).

#### full-text search

The MySQL feature for finding words, phrases, Boolean combinations of words, and so on within table data, in a faster, more convenient, and more flexible way than using the SQL [LIKE](#) operator or writing your own application-level search algorithm. It uses the SQL function [MATCH\(\) \[1229\]](#) and **FULLTEXT indexes**.

See Also [FULLTEXT index](#).

#### FULLTEXT index

The special kind of **index** that holds the **search index** in the MySQL **full-text search** mechanism.

Represents the words from values of a column, omitting any that are specified as **stopwords**. Originally, only available for [MyISAM](#) tables. Starting in MySQL 5.6.4, it is also available for **InnoDB** tables.

See Also [full-text search](#), [index](#), [InnoDB](#), [search index](#), [stopword](#).

#### fuzzy checkpointing

A technique that **flushes** small batches of **dirty pages** from the **buffer pool**, rather than flushing all dirty pages at once which would disrupt database processing.

See Also [buffer pool](#), [dirty page](#), [flush](#).

## G

### GA

"Generally available", the stage when a software product leaves beta and is available for sale, official support, and production use.

See Also [beta](#), [early adopter](#).

### gap

A place in an InnoDB **index** data structure where new values could be inserted. When you lock a set of rows with a statement such as [SELECT ... FOR UPDATE](#), InnoDB can create locks that apply to the gaps as well as the actual values in the index. For example, if you select all values greater than 10 for update, a gap lock prevents another transaction from inserting a new value that is greater than 10. The **supremum record** and **infimum record** represent the gaps containing all values greater than or less than all the current index values.

See Also [concurrency](#), [gap lock](#), [index](#), [infimum record](#), [isolation level](#), [supremum record](#).

### gap lock

A **lock** on a **gap** between index records, or a lock on the gap before the first or after the last index record. For example, [SELECT c1 FOR UPDATE FROM t WHERE c1 BETWEEN 10 and 20;](#) prevents other transactions from inserting a value of 15 into the column `t.c1`, whether or not there was already any such value in the column, because the gaps between all existing values in the range are locked. Contrast with **record lock** and **next-key lock**.

Gap locks are part of the tradeoff between performance and **concurrency**, and are used in some transaction **isolation levels** and not others.

---

See Also [gap](#), [infimum record](#), [lock](#), [next-key lock](#), [record lock](#), [supremum record](#).

general log

See [general query log](#).

general query log

A type of **log** used for diagnosis and troubleshooting of SQL statements processed by the MySQL server. Can be stored in a file or in a database table. You must enable this feature through the [general\\_log](#) [485] configuration option to use it. You can disable it for a specific connection through the [sql\\_log\\_off](#) [544] configuration option.

Records a broader range of queries than the **slow query log**. Unlike the **binary log**, which is used for replication, the general query log contains **SELECT** statements and does not maintain strict ordering. For more information, see [Section 5.2.3, "The General Query Log"](#).

See Also [binary log](#), [general query log](#), [log](#).

global\_transaction

A type of **transaction** involved in **XA** operations. It consists of several actions that are transactional in themselves, but that all must either complete successfully as a group, or all be rolled back as a group. In essence, this extends **ACID** properties "up a level" so that multiple ACID transactions can be executed in concert as components of a global operation that also has ACID properties. For this type of distributed transaction, you must use the **SERIALIZABLE** isolation level to achieve ACID properties.

See Also [ACID](#), [SERIALIZABLE](#), [transaction](#), [XA](#).

group commit

An **InnoDB** optimization that performs some low-level I/O operations (**log write**) once for a set of **commit** operations, rather than flushing and syncing separately for each commit.

When the binlog is enabled, you typically also set the configuration option [sync\\_binlog=0](#), because group commit for the binary log is only supported if it is set to 0.

See Also [commit](#), [plugin](#), [XA](#).

## H

hash index

A type of **index** intended for queries that use equality operators, rather than range operators such as greater-than or [BETWEEN](#). It is available for MEMORY tables. Although hash indexes are the default for MEMORY tables for historic reasons, that storage engine also supports **B-tree** indexes, which are often a better choice for general-purpose queries.

MySQL includes a variant of this index type, the **adaptive hash index**, that is constructed automatically for InnoDB tables if needed based on runtime conditions.

See Also [adaptive hash index](#), [B-tree](#), [index](#), [InnoDB](#).

HDD

Acronym for "hard disk drive". Refers to storage media using spinning platters, usually when comparing and contrasting with **SSD**. Its performance characteristics can influence the throughput of a **disk-based** workload.

See Also [disk-based](#), [SSD](#).

heartbeat

A periodic message that is sent to indicate that a system is functioning properly. In a **replication** context, if the **master** stops sending such messages, one of the **slaves** can take its place. Similar techniques can be used between the servers in a cluster environment, to confirm that all of them are operating properly.

See Also [replication](#).

high-water mark

A value representing an upper limit, either a hard limit that should not be exceeded at runtime, or a record of the maximum value that was actually reached. Contrast with **low-water mark**.

See Also [low-water mark](#).

---

## history list

A list of **transactions** with delete-marked records scheduled to be processed by the **InnoDB purge** operation. Recorded in the **undo log**. The length of the history list is reported by the command `SHOW ENGINE INNODB STATUS`. If the history list grows longer than the value of the `innodb_max_purge_lag [1735]` configuration option, each **DML** operation is delayed slightly to allow the purge operation to finish **flushing** the deleted records.

Also known as **purge lag**.

See Also [flush](#), [purge](#), [purge lag](#), [rollback segment](#), [transaction](#), [undo log](#).

## hot

A condition where a row, table, or internal data structure is accessed so frequently, requiring some form of locking or mutual exclusion, that it results in a performance or scalability issue.

Although "hot" typically indicates an undesirable condition, a **hot backup** is the preferred type of backup.

See Also [hot backup](#).

## hot backup

A backup taken while the database and is running and applications are reading and writing to it. The backup involves more than simply copying data files: it must include any data that was inserted or updated while the backup was in process; it must exclude any data that was deleted while the backup was in process; and it must ignore any changes that were not committed.

The Oracle product that performs hot backups, of InnoDB tables especially but also tables from MyISAM and other storage engines, is known as **MySQL Enterprise Backup**.

The hot backup process consists of two stages. The initial copying of the data files produces a **raw backup**. The **apply** step incorporates any changes to the database that happened while the backup was running. Applying the changes produces a **prepared** backup; these files are ready to be restored whenever necessary. See Also [apply](#), [MySQL Enterprise Backup](#), [prepared backup](#), [raw backup](#).

## |

## .ibd file

Each InnoDB **table** created using the **file-per-table** mode goes into its own **tablespace** file, with a `.ibd` extension, inside the **database** directory. This file contains the table data and any **indexes** for the table. File-per-table mode, controlled by the `innodb_file_per_table` option, affects many aspects of InnoDB storage usage and performance, and is enabled by default in MySQL 5.6.7 and higher.

This extension does not apply to the **system tablespace**, which consists of the **ibdata files**.

When a `.ibd` file is included in a compressed backup by the **MySQL Enterprise Backup** product, the compressed equivalent is a `.ibz` file.

If a table is create with the `DATA DIRECTORY =` clause in MySQL 5.6 and higher, the `.ibd` file is located outside the normal database directory, and is pointed to by a `.isl` file.

See Also [database](#), [file-per-table](#), [ibdata file](#), [.ibz file](#), [index](#), [innodb\\_file\\_per\\_table](#), [.isl file](#), [MySQL Enterprise Backup](#), [system tablespace](#), [table](#), [tablespace](#).

## .ibz file

When the **MySQL Enterprise Backup** product performs a **compressed backup**, it transforms each **tablespace** file that is created using the **file-per-table** setting from a `.ibd` extension to a `.ibz` extension.

The compression applied during backup is distinct from the **compressed row format** that keeps table data compressed during normal operation. A compressed backup operation skips the compression step for a tablespace that is already in compressed row format, as compressing a second time would slow down the backup but produce little or no space savings.

See Also [compressed backup](#), [compressed row format](#), [file-per-table](#), [.ibd file](#), [MySQL Enterprise Backup](#), [tablespace](#).

---

## .isl file

A file that specifies the location of a **.ibd file** for an InnoDB table created with the `DATA DIRECTORY =` clause in MySQL 5.6 and higher. It functions like a symbolic link, without the platform restrictions of the actual symbolic link mechanism. You can store InnoDB **tablespaces** outside the **database** directory, for example, on an especially large or fast storage device depending on the usage of the table. For details, see [Section 5.4.1.2, “Specifying the Location of a Tablespace”](#).  
See Also [database](#), [.ibd file](#), [table](#), [tablespace](#).

## I/O-bound

See [disk-bound](#).

## ib-file set

The set of files managed by InnoDB within a MySQL database: the **system tablespace**, any **file-per-table** tablespaces, and the (typically 2) **redo log** files. Used sometimes in detailed discussions of InnoDB file structures and formats, to avoid ambiguity between the meanings of **database** between different DBMS products, and the non-InnoDB files that may be part of a MySQL database.  
See Also [database](#), [file-per-table](#), [redo log](#), [system tablespace](#).

## ibbackup command

The fundamental command-line tool of the **MySQL Enterprise Backup** product. It performs a **hot backup** operation for InnoDB tables. You use this command directly if all your data is in InnoDB tables, if all the important data that you need to back up is in InnoDB tables, or if you are running on the Windows platform. If you also need to back up tables from MyISAM or other storage engines, you use the **mysqlbackup command** instead (available on UNIX and Linux systems only).  
See Also [hot backup](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

## ibbackup\_logfile

A supplemental backup file created by the **MySQL Enterprise Backup** product during a **hot backup** operation. It contains information about any data changes that occurred while the backup was running. The initial backup files, including [ibbackup\\_logfile](#), are known as a **raw backup**, because the changes that occurred during the backup operation are not yet incorporated. After you perform the **apply** step to the raw backup files, the resulting files do include those final data changes, and are known as a **prepared backup**. At this stage, the [ibbackup\\_logfile](#) file is no longer necessary.  
See Also [apply](#), [hot backup](#), [MySQL Enterprise Backup](#), [prepared backup](#), [raw backup](#).

## ibdata file

A set of files with names such as [ibdata1](#), [ibdata2](#), and so on, that make up the InnoDB **system tablespace**. These files contain metadata about InnoDB tables, (the **data dictionary**), and the storage areas for the **undo log**, the **change buffer**, and the **doublewrite buffer**. They also can contain some or all of the table data also (depending on whether the **file-per-table** mode is in effect when each table is created). When the **innodb\_file\_per\_table** option is enabled, data and indexes for newly created tables are stored in separate **.ibd files** rather than in the system tablespace.

The growth of the [ibdata](#) files is influenced by the [innodb\\_autoextend\\_increment \[1704\]](#) configuration option.

See Also [change buffer](#), [data dictionary](#), [doublewrite buffer](#), [file-per-table](#), [.ibd file](#), [innodb\\_file\\_per\\_table](#), [system tablespace](#), [undo log](#).

## ib\_logfile

A set of files, typically named [ib\\_logfile0](#) and [ib\\_logfile1](#), that form the **redo log**. Also sometimes referred to as the **log group**. These files record statements that attempt to change data in InnoDB tables. These statements are replayed automatically to correct data written by incomplete transactions, on startup following a crash.

This data cannot be used for manual recovery; for that type of operation, use the **binary log**.

See Also [binary log](#), [log group](#), [redo log](#).

## ilist

Within an InnoDB **FULLTEXT index**, the data structure consisting of a document ID and positional information for a token (that is, a particular word).

---

See Also [FULLTEXT index](#).

#### implicit row lock

A row lock that InnoDB acquires to ensure consistency, without you specifically requesting it.  
See Also [row lock](#).

#### in-memory database

A type of database system that maintains data in memory, to avoid overhead due to disk I/O and translation between disk blocks and memory areas. Some in-memory databases sacrifice durability (the "D" in the **ACID** design philosophy) and are vulnerable to hardware, power, and other types of failures, making them more suitable for read-only operations. Other in-memory databases do use durability mechanisms such as logging changes to disk or using non-volatile memory.

MySQL features that address the same kinds of memory-intensive processing include the InnoDB **buffer pool**, **adaptive hash index**, and **read-only transaction** optimization, the MEMORY storage engine, the MyISAM key cache, and the MySQL **query cache**.

See Also [ACID](#), [adaptive hash index](#), [buffer pool](#), [disk-based](#), [read-only transaction](#).

#### incremental backup

A type of **hot backup**, performed by the **MySQL Enterprise Backup** product, that only saves data changed since some point in time. Having a full backup and a succession of incremental backups lets you reconstruct backup data over a long period, without the storage overhead of keeping several full backups on hand. You can restore the full backup and then apply each of the incremental backups in succession, or you can keep the full backup up-to-date by applying each incremental backup to it, then perform a single restore operation.

The granularity of changed data is at the **page** level. A page might actually cover more than one row. Each changed page is included in the backup.

See Also [hot backup](#), [MySQL Enterprise Backup](#), [page](#).

#### index

A data structure that provides a fast lookup capability for **rows** of a **table**, typically by forming a tree structure (**B-tree**) representing all the values of a particular **column** or set of columns.

InnoDB tables always have a **clustered index** representing the **primary key**. They can also have one or more **secondary indexes** defined on one or more columns. Depending on their structure, secondary indexes can be classified as **partial**, **column**, or **composite** indexes.

Indexes are a crucial aspect of **query** performance. Database architects design tables, queries, and indexes to allow fast lookups for data needed by applications. The ideal database design uses a **covering index** where practical; the query results are computed entirely from the index, without reading the actual table data. Each **foreign key** constraint also requires an index, to efficiently check whether values exist in both the **parent** and **child** tables.

Although a B-tree index is the most common, a different kind of data structure is used for **hash indexes**, as in the **MEMORY** storage engine and the InnoDB **adaptive hash index**.

See Also [adaptive hash index](#), [B-tree](#), [child table](#), [clustered index](#), [column index](#), [composite index](#), [covering index](#), [foreign key](#), [hash index](#), [parent table](#), [partial index](#), [primary key](#), [query](#), [row](#), [secondary index](#), [table](#).

#### index cache

A memory area that holds the token data for InnoDB **full-text search**. It buffers the data to minimize disk I/O when data is inserted or updated in columns that are part of a **FULLTEXT index**. The token data is written to disk when the index cache becomes full. Each InnoDB **FULLTEXT** index has its own separate index cache, whose size is controlled by the configuration option `innodb_ft_cache_size` [1724].

See Also [full-text search](#), [FULLTEXT index](#).

#### index hint

Extended SQL syntax for overriding the **indexes** recommended by the optimizer. For example, the **FORCE INDEX**, **USE INDEX**, and **IGNORE INDEX** clauses. Typically used when indexed columns have unevenly distributed values, resulting in inaccurate **cardinality** estimates.

---

See Also [cardinality](#), [index](#).

#### index prefix

In an **index** that applies to multiple columns (known as a **composite index**), the initial or leading columns of the index. A query that references the first 1, 2, 3, and so on columns of a composite index can use the index, even if the query does not reference all the columns in the index.

See Also [composite index](#), [index](#).

#### index statistics

See [statistics](#).

#### infimum record

A **pseudo-record** in an **index**, representing the **gap** below the smallest value in that index. If a transaction has a statement such as `SELECT ... FOR UPDATE ... WHERE col < 10;`, and the smallest value in the column is 5, it is a lock on the infimum record that prevents other transactions from inserting even smaller values such as 0, -10, and so on.

See Also [gap](#), [index](#), [pseudo-record](#), [supremum record](#).

#### INFORMATION\_SCHEMA

The name of the **database** that provides a query interface to the MySQL **data dictionary**. (This name is defined by the ANSI SQL standard.) To examine information (metadata) about the database, you can query tables such as `INFORMATION_SCHEMA.TABLES` and `INFORMATION_SCHEMA.COLUMNS`, rather than using `SHOW` commands that produce unstructured output.

The information schema contains some tables that are specific to **InnoDB**, such as `INNODB_LOCKS` and `INNODB_TRX`. You use these tables not to see how the database is structured, but to get real-time information about the workings of InnoDB tables to help with performance monitoring, tuning, and troubleshooting. In particular, these tables provide information about MySQL features related to **compression**, and **transactions** and their associated **locks**.

See Also [compression](#), [data dictionary](#), [database](#), [InnoDB](#), [lock](#), [transaction](#).

#### innobackup command

A **backup** command from the former **InnoDB Hot Backup** product. It is superseded by the `mysqlbackup` command in the **MySQL Enterprise Backup** product.

See Also [backup](#), [InnoDB Hot Backup](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

#### InnoDB

A MySQL component that combines high performance with **transactional** capability for reliability, robustness, and concurrent access. It embodies the **ACID** design philosophy. Represented as a **storage engine**; it handles tables created or altered with the `ENGINE=INNODB` clause. See [Section 14.2, “The InnoDB Storage Engine”](#) for architectural details and administration procedures, and [Section 8.5, “Optimizing for InnoDB Tables”](#) for performance advice.

In MySQL 5.5 and higher, InnoDB is the default storage engine for new tables and the `ENGINE=INNODB` clause is not required. In MySQL 5.1 only, many of the advanced InnoDB features require enabling the component known as the InnoDB Plugin. See [Section 14.2.1.1, “InnoDB as the Default MySQL Storage Engine”](#) for the considerations involved in transitioning to recent releases where InnoDB tables are the default.

InnoDB tables are ideally suited for **hot backups**. See [Section 23.2, “MySQL Enterprise Backup”](#) for information about the **MySQL Enterprise Backup** product for backing up MySQL servers without interrupting normal processing.

See Also [ACID](#), [hot backup](#), [storage engine](#), [transaction](#).

#### InnoDB Hot Backup

A licensed backup product, superseded in MySQL 5.1 and above by **MySQL Enterprise Backup**.

See Also [MySQL Enterprise Backup](#).

#### innodb\_autoinc\_lock\_mode

The `innodb_autoinc_lock_mode [1705]` option controls the algorithm used for **auto-increment locking**. When you have an auto-incrementing **primary key**, you can use statement-based replication only with



---

the setting `innodb_autoinc_lock_mode=1` [1705]. This setting is known as **consecutive** lock mode, because multi-row inserts within a transaction receive consecutive auto-increment values. If you have `innodb_autoinc_lock_mode=2`, which allows higher concurrency for insert operations, use row-based replication rather than statement-based replication. This setting is known as **interleaved** lock mode, because multiple multi-row insert statements running at the same time can receive autoincrement values that are interleaved. The setting `innodb_autoinc_lock_mode=0` is the previous (traditional) default setting and should not be used except for compatibility purposes.

See Also [auto-increment locking](#), [mixed-mode insert](#), [primary key](#).

#### innodb\_file\_format

The `innodb_file_format` [1716] option determines the **file format** for all InnoDB **tablespaces** created after you specify a value for this option. To create tablespaces other than the **system tablespace**, you must also use the **file-per-table** option. Currently, you can specify the **Antelope** and **Barracuda** file formats. See Also [Antelope](#), [Barracuda](#), [file format](#), [file-per-table](#), [innodb\\_file\\_per\\_table](#), [system tablespace](#), [tablespace](#).

#### innodb\_file\_per\_table

A very important configuration option that affects many aspects of InnoDB file storage, availability of features, and I/O characteristics. In MySQL 5.6.7 and higher, it is enabled by default. Prior to MySQL 5.6.7, it is disabled by default. The `innodb_file_per_table` [1718] option turns on **file-per-table** mode, which stores each newly created InnoDB table and its associated index in its own **.ibd file**, outside the **system tablespace**.

This option affects the performance and storage considerations for a number of SQL statements, such as `DROP TABLE` and `TRUNCATE TABLE`.

This option is needed to take full advantage of many other InnoDB features, such as such as table **compression**, or backups of named tables in **MySQL Enterprise Backup**.

This option was once static, but can now be set using the `SET GLOBAL` command.

For reference information, see `innodb_file_per_table` [1718]. For usage information, see [Section 5.4.1](#), “[Managing InnoDB Tablespaces](#)”.

See Also [compression](#), [file-per-table](#), [.ibd file](#), [MySQL Enterprise Backup](#), [system tablespace](#).

#### innodb\_lock\_wait\_timeout

The `innodb_lock_wait_timeout` [1729] option sets the balance between **waiting** for shared resources to become available, or giving up and handling the error, retrying, or doing alternative processing in your application. Rolls back any InnoDB transaction that waits more than a specified time to acquire a **lock**. Especially useful if **deadlocks** are caused by updates to multiple tables controlled by different storage engines; such deadlocks are not **detected** automatically.

See Also [deadlock](#), [deadlock detection](#), [lock](#), [wait](#).

#### innodb\_strict\_mode

The `innodb_strict_mode` [1749] option controls whether InnoDB operates in **strict mode**, where conditions that are normally treated as warnings, cause errors instead (and the underlying statements fail).

This mode is the default setting in MySQL 5.5.5 and higher.

See Also [strict mode](#).

#### insert

One of the primary **DML** operations in **SQL**. The performance of inserts is a key factor in **data warehouse** systems that load millions of rows into tables, and **OLTP** systems where many concurrent connections might insert rows into the same table, in arbitrary order. If insert performance is important to you, you should learn about **InnoDB** features such as the **insert buffer** used in **change buffering**, and **auto-increment** columns. See Also [auto-increment](#), [change buffering](#), [data warehouse](#), [DML](#), [InnoDB](#), [insert buffer](#), [OLTP](#), [SQL](#).

#### insert buffer

Former name for the **change buffer**. Now that **change buffering** includes delete and update operations as well as inserts, “change buffer” is the preferred term.

---

See Also [change buffer](#), [change buffering](#).

#### insert buffering

The technique of storing secondary index changes due to [INSERT](#) operations in the **insert buffer** rather than writing them immediately, so that the physical writes can be performed to minimize random I/O. It is one of the types of **change buffering**; the others are **delete buffering** and **purge buffering**.

Insert buffering is not used if the secondary index is **unique**, because the uniqueness of new values cannot be verified before the new entries are written out. Other kinds of change buffering do work for unique indexes. See Also [change buffer](#), [change buffering](#), [delete buffering](#), [insert buffer](#), [purge buffering](#), [unique index](#).

#### instance

A single **mysqld** daemon managing a **data directory** representing one or more **databases** with a set of **tables**. It is common in development, testing, and some **replication** scenarios to have multiple instances on the same **server** machine, each managing its own data directory and listening on its own port or socket. With one instance running a **disk-bound** workload, the server might still have extra CPU and memory capacity to run additional instances.

See Also [data directory](#), [database](#), [disk-bound](#), [mysqld](#), [replication](#), [server](#).

#### instrumentation

Modifications at the source code level to collect performance data for tuning and debugging. In MySQL, data collected by instrumentation is exposed through a SQL interface using the [INFORMATION\\_SCHEMA](#) and [PERFORMANCE\\_SCHEMA](#) databases.

See Also [INFORMATION\\_SCHEMA](#), [Performance Schema](#).

#### intention exclusive lock

See [intention lock](#).

#### intention lock

A kind of **lock** that applies to the table level, used to indicate what kind of lock the transaction intends to acquire on rows in the table. Different transactions can acquire different kinds of intention locks on the same table, but the first transaction to acquire an **intention exclusive** (IX) lock on a table prevents other transactions from acquiring any S or X locks on the table. Conversely, the first transaction to acquire an **intention shared** (IS) lock on a table prevents other transactions from acquiring any X locks on the table. The two-phase process allows the lock requests to be resolved in order, without blocking locks and corresponding operations that are compatible. For more details on this locking mechanism, see [Section 14.2.3.2, “InnoDB Lock Modes”](#).

See Also [lock](#), [lock mode](#), [locking](#).

#### intention shared lock

See [intention lock](#).

#### inverted index

A data structure optimized for document retrieval systems, used in the implementation of InnoDB **full-text search**. The InnoDB **FULLTEXT index**, implemented as an inverted index, records the position of each word within a document, rather than the location of a table row. A single column value (a document stored as a text string) is represented by many entries in the inverted index.

See Also [full-text search](#), [FULLTEXT index](#), [ilist](#).

#### IOPS

Acronym for **I/O operations per second**. A common measurement for busy systems, particularly **OLTP** applications. If this value is near the maximum that the storage devices can handle, the application can become **disk-bound**, limiting **scalability**.

See Also [disk-bound](#), [OLTP](#), [scalability](#).

#### isolation level

One of the foundations of database processing. Isolation is the **I** in the acronym **ACID**; the isolation level is the setting that fine-tunes the balance between performance and reliability, consistency, and reproducibility of results when multiple **transactions** are making changes and performing queries at the same time.



---

From highest amount of consistency and protection to the least, the isolation levels supported by InnoDB are: **SERIALIZABLE**, **REPEATABLE READ**, **READ COMMITTED**, and **READ UNCOMMITTED**.

With InnoDB tables, many users can keep the default isolation level (**REPEATABLE READ**) for all operations. Expert users might choose the **read committed** level as they push the boundaries of scalability with OLTP processing, or during data warehousing operations where minor inconsistencies do not affect the aggregate results of large amounts of data. The levels on the edges (**SERIALIZABLE** and **READ UNCOMMITTED**) change the processing behavior to such an extent that they are rarely used. See Also [ACID](#), [READ COMMITTED](#), [READ UNCOMMITTED](#), [REPEATABLE READ](#), [SERIALIZABLE](#), [transaction](#).

## J

join

A **query** that retrieves data from more than one table, by referencing columns in the tables that hold identical values. Ideally, these columns are part of an InnoDB **foreign key** relationship, which ensures **referential integrity** and that the join columns are **indexed**. Often used to save space and improve query performance by replacing repeated strings with numeric IDs, in a **normalized** data design. See Also [foreign key](#), [index](#), [normalized](#), [query](#), [referential integrity](#).

## K

KEY\_BLOCK\_SIZE

An option to specify the size of data pages within an InnoDB table that uses **compressed row format**. The default is 8 kilobytes. Lower values risk hitting internal limits that depend on the combination of row size and compression percentage. See Also [compressed row format](#).

## L

latch

A lightweight structure used by InnoDB to implement a **lock** for its own internal memory structures, typically held for a brief time measured in milliseconds or microseconds. A general term that includes both **mutexes** (for exclusive access) and **rw-locks** (for shared access). Certain latches are the focus of InnoDB performance tuning, such as the **data dictionary** mutex. Statistics about latch use and contention are available through the **Performance Schema** interface. See Also [data dictionary](#), [lock](#), [locking](#), [mutex](#), [Performance Schema](#), [rw-lock](#).

list

The InnoDB **buffer pool** is represented as a list of memory **pages**. The list is reordered as new pages are accessed and enter the buffer pool, as pages within the buffer pool are accessed again and are considered newer, and as pages that are not accessed for a long time are **evicted** from the buffer pool. The buffer pool is actually divided into **sublists**, and the replacement policy is a variation of the familiar **LRU** technique. See Also [buffer pool](#), [eviction](#), [LRU](#), [sublist](#).

lock

The high-level notion of an object that controls access to a resource, such as a table, row, or internal data structure, as part of a **locking** strategy. For intensive performance tuning, you might delve into the actual structures that implement locks, such as **mutexes** and **latches**. See Also [latch](#), [lock mode](#), [locking](#), [mutex](#).

lock escalation

An operation used in some database systems that converts many row locks into a single table lock, saving memory space but reducing concurrent access to the table. InnoDB uses a space-efficient representation for row locks, so that lock escalation is not needed. See Also [locking](#), [row lock](#), [table lock](#).

---

## lock mode

A shared (S) lock allows a transaction to read a row. Multiple transactions can acquire an S lock on that same row at the same time.

An exclusive (X) lock allows a transaction to update or delete a row. No other transaction can acquire any kind of lock on that same row at the same time.

**Intention locks** apply to the table level, and are used to indicate what kind of lock the transaction intends to acquire on rows in the table. Different transactions can acquire different kinds of intention locks on the same table, but the first transaction to acquire an intention exclusive (IX) lock on a table prevents other transactions from acquiring any S or X locks on the table. Conversely, the first transaction to acquire an intention shared (IS) lock on a table prevents other transactions from acquiring any X locks on the table. The two-phase process allows the lock requests to be resolved in order, without blocking locks and corresponding operations that are compatible.

See Also [intention lock](#), [lock](#), [locking](#).

## locking

The system of protecting a **transaction** from seeing or changing data that is being queried or changed by other transactions. The locking strategy must balance reliability and consistency of database operations (the principles of the **ACID** philosophy) against the performance needed for good **concurrency**. Fine-tuning the locking strategy often involves choosing an **isolation level** and ensuring all your database operations are safe and reliable for that isolation level.

See Also [ACID](#), [concurrency](#), [isolation level](#), [latch](#), [lock](#), [mutex](#), [transaction](#).

## locking read

A **SELECT** statement that also performs a **locking** operation on an [InnoDB](#) table. Either `SELECT ... FOR UPDATE` or `SELECT ... LOCK IN SHARE MODE`. It has the potential to produce a **deadlock**, depending on the **isolation level** of the transaction. The opposite of a **non-locking read**. Not allowed for global tables in a **read-only transaction**.

See Also [deadlock](#), [isolation level](#), [locking](#), [non-locking read](#), [read-only transaction](#).

## log

In the [InnoDB](#) context, "log" or "log files" typically refers to the **redo log** represented by the `ib_logfile*` files. Another log area, which is physically part of the **system tablespace**, is the **undo log**.

Other kinds of logs that are important in MySQL are the **error log** (for diagnosing startup and runtime problems), **binary log** (for working with replication and performing point-in-time restores), the **general query log** (for diagnosing application problems), and the **slow query log** (for diagnosing performance problems).

See Also [binary log](#), [error log](#), [general query log](#), [ib\\_logfile](#), [redo log](#), [slow query log](#), [system tablespace](#), [undo log](#).

## log buffer

The memory area that holds data to be written to the **log files** that make up the **redo log**. It is controlled by the `innodb_log_buffer_size` [1732] configuration option.

See Also [log file](#), [redo log](#).

## log file

One of the `ib_logfileN` files that make up the **redo log**. Data is written to these files from the **log buffer** memory area.

See Also [ib\\_logfile](#), [log buffer](#), [redo log](#).

## log group

The set of files that make up the **redo log**, typically named `ib_logfile0` and `ib_logfile1`. (For that reason, sometimes referred to collectively as **ib\_logfile**.)

See Also [ib\\_logfile](#), [redo log](#).

## logical

A type of operation that involves high-level, abstract aspects such as tables, queries, indexes, and other SQL concepts. Typically, logical aspects are important to make database administration and application development convenient and usable. Contrast with **physical**.

See Also [logical backup](#), [physical](#).

---

## logical backup

A **backup** that reproduces table structure and data, without copying the actual data files. For example, the [mysqldump](#) command produces a logical backup, because its output contains statements such as [CREATE TABLE](#) and [INSERT](#) that can re-create the data. Contrast with **physical backup**. A logical backup offers flexibility (for example, you could edit table definitions or insert statements before restoring), but can take substantially longer to **restore** than a physical backup.

See Also [backup](#), [mysqldump](#), [physical backup](#), [restore](#).

## loose\_

In MySQL 5.1, a prefix added to InnoDB configuration options when installing the InnoDB **Plugin** after server startup, so any new configuration options not recognized by the current level of MySQL do not cause a startup failure. MySQL processes configuration options that start with this prefix, but gives a warning rather than a failure if the part after the prefix is not a recognized option.

See Also [plugin](#).

## low-water mark

A value representing a lower limit, typically a threshold value at which some corrective action begins or becomes more aggressive. Contrast with **high-water mark**.

See Also [high-water mark](#).

## LRU

An acronym for "least recently used", a common method for managing storage areas. The items that have not been used recently are **evicted** when space is needed to cache newer items. InnoDB uses the LRU mechanism by default to manage the **pages** within the **buffer pool**, but makes exceptions in cases where a page might be read only a single time, such as during a **full table scan**. This variation of the LRU algorithm is called the **midpoint insertion strategy**. The ways in which the buffer pool management differs from the traditional LRU algorithm is fine-tuned by the options [innodb\\_old\\_blocks\\_pct](#) [1738], [innodb\\_old\\_blocks\\_time](#) [1738], and the new MySQL 5.6 options [innodb\\_lru\\_scan\\_depth](#) [1734] and [innodb\\_flush\\_neighbors](#) [1721].

See Also [buffer pool](#), [eviction](#), [full table scan](#), [midpoint insertion strategy](#).

## LSN

Acronym for "log sequence number". This arbitrary, ever-increasing value represents a point in time corresponding to operations recorded in the **redo log**. (This point in time is regardless of **transaction** boundaries; it can fall in the middle of one or more transactions.) It is used internally by InnoDB during **crash recovery** and for managing the buffer pool.

In the **MySQL Enterprise Backup** product, you can specify an LSN to represent the point in time from which to take an **incremental backup**. The relevant LSN is displayed by the output of the [ibbackup](#) command. Once you have the LSN corresponding to the time of a full backup, you can specify that value to take a subsequent incremental backup, whose output contains another LSN for the next incremental backup.

See Also [crash recovery](#), [incremental backup](#), [MySQL Enterprise Backup](#), [redo log](#), [transaction](#).

# M

## .MRG file

A file containing references to other tables, used by the [MERGE](#) storage engine. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [MySQL Enterprise Backup](#), [mysqlbackup command](#).

## .MYD file

The file that MySQL uses to store data for a MyISAM table. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [.MYI file](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

## .MYI file

A file that MySQL uses to store indexes for a MyISAM table. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [.MYD file](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#).

---

## master server

Frequently shortened to "master". A database server machine in a **replication** scenario that processes the initial insert, update, and delete requests for data. These changes are propagated to, and repeated on, other servers known as **slave servers**.

See Also [replication](#), [slave server](#).

## master thread

An InnoDB **thread** that performs various tasks in the background. Most of these tasks are I/O related, such as writing changes from the **insert buffer** to the appropriate secondary indexes.

To improve **concurrency**, sometimes actions are moved from the master thread to separate background threads. For example, in MySQL 5.6 and higher, **dirty pages** are **flushed** from the **buffer pool** by the **page cleaner** thread rather than the master thread.

See Also [buffer pool](#), [dirty page](#), [flush](#), [insert buffer](#), [page cleaner](#), [thread](#).

## MDL

Acronym for "metadata lock".

See Also [metadata lock](#).

## memcached

A popular component of many MySQL and **NoSQL** software stacks, allowing fast reads and writes for single values and caching the results entirely in memory. Traditionally, applications required extra logic to write the same data to a MySQL database for permanent storage, or to read data from a MySQL database when it was not cached yet in memory. Now, applications can use the simple [memcached](#) protocol, supported by client libraries for many languages, to communicate directly with MySQL servers using **InnoDB** or MySQL Cluster tables. These NoSQL interfaces to MySQL tables allow applications to achieve higher read and write performance than by issuing SQL commands directly, and can simplify application logic and deployment configurations for systems that already incorporated [memcached](#) for in-memory caching.

The [memcached](#) interface to InnoDB tables is available in MySQL 5.6 and higher; see [Section 14.2.9, "InnoDB Integration with memcached"](#) for details. The [memcached](#) interface to MySQL Cluster tables is available in MySQL Cluster 7.2; see <http://dev.mysql.com/doc/ndbapi/en/ndbmemcache.html> for details. See Also [InnoDB](#), [NoSQL](#).

## merge

To apply changes to data cached in memory, such as when a page is brought into the **buffer pool**, and any applicable changes recorded in the **change buffer** are incorporated into the page in the buffer pool. The updated data is eventually written to the **tablespace** by the **flush** mechanism.

See Also [buffer pool](#), [change buffer](#), [flush](#), [tablespace](#).

## metadata lock

A type of **lock** that prevents **DDL** operations on a table that is being used at the same time by another **transaction**. For details, see [Section 8.10.4, "Metadata Locking"](#).

Enhancements to **online** operations, particularly in MySQL 5.6 and higher, are focused on reducing the amount of metadata locking. The objective is for DDL operations that do not change the table structure (such as [CREATE INDEX](#) and [DROP INDEX](#) for [InnoDB](#) tables) to proceed while the table is being queried, updated, and so on by other transactions.

See Also [DDL](#), [lock](#), [online](#), [transaction](#).

## metrics counter

A feature implemented by the [innodb\\_metrics](#) table in the **information\_schema**, in MySQL 5.6 and higher. You can query **counts** and totals for low-level InnoDB operations, and use the results for performance tuning in combination with data from the **performance\_schema**.

See Also [counter](#), [INFORMATION\\_SCHEMA](#), [Performance Schema](#).

## midpoint insertion strategy

The technique of initially bringing **pages** into the InnoDB **buffer pool** not at the "newest" end of the list, but instead somewhere in the middle. The exact location of this point can vary, based on the setting of the [innodb\\_old\\_blocks\\_pct \[1738\]](#) option. The intent is that blocks that are only read once, such as during a **full table scan**, can be aged out of the buffer pool sooner than with a strict **LRU** algorithm.

---

See Also [buffer pool](#), [full table scan](#), [LRU](#), [page](#).

#### mini-transaction

An internal phase of InnoDB processing, when making changes at the **physical** level to internal data structures during **DML** operations. A mini-transaction has no notion of **rollback**; multiple mini-transactions can occur within a single **transaction**. Mini-transactions write information to the **redo log** that is used during **crash recovery**. A mini-transaction can also happen outside the context of a regular transaction, for example during **purge** processing by background threads.

See Also [commit](#), [crash recovery](#), [DML](#), [physical](#), [purge](#), [redo log](#), [rollback](#), [transaction](#).

#### mixed-mode insert

An **INSERT** statement where **auto-increment** values are specified for some but not all of the new rows. For example, a multi-value **INSERT** could specify a value for the auto-increment column in some cases and **NULL** in other cases. InnoDB generates auto-increment values for the rows where the column value was specified as **NULL**. Another example is an **INSERT . . . ON DUPLICATE KEY UPDATE** statement, where auto-increment values might be generated but not used, for any duplicate rows that are processed as **UPDATE** rather than **INSERT** statements.

Can cause consistency issues between **master** and **slave** servers in a **replication** configuration. Can require adjusting the value of the **innodb\_autoinc\_lock\_mode** configuration option.

See Also [auto-increment](#), [innodb\\_autoinc\\_lock\\_mode](#), [master server](#), [replication](#), [slave server](#).

#### multi-core

A type of processor that can take advantage of multi-threaded programs, such as the MySQL server.

#### multiversion concurrency control

See [MVCC](#).

#### mutex

Informal abbreviation for "mutex variable". (Mutex itself is short for "mutual exclusion".) The low-level object that InnoDB uses to represent and enforce exclusive-access **locks** to internal in-memory data structures. Once the lock is acquired, any other process, thread, and so on is prevented from acquiring the same lock. Contrast with **rw-locks**, which allow shared access. Mutexes and rw-locks are known collectively as **latches**. See Also [latch](#), [lock](#), [Performance Schema](#), [Pthreads](#), [rw-lock](#).

#### MVCC

Acronym for "multiversion concurrency control". This technique lets InnoDB **transactions** with certain **isolation levels** to perform **consistent read** operations; that is, to query rows that are being updated by other transactions, and see the values from before those updates occurred. This is a powerful technique to increase **concurrency**, by allowing queries to proceed without waiting due to **locks** held by the other transactions.

This technique is not universal in the database world. Some other database products, and some other MySQL storage engines, do not support it.

See Also [ACID](#), [concurrency](#), [consistent read](#), [isolation level](#), [lock](#), [transaction](#).

#### my.cnf

The name, on UNIX or Linux systems, of the MySQL option file.

See Also [my.ini](#), [option file](#).

#### my.ini

The name, on Windows systems, of the MySQL option file.

See Also [my.cnf](#), [option file](#).

#### mysql

The **mysql** program is the command-line interpreter for the MySQL database. It processes **SQL** statements, and also MySQL-specific commands such as **SHOW TABLES**, by passing requests to the **mysqld** daemon.

See Also [mysqld](#), [SQL](#).

#### MySQL Enterprise Backup

A licensed product, superseding InnoDB Hot Backup, that performs **hot backups** of MySQL databases. It offers the most efficiency and flexibility when backing up **InnoDB** tables, but can also back up MyISAM and other kinds of tables.

---

See Also [hot backup](#), [InnoDB](#).

#### mysqlbackup command

A command-line tool of the **MySQL Enterprise Backup** product. It performs a **hot backup** operation for InnoDB tables, and a **warm backup** for MyISAM and other kinds of tables. See [Section 23.2, “MySQL Enterprise Backup”](#) for more information about this command.

See Also [hot backup](#), [ibbackup command](#), [MySQL Enterprise Backup](#), [warm backup](#).

#### mysqld

The `mysqld` program is the database engine for the MySQL database. It runs as a UNIX daemon or Windows service, constantly waiting for requests and performing maintenance work in the background.

See Also [mysql](#).

#### mysqldump

A command that performs a **logical backup** of some combination of databases, tables, and table data. The results are SQL statements that reproduce the original schema objects, data, or both. For substantial amounts of data, a **physical backup** solution such as **MySQL Enterprise Backup** is faster, particularly for the **restore** operation.

See Also [logical backup](#), [MySQL Enterprise Backup](#), [physical backup](#), [restore](#).

## N

#### natural key

A indexed column, typically a **primary key**, where the values have some real-world significance. Usually advised against because:

- If the value should ever change, there is potentially a lot of index maintenance to re-sort the **clustered index** and update the copies of the primary key value that are repeated in each **secondary index**.
- Even seemingly stable values can change in unpredictable ways that are difficult to represent correctly in the database. For example, one country can change into two or several, making the original country code obsolete. Or, rules about unique values might have exceptions. For example, even if taxpayer IDs are intended to be unique to a single person, a database might have to handle records that violate that rule, such as in cases of identity theft. Taxpayer IDs and other sensitive ID numbers also make poor primary keys, because they may need to be secured, encrypted, and otherwise treated differently than other columns.

Thus, it is typically better to use arbitrary numeric values to form a **synthetic key**, for example using an **auto-increment** column.

See Also [auto-increment](#), [primary key](#), [secondary index](#), [synthetic key](#).

#### neighbor page

Any **page** in the same **extent** as a particular page. When a page is selected to be **flushed**, any neighbor pages that are **dirty** are typically flushed as well, as an I/O optimization for traditional hard disks. In MySQL 5.6 and up, this behavior can be controlled by the configuration variable `innodb_flush_neighbors` [1721]; you might turn that setting off for SSD drives, which do not have the same overhead for writing smaller batches of data at random locations.

See Also [dirty page](#), [extent](#), [flush](#), [page](#).

#### next-key lock

A combination of a **record lock** on the index record and a [gap lock](#) on the gap before the index record.

See Also [gap lock](#), [locking](#), [record lock](#).

#### non-blocking I/O

An industry term that means the same as **asynchronous I/O**.

See Also [asynchronous I/O](#).

#### non-locking read

A **query** that does not use the `SELECT ... FOR UPDATE` or `SELECT ... LOCK IN SHARE MODE` clauses. The only kind of query allowed for global tables in a **read-only transaction**. The opposite of a **locking read**.



---

See Also [locking read](#), [query](#), [read-only transaction](#).

#### non-repeatable read

The situation when a query retrieves data, and a later query within the same **transaction** retrieves what should be the same data, but the queries return different results (changed by another transaction committing in the meantime).

This kind of operation goes against the **ACID** principle of database design. Within a transaction, data should be consistent, with predictable and stable relationships.

Among different **isolation levels**, non-repeatable reads are prevented by the **serializable read** and **repeatable read** levels, and allowed by the **consistent read**, and **read uncommitted** levels.

See Also [ACID](#), [consistent read](#), [isolation level](#), [READ UNCOMMITTED](#), [REPEATABLE READ](#), [SERIALIZABLE](#), [transaction](#).

#### normalized

A database design strategy where data is split into multiple tables, and duplicate values condensed into single rows represented by an ID, to avoid storing, querying, and updating redundant or lengthy values. It is typically used in **OLTP** applications.

For example, an address might be given a unique ID, so that a census database could represent the relationship **lives at this address** by associating that ID with each member of a family, rather than storing multiple copies of a complex value such as **123 Main Street, Anytown, USA**.

For another example, although a simple address book application might store each phone number in the same table as a person's name and address, a phone company database might give each phone number a special ID, and store the numbers and IDs in a separate table. This normalized representation could simplify large-scale updates when area codes split apart.

Normalization is not always recommended. Data that is primarily queried, and only updated by deleting entirely and reloading, is often kept in fewer, larger tables with redundant copies of duplicate values. This data representation is referred to as **denormalized**, and is frequently found in data warehousing applications.

See Also [denormalized](#), [foreign key](#), [OLTP](#), [relational](#).

#### NoSQL

A broad term for a set of data access technologies that do not use the **SQL** language as their primary mechanism for reading and writing data. Some NoSQL technologies act as key-value stores, only accepting single-value reads and writes; some relax the restrictions of the **ACID** methodology; still others do not require a pre-planned **schema**. MySQL users can combine NoSQL-style processing for speed and simplicity with SQL operations for flexibility and convenience, by using the **memcached** API to directly access some kinds of MySQL tables. The **memcached** interface to InnoDB tables is available in MySQL 5.6 and higher; see [Section 14.2.9, "InnoDB Integration with memcached"](#) for details. The **memcached** interface to MySQL Cluster tables is available in MySQL Cluster 7.2; see <http://dev.mysql.com/doc/ndbapi/en/ndbmemcache.html> for details.

See Also [ACID](#), [InnoDB](#), [memcached](#), [schema](#), [SQL](#).

#### NOT NULL constraint

A type of **constraint** that specifies that a **column** cannot contain any **NULL** values. It helps to preserve **referential integrity**, as the database server can identify data with erroneous missing values. It also helps in the arithmetic involved in query optimization, allowing the optimizer to predict the number of entries in an index on that column.

See Also [column](#), [constraint](#), [NULL](#), [primary key](#), [referential integrity](#).

#### NULL

A special value in **SQL**, indicating the absence of data. Any arithmetic operation or equality test involving a **NULL** value, in turn produces a **NULL** result. (Thus it is similar to the IEEE floating-point concept of NaN, "not a number".) Any aggregate calculation such as **AVG()** ignores rows with **NULL** values, when determining how many rows to divide by. The only test that works with **NULL** values uses the SQL idioms **IS NULL** or **IS NOT NULL**.

---

[NULL](#) values play a part in index operations, because for performance a database must minimize the overhead of keeping track of missing data values. Typically, [NULL](#) values are not stored in an index, because a query that tests an indexed column using a standard comparison operator could never match a row with a [NULL](#) value for that column. For the same reason, unique indexes do not prevent [NULL](#) values; those values simply are not represented in the index. Declaring a [NOT NULL](#) constraint on a column provides reassurance that there are no rows left out of the index, allowing for better query optimization (accurate counting of rows and estimation of whether to use the index).

Because the **primary key** must be able to uniquely identify every row in the table, a single-column primary key cannot contain any [NULL](#) values, and a multi-column primary key cannot contain any rows with [NULL](#) values in all columns.

Although the Oracle database allows a [NULL](#) value to be concatenated with a string, InnoDB treats the result of such an operation as [NULL](#).

See Also [index](#), [primary key](#), [SQL](#).

## O

### .OPT file

A file containing database configuration information. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [MySQL Enterprise Backup](#), [mysqlbackup command](#).

### off-page column

A column containing variable-length data (such as [BLOB](#) and [VARCHAR](#)) that is too long to fit on a **B-tree** page. The data is stored in **overflow pages**. The [DYNAMIC](#) row format in the InnoDB **Barracuda** file format is more efficient for such storage than the older [COMPACT](#) row format.

See Also [B-tree](#), [Barracuda](#), [overflow page](#).

### OLTP

Acronym for "Online Transaction Processing". A database system, or a database application, that runs a workload with many **transactions**, with frequent writes as well as reads, typically affecting small amounts of data at a time. For example, an airline reservation system or an application that processes bank deposits. The data might be organized in **normalized** form for a balance between **DML** (insert/update/delete) efficiency and **query** efficiency. Contrast with **data warehouse**.

With its **row-level locking** and **transactional** capability, **InnoDB** is the ideal storage engine for MySQL tables used in OLTP applications.

See Also [data warehouse](#), [DML](#), [InnoDB](#), [query](#), [row lock](#), [transaction](#).

### online

A type of operation that involves no downtime, blocking, or restricted operation for the database. Typically applied to **DDL**. Operations that shorten the periods of restricted operation, such as **fast index creation**, have evolved into a wider set of **online DDL** operations in MySQL 5.6.

In the context of backups, a **hot backup** is an online operation and a **warm backup** is partially an online operation.

See Also [DDL](#), [Fast Index Creation](#), [hot backup](#), [online DDL](#), [warm backup](#).

### online DDL

A feature that improves the performance, concurrency, and availability of InnoDB tables during **DDL** (primarily [ALTER TABLE](#)) operations. See [Section 5.5, "Online DDL for InnoDB Tables"](#) for details.

The details vary according to the type of operation. In some cases, the table can be modified concurrently while the [ALTER TABLE](#) is in progress. The operation might be able to be performed without doing a table copy, or using a specially optimized type of table copy. Space usage is controlled by the [innodb\\_online\\_alter\\_log\\_max\\_size \[1739\]](#) configuration option.

This feature is an enhancement of the **Fast Index Creation** feature in MySQL 5.5 and the InnoDB Plugin for MySQL 5.1.



---

See Also [DDL](#), [Fast Index Creation](#), [online](#).

#### optimistic

A methodology that guides low-level implementation decisions for a relational database system. The requirements of performance and **concurrency** in a relational database mean that operations must be started or dispatched quickly. The requirements of consistency and **referential integrity** mean that any operation could fail: a transaction might be rolled back, a **DML** operation could violate a constraint, a request for a lock could cause a deadlock, a network error could cause a timeout. An optimistic strategy is one that assumes most requests or attempts will succeed, so that relatively little work is done to prepare for the failure case. When this assumption is true, the database does little unnecessary work; when requests do fail, extra work must be done to clean up and undo changes.

InnoDB uses optimistic strategies for operations such as **locking** and **commits**. For example, data changed by a transaction can be written to the data files before the commit occurs, making the commit itself very fast, but requiring more work to undo the changes if the transaction is rolled back.

The opposite of an optimistic strategy is a **pessimistic** one, where a system is optimized to deal with operations that are unreliable and frequently unsuccessful. This methodology is rare in a database system, because so much care goes into choosing reliable hardware, networks, and algorithms.

See Also [commit](#), [concurrency](#), [DML](#), [locking](#), [pessimistic](#).

#### optimizer

The MySQL component that determines the best **indexes** and **join** order to use for a **query**, based on characteristics and data distribution of the relevant **tables**.

See Also [index](#), [join](#), [query](#), [table](#).

#### option

A configuration parameter for MySQL, either stored in the **option file** or passed on the command line.

For the options that apply to **InnoDB** tables, each option name starts with the prefix `innodb_`.

See Also [InnoDB](#), [option file](#).

#### option file

The file that holds the configuration **options** for the MySQL instance. Traditionally, on Linux and UNIX this file is named `my.cnf`, and on Windows it is named `my.ini`.

See Also [configuration file](#), [my.cnf](#), [option](#).

#### overflow page

Separately allocated disk **pages** that hold variable-length columns (such as [BLOB](#) and [VARCHAR](#)) that are too long to fit on a **B-tree** page. The associated columns are known as **off-page columns**.

See Also [B-tree](#), [off-page column](#), [page](#).

## P

#### .PAR file

A file containing partition definitions. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [MySQL Enterprise Backup](#), [mysqlbackup command](#).

#### page

A unit representing how much data InnoDB transfers at any one time between disk (the **data files**) and memory (the **buffer pool**). A page can contain one or more **rows**, depending on how much data is in each row. If a row does not fit entirely into a single page, InnoDB sets up additional pointer-style data structures so that the information about the row can be stored in one page.

One way to fit more data in each page is to use **compressed row format**. For tables that use BLOBs or large text fields, **compact row format** allows those large columns to be stored separately from the rest of the row, reducing I/O overhead and memory usage for queries that do not reference those columns.

When InnoDB reads or writes sets of pages as a batch to increase I/O throughput, it reads or writes an **extent** at a time.

---

All the InnoDB disk data structures within a MySQL instance share the same **page size**.  
See Also [buffer pool](#), [compact row format](#), [compressed row format](#), [data files](#), [extent](#), [page size](#), [row](#).

#### page cleaner

An InnoDB background **thread** that **flushes dirty pages** from the **buffer pool**. Prior to MySQL 5.6, this activity was performed by the **master thread**  
See Also [buffer pool](#), [dirty page](#), [flush](#), [master thread](#), [thread](#).

#### page size

For releases up to and including MySQL 5.5, the size of each InnoDB **page** is fixed at 16 kilobytes. This value represents a balance: large enough to hold the data for most rows, yet small enough to minimize the performance overhead of transferring unneeded data to memory. Other values are not tested or supported.

Starting in MySQL 5.6, the page size for an InnoDB **instance** can be either 4KB, 8KB, or 16KB, controlled by the [innodb\\_page\\_size](#) [1740] configuration option. You set the size when creating the MySQL instance, and it remains constant afterwards. The same page size applies to all InnoDB **tablespaces**, both the **system tablespace** and any separate tablespaces created in **file-per-table** mode.

Smaller page sizes can help performance with storage devices that use small block sizes, particularly for **SSD** devices in **disk-bound** workloads, such as for **OLTP** applications. As individual rows are updated, less data is copied into memory, written to disk, reorganized, locked, and so on.

See Also [disk-bound](#), [file-per-table](#), [instance](#), [OLTP](#), [page](#), [SSD](#), [system tablespace](#), [tablespace](#).

#### parent table

The table in a **foreign key** relationship that holds the initial column values pointed to from the **child table**. The consequences of deleting, or updating rows in the parent table depend on the [ON UPDATE](#) and [ON DELETE](#) clauses in the foreign key definition. Rows with corresponding values in the child table could be automatically deleted or updated in turn, or those columns could be set to [NULL](#), or the operation could be prevented.

See Also [child table](#), [foreign key](#).

#### partial backup

A **backup** that contains some of the **tables** in a MySQL database, or some of the databases in a MySQL instance. Contrast with **full backup**.

See Also [backup](#), [full backup](#), [table](#).

#### partial index

An **index** that represents only part of a column value, typically the first N characters (the **prefix**) of a long [VARCHAR](#) value.

See Also [index](#), [index prefix](#).

#### Performance Schema

The [performance\\_schema](#) schema, in MySQL 5.5 and up, presents a set of tables that you can query to get detailed information about the performance characteristics of many internal parts of the MySQL server.

See Also [latch](#), [mutex](#), [rw-lock](#).

#### persistent statistics

A feature in MySQL 5.6 that stores **index** statistics for InnoDB **tables** on disk, providing better **plan stability** for **queries**.

See Also [index](#), [optimizer](#), [plan stability](#), [query](#), [table](#).

#### pessimistic

A methodology that sacrifices performance or concurrency in favor of safety. It is appropriate if a high proportion of requests or attempts might fail, or if the consequences of a failed request are severe. InnoDB uses what is known as a pessimistic **locking** strategy, to minimize the chance of **deadlocks**. At the application level, you might avoid deadlocks by using a pessimistic strategy of acquiring all locks needed by a transaction at the very beginning.

Many built-in database mechanisms use the opposite **optimistic** methodology.

See Also [deadlock](#), [locking](#), [optimistic](#).

---

## phantom

A row that appears in the result set of a query, but not in the result set of an earlier query. For example, if a query is run twice within a **transaction**, and in the meantime, another transaction commits after inserting a new row or updating a row so that it matches the [WHERE](#) clause of the query.

This occurrence is known as a phantom read. It is harder to guard against than a **non-repeatable read**, because locking all the rows from the first query result set does not prevent the changes that cause the phantom to appear.

Among different **isolation levels**, phantom reads are prevented by the **serializable read** level, and allowed by the **repeatable read**, **consistent read**, and **read uncommitted** levels.

See Also [consistent read](#), [isolation level](#), [non-repeatable read](#), [READ UNCOMMITTED](#), [REPEATABLE READ](#), [SERIALIZABLE](#), [transaction](#).

## physical

A type of operation that involves hardware-related aspects such as disk blocks, memory pages, files, bits, disk reads, and so on. Typically, physical aspects are important during expert-level performance tuning and problem diagnosis. Contrast with **logical**.

See Also [logical](#), [physical backup](#).

## physical backup

A **backup** that copies the actual data files. For example, the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product produces a physical backup, because its output contains data files that can be used directly by the `mysqld` server, resulting in a faster **restore** operation. Contrast with **logical backup**.

See Also [backup](#), [logical backup](#), [MySQL Enterprise Backup](#), [restore](#).

## PITR

Acronym for **point-in-time recovery**.

See Also [point-in-time recovery](#).

## plan stability

A property of a **query execution plan**, where the optimizer makes the same choices each time for a given **query**, so that performance is consistent and predictable.

See Also [query](#), [query execution plan](#).

## plugin

In MySQL 5.1 and earlier, a separately installable form of the **InnoDB** storage engine that includes features and performance enhancements not included in the **built-in** InnoDB for those releases.

For MySQL 5.5 and higher, the MySQL distribution includes the very latest InnoDB features and performance enhancements, known as InnoDB 1.1, and there is no longer a separate InnoDB Plugin.

This distinction is important mainly in MySQL 5.1, where a feature or bug fix might apply to the InnoDB Plugin but not the built-in InnoDB, or vice versa.

See Also [built-in](#), [InnoDB](#).

## point-in-time recovery

The process of restoring a **backup** to recreate the state of the database at a specific date and time.

Commonly abbreviated **PITR**. Because it is unlikely that the specified time corresponds exactly to the time of a backup, this technique usually requires a combination of a **physical backup** and a **logical backup**. For example, with the **MySQL Enterprise Backup** product, you restore the last backup that you took before the specified point in time, then replay changes from the **binary log** between the time of the backup and the PITR time.

See Also [backup](#), [logical backup](#), [MySQL Enterprise Backup](#), [physical backup](#), [PITR](#).

## prefix

See [index prefix](#).

## prepared backup

A set of backup files, produced by the **MySQL Enterprise Backup** product, after all the stages of applying **binary logs** and **incremental backups** are finished. The resulting files are ready to be **restored**. Prior to the apply steps, the files are known as a **raw backup**.

---

See Also [binary log](#), [hot backup](#), [incremental backup](#), [MySQL Enterprise Backup](#), [raw backup](#), [restore](#).

#### primary key

A set of columns -- and by implication, the index based on this set of columns -- that can uniquely identify every row in a table. As such, it must be a unique index that does not contain any [NULL](#) values.

InnoDB requires that every table has such an index (also called the **clustered index** or **cluster index**), and organizes the table storage based on the column values of the primary key.

When choosing primary key values, consider using arbitrary values (a **synthetic key**) rather than relying on values derived from some other source (a **natural key**).

See Also [clustered index](#), [index](#), [natural key](#), [synthetic key](#).

#### process

An instance of an executing program. The operating system switches between multiple running processes, allowing for a certain degree of **concurrency**. On most operating systems, processes can contain multiple **threads** of execution that share resources. Context-switching between threads is faster than the equivalent switching between processes.

See Also [concurrency](#), [thread](#).

#### pseudo-record

An artificial record in an index, used for **locking** key values or ranges that do not currently exist.

See Also [infimum record](#), [locking](#), [supremum record](#).

#### Pthreads

The POSIX threads standard, which defines an API for threading and locking operations on UNIX and Linux systems. On UNIX and Linux systems, InnoDB uses this implementation for **mutexes**.

See Also [mutex](#).

#### purge

A type of garbage collection performed by a separate thread, running on a periodic schedule. The purge includes these actions: removing obsolete values from indexes; physically removing rows that were marked for deletion by previous [DELETE](#) statements.

See Also [crash recovery](#), [delete](#), [doublewrite buffer](#).

#### purge buffering

The technique of storing index changes due to [DELETE](#) operations in the **insert buffer** rather than writing them immediately, so that the physical writes can be performed to minimize random I/O. (Because delete operations are a two-step process, this operation buffers the write that normally purges an index record that was previously marked for deletion.) It is one of the types of **change buffering**; the others are **insert buffering** and **delete buffering**.

See Also [change buffer](#), [change buffering](#), [delete buffering](#), [insert buffer](#), [insert buffering](#).

#### purge lag

Another name for the [InnoDB history list](#). Related to the [innodb\\_max\\_purge\\_lag](#) [1735] configuration option.

See Also [history list](#), [purge](#).

#### purge thread

A **thread** within the InnoDB process that is dedicated to performing the periodic **purge** operation. In MySQL 5.6 and higher, multiple purge threads are enabled by the [innodb\\_purge\\_threads](#) [1742] configuration option.

See Also [purge](#), [thread](#).

## Q

#### query

In **SQL**, an operation that reads information from one or more **tables**. Depending on the organization of data and the parameters of the query, the lookup might be optimized by consulting an **index**. If multiple tables are involved, the query is known as a **join**.

---

For historical reasons, sometimes discussions of internal processing for statements use "query" in a broader sense, including other types of MySQL statements such as **DDL** and **DML** statements.  
See Also [DDL](#), [DML](#), [index](#), [join](#), [SQL](#), [table](#).

#### query execution plan

The set of decisions made by the optimizer about how to perform a **query** most efficiently, including which **index** or indexes to use, and the order in which to **join** tables. **Plan stability** involves the same choices being made consistently for a given query.  
See Also [index](#), [join](#), [plan stability](#), [query](#).

#### query log

See [general query log](#).

#### quiesce

To reduce the amount of database activity, often in preparation for an operation such as an [ALTER TABLE](#), a **backup**, or a **shutdown**. Might or might not involve doing as much **flushing** as possible, so that **InnoDB** does not continue doing background I/O.

In MySQL 5.6 and higher, the syntax [FLUSH TABLES ... FOR EXPORT](#) writes some data to disk for **InnoDB** tables that make it simpler to back up those tables by copying the data files.  
See Also [backup](#), [flush](#), [InnoDB](#), [shutdown](#).

## R

### RAID

Acronym for "Redundant Array of Inexpensive Drives". Spreading I/O operations across multiple drives enables greater **concurrency** at the hardware level, and improves the efficiency of low-level write operations that otherwise would be performed in sequence.  
See Also [concurrency](#).

#### random dive

A technique for quickly estimating the number of different values in a column (the column's cardinality). **InnoDB** samples pages at random from the index and uses that data to estimate the number of different values. This operation occurs when each table is first opened.

Originally, the number of sampled pages was fixed at 8; now, it is determined by the setting of the [innodb\\_stats\\_sample\\_pages \[1748\]](#) parameter.

The way the random pages are picked depends on the setting of the [innodb\\_use\\_legacy\\_cardinality\\_algorithm](#) parameter. The default setting (OFF) has better randomness than in older releases.  
See Also [cardinality](#).

#### raw backup

The initial set of backup files produced by the **MySQL Enterprise Backup** product, before the changes reflected in the **binary log** and any **incremental backups** are applied. At this stage, the files are not ready to **restore**. After these changes are applied, the files are known as a **prepared backup**.  
See Also [binary log](#), [hot backup](#), [ibbackup\\_logfile](#), [incremental backup](#), [MySQL Enterprise Backup](#), [prepared backup](#), [restore](#).

### READ COMMITTED

An **isolation level** that uses a **locking** strategy that relaxes some of the protection between **transactions**, in the interest of performance. Transactions cannot see uncommitted data from other transactions, but they can see data that is committed by another transaction after the current transaction started. Thus, a transaction never sees any bad data, but the data that it does see may depend to some extent on the timing of other transactions.

When a transaction with this isolation level performs [UPDATE ... WHERE](#) or [DELETE ... WHERE](#) operations, other transactions might have to wait. The transaction can perform [SELECT ... FOR UPDATE](#), and [LOCK IN SHARE MODE](#) operations without making other transactions wait.

---

See Also [ACID](#), [isolation level](#), [locking](#), [REPEATABLE READ](#), [SERIALIZABLE](#), [transaction](#).

## READ UNCOMMITTED

The **isolation level** that provides the least amount of protection between transactions. Queries employ a **locking** strategy that allows them to proceed in situations where they would normally wait for another transaction. However, this extra performance comes at the cost of less reliable results, including data that has been changed by other transactions and not committed yet (known as **dirty read**). Use this isolation level only with great caution, and be aware that the results might not be consistent or reproducible, depending on what other transactions are doing at the same time. Typically, transactions with this isolation level do only queries, not insert, update, or delete operations.

See Also [ACID](#), [dirty read](#), [isolation level](#), [locking](#), [transaction](#).

## read view

An internal snapshot used by the **MVCC** mechanism of InnoDB. Certain **transactions**, depending on their **isolation level**, see the data values as they were at the time the transaction (or in some cases, the statement) started. Isolation levels that use a read view are **REPEATABLE READ**, **READ COMMITTED**, and **READ UNCOMMITTED**.

See Also [isolation level](#), [MVCC](#), [READ COMMITTED](#), [READ UNCOMMITTED](#), [REPEATABLE READ](#), [transaction](#).

## read-ahead

A type of I/O request that prefetches a group of **pages** (an entire **extent**) into the **buffer pool** asynchronously, in anticipation that these pages will be needed soon. The linear read-ahead technique prefetches all the pages of one extent based on access patterns for pages in the preceding extent, and is part of all MySQL versions starting with the InnoDB Plugin for MySQL 5.1. The random read-ahead technique prefetches all the pages for an extent once a certain number of pages from the same extent are in the buffer pool. Random read-ahead is not part of MySQL 5.5, but is re-introduced in MySQL 5.6 under the control of the [innodb\\_random\\_read\\_ahead](#) configuration option.

See Also [buffer pool](#), [extent](#), [page](#).

## read-only transaction

A type of transaction that can be optimized for InnoDB tables by eliminating some of the bookkeeping involved with creating a **read view** for each transaction. Can only perform **non-locking read** queries. It can be started explicitly with the syntax `START TRANSACTION READ ONLY`, or automatically under certain conditions. See [Section 14.2.4.2.3, "Optimizations for Read-Only Transactions"](#) for details.

See Also [non-locking read](#), [read view](#), [transaction](#).

## record lock

A **lock** on an index record. For example, `SELECT c1 FOR UPDATE FROM t WHERE c1 = 10;` prevents any other transaction from inserting, updating, or deleting rows where the value of `t.c1` is 10. Contrast with **gap lock** and **next-key lock**.

See Also [gap lock](#), [lock](#), [next-key lock](#).

## redo

The data, in units of records, recorded in the **redo log** when **DML** statements make changes to InnoDB tables. It is used during **crash recovery** to correct data written by incomplete **transactions**. The ever-increasing **LSN** value represents the cumulative amount of redo data that has passed through the redo log.

See Also [crash recovery](#), [DML](#), [LSN](#), [redo log](#), [transaction](#).

## redo log

A disk-based data structure used during **crash recovery**, to correct data written by incomplete **transactions**. During normal operation, it encodes requests to change InnoDB table data, which result from SQL statements or low-level API calls through NoSQL interfaces. Modifications that did not finish updating the **data files** before an unexpected **shutdown** are replayed automatically.

The redo log is physically represented as a set of files, typically named `ib_logfile0` and `ib_logfile1`. The data in the redo log is encoded in terms of records affected; this data is collectively referred to as **redo**. The passage of data through the redo logs is represented by the ever-increasing **LSN** value. The original 4GB limit on maximum size for the redo log is raised to 512GB in MySQL 5.6.



---

The disk layout of the redo log is influenced by the configuration options [innodb\\_log\\_file\\_size \[1732\]](#), [innodb\\_log\\_group\\_home\\_dir \[1733\]](#), and (rarely) [innodb\\_log\\_files\\_in\\_group \[1733\]](#).

The performance of redo log operations is also affected by the **log buffer**, which is controlled by the [innodb\\_log\\_buffer\\_size \[1732\]](#) configuration option.

See Also [crash recovery](#), [data files](#), [ib\\_logfile](#), [log buffer](#), [LSN](#), [redo](#), [shutdown](#), [transaction](#).

#### redundant row format

The oldest InnoDB row format, available for tables using the **Antelope file format**. Prior to MySQL 5.0.3, it was the only row format available in InnoDB. In MySQL 5.0.3 and later, the default is **compact row format**. You can still specify redundant row format for compatibility with older InnoDB tables.

See Also [Antelope](#), [compact row format](#), [file format](#), [row format](#).

#### referential integrity

The technique of maintaining data always in a consistent format, part of the **ACID** philosophy. In particular, data in different tables is kept consistent through the use of **foreign key constraints**, which can prevent changes from happening or automatically propagate those changes to all related tables. Related mechanisms include the **unique constraint**, which prevents duplicate values from being inserted by mistake, and the **NOT NULL constraint**, which prevents blank values from being inserted by mistake.

See Also [ACID](#), [FOREIGN KEY constraint](#), [NOT NULL constraint](#), [unique constraint](#).

#### relational

An important aspect of modern database systems. The database server encodes and enforces relationships such as one-to-one, one-to-many, many-to-one, and uniqueness. For example, a person might have zero, one, or many phone numbers in an address database; a single phone number might be associated with several family members. In a financial database, a person might be required to have exactly one taxpayer ID, and any taxpayer ID could only be associated with one person.

The database server can use these relationships to prevent bad data from being inserted, and to find efficient ways to look up information. For example, if a value is declared to be unique, the server can stop searching as soon as the first match is found, and it can reject attempts to insert a second copy of the same value.

At the database level, these relationships are expressed through SQL features such as **columns** within a table, unique and **NOT NULL constraints**, **foreign keys**, and different kinds of join operations. Complex relationships typically involve data split between more than one table. Often, the data is **normalized**, so that duplicate values in one-to-many relationships are stored only once.

In a mathematical context, the relations within a database are derived from set theory. For example, the **OR** and **AND** operators of a **WHERE** clause represent the notions of union and intersection.

See Also [ACID](#), [constraint](#), [foreign key](#), [normalized](#).

#### relevance

In the **full-text search** feature, a number signifying the similarity between the search string and the data in the **FULLTEXT index**. For example, when you search for a single word, that word is typically more relevant for a row where it occurs several times in the text than a row where it appears only once.

See Also [full-text search](#), [FULLTEXT index](#).

#### REPEATABLE READ

The default **isolation level** for InnoDB. It prevents any rows that are queried from being changed by other transactions, thus blocking **non-repeatable reads** but not **phantom** reads. It uses a moderately strict **locking** strategy so that all queries within a transaction see data from the same snapshot, that is, the data as it was at the time the transaction started.

When a transaction with this isolation level performs **UPDATE ... WHERE**, **DELETE ... WHERE**, **SELECT ... FOR UPDATE**, and **LOCK IN SHARE MODE** operations, other transactions might have to wait.

See Also [ACID](#), [consistent read](#), [isolation level](#), [locking](#), [phantom](#), [SERIALIZABLE](#), [transaction](#).

#### replication

The practice of sending changes from a **master database**, to one or more **slave databases**, so that all databases have the same data. This technique has a wide range of uses, such as load-balancing for better scalability, disaster recovery, and testing software upgrades and configuration changes. The changes can be sent between the database by methods called **row-based replication** and **statement-based replication**.

---

See Also [row-based replication](#), [statement-based replication](#).

#### restore

The process of putting a set of backup files from the **MySQL Enterprise Backup** product in place for use by MySQL. This operation can be performed to fix a corrupted database, to return to some earlier point in time, or (in a **replication** context) to set up a new **slave database**. In the **MySQL Enterprise Backup** product, this operation is performed by the `copy-back` option of the `mysqlbackup` command.

See Also [hot backup](#), [MySQL Enterprise Backup](#), [mysqlbackup command](#), [prepared backup](#), [replication](#).

#### rollback

A **SQL** statement that ends a **transaction**, undoing any changes made by the transaction. It is the opposite of **commit**, which makes permanent any changes made in the transaction.

By default, MySQL uses the **autocommit** setting, which automatically issues a commit following each SQL statement. You must change this setting before you can use the rollback technique.

See Also [ACID](#), [commit](#), [transaction](#).

#### rollback segment

The storage area containing the **undo log**, part of the **system tablespace**.

See Also [system tablespace](#), [undo log](#).

#### row

The logical data structure defined by a set of **columns**. A set of rows makes up a **table**. Within InnoDB **data files**, each **page** can contain one or more rows.

Although InnoDB uses the term **row format** for consistency with MySQL syntax, the row format is a property of each table and applies to all rows in that table.

See Also [column](#), [data files](#), [page](#), [row format](#), [table](#).

#### row format

The disk storage format for a **row** from an InnoDB **table**. As InnoDB gains new capabilities such as compression, new row formats are introduced to support the resulting improvements in storage efficiency and performance.

Each table has its own row format, specified through the `ROW_FORMAT` option. To see the row format for each InnoDB table, issue the command `SHOW TABLE STATUS`. Because all the tables in the system tablespace share the same row format, to take advantage of other row formats typically requires setting the `innodb_file_per_table [1718]` option, so that each table is stored in a separate tablespace.

See Also [compact row format](#), [compressed row format](#), [dynamic row format](#), [fixed row format](#), [redundant row format](#), [row](#), [table](#).

#### row lock

A **lock** that prevents a row from being accessed in an incompatible way by another **transaction**. Other rows in the same table can be freely written to by other transactions. This is the type of **locking** done by **DML** operations on **InnoDB** tables.

Contrast with **table locks** used by MyISAM, or during **DDL** operations on InnoDB tables that cannot be done with **online DDL**; those locks block concurrent access to the table.

See Also [DDL](#), [DML](#), [InnoDB](#), [lock](#), [locking](#), [online DDL](#), [table lock](#), [transaction](#).

#### row-based replication

A form of **replication** where events are propagated from the **master** server specifying how to change individual rows on the **slave** server. It is safe to use for all settings of the `innodb_autoinc_lock_mode [1705]` option.

See Also [auto-increment locking](#), [innodb\\_autoinc\\_lock\\_mode](#), [replication](#), [statement-based replication](#).

#### row-level locking

The **locking** mechanism used for **InnoDB** tables, relying on **row locks** rather than **table locks**. Multiple **transactions** can modify the same table concurrently. Only if two transactions try to modify the same row does one of the transactions wait for the other to complete (and release its row locks).

See Also [InnoDB](#), [locking](#), [row lock](#), [table lock](#), [transaction](#).



---

rw-lock

The low-level object that InnoDB uses to represent and enforce shared-access **locks** to internal in-memory data structures. Once the lock is acquired, any other process, thread, and so on can read the data structure, but no one else can write to it. Contrast with **mutexes**, which enforce exclusive access. Mutexes and rw-locks are known collectively as **latches**.

See Also [latch](#), [lock](#), [mutex](#), [Performance Schema](#).

## S

savepoint

Savepoints help to implement nested **transactions**. They can be used to provide scope to operations on tables that are part of a larger transaction. For example, scheduling a trip in a reservation system might involve booking several different flights; if a desired flight is unavailable, you might **roll back** the changes involved in booking that one leg, without rolling back the earlier flights that were successfully booked.

See Also [rollback](#), [transaction](#).

scalability

The ability to add more work and issue more simultaneous requests to a system, without a sudden drop in performance due to exceeding the limits of system capacity. Software architecture, hardware configuration, application coding, and type of workload all play a part in scalability. When the system reaches its maximum capacity, popular techniques for increasing scalability are **scale up** (increasing the capacity of existing hardware or software) and **scale out** (adding new servers and more instances of MySQL). Often paired with **availability** as critical aspects of a large-scale deployment.

See Also [availability](#), [scale out](#), [scale up](#).

scale out

A technique for increasing **scalability** by adding new servers and more instances of MySQL. For example, setting up replication, MySQL Cluster, connection pooling, or other features that spread work across a group of servers. Contrast with **scale up**.

See Also [scalability](#), [scale up](#).

scale up

A technique for increasing **scalability** by increasing the capacity of existing hardware or software.

For example, increasing the memory on a server and adjusting memory-related parameters such as [innodb\\_buffer\\_pool\\_size](#) [1708] and [innodb\\_buffer\\_pool\\_instances](#) [1706]. Contrast with **scale out**.

See Also [scalability](#), [scale out](#).

schema

Conceptually, a schema is a set of interrelated database objects, such as tables, table columns, data types of the columns, indexes, foreign keys, and so on. These objects are connected through SQL syntax, because the columns make up the tables, the foreign keys refer to tables and columns, and so on. Ideally, they are also connected logically, working together as part of a unified application or flexible framework. For example, the **information\_schema** and **performance\_schema** databases use "schema" in their names to emphasize the close relationships between the tables and columns they contain.

In MySQL, physically, a **schema** is synonymous with a **database**. You can substitute the keyword [SCHEMA](#) instead of [DATABASE](#) in MySQL SQL syntax, for example using [CREATE SCHEMA](#) instead of [CREATE DATABASE](#).

Some other database products draw a distinction. For example, in the Oracle Database product, a **schema** represents only a part of a database: the tables and other objects owned by a single user.

See Also [database](#), [ib-file set](#), [INFORMATION\\_SCHEMA](#), [Performance Schema](#).

search index

In MySQL, **full-text search** queries use a special kind of index, the **FULLTEXT** index. In MySQL 5.6.4 and up, [InnoDB](#) and [MyISAM](#) tables both support [FULLTEXT](#) indexes; formerly, these indexes were only available for [MyISAM](#) tables.

See Also [full-text search](#), [FULLTEXT index](#).

---

## secondary index

A type of InnoDB **index** that represents a subset of table columns. An InnoDB table can have zero, one, or many secondary indexes. (Contrast with the **clustered index**, which is required for each InnoDB table, and stores the data for all the table columns.)

A secondary index can be used to satisfy queries that only require values from the indexed columns. For more complex queries, it can be used to identify the relevant rows in the table, which are then retrieved through lookups using the clustered index.

Creating and dropping secondary indexes has traditionally involved significant overhead from copying all the data in the InnoDB table. The **fast index creation** feature of the InnoDB Plugin makes both `CREATE INDEX` and `DROP INDEX` statements much faster for InnoDB secondary indexes.

See Also [clustered index](#), [Fast Index Creation](#), [index](#).

## segment

A division within an InnoDB **tablespace**. If a tablespace is analogous to a directory, the segments are analogous to files within that directory. A segment can grow. New segments can be created.

For example, within a **file-per-table** tablespace, the table data is in one segment and each associated index is in its own segment. The **system tablespace** contains many different segments, because it can hold many tables and their associated indexes. The system tablespace also includes up to 128 **rollback segments** making up the **undo log**.

Segments grow and shrink as data is inserted and deleted. When a segment needs more room, it is extended by one **extent** (1 megabyte) at a time. Similarly, a segment releases one extent's worth of space when all the data in that extent is no longer needed.

See Also [extent](#), [file-per-table](#), [rollback segment](#), [system tablespace](#), [tablespace](#), [undo log](#).

## selectivity

A property of data distribution, the number of distinct values in a column (its **cardinality**) divided by the number of records in the table. High selectivity means that the column values are relatively unique, and can be retrieved efficiently through an index. If you (or the query optimizer) can predict that a test in a `WHERE` clause only matches a small number (or proportion) of rows in a table, the overall **query** tends to be efficient if it evaluates that test first, using an index.

See Also [cardinality](#), [query](#).

## semi-consistent read

A type of read operation used for `UPDATE` statements, that is a combination of **read committed** and **consistent read**. When an `UPDATE` statement examines a row that is already locked, InnoDB returns the latest committed version to MySQL so that MySQL can determine whether the row matches the `WHERE` condition of the `UPDATE`. If the row matches (must be updated), MySQL reads the row again, and this time InnoDB either locks it or waits for a lock on it. This type of read operation can only happen when the transaction has the read committed **isolation level**, or when the `innodb_locks_unsafe_for_binlog [1729]` option is enabled.

See Also [consistent read](#), [isolation level](#), [READ COMMITTED](#).

## SERIALIZABLE

The **isolation level** that uses the most conservative locking strategy, to prevent any other transactions from inserting or changing data that was read by this transaction, until it is finished. This way, the same query can be run over and over within a transaction, and be certain to retrieve the same set of results each time. Any attempt to change data that was committed by another transaction since the start of the current transaction, cause the current transaction to wait.

This is the default isolation level specified by the SQL standard. In practice, this degree of strictness is rarely needed, so the default isolation level for InnoDB is the next most strict, **repeatable read**.

See Also [ACID](#), [consistent read](#), [isolation level](#), [locking](#), [REPEATABLE READ](#), [transaction](#).

## server

A type of program that runs continuously, waiting to receive and act upon requests from another program (the client). Because often an entire computer is dedicated to running one or more server programs (such as

---

a database server, a web server, an application server, or some combination of these), the term **server** can also refer to the computer that runs the server software.

See Also [client](#), [mysqld](#).

#### shared lock

A kind of **lock** that allows other **transactions** to read the locked object, and to also acquire other shared locks on it, but not to write to it. The opposite of **exclusive lock**.

See Also [exclusive lock](#), [lock](#), [transaction](#).

#### shared tablespace

Another way of referring to the **system tablespace**.

See Also [system tablespace](#).

#### sharp checkpoint

The process of **flushing** to disk all **dirty** buffer pool pages whose redo entries are contained in certain portion of the **redo log**. Occurs before InnoDB reuses a portion of a log file; the log files are used in a circular fashion. Typically occurs with write-intensive **workloads**.

See Also [dirty page](#), [flush](#), [redo log](#), [workload](#).

#### shutdown

The process of stopping the MySQL server. By default, this process does cleanup operations for **InnoDB** tables, so it can **slow** to shut down, but fast to start up later. If you skip the cleanup operations, it is **fast** to shut down but must do the cleanup during the next restart.

The shutdown mode is controlled by the [innodb\\_fast\\_shutdown \[1716\]](#) option.

See Also [fast shutdown](#), [InnoDB](#), [slow shutdown](#), [startup](#).

#### slave server

Frequently shortened to "slave". A database **server** machine in a **replication** scenario that receives changes from another server (the **master**) and applies those same changes. Thus it maintains the same contents as the master, although it might lag somewhat behind.

In MySQL, slave servers are commonly used in disaster recovery, to take the place of a master servers that fails. They are also commonly used for testing software upgrades and new settings, to ensure that database configuration changes do not cause problems with performance or reliability.

Slave servers typically have high workloads, because they process all the **DML** (write) operations relayed from the master, as well as user queries. To ensure that slave servers can apply changes from the master fast enough, they frequently have fast I/O devices and sufficient CPU and memory to run multiple database instances on the same slave server. For example, the master server might use hard drive storage while the slave servers use **SSDs**.

See Also [DML](#), [replication](#), [server](#), [SSD](#).

#### slow query log

A type of **log** used for performance tuning of SQL statements processed by the MySQL server. The log information is stored in a file. You must enable this feature to use it. You control which categories of "slow" SQL statements are logged. For more information, see [Section 5.2.5, "The Slow Query Log"](#).

See Also [general query log](#), [log](#).

#### slow shutdown

A type of shutdown that does additional [InnoDB](#) flushing operations before completing. Also known as a **clean shutdown**. Specified by the configuration parameter [innodb\\_fast\\_shutdown=0 \[1716\]](#) or the command `SET GLOBAL innodb_fast_shutdown=0;`. Although the shutdown itself can take longer, that time will be saved on the subsequent startup.

See Also [clean shutdown](#), [fast shutdown](#), [shutdown](#).

#### snapshot

A representation of data at a particular time, which remains the same even as changes are **committed** by other **transactions**. Used by certain **isolation levels** to allow **consistent reads**.

See Also [commit](#), [consistent read](#), [isolation level](#), [transaction](#).

---

## space ID

An identifier used to uniquely identify an **InnoDB tablespace** within a MySQL instance. The space ID for the **system tablespace** is always zero; this same ID applies to all tables within the system tablespace. Each tablespace file created in **file-per-table** mode also has its own space ID.

Prior to MySQL 5.6, this hardcoded value presented difficulties in moving **InnoDB** tablespace files between MySQL instances. Starting in MySQL 5.6, you can copy tablespace files between instances by using the **transportable tablespace** feature involving the statements `FLUSH TABLES ... FOR EXPORT`, `ALTER TABLE ... DISCARD TABLESPACE`, and `ALTER TABLE ... IMPORT TABLESPACE`. The information needed to adjust the space ID is conveyed in the **.cfg file** which you copy along with the tablespace. See [Section 14.2.4.2.34, "Improved Tablespace Management"](#) for details.

See Also [.cfg file](#), [file-per-table](#), [.ibd file](#), [system tablespace](#), [tablespace](#), [transportable tablespace](#).

## spin

A type of **wait** operation that continuously tests whether a resource becomes available. This technique is used for resources that are typically held only for brief periods, where it is more efficient to wait in a "busy loop" than to put the thread to sleep and perform a context switch. If the resource does not become available within a short time, the spin loop ceases and another wait technique is used.

See Also [latch](#), [lock](#), [mutex](#), [wait](#).

## SQL

The Structured Query Language that is standard for performing database operations. Often divided into the categories **DDL**, **DML**, and **queries**. MySQL includes some additional statement categories such as **replication**. See [Chapter 9, Language Structure](#) for the building blocks of SQL syntax, [Chapter 11, Data Types](#) for the data types to use for MySQL table columns, [Chapter 13, SQL Statement Syntax](#) for details about SQL statements and their associated categories, and [Chapter 12, Functions and Operators](#) for standard and MySQL-specific functions to use in queries.

See Also [DDL](#), [DML](#), [query](#), [replication](#).

## SSD

Acronym for "solid-state drive". A type of storage device with different performance characteristics than a traditional hard disk drive (**HDD**): smaller storage capacity, faster for random reads, no moving parts, and with a number of considerations affecting write performance. Its performance characteristics can influence the throughput of a **disk-bound** workload.

See Also [disk-bound](#), [SSD](#).

## startup

The process of starting the MySQL server. Typically done by one of the programs listed in [Section 4.3, "MySQL Server and Server-Startup Programs"](#). The opposite of **shutdown**.

See Also [shutdown](#).

## statement-based replication

A form of **replication** where SQL statements are sent from the **master** server and replayed on the **slave** server. It requires some care with the setting for the `innodb_autoinc_lock_mode [1705]` option, to avoid potential timing problems with **auto-increment locking**.

See Also [auto-increment locking](#), [innodb\\_autoinc\\_lock\\_mode](#), [replication](#), [row-based replication](#).

## statistics

Estimated values relating to each **InnoDB table** and **index**, used to construct an efficient **query execution plan**. The main values are the **cardinality** (number of distinct values) and the total number of table rows or index entries. The statistics for the table represent the data in its **primary key** index. The statistics for a **secondary index** represent the rows covered by that index.

The values are estimated rather than counted precisely because at any moment, different **transactions** can be inserting and deleting rows from the same table. To keep the values from being recalculated frequently, you can enable **persistent statistics**, where the values are stored in **InnoDB** system tables, and refreshed only when you issue an `ANALYZE TABLE` statement.

You can control how **NULL** values are treated when calculating statistics through the `innodb_stats_method [1746]` configuration option.

---

Other types of statistics are available for database objects and database activity through the **INFORMATION\_SCHEMA** and **PERFORMANCE\_SCHEMA** tables.

See Also [cardinality](#), [index](#), [INFORMATION\\_SCHEMA](#), [NULL](#), [Performance Schema](#), [persistent statistics](#), [primary key](#), [query execution plan](#), [secondary index](#), [table](#), [transaction](#).

#### stemming

The ability to search for different variations of a word based on a common root word, such as singular and plural, or past, present, and future verb tense. This feature is currently supported in MyISAM **full-text search** feature but not in **FULLTEXT indexes** for InnoDB tables.

See Also [full-text search](#), [FULLTEXT index](#).

#### stopword

In a **FULLTEXT index**, a word that is considered common or trivial enough that it is omitted from the **search index** and ignored in search queries. Different configuration settings control stopwords processing for [InnoDB](#) and [MyISAM](#) tables. See [Section 12.9.4, “Full-Text Stopwords”](#) for details.

See Also [FULLTEXT index](#), [search index](#).

#### storage engine

A component of the MySQL database that performs the low-level work of storing, updating, and querying data. In MySQL 5.5 and higher, **InnoDB** is the default storage engine for new tables, superseding MyISAM. Different storage engines are designed with different tradeoffs between factors such as memory usage versus disk usage, read speed versus write speed, and speed versus robustness. Each storage engine manages specific tables, so we refer to [InnoDB tables](#), [MyISAM tables](#), and so on.

The **MySQL Enterprise Backup** product is optimized for backing up InnoDB tables. It can also back up tables handled by MyISAM and other storage engines.

See Also [InnoDB](#), [MySQL Enterprise Backup](#), [table type](#).

#### strict mode

The general name for the setting controlled by the [innodb\\_strict\\_mode \[1749\]](#) option. Turning on this setting causes certain conditions that are normally treated as warnings, to be considered errors. For example, certain invalid combinations of options related to **file format** and **row format**, that normally produce a warning and continue with default values, now cause the `CREATE TABLE` operation to fail.

MySQL also has something called strict mode.

See Also [file format](#), [innodb\\_strict\\_mode](#), [row format](#).

#### sublist

Within the list structure that represents the buffer pool, pages that are relatively old and relatively new are represented by different portions of the list. A set of parameters control the size of these portions and the dividing point between the new and old pages.

See Also [buffer pool](#), [eviction](#), [list](#), [LRU](#).

#### supremum record

A **pseudo-record** in an index, representing the **gap** above the largest value in that index. If a transaction has a statement such as `SELECT ... FOR UPDATE ... WHERE col > 10;`, and the largest value in the column is 20, it is a lock on the supremum record that prevents other transactions from inserting even larger values such as 50, 100, and so on.

See Also [gap](#), [infimum record](#), [pseudo-record](#).

#### surrogate key

Synonym name for **synthetic key**.

See Also [synthetic key](#).

#### synthetic key

A indexed column, typically a **primary key**, where the values are assigned arbitrarily. Often done using an **auto-increment** column. By treating the value as completely arbitrary, you can avoid overly restrictive rules and faulty application assumptions. For example, a numeric sequence representing employee numbers might have a gap if an employee was approved for hiring but never actually joined. Or employee number 100 might have a later hiring date than employee number 500, if they left the company and later rejoined. Numeric

---

values also produce shorter values of predictable length. For example, storing numeric codes meaning "Road", "Boulevard", "Expressway", and so on is more space-efficient than repeating those strings over and over.

Also known as a **surrogate key**. Contrast with **natural key**.

See Also [auto-increment](#), [natural key](#), [primary key](#), [surrogate key](#).

## system tablespace

A small set of data files (the **ibdata** files) containing the metadata for InnoDB-related objects (the **data dictionary**), and the storage areas for the **undo log**, the **change buffer**, and the **doublewrite buffer**. Depending on the setting of the [innodb\\_file\\_per\\_table](#) [1718], when tables are created, it might also contain table and index data for some or all InnoDB tables. The data and metadata in the system tablespace apply to all the **databases** in a MySQL **instance**.

Prior to MySQL 5.6.7, the default was to keep all InnoDB tables and indexes inside the system tablespace, often causing this file to become very large. Because the system tablespace never shrinks, storage problems could arise if large amounts of temporary data were loaded and then deleted. In MySQL 5.6.7 and higher, the default is **file-per-table** mode, where each table and its associated indexes are stored in a separate **.ibd file**. This new default makes it easier to use InnoDB features that rely on the **Barracuda** file format, such as table **compression** and the **DYNAMIC** row format.

In MySQL 5.6 and higher, setting a value for the [innodb\\_undo\\_tablespaces](#) [1753] option splits the **undo log** into one or more separate tablespace files. These files are still considered part of the system tablespace.

Keeping all table data in the system tablespace or in separate **.ibd** files has implications for storage management in general. The **MySQL Enterprise Backup** product might back up a small set of large files, or many smaller files. On systems with thousands of tables, the filesystem operations to process thousands of **.ibd** files can cause bottlenecks.

See Also [Barracuda](#), [change buffer](#), [compression](#), [data dictionary](#), [database](#), [doublewrite buffer](#), [dynamic row format](#), [file-per-table](#), [.ibd file](#), [ibdata file](#), [innodb\\_file\\_per\\_table](#), [instance](#), [MySQL Enterprise Backup](#), [tablespace](#), [undo log](#).

## T

### .TRG file

A file containing **trigger** parameters. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [MySQL Enterprise Backup](#), [mysqlbackup command](#), [.TRN file](#).

### .TRN file

A file containing trigger namespace information. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

See Also [MySQL Enterprise Backup](#), [mysqlbackup command](#), [.TRG file](#).

## table

Each MySQL table is associated with a particular **storage engine**. **InnoDB** tables have particular **physical** and **logical** characteristics that affect performance, **scalability**, **backup**, administration, and application development.

In terms of file storage, each InnoDB table is either part of the single big InnoDB **system tablespace**, or in a separate **.ibd** file if the table is created in **file-per-table** mode. The **.ibd** file holds data for the table and all its **indexes**, and is known as a **tablespace**.

InnoDB tables created in file-per-table mode can use the **Barracuda** file format. Barracuda tables can use the **DYNAMIC row format** or the **COMPRESSED row format**. These relatively new settings enable a number of InnoDB features, such as **compression**, **fast index creation**, and **off-page columns**

For backward compatibility with MySQL 5.1 and earlier, InnoDB tables inside the system tablespace must use the **Antelope** file format, which supports the **compact row format** and the **redundant row format**.



---

The **rows** of an InnoDB table are organized into an index structure known as the **clustered index**, with entries sorted based on the **primary key** columns of the table. Data access is optimized for queries that filter and sort on the primary key columns, and each index contains a copy of the associated primary key columns for each entry. Modifying values for any of the primary key columns is an expensive operation. Thus an important aspect of InnoDB table design is choosing a primary key with columns that are used in the most important queries, and keeping the primary key short, with rarely changing values.

See Also [Antelope](#), [backup](#), [Barracuda](#), [clustered index](#), [compact row format](#), [compressed row format](#), [compression](#), [dynamic row format](#), [Fast Index Creation](#), [file-per-table](#), [.ibd file](#), [index](#), [off-page column](#), [primary key](#), [redundant row format](#), [row](#), [system tablespace](#), [tablespace](#).

#### table lock

A lock that prevents any other **transaction** from accessing a table. InnoDB makes considerable effort to make such locks unnecessary, by using techniques such as **online DDL**, **row locks** and **consistent reads** for processing **DML** statements and **queries**. You can create such a lock through SQL using the `LOCK TABLE` statement; one of the steps in migrating from other database systems or MySQL storage engines is to remove such statements wherever practical.

See Also [consistent read](#), [DML](#), [lock](#), [locking](#), [online DDL](#), [query](#), [row lock](#), [table](#), [transaction](#).

#### table scan

See [full table scan](#).

#### table statistics

See [statistics](#).

#### table type

Obsolete synonym for **storage engine**. We refer to [InnoDB](#) tables, [MyISAM](#) tables, and so on.

See Also [InnoDB](#), [storage engine](#).

#### tablespace

A data file that can hold data for one or more InnoDB **tables** and associated **indexes**. The **system tablespace** contains the tables that make up the **data dictionary**, and prior to MySQL 5.6 holds all the other InnoDB tables by default. Turning on the [innodb\\_file\\_per\\_table \[1718\]](#) option, the default in MySQL 5.6 and higher, allows newly created tables to each have their own tablespace, with a separate **data file** for each table.

Using multiple tablespaces, by turning on the [innodb\\_file\\_per\\_table \[1718\]](#) option, is vital to using many MySQL features such as table compression and transportable tablespaces, and managing disk usage. See [Section 5.4.1, “Managing InnoDB Tablespaces”](#) and [Section 14.2.4.2.34, “Improved Tablespace Management”](#) for details.

Tablespaces created by the built-in InnoDB storage engine are upward compatible with the InnoDB Plugin. Tablespaces created by the InnoDB Plugin are downward compatible with the built-in InnoDB storage engine, if they use the **Antelope** file format.

MySQL Cluster also groups its tables into tablespaces. See [MySQL Cluster Disk Data Objects](#) for details.

See Also [Antelope](#), [Barracuda](#), [compressed row format](#), [data dictionary](#), [data files](#), [file-per-table](#), [index](#), [innodb\\_file\\_per\\_table](#), [system tablespace](#), [table](#).

#### tablespace dictionary

A representation of the **data dictionary** metadata for a table, within the InnoDB **tablespace**. This metadata can be checked against the **.frm file** for consistency when the table is opened, to diagnose errors resulting from out-of-date **.frm** files. This information is present for InnoDB tables that are part of the **system tablespace**, as well as for tables that have their own **.ibd file** because of the **file-per-table** option.

See Also [data dictionary](#), [file-per-table](#), [.frm file](#), [.ibd file](#), [system tablespace](#), [tablespace](#).

#### temporary table

A table whose data does not need to be truly permanent. For example, temporary tables might be used as storage areas for intermediate results in complicated calculations or transformations; this intermediate data would not need to be recovered after a crash. Database products can take various shortcuts to improve the

---

performance of operations on temporary tables, by being less scrupulous about writing data to disk and other measures to protect the data across restarts.

Sometimes, the data itself is removed automatically at a set time, such as when the transaction ends or when the session ends. With some database products, the table itself is removed automatically too.  
See Also [table](#).

#### text collection

The set of columns included in a **FULLTEXT index**.  
See Also [FULLTEXT index](#).

#### thread

A unit of processing that is typically more lightweight than a **process**, allowing for greater **concurrency**.  
See Also [concurrency](#), [master thread](#), [process](#), [Pthreads](#).

#### torn page

An error condition that can occur due to a combination of I/O device configuration and hardware failure. If data is written out in chunks smaller than the InnoDB **page size** (by default, 16KB), a hardware failure while writing could result in only part of a page being stored to disk. The InnoDB **doublewrite buffer** guards against this possibility.  
See Also [doublewrite buffer](#).

#### TPS

Acronym for "**transactions per second**", a unit of measurement sometimes used in benchmarks. Its value depends on the **workload** represented by a particular benchmark test, combined with factors that you control such as the hardware capacity and database configuration.  
See Also [transaction](#), [workload](#).

#### transaction

Transactions are atomic units of work that can be committed or rolled back. When a transaction makes multiple changes to the database, either all the changes succeed when the transaction is committed, or all the changes are undone when the transaction is rolled back.

Database transactions, as implemented by InnoDB, have properties that are collectively known by the acronym **ACID**, for atomicity, consistency, isolation, and durability.  
See Also [ACID](#), [commit](#), [isolation level](#), [lock](#), [rollback](#).

#### transaction ID

An internal field associated with each row. This field is physically changed by INSERT, UPDATE, and DELETE operations to record which transaction has locked the row.  
See Also [implicit row lock](#).

#### transportable tablespace

A feature that allows a **tablespace** to be moved from one instance to another. Traditionally, this has not been possible for InnoDB tablespaces because all table data was part of the **system tablespace**. In MySQL 5.6 and higher, the `FLUSH TABLES ... FOR EXPORT` syntax prepares an InnoDB table for copying to another server; running `ALTER TABLE ... DISCARD TABLESPACE` and `ALTER TABLE ... IMPORT TABLESPACE` on the other server brings the copied data file into the other instance. A separate `.cfg` file, copied along with the `.ibd` file, is used to update the table metadata (for example the **space ID**) as the tablespace is imported. See [Section 14.2.4.2.34, "Improved Tablespace Management"](#) for usage information.  
See Also [.ibd file](#), [space ID](#), [system tablespace](#), [tablespace](#).

#### troubleshooting

Resources for troubleshooting InnoDB reliability and performance issues include: the Information Schema tables.

#### truncate

A **DDL** operation that removes the entire contents of a table, while leaving the table and related indexes intact. Contrast with **drop**. Although conceptually it has the same result as a `DELETE` statement with no



---

**WHERE** clause, it operates differently behind the scenes: InnoDB creates a new empty table, drops the old table, then renames the new table to take the place of the old one. Because this is a DDL operation, it cannot be **rolled back**.

If the table being truncated contains foreign keys that reference another table, the truncation operation uses a slower method of operation, deleting one row at a time so that corresponding rows in the referenced table can be deleted as needed by any **ON DELETE CASCADE** clause. (MySQL 5.5 and higher do not allow this slower form of truncate, and return an error instead if foreign keys are involved. In this case, use a **DELETE** statement instead.

See Also [DDL](#), [drop](#), [foreign key](#), [rollback](#).

#### tuple

A technical term designating an ordered set of elements. It is an abstract notion, used in formal discussions of database theory. In the database field, tuples are usually represented by the columns of a table row. They could also be represented by the result sets of queries, for example, queries that retrieved only some columns of a table, or columns from joined tables.

See Also [cursor](#).

#### two-phase commit

An operation that is part of a distributed **transaction**, under the **XA** specification. (Sometimes abbreviated as 2PC.) When multiple databases participate in the transaction, either all databases **commit** the changes, or all databases **roll back** the changes.

See Also [commit](#), [rollback](#), [transaction](#), [XA](#).

## U

#### undo

Data that is maintained throughout the life of a **transaction**, recording all changes so that they can be undone in case of a **rollback** operation. It is stored in the **undo log**, also known as the **rollback segment**, either within the **system tablespace** or in separate **undo tablespaces**.

See Also [rollback](#), [rollback segment](#), [system tablespace](#), [transaction](#), [undo log](#), [undo tablespace](#).

#### undo buffer

See [undo log](#).

#### undo log

A storage area that holds copies of data modified by active **transactions**. If another transaction needs to see the original data (as part of a **consistent read** operation), the unmodified data is retrieved from this storage area.

By default, this area is physically part of the **system tablespace**. In MySQL 5.6 and higher, you can use the [innodb\\_undo\\_tablespaces \[1753\]](#) and [innodb\\_undo\\_directory \[1752\]](#) configuration options to split it into one or more separate **tablespace** files, the **undo tablespaces**, optionally stored on another storage device such as an **SSD**.

The undo log is split into separate portions, the **insert undo buffer** and the **update undo buffer**. Collectively, these parts are also known as the **rollback segment**, a familiar term for Oracle DBAs.

See Also [consistent read](#), [rollback segment](#), [SSD](#), [system tablespace](#), [transaction](#), [undo tablespace](#).

#### undo tablespace

One of a set of files containing the **undo log**, when the undo log is separated from the **system tablespace** by the [innodb\\_undo\\_tablespaces \[1753\]](#) and [innodb\\_undo\\_directory \[1752\]](#) configuration options. Only applies to MySQL 5.6 and higher.

See Also [system tablespace](#), [undo log](#).

#### unique constraint

A kind of **constraint** that asserts that a column cannot contain any duplicate values. In terms of **relational algebra**, it is used to specify 1-to-1 relationships. For efficiency in checking whether a value can be inserted

---

(that is, the value does not already exist in the column), a unique constraint is supported by an underlying **unique index**.

See Also [constraint](#), [relational](#), [unique index](#).

#### unique index

An index on a column or set of columns that have a **unique constraint**. Because the index is known not to contain any duplicate values, certain kinds of lookups and count operations are more efficient than in the normal kind of index. Most of the lookups against this type of index are simply to determine if a certain value exists or not. The number of values in the index is the same as the number of rows in the table, or at least the number of rows with non-null values for the associated columns.

The **insert buffering** optimization does not apply to unique indexes. As a workaround, you can temporarily set `unique_checks=0` while doing a bulk data load into an InnoDB table.

See Also [cardinality](#), [insert buffering](#), [unique constraint](#), [unique key](#).

#### unique key

The set of columns (one or more) comprising a **unique index**. When you can define a `WHERE` condition that matches exactly one row, and the query can use an associated unique index, the lookup and error handling can be performed very efficiently.

See Also [cardinality](#), [unique constraint](#), [unique index](#).

## V

#### victim

The transaction that is automatically chosen to be rolled back when a **deadlock** is detected. InnoDB rolls back the transaction that has updated the fewest rows.

See Also [deadlock](#), [deadlock detection](#), [innodb\\_lock\\_wait\\_timeout](#).

## W

#### wait

When an operation, such as acquiring a **lock**, **mutex**, or **latch**, cannot be completed immediately, InnoDB pauses and tries again. The mechanism for pausing is elaborate enough that this operation has its own name, the **wait**. Individual threads are paused using a combination of internal InnoDB scheduling, operating system `wait()` calls, and short-duration **spin** loops.

On systems with heavy load and many transactions, you might use the output from the `SHOW INNODB STATUS` command to determine whether threads are spending too much time waiting, and if so, how you can improve **concurrency**.

See Also [concurrency](#), [latch](#), [lock](#), [mutex](#), [spin](#).

#### warm backup

A **backup** taken while the database is running, but that restricts some database operations during the backup process. For example, tables might become read-only. For busy applications and web sites, you might prefer a **hot backup**.

See Also [backup](#), [cold backup](#), [hot backup](#).

#### warm up

To run a system under a typical **workload** for some time after startup, so that the **buffer pool** and other memory regions are filled as they would be under normal conditions.

This process happens naturally over time when a MySQL server is restarted or subjected to a new workload. Starting in MySQL 5.6, you can speed up the warmup process by setting the configuration variables `innodb_buffer_pool_dump_at_shutdown=ON` [1705] and `innodb_buffer_pool_load_at_startup=ON` [1707], to bring the contents of the buffer pool back into memory after a restart. Typically, you run a workload for some time to warm up the buffer pool before running performance tests, to ensure consistent results across multiple runs; otherwise, performance might be artificially low during the first run.

---

See Also [buffer pool](#), [workload](#).

## Windows

The built-in **InnoDB** storage engine and the **InnoDB Plugin** are supported on all the same Microsoft Windows versions as the MySQL server. The **MySQL Enterprise Backup** product has more comprehensive support for Windows systems than the **InnoDB Hot Backup** product that it supercedes.

See Also [InnoDB](#), [InnoDB Hot Backup](#), [MySQL Enterprise Backup](#), [plugin](#).

## workload

The combination and volume of **SQL** and other database operations, performed by a database application during typical or peak usage. You can subject the database to a particular workload during performance testing to identify **bottlenecks**, or during capacity planning.

See Also [bottleneck](#), [disk-bound](#), [disk-bound](#), [SQL](#).

## write combining

An optimization technique that reduces write operations when **dirty pages** are **flushed** from the **InnoDB buffer pool**. If a row in a page is updated multiple times, or multiple rows on the same page are updated, all of those changes are stored to the data files in a single write operation rather than one write for each change.

See Also [buffer pool](#), [dirty page](#), [flush](#).

## X

### XA

A standard interface for coordinating distributed **transactions**, allowing multiple databases to participate in a transaction while maintaining **ACID** compliance. For full details, see [Section 13.3.7, “XA Transactions”](#).

XA Distributed Transaction support is turned on by default. If you are not using this feature, you can disable the `innodb_support_xa` [1749] configuration option, avoiding the performance overhead of an extra fsync for each transaction.

See Also [commit](#), [transaction](#), [two-phase commit](#).

## Y

### young

A characteristic of a **page** in the **InnoDB buffer pool** meaning it has been accessed recently, and so is moved within the buffer pool data structure, so that it will not be **flushed** soon by the **LRU** algorithm. This term is used in some **information schema** column names of tables related to the buffer pool.

See Also [buffer pool](#), [flush](#), [INFORMATION\\_SCHEMA](#), [LRU](#), [page](#).



---

# General Index

## Symbols

! (logical NOT), 1171  
!= (not equal), 1167  
", 1017  
#mysql50 identifier prefix, 1018, 1022  
%, 1201  
% (modulo), 1205  
% (wildcard character), 1012  
& (bitwise AND), 1258  
&& (logical AND), 1171  
( ) (parentheses), 1165  
(Control+Z) \Z, 1012, 1417  
\* (multiplication), 1200  
+ (addition), 1200  
- (subtraction), 1200  
- (unary minus), 1200  
--master-info-repository option, 2028  
--password option, 733  
--relay-log-info-repository option, 2028  
-p option, 733  
.ARM file, 3775  
.ARZ file, 3775  
.cfg file, 3780  
.frm file, 3790  
.ibd file, 3795  
.ibz file, 3795  
.isl file, 3796  
.MRG file, 3803  
.my.cnf file, 212, 215, 216, 711, 734, 769  
.MYD file, 3803  
.MYI file, 3803  
.mylogin.cnf file, 215, 337  
.mysql\_history file, 260, 734  
.NET  
    and ODBC, 2382  
.OPT file, 3808  
.PAR file, 3809  
.pid (process ID) file, 853  
.TRG file, 3822  
.TRN file, 3822  
/ (division), 1200  
/etc/passwd, 747, 1434  
3306 port, 135, 437  
:= (assignment operator), 1172  
:= (assignment), 1029  
< (less than), 1167  
<< (left shift), 198, 1258  
<= (less than or equal), 1167  
<=> (equal to), 1167  
<> (not equal), 1167  
= (assignment operator), 1173  
= (assignment), 1029  
= (equal), 1166  
> (greater than), 1167  
>= (greater than or equal), 1167

>> (right shift), 1258  
\" (double quote), 1012  
' (single quote), 1012  
\. (mysql client command), 193, 262  
\0 (ASCII NUL), 1012, 1417  
\b (backspace), 1012, 1417  
\n (linefeed), 1012, 1417  
\n (newline), 1012, 1417  
\N (NULL), 1417  
\r (carriage return), 1012, 1417  
\t (tab), 1012, 1417  
\Z (Control+Z) ASCII 26, 1012, 1417  
\\ (escape), 1012  
^ (bitwise XOR), 1258  
\_ (wildcard character), 1012  
\_rowid, 1363  
`, 1017  
| (bitwise OR), 1258  
|| (logical OR), 1172  
~ (invert bits), 1259

## A

abort-slave-event-count option  
    mysqld, 1971  
aborted clients, 3731  
aborted connection, 3731  
ABS(), 1202  
access control, 763  
access denied errors, 3723  
access privileges, 751  
account names, 762  
accounts  
    adding privileges, 774  
    anonymous user, 150  
    default, 150  
    root, 150  
accounts table  
    performance\_schema, 2284  
ACID, 30, 1605, 1759, 3775  
ACLs, 751  
ACOS(), 1202  
action option  
    MySQLInstallerConsole, 84  
activating plugins, 609  
Active Server Pages (ASP), 2396  
ActiveState Perl, 168  
adaptive flushing, 3775  
adaptive hash index, 1637, 1653, 1689, 3775  
add-drop-database option  
    mysqldump, 286  
add-drop-table option  
    mysqldump, 286  
add-drop-trigger option  
    mysqldump, 286  
add-locks option  
    mysqldump, 294  
ADDDATE(), 1211  
adding

---

- character sets, 1086
- native functions, 3542
- new account privileges, 774
- new functions, 3532
- new user privileges, 774
- new users, 144
- user-defined functions, 3532
- addition (+), 1200
- ADDTIME(), 1212
- addtodest option
  - mysqlhotcopy, 367
- administration
  - server, 266
- administrative programs, 206
- ADO.NET, 2436
- AES\_DECRYPT(), 1260
- AES\_ENCRYPT(), 1260
- After create
  - thread state, 959
- age
  - calculating, 183
- AHI, 3776
- AIO, 3776
- alias names
  - case sensitivity, 1020
- aliases
  - for expressions, 1289
  - for tables, 1429
  - in GROUP BY clauses, 1289
  - names, 1017
  - on expressions, 1429
- ALL, 1432, 1448
- ALL join type
  - optimizer, 900
- all option
  - mysql\_config\_editor, 342
- all-databases option
  - mysqlcheck, 276
  - mysqldump, 292
- all-in-1 option
  - mysqlcheck, 276
- all-tablespaces option
  - mysqldump, 286
- allocating local table
  - thread state, 965
- allow-keywords option
  - mysqldump, 286
- allow-suspicious-udfs option
  - mysqld, 413
- allowold option
  - mysqlhotcopy, 367
- ALLOW\_INVALID\_DATES SQL mode, 603
- ALTER COLUMN, 1334
- ALTER DATABASE, 1326
- ALTER EVENT, 1327
  - and replication, 2067
- ALTER FUNCTION, 1329
- ALTER PROCEDURE, 1329
- ALTER SCHEMA, 1326
- ALTER SERVER, 1329
- ALTER TABLE, 1330, 1335, 3750
  - and replication log tables, 2028
  - ROW\_FORMAT, 673
- ALTER USER, 1522
- ALTER VIEW, 1343
- altering
  - database, 1326
  - schema, 1326
- altering table
  - thread state, 959
- altering user accounts, 1522
- ANALYSE()
  - PROCEDURE, 880
- analyze option
  - myisamchk, 323
  - mysqlcheck, 276
- ANALYZE TABLE, 1536
  - and partitioning, 2127
- Analyzing
  - thread state, 960
- AND
  - bitwise, 1258
  - logical, 1171
- anonymous user, 150, 150, 764, 766
- ANSI mode
  - running, 27
- ansi option
  - mysqld, 413
- ANSI SQL mode, 603, 608
- ANSI\_QUOTES SQL mode, 603
- answering questions
  - etiquette, 20
- Antelope, 3776
- Antelope file format, 667, 673, 1716
- ANY, 1447
- Apache, 200
- APIs, 2323
  - list of, 44
  - Perl, 3481
- application programming interface (API), 3776
- apply, 3776
- apply-slave-statements option
  - mysqldump, 288
- approximate-value literals, 1315
- ARCHIVE storage engine, 1601, 1801
- Area(), 1306, 1307
- argument processing, 3537
- arithmetic expressions, 1200
- arithmetic functions, 1257
- AS, 1429, 1435
- AsBinary(), 1302
- ASCII(), 1177
- ASIN(), 1202
- ASP.NET, 2444, 2518
- assignment operator
  - :=, 1172

---

---

- =, 1173
- assignment operators, 1172
- AsText(), 1302
- asynchronous I/O, 1656, 3776
- ATAN(), 1202
- ATAN2(), 1203
- atomic, 3776
- atomic instruction, 3777
- attackers
  - security against, 746
- attribute demotion
  - replication, 2063
- attribute promotion
  - replication, 2063
- audit log plugin
  - , 817
- audit plugins, 3489
- audit-log option
  - mysqld, 824
- audit\_log\_buffer\_size system variable, 824
- audit\_log\_file system variable, 825
- audit\_log\_flush system variable, 825
- audit\_log\_policy system variable, 826
- audit\_log\_rotate\_on\_size system variable, 826
- audit\_log\_strategy system variable, 827
- authentication
  - for the InnoDB memcached interface, 1769
- authentication plugins, 3489
- AUTHENTICATION\_PAM\_LOG environment variable, 796
- authentication\_windows\_log\_level system variable, 463
- authentication\_windows\_use\_principal\_name system variable, 464
- auto-generate-sql option
  - mysqlslap, 309
- auto-generate-sql-add-autoincrement option
  - mysqlslap, 309
- auto-generate-sql-execute-number option
  - mysqlslap, 309
- auto-generate-sql-guid-primary option
  - mysqlslap, 309
- auto-generate-sql-load-type option
  - mysqlslap, 309
- auto-generate-sql-secondary-indexes option
  - mysqlslap, 310
- auto-generate-sql-unique-query-number option
  - mysqlslap, 310
- auto-generate-sql-unique-write-number option
  - mysqlslap, 310
- auto-generate-sql-write-number option
  - mysqlslap, 310
- AUTO-INCREMENT
  - ODBC, 2392
- auto-increment, 3777
- auto-increment locking, 3777
- auto-rehash option
  - mysql, 248
- auto-repair option
  - mysqlcheck, 276
- auto-vertical-output option
  - mysql, 248
- auto.cnf file, 1955
  - and SHOW SLAVE HOSTS, 1574
- autocommit, 3777
- autocommit session variable, 464
- automatic\_sp\_privileges system variable, 464
- AUTO\_INCREMENT, 198, 1120
  - and NULL values, 3746
  - and replication, 2058
- auto\_increment\_increment system variable, 1965
- auto\_increment\_offset system variable, 1968
- availability, 3777
- AVG(), 1282
- AVG(DISTINCT), 1282

## B

- B-tree, 3777
- B-tree indexes, 876, 1636
- background threads
  - master, 1657, 1657
  - read, 1656
  - write, 1656
- backslash
  - escape character, 1011
- backspace (\b), 1012, 1417
- backticks, 3778
- backup, 3778
- backup option
  - myisamchk, 321
  - myisampack, 332
- backups, 831, 3554
  - databases and tables, 279, 365
  - InnoDB, 1616
  - with mysqldump, 840
- back\_log system variable, 465
- Barracuda, 3778
- Barracuda file format, 657, 667, 673, 1688, 1716
- base64-output option
  - mysqlbinlog, 348
- BaseCommandInterceptor class, 2531
- basedir option
  - mysql.server, 231
  - mysqld, 414
  - mysqld\_safe, 227
  - mysql\_install\_db, 237
  - mysql\_plugin, 240
  - mysql\_upgrade, 243
- basedir system variable, 466
- BaseExceptionInterceptor class, 2531
- batch mode, 192
- batch option
  - mysql, 248
- batch SQL files, 244
- Batched Key Access
  - optimization, 993, 994
- Bazaar tree, 127

---

---

BdMPolyFromText(), 1299  
BdMPolyFromWKB(), 1299  
BdPolyFromText(), 1299  
BdPolyFromWKB(), 1299  
BEGIN, 1458, 1490  
    labels, 1491  
    XA transactions, 1472  
BENCHMARK(), 1266  
benchmarks, 955, 956  
beta, 3778  
BETWEEN ... AND, 1168  
big-tables option  
    mysqld, 414  
big5, 3640  
BIGINT data type, 1113  
big\_tables session variable, 466  
BIN(), 1177  
BINARY, 1244  
BINARY data type, 1119, 1139  
binary distributions  
    installing, 62  
binary log, 624, 3778  
    event groups, 1481  
binary-mode option  
    mysql, 248  
bind-address option  
    mysql, 248  
    mysqladmin, 270  
    mysqlbinlog, 349  
    mysqlcheck, 276  
    mysqld, 414  
    mysqldump, 284  
    mysqlimport, 299  
    mysqlshow, 304  
bind\_address system variable, 466  
BINLOG, 1588  
binlog, 3779  
Binlog Dump  
    thread command, 957  
BINLOG statement  
    mysqlbinlog output, 358  
binlog-checksum option  
    mysqld, 2006  
binlog-do-db option  
    mysqld, 2004  
binlog-format option  
    mysqld, 415  
binlog-ignore-db option  
    mysqld, 2005  
binlog-row-event-max-size option  
    mysqlbinlog, 349  
    mysqld, 2001  
binlog-rows-query-log-events option  
    mysqld, 2007  
binlog\_cache\_size system variable, 2008  
binlog\_checksum system variable, 2009  
binlog\_direct\_non\_transactional\_updates system variable, 2010  
binlog\_format system variable, 2010  
binlog\_max\_flush\_queue\_time system variable, 2011  
binlog\_order\_commits system variable, 2012  
binlog\_rows\_query\_log\_events system variable, 2014  
binlog\_row\_image system variable, 2012  
binlog\_stmt\_cache\_size system variable, 2014  
BIT data type, 1112  
bit functions, 1258  
    example, 198  
BIT\_AND(), 1282  
BIT\_COUNT, 198  
BIT\_COUNT(), 1259  
BIT\_LENGTH(), 1177  
BIT\_OR, 198  
BIT\_OR(), 1282  
BIT\_XOR(), 1283  
BLACKHOLE storage engine, 1601, 1802  
blind query expansion, 1236, 3779  
BLOB columns  
    default values, 1140  
    indexing, 873, 1363  
    inserting binary data, 1013  
    size, 1150  
    with Connector/Net, 2499  
BLOB data type, 1119, 1140  
Block Nested-Loop  
    optimization, 993, 994  
Block Nested-Loop join algorithm, 982  
block-search option  
    myisamchk, 323  
BOOL data type, 1112  
BOOLEAN data type, 1112  
boolean options, 214  
Boolean search, 1234  
bootstrap option  
    mysqld, 416  
Borland applications, 2396  
bottleneck, 3779  
bounce, 3779  
Boundary(), 1304  
brackets  
    square, 1112  
brief option  
    mysqlaccess, 344  
buddy allocator, 1666, 3779  
buffer, 3779  
buffer pool, 911, 1659, 3779  
    and compressed tables, 664  
buffer pool instance, 3780  
buffer sizes, 911  
    client, 2323  
    mysqld server, 932  
Buffer(), 1308  
bugs  
    known, 3751  
    reporting, 2, 21  
bugs database, 21  
bugs.mysql.com, 21

---



---

building  
   client programs, 2797  
 BUILD\_CONFIG option  
   CMake, 131  
 built-in, 3780  
 bulk loading  
   for InnoDB tables, 885  
   for MyISAM tables, 891  
   with Connector/Net, 2506  
 bulk\_insert\_buffer\_size system variable, 466  
 business rules, 3780

**C**

C API, 2323  
   data types, 2795  
   example programs, 2797  
   functions, 2805  
   linking problems, 2798  
 C prepared statement API  
   functions, 2864, 2865  
   type codes, 2863  
 C#, 2444, 2530  
 C++ Builder, 2397  
 C:\my.cnf file, 711  
 cache, 3780  
 CACHE INDEX, 1588  
   and partitioning, 2141  
 caches  
   clearing, 1590  
 cache\_policies table, 1783  
 calculating  
   aggregate value for a set of rows, 1281  
   cardinality, 1564  
   dates, 183  
 calendar, 1229  
 CALL, 1396  
 calling sequences for aggregate functions  
   UDF, 3536  
 calling sequences for simple functions  
   UDF, 3534  
 can't create/write to file, 3732  
 cardinality, 862, 3780  
 carriage return (\r), 1012, 1417  
 CASE, 1174, 1494  
 case sensitivity  
   in access checking, 761  
   in identifiers, 1020  
   in names, 1020  
   in searches, 3743  
   in string comparisons, 1190  
   of database names, 27  
   of replication filtering options, 2034  
   of table names, 27  
 CAST, 1245  
 cast functions, 1244  
 cast operators, 1244  
 casts, 1161, 1166, 1244  
 catalog option  
   MySQLInstallerConsole, 84  
 CC environment variable, 138, 166  
 CEIL(), 1203  
 CEILING(), 1203  
 Centroid(), 1307  
 certificate file, 2523  
 CFLAGS environment variable, 138, 166  
 cflags option  
   mysql\_config, 373  
 change buffer, 3781  
 change buffering, 3781  
   disabling, 1653  
 CHANGE MASTER TO, 1476  
 Change user  
   thread command, 957  
 ChangeLog, 3755  
 changes  
   release notes, 3755  
 changes to privileges, 768  
 changing  
   column, 1334  
   field, 1334  
   socket location, 147, 3742  
   table, 1330, 1335, 3750  
 Changing master  
   thread state, 969  
 CHAR data type, 1118, 1137  
 CHAR VARYING data type, 1119  
 CHAR(), 1177  
 CHARACTER data type, 1118  
 character set repertoire, 1067  
 character sets, 1037  
   adding, 1086  
   and replication, 2059  
   repertoire, 1060  
   restrictions, 3765  
   with Connector/J, 2666  
   with Connector/Net, 2513  
 CHARACTER VARYING data type, 1119  
 character-set-client-handshake option  
   mysqld, 416  
 character-set-filesystem option  
   mysqld, 416  
 character-set-server option  
   mysqld, 417  
 character-sets-dir option  
   myisamchk, 321  
   myisampack, 332  
   mysql, 248  
   mysqladmin, 270  
   mysqlbinlog, 349  
   mysqlcheck, 276  
   mysqld, 416  
   mysqldump, 287  
   mysqlimport, 299  
   mysqlshow, 304  
 characters  
   multi-byte, 1089

---

CHARACTER\_LENGTH(), 1178  
 CHARACTER\_SETS  
   INFORMATION\_SCHEMA table, 2180  
 character\_sets\_dir system variable, 469  
 character\_set\_client system variable, 467  
 character\_set\_connection system variable, 467  
 character\_set\_database system variable, 468  
 character\_set\_filesystem system variable, 468  
 character\_set\_results system variable, 468  
 character\_set\_server system variable, 468  
 character\_set\_system system variable, 469  
 charset command  
   mysql, 256  
 charset option  
   comp\_err, 235  
 CHARSET(), 1266  
 CHAR\_LENGTH(), 1178  
 check option  
   myisamchk, 320  
   mysqlcheck, 276  
 check options  
   myisamchk, 320  
 CHECK TABLE, 1537  
   and partitioning, 2127  
 check-only-changed option  
   myisamchk, 321  
   mysqlcheck, 276  
 check-upgrade option  
   mysqlcheck, 276  
 checking  
   tables for errors, 849  
 Checking master version  
   thread state, 967  
 checking permissions  
   thread state, 960  
 checking privileges on cached query  
   thread state, 966  
 checking query cache for query  
   thread state, 966  
 Checking table  
   thread state, 960  
 checkpoint, 3781  
 checkpoint option  
   mysqlhotcopy, 367  
 checksum, 3781  
 checksum errors, 118  
 CHECKSUM TABLE, 1539  
 child table, 3782  
 Chinese, Japanese, Korean character sets  
   frequently asked questions, 3640  
 choosing  
   a MySQL version, 50  
   data types, 1151  
 chroot option  
   mysqld, 417  
   mysqlhotcopy, 367  
 CJK (Chinese, Japanese, Korean)  
   Access, PHP, etc., 3640  
   availability of specific characters, 3640  
   big5, 3640  
   character sets available, 3640  
   characters displayed as question marks, 3640  
   CJKV, 3640  
   collations, 3640, 3640  
   conversion problems with Japanese character sets, 3640  
   data truncation, 3640  
   Database and table names, 3640  
   documentation in Chinese, 3640  
   documentation in Japanese, 3640  
   documentation in Korean, 3640  
   FAQ, 3640  
   gb2312, gbk, 3640  
   Japanese character sets, 3640  
   Korean character set, 3640  
   LIKE and FULLTEXT, 3640  
   MySQL 4.0 behavior, 3640  
   ORDER BY treatment, 3640, 3640  
   problems with Access, PHP, etc., 3640  
   problems with Big5 character sets (Chinese), 3640  
   problems with data truncation, 3640  
   problems with euckr character set (Korean), 3640  
   problems with GB character sets (Chinese), 3640  
   problems with LIKE and FULLTEXT, 3640  
   problems with Yen sign (Japanese), 3640  
   rejected characters, 3640  
   sort order problems, 3640, 3640  
   sorting problems, 3640, 3640  
   testing availability of characters, 3640  
   Unicode collations, 3640  
   Vietnamese, 3640  
   Yen sign, 3640  
 clean page, 3782  
 clean shutdown, 618, 707, 2072, 3782  
 cleaning up  
   thread state, 960  
 clear command  
   mysql, 256  
 Clearing  
   thread state, 969  
 clearing  
   caches, 1590  
 client, 3782  
 client connection threads, 946  
 client programs, 205  
   building, 2797  
 client tools, 2323  
 clients  
   debugging, 3550  
   threaded, 2799  
 CLOSE, 1498  
 Close stmt  
   thread command, 958  
 closing  
   tables, 881  
 closing tables

---

---

thread state, 960  
 clustered index, 3782  
   InnoDB, 1635  
 CMake  
   BUILD\_CONFIG option, 131  
   CMAKE\_BUILD\_TYPE option, 131  
   CMAKE\_INSTALL\_PREFIX option, 132  
   COMPILATION\_COMMENT option, 134  
   CPACK\_MONOLITHIC\_INSTALL option, 131  
   DEFAULT\_CHARSET option, 134  
   DEFAULT\_COLLATION option, 134  
   ENABLED\_LOCAL\_INFILE option, 135  
   ENABLED\_PROFILING option, 135  
   ENABLE\_DEBUG\_SYNC option, 134  
   ENABLE\_DOWNLOADS option, 135  
   ENABLE\_DTRACE option, 135  
   ENABLE\_GCOV option, 135  
   ENABLE\_GPROF option, 135  
   HAVE\_EMBEDDED\_PRIVILEGE\_CONTROL option, 137  
   IGNORE\_AIO\_CHECK option, 135  
   INSTALL\_BINDIR option, 132  
   INSTALL\_DOCDIR option, 132  
   INSTALL\_DOCREADMEDIR option, 132  
   INSTALL\_INCLUDEDIR option, 132  
   INSTALL\_INFODIR option, 132  
   INSTALL\_LAYOUT option, 132  
   INSTALL\_LIBDIR option, 132  
   INSTALL\_MANDIR option, 132  
   INSTALL\_MYSQLSHAREDIR option, 132  
   INSTALL\_MYSQLTESTDIR option, 133  
   INSTALL\_PLUGINDIR option, 133  
   INSTALL\_SBINDIR option, 133  
   INSTALL\_SCRIPTDIR option, 133  
   INSTALL\_SHAREDIR option, 133  
   INSTALL\_SQLBENCHDIR option, 133  
   INSTALL\_SUPPORTFILESDIR option, 133  
   MYSQL\_DATADIR option, 133  
   MYSQL\_MAINTAINER\_MODE option, 135  
   MYSQL\_PROJECT\_NAME option, 135  
   MYSQL\_TCP\_PORT option, 135  
   MYSQL\_UNIX\_ADDR option, 135  
   ODBC\_INCLUDES option, 133  
   ODBC\_LIB\_DIR option, 133  
   OPTIMIZER\_TRACE option, 133  
   options, 129  
   running after prior invocation, 126, 137  
   SYSCONFDIR option, 133  
   VERSION file, 138  
   WITH\_DEBUG option, 136  
   WITH\_DEFAULT\_COMPILER\_OPTIONS option, 136  
   WITH\_DEFAULT\_FEATURE\_SET option, 136  
   WITH\_EMBEDDED\_SERVER option, 136  
   WITH\_EXTRA\_CHARSETS option, 136  
   WITH\_LIBEDIT option, 136  
   WITH\_LIBWRAP option, 136  
   WITH\_READLINE option, 136  
   WITH\_SSL option, 136  
   WITH\_UNIXODBC option, 137  
   WITH\_ZLIB option, 137  
 CMakeCache.txt file, 137  
 CMAKE\_BUILD\_TYPE option  
   CMake, 131  
 CMAKE\_INSTALL\_PREFIX option  
   CMake, 132  
 COALESCE(), 1169  
 COERCIBILITY(), 1266  
 cold backup, 3782  
 ColdFusion, 2397  
 collating  
   strings, 1089  
 collation  
   adding, 1089  
   INFORMATION\_SCHEMA, 1058  
   modifying, 1090  
   names, 1051  
 COLLATION(), 1267  
 collation-server option  
   mysqld, 417  
 COLLATIONS  
   INFORMATION\_SCHEMA table, 2180  
 COLLATION\_CHARACTER\_SET\_APPLICABILITY  
   INFORMATION\_SCHEMA table, 2181  
 collation\_connection system variable, 469  
 collation\_database system variable, 469  
 collation\_server system variable, 470  
 column, 3782  
   changing, 1334  
   types, 1111  
 column alias  
   problems, 3746  
   quoting, 1018, 3746  
 column comments, 1362  
 Column Editor, 2418  
 column index, 3783  
 column names  
   case sensitivity, 1020  
 column prefix, 3783  
 column-names option  
   mysql, 249  
 column-type-info option  
   mysql, 249  
 columns  
   displaying, 302  
   indexes, 872  
   names, 1017  
   other types, 1151  
   selecting, 181  
   storage requirements, 1147  
 COLUMNS  
   INFORMATION\_SCHEMA table, 2181  
 columns option  
   mysqlimport, 299  
 columns partitioning, 2097  
 columns per table

---

---

- maximum, 3769
- COLUMN\_PRIVILEGES
  - INFORMATION\_SCHEMA table, 2182
- comma-separated values data, reading, 1416, 1434
- command options
  - mysql, 244
  - mysqladmin, 268
  - mysqld, 412
- command syntax, 4
- command-line history
  - mysql, 260
- command-line tool, 83, 244
- commands
  - for binary distribution, 63
- commands out of sync, 3733
- comment syntax, 1034
- comments
  - adding, 1034
  - starting, 33
- comments option
  - mysql, 249
  - mysqldump, 286
- COMMIT, 30, 1458
  - XA transactions, 1472
- commit, 3783
- commit option
  - mysqlaccess, 344
  - mysqslap, 310
- committing alter table to storage engine
  - thread state, 960
- compact option
  - mysqldump, 290
- compact row format, 673, 3783
- comparison operators, 1165
- compatibility
  - between MySQL versions, 155
  - with mSQL, 1193
  - with ODBC, 543, 1019, 1115, 1162, 1168, 1362, 1437
  - with Oracle, 28, 1285, 1334, 1597
  - with PostgreSQL, 29
  - with standard SQL, 26
- compatible option
  - mysqldump, 290
- COMPILATION\_COMMENT option
  - CMake, 134
- compiling
  - optimizing, 931
  - user-defined functions, 3539
- compiling clients
  - on Unix, 2797
  - on Windows, 2798
- compiling MySQL server
  - problems, 137
- complete-insert option
  - mysqldump, 290
- completion\_type system variable, 470
- composite index, 3783
- composite partitioning, 2108
- compound statements, 1490
- compress option
  - mysql, 249
  - mysqladmin, 270
  - mysqlcheck, 276
  - mysqldump, 284
  - mysqlexport, 299
  - mysqlshow, 304
  - mysqslap, 310
- COMPRESS(), 1261
- compressed backup, 3783
- compressed row format, 673, 3783
- compressed tables, 332, 1793
- compression, 656, 3784
  - application and schema design, 660
  - BLOBs, VARCHAR and TEXT, 664
  - buffer pool considerations, 664
  - compressed page size, 661
  - configuration characteristics, 661
  - data and indexes, 663
  - data characteristics, 658
  - information schema, 1665
  - innodb\_strict\_mode, 1690
  - KEY\_BLOCK\_SIZE, 661
  - log file format, 665
  - monitoring, 661
  - tuning, 658
  - with Connector/Net network traffic, 2523
  - workload characteristics, 660
- compression failure, 3784
- comp\_err, 204, 235
  - charset option, 235
  - debug option, 235
  - debug-info option, 235
  - header\_file option, 235
  - help option, 235
  - in\_file option, 235
  - name\_file option, 235
  - out\_dir option, 235
  - out\_file option, 236
  - statefile option, 236
  - version option, 236
- CONCAT(), 1178
- concatenation
  - string, 1011, 1178
- CONCAT\_WS(), 1179
- concurrency, 1605, 3784
  - of commits, 1712
  - of threads, 1751
  - tickets, 1714
- concurrency option
  - mysqslap, 310
- concurrent inserts, 926, 929
- concurrent\_insert system variable, 471
- Conditions, 1499
- conditions, 1561, 1586
- cond\_instances table

---

- performance\_schema, 2271
- config option
  - MySQLInstallerConsole, 84
- config-file option
  - my\_print\_defaults, 374
- configuration file, 3784
- configuration files, 769
- config\_options table, 1783
- Connect
  - thread command, 958
- connect command
  - mysql, 256
- Connect Out
  - thread command, 958
- connect-expired-password option
  - mysql, 249
- connecting
  - remotely with SSH, 817
  - through Connector/Net, 2484
  - through JDBC and Connector/J, 2642
  - through Visual Studio, 2414
  - to the server, 171, 209
  - verification, 763
- Connecting to master
  - thread state, 967
- connection
  - aborted, 3731
- connection pooling, 2490, 2523, 2684, 2697, 2698
- connection string, 2485, 2523
- CONNECTION\_ID(), 1267
- Connector/C, 2323, 2755
- Connector/C++, 2323, 2713
- Connector/J, 2323, 2635
  - reporting problems, 2712
  - troubleshooting, 2704
- Connector/Net, 2323, 2403
  - reporting problems, 2633
- Connector/ODBC, 2323, 2326
  - reporting problems, 2402, 2402
  - with Borland applications, 2396
- Connector/Python, 2323, 2759
- Connectors, 2323
- connect\_timeout system variable, 472
- connect\_timeout variable, 255, 272
- consistent read, 3785
- consistent reads, 1625
- console option
  - mysqld, 417
- const table
  - optimizer, 898, 1432
- constant table, 859
- constraint, 3785
- constraints, 34
  - foreign keys, 1379
- CONSTRAINTS
  - INFORMATION\_SCHEMA table, 2201
- containers table, 1783
- Contains(), 1311
- contributing companies
  - list of, 45
- contributors
  - list of, 38
- control flow functions, 1173
- CONV(), 1203
- conventions
  - syntax, 3
  - typographical, 3
- CONVERT, 1245
- CONVERT TO, 1336
- converting HEAP to MyISAM
  - thread state, 960
- CONVERT\_TZ(), 1212
- ConvexHull(), 1309
- copy option
  - mysqlaccess, 344
- copy to tmp table
  - thread state, 960
- copying databases, 164
- copying tables, 1376, 1376
- Copying to group table
  - thread state, 960
- Copying to tmp table
  - thread state, 960
- Copying to tmp table on disk
  - thread state, 960
- core-file option
  - mysqld, 418
- core-file-size option
  - mysqld\_safe, 227
- core\_file system variable, 472
- correct-checksum option
  - myisamchk, 321
- correlated subqueries, 1450
- corruption, 1684
- COS(), 1203
- COT(), 1203
- count option
  - myisam\_ftdump, 314
  - mysqladmin, 270
  - mysqlshow, 304
- COUNT(), 1283
- COUNT(DISTINCT), 1283
- counter, 3785
- counting
  - table rows, 188
- covering index, 3785
- CPACK\_MONOLITHIC\_INSTALL option
  - CMake, 131
- crash, 3544, 3785
  - recovery, 848
  - repeated, 3738
  - replication, 2072
- crash recovery, 3786
- crash-me, 956
- crash-me program, 955
- CRC32(), 1203

---

---

CREATE ... IF NOT EXISTS  
  and replication, 2059  
CREATE DATABASE, 1344  
Create DB  
  thread command, 958  
CREATE EVENT, 1344  
  and replication, 2067  
CREATE FUNCTION, 1352, 1544  
CREATE INDEX, 674, 1349  
create option  
  mysqslap, 310  
CREATE PROCEDURE, 1352  
CREATE SCHEMA, 1344  
CREATE SERVER, 1357  
CREATE TABLE, 1358  
  DIRECTORY options  
    and replication, 2067  
  KEY\_BLOCK\_SIZE, 661  
  options for table compression, 657  
  ROW\_FORMAT, 673  
CREATE TABLE ... SELECT  
  and replication, 2060  
CREATE TRIGGER, 1384  
CREATE USER, 1522  
CREATE VIEW, 1388  
create-options option  
  mysqldump, 290  
create-schema option  
  mysqslap, 310  
creating  
  bug reports, 21  
  database, 1344  
  databases, 175  
  default startup options, 215  
  function, 1544  
  schema, 1344  
  tables, 177  
Creating delayed handler  
  thread state, 965  
Creating index  
  thread state, 960  
Creating sort index  
  thread state, 960  
creating table  
  thread state, 960  
Creating tmp table  
  thread state, 960  
creating user accounts, 1522  
CROSS JOIN, 1435  
Crosses(), 1311  
CRUD, 3786  
Crystal Reports, 2514  
CR\_SERVER\_GONE\_ERROR, 3729  
CR\_SERVER\_LOST\_ERROR, 3729  
CSV data, reading, 1416, 1434  
csv option  
  mysqslap, 310  
CSV storage engine, 1601, 1799

CURDATE(), 1212  
CURRENT\_DATE, 1212  
CURRENT\_TIME, 1212  
CURRENT\_TIMESTAMP, 1213  
CURRENT\_USER(), 1267  
cursor, 3786  
Cursors, 1498  
CURTIME(), 1213  
CXX environment variable, 138, 166  
CXXFLAGS environment variable, 138, 166  
cxxflags option  
  mysql\_config, 373

## D

Daemon  
  thread command, 958  
daemon plugins, 3488  
daemon\_memcached\_enable\_binlog system variable, 1698  
daemon\_memcached\_engine\_lib\_name system variable, 1698  
daemon\_memcached\_engine\_lib\_path system variable, 1699  
daemon\_memcached\_option system variable, 1699  
daemon\_memcached\_r\_batch\_size system variable, 1699  
daemon\_memcached\_w\_batch\_size system variable, 1700  
data  
  importing, 262, 297  
  loading into tables, 178  
  retrieving, 179  
  size, 878  
data dictionary, 641, 3786  
DATA DIRECTORY  
  and replication, 2067  
data directory, 3786  
data files, 3786  
Data truncation with CJK characters, 3640  
data type  
  BIGINT, 1113  
  BINARY, 1119, 1139  
  BIT, 1112  
  BLOB, 1119, 1140  
  BOOL, 1112, 1151  
  BOOLEAN, 1112, 1151  
  CHAR, 1118, 1137  
  CHAR VARYING, 1119  
  CHARACTER, 1118  
  CHARACTER VARYING, 1119  
  DATE, 1115, 1126  
  DATETIME, 1115, 1126  
  DEC, 1114  
  DECIMAL, 1114, 1315  
  DOUBLE, 1114  
  DOUBLE PRECISION, 1115  
  ENUM, 1120, 1141  
  FIXED, 1114

---

FLOAT, 1114, 1114, 1115  
 GEOMETRY, 1297  
 GEOMETRYCOLLECTION, 1297  
 INT, 1113  
 INTEGER, 1113  
 LINESTRING, 1297  
 LONG, 1140  
 LONGBLOB, 1120  
 LONGTEXT, 1120  
 MEDIUMBLOB, 1120  
 MEDIUMINT, 1113  
 MEDIUMTEXT, 1120  
 MULTILINESTRING, 1297  
 MULTIPOINT, 1297  
 MULTIPOLYGON, 1297  
 NATIONAL CHAR, 1118  
 NATIONAL VARCHAR, 1119  
 NCHAR, 1118  
 NUMERIC, 1114  
 NVARCHAR, 1119  
 POINT, 1297  
 POLYGON, 1297  
 REAL, 1115  
 SET, 1120, 1144  
 SMALLINT, 1113  
 TEXT, 1119, 1140  
 TIME, 1116, 1127  
 TIMESTAMP, 1116, 1126  
 TINYBLOB, 1119  
 TINYINT, 1112  
 TINYTEXT, 1119  
 VARBINARY, 1119, 1139  
 VARCHAR, 1119, 1137  
 VARCHARACTER, 1119  
 YEAR, 1117, 1128  
 data types, 1111  
   C API, 2795  
   overview, 1112  
 data warehouse, 3787  
 data-file-length option  
   mysamchk, 322  
 database, 3787  
   altering, 1326  
   creating, 1344  
   deleting, 1391  
 Database information  
   obtaining, 1549  
 database metadata, 2178  
 database names  
   case sensitivity, 27, 1020  
 database option  
   mysql, 249  
   mysqlbinlog, 349  
 DATABASE(), 1268  
 databases  
   backups, 831  
   copying, 164  
   creating, 175, 1344  
   defined, 4  
   displaying, 302  
   dumping, 279, 365  
   information about, 191  
   names, 1017  
   replicating, 1929  
   selecting, 176, 1600  
   symbolic links, 941  
   using, 175  
   databases option  
     mysqlcheck, 276  
     mysqldump, 292  
   datadir option  
     mysql.server, 231  
     mysqld, 418  
     mysqld\_safe, 227  
     mysql\_install\_db, 237  
     mysql\_plugin, 240  
     mysql\_upgrade, 243  
   datadir system variable, 472  
   DataJunction, 2397  
   DATE, 3744  
   date and time functions, 1209  
   Date and Time types, 1125  
   date calculations, 183  
   DATE columns  
     problems, 3744  
     with Connector/Net, 2505  
   DATE data type, 1115, 1126  
   date literals, 1014  
   date types, 1149  
   date values  
     problems, 1126  
   DATE(), 1213  
   DATEDIFF(), 1213  
   dates  
     used with partitioning, 2089  
     used with partitioning (examples), 2092, 2104, 2109,  
     2131  
   DATETIME data type, 1115, 1126  
   datetime\_format system variable, 472  
   DATE\_ADD(), 1213  
   date\_format system variable, 472  
   DATE\_FORMAT(), 1215  
   DATE\_SUB(), 1213, 1217  
   DAY(), 1217, 1265  
   DAYNAME(), 1217  
   DAYOFMONTH(), 1217  
   DAYOFWEEK(), 1217  
   DAYOFYEAR(), 1217  
   db option  
     mysqlaccess, 344  
   db table  
     sorting, 767  
   DB2 SQL mode, 608  
   DBI interface, 3481  
   DBI->quote, 1013  
   DBI->trace, 3547

---

---

DBI/DBD interface, 3481  
 DBI\_TRACE environment variable, 166, 3547  
 DBI\_USER environment variable, 166  
 DEBUG package, 3550  
 DCL, 1524, 1534, 3787  
 DDL, 1326, 3787  
 deadlock, 925, 1465, 1623, 1630, 1633, 1741, 2075, 2319, 3787  
 deadlock detection, 1647, 3788  
 DEALLOCATE PREPARE, 1486, 1490  
 Debug  
   thread command, 958  
 debug option  
   comp\_err, 235  
   myisamchk, 319  
   myisampack, 332  
   mysql, 249  
   mysqlaccess, 344  
   mysqladmin, 271  
   mysqlbinlog, 350  
   mysqlcheck, 277  
   mysqld, 418  
   mysqldump, 286  
   mysqldumpslow, 364  
   mysqlhotcopy, 367  
   mysqlimport, 300  
   mysqlshow, 304  
   mysqlslap, 310  
   mysql\_config\_editor, 342  
   my\_print\_defaults, 374  
 debug system variable, 473  
 debug-check option  
   mysql, 249  
   mysqladmin, 271  
   mysqlbinlog, 350  
   mysqlcheck, 277  
   mysqldump, 287  
   mysqlimport, 300  
   mysqlshow, 304  
   mysqlslap, 310  
   mysql\_upgrade, 243  
 debug-info option  
   comp\_err, 235  
   mysql, 249  
   mysqladmin, 271  
   mysqlbinlog, 350  
   mysqlcheck, 277  
   mysqldump, 287  
   mysqlimport, 300  
   mysqlshow, 304  
   mysqlslap, 311  
   mysql\_upgrade, 243  
 debug-sync-timeout option  
   mysqld, 419  
 debugging  
   client, 3550  
   server, 3544  
 debugging support, 129  
 debug\_sync system variable, 473  
 DEC data type, 1114  
 decimal arithmetic, 1315  
 DECIMAL data type, 1114, 1315  
 decimal point, 1111  
 DECLARE, 1492  
 DECODE(), 1261  
 decode\_bits myisamchk variable, 319  
 DEFAULT  
   constraint, 36  
 default  
   privileges, 150  
 default accounts, 150  
 default host name, 209  
 default installation location, 62  
 default options, 215  
 DEFAULT value clause, 1146, 1362  
 default values, 1146, 1362, 1404  
   BLOB and TEXT columns, 1140  
   explicit, 1146  
   implicit, 1146  
   suppression, 36  
 DEFAULT(), 1275  
 default-auth option  
   mysql, 249  
   mysqladmin, 271  
   mysqlbinlog, 351  
   mysqlcheck, 277  
   mysqldump, 284  
   mysqlimport, 300  
   mysqlshow, 305  
   mysqlslap, 311  
   mysql\_upgrade, 243  
 default-authentication-plugin option  
   mysqld, 419  
 default-character-set option  
   mysql, 249  
   mysqladmin, 271  
   mysqlcheck, 277  
   mysqldump, 288  
   mysqlimport, 300  
   mysqlshow, 305  
 default-storage-engine option  
   mysqld, 420  
 default-time-zone option  
   mysqld, 420  
 defaults  
   embedded, 2791  
 defaults-extra-file option, 219  
   mysqld\_multi, 232  
   mysqld\_safe, 227  
   my\_print\_defaults, 374  
 defaults-file option, 219  
   mysqld\_multi, 232  
   mysqld\_safe, 227  
   my\_print\_defaults, 374  
 defaults-group-suffix option, 219  
   my\_print\_defaults, 374

---



---

DEFAULT\_CHARSET option  
     CMake, 134  
 DEFAULT\_COLLATION option  
     CMake, 134  
 default\_storage\_engine system variable, 474  
 default\_tmp\_storage\_engine system variable, 474  
 default\_week\_format system variable, 475  
 DEGREES(), 1204  
 delay-key-write option  
     mysqld, 421, 1790  
 DELAYED, 1408  
     when ignored, 1406  
 Delayed insert  
     thread command, 958  
 delayed inserts  
     thread states, 965  
 delayed replication, 2057  
 delayed-insert option  
     mysqldump, 294  
 delayed\_insert\_limit, 1410  
 delayed\_insert\_limit system variable, 476  
 delayed\_insert\_timeout system variable, 476  
 delayed\_queue\_size system variable, 477  
 delay\_key\_write system variable, 475  
 DELETE, 1398  
 delete, 3788  
 delete buffering, 3788  
 delete option  
     mysqlimport, 300  
 delete-master-logs option  
     mysqldump, 288  
 deleting  
     database, 1391  
     foreign key, 1336, 1381  
     function, 1545  
     index, 1334, 1392  
     primary key, 1335  
     rows, 3747  
     schema, 1391  
     table, 1393  
     user, 778, 1524  
     users, 778, 1524  
 deleting from main table  
     thread state, 961  
 deleting from reference tables  
     thread state, 961  
 deletion  
     mysql.sock, 3742  
 delimiter command  
     mysql, 256  
 delimiter option  
     mysql, 249  
     mysqlslap, 311  
 Delphi, 2396  
 demo\_test table, 1766  
 denormalized, 3788  
 derived tables, 1450  
     optimization, 1004  
 des-key-file option  
     mysqld, 421  
 DESC, 1596  
 descending index, 3788  
 DESCRIBE, 191, 1596  
 description option  
     myisamchk, 323  
 design  
     issues, 3751  
 DES\_DECRYPT(), 1261  
 DES\_ENCRYPT(), 1262  
 detach option  
     mysqlslap, 311  
 development source tree, 127  
 Difference(), 1309  
 digits, 1111  
 Dimension(), 1303  
 directory structure  
     default, 62  
 dirty page, 1657, 1701, 3788  
 dirty read, 3788  
 disable named command  
     mysql, 250  
 --disable option prefix, 214  
 disable-gtid-unsafe-statements option, 2018  
 disable-keys option  
     mysqldump, 294  
 disable-log-bin option  
     mysqlbinlog, 351  
 disable\_gtid\_unsafe\_statements system variable, 2020  
 DISCARD TABLESPACE, 1336, 1619  
 discard\_or\_import\_tablespace  
     thread state, 961  
 disconnect-slave-event-count option  
     mysqld, 1971  
 disconnecting  
     from the server, 171  
 disconnect\_on\_expired\_passwords session variable, 477  
 Disjoint(), 1311  
 disk full, 3740  
 disk performance, 939  
 disk-based, 3788  
 disk-bound, 3789, 3789  
 disks  
     splitting data across, 942  
 display size, 1111  
 display triggers, 1583  
 display width, 1111  
 displaying  
     database information, 302  
     information  
         Cardinality, 1564  
         Collation, 1564  
         SHOW, 1549, 1552, 1563, 1565, 1583  
         table status, 1581  
 DISTINCT, 181, 1000, 1432  
     AVG(), 1282

---

---

COUNT(), 1283  
MAX(), 1284  
MIN(), 1284  
SUM(), 1285  
DISTINCTROW, 1432  
DIV, 1201  
division (/), 1200  
div\_precision\_increment system variable, 478  
DML, 1396, 3789  
    DELETE statement, 1398  
    INSERT statement, 1403  
    UPDATE statement, 1455  
DNS, 947  
DO, 1402  
DocBook XML  
    documentation source format, 2  
document id, 3789  
Documentation  
    in Chinese, 3640  
    in Japanese, 3640  
    in Korean, 3640  
Documenters  
    list of, 42  
DOUBLE data type, 1114  
DOUBLE PRECISION data type, 1115  
double quote ("), 1012  
doublewrite buffer, 592, 638, 1715, 3789  
downgrading, 153, 159  
downloading, 53  
drop, 3789  
DROP ... IF EXISTS  
    and replication, 2061  
DROP DATABASE, 1391  
Drop DB  
    thread command, 958  
DROP EVENT, 1392  
DROP FOREIGN KEY, 1336, 1381  
DROP FUNCTION, 1393, 1545  
DROP INDEX, 674, 1334, 1392  
DROP PREPARE, 1490  
DROP PRIMARY KEY, 1335  
DROP PROCEDURE, 1393  
DROP SCHEMA, 1391  
DROP SERVER, 1393  
DROP TABLE, 1393  
DROP TRIGGER, 1394  
DROP USER, 1524  
DROP VIEW, 1394  
dropping  
    user, 778, 1524  
dryrun option  
    mysqlhotcopy, 367  
DUAL, 1428  
dump option  
    myisam\_ftdump, 315  
dump-date option  
    mysqldump, 287  
dump-slave option

mysqldump, 288  
DUMPFIL, 1435  
dumping  
    databases and tables, 279, 365  
dynamic row format, 673, 3789  
dynamic table characteristics, 1792

## E

early adopter, 3790  
edit command  
    mysql, 257  
ego command  
    mysql, 257  
Eiffel Wrapper, 3482  
ELT(), 1179  
email lists, 18  
embedded MySQL server library, 2790  
embedded option  
    mysql\_config, 373  
--enable option prefix, 214  
enable-clear-text-plugin option  
    mysql, 250  
    mysqladmin, 271  
    mysqslap, 311  
enable-named-pipe option  
    mysqld, 421  
ENABLED\_LOCAL\_INFILE option  
    CMake, 135  
ENABLED\_PROFILING option  
    CMake, 135  
ENABLE\_DEBUG\_SYNC option  
    CMake, 134  
ENABLE\_DOWNLOADS option  
    CMake, 135  
ENABLE\_DTRACE option  
    CMake, 135  
ENABLE\_GCOV option  
    CMake, 135  
ENABLE\_GPROF option  
    CMake, 135  
ENCODE(), 1262  
ENCRYPT(), 1262  
encryption, 746, 806  
encryption functions, 1259  
end  
    thread state, 961  
END, 1490  
EndPoint(), 1305  
end\_markers\_in\_json system variable, 479  
enforce-gtid-consistency option, 2018  
enforce\_gtid\_consistency system variable, 2021  
engine option  
    mysqslap, 311  
ENGINES  
    INFORMATION\_SCHEMA table, 2183  
engine\_condition\_pushdown system variable, 478  
entering  
    queries, 172

---

enterprise components  
   MySQL Enterprise Audit, 817, 3555  
   MySQL Enterprise Backup, 3554  
   MySQL Enterprise Monitor, 3553  
   MySQL Enterprise Security, 790, 797, 3554  
   MySQL Thread Pool, 949, 3555

ENUM  
   size, 1150

ENUM data type, 1120, 1141

Envelope(), 1303

environment variable  
   AUTHENTICATION\_PAM\_LOG, 796  
   CC, 138, 166  
   CFLAGS, 138, 166  
   CXX, 138, 166  
   CXXFLAGS, 138, 166  
   DBI\_TRACE, 166, 3547  
   DBI\_USER, 166  
   HOME, 166, 260  
   LD\_LIBRARY\_PATH, 169  
   LD\_RUN\_PATH, 166, 169  
   LIBMYSQL\_ENABLE\_CLEARTEXT\_PLUGIN, 166  
   LIBMYSQL\_PLUGINS, 166, 2894  
   LIBMYSQL\_PLUGIN\_DIR, 166, 2894  
   MYSQL\_DEBUG, 166, 208, 3550  
   MYSQL\_GROUP\_SUFFIX, 166  
   MYSQL\_HISTFILE, 166, 260  
   MYSQL\_HISTIGNORE, 166, 260  
   MYSQL\_HOME, 166  
   MYSQL\_HOST, 166, 212  
   MYSQL\_PS1, 166  
   MYSQL\_PWD, 166, 208, 212  
   MYSQL\_TCP\_PORT, 166, 208, 710, 711  
   MYSQL\_TEST\_LOGIN\_FILE, 166  
   MYSQL\_UNIX\_PORT, 145, 166, 208, 710, 711  
   PATH, 140, 166, 209  
   TMPDIR, 145, 166, 208, 3741  
   TZ, 166, 3742  
   UMASK, 166, 3735  
   UMASK\_DIR, 166, 3736  
   USER, 166, 212

environment variables, 208, 224, 769  
   list of, 165

equal (=), 1166

Equals(), 1311

eq\_ref join type  
   optimizer, 898

Errcode, 375

errno, 375

Error  
   thread command, 958

error codes, 2671

error log, 3790

error messages  
   can't find file, 3735  
   displaying, 375  
   languages, 1085, 1085

errors  
   access denied, 3723  
   and replication, 2074  
   checking tables for, 849  
   common, 3722  
   directory checksum, 118  
   handling for UDFs, 3539  
   in subqueries, 1452  
   known, 3751  
   linking, 2798  
   list of, 3723  
   lost connection, 3726  
   reporting, 21, 21  
   sources of information, 3659

error\_count session variable, 480

ERROR\_FOR\_DIVISION\_BY\_ZERO SQL mode, 603

ER\_ABORTING\_CONNECTION, 2671

ER\_ACCESS\_DENIED\_ERROR, 2671

ER\_BAD\_FIELD\_ERROR, 2671

ER\_BAD\_HOST\_ERROR, 2671

ER\_BAD\_TABLE\_ERROR, 2671

ER\_BLOBS\_AND\_NO\_TERMINATED, 2671

ER\_BLOB\_CANT\_HAVE\_DEFAULT, 2671

ER\_BLOB\_KEY\_WITHOUT\_LENGTH, 2671

ER\_BLOB\_USED\_AS\_KEY, 2671

ER\_CANT\_DO\_THIS\_DURING\_AN\_TRANSACTION, 2671

ER\_CANT\_DROP\_FIELD\_OR\_KEY, 2671

ER\_CANT\_REMOVE\_ALL\_FIELDS, 2671

ER\_CANT\_USE\_OPTION\_HERE, 2671

ER\_CHECK\_NOT\_IMPLEMENTED, 2671

ER\_CHECK\_NO\_SUCH\_TABLE, 2671

ER\_COLLATION\_CHARSET\_MISMATCH, 2671

ER\_COLUMNACCESS\_DENIED\_ERROR, 2671

ER\_CONNECT\_TO\_MASTER, 2671

ER\_CON\_COUNT\_ERROR, 2671

ER\_DBACCESS\_DENIED\_ERROR, 2671

ER\_DERIVED\_MUST\_HAVE\_ALIAS, 2671

ER\_DUP\_ENTRY, 2671

ER\_DUP\_FIELDNAME, 2671

ER\_DUP\_KEY, 2671

ER\_DUP\_KEYNAME, 2671

ER\_DUP\_UNIQUE, 2671

ER\_EMPTY\_QUERY, 2671

ER\_FIELD\_SPECIFIED\_TWICE, 2671

ER\_FORCING\_CLOSE, 2671

ER\_GRANT\_WRONG\_HOST\_OR\_USER, 2671

ER\_HANDSHAKE\_ERROR, 2671

ER\_HOST\_IS\_BLOCKED, 2671

ER\_HOST\_NOT\_PRIVILEGED, 2671

ER\_ILLEGAL\_GRANT\_FOR\_TABLE, 2671

ER\_ILLEGAL\_REFERENCE, 2671

ER\_INVALID\_DEFAULT, 2671

ER\_INVALID\_USE\_OF\_NULL, 2671

ER\_IPSOCK\_ERROR, 2671

ER\_KEY\_COLUMN\_DOES\_NOT\_EXITS, 2671

ER\_LOCK\_DEADLOCK, 2671

ER\_LOCK\_WAIT\_TIMEOUT, 2671

ER\_MASTER\_NET\_READ, 2671

---

---

ER\_MASTER\_NET\_WRITE, 2671  
 ER\_MIX\_OF\_GROUP\_FUNC\_AND\_FIELDS, 2671  
 ER\_MULTIPLE\_PRI\_KEY, 2671  
 ER\_NET\_ERROR\_ON\_WRITE, 2671  
 ER\_NET\_FCNTL\_ERROR, 2671  
 ER\_NET\_PACKETS\_OUT\_OF\_ORDER, 2671  
 ER\_NET\_PACKET\_TOO\_LARGE, 2671  
 ER\_NET\_READ\_ERROR, 2671  
 ER\_NET\_READ\_ERROR\_FROM\_PIPE, 2671  
 ER\_NET\_READ\_INTERRUPTED, 2671  
 ER\_NET\_UNCOMPRESS\_ERROR, 2671  
 ER\_NET\_WRITE\_INTERRUPTED, 2671  
 ER\_NEW\_ABORTING\_CONNECTION, 2671  
 ER\_NONEXISTING\_GRANT, 2671  
 ER\_NONEXISTING\_TABLE\_GRANT, 2671  
 ER\_NONUNIQ\_TABLE, 2671  
 ER\_NON\_UNIQ\_ERROR, 2671  
 ER\_NOT\_ALLOWED\_COMMAND, 2671  
 ER\_NOT\_SUPPORTED\_AUTH\_MODE, 2671  
 ER\_NOT\_SUPPORTED\_YET, 2671  
 ER\_NO\_DEFAULT, 2671  
 ER\_NO\_PERMISSION\_TO\_CREATE\_USER, 2671  
 ER\_NO\_REFERENCED\_ROW, 2671  
 ER\_NO\_SUCH\_INDEX, 2671  
 ER\_NO\_SUCH\_TABLE, 2671  
 ER\_NULL\_COLUMN\_IN\_INDEX, 2671  
 ER\_OPERAND\_COLUMNS, 2671  
 ER\_OUTOFMEMORY, 2671  
 ER\_OUT\_OF\_SORTMEMORY, 2671  
 ER\_PARSE\_ERROR, 2671  
 ER\_PASSWORD\_ANONYMOUS\_USER, 2671  
 ER\_PASSWORD\_NOT\_ALLOWED, 2671  
 ER\_PASSWORD\_NO\_MATCH, 2671  
 ER\_PRIMARY\_CANT\_HAVE\_NULL, 2671  
 ER\_READ\_ONLY\_TRANSACTION, 2671  
 ER\_REGEX\_ERROR, 2671  
 ER\_REQUIRES\_PRIMARY\_KEY, 2671  
 ER\_ROW\_IS\_REFERENCED, 2671  
 ER\_SELECT\_REDUCED, 2671  
 ER\_SERVER\_SHUTDOWN, 2671  
 ER\_SPATIAL\_CANT\_HAVE\_NULL, 2671  
 ER\_SUBQUERY\_NO\_1\_ROW, 2671  
 ER\_SYNTAX\_ERROR, 2671  
 ER\_TABLEACCESS\_DENIED\_ERROR, 2671  
 ER\_TABLENAME\_NOT\_ALLOWED\_HERE, 2671  
 ER\_TABLE\_CANT\_HANDLE\_AUTO\_INCREMENT,  
 2671  
 ER\_TABLE\_CANT\_HANDLE\_BLOB, 2671  
 ER\_TABLE\_EXISTS\_ERROR, 2671  
 ER\_TABLE\_MUST\_HAVE\_COLUMNS, 2671  
 ER\_TOO\_BIG\_FIELDLENGTH, 2671  
 ER\_TOO\_BIG\_ROWSIZE, 2671  
 ER\_TOO\_BIG\_SELECT, 2671  
 ER\_TOO\_LONG\_IDENT, 2671  
 ER\_TOO\_LONG\_KEY, 2671  
 ER\_TOO\_LONG\_STRING, 2671  
 ER\_TOO\_MANY\_KEYS, 2671  
 ER\_TOO\_MANY\_KEY\_PARTS, 2671  
 ER\_TOO\_MANY\_ROWS, 2671  
 ER\_TOO\_MANY\_USER\_CONNECTIONS, 2671  
 ER\_UNKNOWN\_CHARACTER\_SET, 2671  
 ER\_UNKNOWN\_COM\_ERROR, 2671  
 ER\_UNKNOWN\_PROCEDURE, 2671  
 ER\_UNKNOWN\_STORAGE\_ENGINE, 2671  
 ER\_UNKNOWN\_TABLE, 2671  
 ER\_UNSUPPORTED\_EXTENSION, 2671  
 ER\_USER\_LIMIT\_REACHED, 2671  
 ER\_WARN\_DATA\_OUT\_OF\_RANGE, 2671  
 ER\_WARN\_DATA\_TRUNCATED, 2671  
 ER\_WARN\_NULL\_TO\_NOTNULL, 2671  
 ER\_WARN\_TOO\_FEW\_RECORDS, 2671  
 ER\_WARN\_TOO\_MANY\_RECORDS, 2671  
 ER\_WRONG\_AUTO\_KEY, 2671  
 ER\_WRONG\_COLUMN\_NAME, 2671  
 ER\_WRONG\_DB\_NAME, 2671  
 ER\_WRONG\_FIELD\_SPEC, 2671  
 ER\_WRONG\_FIELD\_TERMINATORS, 2671  
 ER\_WRONG\_FIELD\_WITH\_GROUP, 2671  
 ER\_WRONG\_FK\_DEF, 2671  
 ER\_WRONG\_GROUP\_FIELD, 2671  
 ER\_WRONG\_KEY\_COLUMN, 2671  
 ER\_WRONG\_NAME\_FOR\_CATALOG, 2671  
 ER\_WRONG\_NAME\_FOR\_INDEX, 2671  
 ER\_WRONG\_NUMBER\_OF\_COLUMNS\_IN\_SELECT,  
 2671  
 ER\_WRONG\_OUTER\_JOIN, 2671  
 ER\_WRONG\_PARAMCOUNT\_TO\_PROCEDURE,  
 2671  
 ER\_WRONG\_SUM\_SELECT, 2671  
 ER\_WRONG\_TABLE\_NAME, 2671  
 ER\_WRONG\_TYPE\_FOR\_VAR, 2671  
 ER\_WRONG\_VALUE\_COUNT, 2671  
 ER\_WRONG\_VALUE\_COUNT\_ON\_ROW, 2671  
 ER\_WRONG\_VALUE\_FOR\_VAR, 2671  
 escape (\), 1012  
 escape sequences  
     option files, 217  
     strings, 1011  
 estimating  
     query performance, 908  
 event  
     restrictions, 3757  
 event groups, 1481  
 event scheduler, 2151  
     thread states, 969  
 Event Scheduler, 2158  
     altering events, 1327  
     and MySQL privileges, 2163  
     and mysqladmin debug, 2162  
     and replication, 2067, 2067  
     and SHOW PROCESSLIST, 2160  
     concepts, 2158  
     creating events, 1344  
     dropping events, 1392  
     enabling and disabling, 2159  
     event metadata, 2161

---

---

- obtaining status information, 2162
- SQL statements, 2161
- starting and stopping, 2159
- time representation, 2162
- event-scheduler option
  - mysqld, 422
- events, 2151, 2158
  - altering, 1327
  - creating, 1344
  - dropping, 1392
  - metadata, 2161
  - status variables, 2165
- EVENTS
  - INFORMATION\_SCHEMA table, 2164, 2183
- events option
  - mysqldump, 293
- events\_stages\_current table
  - performance\_schema, 2278
- events\_stages\_history table
  - performance\_schema, 2279
- events\_stages\_history\_long table
  - performance\_schema, 2279
- events\_stages\_summary\_by\_account\_by\_event\_name table
  - performance\_schema, 2296
- events\_stages\_summary\_by\_host\_by\_event\_name table
  - performance\_schema, 2296
- events\_stages\_summary\_by\_thread\_by\_event\_name table
  - performance\_schema, 2289
- events\_stages\_summary\_by\_user\_by\_event\_name table
  - performance\_schema, 2296
- events\_stages\_summary\_global\_by\_event\_name table
  - performance\_schema, 2289
- events\_statements\_current table
  - performance\_schema, 2280
- events\_statements\_history table
  - performance\_schema, 2283
- events\_statements\_history\_long table
  - performance\_schema, 2283
- events\_statements\_summary\_by\_account\_by\_event\_name table
  - performance\_schema, 2296
- events\_statements\_summary\_by\_host\_by\_event\_name table
  - performance\_schema, 2296
- events\_statements\_summary\_by\_thread\_by\_event\_name table
  - performance\_schema, 2290
- events\_statements\_summary\_by\_user\_by\_event\_name table
  - performance\_schema, 2296
- events\_statements\_summary\_global\_by\_event\_name table
  - performance\_schema, 2290
- events\_waits\_current table
  - performance\_schema, 2275
- events\_waits\_history table
  - performance\_schema, 2277
- events\_waits\_history\_long table
  - performance\_schema, 2277
- events\_waits\_summary\_by\_account\_by\_event\_name table
  - performance\_schema, 2296
- events\_waits\_summary\_by\_host\_by\_event\_name table
  - performance\_schema, 2296
- events\_waits\_summary\_by\_instance table
  - performance\_schema, 2288
- events\_waits\_summary\_by\_thread\_by\_event\_name table
  - performance\_schema, 2288
- events\_waits\_summary\_by\_user\_by\_event\_name table
  - performance\_schema, 2296
- events\_waits\_summary\_global\_by\_event\_name table
  - performance\_schema, 2288
- event\_scheduler system variable, 480
- eviction, 3790
- exact-value literals, 1315
- example option
  - mysqld\_multi, 232
- example programs
  - C API, 2797
- EXAMPLE storage engine, 1601, 1814
- examples
  - compressed tables, 333
  - myisamchk output, 324
  - queries, 193
- exclude-gtids option
  - mysqlbinlog, 351
- exclusive lock, 3790
- Execute
  - thread command, 958
- EXECUTE, 1486, 1490
- execute option
  - mysql, 250
- ExecuteNonQuery method, 2445
- ExecuteReader method, 2445
- ExecuteScalar method, 2445
- executing
  - thread state, 961
- executing SQL statements from text files, 192, 262
- Execution of init\_command
  - thread state, 961
- EXISTS
  - with subqueries, 1449
- exit command
  - mysql, 257
- exit-info option
  - mysqld, 422
- EXP(), 1204
- expire\_logs\_days system variable, 480
- expiring passwords, 781
- EXPLAIN, 894, 1596

---

---

EXPLAIN PARTITIONS, 2128, 2129  
EXPLAIN used with partitioned tables, 2128  
explicit default values, 1146  
explicit\_defaults\_for\_timestamp system variable, 481  
EXPORT\_SET(), 1179  
expression aliases, 1289, 1429  
expression syntax, 1032  
expressions  
    extended, 186  
extend-check option  
    myisamchk, 321, 322  
extended option  
    mysqlcheck, 277  
extended-insert option  
    mysqldump, 294  
extensions  
    to standard SQL, 26  
extent, 3790  
ExteriorRing(), 1306  
external locking, 422, 439, 540, 848, 930, 964  
external-locking option  
    mysqld, 422  
external\_user session variable, 482  
extra-file option  
    my\_print\_defaults, 374  
EXTRACT(), 1217  
extracting  
    dates, 183  
ExtractValue(), 1249

## F

failover  
    Java clients, 2643  
FALSE, 1013, 1016  
    testing for, 1167, 1168  
FAQs  
    C API, 2895  
    Connector/ODBC, 2392  
    Connectors and APIs, 3652  
    MySQL Cluster, 3640  
    replication, 3652  
Fast Index Creation, 674, 3791  
fast option  
    myisamchk, 321  
    mysqlcheck, 277  
fast shutdown, 3791  
features of MySQL, 6  
FEDERATED storage engine, 1601, 1809  
Fetch  
    thread command, 958  
FETCH, 1499  
field  
    changing, 1334  
Field List  
    thread command, 958  
FIELD(), 1179  
fields-enclosed-by option  
    mysqldump, 290, 300

fields-escaped-by option  
    mysqldump, 290, 300  
fields-optionally-enclosed-by option  
    mysqldump, 290, 300  
fields-terminated-by option  
    mysqldump, 290, 300  
FILE, 1181  
file format, 667, 3791  
    Antelope, 664  
    Barracuda, 657  
    downgrading, 672  
    identifying, 671  
file format management  
    enabling new file formats, 1688  
file per table, 1688  
file-per-table, 3791  
files  
    binary log, 624  
    error messages, 1085  
    general query log, 623  
    log, 637  
    my.cnf, 2057  
    not found message, 3735  
    permissions, 3735  
    repairing, 321  
    script, 192  
    size limits, 3768  
    slow query log, 635  
    text, 262, 297  
    tmp, 145  
FILES  
    INFORMATION\_SCHEMA table, 2186  
filesort optimization, 997  
file\_instances table  
    performance\_schema, 2271  
file\_summary\_by\_event\_name table  
    performance\_schema, 2292  
file\_summary\_by\_instance table  
    performance\_schema, 2292  
fill factor, 1636, 3791  
FIND\_IN\_SET(), 1179  
Finished reading one binlog; switching to next binlog  
    thread state, 967  
fix-db-names option  
    mysqlcheck, 277  
fix-table-names option  
    mysqlcheck, 277  
FIXED data type, 1114  
fixed row format, 3792  
fixed-point arithmetic, 1315  
FLOAT data type, 1114, 1114, 1115  
floating-point number, 1115  
floating-point values  
    and replication, 2068  
floats, 1013  
FLOOR(), 1204  
FLUSH, 1590  
    and replication, 2069

---

- flush, 3792
- flush list, 3792
- flush list mutex, 1662
- flush option
  - mysqld, 423
- flush system variable, 482
- flush tables, 268
- flush-logs option
  - mysqldump, 294
- flush-privileges option
  - mysqldump, 295
- Flushing tables
  - thread state, 961
- flushlog option
  - mysqlhotcopy, 367
- flush\_time system variable, 482
- FOR UPDATE, 1432
- FORCE INDEX, 1441, 3749
- FORCE KEY, 1441
- force option
  - myisamchk, 321, 322
  - myisampack, 332
  - mysql, 250
  - mysqladmin, 271
  - mysqlcheck, 277
  - mysqldump, 287
  - mysqlimport, 300
  - mysql\_convert\_table\_format, 369
  - mysql\_install\_db, 237
  - mysql\_upgrade, 243
- force-if-open option
  - mysqlbinlog, 351
- force-read option
  - mysqlbinlog, 351
- foreign key, 3792
  - constraint, 35, 35
  - deleting, 1336, 1381
- FOREIGN KEY constraint, 3792
- foreign key constraints, 1379
  - InnoDB, 655
  - restrictions, 655
- FOREIGN KEY constraints
  - and online DDL, 704
  - and TRUNCATE TABLE, 1689
- foreign keys, 32, 196, 1335, 2419
- foreign\_key\_checks session variable, 483
- FORMAT(), 1180
- Forums, 20
- FOUND\_ROWS(), 1268
- fractional seconds
  - and replication, 2068
- fractional seconds precision, 1111, 1115
- FreeBSD troubleshooting, 138
- freeing items
  - thread state, 961
- FROM, 1429
- FROM\_BASE64(), 1180
- FROM\_DAYS(), 1218
- FROM\_UNIXTIME(), 1218
- FTS, 3792
- ft\_boolean\_syntax system variable, 483
- ft\_max\_word\_len myisamchk variable, 319
- ft\_max\_word\_len system variable, 484
- ft\_min\_word\_len myisamchk variable, 319
- ft\_min\_word\_len system variable, 484
- ft\_query\_expansion\_limit system variable, 484
- ft\_stopword\_file myisamchk variable, 319
- ft\_stopword\_file system variable, 485
- full backup, 3793
- full disk, 3740
- full table scan, 3793
- full table scans
  - avoiding, 862
- full-text parser plugins, 3487
- full-text search, 1229, 3793
- FULLTEXT, 1229
- fulltext
  - stopword list, 1241
- FULLTEXT index, 3793
  - InnoDB, 1636
- FULLTEXT initialization
  - thread state, 961
- fulltext join type
  - optimizer, 899
- function
  - creating, 1544
  - deleting, 1545
- function names
  - parsing, 1023
  - resolving ambiguity, 1023
- functions, 1153
  - and replication, 2069
  - arithmetic, 1258
  - bit, 1258
  - C API, 2805
  - C prepared statement API, 2864, 2865
  - cast, 1244
  - control flow, 1173
  - date and time, 1209
  - encryption, 1259
  - for SELECT and WHERE clauses, 1153
  - GROUP BY, 1281
  - grouping, 1165
  - GTIDs, 1273
  - information, 1265
  - mathematical, 1202
  - miscellaneous, 1274
  - native
    - adding, 3542
    - new, 3532
    - stored, 2153
    - string, 1175
    - string comparison, 1190
    - user-defined, 1544, 1545, 3532
      - adding, 3533
  - fuzzy checkpointing, 3793

---

## G

- GA, 3793
- gap, 3793
- gap lock, 3793
  - InnoDB, 1623, 1628, 1629, 1730
- gb2312, gbk, 3640
- gdb
  - using, 3545
- gdb option
  - mysqld, 423
- general information, 1
- General Public License, 5
- general query log, 623, 3794
- general-log option
  - mysqld, 423
- general\_log system variable, 485
- general\_log\_file system variable, 486
- geographic feature, 1290
- GeomCollFromText(), 1298
- GeomCollFromWKB(), 1299
- geometry, 1290
- GEOMETRY data type, 1297
- GEOMETRYCOLLECTION data type, 1297
- GeometryCollection(), 1300
- GeometryCollectionFromText(), 1298
- GeometryCollectionFromWKB(), 1299
- GeometryFromText(), 1298
- GeometryFromWKB(), 1299
- GeometryN(), 1308
- GeometryType(), 1303
- GeomFromText(), 1298, 1303
- GeomFromWKB(), 1299, 1303
- geospatial feature, 1290
- GET DIAGNOSTICS, 1503
- getting MySQL, 53
- GET\_FORMAT(), 1218
- GET\_LOCK(), 1275
- GIS, 1289, 1289
- Glassfish application server, 2698
- GLength(), 1305, 1306
- global privileges, 1524, 1534
- globalization, 1037
- GLOBAL\_STATUS
  - INFORMATION\_SCHEMA table, 2188
- global\_transaction, 3794
- GLOBAL\_VARIABLES
  - INFORMATION\_SCHEMA table, 2188
- go command
  - mysql, 257
- Google Test, 135
- got handler lock
  - thread state, 965
- got old table
  - thread state, 965
- GRANT, 1524
- GRANT statement, 774
- grant tables

- re-creating, 145
- sorting, 765, 767
- structure, 756

- granting
  - privileges, 1524
- GRANTS, 1563
- greater than (>), 1167
- greater than or equal (>=), 1167
- GREATEST(), 1169
- GROUP BY, 998
  - aliases in, 1289
  - extensions to standard SQL, 1288, 1430
- GROUP BY functions, 1281
- group commit, 1656, 3794
- grouping
  - expressions, 1165
- GROUP\_CONCAT(), 1283
- group\_concat\_max\_len system variable, 486
- GTID functions, 1273
- gtid-mode option (mysqld), 2019
- gtid\_done system variable, 2020
- gtid\_executed system variable, 2021
- gtid\_lost system variable, 2022
- gtid\_mode system variable, 2022
- gtid\_next system variable, 2022
- gtid\_owned system variable, 2023
- gtid\_purged system variable, 2023
- GTID\_SUBSET(), 1274
- GTID\_SUBTRACT(), 1274

## H

- HANDLER, 1402
- Handlers, 1500
- handling
  - errors, 3539
- hash index, 3794
- hash indexes, 876
- hash partitioning, 2104
- hash partitions
  - managing, 2121
  - splitting and merging, 2122
- have\_compress system variable, 486
- have\_crypt system variable, 486
- have\_csv system variable, 486
- have\_dynamic\_loading system variable, 487
- HAVE\_EMBEDDED\_PRIVILEGE\_CONTROL option
  - CMake, 137
- have\_geometry system variable, 487
- have\_innodb system variable, 487
- have\_openssl system variable, 487
- have\_partitioning system variable, 487
- have\_profiling system variable, 487
- have\_query\_cache system variable, 487
- have\_rtree\_keys system variable, 487
- have\_ssl system variable, 487
- have\_symlink system variable, 487
- HAVING, 1430
- HDD, 3794



---

header\_file option  
     comp\_err, 235  
 HEAP storage engine, 1601, 1795  
 heartbeat, 3794  
 help command  
     mysql, 256  
 help option  
     comp\_err, 235  
     myisamchk, 318  
     myisampack, 332  
     myisam\_ftdump, 314  
     mysql, 248  
     mysqlaccess, 344  
     mysqladmin, 270  
     mysqlbinlog, 348  
     mysqlcheck, 276  
     mysqld, 413  
     mysqldump, 287  
     mysqldumpslow, 364  
     mysqld\_multi, 232  
     mysqld\_safe, 227  
     mysqlhotcopy, 367  
     mysqlimport, 299  
     MySQLInstallerConsole, 83  
     mysqlshow, 304  
     mysqlslap, 309  
     mysql\_config\_editor, 342  
     mysql\_convert\_table\_format, 369  
     mysql\_find\_rows, 370  
     mysql\_plugin, 240  
     mysql\_setpermission, 370  
     mysql\_upgrade, 243  
     mysql\_waitpid, 371  
     my\_print\_defaults, 374  
     perror, 376  
     resolveip, 377  
     resolve\_stack\_dump, 375  
 HELP option  
     myisamchk, 319  
 HELP statement, 1598  
 HEX(), 1180, 1204  
 hex-blob option  
     mysqldump, 290  
 hexadecimal literals, 1016  
 hexdump option  
     mysqlbinlog, 351  
 high-water mark, 3794  
 HIGH\_NOT\_PRECEDENCE SQL mode, 603  
 HIGH\_PRIORITY, 1432  
 hints, 27  
     index, 1429, 1441  
 histignore option  
     mysql, 250  
 history list, 3795  
 history of MySQL, 8  
 HOME environment variable, 166, 260  
 host name  
     default, 209  
     host name caching, 947  
     host name resolution, 947  
     host names, 209  
         in account names, 762  
         in default accounts, 150  
 host option, 211  
     mysql, 250  
     mysqlaccess, 344  
     mysqladmin, 271  
     mysqlbinlog, 351  
     mysqlcheck, 277  
     mysqldump, 284  
     mysqlhotcopy, 367  
     mysqlimport, 300  
     mysqlshow, 305  
     mysqlslap, 311  
     mysql\_config\_editor, 342  
     mysql\_convert\_table\_format, 369  
     mysql\_setpermission, 371  
 host table  
     sorting, 767  
 host.frm  
     problems finding, 141  
 hostname system variable, 488  
 hosts table  
     performance\_schema, 2285  
 host\_cache table  
     performance\_schema, 2299  
 hot, 3795  
 hot backup, 3795  
 HOUR(), 1219  
 howto option  
     mysqlaccess, 344  
 html option  
     mysql, 250

**I**

i-am-a-dummy option  
     mysql, 252  
 ib-file set, 668, 3796  
 ibbackup command, 3796  
 ibbackup\_logfile, 3796  
 ibdata file, 3796  
 ib\_logfile, 3796  
 icc  
     and MySQL Cluster support, 3543  
     MySQL builds, 62  
 ID  
     unique, 2896  
 identifiers, 1017  
     case sensitivity, 1020  
     quoting, 1017  
 identity session variable, 488  
 IF, 1495  
 IF(), 1174  
 IFNULL(), 1175  
 IGNORE  
     with partitioned tables, 1407

---

---

IGNORE INDEX, 1441  
 IGNORE KEY, 1441  
 ignore option  
   mysqlimport, 300  
 ignore-builtin-innodb option  
   mysqld, 1697  
 ignore-db-dir option  
   mysqld, 424  
 ignore-lines option  
   mysqlimport, 300  
 ignore-spaces option  
   mysql, 250  
 ignore-table option  
   mysqldump, 293  
 IGNORE\_AIO\_CHECK option  
   CMake, 135  
 ignore\_builtin\_innodb system variable, 1700  
 ignore\_db\_dirs system variable, 488  
 IGNORE\_SPACE SQL mode, 604  
 ilist, 3796  
 implicit default values, 1146  
 implicit row lock, 3797  
 IMPORT TABLESPACE, 1336, 1619  
 importing  
   data, 262, 297  
 IN, 1169, 1447  
 in-memory database, 3797  
 include option  
   mysql\_config, 373  
 include-gtids option  
   mysqlbinlog, 351  
 include-master-host-port option  
   mysqldump, 288  
 increasing with replication  
   speed, 1929  
 incremental backup, 3797  
 incremental recovery, 845  
 index, 3797  
   deleting, 1334, 1392  
   rebuilding, 163  
 index cache, 3797  
 INDEX DIRECTORY  
   and replication, 2067  
 index dives (for statistics estimation), 1690  
 index hint, 3797  
 index hints, 1429, 1441  
 index join type  
   optimizer, 899  
 index prefix, 3798  
 index-record lock  
   InnoDB, 1623, 1628, 1629, 1730  
 indexes, 1349  
   and BLOB columns, 873, 1363  
   and IS NULL, 877  
   and LIKE, 876  
   and NULL values, 1363  
   and TEXT columns, 873, 1363  
   assigning to key cache, 1588  
   block size, 492  
   columns, 872  
   creating and dropping, 702  
   leftmost prefix of, 874  
   multi-column, 873  
   multiple-part, 1349  
   names, 1017  
   primary (clustered) and secondary, 702  
   use of, 871  
 Indexes/Keys dialog, 2419  
 index\_merge join type  
   optimizer, 899  
 index\_subquery join type  
   optimizer, 899  
 INET6\_ATON(), 1276  
 INET6\_NTOA(), 1277  
 INET\_ATON(), 1276  
 INET\_NTOA(), 1276  
 infimum record, 3798  
 information functions, 1265  
 information option  
   myisamchk, 321  
 INFORMATION SCHEMA  
   InnoDB tables, 1665  
 INFORMATION\_SCHEMA, 2178, 3798  
   collation and searching, 1058  
   INNODB\_CMP table, 1665  
   INNODB\_CMPMEM table, 1666  
   INNODB\_CMPMEM\_RESET table, 1666  
   INNODB\_CMP\_RESET table, 1665  
   INNODB\_LOCKS table, 1667  
   INNODB\_LOCK\_WAITS table, 1667  
   INNODB\_METRICS table, 2217  
   INNODB\_TRX table, 1667  
 INFORMATION\_SCHEMA plugins, 3489  
 init  
   thread state, 961  
 Init DB  
   thread command, 958  
 init-command option  
   mysql, 250  
 init-file option  
   mysqld, 424  
 Initialized  
   thread state, 969  
 init\_connect system variable, 489  
 init\_file system variable, 489  
 init\_slave system variable, 1987  
 INNER JOIN, 1435  
 innobackup command, 3798  
 innochecksum, 206, 313  
 InnoDB, 1605, 3798  
   adaptive hash index, 1637  
   auto-increment columns, 650  
   autocommit mode, 645, 1633  
   backups, 1616  
   checkpoints, 640  
   clustered index, 1635

---

---

configuration parameters, 1692  
 configuring data files and memory allocation, 1609  
 considerations as default storage engine, 1606  
 consistent reads, 1625  
 crash recovery, 1639  
 data files, 1615  
 deadlock detection, 1633  
 disk I/O, 638  
 file space management, 639  
 file-per-table setting, 641  
 foreign key constraints, 655  
 FULLTEXT index, 1636  
 gap lock, 1623, 1628, 1629, 1730  
 index-record lock, 1623, 1628, 1629, 1730  
 indexes, 1635  
 insert buffering, 1637  
 limits and restrictions, 1756  
 lock modes, 1623  
 locking, 1623  
 locking reads, 1627  
 log files, 1615  
 migrating tables, 1618  
 Monitors, 639, 1618, 1673, 1683, 1685  
 multi-versioning, 1635  
 next-key lock, 1623, 1628, 1629, 1730  
 NFS, 1610, 1756  
 online DDL, 674  
 page size, 1636, 1757  
 raw partitions, 1616  
 record-level locks, 1623, 1628, 1629, 1730  
 replication, 1620  
 row structure, 1638  
 secondary index, 1635  
 semi-consistent read, 1730  
 Solaris 10 x86\_64 issues, 118  
 storage requirements, 1148  
 system tablespace setup, 1614  
 system variables, 1692  
 tables, 640, 1635
 

- converting from other storage engines, 645

 transaction model, 1623  
 troubleshooting, 1683
 

- data dictionary problems, 1685
- deadlocks, 1633
- defragmenting tables, 640
- I/O problems, 1683
- InnoDB error codes, 1640
- online DDL, 704
- OS error codes, 1641
- performance problems, 883
- recovery problems, 1684
- SQL errors, 1640

 InnoDB buffer pool, 911  
 InnoDB Hot Backup, 3798  
 innodb option
 

- mysqld, 1698

 InnoDB parameters, new, 1755
 

- innodb\_adaptive\_flushing, 1657
- innodb\_change\_buffering, 1653
- innodb\_file\_format, 1688
- innodb\_file\_format\_check, 669
- innodb\_io\_capacity, 1657
- innodb\_large\_prefix, 1728
- innodb\_read\_ahead\_threshold, 1655
- innodb\_read\_io\_threads, 1656
- innodb\_spin\_wait\_delay, 1658
- innodb\_stats\_sample\_pages, 1690
- innodb\_strict\_mode, 1690
- innodb\_use\_sys\_malloc, 1652
- innodb\_write\_io\_threads, 1656

 InnoDB parameters, with new defaults, 1756
 

- innodb\_additional\_mem\_pool\_size, 1756
- innodb\_buffer\_pool\_size, 1756
- innodb\_change\_buffering, 1756
- innodb\_file\_format\_check, 1756
- innodb\_log\_buffer\_size, 1756
- innodb\_max\_dirty\_pages\_pct, 1657, 1756
- innodb\_sync\_spin\_loops, 1756
- innodb\_thread\_concurrency, 1756

 InnoDB storage engine, 1601, 1605  
 InnoDB tables, 30  
 innodb-status-file option
 

- mysqld, 1698

 innodb\_adaptive\_flushing, 1657  
 innodb\_adaptive\_flushing system variable, 1700  
 innodb\_adaptive\_flushing\_lwm system variable, 1701  
 innodb\_adaptive\_hash\_index, 1653
 

- and innodb\_thread\_concurrency, 1654
- dynamically changing, 1689

 innodb\_adaptive\_hash\_index system variable, 1701  
 innodb\_adaptive\_max\_sleep\_delay system variable, 1702  
 innodb\_additional\_mem\_pool\_size system variable, 1702
 

- and innodb\_use\_sys\_malloc, 1652

 innodb\_api\_bk\_commit\_interval system variable, 1702  
 innodb\_api\_disable\_rowlock system variable, 1703  
 innodb\_api\_enable\_binlog system variable, 1703  
 innodb\_api\_enable\_mdlog system variable, 1703  
 innodb\_api\_trx\_level system variable, 1704  
 innodb\_autoextend\_increment system variable, 1704  
 innodb\_autoinc\_lock\_mode, 3798  
 innodb\_autoinc\_lock\_mode system variable, 1705  
 INNODB\_BUFFER\_PAGE table, 2214  
 INNODB\_BUFFER\_PAGE\_LRU table, 2215  
 innodb\_buffer\_pool\_dump\_at\_shutdown system variable, 1705  
 innodb\_buffer\_pool\_dump\_now system variable, 1706  
 innodb\_buffer\_pool\_filename system variable, 1706  
 innodb\_buffer\_pool\_instances system variable, 1706  
 innodb\_buffer\_pool\_load\_abort system variable, 1707  
 innodb\_buffer\_pool\_load\_at\_startup system variable, 1707  
 innodb\_buffer\_pool\_load\_now system variable, 1708  
 innodb\_buffer\_pool\_size system variable, 1708  
 INNODB\_BUFFER\_POOL\_STATS table, 2216

---

innodb\_change\_buffering, 1653  
innodb\_change\_buffering system variable, 1709  
innodb\_change\_buffer\_max\_size system variable, 1709  
innodb\_checksums system variable, 1711  
innodb\_checksum\_algorithm system variable, 1710  
INNODB\_CMP table, 2205  
INNODB\_CMPMEM table, 2207  
INNODB\_CMPMEM\_RESET table, 2207  
INNODB\_CMP\_PER\_INDEX table, 2206  
innodb\_cmp\_per\_index\_enabled system variable, 1712  
INNODB\_CMP\_PER\_INDEX\_RESET table, 2206  
INNODB\_CMP\_RESET table, 2205  
innodb\_commit\_concurrency system variable, 1712  
innodb\_compression\_failure\_threshold\_pct system variable, 1713  
innodb\_compression\_level system variable, 1713  
innodb\_compression\_pad\_pct\_max system variable, 1713  
innodb\_concurrency\_tickets, 1654  
innodb\_concurrency\_tickets system variable, 1714  
innodb\_data\_file\_path system variable, 1714  
innodb\_data\_home\_dir system variable, 1715  
innodb\_disable\_sort\_file\_cache system variable, 1715  
innodb\_doublewrite system variable, 1715  
innodb\_fast\_shutdown system variable, 1716  
innodb\_file\_format, 667, 3799  
    Antelope, 664  
    Barracuda, 657  
    enabling new file formats, 1688  
    identifying, 671  
innodb\_file\_format system variable, 1716  
innodb\_file\_format\_check, 669  
innodb\_file\_format\_check system variable, 1717  
innodb\_file\_format\_max system variable, 1717  
innodb\_file\_per\_table, 657, 3799  
    dynamically changing, 1688  
innodb\_file\_per\_table system variable, 1718  
innodb\_flushing\_avg\_loops system variable, 1722  
innodb\_flush\_log\_at\_timeout system variable, 1718  
innodb\_flush\_log\_at\_trx\_commit system variable, 1718  
innodb\_flush\_method system variable, 1719  
innodb\_flush\_neighbors system variable, 1721  
innodb\_force\_load\_corrupted system variable, 1722  
innodb\_force\_recovery system variable, 1722  
innodb\_ft\_aux\_table system variable, 1723  
INNODB\_FT\_BEING\_DELETED table, 2225  
innodb\_ft\_cache\_size system variable, 1724  
INNODB\_FT\_CONFIG table, 2222  
INNODB\_FT\_DEFAULT\_STOPWORD table, 2223  
INNODB\_FT\_DELETED table, 2224  
innodb\_ft\_enable\_stopword system variable, 1724  
INNODB\_FT\_INDEX\_CACHE table, 2224  
INNODB\_FT\_INDEX\_TABLE table, 2223  
innodb\_ft\_max\_token\_size system variable, 1724  
innodb\_ft\_min\_token\_size system variable, 1725  
innodb\_ft\_num\_word\_optimize system variable, 1725  
innodb\_ft\_server\_stopword\_table system variable, 1725  
innodb\_ft\_sort\_pll\_degree system variable, 1726  
innodb\_ft\_user\_stopword\_table system variable, 1726  
innodb\_index\_stats table, 1650  
innodb\_io\_capacity, 1657  
innodb\_io\_capacity system variable, 1727  
innodb\_io\_capacity\_max system variable, 1728  
innodb\_large\_prefix system variable, 1728  
INNODB\_LOCKS table, 2209  
innodb\_locks\_unsafe\_for\_binlog system variable, 1729  
INNODB\_LOCK\_WAITS table, 2210  
innodb\_lock\_wait\_timeout, 3799  
    dynamically changing, 1689  
innodb\_lock\_wait\_timeout system variable, 1729  
innodb\_log\_buffer\_size system variable, 1732  
innodb\_log\_compressed\_pages system variable, 1732  
innodb\_log\_files\_in\_group system variable, 1733  
innodb\_log\_file\_size system variable, 1732  
innodb\_log\_group\_home\_dir system variable, 1733  
innodb\_lru\_scan\_depth system variable, 1734  
innodb\_max\_dirty\_pages\_pct, 1657  
innodb\_max\_dirty\_pages\_pct system variable, 1734  
innodb\_max\_dirty\_pages\_pct\_lwm system variable, 1735  
innodb\_max\_purge\_lag system variable, 1735  
innodb\_max\_purge\_lag\_delay system variable, 1736  
innodb\_memcache database, 1766, 1783  
innodb\_memcached\_config.sql script, 1766  
INNODB\_METRICS table, 2217  
innodb\_mirrored\_log\_groups system variable, 1736  
innodb\_monitor\_disable system variable, 1736  
innodb\_monitor\_enable system variable, 1737  
innodb\_monitor\_reset system variable, 1737  
innodb\_monitor\_reset\_all system variable, 1737  
innodb\_old\_blocks\_pct, 1659  
innodb\_old\_blocks\_pct system variable, 1738  
innodb\_old\_blocks\_time, 1659  
innodb\_old\_blocks\_time system variable, 1738  
innodb\_online\_alter\_log\_max\_size system variable, 1739  
innodb\_open\_files system variable, 1739  
innodb\_optimize\_fulltext\_only system variable, 1740  
innodb\_page\_size system variable, 1740  
innodb\_print\_all\_deadlocks system variable, 1741  
    innodb\_print\_all\_deadlocks, 1741  
innodb\_purge\_batch\_size system variable, 1741  
innodb\_purge\_threads system variable, 1742  
innodb\_random\_read\_ahead system variable, 1742  
innodb\_read\_ahead\_threshold, 1655  
innodb\_read\_ahead\_threshold system variable, 1742  
innodb\_read\_io\_threads, 1656  
innodb\_read\_io\_threads system variable, 1743  
innodb\_read\_only system variable, 1743  
innodb\_replication\_delay system variable, 1744  
innodb\_rollback\_on\_timeout system variable, 1744  
innodb\_rollback\_segments system variable, 1744  
innodb\_sort\_buffer\_size system variable, 1745

---

innodb\_spin\_wait\_delay, 1658  
 innodb\_spin\_wait\_delay system variable, 1745  
 innodb\_stats\_auto\_recalc system variable, 1746  
 innodb\_stats\_method system variable, 1746  
 innodb\_stats\_on\_metadata system variable, 1689, 1747  
 innodb\_stats\_persistent system variable  
     innodb\_stats\_persistent, 1747  
 innodb\_stats\_persistent\_sample\_pages system variable, 1748  
 innodb\_stats\_sample\_pages, 1690  
 innodb\_stats\_sample\_pages system variable, 1748  
 innodb\_stats\_transient\_sample\_pages system variable, 1748  
 innodb\_stat\_persistent system variable, 1747  
 innodb\_strict\_mode, 1690, 3799  
 innodb\_strict\_mode system variable, 1749  
 innodb\_support\_xa system variable, 1749  
 innodb\_sync\_array\_size system variable, 1750  
 innodb\_sync\_spin\_loops system variable, 1750  
 INNODB\_SYS\_COLUMNS table, 2212  
 INNODB\_SYS\_FIELDS table, 2212  
 INNODB\_SYS\_FOREIGN table, 2213  
 INNODB\_SYS\_FOREIGN\_COLS table, 2213  
 INNODB\_SYS\_INDEXES table, 2211  
 INNODB\_SYS\_TABLES table, 2210  
 INNODB\_SYS\_TABLESTATS table, 2214  
 innodb\_table\_locks system variable, 1751  
 innodb\_table\_stats table, 1650  
 innodb\_thread\_concurrency, 1654  
 innodb\_thread\_concurrency system variable, 1751  
 innodb\_thread\_sleep\_delay, 1654  
 innodb\_thread\_sleep\_delay system variable, 1752  
 INNODB\_TRX table, 2207  
 innodb\_undo\_directory system variable, 1752  
 innodb\_undo\_logs system variable, 1752  
 innodb\_undo\_tablespace system variable, 1753  
 innodb\_use\_native\_aio system variable, 1753  
 innodb\_use\_sys\_malloc  
     and innodb\_thread\_concurrency, 1654  
 innodb\_use\_sys\_malloc system variable, 1652, 1754  
 innodb\_version system variable, 1754  
 innodb\_write\_io\_threads, 1656  
 innodb\_write\_io\_threads system variable, 1754  
 INSERT, 862, 1403  
 insert, 3799  
     thread state, 966  
 INSERT ... SELECT, 1407  
 insert buffer, 3799  
 insert buffering, 1637, 3800  
     disabling, 1653  
 INSERT DELAYED, 1408, 1408  
 INSERT statement  
     grant privileges, 776  
 INSERT(), 1180  
 insert-ignore option  
     mysqldump, 294  
 insertable views  
     insertable, 2167  
 inserting  
     speed of, 862  
 inserts  
     concurrent, 926, 929  
 insert\_id session variable, 490  
 install option  
     mysqld, 424  
 INSTALL PLUGIN, 1545  
 install-manual option  
     mysqld, 424  
 Installation, 83  
 installation layouts, 62  
 installation overview, 122  
 installing  
     binary distribution, 62  
     installing on HP-UX, 120  
     Linux RPM packages, 110  
     Mac OS X PKG packages, 100  
     overview, 48  
     Perl, 167  
     Perl on Windows, 168  
     Solaris PKG packages, 118  
     source distribution, 122  
     user-defined functions, 3539  
 installing plugins, 609, 1545  
 INSTALL\_BINDIR option  
     CMake, 132  
 INSTALL\_DOCDIR option  
     CMake, 132  
 INSTALL\_DOCREADMEDIR option  
     CMake, 132  
 INSTALL\_INCLUDEDIR option  
     CMake, 132  
 INSTALL\_INFODIR option  
     CMake, 132  
 INSTALL\_LAYOUT option  
     CMake, 132  
 INSTALL\_LIBDIR option  
     CMake, 132  
 INSTALL\_MANDIR option  
     CMake, 132  
 INSTALL\_MYSQLSHAREDIR option  
     CMake, 132  
 INSTALL\_MYSQLTESTDIR option  
     CMake, 133  
 INSTALL\_PLUGINDIR option  
     CMake, 133  
 INSTALL\_SBINDIR option  
     CMake, 133  
 INSTALL\_SCRIPTDIR option  
     CMake, 133  
 INSTALL\_SHAREDIR option  
     CMake, 133  
 INSTALL\_SQLBENCHDIR option  
     CMake, 133  
 INSTALL\_SUPPORTFILESDIR option  
     CMake, 133

---

---

- instance, 3800
- INSTR(), 1181
- instrumentation, 3800
- INT data type, 1113
- integer arithmetic, 1315
- INTEGER data type, 1113
- integers, 1013
- intention lock, 3800
- interactive\_timeout system variable, 490
- interceptors, 2503, 2523
- InteriorRingN(), 1307
- internal locking, 925
- internal memory allocator
  - disabling, 1652
- internals, 3483
- internationalization, 1037
- Internet Relay Chat, 21
- Intersection(), 1309
- Intersects(), 1312
- INTERVAL(), 1170
- INTO
  - SELECT, 1433
- introducer
  - string literal, 1012, 1043
- invalid data
  - constraint, 36
- invalidating query cache entries
  - thread state, 966
- inverted index, 3800
- in\_file option
  - comp\_err, 235
- IOPS, 3800
- IP addresses
  - in account names, 762
  - in default accounts, 150
- IPv6 addresses
  - for JDBC URLs, 2642
  - in account names, 762
  - in default accounts, 150
- IPv6 connections, 150
- IRC, 21
- IS boolean\_value, 1167
- IS NOT boolean\_value, 1168
- IS NOT NULL, 1168
- IS NULL, 981, 1168
  - and indexes, 877
- IsClosed(), 1306
- IsEmpty(), 1304
- ISNULL(), 1170
- ISOLATION LEVEL, 1468
- isolation level, 1623, 3800, 4091
- IsRing(), 1306
- IsSimple(), 1304
- IS\_FREE\_LOCK(), 1278
- IS\_IPV4(), 1278
- IS\_IPV4\_COMPAT(), 1278
- IS\_IPV4\_MAPPED(), 1278
- IS\_IPV6(), 1279

- IS\_USED\_LOCK(), 1279
- ITERATE, 1496
- iterations option
  - mysqslap, 311

## J

- J2EE
  - connection pooling, 2684
  - load balancing, 2687
- Japanese character sets
  - conversion, 3640
- Japanese, Korean, Chinese character sets
  - frequently asked questions, 3640
- JBoss application server, 2693
- JDBC, 2323
  - and MySQL data types, 2664
  - background information for Connector/J, 2675
  - character sets, 2666
  - CLASSPATH, 2637
  - code examples, 2641
  - compatibility, 2661
  - configuration properties, 2642
  - driver for MySQL, 2635
  - SQLState codes, 2671
  - troubleshooting, 2639, 2704
  - versions supported, 2635
- join, 3801
  - nested-loop algorithm, 986
- JOIN, 1435
- join algorithm
  - Block Nested-Loop, 982
  - Nested-Loop, 982
- join option
  - myisampack, 333
- join type, 3953
  - ALL, 900
  - const, 898
  - eq\_ref, 898
  - fulltext, 899
  - index, 899
  - index\_merge, 899
  - index\_subquery, 899
  - range, 899
  - ref, 898
  - ref\_or\_null, 899
  - system, 898
  - unique\_subquery, 899
- join\_buffer\_size system variable, 490

## K

- keepold option
  - mysqlhotcopy, 367
- keep\_files\_on\_create system variable, 491
- Key cache
  - MyISAM, 914
- key cache
  - assigning indexes to, 1588

---

key partitioning, 2107  
 key partitions  
   managing, 2122  
   splitting and merging, 2122  
 key space  
   MyISAM, 1791  
 key-value store, 877  
 keys, 872  
   foreign, 32, 196  
   multi-column, 873  
   searching on two, 198  
 keys option  
   mysqlshow, 305  
 keys-used option  
   myisamchk, 322  
 keywords, 1026  
 KEY\_BLOCK\_SIZE, 657, 661, 3801  
 key\_buffer\_size myisamchk variable, 319  
 key\_buffer\_size system variable, 492  
 key\_cache\_age\_threshold system variable, 493  
 key\_cache\_block\_size system variable, 493  
 key\_cache\_division\_limit system variable, 494  
 KEY\_COLUMN\_USAGE  
   INFORMATION\_SCHEMA table, 2188  
 Kill  
   thread command, 958  
 KILL, 1594  
 Killed  
   thread state, 961  
 Killing slave  
   thread state, 969  
 known errors, 3751  
 Korean, 3640

**L**

labels  
   stored program block, 1491  
 language option  
   mysqld, 425  
 language support  
   error messages, 1085  
 large page support, 945  
 large-pages option  
   mysqld, 425  
 large\_files\_support system variable, 494  
 large\_pages system variable, 494  
 large\_page\_size system variable, 494  
 last row  
   unique ID, 2896  
 LAST\_DAY(), 1219  
 last\_insert\_id session variable, 495  
 LAST\_INSERT\_ID(), 1269, 1406  
   and replication, 2058  
   and stored routines, 2155  
   and triggers, 2155  
   locking considerations, 32  
 latch, 3801  
 layout of installation, 62

lc-messages option  
   mysqld, 426  
 lc-messages-dir option  
   mysqld, 426  
 LCASE(), 1181  
 lc\_messages system variable, 495  
 lc\_messages\_dir system variable, 495  
 lc\_time\_names system variable, 495  
 ldata option  
   mysql\_install\_db, 237  
 LDML syntax, 1097  
 LD\_LIBRARY\_PATH environment variable, 169  
 LD\_RUN\_PATH environment variable, 166, 169  
 LEAST(), 1170  
 LEAVE, 1496  
 ledir option  
   mysqld\_safe, 227  
 LEFT JOIN, 981, 1435  
 LEFT OUTER JOIN, 1435  
 LEFT(), 1181  
 leftmost prefix of indexes, 874  
 legal names, 1017  
 length option  
   myisam\_ftdump, 315  
 LENGTH(), 1181  
 less than (<), 1167  
 less than or equal (<=), 1167  
 libmysqlclient library, 2323  
 libmysqlclient.a, 2731  
 libmysqlcppconn-static.a, 2731  
 libmysqlcppconn.so, 2731  
 libmysqld, 2790  
   options, 2791  
 libmysqld library, 2323  
 libmysqld-libs option  
   mysql\_config, 373  
 LIBMYSQL\_ENABLE\_CLEARTTEXT\_PLUGIN  
   environment variable, 166  
 LIBMYSQL\_PLUGINS environment variable, 166, 2894  
 LIBMYSQL\_PLUGIN\_DIR environment variable, 166,  
 2894  
 library  
   libmysqlclient, 2323  
   libmysqld, 2323  
 libs option  
   mysql\_config, 373  
 libs\_r option  
   mysql\_config, 373  
 license system variable, 496  
 LIKE, 1190  
   and indexes, 876  
   and wildcards, 876  
 LIMIT, 860, 1268, 1431  
   and replication, 2071  
 limitations  
   MySQL Limitations, 3767  
   replication, 2057  
 limits

---

---

- file-size, 3768
- InnoDB, 1756
- MySQL Limits, limits in MySQL, 3767
- line-numbers option
  - mysql, 250
- linear hash partitioning, 2106
- linear key partitioning, 2108
- linefeed (n), 1012, 1417
- LineFromText(), 1298
- LineFromWKB(), 1299
- lines-terminated-by option
  - mysqldump, 290, 300
- LINESTRING data type, 1297
- LineString(), 1300
- LineStringFromText(), 1298
- LineStringFromWKB(), 1299
- linking, 2797
  - errors, 2798
  - problems, 2798
- links
  - symbolic, 940
- LINQ, 2474
- list, 3801
- list partitioning, 2094, 2097
- list partitions
  - adding and dropping, 2116
  - managing, 2116
- literals, 1011
- LN(), 1204
- load balancing
  - with Connector/J, 2670, 2687
- LOAD DATA
  - and replication, 2071
- LOAD DATA INFILE, 1412, 2506, 3746
- load emulation, 306
- LOAD INDEX INTO CACHE
  - and partitioning, 2141
- LOAD XML, 1420
- loading
  - tables, 178
- LOAD\_FILE(), 1181
- local option
  - mysqlimport, 300
- local-infile option
  - mysql, 250
- local-load option
  - mysqlbinlog, 351
- localhost, 210
- localization, 1037
- LOCALTIME, 1219
- LOCALTIMESTAMP, 1219
- local\_infile system variable, 496
- LOCATE(), 1181
- lock, 3801
- lock escalation, 3801
- LOCK IN SHARE MODE, 1432
- lock mode, 3802
- Lock Monitor
  - InnoDB, 1673
- LOCK TABLES, 1462
- lock wait timeout, 1689
- lock-all-tables option
  - mysqldump, 295
- lock-tables option
  - mysqldump, 295
  - mysqlimport, 300
- locked\_in\_memory system variable, 497
- locking, 931, 3802
  - external, 422, 439, 540, 848, 930, 964
  - information schema, 1667
  - internal, 925
  - row-level, 32, 925
  - table-level, 925
- Locking
  - information schema, 1672
- locking methods, 925
- locking read, 3802
- lock\_wait\_timeout system variable, 496
- log, 3802
- log buf mutex, 1662
- log buffer, 3802
- log file, 3802
- log files
  - maintaining, 637
- log group, 3802
- log option
  - mysqld, 426
  - mysqld\_multi, 232
- log sys mutex, 1662
- log system variable, 497
- LOG(), 1204
- log-bin option
  - mysqld, 2002
- log-bin-index option
  - mysqld, 2002
- log-bin-trust-function-creators option
  - mysqld, 2003
- log-bin-use-v1-row-events option
  - mysqld, 2003
- log-error option
  - mysqld, 427
  - mysqldump, 287
  - mysqld\_safe, 227
- log-isam option
  - mysqld, 427
- log-output option
  - mysqld, 427
- log-queries-not-using-indexes option
  - mysqld, 428
- log-raw option
  - mysqld, 428
- log-short-format option
  - mysqld, 428, 2003
- log-slave-updates option
  - mysqld, 1971
- log-slow-admin-statements option

---



---

- mysqld, 429
- log-slow-queries option
  - mysqld, 429
- log-slow-slave-statements option
  - mysqld, 1972
- log-tc option
  - mysqld, 429
- log-tc-size option
  - mysqld, 429
- log-warnings option
  - mysqld, 430, 1972
- LOG10(), 1205
- LOG2(), 1205
- logging
  - passwords, 735
- logging slow query
  - thread state, 962
- logical, 3802
- logical backup, 3803
- logical operators, 1171
- login
  - thread state, 962
- login-path option, 219, 374
  - mysql\_config\_editor, 342
- logs
  - flushing, 619
  - server, 619
- log\_bin system variable, 2008
- log\_bin\_basename system variable, 2015
- log\_bin\_trust\_function\_creators system variable, 497
- log\_bin\_use\_v1\_row\_events system variable, 2015
- log\_error system variable, 497
- log\_output system variable, 498
- log\_queries\_not\_using\_indexes system variable, 498
- log\_slave\_updates system variable, 2008
- log\_slow\_admin\_statements system variable
  - mysqld, 498
- log\_slow\_queries system variable, 499
- log\_slow\_slave\_statements system variable
  - mysqld, 1988
- log\_throttle\_queries\_not\_using\_indexes system variable, 499
- log\_warnings system variable, 499
- Long Data
  - thread command, 958
- LONG data type, 1140
- LONGBLOB data type, 1120
- LONGTEXT data type, 1120
- long\_query\_time system variable, 500
- LOOP, 1496
  - labels, 1491
- loose option prefix, 214
- loose\_, 3803
- lost connection errors, 3726
- lost+found directory, 424
- low-priority option
  - mysqld, 301
- low-priority-updates option

- mysqld, 430
- low-water mark, 3803
- LOWER(), 1182
- lower\_case\_file\_system system variable, 501
- lower\_case\_table\_names system variable, 501
- low\_priority\_updates system variable, 500
- LPAD(), 1182
- LRU, 3803
- LRU page replacement, 1659
- LSN, 3803
- LTRIM(), 1182
- Lua, 1891

## M

- Mac OS X, 2326
  - installation, 100
- mailing lists, 18
  - archive location, 18
  - guidelines, 20
- main features of MySQL, 6
- maintaining
  - log files, 637
  - tables, 853
- maintenance
  - tables, 272
- MAKEDATE(), 1219
- MAKETIME(), 1219
- MAKE\_SET(), 1182
- Making temp file
  - thread state, 968
- manage keys
  - thread state, 962
- manual
  - available formats, 2
  - online location, 2
  - syntax conventions, 3
  - typographical conventions, 3
- Master has sent all binlog to slave; waiting for binlog to be updated
  - thread state, 967
- master server, 3804
- master thread, 3804
- master-data option
  - mysqldump, 288
- master-info-file option
  - mysqld, 1973
- master-info-repository option
  - mysqld, 2000
- master-retry-count option
  - mysqld, 1973
- master-verify-checksum option
  - mysqld, 2006
- master\_info\_repository system variable, 1988
- MASTER\_POS\_WAIT(), 1279, 1481
- master\_verify\_checksum system variable, 2015
- MATCH ... AGAINST(), 1229
- matching
  - patterns, 186

---

math, 1315  
 mathematical functions, 1202  
 MAX(), 1284  
 MAX(DISTINCT), 1284  
 max-binlog-dump-events option  
   mysqld, 2007  
 max-record-length option  
   mysamchk, 322  
 max-relay-log-size option  
   mysqld, 1973  
 MAXDB SQL mode, 608  
 --maximum option prefix, 214  
 maximums  
   maximum columns per table, 3769  
   maximum number of databases, 3768  
   maximum number of tables, 3768  
   maximum row size, 3769  
   maximum tables per join, 3767  
   table size, 3768  
 max\_allowed\_packet  
   and replication, 2072  
 max\_allowed\_packet system variable, 502  
 max\_allowed\_packet variable, 255  
 max\_binlog\_cache\_size system variable, 2016  
 max\_binlog\_size system variable, 2016  
 max\_binlog\_stmt\_cache\_size system variable, 2017  
 max\_connections system variable, 503  
 MAX\_CONNECTIONS\_PER\_HOUR, 778  
 max\_connect\_errors system variable, 502  
 max\_delayed\_threads system variable, 504  
 max\_error\_count system variable, 504  
 max\_heap\_table\_size system variable, 504  
 max\_insert\_delayed\_threads system variable, 505  
 max\_join\_size system variable, 505  
 max\_join\_size variable, 255  
 max\_length\_for\_sort\_data system variable, 506  
 max\_prepared\_stmt\_count system variable, 506  
 MAX\_QUERIES\_PER\_HOUR, 778  
 max\_relay\_log\_size system variable, 507  
 max\_seeks\_for\_key system variable, 507  
 max\_sort\_length system variable, 508  
 max\_sp\_recursion\_depth system variable, 508  
 max\_tmp\_tables system variable, 508  
 MAX\_UPDATES\_PER\_HOUR, 778  
 MAX\_USER\_CONNECTIONS, 778  
 max\_user\_connections system variable, 509  
 max\_write\_lock\_count system variable, 509  
 MBR, 1309  
 MBRContains(), 1309  
 MBRDisjoint(), 1309  
 MBREqual(), 1309  
 MBRIntersects(), 1309  
 MBROverlaps(), 1310  
 MBRTouches(), 1310  
 MBRWithin(), 1310  
 MD5(), 1263  
 MDL, 3804  
 medium trust, 2520  
 medium-check option  
   mysamchk, 321  
   mysqlcheck, 277  
 MEDIUMBLOB data type, 1120  
 MEDIUMINT data type, 1113  
 MEDIUMTEXT data type, 1120  
 memcached, 1761, 3804  
 MEMCACHED\_SASL\_PWDB environment variable, 1769  
 memcapable command, 1762  
 memlock option  
   mysqld, 431  
 memory allocator  
   innodb\_use\_sys\_malloc, 1652  
 MEMORY storage engine, 1601, 1795  
   and replication, 2072  
 memory usage  
   mysamchk, 330  
 memory use, 943  
   Performance Schema, 2243  
 merge, 3804  
 MERGE storage engine, 1601, 1804  
 MERGE tables  
   defined, 1804  
 metadata  
   database, 2178  
   InnoDB, 2205  
   stored routines, 2154  
   triggers, 2158  
   views, 2169  
 metadata lock, 3804  
 metadata locking  
   transactions, 929  
 metadata\_locks\_cache\_size system variable, 510  
 metadata\_locks\_hash\_instances system variable, 510  
 method option  
   mysqlhotcopy, 367  
 methods  
   locking, 925  
 metrics counter, 3804  
 MICROSECOND(), 1220  
 Microsoft Access, 2393  
 Microsoft ADO, 2395, 2436  
 Microsoft Excel, 2395  
 Microsoft Visual Basic, 2395  
 Microsoft Visual InterDev, 2395  
 Microsoft Visual Studio, 2414  
 MID(), 1183  
 midpoint insertion, 1658  
 midpoint insertion strategy, 3804  
 MIN(), 1284  
 MIN(DISTINCT), 1284  
 min-examined-row-limit option  
   mysqld, 431  
 mini-transaction, 3805  
 Minimum Bounding Rectangle, 1309  
 minus  
   unary (-), 1200

---

---

MINUTE(), 1220  
 min\_examined\_row\_limit system variable, 510  
 mirror sites, 53  
 miscellaneous functions, 1274  
 mixed statements (Replication), 2076  
 mixed-mode insert, 3805  
 MLineFromText(), 1298  
 MLineFromWKB(), 1299  
 MOD (modulo), 1205  
 MOD(), 1205  
 modes  
     batch, 192  
 modulo (%), 1205  
 modulo (MOD), 1205  
 monitor  
     terminal, 171  
 monitoring, 3553  
     threads, 956  
 Monitors  
     InnoDB, 639, 1618, 1673, 1683, 1685  
 Mono, 2403  
 MONTH(), 1220  
 MONTHNAME(), 1220  
 MPointFromText(), 1298  
 MPointFromWKB(), 1299  
 MPolyFromText(), 1298  
 MPolyFromWKB(), 1299  
 mSQL compatibility, 1193  
 msql2mysql, 372  
 MSSQL SQL mode, 608  
 multi mysqld, 231  
 multi-byte character sets, 3734  
 multi-byte characters, 1089  
 multi-column indexes, 873  
 multi-core, 3805  
 Multi-Range Read  
     optimization, 991  
 MULTILINESTRING data type, 1297  
 MultiLineString(), 1300  
 MultiLineStringFromText(), 1298  
 MultiLineStringFromWKB(), 1299  
 multiple buffer pools, 1661  
 multiple rollback segments, 1661  
 multiple servers, 705  
 multiple-part index, 1349  
 multiplication (\*), 1200  
 MULTIPOINT data type, 1297  
 MultiPoint(), 1300  
 MultiPointFromText(), 1298  
 MultiPointFromWKB(), 1299  
 MULTIPOLYGON data type, 1297  
 MultiPolygon(), 1300  
 MultiPolygonFromText(), 1298  
 MultiPolygonFromWKB(), 1299  
 mutex, 3805  
 mutex\_instances table  
     performance\_schema, 2271  
 MVCC, 3805  
 MVCC (multi-version concurrency control), 1635  
 My  
     derivation, 8  
 my-print-defaults option  
     mysql\_plugin, 240  
 my.cnf, 3805  
 my.cnf file, 2057  
 my.ini, 3805  
 MyISAM  
     compressed tables, 332, 1793  
     converting tables to InnoDB, 645  
 MyISAM key cache, 914  
 MyISAM storage engine, 1601, 1787  
 myisam-block-size option  
     mysqld, 432  
 myisam-recover-options option  
     mysqld, 432, 1790  
 myisamchk, 206, 315  
     analyze option, 323  
     backup option, 321  
     block-search option, 323  
     character-sets-dir option, 321  
     check option, 320  
     check-only-changed option, 321  
     correct-checksum option, 321  
     data-file-length option, 322  
     debug option, 319  
     description option, 323  
     example output, 324  
     extend-check option, 321, 322  
     fast option, 321  
     force option, 321, 322  
     help option, 318  
     HELP option, 319  
     information option, 321  
     keys-used option, 322  
     max-record-length option, 322  
     medium-check option, 321  
     no-symlinks option, 322  
     options, 318  
     parallel-recover option, 322  
     quick option, 322  
     read-only option, 321  
     recover option, 322  
     safe-recover option, 322  
     set-auto-increment[ option, 323  
     set-character-set option, 323  
     set-collation option, 323  
     silent option, 319  
     sort-index option, 323  
     sort-records option, 324  
     sort-recover option, 323  
     tmpdir option, 323  
     unpack option, 323  
     update-state option, 321  
     verbose option, 319  
     version option, 319  
     wait option, 319

---

---

mysamlog, 206, 331  
mysampack, 206, 332, 1384, 1793  
  backup option, 332  
  character-sets-dir option, 332  
  debug option, 332  
  force option, 332  
  help option, 332  
  join option, 333  
  silent option, 333  
  test option, 333  
  tmpdir option, 333  
  verbose option, 333  
  version option, 333  
  wait option, 333  
mysam\_block\_size myisamchk variable, 319  
mysam\_data\_pointer\_size system variable, 511  
mysam\_ftdump, 206, 314  
  count option, 314  
  dump option, 315  
  help option, 314  
  length option, 315  
  stats option, 315  
  verbose option, 315  
mysam\_max\_sort\_file\_size system variable, 511  
mysam\_mmap\_size system variable, 512  
mysam\_recover\_options system variable, 512  
mysam\_repair\_threads system variable, 512  
mysam\_sort\_buffer\_size myisamchk variable, 319  
mysam\_sort\_buffer\_size system variable, 513  
mysam\_stats\_method system variable, 513  
mysam\_use\_mmap system variable, 514  
MyODBC, 2326  
MySQL  
  defined, 4  
  introduction, 4  
  pronunciation, 5  
  upgrading, 241  
mysql, 205, 244, 3805  
  auto-rehash option, 248  
  auto-vertical-output option, 248  
  batch option, 248  
  binary-mode option, 248  
  bind-address option, 248  
  character-sets-dir option, 248  
  charset command, 256  
  clear command, 256  
  column-names option, 249  
  column-type-info option, 249  
  comments option, 249  
  compress option, 249  
  connect command, 256  
  connect-expired-password option, 249  
  database option, 249  
  debug option, 249  
  debug-check option, 249  
  debug-info option, 249  
  default-auth option, 249  
  default-character-set option, 249  
  delimiter command, 256  
  delimiter option, 249  
  disable named commands, 250  
  edit command, 257  
  ego command, 257  
  enable-clear-text-plugin option, 250  
  execute option, 250  
  exit command, 257  
  force option, 250  
  go command, 257  
  help command, 256  
  help option, 248  
  histignore option, 250  
  host option, 250  
  html option, 250  
  i-am-a-dummy option, 252  
  ignore-spaces option, 250  
  init-command option, 250  
  line-numbers option, 250  
  local-infile option, 250  
  named-commands option, 250  
  no-auto-rehash option, 251  
  no-beep option, 251  
  nopager command, 257  
  notee command, 257  
  nowarning command, 257  
  one-database option, 251  
  pager command, 257  
  pager option, 251  
  password option, 251  
  pipe option, 251  
  plugin-dir option, 252  
  port option, 252  
  print command, 257  
  prompt command, 257  
  prompt option, 252  
  protocol option, 252  
  quick option, 252  
  quit command, 257  
  raw option, 252  
  reconnect option, 252  
  rehash command, 258  
  safe-updates option, 252  
  secure-auth option, 253  
  server-public-key-path option, 253  
  show-warnings option, 253  
  sigint-ignore option, 253  
  silent option, 253  
  skip-column-names option, 253  
  skip-line-numbers option, 253  
  socket option, 253  
  source command, 258  
  SSL options, 253  
  status command, 258  
  system command, 258  
  table option, 254  
  tee command, 258  
  tee option, 254

---

- unbuffered option, 254
- use command, 258
- user option, 254
- verbose option, 254
- version option, 254
- vertical option, 254
- wait option, 254
- warnings command, 258
- xml option, 254
- MySQL binary distribution, 50
- MYSQL C type, 2800
- MySQL Cluster
  - compiling with icc, 3543
  - FAQ, 3640
  - storage requirements, 1148
- mysql command options, 244
- mysql commands
  - list of, 255
- MySQL Dolphin name, 8
- MySQL Enterprise Audit, 817, 3555
- MySQL Enterprise Backup, 3554, 3805
- MySQL Enterprise Monitor, 3553
- MySQL Enterprise Security, 790, 797, 3554
- MySQL history, 8
- mysql history file, 260
- MySQL Installer, 67, 67
- MySQL mailing lists, 18
- MySQL name, 8
- MySQL Notifier, 88
- mysql prompt command, 259
- MySQL server
  - mysqld, 225, 380
- mysql source (command for reading from text files), 193, 262
- MySQL source distribution, 50
- MySQL storage engines, 1601
- MySQL system tables
  - and replication, 2074
- MySQL Thread Pool, 949, 3555
- MySQL version, 53
- mysql \. (command for reading from text files), 193, 262
- MySql.Data class, 2530
- MySql.Data.dll library, 2411, 2412
- mysql.data.ef.dll library, 2414
- mysql.event table, 2164
- mysql.server, 204, 230
  - basedir option, 231
  - datadir option, 231
  - pid-file option, 231
  - service-startup-timeout option, 231
  - use-mysqld\_safe option, 231
  - user option, 231
- mysql.slave\_master\_info table, 2028
- mysql.slave\_relay\_log\_info table, 2028
- mysql.sock
  - protection, 3742
- mysql.visualstudio.dll library, 2414
- MYSQL323 SQL mode, 608
- MYSQL40 SQL mode, 609
- mysqlaccess, 206, 343
  - brief option, 344
  - commit option, 344
  - copy option, 344
  - db option, 344
  - debug option, 344
  - help option, 344
  - host option, 344
  - howto option, 344
  - old\_server option, 344
  - password option, 344
  - plan option, 345
  - preview option, 345
  - relnotes option, 345
  - rhost option, 345
  - rollback option, 345
  - spassword option, 345
  - superuser option, 345
  - table option, 345
  - user option, 345
  - version option, 345
- mysqladmin, 205, 266, 1344, 1392, 1580, 1584, 1590, 1594
  - bind-address option, 270
  - character-sets-dir option, 270
  - compress option, 270
  - count option, 270
  - debug option, 271
  - debug-check option, 271
  - debug-info option, 271
  - default-auth option, 271
  - default-character-set option, 271
  - enable-clear-text-plugin option, 271
  - force option, 271
  - help option, 270
  - host option, 271
  - no-beep option, 271
  - password option, 271
  - pipe option, 271
  - plugin-dir option, 271
  - port option, 271
  - protocol option, 272
  - relative option, 272
  - silent option, 272
  - sleep option, 272
  - socket option, 272
  - SSL options, 272
  - user option, 272
  - verbose option, 272
  - version option, 272
  - vertical option, 272
  - wait option, 272
- mysqladmin command options, 268
- mysqladmin option
  - mysqld\_multi, 232
- mysqlbackup command, 3806
- mysqlbinlog, 206, 345

---

---

base64-output option, 348  
 bind-address option, 349  
 binlog-row-event-max-size option, 349  
 character-sets-dir option, 349  
 database option, 349  
 debug option, 350  
 debug-check option, 350  
 debug-info option, 350  
 default-auth option, 351  
 disable-log-bin option, 351  
 exclude-gtids option, 351  
 force-if-open option, 351  
 force-read option, 351  
 help option, 348  
 hexdump option, 351  
 host option, 351  
 include-gtids option, 351  
 local-load option, 351  
 offset option, 351  
 password option, 351  
 plugin-dir option, 352  
 port option, 352  
 protocol option, 352  
 raw option, 352  
 read-from-remote-master option, 352  
 read-from-remote-server option, 352  
 result-file option, 353  
 server-id option, 353  
 set-charset option, 353  
 short-form option, 353  
 skip-gtids option, 353  
 socket option, 353  
 start-datetime option, 353  
 start-position option, 353  
 stop-datetime option, 353  
 stop-never option, 354  
 stop-never-slave-server-id option, 354  
 stop-position option, 354  
 to-last-log option, 354  
 user option, 354  
 verbose option, 354  
 verify-binlog-checksum option, 354  
 version option, 354  
 mysqlbug, 236  
 MySQLBulkLoader class, 2506  
 mysqlcheck, 205, 272
 

- all-databases option, 276
- all-in-1 option, 276
- analyze option, 276
- auto-repair option, 276
- bind-address option, 276
- character-sets-dir option, 276
- check option, 276
- check-only-changed option, 276
- check-upgrade option, 276
- compress option, 276
- databases option, 276
- debug option, 277
  - debug-check option, 277
  - debug-info option, 277
  - default-auth option, 277
  - default-character-set option, 277
  - extended option, 277
  - fast option, 277
  - fix-db-names option, 277
  - fix-table-names option, 277
  - force option, 277
  - help option, 276
  - host option, 277
  - medium-check option, 277
  - optimize option, 277
  - password option, 277
  - pipe option, 278
  - plugin-dir option, 278
  - port option, 278
  - protocol option, 278
  - quick option, 278
  - repair option, 278
  - silent option, 278
  - skip-database option, 278
  - socket option, 278
  - SSL options, 278
  - tables option, 278
  - use\_frm option, 279
  - user option, 279
  - verbose option, 279
  - version option, 279
  - write-binlog option, 279

- MySQLClient namespace, 2530
- MySQLCommand class, 2445
- MySQLCommand class, 2489
- mysqlcppconn-static.lib, 2723
- mysqlcppconn.dll, 2723
- mysqld, 204, 3806
- abort-slave-event-count option, 1971
- allow-suspicious-udfs option, 413
- ansi option, 413
- audit-log option, 824
- basedir option, 414
- big-tables option, 414
- bind-address option, 414
- binlog-checksum option, 2006
- binlog-do-db option, 2004
- binlog-format option, 415
- binlog-ignore-db option, 2005
- binlog-row-event-max-size option, 2001
- binlog-rows-query-log-events option, 2007
- bootstrap option, 416
- character-set-client-handshake option, 416
- character-set-filesystem option, 416
- character-set-server option, 417
- character-sets-dir option, 416
- chroot option, 417
- collation-server option, 417
- command options, 412
- console option, 417

---

core-file option, 418  
datadir option, 418  
debug option, 418  
debug-sync-timeout option, 419  
default-authentication-plugin option, 419  
default-storage-engine option, 420  
default-time-zone option, 420  
delay-key-write option, 421, 1790  
des-key-file option, 421  
disconnect-slave-event-count option, 1971  
enable-named-pipe option, 421  
event-scheduler option, 422  
exit-info option, 422  
external-locking option, 422  
flush option, 423  
gdb option, 423  
general-log option, 423  
help option, 413  
ignore-builtin-innodb option, 1697  
ignore-db-dir option, 424  
init-file option, 424  
innodb option, 1698  
innodb-status-file option, 1698  
install option, 424  
install-manual option, 424  
language option, 425  
large-pages option, 425  
lc-messages option, 426  
lc-messages-dir option, 426  
log option, 426  
log-bin option, 2002  
log-bin-index option, 2002  
log-bin-trust-function-creators option, 2003  
log-bin-use-v1-row-events option, 2003  
log-error option, 427  
log-isam option, 427  
log-output option, 427  
log-queries-not-using-indexes option, 428  
log-raw option, 428  
log-short-format option, 428, 2003  
log-slave-updates option, 1971  
log-slow-admin-statements option, 429  
log-slow-queries option, 429  
log-slow-slave-statements option, 1972  
log-tc option, 429  
log-tc-size option, 429  
log-warnings option, 430, 1972  
log\_slow\_admin\_statements system variable, 498  
log\_slow\_slave\_statements system variable, 1988  
low-priority-updates option, 430  
master-info-file option, 1973  
master-info-repository option, 2000  
master-retry-count option, 1973  
master-verify-checksum option, 2006  
max-binlog-dump-events option, 2007  
max-relay-log-size option, 1973  
memlock option, 431  
min-examined-row-limit option, 431  
myisam-block-size option, 432  
myisam-recover-options option, 432, 1790  
MySQL server, 225, 380  
old-alter-table option, 433  
old-style-user-limits option, 433  
one-thread option, 434  
open-files-limit option, 434  
partition option, 434  
performance-schema-consumer-xxx option, 2306  
performance-schema-instrument option, 2306  
pid-file option, 435  
plugin option prefix, 435  
plugin-load option, 435  
plugin-load-add option, 436  
port option, 437  
port-open-timeout option, 437  
read-only option, 1973  
relay-log option, 1974  
relay-log-index option, 1974  
relay-log-info-file option, 1975  
relay-log-info-repository option, 2000  
relay-log-purge option, 1975  
relay-log-recovery option, 1975  
relay-log-space-limit option, 1976  
remove option, 437  
replicate-do-db option, 1976  
replicate-do-table option, 1978  
replicate-ignore-db option, 1977  
replicate-ignore-table option, 1978  
replicate-rewrite-db option, 1979  
replicate-same-server-id option, 1979  
replicate-wild-do-table option, 1979  
replicate-wild-ignore-table option, 1980  
report-host option, 1980  
report-password option, 1981  
report-port option, 1981  
report-user option, 1982  
safe-mode option, 437  
safe-user-create option, 438  
secure-auth option, 438  
secure-file-priv option, 438  
server-id option, 1955  
server\_uuid variable, 1955  
shared-memory option, 439  
shared-memory-base-name option, 439  
show-slave-auth-info option, 1982  
skip-concurrent-insert option, 439  
skip-event-scheduler option, 439  
skip-external-locking option, 439  
skip-grant-tables option, 439  
skip-host-cache option, 439  
skip-innodb option, 440, 1698  
skip-name-resolve option, 440  
skip-networking option, 440  
skip-partition option, 440  
skip-show-database option, 441  
skip-slave-start option, 1984  
skip-stack-trace option, 441

---

skip-symbolic-links option, 441  
 skip-thread-priority option, 441  
 slave-checkpoint-group option, 1982  
 slave-checkpoint-period option, 1983  
 slave-load-tmpdir option, 1985  
 slave-max-allowed-packet, 1964  
 slave-net-timeout option, 1985  
 slave-parallel-workers option, 1983  
 slave-pending-jobs-size-max option, 1984  
 slave-rows-search-algorithms, 1964  
 slave-skip-errors option, 1986  
 slave-sql-verify-checksum option, 1987  
 slave\_compressed\_protocol option, 1985  
 slow-query-log option, 441  
 slow-start-timeout option, 442  
 socket option, 442  
 sporadic-binlog-dump-fail option, 2007  
 sql-mode option, 442  
 SSL options, 440  
 standalone option, 440  
 starting, 748  
 super-large-pages option, 440  
 symbolic-links option, 441  
 sysdate-is-now option, 444  
 tc-heuristic-recover option, 444  
 temp-pool option, 445  
 tmpdir option, 446  
 transaction-isolation option, 445  
 transaction-read-only option, 445  
 user option, 446  
 validate-password option, 742  
 verbose option, 447  
 version option, 447  
 mysqld option  
   malloc-lib, 227  
   mysqld\_multi, 232  
   mysqld\_safe, 228  
   mysql\_plugin, 240  
 mysqld options, 932  
   disable-gtid-unsafe-statements, 2018  
   enforce-gtid-consistency, 2018  
   gtid-mode, 2019  
 mysqld server  
   buffer sizes, 932  
 mysqld-version option  
   mysqld\_safe, 228  
 mysqldump, 165, 205, 279, 3806  
   add-drop-database option, 286  
   add-drop-table option, 286  
   add-drop-trigger option, 286  
   add-locks option, 294  
   all-databases option, 292  
   all-tablespaces option, 286  
   allow-keywords option, 286  
   apply-slave-statements option, 288  
   bind-address option, 284  
   character-sets-dir option, 287  
   comments option, 286  
   compact option, 290  
   compatible option, 290  
   complete-insert option, 290  
   compress option, 284  
   create-options option, 290  
   databases option, 292  
   debug option, 286  
   debug-check option, 287  
   debug-info option, 287  
   default-auth option, 284  
   default-character-set option, 288  
   delayed-insert option, 294  
   delete-master-logs option, 288  
   disable-keys option, 294  
   dump-date option, 287  
   dump-slave option, 288  
   events option, 293  
   extended-insert option, 294  
   fields-enclosed-by option, 290, 300  
   fields-escaped-by option, 290, 300  
   fields-optionally-enclosed-by option, 290, 300  
   fields-terminated-by option, 290, 300  
   flush-logs option, 294  
   flush-privileges option, 295  
   force option, 287  
   help option, 287  
   hex-blob option, 290  
   host option, 284  
   ignore-table option, 293  
   include-master-host-port option, 288  
   insert-ignore option, 294  
   lines-terminated-by option, 290, 300  
   lock-all-tables option, 295  
   lock-tables option, 295  
   log-error option, 287  
   master-data option, 288  
   no-autocommit option, 295  
   no-create-db option, 286  
   no-create-info option, 286  
   no-data option, 293  
   no-set-names option, 288  
   no-tablespaces option, 286  
   opt option, 294  
   order-by-primary option, 295  
   password option, 285  
   pipe option, 285  
   plugin-dir option, 285  
   port option, 285  
   problems, 297, 3764  
   protocol option, 285  
   quick option, 294  
   quote-names option, 291  
   replace option, 286  
   result-file option, 291  
   routines option, 293  
   set-charset option, 288  
   set-gtid-purged option, 289  
   single-transaction option, 295

---



---

skip-comments option, 287  
 skip-opt option, 294  
 socket option, 285  
 SSL options, 285  
 tab option, 291  
 tables option, 293  
 triggers option, 293  
 tz-utc option, 291  
 user option, 285  
 using for backups, 840  
 verbose option, 287  
 version option, 287  
 views, 297, 3764  
 where option, 293  
 workarounds, 297, 3764  
 xml option, 291

**mysqldumpslow**, 206, 364  
 debug option, 364  
 help option, 364  
 verbose option, 365

**mysqld\_multi**, 204, 231  
 defaults-extra-file option, 232  
 defaults-file option, 232  
 example option, 232  
 help option, 232  
 log option, 232  
 mysqladmin option, 232  
 mysqld option, 232  
 no-defaults option, 232  
 no-log option, 232  
 password option, 233  
 silent option, 233  
 tcp-ip option, 233  
 user option, 233  
 verbose option, 233  
 version option, 233

**mysqld\_safe**, 204, 225  
 basedir option, 227  
 core-file-size option, 227  
 datadir option, 227  
 defaults-extra-file option, 227  
 defaults-file option, 227  
 help option, 227  
 ledir option, 227  
 log-error option, 227  
 malloc-lib option, 227  
 mysqld option, 228  
 mysqld-version option, 228  
 nice option, 228  
 no-defaults option, 228  
 open-files-limit option, 228  
 pid-file option, 228  
 plugin-dir option, 228  
 port option, 228  
 skip-kill-mysqld option, 229  
 skip-syslog option, 229  
 socket option, 229  
 syslog option, 229

syslog-tag option, 229  
 timezone option, 229  
 user option, 229

**mysqlhotcopy**, 206, 365  
 addtoexec option, 367  
 allowold option, 367  
 checkpoint option, 367  
 chroot option, 367  
 debug option, 367  
 dryrun option, 367  
 flushlog option, 367  
 help option, 367  
 host option, 367  
 keepold option, 367  
 method option, 367  
 noindices option, 367  
 old\_server option, 368  
 password option, 367  
 port option, 368  
 quiet option, 368  
 record\_log\_pos option, 368  
 regexp option, 368  
 resetmaster option, 368  
 resetslave option, 368  
 socket option, 368  
 suffix option, 368  
 tmpdir option, 368  
 user option, 368

**mysqlimport**, 165, 205, 297, 1412  
 bind-address option, 299  
 character-sets-dir option, 299  
 columns option, 299  
 compress option, 299  
 debug option, 300  
 debug-check option, 300  
 debug-info option, 300  
 default-auth option, 300  
 default-character-set option, 300  
 delete option, 300  
 force option, 300  
 help option, 299  
 host option, 300  
 ignore option, 300  
 ignore-lines option, 300  
 local option, 300  
 lock-tables option, 300  
 low-priority option, 301  
 password option, 301  
 pipe option, 301  
 plugin-dir option, 301  
 port option, 301  
 protocol option, 301  
 replace option, 301  
 silent option, 301  
 socket option, 301  
 SSL options, 301  
 use-threads option, 302  
 user option, 302

---

---

verbose option, 302  
 version option, 302  
**MySQLInstallerConsole**, 83  
   action option, 84  
   catalog option, 84  
   config option, 84  
   help option, 83  
   install action option, 84  
   list action option, 84  
   nowait option, 84  
   product option, 84  
   remove action option, 84  
   status action option, 84  
   type option, 84  
   updates option, 83  
   upgrade action option, 84  
   user option, 84  
**mysqlshow**, 205, 302  
   bind-address option, 304  
   character-sets-dir option, 304  
   compress option, 304  
   count option, 304  
   debug option, 304  
   debug-check option, 304  
   debug-info option, 304  
   default-auth option, 305  
   default-character-set option, 305  
   help option, 304  
   host option, 305  
   keys option, 305  
   password option, 305  
   pipe option, 305  
   plugin-dir option, 305  
   port option, 305  
   protocol option, 305  
   show-table-type option, 305  
   socket option, 305  
   SSL options, 306  
   status option, 306  
   user option, 306  
   verbose option, 306  
   version option, 306  
**mysqlslap**, 206, 306  
   auto-generate-sql option, 309  
   auto-generate-sql-add-autoincrement option, 309  
   auto-generate-sql-execute-number option, 309  
   auto-generate-sql-guid-primary option, 309  
   auto-generate-sql-load-type option, 309  
   auto-generate-sql-secondary-indexes option, 310  
   auto-generate-sql-unique-query-number option, 310  
   auto-generate-sql-unique-write-number option, 310  
   auto-generate-sql-write-number option, 310  
   commit option, 310  
   compress option, 310  
   concurrency option, 310  
   create option, 310  
   create-schema option, 310  
   csv option, 310  
   debug option, 310  
   debug-check option, 310  
   debug-info option, 311  
   default-auth option, 311  
   delimiter option, 311  
   detach option, 311  
   enable-clear-text-plugin option, 311  
   engine option, 311  
   help option, 309  
   host option, 311  
   iterations option, 311  
   no-drop option, 311  
   number-char-cols option, 311  
   number-int-cols option, 311  
   number-of-queries option, 311  
   only-print option, 311  
   password option, 311  
   pipe option, 312  
   plugin-dir option, 312  
   port option, 312  
   post-query option, 312  
   post-system option, 312  
   pre-query option, 312  
   pre-system option, 312  
   protocol option, 312  
   query option, 312  
   shared-memory-base-name option, 312  
   silent option, 312  
   socket option, 313  
   SSL options, 313  
   user option, 313  
   verbose option, 313  
   version option, 313  
**mysqltest**  
   MySQL Test Suite, 3484  
**mysql\_affected\_rows()**, 2809  
**mysql\_autocommit()**, 2810  
**MYSQL\_BIND** C type, 2860  
**mysql\_change\_user()**, 2810  
**mysql\_character\_set\_name()**, 2811  
**mysql\_client\_find\_plugin()**, 2893  
**mysql\_client\_register\_plugin()**, 2893  
**mysql\_close()**, 2812  
**mysql\_commit()**, 2812  
**mysql\_config**, 373  
   cflags option, 373  
   cxxflags option, 373  
   embedded option, 373  
   include option, 373  
   libmysqld-libs option, 373  
   libs option, 373  
   libs\_r option, 373  
   plugindir option, 373  
   port option, 373  
   socket option, 373  
   version option, 373  
**mysql\_config\_editor**, 206, 337  
   all option, 342

---

---

debug option, 342  
 help option, 342  
 host option, 342  
 login-path option, 342  
 password option, 342  
 port option, 342  
 socket option, 343  
 user option, 343  
 verbose option, 343  
 version option, 343  
 warn option, 343  
 mysql\_connect(), 2812  
 mysql\_convert\_table\_format, 206, 368  
   force option, 369  
   help option, 369  
   host option, 369  
   password option, 369  
   port option, 369  
   socket option, 369  
   type option, 369  
   user option, 369  
   verbose option, 369  
   version option, 369  
 mysql\_create\_db(), 2813  
 MYSQL\_DATADIR option  
   CMake, 133  
 mysql\_data\_seek(), 2813  
 MYSQL\_DEBUG environment variable, 166, 208, 3550  
 mysql\_debug(), 2814  
 mysql\_drop\_db(), 2814  
 mysql\_dump\_debug\_info(), 2815  
 mysql\_eof(), 2815  
 mysql\_errno(), 2816  
 mysql\_error(), 2817  
 mysql\_escape\_string(), 2817  
 mysql\_fetch\_field(), 2817  
 mysql\_fetch\_fields(), 2818  
 mysql\_fetch\_field\_direct(), 2818  
 mysql\_fetch\_lengths(), 2819  
 mysql\_fetch\_row(), 2820  
 MYSQL\_FIELD C type, 2800  
 mysql\_field\_count(), 2820, 2834  
 MYSQL\_FIELD\_OFFSET C type, 2800  
 mysql\_field\_seek(), 2821  
 mysql\_field\_tell(), 2822  
 mysql\_find\_rows, 207, 369  
   help option, 370  
   regexp option, 370  
   rows option, 370  
   skip-use-db option, 370  
   start\_row option, 370  
 mysql\_fix\_extensions, 207, 370  
 mysql\_free\_result(), 2822  
 mysql\_get\_character\_set\_info(), 2822  
 mysql\_get\_client\_info(), 2823  
 mysql\_get\_client\_version(), 2823  
 mysql\_get\_host\_info(), 2823  
 mysql\_get\_proto\_info(), 2823  
 mysql\_get\_server\_info(), 2824  
 mysql\_get\_server\_version(), 2824  
 mysql\_get\_ssl\_cipher(), 2824  
 MYSQL\_GROUP\_SUFFIX environment variable, 166  
 mysql\_hex\_string(), 2824  
 MYSQL\_HISTFILE environment variable, 166, 260  
 MYSQL\_HISTIGNORE environment variable, 166, 260  
 MYSQL\_HOME environment variable, 166  
 MYSQL\_HOST environment variable, 166, 212  
 mysql\_info(), 1337, 1406, 1420, 1457, 2825  
 mysql\_init(), 2826  
 mysql\_insert\_id(), 32, 1406, 2826  
 mysql\_install\_db, 144, 204, 236  
   basedir option, 237  
   datadir option, 237  
   force option, 237  
   ldata option, 237  
   random-passwords option, 237  
   rpm option, 238  
   skip-name-resolve option, 238  
   srcdir option, 238  
   user option, 238  
   verbose option, 238  
   windows option, 238  
 mysql\_kill(), 2828  
 mysql\_library\_end(), 2828  
 mysql\_library\_init(), 2829  
 mysql\_list\_dbs(), 2830  
 mysql\_list\_fields(), 2831  
 mysql\_list\_processes(), 2831  
 mysql\_list\_tables(), 2832  
 mysql\_load\_plugin(), 2894  
 mysql\_load\_plugin\_v(), 2895  
 MYSQL\_MAINTAINER\_MODE option  
   CMake, 135  
 mysql\_more\_results(), 2832  
 mysql\_next\_result(), 2833  
 mysql\_num\_fields(), 2834  
 mysql\_num\_rows(), 2835  
 mysql\_options(), 2835  
 mysql\_options4(), 2840  
 mysql\_ping(), 2841  
 mysql\_plugin, 205, 238  
   basedir option, 240  
   datadir option, 240  
   help option, 240  
   my-print-defaults option, 240  
   mysqld option, 240  
   no-defaults option, 240  
   plugin-dir option, 240  
   plugin-ini option, 240  
   print-defaults option, 240  
   verbose option, 240  
   version option, 240  
 mysql\_plugin\_options(), 2895  
 MYSQL\_PROJECT\_NAME option  
   CMake, 135  
 MYSQL\_PS1 environment variable, 166

---

---

MYSQL\_PWD environment variable, 166, 208, 212  
 mysql\_query(), 2842, 2895  
 mysql\_real\_connect(), 2843  
 mysql\_real\_escape\_string(), 1013, 1183, 2846  
 mysql\_real\_query(), 2847  
 mysql\_refresh(), 2848  
 mysql\_reload(), 2849  
 MYSQL\_RES C type, 2800  
 mysql\_rollback(), 2850  
 MYSQL\_ROW C type, 2800  
 mysql\_row\_seek(), 2850  
 mysql\_row\_tell(), 2850  
 mysql\_secure\_installation, 205, 240  
 mysql\_select\_db(), 2851  
 MYSQL\_SERVER\_AUTH\_INFO plugin structure, 3522  
 mysql\_server\_end(), 2892  
 mysql\_server\_init(), 2892  
 mysql\_setpermission, 207, 370  
     help option, 370  
     host option, 371  
     password option, 371  
     port option, 371  
     socket option, 371  
     user option, 371  
 mysql\_set\_character\_set(), 2851  
 mysql\_set\_local\_infile\_default(), 2852, 2852  
 mysql\_set\_server\_option(), 2853  
 mysql\_shutdown(), 2854  
 mysql\_sqlstate(), 2854  
 mysql\_ssl\_set(), 2855  
 mysql\_stat(), 2856  
 MYSQL\_STMT C type, 2860  
 mysql\_stmt\_affected\_rows(), 2868  
 mysql\_stmt\_attr\_get(), 2868  
 mysql\_stmt\_attr\_set(), 2869  
 mysql\_stmt\_bind\_param(), 2870  
 mysql\_stmt\_bind\_result(), 2870  
 mysql\_stmt\_close(), 2871  
 mysql\_stmt\_data\_seek(), 2872  
 mysql\_stmt\_errno(), 2872  
 mysql\_stmt\_error(), 2872  
 mysql\_stmt\_execute(), 2873  
 mysql\_stmt\_fetch(), 2876  
 mysql\_stmt\_fetch\_column(), 2881  
 mysql\_stmt\_field\_count(), 2881  
 mysql\_stmt\_free\_result(), 2881  
 mysql\_stmt\_init(), 2882  
 mysql\_stmt\_insert\_id(), 2882  
 mysql\_stmt\_next\_result(), 2882  
 mysql\_stmt\_num\_rows(), 2883  
 mysql\_stmt\_param\_count(), 2884  
 mysql\_stmt\_param\_metadata(), 2884  
 mysql\_stmt\_prepare(), 2884  
 mysql\_stmt\_reset(), 2885  
 mysql\_stmt\_result\_metadata, 2886  
 mysql\_stmt\_row\_seek(), 2887  
 mysql\_stmt\_row\_tell(), 2887  
 mysql\_stmt\_send\_long\_data(), 2887  
 mysql\_stmt\_sqlstate(), 2889  
 mysql\_stmt\_store\_result(), 2889  
 mysql\_store\_result(), 2856, 2895  
 MYSQL\_TCP\_PORT environment variable, 166, 208, 710, 711  
 MYSQL\_TCP\_PORT option  
     CMake, 135  
 MYSQL\_TEST\_LOGIN\_FILE environment variable, 166  
 mysql\_thread\_end(), 2891  
 mysql\_thread\_id(), 2857  
 mysql\_thread\_init(), 2891  
 mysql\_thread\_safe(), 2891  
 MYSQL\_TIME C type, 2862  
 mysql\_tzinfo\_to\_sql, 205, 241  
 MYSQL\_UNIX\_ADDR option  
     CMake, 135  
 MYSQL\_UNIX\_PORT environment variable, 145, 166, 208, 710, 711  
 mysql\_upgrade, 205, 241, 769  
     basedir option, 243  
     datadir option, 243  
     debug-check option, 243  
     debug-info option, 243  
     default-auth option, 243  
     force option, 243  
     help option, 243  
     mysql\_upgrade\_info file, 242  
     plugin-dir option, 243  
     tmpdir option, 243  
     upgrade-system-tables option, 243  
     user option, 243  
     verbose option, 243  
     version-check option, 244  
     write-binlog option, 244  
 mysql\_upgrade\_info file  
     mysql\_upgrade, 242  
 mysql\_use\_result(), 2858  
 mysql\_waitpid, 207, 371  
     help option, 371  
     verbose option, 371  
     version option, 371  
 mysql\_warning\_count(), 2859  
 mysql\_zap, 207, 371  
 my\_bool C type, 2801  
 my\_bool values  
     printing, 2801  
 my\_init(), 2890  
 my\_print\_defaults, 207, 374  
     config-file option, 374  
     debug option, 374  
     defaults-extra-file option, 374  
     defaults-file option, 374  
     defaults-group-suffix option, 374  
     extra-file option, 374  
     help option, 374  
     no-defaults option, 374  
     verbose option, 375

---

---

- version option, 375
- my\_ulonglong C type, 2800
- my\_ulonglong values
  - printing, 2800

**N**

- named pipes, 86, 95
- named-commands option
  - mysql, 250
- named\_pipe system variable, 514
- names, 1017
  - case sensitivity, 1020
  - variables, 1029
- NAME\_CONST(), 1279, 2175
- name\_file option
  - comp\_err, 235
- naming
  - releases of MySQL, 51
- NATIONAL CHAR data type, 1118
- NATIONAL VARCHAR data type, 1119
- native functions
  - adding, 3542
- native thread support, 49
- natural key, 3806
- NATURAL LEFT JOIN, 1435
- NATURAL LEFT OUTER JOIN, 1435
- NATURAL RIGHT JOIN, 1435
- NATURAL RIGHT OUTER JOIN, 1435
- NCHAR data type, 1118
- ndb option
  - perror, 376
- NDB storage engine
  - FAQ, 3640
- ndb\_size.pl script, 1148
- negative values, 1013
- neighbor page, 3806
- nested queries, 1445
- Nested-Loop join algorithm, 982
- nested-loop join algorithm, 986
- net etiquette, 20
- NetBeans, 2731
- netmask notation
  - in account names, 763
- net\_buffer\_length system variable, 515
- net\_buffer\_length variable, 255
- net\_read\_timeout system variable, 515
- net\_retry\_count system variable, 515
- net\_write\_timeout system variable, 516
- new features in MySQL 5.6, 9
- new system variable, 516
- new users
  - adding, 144
- newline (`\n`), 1012, 1417
- next-key lock, 3806
  - InnoDB, 1623, 1628, 1629, 1730
- NFS
  - InnoDB, 1610, 1756
- nice option
  - mysqld\_safe, 228
- no matching rows, 3747
- no-auto-rehash option
  - mysql, 251
- no-autocommit option
  - mysqldump, 295
- no-beep option
  - mysql, 251
  - mysqladmin, 271
- no-create-db option
  - mysqldump, 286
- no-create-info option
  - mysqldump, 286
- no-data option
  - mysqldump, 293
- no-defaults option, 220
  - mysqld\_multi, 232
  - mysqld\_safe, 228
  - mysql\_plugin, 240
  - my\_print\_defaults, 374
- no-drop option
  - mysqlslap, 311
- no-log option
  - mysqld\_multi, 232
- no-set-names option
  - mysqldump, 288
- no-symlinks option
  - myisamchk, 322
- no-tablespaces option
  - mysqldump, 286
- noindices option
  - mysqlhotcopy, 367
- non-blocking I/O, 3806
- non-locking read, 3806
- non-repeatable read, 3807
- nondelimited strings, 1015
- Nontransactional tables, 3747
- nopager command
  - mysql, 257
- normalized, 3807
- NoSQL, 3807
- NOT
  - logical, 1171
- NOT BETWEEN, 1169
- not equal (`!=`), 1167
- not equal (`<>`), 1167
- NOT EXISTS
  - with subqueries, 1449
- NOT IN, 1170
- NOT LIKE, 1192
- NOT NULL
  - constraint, 36
- NOT NULL constraint, 3807
- NOT REGEXP, 1193
- notee command
  - mysql, 257
- Notifier, 88
- NOW(), 1220

---

---

nowait option  
   MySQLInstallerConsole, 84  
 nowarning command  
   mysql, 257  
 NO\_AUTO\_CREATE\_USER SQL mode, 604  
 NO\_AUTO\_VALUE\_ON\_ZERO SQL mode, 604  
 NO\_BACKSLASH\_ESCAPES SQL mode, 604  
 NO\_DIR\_IN\_CREATE SQL mode, 604  
 NO\_ENGINE\_SUBSTITUTION SQL mode, 605  
 NO\_FIELD\_OPTIONS SQL mode, 605  
 NO\_KEY\_OPTIONS SQL mode, 605  
 NO\_TABLE\_OPTIONS SQL mode, 605  
 NO\_UNSIGNED\_SUBTRACTION SQL mode, 605  
 NO\_ZERO\_DATE SQL mode, 606  
 NO\_ZERO\_IN\_DATE SQL mode, 606  
 NUL, 1012, 1417  
 NULL, 185, 3745, 3807  
   ORDER BY, 996, 1430  
   testing for null, 1167, 1168, 1168, 1169, 1175  
   thread state, 962  
 NULL value, 185, 1017  
 NULL values  
   and AUTO\_INCREMENT columns, 3746  
   and indexes, 1363  
   and TIMESTAMP columns, 3746  
   vs. empty values, 3745  
 NULLIF(), 1175  
 number-char-cols option  
   mysqslap, 311  
 number-int-cols option  
   mysqslap, 311  
 number-of-queries option  
   mysqslap, 311  
 numbers, 1013  
 NUMERIC data type, 1114  
 numeric precision, 1111  
 numeric scale, 1111  
 numeric types, 1148  
 numeric-dump-file option  
   resolve\_stack\_dump, 375  
 NumGeometries(), 1308  
 NumInteriorRings(), 1307  
 NumPoints(), 1305  
 NVARCHAR data type, 1119

**O**

objects\_summary\_global\_by\_type table  
   performance\_schema, 2291  
 obtaining information about partitions, 2128  
 OCT(), 1183  
 OCTET\_LENGTH(), 1183  
 ODBC, 2326  
 ODBC compatibility, 543, 1019, 1115, 1162, 1168, 1362, 1437  
 ODBC.NET, 2382  
 ODBC\_INCLUDES= option  
   CMake, 133  
 ODBC\_LIB\_DIR option  
   CMake, 133  
 off-page column, 3808  
 offset option  
   mysqlbinlog, 351  
 OLAP, 1285  
 old system variable, 517  
 old-alter-table option  
   mysqld, 433  
 old-style-user-limits option  
   mysqld, 433  
 old\_alter\_table system variable, 517  
 OLD\_PASSWORD(), 1263  
 old\_passwords system variable, 517  
 old\_server option  
   mysqlaccess, 344  
   mysqlhotcopy, 368  
 OLTP, 3808  
 ON DUPLICATE KEY UPDATE, 1403  
 one-database option  
   mysql, 251  
 one-thread option  
   mysqld, 434  
 online, 3808  
 online DDL, 674, 3808  
   concurrency, 679  
   crash recovery, 703  
   examples, 682  
   limitations, 704  
 online location of manual, 2  
 only-print option  
   mysqslap, 311  
 ONLY\_FULL\_GROUP\_BY  
   SQL mode, 1288  
 ONLY\_FULL\_GROUP\_BY SQL mode, 606  
 OPEN, 1499  
 Open Source  
   defined, 5  
 open tables, 268, 881  
 open-files-limit option  
   mysqld, 434  
   mysqld\_safe, 228  
 OpenGIS, 1289  
 opening  
   tables, 881  
 Opening master dump table  
   thread state, 969  
 Opening table  
   thread state, 962  
 Opening tables  
   thread state, 962  
 opens, 268  
 OpenSSL, 806, 807  
 open\_files\_limit system variable, 518  
 open\_files\_limit variable, 354  
 operating systems  
   file-size limits, 3768  
   supported, 49  
 operations

---

---

- arithmetic, 1200
- operators, 1153
  - assignment, 1029, 1172
  - cast, 1199, 1244
  - logical, 1171
  - precedence, 1165
- opt option
  - mysqldump, 294
- optimistic, 3809
- optimization, 856
  - Batched Key Access, 993, 994
  - benchmarking, 954
  - BLOB types, 880
  - Block Nested-Loop, 993, 994
  - character and string types, 879
  - data size, 878
  - DELETE statements, 863
  - disk I/O, 939
  - DML statements, 862
  - foreign keys, 872
  - full table scans, 862
  - indexes, 871
  - InnoDB tables, 883
  - INSERT statements, 862
  - internal details, 969
  - many tables, 881
  - MEMORY tables, 893
  - memory usage, 943
  - Multi-Range Read, 991
  - MyISAM tables, 889
  - network usage, 946
  - numeric types, 879
  - PERFORMANCE\_SCHEMA, 956
  - primary keys, 872
  - SQL statements, 858
  - subquery, 1001, 1005
  - subquery materialization, 1003
  - tips, 868
  - UPDATE statements, 863
- optimizations, 859, 973
  - LIMIT clause, 860
- optimize option
  - mysqlcheck, 277
- OPTIMIZE TABLE, 1540
  - and partitioning, 2127
- optimizer, 3809
  - and replication, 2074
  - controlling, 908
  - query plan evaluation, 908
  - switchable optimizations, 909
- optimizer statistics
  - for InnoDB tables, 1650
- Optimizer Statistics Estimation, 1689, 1690
- optimizer\_join\_cache\_level system variable, 519
- optimizer\_prune\_level system variable, 519
- optimizer\_search\_depth system variable, 520
- optimizer\_switch system variable, 520, 909
- OPTIMIZER\_TRACE
  - INFORMATION\_SCHEMA table, 2189
- OPTIMIZER\_TRACE option
  - CMake, 133
- optimizer\_trace system variable, 523
- optimizer\_trace\_features system variable, 523
- optimizer\_trace\_limit system variable, 524
- optimizer\_trace\_max\_mem\_size system variable, 524
- optimizer\_trace\_offset system variable, 524
- optimizing
  - DISTINCT, 1000
  - filesort, 997
  - GROUP BY, 998
  - LEFT JOIN, 981
  - tables, 852
  - thread state, 962
- option, 3809
- option file, 3809
- option files, 215, 769
  - escape sequences, 217
- option prefix
  - disable, 214
  - enable, 214
  - loose, 214
  - maximum, 214
  - skip, 214
- options
  - boolean, 214
  - CMake, 129
  - command-line
    - mysql, 244
    - mysqladmin, 268
  - embedded server, 2791
  - libmysqld, 2791
  - myisamchk, 318
  - provided by MySQL, 171
  - replication, 2057
- OR, 197, 973
  - bitwise, 1258
  - logical, 1172
- OR Index Merge optimization, 973
- Oracle compatibility, 28, 1285, 1334, 1597
- ORACLE SQL mode, 609
- ORD(), 1183
- ORDER BY, 182, 1335, 1429
  - NULL, 996, 1430
- order-by-primary option
  - mysqldump, 295
- Out of resources error
  - and partitioned tables, 2141
- out-of-range handling, 1123
- OUTFILE, 1434
- out\_dir option
  - comp\_err, 235
- out\_file option
  - comp\_err, 236
- overflow handling, 1123
- overflow page, 3809
- Overlaps(), 1312

---

---

overview, 1

## P

packages

list of, 44

PAD\_CHAR\_TO\_FULL\_LENGTH SQL mode, 606

page, 3809

page cleaner, 3810

page size, 3810

InnoDB, 1636, 1757

pager command

mysql, 257

pager option

mysql, 251

PAM

pluggable authentication, 790

PAM authentication

with Connector/J, 2670

parallel-recover option

myisamchk, 322

parameters

server, 932

PARAMETERS

INFORMATION\_SCHEMA table, 2189

parent table, 3810

parentheses ( and ), 1165

partial backup, 3810

partial index, 3810

partial trust, 2520

partial updates

and replication, 2074

PARTITION, 2085

PARTITION BY LIST COLUMNS, 2097

PARTITION BY RANGE COLUMNS, 2097

partition management, 2115

partition option

mysqld, 434

partition pruning, 2130

partitioning, 2085

advantages, 2088

and dates, 2089

and foreign keys, 2142

and FULLTEXT indexes, 2142

and key cache, 2141

and query cache, 2141

and replication, 2075

and SQL mode, 2075, 2139

and subqueries, 2143

and temporary tables, 2142, 2144

by hash, 2104

by key, 2107

by linear hash, 2106

by linear key, 2108

by list, 2094

by range, 2091

COLUMNS, 2097

concepts, 2086

data type of partitioning key, 2142

enabling, 2085

functions supported in partitioning expressions, 2148

limitations, 2139

operators not permitted in partitioning expressions, 2139

operators supported in partitioning expressions, 2139

optimization, 2129, 2130

resources, 2086

storage engines (limitations), 2147

subpartitioning, 2143

support, 2085

types, 2089

Partitioning

maximum number of partitions, 2141

partitioning information statements, 2128

partitioning keys and primary keys, 2144

partitioning keys and unique keys, 2144

partitions

adding and dropping, 2115

analyzing, 2127

checking, 2127

managing, 2115

modifying, 2115

optimizing, 2127

repairing, 2127

splitting and merging, 2115

truncating, 2115

PARTITIONS

INFORMATION\_SCHEMA table, 2190

password

root user, 150

password encryption

reversibility of, 1264

password option, 211

mysql, 251

mysqlaccess, 344

mysqladmin, 271

mysqlbinlog, 351

mysqlcheck, 277

mysqldump, 285

mysqld\_multi, 233

mysqlhotcopy, 367

mysqlimport, 301

mysqlshow, 305

mysqlslap, 311

mysql\_config\_editor, 342

mysql\_convert\_table\_format, 369

mysql\_setpermission, 371

password policy, 742

password validation, 741

PASSWORD(), 764, 780, 1263, 3733

passwords

administrator guidelines, 735

expiration, 781

for the InnoDB memcached interface, 1769

for users, 773

forgotten, 3736



---

hashing, 736  
 logging, 735  
 lost, 3736  
 resetting, 3736  
 security, 733, 751  
 setting, 780, 1530, 1535  
 user guidelines, 733  
 PATH environment variable, 140, 166, 209  
 path name separators  
     Windows, 217  
 pattern matching, 186, 1193  
 PEP-249 specification for Python, 2759  
 performance, 856  
     benchmarks, 956  
     disk I/O, 939  
     estimating, 908  
 Performance Schema, 1660, 2233, 3810  
     event filtering, 2247  
     memory use, 2243  
 performance-schema-consumer-xxx option  
     mysqld, 2306  
 performance-schema-instrument option  
     mysqld, 2306  
 performance\_schema  
     accounts table, 2284  
     cond\_instances table, 2271  
     events\_stages\_current table, 2278  
     events\_stages\_history table, 2279  
     events\_stages\_history\_long table, 2279  
     events\_stages\_summary\_by\_account\_by\_event\_name table, 2296  
     events\_stages\_summary\_by\_host\_by\_event\_name table, 2296  
     events\_stages\_summary\_by\_thread\_by\_event\_name table, 2289  
     events\_stages\_summary\_by\_user\_by\_event\_name table, 2296  
     events\_stages\_summary\_global\_by\_event\_name table, 2289  
     events\_statements\_current table, 2280  
     events\_statements\_history table, 2283  
     events\_statements\_history\_long table, 2283  
     events\_statements\_summary\_by\_account\_by\_event\_name table, 2296  
     events\_statements\_summary\_by\_digest table, 2290  
     events\_statements\_summary\_by\_host\_by\_event\_name table, 2296  
     events\_statements\_summary\_by\_thread\_by\_event\_name table, 2290  
     events\_statements\_summary\_by\_user\_by\_event\_name table, 2296  
     events\_statements\_summary\_global\_by\_event\_name table, 2290  
     events\_waits\_current table, 2275  
     events\_waits\_history table, 2277  
     events\_waits\_history\_long table, 2277  
     events\_waits\_summary\_by\_account\_by\_event\_name table, 2296  
     events\_waits\_summary\_by\_host\_by\_event\_name table, 2296  
     events\_waits\_summary\_by\_instance table, 2288  
     events\_waits\_summary\_by\_thread\_by\_event\_name table, 2288  
     events\_waits\_summary\_by\_user\_by\_event\_name table, 2296  
     events\_waits\_summary\_global\_by\_event\_name table, 2288  
     file\_instances table, 2271  
     file\_summary\_by\_event\_name table, 2292  
     file\_summary\_by\_instance table, 2292  
     hosts table, 2285  
     host\_cache table, 2299  
     mutex\_instances table, 2271  
     objects\_summary\_global\_by\_type table, 2291  
     performance\_timers table, 2301  
     rwlock\_instances table, 2273  
     session\_account\_connect\_attrs table, 2286  
     session\_connect\_attrs table, 2286  
     setup\_actors table, 2267  
     setup\_consumers table, 2267  
     setup\_instruments table, 2268  
     setup\_objects table, 2269  
     setup\_timers table, 2270  
     socket\_instances table, 2273  
     socket\_summary\_by\_event\_name table, 2298  
     socket\_summary\_by\_instance table, 2298  
     table\_io\_waits\_summary\_by\_index\_usage table, 2294  
     table\_io\_waits\_summary\_by\_table table, 2293  
     table\_lock\_waits\_summary\_by\_table table, 2295  
     thread table, 2302  
     users table, 2285  
 performance\_schema database, 2233  
     restrictions, 3765  
     TRUNCATE TABLE, 2266, 3765  
 PERFORMANCE\_SCHEMA storage engine, 2233  
 performance\_schema system variable, 2307  
 performance\_schema\_accounts\_size system variable, 2308  
 performance\_schema\_digests\_size system variable, 2308  
 performance\_schema\_events\_stages\_history\_long\_size system variable, 2308  
 performance\_schema\_events\_stages\_history\_size system variable, 2309  
 performance\_schema\_events\_statements\_history\_long\_size system variable, 2309  
 performance\_schema\_events\_statements\_history\_size system variable, 2309  
 performance\_schema\_events\_waits\_history\_long\_size system variable, 2310  
 performance\_schema\_events\_waits\_history\_size system variable, 2310  
 performance\_schema\_hosts\_size system variable, 2310

---

---

performance\_schema\_max\_cond\_classes system variable, 2311  
 performance\_schema\_max\_cond\_instances system variable, 2311  
 performance\_schema\_max\_file\_classes system variable, 2311  
 performance\_schema\_max\_file\_handles system variable, 2312  
 performance\_schema\_max\_file\_instances system variable, 2312  
 performance\_schema\_max\_mutex\_classes system variable, 2312  
 performance\_schema\_max\_mutex\_instances system variable, 2313  
 performance\_schema\_max\_rwlock\_classes system variable, 2313  
 performance\_schema\_max\_rwlock\_instances system variable, 2313  
 performance\_schema\_max\_socket\_classes system variable, 2314  
 performance\_schema\_max\_socket\_instances system variable, 2314  
 performance\_schema\_max\_stage\_classes system variable, 2314  
 performance\_schema\_max\_statement\_classes system variable, 2314  
 performance\_schema\_max\_table\_handles system variable, 2315  
 performance\_schema\_max\_table\_instances system variable, 2315  
 performance\_schema\_max\_thread\_classes system variable, 2316  
 performance\_schema\_max\_thread\_instances system variable, 2316  
 performance\_schema\_session\_connect\_attrs\_size system variable, 2316  
 performance\_schema\_setup\_actors\_size system variable, 2317  
 performance\_schema\_setup\_objects\_size system variable, 2317  
 performance\_schema\_users\_size system variable, 2317  
 performance\_timers table  
     performance\_schema, 2301  
 PERIOD\_ADD(), 1221  
 PERIOD\_DIFF(), 1221  
 Perl  
     installing, 167  
     installing on Windows, 168  
 Perl API, 3481  
 Perl DBI/DBD  
     installation problems, 168  
 permission checks  
     effect on speed, 863  
 perror, 207, 375  
     help option, 376  
     ndb option, 376  
     silent option, 376  
     verbose option, 376  
     version option, 376  
 persistent statistics, 3810  
 pessimistic, 3810  
 phantom, 3811  
 phantom rows, 1629  
 PHP API, 2905  
 physical, 3811  
 physical backup, 3811  
 PI(), 1205  
 pid-file option  
     mysql.server, 231  
     mysqld, 435  
     mysqld\_safe, 228  
 pid\_file system variable, 524  
 Ping  
     thread command, 958  
 pipe option, 211  
     mysql, 251, 278  
     mysqladmin, 271  
     mysqldump, 285  
     mysqlexport, 301  
     mysqlshow, 305  
     mysqslap, 312  
 PIPES\_AS\_CONCAT SQL mode, 607  
 PITR, 3811  
 plan option  
     mysqlaccess, 345  
 plan stability, 3811  
 pluggable authentication  
     PAM, 790  
     restrictions, 3766  
     Windows, 797  
 plugin, 3811  
     audit log, 817  
 plugin API, 609, 3485  
 plugin option prefix  
     mysqld, 435  
 plugin services, 3530  
 plugin-dir option  
     mysql, 252  
     mysqladmin, 271  
     mysqlbinlog, 352  
     mysqlcheck, 278  
     mysqldump, 285  
     mysqld\_safe, 228  
     mysqlexport, 301  
     mysqlshow, 305  
     mysqslap, 312  
     mysql\_plugin, 240  
     mysql\_upgrade, 243  
 plugin-ini option  
     mysql\_plugin, 240  
 plugin-load option  
     mysqld, 435  
 plugin-load-add option  
     mysqld, 436  
 plugindir option

---

- mysql\_config, 373
- plugins
  - activating, 609
  - adding, 3485
  - audit, 3489
  - authentication, 3489
  - daemon, 3488
  - full-text parser, 3487
  - INFORMATION\_SCHEMA, 3489
  - installing, 609, 1545
  - semisynchronous replication, 3489
  - server, 609
  - storage engine, 3487
  - uninstalling, 609, 1546
- PLUGINS
  - INFORMATION\_SCHEMA table, 2193
- plugin\_dir system variable, 525
- POINT data type, 1297
- Point(), 1300
- point-in-time recovery, 845, 3811
- PointFromText(), 1298
- PointFromWKB(), 1299
- PointN(), 1305
- PointOnSurface(), 1307
- PolyFromText(), 1298
- PolyFromWKB(), 1299
- POLYGON data type, 1297
- Polygon(), 1300
- PolygonFromText(), 1298
- PolygonFromWKB(), 1299
- port option, 211
  - mysql, 252
  - mysqladmin, 271
  - mysqlbinlog, 352
  - mysqlcheck, 278
  - mysqld, 437
  - mysqldump, 285
  - mysqld\_safe, 228
  - mysqlhotcopy, 368
  - mysqlimport, 301
  - mysqlshow, 305
  - mysqlslap, 312
  - mysql\_config, 373
  - mysql\_config\_editor, 342
  - mysql\_convert\_table\_format, 369
  - mysql\_setpermission, 371
- port system variable, 525
- port-open-timeout option
  - mysqld, 437
- portability, 857
  - types, 1151
- porting
  - to other systems, 3543
- ports, 135, 148, 165, 209, 352, 710, 732, 768, 817, 1981, 2273, 3723
- POSITION(), 1183
- post-filtering
  - Performance Schema, 2247
- post-query option
  - mysqlslap, 312
- post-system option
  - mysqlslap, 312
- PostgreSQL compatibility, 29
- POSTGRESQL SQL mode, 609
- postinstall
  - multiple servers, 705
- postinstallation
  - setup and testing, 139
- POW(), 1206
- POWER(), 1206
- pre-filtering
  - Performance Schema, 2247
- pre-query option
  - mysqlslap, 312
- pre-system option
  - mysqlslap, 312
- precedence
  - operator, 1165
- precision
  - arithmetic, 1315
  - fractional seconds, 1111, 1115
  - numeric, 1111
- precision math, 1315
- preload\_buffer\_size system variable, 525
- Prepare
  - thread command, 958
- PREPARE, 1486, 1489
  - XA transactions, 1472
- prepared backup, 3811
- prepared statements, 1486, 1489, 1490, 1490, 2859
  - reparation, 924
  - with Connector/Net, 2495
- preparing
  - thread state, 962
- preparing for alter table
  - thread state, 962
- preview option
  - mysqlaccess, 345
- primary key, 3812
  - constraint, 35
  - deleting, 1335
- PRIMARY KEY, 1335, 1362
- primary keys
  - and partitioning keys, 2144
- print command
  - mysql, 257
- print-defaults option, 220
  - mysql\_plugin, 240
- privilege
  - changes, 768
- privilege checks
  - effect on speed, 863
- privilege information
  - location, 756
- privilege system, 751
- privileges

---

---

- access, 751
- adding, 774
- and replication, 2074
- default, 150
- deleting, 778, 1524
- display, 1563
- dropping, 778, 1524
- granting, 1524
- revoking, 1534
- problems
  - access denied errors, 3723
  - common errors, 3722
  - compiling MySQL server, 137
  - DATE columns, 3744
  - date values, 1126
  - installing on Solaris, 118
  - installing Perl, 168
  - linking, 2798
  - lost connection errors, 3726
  - ODBC, 2402
  - reporting, 2, 21
  - starting the server, 148
  - table locking, 927
  - time zone, 3742
- PROCEDURE, 1432
- PROCEDURE ANALYSE(), 880
- procedures
  - stored, 2153
- process, 3812
- process support, 49
- processes
  - display, 1568
- processing
  - arguments, 3537
- Processlist
  - thread command, 959
- PROCESSLIST, 1568
  - INFORMATION\_SCHEMA table, 2194
  - possible inconsistency with INFORMATION\_SCHEMA tables, 1672
- product option
  - MySQLInstallerConsole, 84
- PROFILING
  - INFORMATION\_SCHEMA table, 2195
- profiling session variable, 526
- profiling\_history\_size session variable, 526
- program variables
  - setting, 220
- program-development utilities, 207
- programs
  - administrative, 206
  - client, 205, 2797
  - stored, 1490, 2151
  - utility, 206
- prompt command
  - mysql, 257
- prompt option
  - mysql, 252

- prompts
  - meanings, 174
- pronunciation
  - MySQL, 5
- protocol option, 211
  - mysql, 252
  - mysqladmin, 272
  - mysqlbinlog, 352
  - mysqlcheck, 278
  - mysqldump, 285
  - mysqlexport, 301
  - mysqlshow, 305
  - mysqslap, 312
- protocol\_version system variable, 526
- proximity search, 1234
- Proxy, 1891
- proxy\_user session variable, 526
- pseudo-record, 3812
- pseudo\_slave\_mode system variable, 526
- pseudo\_thread\_id system variable, 527
- Pthreads, 3812
- purge, 3812
- PURGE BINARY LOGS, 1474
- purge buffering, 3812
- purge lag, 3812
- PURGE MASTER LOGS, 1474
- purge scheduling, 1662
- purge thread, 3812
- Purging old relay logs
  - thread state, 962
- Python, 2759
  - third-party driver, 3482

## Q

- QUARTER(), 1221
- queries
  - entering, 172
  - estimating performance, 908
  - examples, 193
  - speed of, 858
- Query
  - thread command, 959
- query, 3812
- Query Cache, 918
- query cache
  - and partitioned tables, 2141
  - thread states, 966
- query end
  - thread state, 962
- query execution plan, 3813
- query expansion, 1236
- query option
  - mysqslap, 312
- query\_alloc\_block\_size system variable, 527
- query\_cache\_limit system variable, 527
- query\_cache\_min\_res\_unit system variable, 528
- query\_cache\_size system variable, 528
- query\_cache\_type system variable, 529

---

query\_cache\_wlock\_invalidate system variable, 530  
 query\_prealloc\_size system variable, 531  
 questions, 268  
   answering, 20  
 Queuing master event to the relay log  
   thread state, 968  
 quick option  
   myisamchk, 322  
   mysql, 252  
   mysqlcheck, 278  
   mysqldump, 294  
 quiesce, 3813  
 quiet option  
   mysqlhotcopy, 368  
 Quit  
   thread command, 959  
 quit command  
   mysql, 257  
 quotation marks  
   in strings, 1012  
 QUOTE(), 1013, 1183, 2846  
 quote-names option  
   mysqldump, 291  
 quoting, 1013  
   column alias, 1018, 3746  
 quoting binary data, 1013  
 quoting of identifiers, 1017

**R**

RADIANS(), 1206  
 RAID, 3813  
 RAND(), 1206  
 random dive, 3813  
 random-passwords option  
   mysql\_install\_db, 237  
 rand\_seed1 session variable, 531  
 rand\_seed2 session variable, 531  
 range join type  
   optimizer, 899  
 range partitioning, 2091, 2097  
 range partitions  
   adding and dropping, 2116  
   managing, 2116  
 range\_alloc\_block\_size system variable, 531  
 raw backup, 3813  
 raw option  
   mysql, 252  
   mysqlbinlog, 352  
 raw partitions, 1616  
 re-creating  
   grant tables, 145  
 READ COMMITTED, 3813  
   transaction isolation level, 1469  
 READ UNCOMMITTED, 3814  
   transaction isolation level, 1470  
 read view, 3814  
 read-ahead, 1692, 3814  
   linear, 1655  
   random, 1655  
 read-from-remote-master option  
   mysqlbinlog, 352  
 read-from-remote-server option  
   mysqlbinlog, 352  
 read-only option  
   myisamchk, 321  
   mysqld, 1973  
 read-only transaction, 3814  
 Reading event from the relay log  
   thread state, 968  
 Reading from net  
   thread state, 962  
 Reading master dump table data  
   thread state, 969  
 read\_buffer\_size myisamchk variable, 319  
 read\_buffer\_size system variable, 532  
 read\_only system variable, 532  
 read\_rnd\_buffer\_size system variable, 533  
 REAL data type, 1115  
 REAL\_AS\_FLOAT SQL mode, 607  
 Rebuilding the index on master dump table  
   thread state, 969  
 reconfiguring, 137  
 reconnect option  
   mysql, 252  
 Reconnecting after a failed binlog dump request  
   thread state, 967  
 Reconnecting after a failed master event read  
   thread state, 968  
 reconnection  
   automatic, 2304, 2897  
 record lock, 3814  
 record-level locks  
   InnoDB, 1623, 1628, 1629, 1730  
 record\_log\_pos option  
   mysqlhotcopy, 368  
 RECOVER  
   XA transactions, 1472  
 recover option  
   myisamchk, 322  
 recovery  
   from crash, 848  
   incremental, 845  
   point in time, 845  
 redo, 3814  
 redo log, 1615, 3814  
 reducing  
   data size, 878  
 redundant row format, 673, 3815  
 ref join type  
   optimizer, 898  
 references, 1335  
 referential integrity, 1605, 1606, 3815  
 REFERENTIAL\_CONSTRAINTS  
   INFORMATION\_SCHEMA table, 2196  
 Refresh  
   thread command, 959

---

---

- ref\_or\_null, 981
- ref\_or\_null join type
  - optimizer, 899
- REGEXP, 1193
- REGEXP operator, 1193
- regex option
  - mysqlhotcopy, 368
  - mysql\_find\_rows, 370
- Register Slave
  - thread command, 959
- Registering slave on master
  - thread state, 967
- regular expression syntax, 1193
- rehash command
  - mysql, 258
- relational, 3815
- relational databases
  - defined, 5
- relative option
  - mysqladmin, 272
- relay logs (replication), 2028
- relay-log option
  - mysqld, 1974
- relay-log-index option
  - mysqld, 1974
- relay-log-info-file option
  - mysqld, 1975
- relay-log-info-repository option
  - mysqld, 2000
- relay-log-purge option
  - mysqld, 1975
- relay-log-recovery option
  - mysqld, 1975
- relay-log-space-limit option
  - mysqld, 1976
- relay\_log system variable, 1989
- relay\_log\_basename system variable, 1989
- relay\_log\_index system variable, 1989
- relay\_log\_info\_file system variable, 1989
- relay\_log\_info\_repository system variable, 1990
- relay\_log\_purge system variable, 534
- relay\_log\_recovery system variable, 1990
- relay\_log\_space\_limit system variable, 534
- release notes, 3755
- release numbers, 50
- RELEASE SAVEPOINT, 1462
- releases
  - naming scheme, 51
  - testing, 52
  - updating, 53
- RELEASE\_LOCK(), 1280
- relevance, 3815
- relnotes option
  - mysqlaccess, 345
- remove option
  - mysqld, 437
- Removing duplicates
  - thread state, 962
- removing tmp table
  - thread state, 962
- rename
  - thread state, 962
- rename result table
  - thread state, 963
- RENAME TABLE, 1394
- RENAME USER, 1534
- renaming user accounts, 1534
- Reopen tables
  - thread state, 963
- repair
  - tables, 272
- Repair by sorting
  - thread state, 963
- Repair done
  - thread state, 963
- repair option
  - mysqlcheck, 278
- repair options
  - myisamchk, 321
- REPAIR TABLE, 1542
  - and partitioning, 2127
  - and replication, 1544, 2071
- Repair with keycache
  - thread state, 963
- repairing
  - tables, 850
- REPEAT, 1496
  - labels, 1491
- REPEAT(), 1183
- REPEATABLE READ, 3815
  - transaction isolation level, 1469
- repertoire
  - character set, 1060, 1067
- replace, 207
- REPLACE, 1426
- replace option
  - mysqldump, 286
  - mysqlimport, 301
- replace utility, 376
- REPLACE(), 1183
- replicate-do-db option
  - mysqld, 1976
- replicate-do-table option
  - mysqld, 1978
- replicate-ignore-db option
  - mysqld, 1977
- replicate-ignore-table option
  - mysqld, 1978
- replicate-rewrite-db option
  - mysqld, 1979
- replicate-same-server-id option
  - mysqld, 1979
- replicate-wild-do-table option
  - mysqld, 1979
- replicate-wild-ignore-table option
  - mysqld, 1980

---

---

- replication, 1929, 3815
  - and AUTO\_INCREMENT, 2058
  - and character sets, 2059
  - and CREATE ... IF NOT EXISTS, 2059
  - and CREATE TABLE ... SELECT, 2060
  - and DATA DIRECTORY, 2067
  - and DROP ... IF EXISTS, 2061
  - and errors on slave, 2074
  - and floating-point values, 2068
  - and FLUSH, 2069
  - and fractional seconds, 2068
  - and functions, 2069
  - and INDEX DIRECTORY, 2067
  - and invoked features, 2067
  - and LAST\_INSERT\_ID(), 2058
  - and LIMIT, 2071
  - and LOAD DATA, 2071
  - and max\_allowed\_packet, 2072
  - and MEMORY tables, 2072
  - and mysql (system) database, 2074
  - and partial updates, 2074
  - and partitioning, 2075
  - and privileges, 2074
  - and query optimizer, 2074
  - and REPAIR TABLE statement, 1544, 2071
  - and reserved words, 2074
  - and scheduled events, 2067, 2067
  - and SQL mode, 2075
  - and stored routines, 2067
  - and temporary tables, 2073
  - and time zones, 2075
  - and TIMESTAMP, 2058
  - and transactions, 2075, 2076
  - and triggers, 2067, 2077
  - and TRUNCATE TABLE, 2077
  - and variables, 2078
  - and views, 2079
  - attribute demotion, 2063
  - attribute promotion, 2063
  - crashes, 2072
  - delayed, 2057
  - relay logs, 2028
  - row-based vs statement-based, 1941
  - safe and unsafe statements, 1945
  - semisynchronous, 2052
  - shutdown and restart, 2072, 2073
  - statements incompatible with STATEMENT format, 1941
  - status logs, 2028
  - timeouts, 2075
  - with differing tables on master and slave, 2061
  - with ZFS, 1832
- replication filtering options
  - and case sensitivity, 2034
- replication formats
  - compared, 1941
- replication implementation, 2026
- replication limitations, 2057
- replication log tables, 2028
- replication master
  - thread states, 967
- replication masters
  - statements, 1474
- replication options, 2057
- replication slave
  - thread states, 967, 968, 969
- replication slaves
  - statements, 1476
- report-host option
  - mysqld, 1980
- report-password option
  - mysqld, 1981
- report-port option
  - mysqld, 1981
- report-user option
  - mysqld, 1982
- reporting
  - bugs, 2, 21
  - Connector/Net problems, 2633
  - Connector/ODBC problems, 2402, 2402
  - errors, 21
  - problems, 2
- report\_host system variable, 535
- report\_password system variable, 535
- report\_port system variable, 535
- report\_user system variable, 536
- Requesting binlog dump
  - thread state, 967
- REQUIRE GRANT option, 1532
- reschedule
  - thread state, 966
- reserved words, 1026
  - and replication, 2074
- RESET MASTER, 1475
- RESET SLAVE, 1481
- RESET SLAVE ALL, 1481
- Reset stmt
  - thread command, 959
- resetmaster option
  - mysqlhotcopy, 368
- resetslave option
  - mysqlhotcopy, 368
- RESIGNAL, 1506
- resolveip, 208, 377
  - help option, 377
  - silent option, 377
  - version option, 377
- resolve\_stack\_dump, 207, 375
  - help option, 375
  - numeric-dump-file option, 375
  - symbols-file option, 375
  - version option, 375
- resource limits
  - user accounts, 509, 778, 1531
- restarting
  - the server, 142

---

---

- restore, 3816
- restrictions
  - character sets, 3765
  - events, 3757
  - InnoDB, 1756
  - performance\_schema database, 3765
  - pluggable authentication, 3766
  - server-side cursors, 3760
  - signal, 3760
  - stored routines, 3757
  - subqueries, 3761
  - triggers, 3757
  - views, 3763
  - XA transactions, 3764
- result-file option
  - mysqlbinlog, 353
  - mysqldump, 291
- retrieving
  - data from tables, 179
- RETURN, 1497
- return (\r), 1012, 1417
- return values
  - UDFs, 3539
- REVERSE(), 1184
- REVOKE, 1534
- revoking
  - privileges, 1534
- rhost option
  - mysqlaccess, 345
- RIGHT JOIN, 1435
- RIGHT OUTER JOIN, 1435
- RIGHT(), 1184
- RLIKE, 1193
- ROLLBACK, 30, 1458
  - XA transactions, 1472
- rollback, 3816
- rollback option
  - mysqlaccess, 345
- rollback segment, 1645, 3816
- ROLLBACK TO SAVEPOINT, 1462
- Rolling back
  - thread state, 963
- ROLLUP, 1285
- root password, 150
- root user, 732
  - password resetting, 3736
- ROUND(), 1207
- rounding, 1315
- rounding errors, 1113
- ROUTINES
  - INFORMATION\_SCHEMA table, 2196
- routines option
  - mysqldump, 293
- ROW, 1448
  - maximum, 3769
- row, 3816
- row format, 3816
- row lock, 3816
- row size
  - row subqueries, 1448
- row-based replication, 3816
  - advantages, 1942
  - disadvantages, 1943
- row-level locking, 925, 3816
- rows
  - counting, 188
  - deleting, 3747
  - locking, 32
  - matching problems, 3747
  - selecting, 180
  - sorting, 182
- rows option
  - mysql\_find\_rows, 370
- ROW\_COUNT(), 1272
- ROW\_FORMAT
  - COMPACT, 673
  - COMPRESSED, 657, 673
  - DYNAMIC, 673
  - REDUNDANT, 673
- RPAD(), 1184
- Rpl\_semi\_sync\_master\_clients status variable, 596
- rpl\_semi\_sync\_master\_enabled system variable, 536
- Rpl\_semi\_sync\_master\_net\_avg\_wait\_time status variable, 596
- Rpl\_semi\_sync\_master\_net\_waits status variable, 596
- Rpl\_semi\_sync\_master\_net\_wait\_time status variable, 596
- Rpl\_semi\_sync\_master\_no\_times status variable, 596
- Rpl\_semi\_sync\_master\_no\_tx status variable, 596
- Rpl\_semi\_sync\_master\_status status variable, 596
- Rpl\_semi\_sync\_master\_timefunc\_failures status variable, 597
- rpl\_semi\_sync\_master\_timeout system variable, 536
- rpl\_semi\_sync\_master\_trace\_level system variable, 536
- Rpl\_semi\_sync\_master\_tx\_avg\_wait\_time status variable, 597
- Rpl\_semi\_sync\_master\_tx\_waits status variable, 597
- Rpl\_semi\_sync\_master\_tx\_wait\_time status variable, 597
- rpl\_semi\_sync\_master\_wait\_no\_slave system variable, 537
- Rpl\_semi\_sync\_master\_wait\_pos\_backtraverse status variable, 597
- Rpl\_semi\_sync\_master\_wait\_sessions status variable, 597
- Rpl\_semi\_sync\_master\_yes\_tx status variable, 597
- rpl\_semi\_sync\_slave\_enabled system variable, 537
- Rpl\_semi\_sync\_slave\_status status variable, 597
- rpl\_semi\_sync\_slave\_trace\_level system variable, 538
- RPM file, 110
- rpm option
  - mysql\_install\_db, 238
- RPM Package Manager, 110
- RTRIM(), 1184
- Ruby API, 3482

---



---

running  
   ANSI mode, 27  
   batch mode, 192  
   multiple servers, 705  
   queries, 172  
 running CMake after prior invocation, 126, 137  
 rw-lock, 3817  
 rwlock\_instances table  
   performance\_schema, 2273

**S**

safe statement (replication)  
   defined, 1945  
 safe-mode option  
   mysqld, 437  
 safe-recover option  
   myisamchk, 322  
 safe-updates option, 265  
   mysql, 252  
 safe-user-create option  
   mysqld, 438  
 Sakila, 8  
 sandbox mode, 781  
 SASL, 1769  
 SAVEPOINT, 1462  
 savepoint, 3817  
 Saving state  
   thread state, 963  
 scalability, 3817  
 scale  
   arithmetic, 1315  
   numeric, 1111  
 scale out, 3817  
 scale up, 3817  
 schema, 3817  
   altering, 1326  
   creating, 1344  
   deleting, 1391  
 SCHEMA(), 1273  
 SCHEMATA  
   INFORMATION\_SCHEMA table, 2198  
 SCHEMA\_PRIVILEGES  
   INFORMATION\_SCHEMA table, 2198  
 script files, 192  
 scripts, 225, 231  
   mysql\_install\_db, 144  
   SQL, 244  
 search index, 3817  
 searching  
   and case sensitivity, 3743  
   full-text, 1229  
   MySQL Web pages, 21  
   two keys, 198  
 Searching rows for update  
   thread state, 963  
 SECOND(), 1221  
 secondary index, 3818  
   InnoDB, 1635  
   secure-auth option  
     mysql, 253  
     mysqld, 438  
   secure-file-priv option  
     mysqld, 438  
   secure\_auth system variable, 538  
   secure\_file\_priv system variable, 538  
   security  
     against attackers, 746  
     for the InnoDB memcached interface, 1769  
   security system, 751  
   SEC\_TO\_TIME(), 1221  
   segment, 3818  
   SELECT  
     INTO, 1433  
     LIMIT, 1427  
     optimizing, 894, 1596  
     Query Cache, 918  
   SELECT INTO TABLE, 30  
   SELECT speed, 858  
   selecting  
     databases, 176  
   selectivity, 3818  
   select\_limit variable, 255  
   semi-consistent read, 3818  
     InnoDB, 1730  
   semi-joins, 1001  
   semisynchronous replication, 2052  
     administrative interface, 2053  
     configuration, 2054  
     installation, 2054  
     monitoring, 2056  
   semisynchronous replication plugins, 3489  
   Sending binlog event to slave  
     thread state, 967  
   sending cached result to client  
     thread state, 966  
   SEQUENCE, 198  
   sequence emulation, 1271  
   sequences, 198  
   SERIAL, 1112, 1113  
   SERIAL DEFAULT VALUE, 1147  
   SERIALIZABLE, 3818  
     transaction isolation level, 1470  
   server, 3818  
     connecting, 171, 209  
     debugging, 3544  
     disconnecting, 171  
     logs, 619  
     restart, 142  
     shutdown, 142  
     signal handling, 617  
     starting, 140  
     starting and stopping, 146  
     starting problems, 148  
   server administration, 266  
   server plugins, 609  
   server variables, 1584 (see system variables)

---

---

- server-id option
  - mysqlbinlog, 353
  - mysqld, 1955
- server-public-key-path option
  - mysql, 253
- server-side cursor
  - restrictions, 3760
- servers
  - multiple, 705
- server\_id system variable, 539
- server\_uuid system variable
  - mysqld, 1955
- service-startup-timeout option
  - mysql.server, 231
- services
  - for plugins, 3530
- servlet, 2701
- session variable
  - autocommit, 464
  - big\_tables, 466
  - disconnect\_on\_expired\_passwords, 477
  - error\_count, 480
  - external\_user, 482
  - foreign\_key\_checks, 483
  - identity, 488
  - insert\_id, 490
  - last\_insert\_id, 495
  - profiling, 526
  - profiling\_history\_size, 526
  - proxy\_user, 526
  - rand\_seed1, 531
  - rand\_seed2, 531
  - sql\_auto\_is\_null, 543
  - sql\_big\_selects, 543
  - sql\_buffer\_result, 544
  - sql\_notes, 546
  - sql\_quote\_show\_create, 546
  - sql\_safe\_updates, 547
  - sql\_warnings, 547
  - timestamp, 558
  - unique\_checks, 561
  - warning\_count, 564
- session variables
  - and replication, 2078
- session\_account\_connect\_attrs table
  - performance\_schema, 2286
- session\_connect\_attrs table
  - performance\_schema, 2286
- SESSION\_STATUS
  - INFORMATION\_SCHEMA table, 2188
- SESSION\_USER(), 1273
- SESSION\_VARIABLES
  - INFORMATION\_SCHEMA table, 2188
- SET, 1547
  - CHARACTER SET, 1046, 1549
  - NAMES, 1046, 1049, 1549
  - ONE\_SHOT, 1549
  - size, 1150
  - SET data type, 1120, 1144
  - SET GLOBAL sql\_slave\_skip\_counter, 1481
  - Set option
    - thread command, 959
  - SET PASSWORD, 1535
  - SET PASSWORD statement, 780
  - SET sql\_log\_bin, 1476
  - SET statement
    - assignment operator, 1173
  - SET TRANSACTION, 1468
  - set-auto-increment[ option
    - myisamchk, 323
  - set-character-set option
    - myisamchk, 323
  - set-charset option
    - mysqlbinlog, 353
    - mysqldump, 288
  - set-collation option
    - myisamchk, 323
  - set-gtid-purged option
    - mysqldump, 289
  - setting
    - passwords, 780
  - setting passwords, 1535
  - setting program variables, 220
  - setup
    - postinstallation, 139
    - thread state, 963
  - setup\_actors table
    - performance\_schema, 2267
  - setup\_consumers table
    - performance\_schema, 2267
  - setup\_instruments table
    - performance\_schema, 2268
  - setup\_objects table
    - performance\_schema, 2269
  - setup\_timers table
    - performance\_schema, 2270
  - SHA(), 1264
  - SHA1(), 1264
  - SHA2(), 1264
  - sha256\_password\_private\_key\_path system variable, 539
  - sha256\_password\_public\_key\_path system variable, 539
  - shared lock, 3819
  - shared tablespace, 3819
  - shared-memory option
    - mysqld, 439
  - shared-memory-base-name option, 211
    - mysqld, 439
    - mysqldslap, 312
  - shared\_memory system variable, 540
  - shared\_memory\_base\_name system variable, 540
  - sharp checkpoint, 3819
  - shell syntax, 4
  - short-form option
    - mysqlbinlog, 353

---

---

SHOW AUTHORS, 1549, 1550  
 SHOW BINARY LOGS, 1549, 1550  
 SHOW BINLOG EVENTS, 1549, 1551  
 SHOW CHARACTER SET, 1549, 1551  
 SHOW COLLATION, 1549, 1552  
 SHOW COLUMNS, 1549, 1552  
 SHOW CONTRIBUTORS, 1549, 1554  
 SHOW CREATE DATABASE, 1549, 1554  
 SHOW CREATE EVENT, 1549  
 SHOW CREATE FUNCTION, 1549, 1555  
 SHOW CREATE PROCEDURE, 1549, 1555  
 SHOW CREATE SCHEMA, 1549, 1554  
 SHOW CREATE TABLE, 1549, 1556  
 SHOW CREATE TRIGGER, 1549, 1556  
 SHOW CREATE VIEW, 1549, 1556  
 SHOW DATABASES, 1549, 1557  
 SHOW ENGINE, 1549, 1557  
 SHOW ENGINE INNODB MUTEX, 1692  
 SHOW ENGINE INNODB STATUS, 1557  
     and innodb\_adaptive\_hash\_index, 1653  
     and innodb\_use\_sys\_malloc, 1652  
 SHOW ENGINES, 1549, 1559  
 SHOW ERRORS, 1549, 1561  
 SHOW EVENTS, 1549, 1561  
 SHOW extensions, 2229  
 SHOW FIELDS, 1549, 1554  
 SHOW FUNCTION CODE, 1549, 1563  
 SHOW FUNCTION STATUS, 1549, 1563  
 SHOW GRANTS, 1549, 1563  
 SHOW INDEX, 1549, 1563  
 SHOW KEYS, 1549, 1563  
 SHOW MASTER LOGS, 1549, 1550  
 SHOW MASTER STATUS, 1549, 1565  
 SHOW OPEN TABLES, 1549, 1565  
 SHOW PLUGINS, 1549, 1566  
 SHOW PRIVILEGES, 1549, 1566  
 SHOW PROCEDURE CODE, 1549, 1567  
 SHOW PROCEDURE STATUS, 1549, 1568  
 SHOW PROCESSLIST, 1549, 1568  
 SHOW PROFILE, 1549, 1570  
 SHOW PROFILES, 1549, 1570, 1572  
 SHOW RELAYLOG EVENTS, 1573  
 SHOW SCHEDULER STATUS, 2162  
 SHOW SCHEMAS, 1557  
 SHOW SLAVE HOSTS, 1549, 1573  
 SHOW SLAVE STATUS, 1549, 1574  
 SHOW STATUS, 1549  
 SHOW STORAGE ENGINES, 1559  
 SHOW TABLE STATUS, 1549  
 SHOW TABLES, 1549, 1583  
 SHOW TRIGGERS, 1549, 1583  
 SHOW VARIABLES, 1549  
 SHOW WARNINGS, 1549, 1586  
 SHOW with WHERE, 2178, 2229  
 show-slave-auth-info option  
     mysqld, 1982  
 show-table-type option  
     mysqlshow, 305  
 show-warnings option  
     mysql, 253  
 showing  
     database information, 302  
 shutdown, 618, 3819  
 Shutdown  
     thread command, 959  
 shutdown\_timeout variable, 272  
 shutting down  
     the server, 142  
 sigint-ignore option  
     mysql, 253  
 SIGN(), 1208  
 SIGNAL, 1510  
 signal  
     restrictions, 3760  
 signals  
     server response, 617  
 silent column changes, 1383  
 silent option  
     myisamchk, 319  
     myisampack, 333  
     mysql, 253  
     mysqladmin, 272  
     mysqlcheck, 278  
     mysqld\_multi, 233  
     mysqlimport, 301  
     mysqlslap, 312  
     perror, 376  
     resolveip, 377  
 SIN(), 1208  
 single quote ('), 1012  
 single-transaction option  
     mysqldump, 295  
 size of tables, 3768  
 sizes  
     display, 1111  
 --skip option prefix, 214  
 skip-column-names option  
     mysql, 253  
 skip-comments option  
     mysqldump, 287  
 skip-concurrent-insert option  
     mysqld, 439  
 skip-database option  
     mysqlcheck, 278  
 skip-event-scheduler option  
     mysqld, 439  
 skip-external-locking option  
     mysqld, 439  
 skip-grant-tables option  
     mysqld, 439  
 skip-gtids option  
     mysqlbinlog, 353  
 skip-host-cache option  
     mysqld, 439  
 skip-innodb option  
     mysqld, 440, 1698

---

---

skip-kill-mysqld option  
   mysqld\_safe, 229  
 skip-line-numbers option  
   mysql, 253  
 skip-name-resolve option  
   mysqld, 440  
   mysql\_install\_db, 238  
 skip-networking option  
   mysqld, 440  
 skip-opt option  
   mysqldump, 294  
 skip-partition option  
   mysqld, 440  
 skip-show-database option  
   mysqld, 441  
 skip-slave-start option  
   mysqld, 1984  
 skip-ssl option, 810  
 skip-stack-trace option  
   mysqld, 441  
 skip-symbolic-links option  
   mysqld, 441  
 skip-syslog option  
   mysqld\_safe, 229  
 skip-thread-priority option  
   mysqld, 441  
 skip-use-db option  
   mysql\_find\_rows, 370  
 skip\_external\_locking system variable, 540  
 skip\_name\_resolve system variable, 540  
 skip\_networking system variable, 540  
 skip\_show\_database system variable, 540  
 Slave has read all relay log; waiting for the slave I/O  
 thread to update it  
   thread state, 968  
 slave server, 3819  
 slave-checkpoint-group option  
   mysqld, 1982  
 slave-checkpoint-period option  
   mysqld, 1983  
 slave-load-tmpdir option  
   mysqld, 1985  
 slave-max-allowed-packet (mysqld), 1964  
 slave-net-timeout option  
   mysqld, 1985  
 slave-parallel-workers option  
   mysqld, 1983  
 slave-pending-jobs-size-max option  
   mysqld, 1984  
 slave-rows-search-algorithms (mysqld), 1964  
 slave-skip-errors option  
   mysqld, 1986  
 slave-sql-verify-checksum option  
   mysqld, 1987  
 slave\_checkpoint\_group system variable, 1991  
 slave\_checkpoint\_period system variable, 1991  
 slave\_compressed\_protocol option  
   mysqld, 1985  
 slave\_compressed\_protocol system variable, 1992  
 slave\_exec\_mode system variable, 1992  
 slave\_load\_tmpdir system variable, 1992  
 slave\_max\_allowed\_packet system variable, 1969  
 slave\_net\_timeout system variable, 1993  
 slave\_parallel\_workers system variable, 1993  
 slave\_pending\_jobs\_size\_max system variable, 1994  
 slave\_rows\_search\_algorithms system variable, 1969  
 slave\_skip\_errors system variable, 1994  
 slave\_sql\_verify\_checksum system variable, 1995  
 slave\_transaction\_retries system variable, 1995  
 slave\_type\_conversions system variable, 1996  
 Sleep  
   thread command, 959  
 sleep option  
   mysqladmin, 272  
 SLEEP(), 1280  
 slow queries, 268  
 slow query log, 635, 3819  
 slow shutdown, 3819  
 slow-query-log option  
   mysqld, 441  
 slow-start-timeout option  
   mysqld, 442  
 slow\_launch\_time system variable, 541  
 slow\_query\_log system variable, 541  
 slow\_query\_log\_file system variable, 541  
 SMALLINT data type, 1113  
 snapshot, 3819  
 socket option, 211  
   mysql, 253  
   mysqladmin, 272  
   mysqlbinlog, 353  
   mysqlcheck, 278  
   mysqld, 442  
   mysqldump, 285  
   mysqld\_safe, 229  
   mysqlhotcopy, 368  
   mysqlimport, 301  
   mysqlshow, 305  
   mysqslap, 313  
   mysql\_config, 373  
   mysql\_config\_editor, 343  
   mysql\_convert\_table\_format, 369  
   mysql\_setpermission, 371  
 socket system variable, 542  
 socket\_instances table  
   performance\_schema, 2273  
 socket\_summary\_by\_event\_name table  
   performance\_schema, 2298  
 socket\_summary\_by\_instance table  
   performance\_schema, 2298  
 Solaris  
   installation, 118  
 Solaris installation problems, 118  
 Solaris troubleshooting, 138  
 Solaris x86\_64 issues, 887  
 SOME, 1447

---

---

- sort-index option
  - myisamchk, 323
- sort-records option
  - myisamchk, 324
- sort-recover option
  - myisamchk, 323
- sorting
  - data, 182
  - grant tables, 765, 767
  - table rows, 182
- Sorting for group
  - thread state, 963
- Sorting for order
  - thread state, 963
- Sorting index
  - thread state, 963
- Sorting result
  - thread state, 963
- sort\_buffer\_size myisamchk variable, 319
- sort\_buffer\_size system variable, 542
- sort\_key\_blocks myisamchk variable, 319
- SOUNDEX(), 1184
- SOUNDS LIKE, 1185
- source (mysql client command), 193, 262
- source command
  - mysql, 258
- source distribution
  - installing, 122
- space ID, 3820
- SPACE(), 1185
- spassword option
  - mysqlaccess, 345
- Spatial Extensions in MySQL, 1289
- speed
  - increasing with replication, 1929
  - inserting, 862
  - of queries, 858, 858
- spin, 3820
- sporadic-binlog-dump-fail option
  - mysqld, 2007
- Spring framework, 2694
- SQL, 3820
  - defined, 5
- SQL Editor, 2441
- SQL mode, 602
  - ALLOW\_INVALID\_DATES, 603
  - and partitioning, 2075, 2139
  - and replication, 2075
  - ANSI, 603, 608
  - ANSI\_QUOTES, 603
  - DB2, 608
  - ERROR\_FOR\_DIVISION\_BY\_ZERO, 603
  - HIGH\_NOT\_PRECEDENCE, 603
  - IGNORE\_SPACE, 604
  - MAXDB, 608
  - MSSQL, 608
  - MYSQL323, 608
  - MYSQL40, 609
  - NO\_AUTO\_CREATE\_USER, 604
  - NO\_AUTO\_VALUE\_ON\_ZERO, 604
  - NO\_BACKSLASH\_ESCAPES, 604
  - NO\_DIR\_IN\_CREATE, 604
  - NO\_ENGINE\_SUBSTITUTION, 605
  - NO\_FIELD\_OPTIONS, 605
  - NO\_KEY\_OPTIONS, 605
  - NO\_TABLE\_OPTIONS, 605
  - NO\_UNSIGNED\_SUBTRACTION, 605
  - NO\_ZERO\_DATE, 606
  - NO\_ZERO\_IN\_DATE, 606
  - ONLY\_FULL\_GROUP\_BY, 606, 1288
  - ORACLE, 609
  - PAD\_CHAR\_TO\_FULL\_LENGTH, 606
  - PIPES\_AS\_CONCAT, 607
  - POSTGRESQL, 609
  - REAL\_AS\_FLOAT, 607
  - strict, 603
  - STRICT\_ALL\_TABLES, 607
  - STRICT\_TRANS\_TABLES, 603, 607
  - TRADITIONAL, 603, 609
- SQL scripts, 244
- SQL statements
  - replication masters, 1474
  - replication slaves, 1476
- SQL-92
  - extensions to, 26
- sql-mode option
  - mysqld, 442
- SQLState error codes, 2671
- sql\_auto\_is\_null session variable, 543
- SQL\_BIG\_RESULT, 1432
- sql\_big\_selects session variable, 543
- SQL\_BUFFER\_RESULT, 1433
- sql\_buffer\_result session variable, 544
- SQL\_CACHE, 920, 1433
- SQL\_CALC\_FOUND\_ROWS, 1433
- sql\_log\_bin system variable, 544
- sql\_log\_off system variable, 544
- sql\_mode system variable, 545
- sql\_notes session variable, 546
- SQL\_NO\_CACHE, 921, 1433
- sql\_quote\_show\_create session variable, 546
- sql\_safe\_updates session variable, 547
- sql\_select\_limit system variable, 547
- sql\_slave\_skip\_counter, 1481
- sql\_slave\_skip\_counter system variable, 1996
- SQL\_SMALL\_RESULT, 1432
- SQL\_THREAD\_WAIT\_AFTER\_GTIDS(), 1274
- sql\_warnings session variable, 547
- SQRT(), 1208
- square brackets, 1112
- srcdir option
  - mysql\_install\_db, 238
- SRID(), 1304
- SSD, 656, 3820
- SSH, 746, 817
- SSL, 806, 2523, 2667

---

---

- command options, 809
- configuring, 807
- establishing connections, 807
  - X509 Basics, 806
- SSL certificates, 2523
- ssl option, 810
- SSL options, 212
  - mysql, 253
  - mysqladmin, 272
  - mysqlcheck, 278
  - mysqld, 440
  - mysqldump, 285
  - mysqlimport, 301
  - mysqlshow, 306
  - mysqslap, 313
- SSL related options, 1532
- ssl-ca option, 810
- ssl-capath option, 810
- ssl-cert option, 811
- ssl-cipher option, 811
- ssl-crl option, 811
- ssl-crlpath option, 811
- ssl-key option, 811
- ssl-verify-server-cert option, 811
- ssl\_ca system variable, 547
- ssl\_capath system variable, 547
- ssl\_cert system variable, 548
- ssl\_cipher system variable, 548
- ssl\_crl system variable, 548
- ssl\_crlpath system variable, 548
- ssl\_key system variable, 549
- standalone option
  - mysqld, 440
- Standard Monitor
  - InnoDB, 1673
- Standard SQL
  - differences from, 30, 1533
  - extensions to, 26, 27
- standards compatibility, 26
- START
  - XA transactions, 1472
- START SLAVE, 1482
- START TRANSACTION, 1458
- start-datetime option
  - mysqlbinlog, 353
- start-position option
  - mysqlbinlog, 353
- starting
  - comments, 33
  - mysqld, 748
  - the server, 140
  - the server automatically, 146
- Starting many servers, 705
- StartPoint(), 1305
- startup, 3820
- startup options
  - default, 215
- startup parameters, 932
  - mysql, 244
  - mysqladmin, 268
  - tuning, 931
- start\_row option
  - mysql\_find\_rows, 370
- statefile option
  - comp\_err, 236
- statement-based replication, 3820
  - advantages, 1941
  - disadvantages, 1941
  - unsafe statements, 1941
- statements
  - compound, 1490
  - GRANT, 774
  - INSERT, 776
  - replication masters, 1474
  - replication slaves, 1476
- Statistics
  - thread command, 959
- statistics, 3820
  - thread state, 964
- STATISTICS
  - INFORMATION\_SCHEMA table, 2198
- stats option
  - myisam\_ftdump, 315
- stats\_method myisamchk variable, 319
- status
  - tables, 1581
- status command
  - mysql, 258
  - results, 268
- status logs (replication), 2028
- status option
  - mysqlshow, 306
- status variable
  - Rpl\_semi\_sync\_master\_clients, 596
  - Rpl\_semi\_sync\_master\_net\_avg\_wait\_time, 596
  - Rpl\_semi\_sync\_master\_net\_waits, 596
  - Rpl\_semi\_sync\_master\_net\_wait\_time, 596
  - Rpl\_semi\_sync\_master\_no\_times, 596
  - Rpl\_semi\_sync\_master\_no\_tx, 596
  - Rpl\_semi\_sync\_master\_status, 596
  - Rpl\_semi\_sync\_master\_timefunc\_failures, 597
  - Rpl\_semi\_sync\_master\_tx\_avg\_wait\_time, 597
  - Rpl\_semi\_sync\_master\_tx\_waits, 597
  - Rpl\_semi\_sync\_master\_tx\_wait\_time, 597
  - Rpl\_semi\_sync\_master\_wait\_pos\_backtraverse, 597
  - Rpl\_semi\_sync\_master\_wait\_sessions, 597
  - Rpl\_semi\_sync\_master\_yes\_tx, 597
  - Rpl\_semi\_sync\_slave\_status, 597
- status variables, 576, 1580
- STD(), 1284
- STDDEV(), 1285
- STDDEV\_POP(), 1285
- STDDEV\_SAMP(), 1285
- stemming, 3821
- STOP SLAVE, 1485
- stop-datetime option

---

---

- mysqlbinlog, 353
- stop-never option
  - mysqlbinlog, 354
- stop-never-slave-server-id option
  - mysqlbinlog, 354
- stop-position option
  - mysqlbinlog, 354
- stopping
  - the server, 146
- stopword, 3821
- stopword list
  - user-defined, 1241
- stopwords, 1237
- storage engine, 3821
  - ARCHIVE, 1801
  - InnoDB, 1605
  - PERFORMANCE\_SCHEMA, 2233
- storage engine plugins, 3487
- storage engines
  - choosing, 1601
  - InnoDB as default, 1606
- storage requirements
  - data type, 1147
- storage space
  - minimizing, 878
- storage\_engine system variable, 549
- stored functions, 2153
  - and INSERT DELAYED, 1406
- stored procedures, 2153
  - calling from Connector/Net, 2451
  - editing with Visual Studio, 2423
  - with Connector/Net, 2496
- stored programs, 1490, 2151
  - reparsing, 924
- stored routines
  - and replication, 2067
  - LAST\_INSERT\_ID(), 2155
  - metadata, 2154
  - restrictions, 3757
- storing result in query cache
  - thread state, 966
- storing row into queue
  - thread state, 965
- STRAIGHT\_JOIN, 894, 905, 982, 982, 1432, 1435
- STRCMP(), 1192
- strict mode, 1690, 3821
- strict SQL mode, 603
- STRICT\_ALL\_TABLES SQL mode, 607
- STRICT\_TRANS\_TABLES SQL mode, 603, 607
- string collating, 1089
- string comparison functions, 1190
- string comparisons
  - case sensitivity, 1190
- string concatenation, 1011, 1178
- string functions, 1175
- string literal introducer, 1012, 1043
- string replacement
  - replace utility, 376
- string types, 1137, 1149
- strings
  - defined, 1011
  - escape sequences, 1011
  - nondelimited, 1015
- striping
  - defined, 940
- STR\_TO\_DATE(), 1221
- ST\_Contains(), 1310
- ST\_Crosses(), 1310
- ST\_Disjoint(), 1311
- ST\_Equals(), 1311
- ST\_Intersects(), 1311
- ST\_Overlaps(), 1311
- ST\_Touches(), 1311
- ST\_Within(), 1311
- SUBDATE(), 1222
- sublist, 3821
- SUBPARTITION BY KEY
  - known issues, 2143
- subpartitioning, 2108
- subpartitions, 2108
  - known issues, 2143
- subqueries, 1445
  - correlated, 1450
  - errors, 1452
  - optimization, 1005
  - restrictions, 3761
  - rewriting as joins, 1455
  - with ALL, 1448
  - with ANY, IN, SOME, 1447
  - with EXISTS, 1449
  - with NOT EXISTS, 1449
  - with ROW, 1448
- subquery (see subqueries)
- subquery materialization, 1003, 1003
- subquery optimization, 1001
- subselects, 1445
- SUBSTR(), 1185
- SUBSTRING(), 1185
- SUBSTRING\_INDEX(), 1185
- SUBTIME(), 1223
- subtraction (-), 1200
- suffix option
  - mysqlhotcopy, 368
- SUM(), 1285
- SUM(DISTINCT), 1285
- super-large-pages option
  - mysqld, 440
- superuser, 150
- superuser option
  - mysqlaccess, 345
- support
  - for operating systems, 49
- suppression
  - default values, 36
- supremum record, 3821
- surrogate key, 3821

---

---

symbolic links, 940, 942  
symbolic-links option  
  mysqld, 441  
symbols-file option  
  resolve\_stack\_dump, 375  
SymDifference(), 1309  
sync\_binlog system variable, 2017  
sync\_frm system variable, 550  
sync\_master\_info system variable, 1997  
sync\_relay\_log system variable, 1998  
sync\_relay\_log\_info system variable, 1999  
syntax  
  regular expression, 1193  
syntax conventions, 3  
synthetic key, 3821  
SYSCONFDIR option  
  CMake, 133  
SYSDATE(), 1223  
sysdate-is-now option  
  mysqld, 444  
syslog option  
  mysqld\_safe, 229  
syslog-tag option  
  mysqld\_safe, 229  
system  
  privilege, 751  
  security, 732  
system command  
  mysql, 258  
System lock  
  thread state, 964  
system optimization, 931  
system table  
  optimizer, 898, 1432  
system tablespace, 3822  
system variable  
  audit\_log\_buffer\_size, 824  
  audit\_log\_file, 825  
  audit\_log\_flush, 825  
  audit\_log\_policy, 826  
  audit\_log\_rotate\_on\_size, 826  
  audit\_log\_strategy, 827  
  authentication\_windows\_log\_level, 463  
  authentication\_windows\_use\_principal\_name, 464  
  automatic\_sp\_privileges, 464  
  auto\_increment\_increment, 1965  
  auto\_increment\_offset, 1968  
  back\_log, 465  
  basedir, 466  
  bind\_address, 466  
  binlog\_cache\_size, 2008  
  binlog\_checksum, 2009  
  binlog\_direct\_non\_transactional\_updates, 2010  
  binlog\_format, 2010  
  binlog\_max\_flush\_queue\_time, 2011  
  binlog\_order\_commits, 2012  
  binlog\_rows\_query\_log\_events, 2014  
  binlog\_row\_image, 2012  
  binlog\_stmt\_cache\_size, 2014  
  bulk\_insert\_buffer\_size, 466  
  character\_sets\_dir, 469  
  character\_set\_client, 467  
  character\_set\_connection, 467  
  character\_set\_database, 468  
  character\_set\_filesystem, 468  
  character\_set\_results, 468  
  character\_set\_server, 468  
  character\_set\_system, 469  
  collation\_connection, 469  
  collation\_database, 469  
  collation\_server, 470  
  completion\_type, 470  
  concurrent\_insert, 471  
  connect\_timeout, 472  
  core\_file, 472  
  datadir, 472  
  datetime\_format, 472  
  date\_format, 472  
  debug, 473  
  debug\_sync, 473  
  default\_storage\_engine, 474  
  default\_tmp\_storage\_engine, 474  
  default\_week\_format, 475  
  delayed\_insert\_limit, 476  
  delayed\_insert\_timeout, 476  
  delayed\_queue\_size, 477  
  delay\_key\_write, 475  
  div\_precision\_increment, 478  
  end\_markers\_in\_json, 479  
  engine\_condition\_pushdown, 478  
  event\_scheduler, 480  
  expire\_logs\_days, 480  
  explicit\_defaults\_for\_timestamp, 481  
  flush, 482  
  flush\_time, 482  
  ft\_boolean\_syntax, 483  
  ft\_max\_word\_len, 484  
  ft\_min\_word\_len, 484  
  ft\_query\_expansion\_limit, 484  
  ft\_stopword\_file, 485  
  general\_log, 485  
  general\_log\_file, 486  
  group\_concat\_max\_len, 486  
  gtid\_done, 2020  
  gtid\_executed, 2021  
  gtid\_lost, 2022  
  gtid\_purged, 2023  
  have\_compress, 486  
  have\_crypt, 486  
  have\_csv, 486  
  have\_dynamic\_loading, 487  
  have\_geometry, 487  
  have\_innodb, 487  
  have\_openssl, 487  
  have\_partitioning, 487  
  have\_profiling, 487



---

have\_query\_cache, 487  
have\_rtree\_keys, 487  
have\_ssl, 487  
have\_symlink, 487  
hostname, 488  
ignore\_builtin\_innodb, 1700  
ignore\_db\_dirs, 488  
init\_connect, 489  
init\_file, 489  
init\_slave, 1987  
innodb\_adaptive\_flushing, 1700  
innodb\_adaptive\_hash\_index, 1701  
innodb\_additional\_mem\_pool\_size, 1702  
innodb\_autoextend\_increment, 1704  
innodb\_autoinc\_lock\_mode, 1705  
innodb\_buffer\_pool\_instances, 1706  
innodb\_buffer\_pool\_size, 1708  
innodb\_change\_buffering, 1709  
innodb\_checksums, 1711  
innodb\_commit\_concurrency, 1712  
innodb\_concurrency\_tickets, 1714  
innodb\_data\_file\_path, 1714  
innodb\_data\_home\_dir, 1715  
innodb\_disable\_sort\_file\_cache, 1715  
innodb\_doublewrite, 1715  
innodb\_fast\_shutdown, 1716  
innodb\_file\_format, 1716  
innodb\_file\_format\_check, 1717  
innodb\_file\_format\_max, 1717  
innodb\_file\_per\_table, 1718  
innodb\_flush\_log\_at\_timeout, 1718  
innodb\_flush\_log\_at\_trx\_commit, 1718  
innodb\_flush\_method, 1719  
innodb\_force\_recovery, 1722  
innodb\_io\_capacity, 1727  
innodb\_locks\_unsafe\_for\_binlog, 1729  
innodb\_lock\_wait\_timeout, 1729  
innodb\_log\_buffer\_size, 1732  
innodb\_log\_files\_in\_group, 1733  
innodb\_log\_file\_size, 1732  
innodb\_log\_group\_home\_dir, 1733  
innodb\_max\_dirty\_pages\_pct, 1734  
innodb\_max\_purge\_lag, 1735  
innodb\_max\_purge\_lag\_delay, 1736  
innodb\_mirrored\_log\_groups, 1736  
innodb\_old\_blocks\_pct, 1738  
innodb\_old\_blocks\_time, 1738  
innodb\_open\_files, 1739  
innodb\_purge\_batch\_size, 1741  
innodb\_purge\_threads, 1742  
innodb\_read\_ahead\_threshold, 1742  
innodb\_read\_io\_threads, 1743  
innodb\_replication\_delay, 1744  
innodb\_rollback\_on\_timeout, 1744  
innodb\_spin\_wait\_delay, 1745  
innodb\_stats\_method, 1746  
innodb\_stats\_on\_metadata, 1747  
innodb\_stats\_sample\_pages, 1748  
innodb\_strict\_mode, 1749  
innodb\_support\_xa, 1749  
innodb\_sync\_spin\_loops, 1750  
innodb\_table\_locks, 1751  
innodb\_thread\_concurrency, 1751  
innodb\_thread\_sleep\_delay, 1752  
innodb\_use\_native\_aio, 1753  
innodb\_use\_sys\_malloc, 1754  
innodb\_version, 1754  
innodb\_write\_io\_threads, 1754  
interactive\_timeout, 490  
join\_buffer\_size, 490  
keep\_files\_on\_create, 491  
key\_buffer\_size, 492  
key\_cache\_age\_threshold, 493  
key\_cache\_block\_size, 493  
key\_cache\_division\_limit, 494  
large\_files\_support, 494  
large\_pages, 494  
large\_page\_size, 494  
lc\_messages, 495  
lc\_messages\_dir, 495  
lc\_time\_names, 495  
license, 496  
local\_infile, 496  
locked\_in\_memory, 497  
lock\_wait\_timeout, 496  
log, 497  
log\_bin, 2008  
log\_bin\_basename, 2015  
log\_bin\_trust\_function\_creators, 497  
log\_bin\_use\_v1\_row\_events, 2015  
log\_error, 497  
log\_output, 498  
log\_queries\_not\_using\_indexes, 498  
log\_slave\_updates, 2008  
log\_slow\_queries, 499  
log\_throttle\_queries\_not\_using\_indexes, 499  
log\_warnings, 499  
long\_query\_time, 500  
lower\_case\_file\_system, 501  
lower\_case\_table\_names, 501  
low\_priority\_updates, 500  
master\_info\_repository, 1988  
master\_verify\_checksum, 2015  
max\_allowed\_packet, 502  
max\_binlog\_cache\_size, 2016  
max\_binlog\_size, 2016  
max\_binlog\_stmt\_cache\_size, 2017  
max\_connections, 503  
max\_connect\_errors, 502  
max\_delayed\_threads, 504  
max\_error\_count, 504  
max\_heap\_table\_size, 504  
max\_insert\_delayed\_threads, 505  
max\_join\_size, 505  
max\_length\_for\_sort\_data, 506  
max\_prepared\_stmt\_count, 506

---

---

max\_relay\_log\_size, 507  
max\_seeks\_for\_key, 507  
max\_sort\_length, 508  
max\_sp\_recursion\_depth, 508  
max\_tmp\_tables, 508  
max\_user\_connections, 509  
max\_write\_lock\_count, 509  
metadata\_locks\_cache\_size, 510  
metadata\_locks\_hash\_instances, 510  
min\_examined\_row\_limit, 510  
mysam\_data\_pointer\_size, 511  
mysam\_max\_sort\_file\_size, 511  
mysam\_mmap\_size, 512  
mysam\_recover\_options, 512  
mysam\_repair\_threads, 512  
mysam\_sort\_buffer\_size, 513  
mysam\_stats\_method, 513  
mysam\_use\_mmap, 514  
named\_pipe, 514  
net\_buffer\_length, 515  
net\_read\_timeout, 515  
net\_retry\_count, 515  
net\_write\_timeout, 516  
new, 516  
old, 517  
old\_alter\_table, 517  
old\_passwords, 517  
open\_files\_limit, 518  
optimizer\_join\_cache\_level, 519  
optimizer\_prune\_level, 519  
optimizer\_search\_depth, 520  
optimizer\_switch, 520  
optimizer\_trace, 523  
optimizer\_trace\_features, 523  
optimizer\_trace\_limit, 524  
optimizer\_trace\_max\_mem\_size, 524  
optimizer\_trace\_offset, 524  
performance\_schema, 2307  
performance\_schema\_accounts\_size, 2308  
performance\_schema\_digests\_size, 2308  
performance\_schema\_events\_stages\_history\_long\_size, 2308  
performance\_schema\_events\_stages\_history\_size, 2309  
performance\_schema\_events\_statements\_history\_long\_size, 2309  
performance\_schema\_events\_statements\_history\_size, 2309  
performance\_schema\_events\_waits\_history\_long\_size, 2310  
performance\_schema\_events\_waits\_history\_size, 2310  
performance\_schema\_hosts\_size, 2310  
performance\_schema\_max\_cond\_classes, 2311  
performance\_schema\_max\_cond\_instances, 2311  
performance\_schema\_max\_file\_classes, 2311  
performance\_schema\_max\_file\_handles, 2312  
performance\_schema\_max\_file\_instances, 2312  
performance\_schema\_max\_mutex\_classes, 2312  
performance\_schema\_max\_mutex\_instances, 2313  
performance\_schema\_max\_rwlock\_classes, 2313  
performance\_schema\_max\_rwlock\_instances, 2313  
performance\_schema\_max\_socket\_classes, 2314  
performance\_schema\_max\_socket\_instances, 2314  
performance\_schema\_max\_stage\_classes, 2314  
performance\_schema\_max\_statement\_classes, 2314  
performance\_schema\_max\_table\_handles, 2315  
performance\_schema\_max\_table\_instances, 2315  
performance\_schema\_max\_thread\_classes, 2316  
performance\_schema\_max\_thread\_instances, 2316  
performance\_schema\_session\_connect\_attrs\_size, 2316  
performance\_schema\_setup\_actors\_size, 2317  
performance\_schema\_setup\_objects\_size, 2317  
performance\_schema\_users\_size, 2317  
pid\_file, 524  
plugin\_dir, 525  
port, 525  
preload\_buffer\_size, 525  
protocol\_version, 526  
pseudo\_slave\_mode, 526  
pseudo\_thread\_id, 527  
query\_alloc\_block\_size, 527  
query\_cache\_limit, 527  
query\_cache\_min\_res\_unit, 528  
query\_cache\_size, 528  
query\_cache\_type, 529  
query\_cache\_wlock\_invalidate, 530  
query\_prealloc\_size, 531  
range\_alloc\_block\_size, 531  
read\_buffer\_size, 532  
read\_only, 532  
read\_rnd\_buffer\_size, 533  
relay\_log, 1989  
relay\_log\_basename, 1989  
relay\_log\_index, 1989  
relay\_log\_info\_file, 1989  
relay\_log\_info\_repository, 1990  
relay\_log\_purge, 534  
relay\_log\_recovery, 1990  
relay\_log\_space\_limit, 534  
report\_host, 535  
report\_password, 535  
report\_port, 535  
report\_user, 536  
rpl\_semi\_sync\_master\_enabled, 536  
rpl\_semi\_sync\_master\_timeout, 536  
rpl\_semi\_sync\_master\_trace\_level, 536  
rpl\_semi\_sync\_master\_wait\_no\_slave, 537  
rpl\_semi\_sync\_slave\_enabled, 537  
rpl\_semi\_sync\_slave\_trace\_level, 538  
secure\_auth, 538  
secure\_file\_priv, 538  
server\_id, 539  
sha256\_password\_private\_key\_path, 539

---

sha256\_password\_public\_key\_path, 539  
 shared\_memory, 540  
 shared\_memory\_base\_name, 540  
 skip\_external\_locking, 540  
 skip\_name\_resolve, 540  
 skip\_networking, 540  
 skip\_show\_database, 540  
 slave\_checkpoint\_group, 1991  
 slave\_checkpoint\_period, 1991  
 slave\_compressed\_protocol, 1992  
 slave\_exec\_mode, 1992  
 slave\_load\_tmpdir, 1992  
 slave\_max\_allowed\_packet, 1969  
 slave\_net\_timeout, 1993  
 slave\_parallel\_workers, 1993  
 slave\_pending\_jobs\_size\_max, 1994  
 slave\_rows\_search\_algorithms, 1969  
 slave\_skip\_errors, 1994  
 slave\_sql\_verify\_checksum, 1995  
 slave\_transaction\_retries, 1995  
 slave\_type\_conversions, 1996  
 slow\_launch\_time, 541  
 slow\_query\_log, 541  
 slow\_query\_log\_file, 541  
 socket, 542  
 sort\_buffer\_size, 542  
 sql\_log\_bin, 544  
 sql\_log\_off, 544  
 sql\_mode, 545  
 sql\_select\_limit, 547  
 sql\_slave\_skip\_counter, 1996  
 ssl\_ca, 547  
 ssl\_capath, 547  
 ssl\_cert, 548  
 ssl\_cipher, 548  
 ssl\_crl, 548  
 ssl\_crlpath, 548  
 ssl\_key, 549  
 storage\_engine, 549  
 sync\_binlog, 2017  
 sync\_frm, 550  
 sync\_master\_info, 1997  
 sync\_relay\_log, 1998  
 sync\_relay\_log\_info, 1999  
 system\_time\_zone, 550  
 sysvar\_stored\_program\_cache, 549  
 table\_definition\_cache, 550  
 table\_open\_cache, 551  
 table\_open\_cache\_instances, 552  
 thread\_cache\_size, 552  
 thread\_concurrency, 553  
 thread\_handling, 553  
 thread\_pool\_algorithm, 554  
 thread\_pool\_high\_priority\_connection, 554  
 thread\_pool\_max\_unused\_threads, 555  
 thread\_pool\_prio\_kickup\_timer, 555  
 thread\_pool\_size, 556  
 thread\_pool\_stall\_limit, 556  
 thread\_stack, 556  
 timed\_mutexes, 557  
 time\_format, 557  
 time\_zone, 557  
 tmpdir, 559  
 tmp\_table\_size, 558  
 transaction\_alloc\_block\_size, 559  
 transaction\_prealloc\_size, 560  
 tx\_isolation, 560  
 tx\_read\_only, 561  
 updatable\_views\_with\_limit, 562  
 validate\_password\_dictionary\_file, 743  
 validate\_password\_length, 744  
 validate\_password\_mixed\_case\_count, 744  
 validate\_password\_number\_count, 745  
 validate\_password\_policy, 745  
 validate\_password\_special\_char\_count, 745  
 validate\_user\_plugins, 562  
 version, 563  
 version\_comment, 563  
 version\_compile\_machine, 563  
 version\_compile\_os, 563  
 wait\_timeout, 563  
 system variables, 447, 564, 1584  
   and replication, 2078  
   disable\_gtid\_unsafe\_statements, 2020  
   enforce\_gtid\_consistency, 2021  
   gtid\_mode, 2022  
   gtid\_next, 2022  
   gtid\_owned, 2023  
 system\_time\_zone system variable, 550  
 SYSTEM\_USER(), 1273  
 sysvar\_stored\_program\_cache system variable, 549

## T

tab (t), 1012, 1417  
 tab option  
   mysqldump, 291  
 table, 3822  
   changing, 1330, 1335, 3750  
   deleting, 1393  
   rebuilding, 163  
   repair, 163  
   row size, 1147  
 table aliases, 1429  
 table cache, 881, 2494  
 table description  
   myisamchk, 324  
 Table Designer, 2417  
 Table Dump  
   thread command, 959  
 table is full, 466, 3732  
 table lock, 3823  
 Table Monitor  
   InnoDB, 1673, 1685  
 table names  
   case sensitivity, 27, 1020  
 table option

---

---

- mysql, 254
- mysqlaccess, 345
- table scan, 1659
- table type, 3823
  - choosing, 1601
- table-level locking, 925
- TableDirect commands, 2523
- tables
  - BLACKHOLE, 1802
  - checking, 320
  - closing, 881
  - compressed, 332
  - compressed format, 1793
  - const, 898
  - constant, 859
  - copying, 1376, 1376
  - counting rows, 188
  - creating, 177
  - CSV, 1799
  - defragment, 1792
  - defragmenting, 853, 1540
  - deleting rows, 3747
  - displaying, 302
  - displaying status, 1581
  - dumping, 279, 365
  - dynamic, 1792
  - error checking, 849
  - EXAMPLE, 1814
  - FEDERATED, 1809
  - flush, 268
  - fragmentation, 1540
  - HEAP, 1795
  - improving performance, 878
  - information, 324
  - information about, 191
  - InnoDB, 1605
  - loading data, 178
  - maintenance, 272
  - maintenance schedule, 853
  - maximum size, 3768
  - MEMORY, 1795
  - MERGE, 1804
  - merging, 1804
  - multiple, 190
  - MyISAM, 1787
  - names, 1017
  - open, 881
  - opening, 881
  - optimizing, 852
  - partitioning, 1804
  - repair, 272
  - repairing, 850
  - retrieving data, 179
  - selecting columns, 181
  - selecting rows, 180
  - sorting rows, 182
  - symbolic links, 941
  - system, 898
  - too many, 882
  - unique ID for last row, 2896
- TABLES
  - INFORMATION\_SCHEMA table, 2199
- tables option
  - mysqlcheck, 278
  - mysqldump, 293
- TABLESPACE
  - INFORMATION\_SCHEMA table, 2200
- tablespace, 3823
- tablespace dictionary, 3823
- Tablespace Monitor
  - InnoDB, 639, 1618, 1673
- table\_definition\_cache system variable, 550
- table\_io\_waits\_summary\_by\_index\_usage table
  - performance\_schema, 2294
- table\_io\_waits\_summary\_by\_table table
  - performance\_schema, 2293
- table\_lock\_waits\_summary\_by\_table table
  - performance\_schema, 2295
- table\_open\_cache, 881
- table\_open\_cache system variable, 551
- table\_open\_cache\_instances system variable, 552
- TABLE\_PRIVILEGES
  - INFORMATION\_SCHEMA table, 2201
- TAN(), 1208
- tar
  - problems on Solaris, 118, 118
- tc-heuristic-recover option
  - mysqld, 444
- Tcl API, 3482
- tcp-ip option
  - mysqld\_multi, 233
- TCP/IP, 86, 95, 135, 165, 209, 228, 252, 352, 373, 437, 705, 746, 768, 946, 1981, 2273, 3723
- tee command
  - mysql, 258
- tee option
  - mysql, 254
- temp-pool option
  - mysqld, 445
- temporary file
  - write access, 145
- temporary files, 3741
- temporary table, 3823
- temporary tables
  - and replication, 2073
  - internal, 882
  - problems, 3750
- terminal monitor
  - defined, 171
- test option
  - myisampack, 333
- testing
  - connection to the server, 763
  - installation, 140
  - of MySQL releases, 52
  - postinstallation, 139

---

---

testing mysqld  
   mysqltest, 3484  
 TEXT  
   size, 1150  
 text collection, 3824  
 TEXT columns  
   default values, 1140  
   indexing, 873, 1363  
 TEXT data type, 1119, 1140  
 text files  
   importing, 262, 297, 1412  
 thread, 3824  
 thread cache, 946  
 thread command  
   Binlog Dump, 957  
   Change user, 957  
   Close stmt, 958  
   Connect, 958  
   Connect Out, 958  
   Create DB, 958  
   Daemon, 958  
   Debug, 958  
   Delayed insert, 958  
   Drop DB, 958  
   Error, 958  
   Execute, 958  
   Fetch, 958  
   Field List, 958  
   Init DB, 958  
   Kill, 958  
   Long Data, 958  
   Ping, 958  
   Prepare, 958  
   Processlist, 959  
   Query, 959  
   Quit, 959  
   Refresh, 959  
   Register Slave, 959  
   Reset stmt, 959  
   Set option, 959  
   Shutdown, 959  
   Sleep, 959  
   Statistics, 959  
   Table Dump, 959  
   Time, 959  
 thread commands, 957  
 thread state  
   After create, 959  
   allocating local table, 965  
   altering table, 959  
   Analyzing, 960  
   Changing master, 969  
   Checking master version, 967  
   checking permissions, 960  
   checking privileges on cached query, 966  
   checking query cache for query, 966  
   Checking table, 960  
   cleaning up, 960  
   Clearing, 969  
   closing tables, 960  
   committing alter table to storage engine, 960  
   Connecting to master, 967  
   converting HEAP to MyISAM, 960  
   copy to tmp table, 960  
   Copying to group table, 960  
   Copying to tmp table, 960  
   Copying to tmp table on disk, 960  
   Creating delayed handler, 965  
   Creating index, 960  
   Creating sort index, 960  
   creating table, 960  
   Creating tmp table, 960  
   deleting from main table, 961  
   deleting from reference tables, 961  
   discard\_or\_import\_tablespace, 961  
   end, 961  
   executing, 961  
   Execution of init\_command, 961  
   Finished reading one binlog; switching to next binlog, 967  
   Flushing tables, 961  
   freeing items, 961  
   FULLTEXT initialization, 961  
   got handler lock, 965  
   got old table, 965  
   init, 961  
   Initialized, 969  
   insert, 966  
   invalidating query cache entries, 966  
   Killed, 961  
   Killing slave, 969  
   logging slow query, 962  
   login, 962  
   Making temp file, 968  
   manage keys, 962  
   Master has sent all binlog to slave; waiting for binlog to be updated, 967  
   NULL, 962  
   Opening master dump table, 969  
   Opening table, 962  
   Opening tables, 962  
   optimizing, 962  
   preparing, 962  
   preparing for alter table, 962  
   Purging old relay logs, 962  
   query end, 962  
   Queueing master event to the relay log, 968  
   Reading event from the relay log, 968  
   Reading from net, 962  
   Reading master dump table data, 969  
   Rebuilding the index on master dump table, 969  
   Reconnecting after a failed binlog dump request, 967  
   Reconnecting after a failed master event read, 968  
   Registering slave on master, 967  
   Removing duplicates, 962  
   removing tmp table, 962

---

---

- rename, 962
- rename result table, 963
- Reopen tables, 963
- Repair by sorting, 963
- Repair done, 963
- Repair with keycache, 963
- Requesting binlog dump, 967
- reschedule, 966
- Rolling back, 963
- Saving state, 963
- Searching rows for update, 963
- Sending binlog event to slave, 967
- sending cached result to client, 966
- setup, 963
- Slave has read all relay log; waiting for the slave I/O thread to update it, 968
- Sorting for group, 963
- Sorting for order, 963
- Sorting index, 963
- Sorting result, 963
- statistics, 964
- storing result in query cache, 966
- storing row into queue, 965
- System lock, 964
- update, 964
- Updating, 964
- updating main table, 964
- updating reference tables, 964
- upgrading lock, 966
- User lock, 964
- User sleep, 964
- Waiting for commit lock, 964
- waiting for delay\_list, 965
- Waiting for global read lock, 964, 964
- waiting for handler insert, 966
- waiting for handler lock, 966
- waiting for handler open, 966
- Waiting for INSERT, 966
- Waiting for master to send event, 968
- Waiting for master update, 967
- Waiting for next activation, 969
- Waiting for query cache lock, 966
- Waiting for scheduler to stop, 969
- Waiting for schema metadata lock, 964
- Waiting for slave mutex on exit, 968, 968
- Waiting for stored function metadata lock, 964
- Waiting for stored procedure metadata lock, 964
- Waiting for table, 964
- Waiting for table level lock, 964
- Waiting for table metadata lock, 964
- Waiting for tables, 964
- Waiting for the next event in relay log, 968
- Waiting for the slave SQL thread to free enough relay log space, 968
- Waiting for trigger metadata lock, 964
- Waiting on cond, 965
- Waiting on empty queue, 969
- Waiting to finalize termination, 967
- Waiting to reconnect after a failed binlog dump request, 967
- Waiting to reconnect after a failed master event read, 968
- Waiting until MASTER\_DELAY seconds after master executed event, 968
- Writing to net, 965
- thread states, 956
  - delayed inserts, 965
  - event scheduler, 969
  - general, 959
  - query cache, 966
  - replication master, 967
  - replication slave, 967, 968, 969
- thread support, 49
- thread table
  - performance\_schema, 2302
- threaded clients, 2799
- threads, 268, 1568, 3483
  - display, 1568
  - monitoring, 956, 1568, 2194, 2302
- thread\_cache\_size system variable, 552
- thread\_concurrency system variable, 553
- thread\_handling system variable, 553
- thread\_pool\_algorithm system variable, 554
- thread\_pool\_high\_priority\_connection system variable, 554
- thread\_pool\_max\_unused\_threads system variable, 555
- thread\_pool\_prio\_kickup\_timer system variable, 555
- thread\_pool\_size system variable, 556
- thread\_pool\_stall\_limit system variable, 556
- thread\_stack system variable, 556
- Time
  - thread command, 959
- TIME columns
  - with Connector/Net, 2505
- TIME data type, 1116, 1127
- time literals, 1014
- time representation
  - Event Scheduler, 2162
- time types, 1149
- time zone problems, 3742
- time zone tables, 241
- time zones
  - and replication, 2075
  - leap seconds, 1105
  - support, 1102
  - upgrading, 1104
- TIME(), 1223
- TIMEDIFF(), 1223
- timed\_mutexes system variable, 557
- timeout, 472, 1275, 1410
  - connect\_timeout variable, 255, 272
  - shutdown\_timeout variable, 272
- timeouts (replication), 2075
- TIMESTAMP
  - and NULL values, 3746

---

---

- and replication, 2058
- initialization and updating, 1131
- TIMESTAMP data type, 1116, 1126
- timestamp session variable, 558
- TIMESTAMP(), 1224
- TIMESTAMPADD(), 1224
- TIMESTAMPDIFF(), 1224
- timezone option
  - mysqld\_safe, 229
- time\_format system variable, 557
- TIME\_FORMAT(), 1224
- TIME\_TO\_SEC(), 1225
- time\_zone system variable, 557
- TINYBLOB data type, 1119
- TINYINT data type, 1112
- TINYTEXT data type, 1119
- tips
  - optimization, 868
- TMPDIR environment variable, 145, 166, 208, 3741
- tmpdir option
  - myisamchk, 323
  - myisampack, 333
  - mysqld, 446
  - mysqlhotcopy, 368
  - mysql\_upgrade, 243
- tmpdir system variable, 559
- tmp\_table\_size system variable, 558
- to-last-log option
  - mysqlbinlog, 354
- Tomcat application server, 2691
- tools
  - command-line, 83, 244
  - list of, 44
  - mysqld\_multi, 231
  - mysqld\_safe, 225
- torn page, 638, 3824
- Touches(), 1312
- TO\_BASE64(), 1186
- TO\_DAYS(), 1225
- TO\_SECONDS(), 1226
- TPS, 3824
- trace DBI method, 3547
- tracing
  - with Connector/Net, 2508
- TRADITIONAL SQL mode, 603, 609
- transaction, 3824
- transaction access mode, 1468
- transaction ID, 3824
- transaction isolation level, 1468
  - READ COMMITTED, 1469
  - READ UNCOMMITTED, 1470
  - REPEATABLE READ, 1469
  - SERIALIZABLE, 1470
- transaction-isolation option
  - mysqld, 445
- transaction-read-only option
  - mysqld, 445
- transaction-safe tables, 30, 1605
- transactions, 1623
  - and replication, 2075, 2076
  - isolation levels, 1623, 4091
  - metadata locking, 929
  - support, 30, 1605
- transaction\_alloc\_block\_size system variable, 559
- transaction\_prealloc\_size system variable, 560
- Translators
  - list of, 42
- transportable tablespace, 644, 3824
- trigger
  - restrictions, 3757
- triggers, 1384, 1394, 1583, 2151, 2155
  - and INSERT DELAYED, 1406
  - and replication, 2067, 2077
  - editing with Visual Studio, 2425
  - LAST\_INSERT\_ID(), 2155
  - metadata, 2158
- TRIGGERS
  - INFORMATION\_SCHEMA table, 2202
- triggers option
  - mysqldump, 293
- TRIM(), 1186
- troubleshooting, 3659, 3824
  - ALTER TABLE problems, 3750
  - C API, 2895
  - compiling MySQL server, 137
  - connection problems, 768
  - Connector/J, 2704
  - Connector/ODBC connection problems, 2352
  - InnoDB deadlocks, 1633
  - InnoDB errors, 1640
  - InnoDB recovery problems, 1684
  - InnoDB table fragmentation, 640
  - JDBC, 2639
  - JDBC SQLState codes, 2671
  - OS error codes, 1641
  - PHP, 3480
  - replication, 2082
  - startup problems, 148
  - with MySQL Enterprise Monitor, 3553
  - with MySQL Performance Schema, 2319
- TRUE, 1013, 1016
  - testing for, 1167, 1168
- truncate, 3824
- TRUNCATE TABLE, 1395, 1689
  - and replication, 2077
  - performance\_schema database, 2266, 3765
- TRUNCATE(), 1209
- tuning, 856
  - InnoDB compressed tables, 658
- tuple, 3825
- tutorial, 171
- two-phase commit, 589, 589, 1749, 3825
- tx\_isolation system variable, 560
- tx\_read\_only system variable, 561
- type codes
  - C prepared statement API, 2863

---

---

type conversions, 1161, 1166  
 type option  
   MySQLInstallerConsole, 84  
   mysql\_convert\_table\_format, 369  
 types  
   columns, 1111, 1151  
   data, 1111  
   date, 1149  
   Date and Time, 1124  
   numeric, 1148  
   of tables, 1601  
   portability, 1151  
   string, 1149  
   strings, 1137  
   time, 1149  
 Types namespace, 2623  
 typographical conventions, 3  
 TZ environment variable, 166, 3742  
 tz-utc option  
   mysqldump, 291

**U**

UCASE(), 1186  
 UCS-2, 1037  
 ucs2 character set, 1066  
 UDFs, 1544, 1545  
   compiling, 3539  
   defined, 3532  
   return values, 3539  
 ulimit, 3735  
 UMASK environment variable, 166, 3735  
 UMASK\_DIR environment variable, 166, 3736  
 unary minus (-), 1200  
 unbuffered option  
   mysql, 254  
 UNCOMPRESS(), 1265  
 UNCOMPRESSED\_LENGTH(), 1265  
 undo, 3825  
 undo log, 641, 1645, 3825  
 undo tablespace, 3825  
 UNHEX(), 1186  
 Unicode, 1037  
   with Connector/J, 2666  
 Unicode Collation Algorithm, 1075  
 UNINSTALL PLUGIN, 1546  
 uninstalling plugins, 609, 1546  
 UNION, 198, 1443  
 Union(), 1309  
 UNIQUE, 1335  
 unique constraint, 3825  
 unique ID, 2896  
 unique index, 3826  
 unique key, 3826  
   constraint, 35  
 unique keys  
   and partitioning keys, 2144  
 unique\_checks session variable, 561  
 unique\_subquery join type  
   optimizer, 899  
 Unix, 2326, 2403  
   compiling clients on, 2797  
 UNIX\_TIMESTAMP(), 1226  
 UNKNOWN  
   testing for, 1167, 1168  
 unloading  
   tables, 179  
 UNLOCK TABLES, 1462  
 unnamed views, 1450  
 unpack option  
   myisamchk, 323  
 unsafe statement (replication)  
   defined, 1945  
 unsafe statements (replication), 1946  
 UNSIGNED, 1112, 1120  
 UNTIL, 1497  
 updatable views, 2167  
 updatable\_views\_with\_limit system variable, 562  
 UPDATE, 30, 1455  
 update  
   thread state, 964  
 update-state option  
   myisamchk, 321  
 updates option  
   MySQLInstallerConsole, 83  
 UpdateXML(), 1251  
 updating  
   releases of MySQL, 53  
 Updating  
   thread state, 964  
 updating main table  
   thread state, 964  
 updating reference tables  
   thread state, 964  
 upgrade-system-tables option  
   mysql\_upgrade, 243  
 upgrading, 153, 154  
   different architecture, 164  
   to &current-series;, 155  
 upgrading lock  
   thread state, 966  
 upgrading MySQL, 241  
 UPPER(), 1187  
 uptime, 268  
 URLs for downloading MySQL, 53  
 USE, 1599  
 use command  
   mysql, 258  
 USE INDEX, 1441  
 USE KEY, 1441  
 use\_frm option  
   mysqlcheck, 279  
 use-mysqld\_safe option  
   mysql.server, 231  
 use-threads option  
   mysqlimport, 302  
 useConfigs connection properties, 2660

---



---

- user accounts
  - altering, 1522
  - creating, 1522
  - renaming, 1534
  - resource limits, 509, 778, 1531
- USER environment variable, 166, 212
- User lock
  - thread state, 964
- user names
  - and passwords, 773
  - in account names, 762
  - in default accounts, 150
- user option, 212
  - mysql, 254
  - mysql.server, 231
  - mysqlaccess, 345
  - mysqladmin, 272
  - mysqlbinlog, 354
  - mysqlcheck, 279
  - mysqld, 446
  - mysqldump, 285
  - mysqld\_multi, 233
  - mysqld\_safe, 229
  - mysqlhotcopy, 368
  - mysqlimport, 302
  - MySQLInstallerConsole, 84
  - mysqlshow, 306
  - mysqlslap, 313
  - mysql\_config\_editor, 343
  - mysql\_convert\_table\_format, 369
  - mysql\_install\_db, 238
  - mysql\_setpermission, 371
  - mysql\_upgrade, 243
- user privileges
  - adding, 774
  - deleting, 778, 1524
  - dropping, 778, 1524
- User sleep
  - thread state, 964
- user table
  - sorting, 765
- user variables, 1029
  - and replication, 2078
- USER(), 1273
- User-defined functions, 1544, 1545
- user-defined functions
  - adding, 3532, 3533
  - editing with Visual Studio, 2426
- users
  - adding, 144
  - deleting, 778, 1524
  - root, 150
- users table
  - performance\_schema, 2285
- USER\_PRIVILEGES
  - INFORMATION\_SCHEMA table, 2203
- using multiple disks to start data, 942
- UTC\_DATE(), 1227
- UTC\_TIME(), 1227
- UTC\_TIMESTAMP(), 1227
- UTF-8, 1037
- utf16 character set, 1066
- utf16le character set, 1066
- utf16\_bin collation, 1077
- utf32 character set, 1066
- utf8 character set, 1067
- utf8mb3 character set, 1067
- utf8mb4 character set, 1068
- utilities
  - program-development, 207
- utility programs, 206
- UUID(), 1280
- UUID\_SHORT(), 1281

**V**

- valid numbers
  - examples, 1013
- validate-password option
  - mysqld, 742
- validate\_password plugin, 741
  - configuring, 742
  - installing, 742
  - system variables, 742
- validate\_password\_dictionary\_file system variable, 743
- validate\_password\_length system variable, 744
- validate\_password\_mixed\_case\_count system variable, 744
- validate\_password\_number\_count system variable, 745
- validate\_password\_policy system variable, 745
- validate\_password\_special\_char\_count system variable, 745
- validate\_user\_plugins system variable, 562
- VALUES(), 1281
- VARBINARY data type, 1119, 1139
- VARCHAR
  - size, 1150
- VARCHAR data type, 1119, 1137
- VARCHARACTER data type, 1119
- variables
  - and replication, 2078
  - environment, 208
  - mysqld, 932
  - server, 1584
  - status, 576, 1580
  - system, 447, 564, 1584
  - user, 1029
- VARIANCE(), 1285
- VAR\_POP(), 1285
- VAR\_SAMP(), 1285
- verbose option
  - myisamchk, 319
  - myisampack, 333
  - myisam\_ftdump, 315
  - mysql, 254
  - mysqladmin, 272

---

- mysqlbinlog, 354
- mysqlcheck, 279
- mysqld, 447
- mysqldump, 287
- mysqldumpslow, 365
- mysqld\_multi, 233
- mysqlimport, 302
- mysqlshow, 306
- mysqslap, 313
- mysql\_config\_editor, 343
- mysql\_convert\_table\_format, 369
- mysql\_install\_db, 238
- mysql\_plugin, 240
- mysql\_upgrade, 243
- mysql\_waitpid, 371
- my\_print\_defaults, 375
- perror, 376
- verify-binlog-checksum option
  - mysqlbinlog, 354
- version
  - choosing, 50
  - latest, 53
- VERSION file
  - CMake, 138
- version option
  - comp\_err, 236
  - myisamchk, 319
  - myisampack, 333
  - mysql, 254
  - mysqlaccess, 345
  - mysqladmin, 272
  - mysqlbinlog, 354
  - mysqlcheck, 279
  - mysqld, 447
  - mysqldump, 287
  - mysqld\_multi, 233
  - mysqlimport, 302
  - mysqlshow, 306
  - mysqslap, 313
  - mysql\_config, 373
  - mysql\_config\_editor, 343
  - mysql\_convert\_table\_format, 369
  - mysql\_plugin, 240
  - mysql\_waitpid, 371
  - my\_print\_defaults, 375
  - perror, 376
  - resolveip, 377
  - resolve\_stack\_dump, 375
- version system variable, 563
- VERSION(), 1273
- version-check option
  - mysql\_upgrade, 244
- version\_comment system variable, 563
- version\_compile\_machine system variable, 563
- version\_compile\_os system variable, 563
- vertical option
  - mysql, 254
  - mysqladmin, 272

- victim, 3826
- Vietnamese, 3640
- view
  - restrictions, 3763
- views, 1388, 2151, 2165
  - algorithms, 2166
  - and replication, 2079
  - metadata, 2169
  - updatable, 1388, 2167
- VIEWS
  - INFORMATION\_SCHEMA table, 2204
- Views
  - limitations, 3764
  - privileges, 3764
  - problems, 3764
- Vision, 2397
- Visual Objects, 2395
- Visual Studio, 2414

## W

- wait, 3826
- wait option
  - myisamchk, 319
  - myisampack, 333
  - mysql, 254
  - mysqladmin, 272
- Waiting for commit lock
  - thread state, 964
- waiting for delay\_list
  - thread state, 965
- Waiting for event metadata lock
  - thread state, 964
- Waiting for event read lock
  - thread state, 964
- Waiting for global read lock
  - thread state, 964
- waiting for handler insert
  - thread state, 966
- waiting for handler lock
  - thread state, 966
- waiting for handler open
  - thread state, 966
- Waiting for INSERT
  - thread state, 966
- Waiting for master to send event
  - thread state, 968
- Waiting for master update
  - thread state, 967
- Waiting for next activation
  - thread state, 969
- Waiting for query cache lock
  - thread state, 966
- Waiting for scheduler to stop
  - thread state, 969
- Waiting for schema metadata lock
  - thread state, 964
- Waiting for slave mutex on exit
  - thread state, 968, 968

---

Waiting for stored function metadata lock  
   thread state, 964  
 Waiting for stored procedure metadata lock  
   thread state, 964  
 Waiting for table  
   thread state, 964  
 Waiting for table level lock  
   thread state, 964  
 Waiting for table metadata lock  
   thread state, 964  
 Waiting for tables  
   thread state, 964  
 Waiting for the next event in relay log  
   thread state, 968  
 Waiting for the slave SQL thread to free enough relay  
 log space  
   thread state, 968  
 Waiting for trigger metadata lock  
   thread state, 964  
 Waiting on cond  
   thread state, 965  
 Waiting on empty queue  
   thread state, 969  
 Waiting to finalize termination  
   thread state, 967  
 Waiting to reconnect after a failed binlog dump request  
   thread state, 967  
 Waiting to reconnect after a failed master event read  
   thread state, 968  
 Waiting until MASTER\_DELAY seconds after master  
 executed event  
   thread state, 968  
 wait\_timeout system variable, 563  
 WAIT\_UNTIL\_SQL\_THREAD\_AFTER\_GTIDS(), 1274  
 warm backup, 3826  
 warm up, 3826  
 warn option  
   mysql\_config\_editor, 343  
 warnings command  
   mysql, 258  
 warning\_count session variable, 564  
 Website Configuration Tool, 2437  
 WEEK(), 1228  
 WEEKDAY(), 1228  
 WEEKOFYEAR(), 1229  
 WEIGHT\_STRING(), 1187  
 Well-Known Binary format, 1297  
 Well-Known Text format, 1296  
 WHERE, 859  
   with SHOW, 2178, 2229  
 where option  
   mysqldump, 293  
 WHILE, 1497  
   labels, 1491  
 widths  
   display, 1111  
 Wildcard character (%), 1012  
 Wildcard character (\_), 1012

wildcards  
   and LIKE, 876  
   in account names, 762  
   in mysql.columns\_priv table, 767  
   in mysql.db table, 766  
   in mysql.procs\_priv table, 767  
   in mysql.tables\_priv table, 767  
 Windows, 2326, 2403, 3827  
   compiling clients on, 2798  
   MySQL limitations, 3770, 3771  
   path name separators, 217  
   pluggable authentication, 797  
   upgrading, 97  
 Windows Failover Clustering, 1825  
 windows option  
   mysql\_install\_db, 238  
 Windows Server Failover Clustering, 1825  
 Within(), 1312  
 WITH\_DEBUG option  
   CMake, 136  
 WITH\_DEFAULT\_COMPILER\_OPTIONS option  
   CMake, 136  
 WITH\_DEFAULT\_FEATURE\_SET option  
   CMake, 136  
 WITH\_EMBEDDED\_SERVER option  
   CMake, 136  
 WITH\_EXTRA\_CHARSETS option  
   CMake, 136  
 WITH\_LIBEDIT option  
   CMake, 136  
 WITH\_LIBWRAP option  
   CMake, 136  
 WITH\_READLINE option  
   CMake, 136  
 WITH\_SSL option  
   CMake, 136  
 WITH\_UNIXODBC option  
   CMake, 137  
 WITH\_ZLIB option  
   CMake, 137  
 WKB format, 1297  
 WKT format, 1296  
 workload, 3827  
 wrappers  
   Eiffel, 3482  
 write access  
   tmp, 145  
 write combining, 3827  
 write-binlog option  
   mysqlcheck, 279  
   mysql\_upgrade, 244  
 write\_buffer\_size myisamchk variable, 319  
 Writing to net  
   thread state, 965

**X**  
 X(), 1304  
 X509/Certificate, 806

---

---

- XA, 3827
- XA BEGIN, 1472
- XA COMMIT, 1472
- XA PREPARE, 1472
- XA RECOVER, 1472
- XA ROLLBACK, 1472
- XA START, 1472
- XA transactions, 1470
  - restrictions, 3764
  - transaction identifiers, 1472
- xid
  - XA transaction identifier, 1472
- xml option
  - mysql, 254
  - mysqldump, 291
- XOR
  - bitwise, 1258
  - logical, 1172

## Y

- Y(), 1304
- yaSSL, 807
- YEAR data type, 1117, 1128
- YEAR(), 1229
- YEARWEEK(), 1229
- Yen sign (Japanese), 3640
- young, 3827
- Your password does not satisfy the current policy requirements
  - password error, 741

## Z

- ZEROFILL, 1112, 1120, 2900
- ZFS, 1832

---

## C Function Index

### **my\_init()**

Section 21.8.11.1, “my\_init()”  
Section 21.8.11.3, “mysql\_thread\_init()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_affected\_rows()**

Section 13.2.1, “CALL Syntax”  
Section 13.2.5, “INSERT Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.48, “mysql\_num\_rows()”  
Section 21.8.10.1, “mysql\_stmt\_affected\_rows()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 13.2.8, “REPLACE Syntax”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 12.14, “Information Functions”  
Section 21.8.14.2, “What Results You Can Get from a Query”

### **mysql\_autocommit()**

Section 21.8.6.2, “mysql\_autocommit()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_change\_user()**

Section 21.8.6.3, “mysql\_change\_user()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_character\_set\_name()**

Section 21.8.6.4, “mysql\_character\_set\_name()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_client\_find\_plugin()**

Section 21.8.13.1, “mysql\_client\_find\_plugin()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_client\_register\_plugin()**

Section 21.8.13.2,  
“mysql\_client\_register\_plugin()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_close()**

Section 21.8.6.5, “mysql\_close()”  
Section 21.8.6.6, “mysql\_commit()”  
Section 21.8.6.7, “mysql\_connect()”  
Section 21.8.6.36, “mysql\_init()”  
Section 21.8.6.58, “mysql\_rollback()”  
Section 21.8.5, “C API Function Overview”  
Section C.5.2.11, “Communication Errors and Aborted Connections”

### **mysql\_commit()**

Section 21.8.6.6, “mysql\_commit()”

Section 21.8.5, “C API Function Overview”

### **mysql\_connect()**

Section 21.8.11.1, “my\_init()”  
Section 21.8.6.5, “mysql\_close()”  
Section 21.8.6.7, “mysql\_connect()”  
Section 21.8.6.49, “mysql\_options()”  
Section 21.8.11.3, “mysql\_thread\_init()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_create\_db()**

Section 21.8.6.8, “mysql\_create\_db()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_data\_seek()**

Section 21.8.6.9, “mysql\_data\_seek()”  
Section 21.8.6.59, “mysql\_row\_seek()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_debug()**

Section 21.8.6.10, “mysql\_debug()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_drop\_db()**

Section 21.8.6.11, “mysql\_drop\_db()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_dump\_debug\_info()**

Section 21.8.6.12, “mysql\_dump\_debug\_info()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_eof()**

Section 21.8.6.13, “mysql\_eof()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_errno()**

Section 21.8.13.1, “mysql\_client\_find\_plugin()”  
Section 21.8.13.2,  
“mysql\_client\_register\_plugin()”  
Section 21.8.6.7, “mysql\_connect()”  
Section 21.8.6.13, “mysql\_eof()”  
Section 21.8.6.14, “mysql\_errno()”  
Section 21.8.6.22, “mysql\_field\_count()”  
Section 21.8.13.3, “mysql\_load\_plugin()”  
Section 21.8.6.47, “mysql\_num\_fields()”  
Section 21.8.6.67, “mysql\_sqlstate()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.6, “C API Function Descriptions”  
Section 21.8.5, “C API Function Overview”  
Section 13.6.7.5.1, “Signal Condition Information Items”  
Section 6.3.11.2, “The Audit Log File”

---

Section C.2, “Types of Error Values”  
Section 21.8.14.1, “Why `mysql_store_result()`  
Sometimes Returns `NULL` After `mysql_query()`  
Returns Success”  
Section 22.2.4.8, “Writing Audit Plugins”

## **mysql\_error()**

Section 21.8.13.1, “`mysql_client_find_plugin()`”  
Section 21.8.13.2,  
“`mysql_client_register_plugin()`”  
Section 21.8.6.7, “`mysql_connect()`”  
Section 21.8.6.13, “`mysql_eof()`”  
Section 21.8.6.15, “`mysql_error()`”  
Section 21.8.13.3, “`mysql_load_plugin()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section 21.8.6, “C API Function Descriptions”  
Section 21.8.5, “C API Function Overview”  
Section 13.6.7.5.1, “Signal Condition Information Items”  
Section C.2, “Types of Error Values”  
Section 21.8.14.1, “Why `mysql_store_result()`  
Sometimes Returns `NULL` After `mysql_query()`  
Returns Success”  
Section 22.2.4.8, “Writing Audit Plugins”

## **mysql\_escape\_string()**

Section 21.8.6.16, “`mysql_escape_string()`”  
Section 21.8.5, “C API Function Overview”  
Section 6.1.7, “Client Programming Security  
Guidelines”

## **mysql\_fetch\_field()**

Section 21.8.6.17, “`mysql_fetch_field()`”  
Section 21.8.6.23, “`mysql_field_seek()`”  
Section 21.8.6.24, “`mysql_field_tell()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”

## **mysql\_fetch\_field\_direct()**

Section 21.8.6.18, “`mysql_fetch_field_direct()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.5, “C API Function Overview”

## **mysql\_fetch\_fields()**

Section 21.8.6.19, “`mysql_fetch_fields()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.5, “C API Function Overview”

## **mysql\_fetch\_lengths()**

Section 21.8.6.20, “`mysql_fetch_lengths()`”  
Section 21.8.6.21, “`mysql_fetch_row()`”  
Section 21.8.5, “C API Function Overview”

## **mysql\_fetch\_row()**

Section 14.9.1, “FEDERATED Storage Engine Overview”  
Section 21.8.6.13, “`mysql_eof()`”  
Section 21.8.6.14, “`mysql_errno()`”  
Section 21.8.6.20, “`mysql_fetch_lengths()`”  
Section 21.8.6.21, “`mysql_fetch_row()`”  
Section 21.8.6.60, “`mysql_row_tell()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.14.2, “What Results You Can Get from a  
Query”

## **mysql\_field\_count()**

Section 21.8.6.22, “`mysql_field_count()`”  
Section 21.8.6.47, “`mysql_num_fields()`”  
Section 21.8.6.52, “`mysql_query()`”  
Section 21.8.6.55, “`mysql_real_query()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.14.1, “Why `mysql_store_result()`  
Sometimes Returns `NULL` After `mysql_query()`  
Returns Success”

## **mysql\_field\_seek()**

Section 21.8.6.17, “`mysql_fetch_field()`”  
Section 21.8.6.23, “`mysql_field_seek()`”  
Section 21.8.6.24, “`mysql_field_tell()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”

## **mysql\_field\_tell()**

Section 21.8.6.24, “`mysql_field_tell()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.5, “C API Function Overview”

## **mysql\_free\_result()**

Section C.5.2.14, “Commands out of sync”  
Section 21.8.6.25, “`mysql_free_result()`”  
Section 21.8.6.41, “`mysql_list_dbs()`”  
Section 21.8.6.42, “`mysql_list_fields()`”  
Section 21.8.6.43, “`mysql_list_processes()`”  
Section 21.8.6.44, “`mysql_list_tables()`”  
Section 21.8.6.46, “`mysql_next_result()`”  
Section 21.8.10.23,  
“`mysql_stmt_result_metadata()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.9, “C API Prepared Statement Function  
Overview”

---

**mysql\_get\_character\_set\_info()** Section 21.8.6.5, “mysql\_close()”  
 Section 21.8.6.26, “mysql\_get\_ssl\_cipher()”  
 “mysql\_get\_character\_set\_info()”  
 Section 21.8.5, “C API Function Overview”  
 Section 10.4.2, “Choosing a Collation ID”

**mysql\_get\_client\_info()** Section 21.8.6.7, “mysql\_connect()”  
 Section 21.8.6.27, “mysql\_get\_client\_info()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_client\_version()** Section 21.8.6.28, “mysql\_get\_client\_version()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_host\_info()** Section 21.8.6.29, “mysql\_get\_host\_info()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_proto\_info()** Section 21.8.6.30, “mysql\_get\_proto\_info()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_server\_info()** Section 21.8.6.31, “mysql\_get\_server\_info()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_server\_version()** Section 21.8.6.32, “mysql\_get\_server\_version()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_get\_ssl\_cipher()** Section 21.8.6.33, “mysql\_get\_ssl\_cipher()”  
 Section 21.8.5, “C API Function Overview”  
 Section 6.3.9.3, “Using SSL Connections”

**mysql\_hex\_string()** Section 21.8.6.34, “mysql\_hex\_string()”  
 Section 21.8.5, “C API Function Overview”

**mysql\_info()** Section 13.1.6, “ALTER TABLE Syntax”  
 Section 13.2.5.2, “INSERT DELAYED Syntax”  
 Section 13.2.5, “INSERT Syntax”  
 Section 13.2.6, “LOAD DATA INFILE Syntax”  
 Section 21.8.6.35, “mysql\_info()”  
 Section 21.8.6.49, “mysql\_options()”  
 Section 1.8.6.1, “PRIMARY KEY and UNIQUE Index Constraints”  
 Section 13.2.11, “UPDATE Syntax”  
 Section 21.8.5, “C API Function Overview”  
 Section 21.8.14.2, “What Results You Can Get from a Query”

**mysql\_init()** Section 21.8.11.1, “my\_init()”

**mysql\_insert\_id()** Section 13.1.14, “CREATE TABLE Syntax”  
 Section 13.2.5, “INSERT Syntax”  
 Section 21.8.6.37, “mysql\_insert\_id()”  
 Section 21.8.4, “C API Data Structures”  
 Section 21.8.5, “C API Function Overview”  
 Section 21.8.14.3, “How to Get the Unique ID for the Last Inserted Row”  
 Section 12.14, “Information Functions”  
 Section 5.1.4, “Server System Variables”  
 Section 1.8.5.3, “Transaction and Atomic Operation Differences”  
 Section 3.6.9, “Using AUTO\_INCREMENT”  
 Section 21.8.14.2, “What Results You Can Get from a Query”

**mysql\_kill()** Section 21.8.6.38, “mysql\_kill()”  
 Section 21.8.6.71, “mysql\_thread\_id()”  
 Section 21.8.5, “C API Function Overview”  
 Section 21.8.15, “Controlling Automatic Reconnection Behavior”

**mysql\_library\_end()** Section 21.8.6.39, “mysql\_library\_end()”  
 Section 21.8.6.40, “mysql\_library\_init()”  
 Section 21.8.12.2, “mysql\_server\_end()”  
 Section 21.8.12, “C API Embedded Server Function Descriptions”  
 Section 21.8.5, “C API Function Overview”  
 Section 21.7, “libmysqld, the Embedded MySQL Server Library”

**mysql\_library\_init()** Section 21.8.11.1, “my\_init()”  
 Section 21.8.6.40, “mysql\_library\_init()”  
 Section 21.8.13.3, “mysql\_load\_plugin()”  
 Section 21.8.12.1, “mysql\_server\_init()”  
 Section 21.8.11.3, “mysql\_thread\_init()”  
 Section 21.8.12, “C API Embedded Server Function Descriptions”  
 Section 21.8.5, “C API Function Overview”  
 Section 21.7, “libmysqld, the Embedded MySQL Server Library”  
 Section 21.7.3, “Options with the Embedded Server”

---



---

Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_list\_dbs()**

Section 21.8.6.25, “mysql\_free\_result()”  
Section 21.8.6.41, “mysql\_list\_dbs()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_list\_fields()**

Section 21.8.6.42, “mysql\_list\_fields()”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”

### **mysql\_list\_processes()**

Section 21.8.6.43, “mysql\_list\_processes()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_list\_tables()**

Section 21.8.6.44, “mysql\_list\_tables()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_load\_plugin()**

Section 21.8.13.3, “mysql\_load\_plugin()”  
Section 21.8.13.4, “mysql\_load\_plugin\_v()”  
Section 21.8.5, “C API Function Overview”  
Section 22.2.4.2.3, “Client Plugin Descriptors”

### **mysql\_load\_plugin\_v()**

Section 21.8.13.3, “mysql\_load\_plugin()”  
Section 21.8.13.4, “mysql\_load\_plugin\_v()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_more\_results()**

Section 21.8.6.45, “mysql\_more\_results()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.10.17, “mysql\_stmt\_next\_result()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”

### **mysql\_next\_result()**

Section 13.2.1, “CALL Syntax”  
Section 21.8.6.45, “mysql\_more\_results()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.6.65, “mysql\_set\_server\_option()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”

### **mysql\_num\_fields()**

Section 21.8.6.18, “mysql\_fetch\_field\_direct()”  
Section 21.8.6.21, “mysql\_fetch\_row()”  
Section 21.8.6.47, “mysql\_num\_fields()”

Section 21.8.10.23,  
“mysql\_stmt\_result\_metadata()”  
Section 21.8.5, “C API Function Overview”

### **mysql\_num\_rows()**

Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.9, “mysql\_data\_seek()”  
Section 21.8.6.48, “mysql\_num\_rows()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.14.2, “What Results You Can Get from a Query”

### **mysql\_options()**

Section C.5.2.9, “MySQL server has gone away”  
Section 21.8.6.49, “mysql\_options()”  
Section 21.8.6.50, “mysql\_options4()”  
Section 21.8.6.51, “mysql\_ping()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.13, “C API Client Plugin Functions”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.8, “C API Prepared Statement Data Structures”  
Section 22.2.4.2.3, “Client Plugin Descriptors”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 21.8.15, “Controlling Automatic Reconnection Behavior”  
Section 6.3.6, “Password Expiration and Sandbox Mode”  
Section 20.9.7, “Performance Schema Connection Attribute Tables”  
Section 22.2.2, “Plugin API Components”  
Section 6.1.6, “Security Issues with LOAD DATA LOCAL”  
Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”  
Section 6.3.1, “User Names and Passwords”  
Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”  
Section 22.2.4.9.3, “Using the Authentication Plugins”

### **mysql\_options4()**

Section 21.8.6.49, “mysql\_options()”  
Section 21.8.6.50, “mysql\_options4()”  
Section 21.8.5, “C API Function Overview”  
Section 20.9.7, “Performance Schema Connection Attribute Tables”

### **mysql\_ping()**

Section C.5.2.9, “MySQL server has gone away”  
Section 21.8.6.51, “mysql\_ping()”  
Section 21.8.6.71, “mysql\_thread\_id()”  
Section 21.8.5, “C API Function Overview”



---

Section 21.8.15, “Controlling Automatic Reconnection Behavior”

## **mysql\_plugin\_options()**

Section 21.8.13.5, “mysql\_plugin\_options()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_query()**

Section 13.2.1, “CALL Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.8, “mysql\_create\_db()”  
Section 21.8.6.11, “mysql\_drop\_db()”  
Section 21.8.6.17, “mysql\_fetch\_field()”  
Section 21.8.6.38, “mysql\_kill()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.52, “mysql\_query()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.6.55, “mysql\_real\_query()”  
Section 21.8.6.57, “mysql\_reload()”  
Section 21.8.6.64, “mysql\_set\_local\_infile\_handler()”  
Section 21.8.6.65, “mysql\_set\_server\_option()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”  
Section 21.8.14.3, “How to Get the Unique ID for the Last Inserted Row”  
Section 21.8.14.1, “Why mysql\_store\_result() Sometimes Returns NULL After mysql\_query() Returns Success”  
Section 21.8.3.2, “Writing C API Threaded Client Programs”

## **mysql\_real\_connect()**

Section 13.2.1, “CALL Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.3, “mysql\_change\_user()”  
Section 21.8.6.7, “mysql\_connect()”  
Section 21.8.6.36, “mysql\_init()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.49, “mysql\_options()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.6.65, “mysql\_set\_server\_option()”  
Section 21.8.6.67, “mysql\_sqlstate()”  
Section 21.8.6.68, “mysql\_ssl\_set()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”  
Chapter 12, *Functions and Operators*  
Section 12.14, “Information Functions”  
Section 6.3.6, “Password Expiration and Sandbox Mode”  
Section 5.1.4, “Server System Variables”  
Section 13.5, “SQL Syntax for Prepared Statements”  
Section 18.2.1, “Stored Routine Syntax”

Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”  
Section 6.3.9.3, “Using SSL Connections”

## **mysql\_real\_escape\_string()**

Section 21.8.6.16, “mysql\_escape\_string()”  
Section 21.8.6.54, “mysql\_real\_escape\_string()”  
Section 21.8.6.62, “mysql\_set\_character\_set()”  
Section 21.8.5, “C API Function Overview”  
Section 6.1.7, “Client Programming Security Guidelines”  
Section 12.18.4.4, “Populating Spatial Columns”  
Section 9.1.1, “String Literals”

## **mysql\_real\_query()**

Section 13.2.1, “CALL Syntax”  
Section 14.9.1, “FEDERATED Storage Engine Overview”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.52, “mysql\_query()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.6.55, “mysql\_real\_query()”  
Section 21.8.6.65, “mysql\_set\_server\_option()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”

## **mysql\_refresh()**

Section 21.8.6.56, “mysql\_refresh()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_reload()**

Section 21.8.6.57, “mysql\_reload()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_rollback()**

Section 21.8.6.58, “mysql\_rollback()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_row\_seek()**

Section 21.8.6.59, “mysql\_row\_seek()”  
Section 21.8.6.60, “mysql\_row\_tell()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_row\_tell()**

Section 21.8.6.59, “mysql\_row\_seek()”  
Section 21.8.6.60, “mysql\_row\_tell()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.5, “C API Function Overview”

## **mysql\_select\_db()**

Section 21.8.6.61, “mysql\_select\_db()”

---

Section 21.8.5, “C API Function Overview”

## **mysql\_server\_end()**

Section 21.8.12.2, “mysql\_server\_end()”

Section 21.8.5, “C API Function Overview”

## **mysql\_server\_init()**

Section 21.8.11.1, “my\_init()”

Section 21.8.12.1, “mysql\_server\_init()”

Section 21.8.11.3, “mysql\_thread\_init()”

Section 21.8.5, “C API Function Overview”

## **mysql\_set\_character\_set()**

Section 21.8.6.26,

“mysql\_get\_character\_set\_info()”

Section 21.8.6.54, “mysql\_real\_escape\_string()”

Section 21.8.6.62, “mysql\_set\_character\_set()”

Section 21.8.5, “C API Function Overview”

## **mysql\_set\_local\_infile\_defaults()**

Section 21.8.6.63,

“mysql\_set\_local\_infile\_default()”

Section 21.8.5, “C API Function Overview”

## **mysql\_set\_local\_infile\_handler()**

Section 21.8.6.63,

“mysql\_set\_local\_infile\_default()”

Section 21.8.6.64,

“mysql\_set\_local\_infile\_handler()”

Section 21.8.5, “C API Function Overview”

## **mysql\_set\_server\_option()**

Section 21.8.6.65, “mysql\_set\_server\_option()”

Section 21.8.5, “C API Function Overview”

Section 21.8.16, “C API Support for Multiple Statement Execution”

## **mysql\_shutdown()**

Section 21.8.6.66, “mysql\_shutdown()”

Section 21.8.5, “C API Function Overview”

## **mysql\_sqlstate()**

Section 21.8.6.14, “mysql\_errno()”

Section 21.8.6.67, “mysql\_sqlstate()”

Section 21.8.5, “C API Function Overview”

Section 13.6.7.5.1, “Signal Condition Information Items”

Section C.2, “Types of Error Values”

## **mysql\_ssl\_set()**

Section 21.8.6.53, “mysql\_real\_connect()”

Section 21.8.6.68, “mysql\_ssl\_set()”

Section 21.8.5, “C API Function Overview”

Section 6.3.9.3, “Using SSL Connections”

## **mysql\_stat()**

Section 21.8.6.69, “mysql\_stat()”

Section 21.8.5, “C API Function Overview”

## **mysql\_stmt\_affected\_rows()**

Section 21.8.10.1, “mysql\_stmt\_affected\_rows()”

Section 21.8.10.10, “mysql\_stmt\_execute()”

Section 21.8.10.17, “mysql\_stmt\_next\_result()”

Section 21.8.10.18, “mysql\_stmt\_num\_rows()”

Section 21.8.9, “C API Prepared Statement Function Overview”

## **mysql\_stmt\_attr\_get()**

Section 21.8.10.2, “mysql\_stmt\_attr\_get()”

Section 21.8.9, “C API Prepared Statement Function Overview”

## **mysql\_stmt\_attr\_set()**

Section 21.8.10.3, “mysql\_stmt\_attr\_set()”

Section 21.8.10.10, “mysql\_stmt\_execute()”

Section 21.8.10.11, “mysql\_stmt\_fetch()”

Section 21.8.10.28, “mysql\_stmt\_store\_result()”

Section 21.8.4, “C API Data Structures”

Section 21.8.9, “C API Prepared Statement Function Overview”

Section 21.8.8.2, “C API Prepared Statement Type Conversions”

Section E.3, “Restrictions on Server-Side Cursors”

## **mysql\_stmt\_bind\_param()**

Section 21.8.10.4, “mysql\_stmt\_bind\_param()”

Section 21.8.10.10, “mysql\_stmt\_execute()”

Section 21.8.10.21, “mysql\_stmt\_prepare()”

Section 21.8.10.26,

“mysql\_stmt\_send\_long\_data()”

Section 21.8.8, “C API Prepared Statement Data Structures”

Section 21.8.9, “C API Prepared Statement Function Overview”

Section 21.8.18, “C API Prepared Statement Handling of Date and Time Values”

## **mysql\_stmt\_bind\_result()**

Section 21.8.10.5, “mysql\_stmt\_bind\_result()”

Section 21.8.10.11, “mysql\_stmt\_fetch()”

Section 21.8.10.12, “mysql\_stmt\_fetch\_column()”

Section 21.8.10.17, “mysql\_stmt\_next\_result()”

Section 21.8.10.28, “mysql\_stmt\_store\_result()”

Section 21.8.8, “C API Prepared Statement Data Structures”

Section 21.8.9, “C API Prepared Statement Function Overview”

Section 21.8.18, “C API Prepared Statement Handling of Date and Time Values”

## **mysql\_stmt\_close()**

Section 21.8.10.6, “mysql\_stmt\_close()”

Section 21.8.10.15, “mysql\_stmt\_init()”

---

Section 21.8.8, “C API Prepared Statement Data Structures”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_data\_seek()**

Section 21.8.10.7, “mysql\_stmt\_data\_seek()”  
Section 21.8.10.24, “mysql\_stmt\_row\_seek()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_errno()**

Section 21.8.10.8, “mysql\_stmt\_errno()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section C.2, “Types of Error Values”

### **mysql\_stmt\_error()**

Section 21.8.10.9, “mysql\_stmt\_error()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.21, “mysql\_stmt\_prepare()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section C.2, “Types of Error Values”

### **mysql\_stmt\_execute()**

Section 21.8.10.1, “mysql\_stmt\_affected\_rows()”  
Section 21.8.10.3, “mysql\_stmt\_attr\_set()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.17, “mysql\_stmt\_next\_result()”  
Section 21.8.10.26,  
“mysql\_stmt\_send\_long\_data()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.8, “C API Prepared Statement Data Structures”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.18, “C API Prepared Statement Handling of Date and Time Values”  
Section 21.8.8.2, “C API Prepared Statement Type Conversions”  
Section 8.9.3.1, “How the Query Cache Operates”

### **mysql\_stmt\_fetch()**

Section 21.8.10.5, “mysql\_stmt\_bind\_result()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.23,  
“mysql\_stmt\_result\_metadata()”  
Section 21.8.10.25, “mysql\_stmt\_row\_tell()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.8, “C API Prepared Statement Data Structures”

Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.8.2, “C API Prepared Statement Type Conversions”

### **mysql\_stmt\_fetch\_column()**

Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.12, “mysql\_stmt\_fetch\_column()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_field\_count()**

Section 21.8.10.13, “mysql\_stmt\_field\_count()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_free\_result()**

Section 21.8.10.3, “mysql\_stmt\_attr\_set()”  
Section 21.8.10.14, “mysql\_stmt\_free\_result()”  
Section 21.8.10.17, “mysql\_stmt\_next\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_init()**

Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.15, “mysql\_stmt\_init()”  
Section 21.8.10.21, “mysql\_stmt\_prepare()”  
Section 21.8.8, “C API Prepared Statement Data Structures”  
Section 21.8.10, “C API Prepared Statement Function Descriptions”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.7, “C API Prepared Statements”

### **mysql\_stmt\_insert\_id()**

Section 21.8.10.16, “mysql\_stmt\_insert\_id()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_next\_result()**

Section 13.2.1, “CALL Syntax”  
Section 21.8.10.17, “mysql\_stmt\_next\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.19, “C API Support for Prepared CALL Statements”

### **mysql\_stmt\_num\_rows()**

Section 21.8.10.7, “mysql\_stmt\_data\_seek()”  
Section 21.8.10.18, “mysql\_stmt\_num\_rows()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_param\_count()**

Section 21.8.10.10, “mysql\_stmt\_execute()”

---

Section 21.8.10.19, “mysql\_stmt\_param\_count()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_param\_metadata()**

Section 21.8.10.20,  
“mysql\_stmt\_param\_metadata()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_prepare()**

Section 21.8.10.4, “mysql\_stmt\_bind\_param()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.13, “mysql\_stmt\_field\_count()”  
Section 21.8.10.21, “mysql\_stmt\_prepare()”  
Section 21.8.10.22, “mysql\_stmt\_reset()”  
Section 21.8.10.23,  
“mysql\_stmt\_result\_metadata()”  
Section 21.8.8, “C API Prepared Statement Data Structures”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.18, “C API Prepared Statement Handling of Date and Time Values”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 13.5, “SQL Syntax for Prepared Statements”

### **mysql\_stmt\_reset()**

Section 21.8.10.3, “mysql\_stmt\_attr\_set()”  
Section 21.8.10.22, “mysql\_stmt\_reset()”  
Section 21.8.10.26,  
“mysql\_stmt\_send\_long\_data()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_result\_metadata()**

Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.23,  
“mysql\_stmt\_result\_metadata()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.8.2, “C API Prepared Statement Type Conversions”

### **mysql\_stmt\_row\_seek()**

Section 21.8.10.24, “mysql\_stmt\_row\_seek()”  
Section 21.8.10.25, “mysql\_stmt\_row\_tell()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_row\_tell()**

Section 21.8.10.24, “mysql\_stmt\_row\_seek()”

Section 21.8.10.25, “mysql\_stmt\_row\_tell()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_stmt\_send\_long\_data()**

Section 21.8.10.22, “mysql\_stmt\_reset()”  
Section 21.8.10.26,  
“mysql\_stmt\_send\_long\_data()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 5.1.4, “Server System Variables”

### **mysql\_stmt\_sqlstate()**

Section 21.8.10.27, “mysql\_stmt\_sqlstate()”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section C.2, “Types of Error Values”

### **mysql\_stmt\_store\_result()**

Section 21.8.10.3, “mysql\_stmt\_attr\_set()”  
Section 21.8.10.7, “mysql\_stmt\_data\_seek()”  
Section 21.8.10.11, “mysql\_stmt\_fetch()”  
Section 21.8.10.18, “mysql\_stmt\_num\_rows()”  
Section 21.8.10.24, “mysql\_stmt\_row\_seek()”  
Section 21.8.10.25, “mysql\_stmt\_row\_tell()”  
Section 21.8.10.28, “mysql\_stmt\_store\_result()”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.9, “C API Prepared Statement Function Overview”

### **mysql\_store\_result()**

Section C.5.2.14, “Commands out of sync”  
Section 14.9.1, “FEDERATED Storage Engine Overview”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.9, “mysql\_data\_seek()”  
Section 21.8.6.13, “mysql\_eof()”  
Section 21.8.6.17, “mysql\_fetch\_field()”  
Section 21.8.6.21, “mysql\_fetch\_row()”  
Section 21.8.6.22, “mysql\_field\_count()”  
Section 21.8.6.25, “mysql\_free\_result()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.47, “mysql\_num\_fields()”  
Section 21.8.6.48, “mysql\_num\_rows()”  
Section 21.8.6.59, “mysql\_row\_seek()”  
Section 21.8.6.60, “mysql\_row\_tell()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.23,  
“mysql\_stmt\_result\_metadata()”  
Section 21.8.6.70, “mysql\_store\_result()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 4.5.1, “mysql — The MySQL Command-Line Tool”  
Section 21.8.14.2, “What Results You Can Get from a Query”

---

Section 21.8.14.1, “Why `mysql_store_result()` Sometimes Returns `NULL` After `mysql_query()` Returns Success”  
Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_thread\_end()**

Section 21.8.11.2, “`mysql_thread_end()`”  
Section 21.8.5, “C API Function Overview”  
Section 21.7, “libmysqld, the Embedded MySQL Server Library”  
Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_thread\_id()**

Section 21.8.6.51, “`mysql_ping()`”  
Section 21.8.6.71, “`mysql_thread_id()`”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.15, “Controlling Automatic Reconnection Behavior”

### **mysql\_thread\_init()**

Section 21.8.11.1, “`my_init()`”  
Section 21.8.11.2, “`mysql_thread_end()`”  
Section 21.8.11.3, “`mysql_thread_init()`”  
Section 21.8.5, “C API Function Overview”  
Section 21.7, “libmysqld, the Embedded MySQL Server Library”  
Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_thread\_safe()**

Section 21.8.11.4, “`mysql_thread_safe()`”  
Section 21.8.5, “C API Function Overview”

### **mysql\_use\_result()**

Section C.5.2.14, “Commands out of sync”  
Section 21.8.6.9, “`mysql_data_seek()`”  
Section 21.8.6.13, “`mysql_eof()`”  
Section 21.8.6.21, “`mysql_fetch_row()`”  
Section 21.8.6.25, “`mysql_free_result()`”  
Section 21.8.6.46, “`mysql_next_result()`”  
Section 21.8.6.47, “`mysql_num_fields()`”  
Section 21.8.6.48, “`mysql_num_rows()`”  
Section 21.8.6.59, “`mysql_row_seek()`”  
Section 21.8.6.60, “`mysql_row_tell()`”  
Section 21.8.10.10, “`mysql_stmt_execute()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section C.5.2.8, “Out of memory”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 4.5.1, “`mysql` — The MySQL Command-Line Tool”  
Section 21.8.14.2, “What Results You Can Get from a Query”

Section 21.8.3.2, “Writing C API Threaded Client Programs”

### **mysql\_warning\_count()**

Section 21.8.6.46, “`mysql_next_result()`”  
Section 21.8.6.73, “`mysql_warning_count()`”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 21.8.5, “C API Function Overview”  
Section C.2, “Types of Error Values”





---

# Command Index

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [G](#) | [H](#) | [I](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [R](#) | [S](#) | [T](#)  
[U](#) | [V](#) | [W](#) | [Y](#) | [Z](#)

## A

[\[index top \[3911\]\]](#)

### Access

Section 13.2.2, “`DELETE` Syntax”

### addgroup

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

### addr2line

Section 22.4.1.5, “Using a Stack Trace”

### adduser

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

### apt-get

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

## B

[\[index top \[3911\]\]](#)

### bash

Section 6.1.2.1, “End-User Guidelines for Password Security”  
Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”  
Section 4.2.1, “Invoking MySQL Programs”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 4.2.4, “Setting Environment Variables”  
Section 1.2, “Typographical and Syntax Conventions”

### bison

Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 2.9.3, “Installing MySQL from a Development Source Tree”  
Section 2.9, “Installing MySQL from Source”

### bzip

Section 2.9.3, “Installing MySQL from a Development Source Tree”

## C

[\[index top \[3911\]\]](#)

### c++filt

Section 22.4.1.5, “Using a Stack Trace”

### cat

Section 4.5.1.1, “`mysql` Options”

### chkconfig

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”  
Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

### chroot

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

### CMake

Section C.5.2.17, “Can't initialize character set”  
Section 10.3, “Adding a Character Set”  
Section 22.2.4.3, “Compiling and Installing Plugin Libraries”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 6.3.9.2, “Configuring MySQL for SSL”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 2.12, “Environment Variables”  
Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”  
Section 2.9.3, “Installing MySQL from a Development Source Tree”  
Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.9, “Installing MySQL from Source”  
Section 21.7, “`libmysqld`, the Embedded MySQL Server Library”  
Section 22.2.5, “MySQL Services for Plugins”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 2.1.1, “Operating Systems Supported by MySQL Community Server”  
Section 20.2.1, “Performance Schema Build Configuration”  
Section 5.6.3, “Running Multiple MySQL Instances on Unix”  
Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 5.1.4, “Server System Variables”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 14.10, “The `EXAMPLE` Storage Engine”  
Section 14.9, “The `FEDERATED` Storage Engine”

---

Section 5.7, “Tracing `mysqld` Using DTrace”  
Section 4.2.3.3, “Using Option Files”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## **cmake**

Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 14.2.9.4.1, “Password-Protecting the `memcached` Interface through SASL”

## **cmd**

Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”

## **cmd.exe**

Section 4.2.1, “Invoking MySQL Programs”  
Section 1.2, “Typographical and Syntax Conventions”

## **command.com**

Section 4.2.1, “Invoking MySQL Programs”  
Section 1.2, “Typographical and Syntax Conventions”

## **comp\_err**

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”  
Section 4.1, “Overview of MySQL Programs”

## **configure**

Section 2.12, “Environment Variables”  
Section 1.7, “How to Report Bugs or Problems”  
Section 1.2, “Typographical and Syntax Conventions”

## **copy**

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

## **coreadm**

Section 2.6, “Installing MySQL on Solaris and OpenSolaris”  
Section 5.1.3, “Server Command Options”

## **cp**

Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 7.1, “Backup and Recovery Types”  
Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”

## **crash-me**

Section 8.12.2, “The MySQL Benchmark Suite”

## **crc32**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **cron**

Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section 13.7.2.2, “CHECK TABLE Syntax”  
Section 14.3.1, “MyISAM Startup Options”  
Section 5.2.6, “Server Log Maintenance”  
Section 7.6.5, “Setting Up a MyISAM Table Maintenance Schedule”  
Section 3.5, “Using `mysql` in Batch Mode”

## **csh**

Section 4.2.1, “Invoking MySQL Programs”  
Section 4.2.4, “Setting Environment Variables”  
Section 1.2, “Typographical and Syntax Conventions”

## **D**

[index top [3911]]

## **date**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **df**

Section C.5.1, “How to Determine What Is Causing a Problem”

## **Directory Utility**

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

## **drwtsn32.exe**

Section 22.4.1.3, “Using `pdb` to create a Windows crashdump”

## **dump**

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

## **E**

[index top [3911]]

## **emerge**

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

## **G**

[index top [3911]]



---

## gcc

Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 21.7.1, “Compiling Programs with `libmysqld`”  
Section 2.13.3, “Problems Using the Perl `DBI/DBD` Interface”

## gdb

Section 22.4.1.1, “Compiling MySQL for Debugging”  
Section 22.4.1.4, “Debugging `mysqld` under `gdb`”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## gmake

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.9, “Installing MySQL from Source”  
Section 2.8, “Installing MySQL on FreeBSD”

## GnuPG

Section 2.1.4.2, “Signature Checking Using `GnuPG`”

## gnutar

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”  
Section 2.9, “Installing MySQL from Source”

## gogoc

Section 5.1.9.5, “Obtaining an IPv6 Address from a Broker”

## gpg

Section 2.1.4.2, “Signature Checking Using `GnuPG`”

## grep

Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 3.3.4.7, “Pattern Matching”

## groupadd

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”  
Section 2.5.1, “Installing MySQL from RPM Packages on Linux”  
Section 2.6, “Installing MySQL on Solaris and OpenSolaris”

## gtar

Section 2.7.1, “General Notes on Installing MySQL on HP-UX”  
Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”  
Section 2.9, “Installing MySQL from Source”

Section 2.6, “Installing MySQL on Solaris and OpenSolaris”

## gunzip

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

## gzip

Section 1.7, “How to Report Bugs or Problems”  
Section 2.4, “Installing MySQL on Mac OS X”

## H

[[index top \[3911\]](#)]

## hdparm

Section 14.2.6, “`InnoDB` Startup Options and System Variables”

## help contents

Section 4.5.1.4, “`mysql` Server-Side Help”

## hostname

Section C.5.2.2, “Can't connect to [local] MySQL server”

## I

[[index top \[3911\]](#)]

## icc

Section 2.1.6, “Compiler-Specific Build Characteristics”  
Section 22.4, “Debugging and Porting MySQL”

## ifconfig

Section 5.1.9.1, “Verifying System Support for IPv6”

## innochecksum

Section 4.6.1, “`innochecksum` — Offline `InnoDB` File Checksum Utility”  
Section 13.7.2.2, “`CHECK TABLE` Syntax”  
Section 4.1, “Overview of MySQL Programs”

## K

[[index top \[3911\]](#)]

## kill

Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section E.6, “Restrictions on XA Transactions”

---

## **ksh**

Section 4.2.1, “Invoking MySQL Programs”  
Section 4.2.4, “Setting Environment Variables”

## **L**

[index top [3911]]

## **ldconfig**

Section 22.3.2.5, “Compiling and Installing User-Defined Functions”

## **less**

Section 4.5.1.2, “mysql Commands”  
Section 4.5.1.1, “mysql Options”

## **ln**

Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”

## **logger**

Section 4.3.2, “mysqld\_safe — MySQL Server Startup Script”

## **ls +Ll**

Section C.5.4.4, “Where MySQL Stores Temporary Files”

## **M**

[index top [3911]]

## **m4**

Section 2.9, “Installing MySQL from Source”

## **make**

Section 22.2.4.3, “Compiling and Installing Plugin Libraries”  
Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.9, “Installing MySQL from Source”  
Section 2.8, “Installing MySQL on FreeBSD”  
Section 2.13.3, “Problems Using the Perl DBI/DBD Interface”

## **make install**

Section 22.2.4.3, “Compiling and Installing Plugin Libraries”

## **make package**

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”

Section 2.9.4, “MySQL Source-Configuration Options”

## **make test**

Section 2.9.3, “Installing MySQL from a Development Source Tree”  
Section 2.13.1, “Installing Perl on Unix”  
Section 22.1.2, “The MySQL Test Suite”

## **md5**

Section 2.1.4.1, “Verifying the MD5 Checksum”

## **md5.exe**

Section 2.1.4.1, “Verifying the MD5 Checksum”

## **md5sum**

Section 2.1.4.1, “Verifying the MD5 Checksum”

## **memcached**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.2.9.5.2, “Adapting an Existing memcached Application for the Integrated memcached Daemon”  
Section 14.2.9.5.1, “Adapting an Existing MySQL Schema for a memcached Application”  
Section 14.2.9.5.5, “Adapting DML Statements to memcached Operations”  
Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 14.2.9.1, “Benefits of the InnoDB / memcached Combination”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”  
Section 14.2.9.3, “Getting Started with InnoDB Memcached Plugin”  
Section 14.2.9, “InnoDB Integration with memcached”  
Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”  
Section 14.2.9.7, “Internals of the InnoDB memcached Plugin”  
Section 14.2.9.4.1, “Password-Protecting the memcached Interface through SASL”  
Section 14.2.9.5.6, “Performing DML and DDL Statements on the Underlying InnoDB Table”  
Section 14.2.9.3.1, “Prerequisites for the InnoDB memcached Plugin”  
Section 14.2.9.4, “Security Considerations for the InnoDB memcached Plugin”  
Section 14.2.9.8, “Troubleshooting the InnoDB memcached Plugin”  
Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”  
Section 15.6, “Using MySQL with memcached”  
Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”  
Section 14.2.9.3.3, “Verifying the InnoDB and memcached Setup”  
Section 1.4, “What Is New in MySQL 5.6”

---

Section 14.2.9.5, “Writing Applications for the InnoDB memcached Interface”

## **memcapable**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

## **memcat**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 14.2.9.3.3, “Verifying the InnoDB and memcached Setup”

## **memcp**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

## **memflush**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

## **memrm**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

## **memslap**

Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

## **mkdir**

Section 13.1.8, “CREATE DATABASE Syntax”

## **mklink**

Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”

## **more**

Section 4.5.1.2, “mysql Commands”  
Section 4.5.1.1, “mysql Options”

## **mysql2mysql**

Section 4.7.1, “mysql2mysql — Convert mSQL Programs for Use with MySQL”  
Section 4.8.2, “replace — A String-Replacement Utility”  
Section 4.1, “Overview of MySQL Programs”

## **mysql2mysqlC-source-file ...**

Section 4.7.1, “mysql2mysql — Convert mSQL Programs for Use with MySQL”

## **mv**

Section 5.2.6, “Server Log Maintenance”  
Section 5.2.2, “The Error Log”  
Section 5.2.3, “The General Query Log”

## **my\_print\_defaults**

Section 4.7.3, “my\_print\_defaults — Display Options from Option Files”  
Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”  
Section 4.7, “MySQL Program Development Utilities”  
Section 4.1, “Overview of MySQL Programs”

## **myisam\_ftdump**

Section 4.6.2, “myisam\_ftdump — Display Full-Text Index information”  
Section 12.9, “Full-Text Search Functions”  
Section 4.1, “Overview of MySQL Programs”

## **myisamchk**

Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”  
Section 4.6.3.2, “myisamchk Check Options”  
Section 4.6.3.1, “myisamchk General Options”  
Section 4.6.3.6, “myisamchk Memory Usage”  
Section 4.6.3.3, “myisamchk Repair Options”  
Section 13.7.2.1, “ANALYZE TABLE Syntax”  
Section 13.7.2.2, “CHECK TABLE Syntax”  
Section 13.2.2, “DELETE Syntax”  
Section 8.8.2, “EXPLAIN Output Format”  
Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”  
Section 13.7.6.5, “LOAD INDEX INTO CACHE Syntax”  
Section 14.3.1, “MyISAM Startup Options”  
Section 7.6, “MyISAM Table Maintenance and Crash Recovery”  
Section 7.6.4, “MyISAM Table Optimization”  
Section 14.3.3, “MyISAM Table Storage Formats”  
Section 13.7.2.5, “REPAIR TABLE Syntax”  
Section 13.7.5.23, “SHOW INDEX Syntax”  
Section 13.7.5.37, “SHOW TABLE STATUS Syntax”  
Section 8.6.2, “Bulk Data Loading for MyISAM Tables”  
Section 10.5, “Character Set Configuration”  
Section 14.3.3.3, “Compressed Table Characteristics”  
Section 14.3.4.1, “Corrupted MyISAM Tables”  
Section 7.2, “Database Backup Methods”  
Section 22.4.1, “Debugging a MySQL Server”  
Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.10, “mysqlhotcopy — A Database Backup Program”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 14.3.3.2, “Dynamic Table Characteristics”  
Section 8.10.5, “External Locking”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 7.6.2, “How to Check MyISAM Tables for Errors”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 1.7, “How to Report Bugs or Problems”  
Section E.10.3, “Limits on Table Size”  
Section 17.3.4, “Maintenance of Partitions”

---

Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”  
Section 4.6.3.5, “Obtaining Table Information with `myisamchk`”  
Section 8.6.1, “Optimizing `MyISAM` Queries”  
Section 4.6.3.4, “Other `myisamchk` Options”  
Section 8.2.5, “Other Optimization Tips”  
Section 4.1, “Overview of MySQL Programs”  
Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.3, “Server Command Options”  
Section 7.6.5, “Setting Up a `MyISAM` Table Maintenance Schedule”  
Section 8.6.3, “Speed of `REPAIR TABLE` Statements”  
Section 14.3.3.1, “Static (Fixed-Length) Table Characteristics”  
Section 8.11.1, “System Factors and Startup Parameter Tuning”  
Section 14.3, “The `MyISAM` Storage Engine”  
Section 7.6.1, “Using `myisamchk` for Crash Recovery”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in `mysqld`”  
Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## `myisamchk *.MYI`

Section 7.6.3, “How to Repair `MyISAM` Tables”

## `myisamchk tbl_name`

Section 7.6.2, “How to Check `MyISAM` Tables for Errors”

## `myisamlog`

Section 4.6.4, “`myisamlog` — Display `MyISAM` Log File Contents”  
Section 4.1, “Overview of MySQL Programs”

## `myisampack`

Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only `MyISAM` Tables”  
Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 14.8.1, “`MERGE` Table Advantages and Disadvantages”  
Section 14.3.3, “`MyISAM` Table Storage Formats”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”  
Section 14.3.3.3, “Compressed Table Characteristics”  
Section 8.10.5, “External Locking”  
Section E.10.3, “Limits on Table Size”  
Section 4.6.3.5, “Obtaining Table Information with `myisamchk`”  
Section 8.4.1, “Optimizing Data Size”  
Section 4.1, “Overview of MySQL Programs”

Section 13.1.14.3, “Silent Column Specification Changes”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 14.3, “The `MyISAM` Storage Engine”

## `mysql`

Section 1.8.5.5, “`--` as the Start of a Comment”  
Section 4.5.1, “`mysql` — The MySQL Command-Line Tool”  
Section 4.5.1.2, “`mysql` Commands”  
Section 4.5.1.3, “`mysql` Logging”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.5.1.4, “`mysql` Server-Side Help”  
Section 4.5.1.6, “`mysql` Tips”  
Section 13.1.6.1, “`ALTER TABLE` Partition Operations”  
Section 13.6.1, “`BEGIN ... END` Compound-Statement Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.6.7.3, “`GET DIAGNOSTICS` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.8.3, “`HELP` Syntax”  
Section C.5.2.15, “Ignoring `user`”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 13.2.7, “`LOAD XML` Syntax”  
Section 21.8.6.14, “`mysql_errno()`”  
Section 21.8.6.67, “`mysql_sqlstate()`”  
Section C.5.2.8, “Out of memory”  
Section 17.2.3.1, “`RANGE COLUMNS` partitioning”  
Section 13.7.1.6, “`REVOKE` Syntax”  
Section 13.2.9.1, “`SELECT ... INTO` Syntax”  
Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 13.6.7.5, “`SIGNAL` Syntax”  
Section 6.3.2, “Adding User Accounts”  
Section 14.2.2.4, “Backing Up and Recovering an `InnoDB` Database”  
Section 7.1, “Backup and Recovery Types”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”  
Section 9.6, “Comment Syntax”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 3.1, “Connecting to and Disconnecting from the Server”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 5.1.9.4, “Connecting Using IPv6 Nonlocal Host Addresses”  
Section 5.1.9.3, “Connecting Using the IPv6 Local Host Address”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 21.8.15, “Controlling Automatic Reconnection Behavior”  
Section 14.2.5.8, “Controlling Optimizer Statistics Estimation”

---

Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”

Section 2.11.5, “Copying MySQL Databases to Another Machine”

Section 3.3.1, “Creating and Selecting a Database”

Section 14.2.2.1, “Creating the `InnoDB` Tablespace”

Section 2.3.4.7, “Customizing the `PATH` for MySQL Tools”

Section 22.4.2, “Debugging a MySQL Client”

Section 18.1, “Defining Stored Programs”

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

Section 4.4.6, “`mysql_tzinfo_to_sql` — Load the Time Zone Tables”

Section 4.5.1.6.5, “Disabling `mysql` Auto-Reconnect”

Section 2.11.2, “Downgrading MySQL”

Section 6.1.2.1, “End-User Guidelines for Password Security”

Section 3.2, “Entering Queries”

Section 2.12, “Environment Variables”

Section 18.4.2, “Event Scheduler Configuration”

Section 7.3, “Example Backup and Recovery Strategy”

Section 21.8.2, “Example C API Client Programs”

Section 3.6, “Examples of Common Queries”

Section 5.5.5, “Examples of Online DDL”

Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”

Section 4.5.1.5, “Executing SQL Statements from a Text File”

Chapter 12, *Functions and Operators*

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

Section 14.2.3.10, “How to Cope with Deadlocks”

Section C.5.1, “How to Determine What Is Causing a Problem”

Section 1.7, “How to Report Bugs or Problems”

Section 6.1.5, “How to Run MySQL as a Normal User”

Section 12.14, “Information Functions”

Section 4.5.1.6.1, “Input-Line Editing”

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 4.2.1, “Invoking MySQL Programs”

Section 7.4.5.1, “Making a Copy of a Database”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 8.12.1, “Measuring the Speed of Expressions and Functions”

Section 15.7, “MySQL Proxy”

Section 10.6, “MySQL Server Time Zone Support”

Section 8.2.1.3, “Optimizing `LIMIT` Queries”

Section 4.2.3.6, “Option Defaults, Options Expecting Values, and the = Sign”

Section 4.1, “Overview of MySQL Programs”

Section C.5.2.10, “Packet Too Large”

Section 6.3.6, “Password Expiration and Sandbox Mode”

Section 6.1.2.4, “Password Hashing in MySQL”

Section 6.3.7, “Pluggable Authentication”

Section 7.5, “Point-in-Time (Incremental) Recovery Using the Binary Log”

Section 2.10.1.1, “Problems Running `mysql_install_db`”

Section 4.2.3.2, “Program Option Modifiers”

Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”

Section 7.4.4, “Reloading Delimited-Text Format Backups”

Section 7.4.2, “Reloading SQL-Format Backups”

Section C.5.4.1.3, “Resetting the Root Password: Generic Instructions”

Section E.9, “Restrictions on Pluggable Authentication”

Section 2.10.2, “Securing the Initial MySQL Accounts”

Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”

Section 5.1.3, “Server Command Options”

Section C.3, “Server Error Codes and Messages”

Section 5.1.4, “Server System Variables”

Section 5.1.10, “Server-Side Help”

Section 12.18.5, “Spatial Analysis Functions”

Section 4.2.3, “Specifying Program Options”

Section 13.5, “SQL Syntax for Prepared Statements”

Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

Section 2.3.4.8, “Starting MySQL as a Windows Service”

Section 9.1.1, “String Literals”

Section 11.4.3, “The `BLOB` and `TEXT` Types”

Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

Section 6.3.7.6, “The Socket Peer-Credential Authentication Plugin”

Section 18.3.1, “Trigger Syntax”

Section 14.2.4.7, “Troubleshooting `InnoDB` Data Dictionary Operations”

Chapter 3, *Tutorial*

Section 21.2.4.8, “Tutorial: Using `MySqlScript`”

Section 1.2, “Typographical and Syntax Conventions”

Section 4.5.1.6.2, “Unicode Support on Windows”

Section 6.3.7.3.2.3, “Unix Password Authentication with Proxy Users and Group Mapping”

Section 6.3.7.3.2.1, “Unix Password Authentication without Proxy Users”

Section 2.10.1, “Unix Postinstallation Procedures”

Section 3.5, “Using `mysql` in Batch Mode”

---



---

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”  
Section 7.4, “Using `mysqldump` for Backups”  
Section 7.3.2, “Using Backups for Recovery”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 4.2.3.3, “Using Option Files”  
Section 4.2.3.1, “Using Options on the Command Line”  
Section 4.2.3.5, “Using Options to Set Program Variables”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in `mysqld`”  
Section 6.3.9.3, “Using SSL Connections”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”  
Section 1.4, “What Is New in MySQL 5.6”  
Section 2.3.7, “Windows Postinstallation Procedures”  
Section 12.11, “XML Functions”

## **mysql ...**

Section 22.4.1.1, “Compiling MySQL for Debugging”

## **mysql {start|stop}**

Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”

## **mysql-server**

Section 2.8, “Installing MySQL on FreeBSD”

## **mysql-test-run.pl**

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 22.1.2, “The MySQL Test Suite”

## **mysql-test-run.pl test\_name**

Section 22.1.2, “The MySQL Test Suite”

## **mysql.exe**

Section 4.5.1.6.2, “Unicode Support on Windows”

## **mysql.script**

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

## **mysql.server**

Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 2.5, “Installing MySQL on Linux”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 4.1, “Overview of MySQL Programs”  
Section 5.1.3, “Server Command Options”  
Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

Section C.5.4.6, “Time Zone Problems”

## **mysql.server stop**

Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

## **mysql\_config**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”  
Section 21.8.3.1, “Building C API Client Programs”  
Section 21.7.1, “Compiling Programs with `libmysqld`”  
Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 21.8.1, “MySQL C API Implementations”  
Section 4.1, “Overview of MySQL Programs”  
Section 22.2.2, “Plugin API Components”

## **mysql\_config options**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **mysql\_config\_editor**

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”  
Section 6.1.2.1, “End-User Guidelines for Password Security”  
Section 2.12, “Environment Variables”  
Section 4.1, “Overview of MySQL Programs”  
Section 4.2.3.3, “Using Option Files”  
Section 1.4, “What Is New in MySQL 5.6”

## **mysql\_config\_editor options command**

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

## **mysql\_convert\_table\_format**

Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.1, “Overview of MySQL Programs”

## **mysql\_find\_rows**

Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”  
Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”  
Section 4.1, “Overview of MySQL Programs”

## **mysql\_fix\_extensions**

Section 4.6.13, “`mysql_fix_extensions` — Normalize Table File Name Extensions”  
Section 4.1, “Overview of MySQL Programs”

---

## **mysql\_fix\_extensions**

### **data\_dir**

Section 4.6.13, “[mysql\\_fix\\_extensions](#) — Normalize Table File Name Extensions”

## **mysql\_install\_db**

Section 4.4.3, “[mysql\\_install\\_db](#) — Initialize MySQL Data Directory”

Section 6.3.2, “Adding User Accounts”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

Section 2.4.2, “Installing MySQL on Mac OS X Using Native Packages”

Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”

Section 14.2.7, “Limits on InnoDB Tables”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 4.1, “Overview of MySQL Programs”

Section 2.10.1.1, “Problems Running

[mysql\\_install\\_db](#)”

Section 2.10.2, “Securing the Initial MySQL Accounts”

Section 5.1.3, “Server Command Options”

Section 5.1.10, “Server-Side Help”

Section 5.6.1, “Setting Up Multiple Data Directories”

Section 2.10.1, “Unix Postinstallation Procedures”

Section 5.1.2.2, “Using a Sample Default Server Configuration File”

Section 1.4, “What Is New in MySQL 5.6”

## **mysql\_plugin**

Section 4.4.4, “[mysql\\_plugin](#) — Configure MySQL Server Plugins”

Section 4.1, “Overview of MySQL Programs”

## **mysql\_secure\_installation**

Section 4.4.5, “[mysql\\_secure\\_installation](#) — Improve MySQL Installation Security”

Section 4.4.3, “[mysql\\_install\\_db](#) — Initialize MySQL Data Directory”

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”

Section 2.6.1, “Installing MySQL on Solaris using a Solaris `PKG`”

Section 4.1, “Overview of MySQL Programs”

Section 2.10.2, “Securing the Initial MySQL Accounts”

## **mysql\_setpermission**

Section 4.6.14, “[mysql\\_setpermission](#) — Interactively Set Permissions in Grant Tables”

Section 4.1, “Overview of MySQL Programs”

Section 2.10.1, “Unix Postinstallation Procedures”

## **mysql\_setpermissions**

Section 4.6.14, “[mysql\\_setpermission](#) — Interactively Set Permissions in Grant Tables”

## **mysql\_stmt\_execute()**

Section 5.1.6, “Server Status Variables”

## **mysql\_stmt\_prepare()**

Section 5.1.6, “Server Status Variables”

## **mysql\_tzinfo\_to\_sql**

Section 4.4.6, “[mysql\\_tzinfo\\_to\\_sql](#) — Load the Time Zone Tables”

Section 10.6, “MySQL Server Time Zone Support”

Section 4.1, “Overview of MySQL Programs”

## **mysql\_tzinfo\_to\_sql**

### **arguments**

Section 4.4.6, “[mysql\\_tzinfo\\_to\\_sql](#) — Load the Time Zone Tables”

## **mysql\_upgrade**

Section 4.4.7, “[mysql\\_upgrade](#) — Check and Upgrade MySQL Tables”

Section 13.1.1, “ALTER DATABASE Syntax”

Section 13.7.3.1, “CREATE FUNCTION Syntax for User-Defined Functions”

Section 11.3.4, “YEAR(2) Limitations and Migrating to YEAR(4)”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”

Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”

Section 2.11.2, “Downgrading MySQL”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 4.1, “Overview of MySQL Programs”

Section 6.1.2.4, “Password Hashing in MySQL”

Section 20.2.1, “Performance Schema Build Configuration”

Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

Section 5.1.3, “Server Command Options”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

Section 2.11.1, “Upgrading MySQL”

Section 2.3.6, “Upgrading MySQL on Windows”

Section 1.4, “What Is New in MySQL 5.6”

---

## mysql\_waitpid

Section 4.6.15, “mysql\_waitpid — Kill Process and Wait for Its Termination”

Section 4.1, “Overview of MySQL Programs”

## mysql\_waitpid()

Section 4.6.15, “mysql\_waitpid — Kill Process and Wait for Its Termination”

## mysql\_zap

Section 4.6.16, “mysql\_zap — Kill Processes That Match a Pattern”

Section C.5.2.2, “Can't connect to [local] MySQL server”

Section 4.1, “Overview of MySQL Programs”

## mysqlaccess

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 1.7, “How to Report Bugs or Problems”

Section 4.1, “Overview of MySQL Programs”

Section 2.10.1, “Unix Postinstallation Procedures”

## mysqladmin

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section C.5.2.2, “Can't connect to [local] MySQL server”

Section 13.1.8, “CREATE DATABASE Syntax”

Section 13.1.17, “DROP DATABASE Syntax”

Section 13.7.6.3, “FLUSH Syntax”

Section 6.3.5, “Assigning Account Passwords”

Section 16.3.1.1, “Backing Up a Slave Using

mysqldump”

Section 4.2.2, “Connecting to the MySQL Server”

Section 2.3.4.7, “Customizing the PATH for MySQL Tools”

Section 22.4.1, “Debugging a MySQL Server”

Section 4.6.6, “mysql\_config\_editor — MySQL Configuration Utility”

Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”

Section 4.3.4, “mysqld\_multi — Manage Multiple MySQL Servers”

Section 21.8.2, “Example C API Client Programs”

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

Section C.5.1, “How to Determine What Is Causing a Problem”

Section 7.6.3, “How to Repair MyISAM Tables”

Section 1.7, “How to Report Bugs or Problems”

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 5.2, “MySQL Server Logs”

Section 4.1, “Overview of MySQL Programs”

Section 6.3.6, “Password Expiration and Sandbox Mode”

Section 6.2.1, “Privileges Provided by MySQL”

Section E.9, “Restrictions on Pluggable Authentication”

Section 5.6.3, “Running Multiple MySQL Instances on Unix”

Section 2.10.2, “Securing the Initial MySQL Accounts”

Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 2.3.4.8, “Starting MySQL as a Windows Service”

Section 2.3.4.5, “Starting MySQL from the Windows Command Line”

Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”

Section 5.1.12, “The Shutdown Process”

Section 8.11.2, “Tuning Server Parameters”

Section 2.10.1, “Unix Postinstallation Procedures”

Section 2.3.6, “Upgrading MySQL on Windows”

Section 4.2.3.3, “Using Option Files”

Section 4.2.3.1, “Using Options on the Command Line”

Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## mysqladmin debug

Section 13.7.1.4, “GRANT Syntax”

Section 22.4.1, “Debugging a MySQL Server”

Section 18.4.5, “Event Scheduler Status”

Section 6.2.1, “Privileges Provided by MySQL”

## mysqladmin extended-status

Section 13.2.5.2, “INSERT DELAYED Syntax”

Section 13.7.5.36, “SHOW STATUS Syntax”

## mysqladmin flush-hosts

Section C.5.2.6, “Host '*host\_name*' is blocked”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”

Section 5.1.4, “Server System Variables”

## mysqladmin flush-logs

Section 7.3.3, “Backup Strategy Summary”

Section 7.3.1, “Establishing a Backup Policy”

Section 5.2.6, “Server Log Maintenance”

Section 5.2.4, “The Binary Log”

Section 5.2.2, “The Error Log”

Section 16.2.2.1, “The Slave Relay Log”

## mysqladmin flush-privileges

Section 6.2.7, “Causes of Access-Denied Errors”

Section 2.11.5, “Copying MySQL Databases to Another Machine”

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 6.2.2, “Privilege System Grant Tables”

Section 2.10.1.1, “Problems Running

mysql\_install\_db”



---

Section 5.1.3, “Server Command Options”  
Section 6.2.6, “When Privilege Changes Take Effect”

### **mysqladmin flush-tables**

Section 8.6.2, “Bulk Data Loading for MyISAM Tables”  
Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”  
Section 8.10.5, “External Locking”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 7.6.1, “Using myisamchk for Crash Recovery”

### **mysqladmin flush-xxx**

Section 6.3.2, “Adding User Accounts”

### **mysqladmin kill**

Section 13.7.6.4, “KILL Syntax”  
Section C.5.2.9, “MySQL server has gone away”  
Section C.5.4.3, “How MySQL Handles a Full Disk”  
Section 12.16, “Miscellaneous Functions”  
Section 6.2.1, “Privileges Provided by MySQL”

### **mysqladmin password**

Section 6.3.5, “Assigning Account Passwords”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 6.3.6, “Password Expiration and Sandbox Mode”

### **mysqladmin processlist**

Section 13.7.6.4, “KILL Syntax”  
Section 21.8.6.43, “mysql\_list\_processes()”  
Section 13.7.5.30, “SHOW PROCESSLIST Syntax”  
Section 6.3.2, “Adding User Accounts”  
Section 8.12.5, “Examining Thread Information”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 22.1.1, “MySQL Threads”  
Section 6.2.1, “Privileges Provided by MySQL”

### **mysqladmin processlist status**

Section 22.4.1, “Debugging a MySQL Server”

### **mysqladmin refresh**

Section 6.3.2, “Adding User Accounts”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”  
Section 5.2.6, “Server Log Maintenance”

### **mysqladmin reload**

Section 6.3.2, “Adding User Accounts”  
Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

Section 1.7, “How to Report Bugs or Problems”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 2.10.1.1, “Problems Running mysql\_install\_db”  
Section 5.1.3, “Server Command Options”  
Section 6.3.4, “Setting Account Resource Limits”  
Section 6.2.6, “When Privilege Changes Take Effect”

### **mysqladmin reload version**

Section 1.7, “How to Report Bugs or Problems”

### **mysqladmin shutdown**

Section 13.7.1.4, “GRANT Syntax”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 14.2.2.1, “Creating the InnoDB Tablespace”  
Section 22.4.1.2, “Creating Trace Files”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 6.1.5, “How to Run MySQL as a Normal User”  
Section 2.4.2, “Installing MySQL on Mac OS X Using Native Packages”  
Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.21, “Replication and Temporary Tables”  
Section 2.10.1.2, “Starting and Stopping MySQL Automatically”  
Section 2.3.4.8, “Starting MySQL as a Windows Service”  
Section 5.1.12, “The Shutdown Process”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

### **mysqladmin status**

Section 21.8.6.69, “mysql\_stat()”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”

### **mysqladmin variables**

Section C.5.2.9, “MySQL server has gone away”  
Section 13.7.5.40, “SHOW VARIABLES Syntax”

### **mysqladmin variables extended-status processlist**

Section 1.7, “How to Report Bugs or Problems”

### **mysqladmin ver**

Section 22.4.1.1, “Compiling MySQL for Debugging”

### **mysqladmin version**

Section C.5.2.2, “Can't connect to [local] MySQL server”

---

Section C.5.2.9, “MySQL server has gone away”  
Section 1.7, “How to Report Bugs or Problems”  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## mysqlanalyze

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

## mysqlbackup

Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”  
Section 7.1, “Backup and Recovery Types”  
Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 4.5.4, “mysqldump — A Database Backup Program”

## mysqlbinlog

Section 4.5.1.1, “mysql Options”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”  
Section 4.6.8.1, “mysqlbinlog Hex Dump Format”  
Section 4.6.8.2, “mysqlbinlog Row Event Display”  
Section 13.7.6.1, “BINLOG Syntax”  
Section 13.7.5.3, “SHOW BINLOG EVENTS Syntax”  
Section 13.7.5.33, “SHOW RELAYLOG EVENTS Syntax”  
Section 13.4.2.5, “START SLAVE Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.7.1.2, “Command Probes”  
Section 16.1.3.1, “GTID Concepts”  
Section 16.4.5, “How to Report Replication Bugs or Problems”  
Section C.5.8, “Known Issues in MySQL”  
Section 12.16, “Miscellaneous Functions”  
Section 4.1, “Overview of MySQL Programs”  
Section 7.5, “Point-in-Time (Incremental) Recovery Using the Binary Log”  
Section 7.5.2, “Point-in-Time Recovery Using Event Positions”  
Section 7.5.1, “Point-in-Time Recovery Using Event Times”  
Section 16.4.1.32, “Replication and Variables”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section 4.6.8.4, “Specifying the mysqlbinlog Server ID”  
Section 5.2.4, “The Binary Log”  
Section 16.2.2.1, “The Slave Relay Log”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

Section 4.6.8.3, “Using mysqlbinlog to Back Up Binary Log Files”  
Section 7.3.2, “Using Backups for Recovery”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 1.4, “What Is New in MySQL 5.6”

## mysqlbinlog binary-log-file | mysql

Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”

## mysqlbinlog |mysql

Section C.5.8, “Known Issues in MySQL”

## mysqlbug

Section 4.4.2, “mysqlbug — Generate Bug Report”

## mysqlcheck

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 13.1.1, “ALTER DATABASE Syntax”  
Section 7.6, “MyISAM Table Maintenance and Crash Recovery”  
Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”  
Section 4.7.3, “my\_print\_defaults — Display Options from Option Files”  
Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”  
Section 17.3.4, “Maintenance of Partitions”  
Section 9.2.3, “Mapping of Identifiers to File Names”  
Section 4.1, “Overview of MySQL Programs”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.4, “Server System Variables”  
Section 14.3, “The MyISAM Storage Engine”

## mysqld

Section 4.6.3.2, “myisamchk Check Options”  
Section 4.6.3.1, “myisamchk General Options”  
Section 4.3.1, “mysqld — The MySQL Server”  
Section 5.7.1, “mysqld DTrace Probe Reference”  
Section C.5.2.18, “'File' Not Found and Similar Errors”  
Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section C.5.2.13, “Can't create/write to file”  
Section C.5.2.17, “Can't initialize character set”  
Section C.5.2.4, “Client does not support authentication protocol”  
Section 13.7.3.1, “CREATE FUNCTION Syntax for User-Defined Functions”  
Section 13.7.6.3, “FLUSH Syntax”

---

Section C.5.2.6, “Host ‘*host\_name*’ is blocked”

Section C.5.2.15, “Ignoring user”

Section 14.2.4.5, “InnoDB General Troubleshooting”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 13.2.5.2, “INSERT DELAYED Syntax”

Section 13.7.6.4, “KILL Syntax”

Section 13.2.6, “LOAD DATA INFILE Syntax”

Section 14.3.1, “MyISAM Startup Options”

Section C.5.2.9, “MySQL server has gone away”

Section 21.8.6.1, “mysql\_affected\_rows()”

Section 21.8.6.49, “mysql\_options()”

Section 13.7.2.4, “OPTIMIZE TABLE Syntax”

Section 13.7.2.5, “REPAIR TABLE Syntax”

Section 13.4.2.3, “RESET SLAVE Syntax”

Section 13.2.9.1, “SELECT ... INTO Syntax”

Section 13.3.6, “SET TRANSACTION Syntax”

Section 14.2.4.4, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”

Section 13.7.5.40, “SHOW VARIABLES Syntax”

Section C.5.2.7, “Too many connections”

Section 22.3.2, “Adding a New User-Defined Function”

Section 22.3, “Adding New Functions to MySQL”

Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”

Section 16.1.4.4, “Binary Log Options and Variables”

Section 5.2.4.1, “Binary Logging Formats”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 2.1.2.2, “Choosing a Distribution Format”

Section 9.6, “Comment Syntax”

Section C.5.2.11, “Communication Errors and Aborted Connections”

Section 22.3.2.5, “Compiling and Installing User-Defined Functions”

Section 22.4.1.1, “Compiling MySQL for Debugging”

Section 14.2.1.2, “Configuring InnoDB”

Section 6.3.9.2, “Configuring MySQL for SSL”

Section 14.3.4.1, “Corrupted MyISAM Tables”

Section 14.2.2.1, “Creating the InnoDB Tablespace”

Section 22.4.1.2, “Creating Trace Files”

Section 22.4.1.4, “Debugging mysqld under gdb”

Section 22.4.1, “Debugging a MySQL Server”

Section 22.4, “Debugging and Porting MySQL”

Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”

Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.7.4, “resolve\_stack\_dump — Resolve Numeric Stack Trace Dump to Symbols”

Section 4.4.1, “comp\_err — Compile MySQL Error Message File”

Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”

Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”

Section 4.3.3, “mysql.server — MySQL Server Startup Script”

Section 4.3.4, “mysqld\_multi — Manage Multiple MySQL Servers”

Section 4.3.2, “mysqld\_safe — MySQL Server Startup Script”

Section 2.12, “Environment Variables”

Section 8.10.5, “External Locking”

Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

Section 8.12.5.2, “General Thread States”

Section 8.4.3.1, “How MySQL Opens and Closes Tables”

Section 8.11.4.1, “How MySQL Uses Memory”

Section 8.2.1.4, “How to Avoid Full Table Scans”

Section C.5.1, “How to Determine What Is Causing a Problem”

Section 7.6.3, “How to Repair MyISAM Tables”

Section 1.7, “How to Report Bugs or Problems”

Section 6.1.5, “How to Run MySQL as a Normal User”

Section 9.2.2, “Identifier Case Sensitivity”

Section 12.14, “Information Functions”

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 2.4.2, “Installing MySQL on Mac OS X Using Native Packages”

Section 2.6, “Installing MySQL on Solaris and OpenSolaris”

Section 22.4.1.7, “Making a Test Case If You Experience Table Corruption”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 12.16, “Miscellaneous Functions”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 14.2.2.5, “Moving or Copying InnoDB Tables to Another Machine”

Section 2.3.1, “MySQL Installation Layout on Microsoft Windows”

Chapter 5, *MySQL Server Administration*

Section 4.3, “MySQL Server and Server-Startup Programs”

Section 5.2, “MySQL Server Logs”

Section 10.6, “MySQL Server Time Zone Support”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 22.1.1, “MySQL Threads”

---

---

Section C.5.6, “Optimizer-Related Issues”  
Section 21.7.3, “Options with the Embedded Server”  
Section 4.1, “Overview of MySQL Programs”  
Section C.5.2.10, “Packet Too Large”  
Section 20.2.2, “Performance Schema Startup Configuration”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”  
Section 2.10.1.1, “Problems Running `mysql_install_db`”  
Section C.5.3.1, “Problems with File Permissions”  
Section 4.2.3.2, “Program Option Modifiers”  
Section 8.9.3.3, “Query Cache Configuration”  
Section 16.1.4.1, “Replication and Binary Logging Option and Variable Reference”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 16.2.2, “Replication Relay and Status Logs”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section C.5.4.1.3, “Resetting the Root Password: Generic Instructions”  
Section C.5.4.1.2, “Resetting the Root Password: Unix Systems”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section C.5.5.5, “Rollback Failure for Nontransactional Tables”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 1.8.3, “Running MySQL in ANSI Mode”  
Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”  
Section 6.1.4, “Security-Related `mysqld` Options and Variables”  
Section 2.3.4.3, “Selecting a MySQL Server Type”  
Section 1.8.2, “Selecting SQL Modes”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 5.1.3, “Server Command Options”  
Section 22.2.4.2.2, “Server Plugin Status and System Variables”  
Section 5.1.11, “Server Response to Signals”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 10.2, “Setting the Error Message Language”  
Section 7.6.5, “Setting Up a `MyISAM` Table Maintenance Schedule”  
Section 16.1.3.2, “Setting Up Replication Using GTIDs”  
Section 14.2.4.6, “Starting `InnoDB` on a Corrupted Database”  
Section 2.10.1.2, “Starting and Stopping MySQL Automatically”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”  
Section 5.6.2.1, “Starting Multiple MySQL Instances at the Windows Command Line”  
Section 2.3.4.8, “Starting MySQL as a Windows Service”  
Section 2.3.4.5, “Starting MySQL from the Windows Command Line”  
Section 10.6.1, “Staying Current with Time Zone Changes”  
Section 8.10.2, “Table Locking Issues”  
Section C.5.2.19, “Table-Corruption Issues”  
Section 2.3.4.9, “Testing The MySQL Installation”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 14.3, “The `MyISAM` Storage Engine”  
Section 5.2.4, “The Binary Log”  
Section 5.2.2, “The Error Log”  
Section 5.2.3, “The General Query Log”  
Section 8.9.3, “The MySQL Query Cache”  
Section 5.1, “The MySQL Server”  
Section 22.1.2, “The MySQL Test Suite”  
Section 5.2.5, “The Slow Query Log”  
Section C.5.4.6, “Time Zone Problems”  
Section 5.7, “Tracing `mysqld` Using DTrace”  
Section 14.2.4.5.1, “Troubleshooting `InnoDB` I/O Problems”  
Section 2.3.5, “Troubleshooting a Microsoft Windows MySQL Server Installation”  
Section 8.11.2, “Tuning Server Parameters”  
Section 1.2, “Typographical and Syntax Conventions”  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section 2.11.1, “Upgrading MySQL”  
Section 2.3.6, “Upgrading MySQL on Windows”  
Section 22.3.2.6, “User-Defined Function Security Precautions”  
Section 7.6.1, “Using `myisamchk` for Crash Recovery”  
Section 22.4.1.3, “Using `pdb` to create a Windows crashdump”  
Section 22.4.1.5, “Using a Stack Trace”  
Section 4.2.3.3, “Using Option Files”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in `mysqld`”  
Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”  
Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”  
Section 2.4.5, “Using the Bundled MySQL on Mac OS X Server”  
Section 14.2.9.6, “Using the `InnoDB` memcached Plugin with Replication”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”  
Section 6.2.6, “When Privilege Changes Take Effect”  
Section C.5.4.4, “Where MySQL Stores Temporary Files”  
Section 22.2.4, “Writing Plugins”

---



---

## **mysqld mysqld.trace**

Section 22.4.1.2, “Creating Trace Files”

## **mysqld-**

Section 2.11.1, “Upgrading MySQL”

## **mysqld-debug**

Section 2.1.2.2, “Choosing a Distribution Format”

Section 22.4.1.2, “Creating Trace Files”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

Section 2.3.4.3, “Selecting a MySQL Server Type”

Section 5.6.2.1, “Starting Multiple MySQL Instances at the Windows Command Line”

## **mysqld\_multi**

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 4.1, “Overview of MySQL Programs”

Section 5.6.3, “Running Multiple MySQL Instances on Unix”

## **mysqld\_safe**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section C.5.2.18, “`'File' Not Found` and Similar Errors”

Section 14.2.4.5, “InnoDB General Troubleshooting”

Section 22.4.1.1, “Compiling MySQL for Debugging”

Section 14.2.2.1, “Creating the InnoDB Tablespace”

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 8.11.4.2, “Enabling Large Page Support”

Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 10.6, “MySQL Server Time Zone Support”

Section 4.2.3.6, “Option Defaults, Options Expecting Values, and the = Sign”

Section 4.1, “Overview of MySQL Programs”

Section C.5.2.10, “Packet Too Large”

Section C.5.3.1, “Problems with File Permissions”

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 5.6.3, “Running Multiple MySQL Instances on Unix”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”

Section 5.2.2, “The Error Log”

Section C.5.4.6, “Time Zone Problems”

Section 8.11.2, “Tuning Server Parameters”

Section 2.10.1, “Unix Postinstallation Procedures”

Section 5.1.2.2, “Using a Sample Default Server Configuration File”

Section 4.2.3.3, “Using Option Files”

## **mysqld\_safe options**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **mysqldump**

Section 4.5.1.1, “`mysql` Options”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.5, “`mysqldump` Tips”

Section 13.1.14, “`CREATE TABLE` Syntax”

Section 13.2.6, “`LOAD DATA INFILE` Syntax”

Section 13.2.7, “`LOAD XML` Syntax”

Section 13.7.3.4, “`UNINSTALL PLUGIN` Syntax”

Section 11.3.4, “`YEAR(2)` Limitations and Migrating to `YEAR(4)`”

Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”

Section 16.3.1.3, “Backing Up a Master or Slave by Making It Read Only”

Section 16.3.1.1, “Backing Up a Slave Using `mysqldump`”

Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”

Chapter 7, *Backup and Recovery*

Section 7.1, “Backup and Recovery Types”

Section 7.3.3, “Backup Strategy Summary”

Section 8.5.4, “Bulk Data Loading for InnoDB Tables”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”

Section 4.2.2, “Connecting to the MySQL Server”

Section 7.4.5.2, “Copy a Database from one Server to Another”

Section 2.11.5, “Copying MySQL Databases to Another Machine”

Section 5.4.1.3, “Copying Tablespaces to Another Server (Transportable Tablespaces)”

Section 16.1.1.5, “Creating a Data Snapshot Using `mysqldump`”

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

Section 5.4, “Creating and Using InnoDB Tables and Indexes”

Section 2.3.4.7, “Customizing the PATH for MySQL Tools”

---

Section 7.2, “Database Backup Methods”  
Section 5.3.4, “Defragmenting a Table”  
Section 4.6.1, “`innochecksum` — Offline InnoDB File Checksum Utility”  
Section 2.11.2, “Downgrading MySQL”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”  
Section 7.4.1, “Dumping Data in SQL Format with `mysqldump`”  
Section 7.4.5.3, “Dumping Stored Programs”  
Section 7.4.5.4, “Dumping Table Definitions and Content Separately”  
Section 5.4.1.1, “Enabling and Disabling File-Per-Table Mode”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 7.3, “Example Backup and Recovery Strategy”  
Section 1.8.5.4, “Foreign Key Differences”  
Section 1.7, “How to Report Bugs or Problems”  
Section 16.1.1, “How to Set Up Replication”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 7.4.5.1, “Making a Copy of a Database”  
Section 9.2.3, “Mapping of Identifiers to File Names”  
Section 14.2.2.5, “Moving or Copying InnoDB Tables to Another Machine”  
Section 5.2, “MySQL Server Logs”  
Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”  
Section 4.1, “Overview of MySQL Programs”  
Section C.5.5.8, “Problems with Floating-Point Values”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 7.4.4, “Reloading Delimited-Text Format Backups”  
Section 7.4.2, “Reloading SQL-Format Backups”  
Section 16.3.4, “Replicating Different Databases to Different Slaves”  
Section 16.2.2, “Replication Relay and Status Logs”  
Section E.8, “Restrictions on Performance Schema”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section E.5, “Restrictions on Views”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.2.6, “Server Log Maintenance”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 16.1.1.8, “Setting Up Replication with Existing Data”  
Section C.5.5.7, “Solving Problems with No Matching Rows”  
Section 4.2.3, “Specifying Program Options”  
Section 2.3.4.8, “Starting MySQL as a Windows Service”  
Section 11.4.3, “The `BLOB` and `TEXT` Types”  
Section 8.9.1, “The InnoDB Buffer Pool”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 10.1.11, “Upgrading from Previous to Current Unicode Support”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”  
Section 7.4, “Using `mysqldump` for Backups”  
Section 7.4.5.5, “Using `mysqldump` to Test for Upgrade Incompatibilities”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 16.3.1, “Using Replication for Backups”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”  
Section 14.2.9.6, “Using the InnoDB `memcached` Plugin with Replication”  
Section 12.11, “XML Functions”

## `mysqldump mysql`

Section 6.2.7, “Causes of Access-Denied Errors”

## `mysqldumpslow`

Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.1, “Overview of MySQL Programs”  
Section 5.2.5, “The Slow Query Log”

## `mysqlhotcopy`

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Chapter 7, *Backup and Recovery*  
Section 7.1, “Backup and Recovery Types”  
Section 7.2, “Database Backup Methods”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.1, “Overview of MySQL Programs”

## `mysqlhotcopy arguments`

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## `mysqlimport`

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 7.1, “Backup and Recovery Types”  
Section 2.11.5, “Copying MySQL Databases to Another Machine”  
Section 7.2, “Database Backup Methods”  
Section 2.11.2, “Downgrading MySQL”  
Section 4.1, “Overview of MySQL Programs”  
Section 7.4.4, “Reloading Delimited-Text Format Backups”  
Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”

## `mysqloptimize`

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

---

## mysqlrepair

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

## mysqlshow

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

Section 13.7.5.15, “SHOW DATABASES Syntax”

Section 13.7.5.23, “SHOW INDEX Syntax”

Section 13.7.5.37, “SHOW TABLE STATUS Syntax”

Section 4.2.2, “Connecting to the MySQL Server”

Section 21.8.2, “Example C API Client Programs”

Section 4.1, “Overview of MySQL Programs”

Section 2.3.4.9, “Testing The MySQL Installation”

Section 2.3.7, “Windows Postinstallation Procedures”

## mysqlshow db\_name

Section 13.7.5.38, “SHOW TABLES Syntax”

## mysqlshow db\_name tbl\_name

Section 13.7.5.6, “SHOW COLUMNS Syntax”

## mysqlshow mysql user

Section C.5.2.15, “Ignoring user”

## mysqlslap

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 4.1, “Overview of MySQL Programs”

Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”

Section 8.12.3, “Using Your Own Benchmarks”

## mysqltest

Section 6.3.6, “Password Expiration and Sandbox Mode”

Section 22.1.2, “The MySQL Test Suite”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

## N

[index top [3911]]

## ndb\_restore

Section 7.1, “Backup and Recovery Types”

## ndb\_size.pl

Section 11.6, “Data Type Storage Requirements”

## NET

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## NET START

Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”

## NET START MySQL

Section 2.3.4.8, “Starting MySQL as a Windows Service”

Section 2.3.5, “Troubleshooting a Microsoft Windows MySQL Server Installation”

Section 2.3.6, “Upgrading MySQL on Windows”

## NET STOP

Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”

## NET STOP MySQL

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## Netinfo Manager

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

## nm

Section 4.7.4, “resolve\_stack\_dump — Resolve Numeric Stack Trace Dump to Symbols”

Section 22.4.1.5, “Using a Stack Trace”

## O

[index top [3911]]

## openssl

Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”

## openssl md5 package\_name

Section 2.1.4.1, “Verifying the MD5 Checksum”

## P

[index top [3911]]

## perror

Section 4.8.1, “perror — Explain Error Codes”

Section C.5.2.18, “'File' Not Found and Similar Errors”

Section C.5.2.13, “Can't create/write to file”

Section 7.6.3, “How to Repair MyISAM Tables”

Section 14.2.3.14.2, “Operating System Error Codes”

Section 4.1, “Overview of MySQL Programs”

Section C.1, “Sources of Error Information”

## pfexec

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”

---

## PGP

Section 2.1.4.2, “Signature Checking Using `GnuPG`”

## ping6

Section 5.1.9.5, “Obtaining an IPv6 Address from a Broker”

## pkg

Section 2.6.2, “Installing MySQL on OpenSolaris using IPS”

## pkgadd

Section 2.6.1, “Installing MySQL on Solaris using a Solaris `PKG`”

## pkgrm

Section 2.6.1, “Installing MySQL on Solaris using a Solaris `PKG`”

## ppm

Section 2.13, “Perl Installation Notes”

## ps

Section 4.6.16, “`mysql_zap` — Kill Processes That Match a Pattern”

Section 6.1.2.1, “End-User Guidelines for Password Security”

Section 8.11.4.1, “How MySQL Uses Memory”

Section C.5.1, “How to Determine What Is Causing a Problem”

Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”

## ps auxw

Section 4.2.2, “Connecting to the MySQL Server”

## ps xa | grep mysqld

Section C.5.2.2, “Can't connect to [local] MySQL server”

## R

[index top [3911]]

## rename

Section 5.2.6, “Server Log Maintenance”

Section 5.2.2, “The Error Log”

Section 5.2.3, “The General Query Log”

## replace

Section 1.8.5.5, “`--` as the Start of a Comment”

Section 4.8.2, “`replace` — A String-Replacement Utility”

Section 4.7.1, “`mysql2mysql` — Convert mSQL Programs for Use with MySQL”

Section 4.1, “Overview of MySQL Programs”

Section 16.3.3, “Using Replication for Scale-Out”

## replace arguments

Section 4.8.2, “`replace` — A String-Replacement Utility”

## resolve\_stack\_dump

Section 4.7.4, “`resolve_stack_dump` — Resolve Numeric Stack Trace Dump to Symbols”

Section 4.1, “Overview of MySQL Programs”

Section 22.4.1.5, “Using a Stack Trace”

## resolveip

Section 4.8.3, “`resolveip` — Resolve Host name to IP Address or Vice Versa”

Section 4.1, “Overview of MySQL Programs”

## rm

Section 13.4.1.1, “`PURGE BINARY LOGS` Syntax”

## rpm

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”

Section 2.1.4.4, “Signature Checking Using `RPM`”

## rpmbuild

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”

Section 2.9, “Installing MySQL from Source”

## rsync

Section 7.1, “Backup and Recovery Types”

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”

## S

[index top [3911]]

## scp

Section 7.1, “Backup and Recovery Types”

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

## sed

Section 3.3.4.7, “Pattern Matching”

## service

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

## Service Control Manager

Section 2.3, “Installing MySQL on Microsoft Windows”



---

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## Services

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## setenv

Section 4.2.4, “Setting Environment Variables”

## sh

Section C.5.2.18, “`'File' Not Found` and Similar Errors”

Section 4.2.1, “Invoking MySQL Programs”

Section 4.2.4, “Setting Environment Variables”

Section 1.2, “Typographical and Syntax Conventions”

## sleep

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## Start>Run>cmd.exe

Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”

## strings

Section 6.1.1, “Security Guidelines”

## sudo

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

## swinstall

Section 2.7.2, “Installing MySQL on HP-UX using DEPOT”

## System Preferences...

Section 2.4.4, “Installing and Using the MySQL Preference Pane”

## T

[index top [3911]]

## tar

Section 16.3.1.2, “Backing Up Raw Data from a Slave”

Section 7.1, “Backup and Recovery Types”

Section 2.1.2.2, “Choosing a Distribution Format”

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”

Section 3.3, “Creating and Using a Database”

Section 2.7.1, “General Notes on Installing MySQL on HP-UX”

Section 1.7, “How to Report Bugs or Problems”

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”

Section 2.9, “Installing MySQL from Source”

Section 2.4, “Installing MySQL on Mac OS X”

Section 2.6, “Installing MySQL on Solaris and OpenSolaris”

Section 2.13.1, “Installing Perl on Unix”

Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”

## tcpdump

Section 6.1.1, “Security Guidelines”

## tcsh

Section C.5.2.18, “`'File' Not Found` and Similar Errors”

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

Section 4.2.1, “Invoking MySQL Programs”

Section 4.2.4, “Setting Environment Variables”

Section 1.2, “Typographical and Syntax Conventions”

## tee

Section 4.5.1.2, “`mysql` Commands”

## telnet

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

Section 6.1.1, “Security Guidelines”

Section 14.2.9.3.3, “Verifying the InnoDB and memcached Setup”

## Terminal

Section 2.4, “Installing MySQL on Mac OS X”

## Text in this style

Section 1.2, “Typographical and Syntax Conventions”

## top

Section C.5.1, “How to Determine What Is Causing a Problem”

## U

[index top [3911]]

## ulimit

Section C.5.2.18, “`'File' Not Found` and Similar Errors”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 8.11.4.2, “Enabling Large Page Support”

Section C.5.2.10, “Packet Too Large”

---

Section 14.2.9.3.1, “Prerequisites for the InnoDB memcached Plugin”  
Section 5.1.3, “Server Command Options”

## **useradd**

Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”  
Section 2.5.1, “Installing MySQL from RPM Packages on Linux”  
Section 2.6, “Installing MySQL on Solaris and OpenSolaris”

## **usermod**

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

## **V**

[index top [3911]]

## **vi**

Section 4.5.1.2, “mysql Commands”  
Section 3.3.4.7, “Pattern Matching”

## **W**

[index top [3911]]

## **WinDbg**

Section 22.4.1.3, “Using pdb to create a Windows crashdump”

## **winMd5Sum**

Section 2.1.4.1, “Verifying the MD5 Checksum”

## **WinZip**

Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.9, “Installing MySQL from Source”

## **WordPad**

Section 13.2.6, “LOAD DATA INFILE Syntax”

## **Y**

[index top [3911]]

## **yacc**

Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 9.3, “Reserved Words”

## **yum**

Section 2.5.1, “Installing MySQL from RPM Packages on Linux”

Section 2.5.2, “Installing MySQL on Linux using Native Package Manager”

## **Z**

[index top [3911]]

## **zip**

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 1.7, “How to Report Bugs or Problems”

## **zsh**

Section 4.2.4, “Setting Environment Variables”

---

# Function Index

[Symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#) | [Y](#)

## Symbols

[[index top \[3931\]](#)]

%

Section 1.8.4, “MySQL Extensions to Standard SQL”

## A

[[index top \[3931\]](#)]

### **ABS ( )**

Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”

Section 22.3, “Adding New Functions to MySQL”

Section 12.6.2, “Mathematical Functions”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **ACOS ( )**

Section 12.6.2, “Mathematical Functions”

### **ADDDATE ( )**

Section 12.7, “Date and Time Functions”

### **addslashes ( )**

Section 6.1.7, “Client Programming Security Guidelines”

### **ADDTIME ( )**

Section 12.7, “Date and Time Functions”

### **AES\_DECRYPT ( )**

Section 12.13, “Encryption and Compression Functions”

### **AES\_ENCRYPT ( )**

Section 12.13, “Encryption and Compression Functions”

### **Area ( )**

Section 12.18.5.2, “Geometry Property Functions”

Section 12.18.5.2.6, “MultiPolygon Functions”

Section 12.18.5.2.5, “Polygon Functions”

### **AsBinary ( )**

Section 12.18.4.5, “Fetching Spatial Data”

Section 12.18.5.1, “Geometry Format Conversion Functions”

### **ASCII ( )**

Section 13.8.3, “`HELP` Syntax”

Section 12.5, “String Functions”

### **ASIN ( )**

Section 12.6.2, “Mathematical Functions”

### **AsText ( )**

Section 12.18.4.5, “Fetching Spatial Data”

Section 12.18.5.1, “Geometry Format Conversion Functions”

### **AsWKB ( )**

Section 12.18.5.1, “Geometry Format Conversion Functions”

### **AsWKT ( )**

Section 12.18.5.1, “Geometry Format Conversion Functions”

### **ATAN ( )**

Section 12.6.2, “Mathematical Functions”

### **ATAN2 ( )**

Section 12.6.2, “Mathematical Functions”

### **AVG ( )**

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”

Section 11.1.2, “Date and Time Type Overview”

Section 8.13.14.1, “Loose Index Scan”

Section 11.4.4, “The `ENUM` Type”

Section 11.4.5, “The `SET` Type”

## B

[[index top \[3931\]](#)]

### **BdMPolyFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

### **BdMPolyFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

### **BdPolyFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

### **BdPolyFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

### **BENCHMARK ( )**

Section 8.9.3.1, “How the Query Cache Operates”

Section 12.14, “Information Functions”

---

Section 8.12.1, “Measuring the Speed of Expressions and Functions”

Section 13.2.10.10, “Optimizing Subqueries”

Section 13.2.10.8, “Subqueries in the `FROM` Clause”

## **BIN( )**

Section 9.1.6, “Bit-Field Literals”

Section 12.5, “String Functions”

## **BIT\_AND( )**

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

## **BIT\_COUNT( )**

Section 12.12, “Bit Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

## **BIT\_LENGTH( )**

Section 12.5, “String Functions”

## **BIT\_OR( )**

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

## **BIT\_XOR( )**

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

## **Boundary( )**

Section 12.18.5.2.1, “General Geometry Functions”

## **Buffer( )**

Section 12.18.5.3.2, “Spatial Operators”

## **C**

[index top [3931]]

## **CAST( )**

Section 10.1.9.2, “`CONVERT( )` and `CAST( )`”

Section 9.1.6, “Bit-Field Literals”

Section 12.10, “Cast Functions and Operators”

Section 12.3.2, “Comparison Functions and Operators”

Section 11.3.7, “Conversion Between Date and Time Types”

Section 12.7, “Date and Time Functions”

Section 9.1.4, “Hexadecimal Literals”

Section 1.8.5, “MySQL Differences from Standard SQL”

Section 10.1.9.1, “Result Strings”

Section 10.1.7.7, “The `BINARY` Operator”

Section 12.2, “Type Conversion in Expression Evaluation”

Section 9.4, “User-Defined Variables”

## **CEIL( )**

Section 12.6.2, “Mathematical Functions”

## **CEILING( )**

Section 17.2.4.1, “`LINEAR HASH` Partitioning”

Section 12.6.2, “Mathematical Functions”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **Centroid( )**

Section 12.18.5.2.6, “`MultiPolygon` Functions”

## **CHAR( )**

Section 12.10, “Cast Functions and Operators”

Section 12.13, “Encryption and Compression Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 12.5, “String Functions”

## **CHAR\_LENGTH( )**

Section 12.5, “String Functions”

## **CHARACTER\_LENGTH( )**

Section 12.5, “String Functions”

## **CHARSET( )**

Section 12.14, “Information Functions”

Section 10.1.9.1, “Result Strings”

## **COALESCE( )**

Section 13.2.9.2, “`JOIN` Syntax”

Section 12.3.2, “Comparison Functions and Operators”

## **COERCIBILITY( )**

Section 10.1.7.5, “Collation of Expressions”

Section 12.14, “Information Functions”

## **COLLATION( )**

Section C.5.5.1, “Case Sensitivity in String Searches”

Section 12.14, “Information Functions”

Section 10.1.9.1, “Result Strings”

## **COMPRESS( )**

Section 12.13, “Encryption and Compression Functions”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 5.1.4, “Server System Variables”

## **CONCAT( )**

Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”

Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”

Section 12.10, “Cast Functions and Operators”

Section 10.1.7.5, “Collation of Expressions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 10.1.9.1, “Result Strings”

Section 5.1.7, “Server SQL Modes”

Section 12.5, “String Functions”

---

Section 10.1.8, “String Repertoire”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 12.2, “Type Conversion in Expression Evaluation”  
Section 12.11, “XML Functions”

## **CONCAT\_WS ( )**

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”  
Section 12.5, “String Functions”

## **CONNECTION\_ID ( )**

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section 20.9.9.3, “The `threads` Table”  
Section 6.3.11.2, “The Audit Log File”

## **Contains ( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”  
Section 1.4, “What Is New in MySQL 5.6”

## **CONV ( )**

Section 12.6.2, “Mathematical Functions”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **CONVERT ( )**

Section 10.1.9.2, “`CONVERT ( )` and `CAST ( )`”  
Section 12.10, “Cast Functions and Operators”  
Section 10.1.3.5, “Character String Literal Character Set and Collation”  
Section 12.3.2, “Comparison Functions and Operators”

## **CONVERT\_TZ ( )**

Section 12.7, “Date and Time Functions”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 16.4.1.28, “Replication and Time Zones”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

## **ConvexHull ( )**

Section 12.18.5.3.2, “Spatial Operators”

## **COS ( )**

Section 12.6.2, “Mathematical Functions”

## **COT ( )**

Section 12.6.2, “Mathematical Functions”

## **COUNT ( )**

Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”  
Section 8.8.2, “`EXPLAIN` Output Format”

Section 12.17.1, “`GROUP BY` (Aggregate) Functions”  
Section 3.3.4.8, “Counting Rows”  
Section 8.2.1.2, “How MySQL Optimizes `WHERE` Clauses”  
Section 8.13.14.1, “Loose Index Scan”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section C.5.5.3, “Problems with `NULL` Values”  
Section 5.1.7, “Server SQL Modes”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 18.5.2, “View Processing Algorithms”

## **CRC32 ( )**

Section 12.6.2, “Mathematical Functions”

## **Crosses ( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **crypt ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 5.1.4, “Server System Variables”

## **CURDATE ( )**

Section 12.7, “Date and Time Functions”  
Section 3.3.4.5, “Date Calculations”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”

## **CURRENT\_DATE**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 11.5, “Data Type Default Values”  
Section 12.7, “Date and Time Functions”

## **CURRENT\_DATE ( )**

Section 11.3.7, “Conversion Between Date and Time Types”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”

## **CURRENT\_TIME**

Section 12.7, “Date and Time Functions”

## **CURRENT\_TIME ( )**

Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”

## **CURRENT\_TIMESTAMP**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.1.14, “`CREATE TABLE` Syntax”

---

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 11.5, “Data Type Default Values”  
Section 12.7, “Date and Time Functions”

### **CURRENT\_TIMESTAMP ( )**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”

### **CURRENT\_USER**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.1.15, “`CREATE TRIGGER` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 18.6, “Access Control for Stored Programs and Views”  
Section 12.14, “Information Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 16.4.1.14, “Replication and System Functions”  
Section 16.4.1.6, “Replication of `CURRENT_USER ( )`”  
Section 6.2.3, “Specifying Account Names”

### **CURRENT\_USER ( )**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.1.15, “`CREATE TRIGGER` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 6.2.4, “Access Control, Stage 1: Connection Verification”  
Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”  
Section 12.14, “Information Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 6.3.8, “Proxy Users”  
Section 16.4.1.14, “Replication and System Functions”  
Section 16.4.1.6, “Replication of `CURRENT_USER ( )`”  
Section 6.2.3, “Specifying Account Names”  
Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”  
Section 10.1.12, “UTF-8 for Metadata”  
Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

### **CURTIME ( )**

Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 11.3.6, “Fractional Seconds in Time Values”  
Section 8.9.3.1, “How the Query Cache Operates”

Section 10.6, “MySQL Server Time Zone Support”  
Section 16.4.1.12, “Replication and Fractional Seconds Support”

## **D**

[index top [3931]]

### **DATABASE ( )**

Section 13.1.17, “`DROP DATABASE` Syntax”  
Section 3.3.1, “Creating and Selecting a Database”  
Section 3.4, “Getting Information About Databases and Tables”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section C.5.8, “Known Issues in MySQL”  
Section 10.1.12, “UTF-8 for Metadata”

### **DATE ( )**

Section 12.7, “Date and Time Functions”

### **DATE\_ADD ( )**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 12.6.1, “Arithmetic Operators”  
Section 12.7, “Date and Time Functions”  
Section 11.3, “Date and Time Types”  
Section 3.3.4.5, “Date Calculations”  
Section 9.5, “Expression Syntax”

### **DATE\_FORMAT ( )**

Section 21.8.17, “C API Prepared Statement Problems”  
Section 12.7, “Date and Time Functions”  
Section 10.7, “MySQL Server Locale Support”  
Section 5.1.4, “Server System Variables”

### **DATE\_SUB ( )**

Section 12.7, “Date and Time Functions”  
Section 11.3, “Date and Time Types”

### **DATEDIFF ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **DAY ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **DAYNAME ( )**

Section 12.7, “Date and Time Functions”  
Section 10.7, “MySQL Server Locale Support”  
Section 5.1.4, “Server System Variables”

### **DAYOFMONTH ( )**

Section 12.7, “Date and Time Functions”  
Section 3.3.4.5, “Date Calculations”



---

Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **DAYOFWEEK ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **DAYOFYEAR ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 17.2, “Partitioning Types”

### **DECODE ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

### **DEFAULT ( )**

Section 13.2.5, “INSERT Syntax”  
Section 13.2.8, “REPLACE Syntax”  
Section 11.5, “Data Type Default Values”  
Section 12.16, “Miscellaneous Functions”

### **DEGREES ( )**

Section 12.6.2, “Mathematical Functions”

### **DES\_DECRYPT ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 5.1.3, “Server Command Options”

### **DES\_ENCRYPT ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 5.1.3, “Server Command Options”

### **Difference ( )**

Section 12.18.5.3.2, “Spatial Operators”

### **Dimension ( )**

Section 12.18.5.2.1, “General Geometry Functions”

### **Disjoint ( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **E**

[index top [3931]]

### **ELT ( )**

Section C.5.8, “Known Issues in MySQL”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 10.1.9.1, “Result Strings”

Section 12.5, “String Functions”

### **ENCODE ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

### **ENCRYPT ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section E.7, “Restrictions on Character Sets”  
Section 5.1.4, “Server System Variables”  
Section 6.3.1, “User Names and Passwords”

### **EndPoint ( )**

Section 12.18.5.2.3, “LineString Functions”  
Section 12.18.5.2.4, “MultiLineString Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

### **Envelope ( )**

Section 12.18.5.2.1, “General Geometry Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

### **Equals ( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

### **EXP ( )**

Section 12.6.2, “Mathematical Functions”

### **EXPORT\_SET ( )**

Section 12.5, “String Functions”

### **expr IN ( )**

Section 12.3.2, “Comparison Functions and Operators”

### **expr NOT IN ( )**

Section 12.3.2, “Comparison Functions and Operators”

### **ExteriorRing ( )**

Section 12.18.5.2.5, “Polygon Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

### **EXTRACT ( )**

Section 12.10, “Cast Functions and Operators”  
Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **ExtractValue ( )**

Section 12.11, “XML Functions”

---

## F

[index top [3931]]

### **FIELD()**

Section 12.5, “String Functions”

### **FIND\_IN\_SET()**

Section 12.5, “String Functions”  
Section 11.4.5, “The `SET` Type”

### **FLOOR()**

Section 12.6.1, “Arithmetic Operators”  
Section 12.6.2, “Mathematical Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **FORMAT()**

Section 12.6.2, “Mathematical Functions”  
Section 12.16, “Miscellaneous Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 10.7, “MySQL Server Locale Support”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

### **FOUND\_ROWS()**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.14, “Replication and System Functions”

### **FROM\_BASE64()**

Section 12.5, “String Functions”

### **FROM\_DAYS()**

Section 12.7, “Date and Time Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

### **FROM\_UNIXTIME()**

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 12.7, “Date and Time Functions”  
Section 16.4.1.28, “Replication and Time Zones”

## G

[index top [3931]]

### **GeomCollFromText()**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

### **GeomCollFromWKB()**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

### **GeometryCollection()**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

### **GeometryCollectionFromText()**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

### **GeometryCollectionFromWKB()**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

### **GeometryFromText()**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

### **GeometryFromWKB()**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

### **GeometryN()**

Section 12.18.5.2.7, “GeometryCollection Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

### **GeometryType()**

Section 12.18.5.2.1, “General Geometry Functions”

### **GeomFromText()**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”  
Section 12.18.5.1, “Geometry Format Conversion Functions”  
Section 12.18.4.4, “Populating Spatial Columns”

### **GeomFromWKB()**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”  
Section 12.18.5.1, “Geometry Format Conversion Functions”

### **GET\_FORMAT()**

Section 12.7, “Date and Time Functions”  
Section 10.7, “MySQL Server Locale Support”

### **GET\_LOCK()**

Section 13.1.9, “CREATE EVENT Syntax”  
Section 13.7.6.4, “KILL Syntax”  
Section 21.8.6.3, “mysql\_change\_user()”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”



---

Section 21.8.15, “Controlling Automatic Reconnection Behavior”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 18.4.1, “Event Scheduler Overview”  
Section 8.12.5.2, “General Thread States”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.16, “Miscellaneous Functions”  
Section 16.4.1.14, “Replication and System Functions”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

### **gethostbyaddr ( )**

Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”

### **gethostbyaddr\_r ( )**

Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”

### **gethostbyname ( )**

Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”

### **gethostbyname\_r ( )**

Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”

### **GLength ( )**

Section 12.18.5.2.3, “LineString Functions”  
Section 12.18.5.2.4, “MultiLineString Functions”  
Section 12.18.7, “MySQL Conformance and Compatibility”

### **GREATEST ( )**

Section 12.3.2, “Comparison Functions and Operators”  
Section 10.1.9.1, “Result Strings”

### **GROUP\_CONCAT ( )**

Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section C.5.8, “Known Issues in MySQL”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 5.1.4, “Server System Variables”

### **GTID\_SUBSET ( )**

Section 16.1.3.1, “GTID Concepts”  
Section 12.15, “GTID Functions”

### **GTID\_SUBTRACT ( )**

Section 16.1.3.1, “GTID Concepts”  
Section 12.15, “GTID Functions”

## **H**

[index top [3931]]

### **HEX ( )**

Section 10.1.3.5, “Character String Literal Character Set and Collation”  
Section 9.1.4, “Hexadecimal Literals”  
Section 12.6.2, “Mathematical Functions”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

### **HOUR ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **I**

[index top [3931]]

### **IF ( )**

Section 13.6.5.2, “IF Syntax”  
Section 12.4, “Control Flow Functions”  
Section C.5.8, “Known Issues in MySQL”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 10.1.9.1, “Result Strings”

### **IFNULL ( )**

Section 12.4, “Control Flow Functions”  
Section C.5.5.3, “Problems with NULL Values”

## **IN**

Section 12.3.1, “Operator Precedence”

### **IN ( )**

Section 8.8.2, “EXPLAIN Output Format”  
Section E.4, “Restrictions on Subqueries”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 12.2, “Type Conversion in Expression Evaluation”

### **INET6\_ATON ( )**

Section 5.1.9, “IPv6 Support”  
Section 12.16, “Miscellaneous Functions”

### **INET6\_NTOA ( )**

Section 5.1.9, “IPv6 Support”  
Section 12.16, “Miscellaneous Functions”

### **INET\_ATON ( )**

Section 5.1.9, “IPv6 Support”  
Section 12.16, “Miscellaneous Functions”

### **INET\_NTOA ( )**

Section 5.1.9, “IPv6 Support”

---

Section 12.16, “Miscellaneous Functions”

## **INSERT()**

Section 12.5, “String Functions”

## **INSTR()**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **InteriorRingN()**

Section 12.18.5.2.5, “Polygon Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

## **Intersection()**

Section 12.18.5.3.2, “Spatial Operators”

## **Intersects()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **INTERVAL()**

Section 12.3.2, “Comparison Functions and Operators”

## **IS\_FREE\_LOCK()**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.16, “Miscellaneous Functions”  
Section 16.4.1.14, “Replication and System Functions”

## **IS\_IPV4()**

Section 12.16, “Miscellaneous Functions”

## **IS\_IPV4\_COMPAT()**

Section 12.16, “Miscellaneous Functions”

## **IS\_IPV4\_MAPPED()**

Section 12.16, “Miscellaneous Functions”

## **IS\_IPV6()**

Section 12.16, “Miscellaneous Functions”

## **IS\_USED\_LOCK()**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.16, “Miscellaneous Functions”  
Section 16.4.1.14, “Replication and System Functions”

## **IsClosed()**

Section 12.18.5.2.4, “MultiLineString Functions”

## **IsEmpty()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”  
Section 12.18.5.2.1, “General Geometry Functions”

## **ISNULL()**

Section 12.3.2, “Comparison Functions and Operators”

## **IsRing()**

Section 12.18.5.2.3, “LineString Functions”

## **IsSimple()**

Section 12.18.5.2.1, “General Geometry Functions”

## **L**

[index top [3931]]

## **LAST\_DAY()**

Section 12.7, “Date and Time Functions”

## **LAST\_INSERT\_ID()**

Section 4.6.8.1, “mysqlbinlog Hex Dump Format”  
Section 13.1.14, “CREATE TABLE Syntax”  
Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”  
Section 13.2.5.2, “INSERT DELAYED Syntax”  
Section 13.2.5, “INSERT Syntax”  
Section 21.8.6.37, “mysql\_insert\_id()”  
Section 21.8.10.16, “mysql\_stmt\_insert\_id()”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 21.8.15, “Controlling Automatic Reconnection Behavior”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 21.8.14.3, “How to Get the Unique ID for the Last Inserted Row”  
Section 12.14, “Information Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 16.4.1.1, “Replication and AUTO\_INCREMENT”  
Section 5.1.4, “Server System Variables”  
Section 18.2.4, “Stored Procedures, Functions, Triggers, and LAST\_INSERT\_ID()”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”  
Section 16.4.4, “Troubleshooting Replication”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 3.6.9, “Using AUTO\_INCREMENT”

## **LCASE()**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

---

## **LEAST ( )**

Section 12.3.2, “Comparison Functions and Operators”  
Section 10.1.9.1, “Result Strings”

## **LEFT ( )**

Section 12.10, “Cast Functions and Operators”  
Section 12.5, “String Functions”

## **LENGTH ( )**

Section 12.5, “String Functions”

## **Length ( )**

Section 12.18.5.2.3, “`LineString` Functions”  
Section 12.18.5.2.4, “`MultiLineString` Functions”  
Section 12.18.7, “MySQL Conformance and Compatibility”

## **LineFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”  
Section 12.18.5.1, “Geometry Format Conversion Functions”

## **LineFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”  
Section 12.18.5.1, “Geometry Format Conversion Functions”

## **LineString ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **LineStringFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **LineStringFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **LN ( )**

Section 12.6.2, “Mathematical Functions”

## **LOAD\_FILE ( )**

Section 13.2.7, “`LOAD XML` Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.14, “Replication and System Functions”

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 12.5, “String Functions”

## **LOCALTIME**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.7, “Date and Time Functions”

## **LOCALTIME ( )**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

## **LOCALTIMESTAMP**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.7, “Date and Time Functions”

## **LOCALTIMESTAMP ( )**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

## **LOCATE ( )**

Section 12.5, “String Functions”

## **LOG ( )**

Section 17.2.4.1, “`LINEAR HASH` Partitioning”  
Section 12.6.2, “Mathematical Functions”

## **LOG10 ( )**

Section 12.6.2, “Mathematical Functions”

## **LOG2 ( )**

Section 12.6.2, “Mathematical Functions”

## **LOWER ( )**

Section 12.10, “Cast Functions and Operators”  
Section 10.1.7.9, “Collation and `INFORMATION_SCHEMA` Searches”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”  
Section 10.1.14.1, “Unicode Character Sets”

## **LPAD ( )**

Section 12.5, “String Functions”

## **LTRIM ( )**

Section 10.1.9.1, “Result Strings”

---

Section 12.5, “String Functions”

## M

[index top [3931]]

### **MAKE\_SET ( )**

Section 12.5, “String Functions”

### **MAKEDATE ( )**

Section 12.7, “Date and Time Functions”

### **MAKETIME ( )**

Section 12.7, “Date and Time Functions”

### **MASTER\_POS\_WAIT ( )**

Section 13.4.2.1, “CHANGE MASTER TO Syntax”  
Section 13.4.2.2, “MASTER\_POS\_WAIT ( ) Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.16, “Miscellaneous Functions”

### **MATCH**

Section 9.5, “Expression Syntax”

### **MATCH ( )**

Section 12.9, “Full-Text Search Functions”

### **MATCH ( )**

Section 14.2.3.12.3, “FULLTEXT Indexes”  
Section 12.9.2, “Boolean Full-Text Searches”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 12.9.5, “Full-Text Restrictions”  
Section 12.9, “Full-Text Search Functions”  
Section 12.9.1, “Natural Language Full-Text Searches”  
Section 1.4, “What Is New in MySQL 5.6”

### **MAX ( )**

Section 8.8.2, “EXPLAIN Output Format”  
Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section 8.3.1, “How MySQL Uses Indexes”  
Section C.5.8, “Known Issues in MySQL”  
Section 8.13.14.1, “Loose Index Scan”  
Section 12.17.3, “MySQL Extensions to GROUP BY”  
Section 11.1.1, “Numeric Type Overview”  
Section 13.2.10.10, “Optimizing Subqueries”  
Section 5.1.7, “Server SQL Modes”  
Section 19.31.2, “The INFORMATION\_SCHEMA TP\_THREAD\_GROUP\_STATE Table”  
Section 11.3.8, “Two-Digit Years in Dates”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 8.13.5, “Use of Index Extensions”

Section 3.6.9, “Using AUTO\_INCREMENT”  
Section 18.5.2, “View Processing Algorithms”

### **MBRContains ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”  
Section 12.18.6.2, “Using a Spatial Index”

### **MBRDisjoint ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”

### **MBREqual ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”

### **MBRIntersects ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”

### **MBROverlaps ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”

### **MBRTouches ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”

### **MBRWithin ( )**

Section 12.18.5.4.1, “Relations on Geometry Minimal Bounding Rectangles (MBRs)”  
Section 12.18.6.2, “Using a Spatial Index”

### **MD5 ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 9.2, “Schema Object Names”  
Section 6.1.1, “Security Guidelines”

### **MICROSECOND ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

### **MID ( )**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

### **MIN ( )**

Section 8.8.2, “EXPLAIN Output Format”  
Section 12.17.1, “GROUP BY (Aggregate) Functions”

---

Section 21.8.17, “C API Prepared Statement Problems”  
Section 8.2.1.2, “How MySQL Optimizes `WHERE` Clauses”  
Section 8.3.1, “How MySQL Uses Indexes”  
Section C.5.8, “Known Issues in MySQL”  
Section 8.13.14.1, “Loose Index Scan”  
Section 12.17.3, “MySQL Extensions to `GROUP BY`”  
Section 11.1.1, “Numeric Type Overview”  
Section 13.2.10.10, “Optimizing Subqueries”  
Section C.5.5.3, “Problems with `NULL` Values”  
Section 11.3.8, “Two-Digit Years in Dates”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 8.13.5, “Use of Index Extensions”  
Section 18.5.2, “View Processing Algorithms”

## **MINUTE ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **MLineFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MLineFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **MOD ( )**

Section 12.6.1, “Arithmetic Operators”  
Section 3.3.4.5, “Date Calculations”  
Section 12.6.2, “Mathematical Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 5.1.7, “Server SQL Modes”

## **MONTH ( )**

Section 12.7, “Date and Time Functions”  
Section 3.3.4.5, “Date Calculations”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 17.2, “Partitioning Types”

## **MONTHNAME ( )**

Section 12.7, “Date and Time Functions”  
Section 10.7, “MySQL Server Locale Support”  
Section 5.1.4, “Server System Variables”

## **MPointFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MPointFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **MPolyFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MPolyFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **MultiLineString ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **MultiLineStringFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MultiLineStringFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **MultiPoint ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **MultiPointFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MultiPointFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **MultiPolygon ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **MultiPolygonFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **MultiPolygonFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **my\_open ( )**

Section 5.1.6, “Server Status Variables”

## **N**

[[index top \[3931\]](#)]

## **NAME\_CONST ( )**

Section 18.7, “Binary Logging of Stored Programs”  
Section 12.16, “Miscellaneous Functions”

---

## **NOW( )**

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 11.5, “Data Type Default Values”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 11.3.6, “Fractional Seconds in Time Values”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 10.6, “MySQL Server Time Zone Support”  
Section 16.4.1.14, “Replication and System Functions”  
Section 16.4.1.28, “Replication and Time Zones”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 11.3.3, “The `YEAR` Type”  
Section 10.6.2, “Time Zone Leap Second Support”

## **NULLIF( )**

Section 12.4, “Control Flow Functions”

## **NumGeometries( )**

Section 12.18.5.2.7, “`GeometryCollection` Functions”

## **NumInteriorRings( )**

Section 12.18.5.2.5, “`Polygon` Functions”

## **NumPoints( )**

Section 12.18.5.2.3, “`LineString` Functions”

## **O**

[index top [3931]]

## **OCT( )**

Section 12.5, “String Functions”

## **OCTET\_LENGTH( )**

Section 12.5, “String Functions”

## **OLD\_PASSWORD( )**

Section C.5.2.4, “Client does not support authentication protocol”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 12.13, “Encryption and Compression Functions”  
Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”  
Section 6.1.2.4, “Password Hashing in MySQL”  
Section 22.2.3.8, “Password-Validation Plugins”  
Section 5.1.4, “Server System Variables”

Section 6.1.2.6, “The Password Validation Plugin”  
Section 1.4, “What Is New in MySQL 5.6”

## **ORD( )**

Section 12.5, “String Functions”

## **Overlaps( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **P**

[index top [3931]]

## **PASSWORD( )**

Section 13.7.1.2, “`CREATE USER` Syntax”  
Section C.5.2.15, “Ignoring user”  
Section 17.2.5, “`KEY` Partitioning”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 6.2.4, “Access Control, Stage 1: Connection Verification”  
Section 6.3.2, “Adding User Accounts”  
Section 6.3.5, “Assigning Account Passwords”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.13, “Encryption and Compression Functions”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 6.1.2.4, “Password Hashing in MySQL”  
Section 22.2.3.8, “Password-Validation Plugins”  
Section 6.1.2.3, “Passwords and Logging”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 6.1.2.6, “The Password Validation Plugin”  
Section 6.3.7.2, “The SHA-256 Authentication Plugin”  
Section 6.3.1, “User Names and Passwords”  
Section 1.4, “What Is New in MySQL 5.6”

## **PERIOD\_ADD( )**

Section 12.7, “Date and Time Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

## **PERIOD\_DIFF( )**

Section 12.7, “Date and Time Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

## **PI( )**

Section 9.2.4, “Function Name Parsing and Resolution”  
Section 12.6.2, “Mathematical Functions”

## **POINT( )**

Section 12.18.3.1, “Well-Known Text (WKT) Format”



---

## **Point ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **PointFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

Section 12.18.5.1, “Geometry Format Conversion Functions”

## **PointFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

Section 12.18.5.1, “Geometry Format Conversion Functions”

## **PointN ( )**

Section 12.18.5.2.3, “`LineString` Functions”

Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

## **PolyFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **PolyFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **Polygon ( )**

Section 12.18.4.2.3, “Creating Geometry Values Using MySQL-Specific Functions”

## **PolygonFromText ( )**

Section 12.18.4.2.1, “Creating Geometry Values Using WKT Functions”

## **PolygonFromWKB ( )**

Section 12.18.4.2.2, “Creating Geometry Values Using WKB Functions”

## **POSITION ( )**

Section 12.5, “String Functions”

## **POW ( )**

Section 17.2.4, “`HASH` Partitioning”

Section 12.6.2, “Mathematical Functions”

## **POWER ( )**

Section 17.2.4.1, “`LINEAR HASH` Partitioning”

Section 12.6.2, “Mathematical Functions”

## **Q**

[index top [3931]]

## **QUARTER ( )**

Section 12.7, “Date and Time Functions”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **QUOTE ( )**

Section 21.8.6.54, “`mysql_real_escape_string ( )`”

Section 12.5, “String Functions”

Section 9.1.1, “String Literals”

## **R**

[index top [3931]]

## **RADIANS ( )**

Section 12.6.2, “Mathematical Functions”

## **RAND ( )**

Section 4.6.8.1, “`mysqlbinlog` Hex Dump Format”

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 8.9.3.1, “How the Query Cache Operates”

Section 12.6.2, “Mathematical Functions”

Section 16.4.1.14, “Replication and System Functions”

Section 5.1.4, “Server System Variables”

## **RELEASE\_LOCK ( )**

Section 13.2.3, “`DO` Syntax”

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 8.9.3.1, “How the Query Cache Operates”

Section 12.16, “Miscellaneous Functions”

Section 16.4.1.14, “Replication and System Functions”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

## **REPEAT ( )**

Section 10.1.9.1, “Result Strings”

Section 12.5, “String Functions”

## **REPLACE ( )**

Section 10.1.9.1, “Result Strings”

Section 12.5, “String Functions”

## **REVERSE ( )**

Section 10.1.9.1, “Result Strings”

Section 12.5, “String Functions”

## **RIGHT ( )**

Section 10.1.9.1, “Result Strings”

---

Section 12.5, “String Functions”

## **ROUND ( )**

Section 12.6.2, “Mathematical Functions”  
Section 12.19, “Precision Math”  
Section 12.19.5, “Precision Math Examples”  
Section 12.19.4, “Rounding Behavior”

## **ROW\_COUNT ( )**

Section 13.2.1, “CALL Syntax”  
Section 13.2.2, “DELETE Syntax”  
Section 13.2.5, “INSERT Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 13.6.7.7.2, “Diagnostics Area Information Items”  
Section 12.14, “Information Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.14, “Replication and System Functions”

## **RPAD ( )**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **RTRIM ( )**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **S**

[index top [3931]]

## **SCHEMA ( )**

Section 12.14, “Information Functions”

## **SEC\_TO\_TIME ( )**

Section 12.7, “Date and Time Functions”

## **SECOND ( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **SESSION\_USER ( )**

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.14, “Information Functions”  
Section 10.1.12, “UTF-8 for Metadata”

## **setrlimit ( )**

Section 5.1.3, “Server Command Options”

## **SHA ( )**

Section 12.13, “Encryption and Compression Functions”

## **SHA1 ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”  
Section 6.1.1, “Security Guidelines”

## **SHA2 ( )**

Section 12.13, “Encryption and Compression Functions”  
Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”  
Section 6.1.1, “Security Guidelines”

## **SIGN ( )**

Section 12.6.2, “Mathematical Functions”

## **SIN ( )**

Section 12.6.2, “Mathematical Functions”  
Section 22.3.2.3, “UDF Argument Processing”

## **SLEEP ( )**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.12.5.2, “General Thread States”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.16, “Miscellaneous Functions”  
Section 16.4.1, “Replication Features and Issues”  
Section 19.31.3, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATS` Table”

## **SOUNDEX ( )**

Section 22.3, “Adding New Functions to MySQL”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **SPACE ( )**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **SQL\_THREAD\_WAIT\_AFTER\_GTIDS ( )**

Section 12.15, “GTID Functions”

## **SQRT ( )**

Section 12.6.2, “Mathematical Functions”

## **SRID ( )**

Section 12.18.5.2.1, “General Geometry Functions”

## **ST\_Contains ( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”  
Section 1.4, “What Is New in MySQL 5.6”



---

## **ST\_Crosses()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Disjoint()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Equals()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Intersects()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Overlaps()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Touches()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **ST\_Within()**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **StartPoint()**

Section 12.18.5.2.3, “LineString Functions”  
Section 12.18.5.2.4, “MultiLineString Functions”  
Section 12.18.5.3.1, “Geometry Functions That Produce New Geometries”

## **STD()**

Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

## **STDDEV()**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **STDDEV\_POP()**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **STDDEV\_SAMP()**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **STR\_TO\_DATE()**

Section 12.7, “Date and Time Functions”  
Section 10.7, “MySQL Server Locale Support”

## **STRCMP()**

Section C.5.5.2, “Problems Using DATE Columns”  
Section 12.5.1, “String Comparison Functions”

## **SUBDATE()**

Section 12.7, “Date and Time Functions”

## **SUBSTR()**

Section 12.5, “String Functions”

## **SUBSTRING()**

Section 10.1.9.1, “Result Strings”  
Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”  
Section 12.5, “String Functions”

## **SUBSTRING\_INDEX()**

Section 12.5, “String Functions”

## **SUBTIME()**

Section 12.7, “Date and Time Functions”

## **SUM()**

Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section 22.3.2, “Adding a New User-Defined Function”  
Section 11.1.2, “Date and Time Type Overview”  
Section 8.13.14.1, “Loose Index Scan”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section C.5.5.3, “Problems with NULL Values”  
Section 11.4.4, “The ENUM Type”  
Section 11.4.5, “The SET Type”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 18.5.2, “View Processing Algorithms”

## **SymDifference()**

Section 12.18.5.3.2, “Spatial Operators”

## **SYSDATE()**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 11.3.6, “Fractional Seconds in Time Values”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 16.4.1.12, “Replication and Fractional Seconds Support”  
Section 16.4.1.14, “Replication and System Functions”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **SYSTEM\_USER()**

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.14, “Information Functions”  
Section 10.1.12, “UTF-8 for Metadata”

## **T**

[index top [3931]]

---

## **TAN( )**

Section 12.6.2, “Mathematical Functions”

## **thr\_setconcurrency( )**

Section 5.1.4, “Server System Variables”

## **TIME( )**

Section 12.7, “Date and Time Functions”

## **TIME\_FORMAT( )**

Section 12.7, “Date and Time Functions”

## **TIME\_TO\_SEC( )**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **TIMEDIFF( )**

Section 12.7, “Date and Time Functions”

## **TIMESTAMP( )**

Section 12.7, “Date and Time Functions”

## **TIMESTAMPADD( )**

Section 12.7, “Date and Time Functions”

## **TIMESTAMPDIFF( )**

Section 12.7, “Date and Time Functions”  
Section 3.3.4.5, “Date Calculations”

## **TO\_BASE64( )**

Section 12.5, “String Functions”

## **TO\_DAYS( )**

Section 17.2.4, “HASH Partitioning”  
Section 12.7, “Date and Time Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.4, “Partition Pruning”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 17.2, “Partitioning Types”

## **TO\_SECONDS( )**

Section 12.7, “Date and Time Functions”  
Section 17.4, “Partition Pruning”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 17.2, “Partitioning Types”

## **Touches( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **TRIM( )**

Section 10.1.13, “Column Character Set Conversion”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **TRUNCATE( )**

Section 12.6.2, “Mathematical Functions”

## **U**

[[index top \[3931\]](#)]

## **UCASE( )**

Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”

## **UNCOMPRESS( )**

Section 12.13, “Encryption and Compression Functions”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.1.4, “Server System Variables”

## **UNCOMPRESSED\_LENGTH( )**

Section 12.13, “Encryption and Compression Functions”

## **UNHEX( )**

Section 12.13, “Encryption and Compression Functions”  
Section 12.5, “String Functions”

## **Union( )**

Section 12.18.5.3.2, “Spatial Operators”

## **UNIX\_TIMESTAMP( )**

Section 17.2.1, “RANGE Partitioning”  
Section 12.7, “Date and Time Functions”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section C.5.4.6, “Time Zone Problems”

## **UpdateXML( )**

Section 12.11, “XML Functions”

## **UPPER( )**

Section 12.10, “Cast Functions and Operators”  
Section 10.1.7.9, “Collation and INFORMATION\_SCHEMA Searches”  
Section 10.1.9.1, “Result Strings”  
Section 12.5, “String Functions”  
Section 10.1.8, “String Repertoire”

---

Section 10.1.14.1, “Unicode Character Sets”

## **USER ( )**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 10.1.7.5, “Collation of Expressions”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 8.9.3.1, “How the Query Cache Operates”

Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”

Section 12.14, “Information Functions”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 6.3.8, “Proxy Users”

Section 16.4.1.14, “Replication and System Functions”

Section 6.3.12, “SQL-Based MySQL Account Activity Auditing”

Section 10.1.12, “UTF-8 for Metadata”

Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

## **UTC\_DATE**

Section 12.7, “Date and Time Functions”

## **UTC\_DATE ( )**

Section 12.7, “Date and Time Functions”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

## **UTC\_TIME**

Section 12.7, “Date and Time Functions”

## **UTC\_TIME ( )**

Section 12.7, “Date and Time Functions”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

## **UTC\_TIMESTAMP**

Section 12.7, “Date and Time Functions”

## **UTC\_TIMESTAMP ( )**

Section 12.7, “Date and Time Functions”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 11.3.6, “Fractional Seconds in Time Values”

Section 10.6, “MySQL Server Time Zone Support”

Section 16.4.1.12, “Replication and Fractional Seconds Support”

## **UUID ( )**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 16.1.4.4, “Binary Log Options and Variables”

Section 18.7, “Binary Logging of Stored Programs”

Section 21.2.6, “Connector/Net Connection String Options Reference”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 8.9.3.1, “How the Query Cache Operates”

Section 12.16, “Miscellaneous Functions”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.4.1.14, “Replication and System Functions”

Section 5.2.4.2, “Setting The Binary Log Format”

## **UUID\_SHORT ( )**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”

Section 8.9.3.1, “How the Query Cache Operates”

Section 12.16, “Miscellaneous Functions”

## **V**

[index top [3931]]

## **VALIDATE\_PASSWORD\_STRENGTH ( )**

Section 12.13, “Encryption and Compression Functions”

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

Section 22.2.3.8, “Password-Validation Plugins”

Section 6.1.2.6, “The Password Validation Plugin”

Section 1.4, “What Is New in MySQL 5.6”

## **VALUES ( )**

Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”

Section 12.16, “Miscellaneous Functions”

## **VAR\_POP ( )**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **VAR\_SAMP ( )**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **VARIANCE ( )**

Section 12.17.1, “GROUP BY (Aggregate) Functions”

## **VERSION ( )**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section C.5.5.1, “Case Sensitivity in String Searches”

Section 10.1.7.5, “Collation of Expressions”

Section 12.14, “Information Functions”

Section 16.4.1.14, “Replication and System Functions”

Section 6.3.11.2, “The Audit Log File”

Section 10.1.12, “UTF-8 for Metadata”

## **W**

[index top [3931]]

---

## **WAIT\_UNTIL\_SQL\_THREAD\_AFTER\_GTIDS( )**

Section 12.15, “GTID Functions”

## **YEARWEEK( )**

Section 12.7, “Date and Time Functions”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

## **WEEK( )**

Section 12.7, “Date and Time Functions”

Section 5.1.4, “Server System Variables”

## **WEEKDAY( )**

Section 12.7, “Date and Time Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

Section 17.2, “Partitioning Types”

## **WEEKOFYEAR( )**

Section 12.7, “Date and Time Functions”

## **WEIGHT\_STRING( )**

Section 10.4, “Adding a Collation to a Character Set”

Section C.5.5.1, “Case Sensitivity in String Searches”

Section 12.5, “String Functions”

Section 10.1.14.1, “Unicode Character Sets”

## **Within( )**

Section 12.18.5.4.2, “Functions That Test Spatial Relationships Between Geometries”

## **X**

[[index top \[3931\]](#)]

## **X( )**

Section 12.18.5.2.2, “Point Functions”

## **Y**

[[index top \[3931\]](#)]

## **Y( )**

Section 12.18.5.2.2, “Point Functions”

## **YEAR( )**

Section 17.2.4, “HASH Partitioning”

Section 17.2.1, “RANGE Partitioning”

Section 12.7, “Date and Time Functions”

Section 3.3.4.5, “Date Calculations”

Section 17.2.7, “How MySQL Partitioning Handles NULL”

Section 17.3.1, “Management of RANGE and LIST Partitions”

Section 17.4, “Partition Pruning”

Section 17.6.3, “Partitioning Limitations Relating to Functions”

Section 17.2, “Partitioning Types”

---

## INFORMATION\_SCHEMA Index

[C](#) | [E](#) | [F](#) | [G](#) | [I](#) | [K](#) | [O](#) | [P](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#)

### C

[\[index top \[3949\]\]](#)

#### CHARACTER\_SETS

Section 10.1.9.3, “SHOW Statements and INFORMATION\_SCHEMA”

Section 19.1, “The INFORMATION\_SCHEMA CHARACTER\_SETS Table”

#### COLLATION\_CHARACTER\_SET\_APPLICABILITY

Section 19.3, “The INFORMATION\_SCHEMA COLLATION\_CHARACTER\_SET\_APPLICABILITY Table”

#### COLLATIONS

Section 10.1.9.3, “SHOW Statements and INFORMATION\_SCHEMA”

Section 21.8.4, “C API Data Structures”

Section 8.2.4, “Optimizing INFORMATION\_SCHEMA Queries”

Section 19.2, “The INFORMATION\_SCHEMA COLLATIONS Table”

#### COLUMN\_PRIVILEGES

Section 19.5, “The INFORMATION\_SCHEMA COLUMN\_PRIVILEGES Table”

#### COLUMNS

Section 8.2.4, “Optimizing INFORMATION\_SCHEMA Queries”

Section 19.4, “The INFORMATION\_SCHEMA COLUMNS Table”

Section 19.13, “The INFORMATION\_SCHEMA PARAMETERS Table”

Section 19.19, “The INFORMATION\_SCHEMA ROUTINES Table”

### E

[\[index top \[3949\]\]](#)

#### ENGINES

Section 19.6, “The INFORMATION\_SCHEMA ENGINES Table”

#### EVENTS

Section 18.4.4, “Event Metadata”

Section 18.4.2, “Event Scheduler Configuration”

Section 16.4.1.10, “Replication of Invoked Features”

Section 19.7, “The INFORMATION\_SCHEMA EVENTS Table”

### F

[\[index top \[3949\]\]](#)

#### FILES

Section 19.8, “The INFORMATION\_SCHEMA FILES Table”

### G

[\[index top \[3949\]\]](#)

#### GLOBAL\_STATUS

Section 19.9, “The INFORMATION\_SCHEMA GLOBAL\_STATUS and SESSION\_STATUS Tables”

#### GLOBAL\_VARIABLES

Section 19.10, “The INFORMATION\_SCHEMA GLOBAL\_VARIABLES and SESSION\_VARIABLES Tables”

### I

[\[index top \[3949\]\]](#)

#### INFORMATION\_SCHEMA

##### GLOBAL\_STATUS

Section 19.9, “The INFORMATION\_SCHEMA GLOBAL\_STATUS and SESSION\_STATUS Tables”

##### INFORMATION\_SCHEMA

##### GLOBAL\_VARIABLES

Section 19.10, “The INFORMATION\_SCHEMA GLOBAL\_VARIABLES and SESSION\_VARIABLES Tables”

##### INFORMATION\_SCHEMA.COLLATIONS

Section 10.4.2, “Choosing a Collation ID”

##### INFORMATION\_SCHEMA.COLUMNS

Section 20.1, “Performance Schema Quick Start”

##### INFORMATION\_SCHEMA.ENGINES

Section 20.1, “Performance Schema Quick Start”

##### INFORMATION\_SCHEMA.EVENTS

Section 13.7.5.19, “SHOW EVENTS Syntax”

Section 18.4.4, “Event Metadata”

Section 16.4.1.10, “Replication of Invoked Features”

Section E.1, “Restrictions on Stored Programs”

Section 18.4.6, “The Event Scheduler and MySQL Privileges”

##### INFORMATION\_SCHEMA.INNOODB\_CMP

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

---

## **INFORMATION\_SCHEMA.INNODB\_CMP**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

## **information\_schema.innodb\_ft\_default\_stopword**

Section 12.9.4, “Full-Text Stopwords”

## **INFORMATION\_SCHEMA.INNODB\_METRICS**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **INFORMATION\_SCHEMA.KEY\_COLUMN\_USAGE**

Section 1.8.6.2, “FOREIGN KEY Constraints”

Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **INFORMATION\_SCHEMA.OPTIMIZER\_TRACE**

Section 1.4, “What Is New in MySQL 5.6”

## **INFORMATION\_SCHEMA.PARTITIONS**

Section 17.2.5, “KEY Partitioning”

Section 17.2.3.1, “RANGE COLUMNS partitioning”

Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”

Section 17.2.7, “How MySQL Partitioning Handles NULL”

Section 17.3.5, “Obtaining Information About Partitions”

Section 5.1.3, “Server Command Options”

Section 19.14, “The INFORMATION\_SCHEMA PARTITIONS Table”

## **INFORMATION\_SCHEMA.PLUGINS**

Section 13.7.3.3, “INSTALL PLUGIN Syntax”

Section 5.1.8.1, “Installing and Uninstalling Plugins”

Section 6.3.11.1, “Installing the Audit Log Plugin”

Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”

Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”

Section 5.1.8.2, “Obtaining Server Plugin Information”

Section 6.1.2.6.1, “Password Validation Plugin Installation”

Section 22.2.1, “Plugin API Characteristics”

Section 22.2.2, “Plugin API Components”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”

Section 8.11.6.1, “Thread Pool Components and Installation”

Section 22.2.4.6, “Writing INFORMATION\_SCHEMA Plugins”

Section 22.2.4.8, “Writing Audit Plugins”

Section 22.2.4.5, “Writing Daemon Plugins”

Section 22.2.4.4, “Writing Full-Text Parser Plugins”

Section 22.2.4.10, “Writing Password-Validation Plugins”

Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

## **INFORMATION\_SCHEMA.PROCESSLIST**

Section 13.7.5.30, “SHOW PROCESSLIST Syntax”

Section 8.12.5, “Examining Thread Information”

Section 20.4, “Performance Schema Instrument Naming Conventions”

Section 20.9.4, “Performance Schema Stage Event Tables”

Section 20.16, “The INFORMATION\_SCHEMA PROCESSLIST Table”

Section 20.9.9.3, “The threads Table”

## **INFORMATION\_SCHEMA.ROUTINES**

Chapter 19, *INFORMATION\_SCHEMA Tables*

Section 21.2.6, “Connector/Net Connection String Options Reference”

Section 19.19, “The INFORMATION\_SCHEMA ROUTINES Table”

## **INFORMATION\_SCHEMA.TABLE\_CONSTRAINTS**

Section 19.18, “The INFORMATION\_SCHEMA REFERENTIAL\_CONSTRAINTS Table”

## **information\_schema.table\_constraint**

Section 5.5.1, “Overview of Online DDL”

## **INFORMATION\_SCHEMA.TABLES**

Chapter 19, *INFORMATION\_SCHEMA Tables*

Section 5.1.3, “Server Command Options”

## **INFORMATION\_SCHEMA.TRIGGERS**

Section 13.7.5.39, “SHOW TRIGGERS Syntax”

## **INFORMATION\_SCHEMA.VIEWS**

Section 18.5.3, “Updatable and Insertable Views”

## **INNODB\_CMP**

Section 5.4.6.4, “Monitoring Compression at Runtime”

## **INNODB\_CMP\_PER\_INDEX**

Section 5.4.6.4, “Monitoring Compression at Runtime”

## **INNODB\_FT\_BEING\_DELETED**

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **INNODB\_FT\_CONFIG**

Section 14.2.3.12.3, “FULLTEXT Indexes”



---

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **INNODB\_FT\_DEFAULT\_STOPWORD**

Section 14.2.3.12.3, “FULLTEXT Indexes”

## **innodb\_ft\_default\_stopword**

Section 12.9.4, “Full-Text Stopwords”

## **INNODB\_FT\_DELETED**

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 19.30.23, “The `INFORMATION_SCHEMA` `INNODB_FT_BEING_DELETED` Table”

## **INNODB\_FT\_INDEX\_CACHE**

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **INNODB\_FT\_INDEX\_TABLE**

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **INNODB\_SYS\_FOREIGN**

Section 1.8.6.2, “FOREIGN KEY Constraints”

Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **INNODB\_SYS\_FOREIGN\_COLS**

Section 1.8.6.2, “FOREIGN KEY Constraints”

Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **INNODB\_SYS\_INDEXES**

Section 14.2.3.12.3, “FULLTEXT Indexes”

## **INNODB\_SYS\_TABLES**

Section 14.2.3.12.3, “FULLTEXT Indexes”

## **K**

[index top [3949]]

## **KEY\_COLUMN\_USAGE**

Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”

Section 19.11, “The `INFORMATION_SCHEMA` `KEY_COLUMN_USAGE` Table”

## **O**

[index top [3949]]

## **OPTIMIZER\_TRACE**

Section 19.12, “The `INFORMATION_SCHEMA` `OPTIMIZER_TRACE` Table”

## **P**

[index top [3949]]

## **PARAMETERS**

Section 19.13, “The `INFORMATION_SCHEMA` `PARAMETERS` Table”

Section 19.19, “The `INFORMATION_SCHEMA` `ROUTINES` Table”

## **PARTITIONS**

Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”

Section 17.2.7, “How MySQL Partitioning Handles `NULL`”

Section 17.3.5, “Obtaining Information About Partitions”

Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”

Chapter 17, *Partitioning*

Section 19.14, “The `INFORMATION_SCHEMA` `PARTITIONS` Table”

## **PLUGINS**

Section 5.1.8.2, “Obtaining Server Plugin Information”

Section 19.15, “The `INFORMATION_SCHEMA` `PLUGINS` Table”

## **PROCESSLIST**

Section 13.7.5.30, “`SHOW PROCESSLIST` Syntax”

Section 8.12.5, “Examining Thread Information”

Section 19.16, “The `INFORMATION_SCHEMA` `PROCESSLIST` Table”

## **PROFILING**

Section 13.7.5.31, “`SHOW PROFILE` Syntax”

Section 19.17, “The `INFORMATION_SCHEMA` `PROFILING` Table”

## **R**

[index top [3949]]

## **REFERENTIAL\_CONSTRAINTS**

Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”

Section 19.18, “The `INFORMATION_SCHEMA` `REFERENTIAL_CONSTRAINTS` Table”

## **ROUTINES**

Section 13.7.5.29, “`SHOW PROCEDURE STATUS` Syntax”

---

Section 18.2.3, “Stored Routine Metadata”  
Section 19.13, “The `INFORMATION_SCHEMA` `PARAMETERS` Table”  
Section 19.19, “The `INFORMATION_SCHEMA` `ROUTINES` Table”

## S

[[index top \[3949\]](#)]

### **SCHEMA\_PRIVILEGES**

Section 19.21, “The `INFORMATION_SCHEMA` `SCHEMA_PRIVILEGES` Table”

### **SCHEMATA**

Section 6.2.2, “Privilege System Grant Tables”  
Section 19.20, “The `INFORMATION_SCHEMA` `SCHEMATA` Table”

### **SESSION\_STATUS**

Section 19.9, “The `INFORMATION_SCHEMA` `GLOBAL_STATUS` and `SESSION_STATUS` Tables”

### **SESSION\_VARIABLES**

Section 19.10, “The `INFORMATION_SCHEMA` `GLOBAL_VARIABLES` and `SESSION_VARIABLES` Tables”

### **STATISTICS**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”  
Section 19.22, “The `INFORMATION_SCHEMA` `STATISTICS` Table”

## T

[[index top \[3949\]](#)]

### **TABLE\_CONSTRAINTS**

Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”  
Section 19.25, “The `INFORMATION_SCHEMA` `TABLE_CONSTRAINTS` Table”

### **TABLE\_PRIVILEGES**

Section 19.26, “The `INFORMATION_SCHEMA` `TABLE_PRIVILEGES` Table”

### **TABLES**

Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”

Section 19.23, “The `INFORMATION_SCHEMA` `TABLES` Table”

### **TABLESPACES**

Section 19.24, “The `INFORMATION_SCHEMA` `TABLESPACES` Table”

### **TP\_THREAD\_GROUP\_STATE**

Section 19.31.2, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATE` Table”  
Section 19.31, “Thread Pool `INFORMATION_SCHEMA` Tables”  
Section 8.11.6.1, “Thread Pool Components and Installation”

### **TP\_THREAD\_GROUP\_STATS**

Section 19.31.1, “The `INFORMATION_SCHEMA` `TP_THREAD_STATE` Table”  
Section 19.31, “Thread Pool `INFORMATION_SCHEMA` Tables”  
Section 8.11.6.1, “Thread Pool Components and Installation”

### **TP\_THREAD\_STATE**

Section 19.31, “Thread Pool `INFORMATION_SCHEMA` Tables”  
Section 8.11.6.1, “Thread Pool Components and Installation”

### **TRIGGERS**

Section 13.7.5.13, “`SHOW CREATE TRIGGER` Syntax”  
Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”  
Section 19.27, “The `INFORMATION_SCHEMA` `TRIGGERS` Table”  
Section 18.3.2, “Trigger Metadata”

## U

[[index top \[3949\]](#)]

### **USER\_PRIVILEGES**

Section 19.28, “The `INFORMATION_SCHEMA` `USER_PRIVILEGES` Table”

## V

[[index top \[3949\]](#)]

### **VIEWS**

Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”  
Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”  
Section 19.29, “The `INFORMATION_SCHEMA` `VIEWS` Table”  
Section 18.5.4, “View Metadata”



---

# Join Types Index

[A](#) | [C](#) | [E](#) | [F](#) | [I](#) | [R](#) | [S](#) | [U](#)

## A

[[index top \[3953\]](#)]

### ALL

Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”  
Section 8.13.8, “Nested-Loop Join Algorithms”

## C

[[index top \[3953\]](#)]

### const

Section 8.8.2, “EXPLAIN Output Format”  
Section 8.13.13, “ORDER BY Optimization”  
Section 13.2.9, “SELECT Syntax”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

## E

[[index top \[3953\]](#)]

### eq\_ref

Section 8.8.2, “EXPLAIN Output Format”  
Section 14.8.1, “MERGE Table Advantages and Disadvantages”  
Section 8.13.12.3, “Batched Key Access Joins”  
Section 8.13.4, “Index Condition Pushdown Optimization”  
Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”

## F

[[index top \[3953\]](#)]

### fulltext

Section 8.8.2, “EXPLAIN Output Format”

## I

[[index top \[3953\]](#)]

### index

Section 8.8.2, “EXPLAIN Output Format”  
Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”  
Section 8.13.8, “Nested-Loop Join Algorithms”

### index\_merge

Section 8.8.2, “EXPLAIN Output Format”

Section 8.13.2, “Index Merge Optimization”

### index\_subquery

Section 8.8.2, “EXPLAIN Output Format”  
Section 13.2.10.10, “Optimizing Subqueries”  
Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”

## R

[[index top \[3953\]](#)]

### range

Section 8.8.2, “EXPLAIN Output Format”  
Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”  
Section 8.13.4, “Index Condition Pushdown Optimization”  
Section 8.13.2, “Index Merge Optimization”  
Section 8.13.14.1, “Loose Index Scan”  
Section 8.13.8, “Nested-Loop Join Algorithms”  
Section 8.13.1, “Range Optimization”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

### ref

Section 8.8.2, “EXPLAIN Output Format”  
Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”  
Section 14.8.1, “MERGE Table Advantages and Disadvantages”  
Section 8.13.12.3, “Batched Key Access Joins”  
Section 8.13.4, “Index Condition Pushdown Optimization”  
Section 8.13.16.3, “Optimizing Subqueries in the FROM Clause (Derived Tables)”  
Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”

### ref\_or\_null

Section 8.8.2, “EXPLAIN Output Format”  
Section 8.13.6, “IS NULL Optimization”  
Section 8.13.4, “Index Condition Pushdown Optimization”  
Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”

## S

[[index top \[3953\]](#)]

### system

Section 8.8.2, “EXPLAIN Output Format”  
Section 13.2.9, “SELECT Syntax”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

---

## U

[\[index top \[3953\]\]](#)

### **unique\_subquery**

[Section 8.8.2, “EXPLAIN Output Format”](#)

[Section 13.2.10.10, “Optimizing Subqueries”](#)

[Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”](#)

---

# Operator Index

[Symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [I](#) | [L](#) | [N](#) | [O](#) | [R](#) | [X](#)

## Symbols

[\[index top \[3955\]\]](#)

–

Section 12.6.1, “Arithmetic Operators”  
Section 12.10, “Cast Functions and Operators”  
Section 12.7, “Date and Time Functions”  
Section 11.1.1, “Numeric Type Overview”  
Section 17.6, “Restrictions and Limitations on Partitioning”

!

Section 9.5, “Expression Syntax”  
Section 12.3.3, “Logical Operators”  
Section 12.3.1, “Operator Precedence”

!=

Section 12.3.2, “Comparison Functions and Operators”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

%

Section 12.6.1, “Arithmetic Operators”

&

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 12.12, “Bit Functions”  
Section 17.6, “Restrictions and Limitations on Partitioning”

&&

Section 12.3.3, “Logical Operators”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

>

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 8.3.8, “Comparison of B-Tree and Hash Indexes”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

>>

Section 12.12, “Bit Functions”

Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.6, “Restrictions and Limitations on Partitioning”

>=

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 8.3.8, “Comparison of B-Tree and Hash Indexes”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

<

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 8.3.8, “Comparison of B-Tree and Hash Indexes”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 3.3.4.6, “Working with `NULL` Values”

<>

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 3.3.4.6, “Working with `NULL` Values”

<<

Section 12.12, “Bit Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.6, “Restrictions and Limitations on Partitioning”

<=

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 8.3.8, “Comparison of B-Tree and Hash Indexes”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

---

## <=>

Section 8.8.2, “[EXPLAIN Output Format](#)”  
Section 12.3.2, “[Comparison Functions and Operators](#)”  
Section 1.8.4, “[MySQL Extensions to Standard SQL](#)”  
Section 12.3.1, “[Operator Precedence](#)”  
Section 8.13.1.2, “[The Range Access Method for Multiple-Part Indexes](#)”  
Section 8.13.1.1, “[The Range Access Method for Single-Part Indexes](#)”  
Section 12.2, “[Type Conversion in Expression Evaluation](#)”

## \*

Section 12.6.1, “[Arithmetic Operators](#)”  
Section 11.1.1, “[Numeric Type Overview](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”

## +

Section 12.6.1, “[Arithmetic Operators](#)”  
Section 12.10, “[Cast Functions and Operators](#)”  
Section 12.7, “[Date and Time Functions](#)”  
Section 11.1.1, “[Numeric Type Overview](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”

## /

Section 12.6.1, “[Arithmetic Operators](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”  
Section 5.1.4, “[Server System Variables](#)”

## =

Section 8.8.2, “[EXPLAIN Output Format](#)”  
Section 13.7.4, “[SET Syntax](#)”  
Section 12.3.4, “[Assignment Operators](#)”  
Section 12.3.2, “[Comparison Functions and Operators](#)”  
Section 8.3.8, “[Comparison of B-Tree and Hash Indexes](#)”  
Section 1.8.4, “[MySQL Extensions to Standard SQL](#)”  
Section 12.3.1, “[Operator Precedence](#)”  
Section E.4, “[Restrictions on Subqueries](#)”  
Section 12.5.1, “[String Comparison Functions](#)”  
Section 8.13.1.2, “[The Range Access Method for Multiple-Part Indexes](#)”  
Section 8.13.1.1, “[The Range Access Method for Single-Part Indexes](#)”  
Section 9.4, “[User-Defined Variables](#)”  
Section 3.3.4.6, “[Working with NULL Values](#)”

## ^

Section 12.12, “[Bit Functions](#)”  
Section 9.5, “[Expression Syntax](#)”  
Section 12.3.1, “[Operator Precedence](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”

## |

Section 12.12, “[Bit Functions](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”

## ||

Section 10.1.7.3, “[COLLATE Clause Precedence](#)”  
Section 9.5, “[Expression Syntax](#)”  
Section 12.3.3, “[Logical Operators](#)”  
Section 1.8.4, “[MySQL Extensions to Standard SQL](#)”  
Section 12.3.1, “[Operator Precedence](#)”  
Section 10.1.9.1, “[Result Strings](#)”  
Section 5.1.7, “[Server SQL Modes](#)”

## ~

Section 12.12, “[Bit Functions](#)”  
Section 17.6, “[Restrictions and Limitations on Partitioning](#)”

## A

[\[index top \[3955\]\]](#)

## AND

Section 13.1.14, “[CREATE TABLE Syntax](#)”  
Section 8.3.8, “[Comparison of B-Tree and Hash Indexes](#)”  
Section 8.13.2, “[Index Merge Optimization](#)”  
Section 12.3.3, “[Logical Operators](#)”  
Section 1.8.4, “[MySQL Extensions to Standard SQL](#)”  
Section 8.13.16.4, “[Optimizing Subqueries with EXISTS Strategy](#)”  
Section E.4, “[Restrictions on Subqueries](#)”  
Section 3.6.7, “[Searching on Two Keys](#)”  
Section 3.3.4.2, “[Selecting Particular Rows](#)”  
Section 12.5.1, “[String Comparison Functions](#)”  
Section 8.13.2.1, “[The Index Merge Intersection Access Algorithm](#)”  
Section 8.13.1.2, “[The Range Access Method for Multiple-Part Indexes](#)”  
Section 8.13.1.1, “[The Range Access Method for Single-Part Indexes](#)”  
Section 18.5.2, “[View Processing Algorithms](#)”

## B

[\[index top \[3955\]\]](#)

## BETWEEN

Section 8.8.2, “[EXPLAIN Output Format](#)”  
Section 12.3.2, “[Comparison Functions and Operators](#)”  
Section 8.3.8, “[Comparison of B-Tree and Hash Indexes](#)”  
Section 8.13.1.2, “[The Range Access Method for Multiple-Part Indexes](#)”  
Section 8.13.1.1, “[The Range Access Method for Single-Part Indexes](#)”

---

Section 12.2, “Type Conversion in Expression Evaluation”

## **BINARY**

Section 12.10, “Cast Functions and Operators”  
Section 3.3.4.7, “Pattern Matching”  
Section 3.3.4.4, “Sorting Rows”  
Section 10.1.7.7, “The `BINARY` Operator”

## **BINARY str**

Section 12.10, “Cast Functions and Operators”

## **C**

[[index top \[3955\]](#)]

## **CASE**

Section 13.6.5.1, “`CASE` Syntax”  
Section 12.4, “Control Flow Functions”  
Section 9.5, “Expression Syntax”  
Section 1.8.4, “MySQL Extensions to Standard SQL”

## **CASE value WHEN END**

Section 12.4, “Control Flow Functions”

## **CASE WHEN END**

Section 12.4, “Control Flow Functions”

## **CASE WHEN expr1 = expr2 THEN NULL ELSE expr1 END**

Section 12.4, “Control Flow Functions”

## **D**

[[index top \[3955\]](#)]

## **DIV**

Section 12.6.1, “Arithmetic Operators”  
Section 17.6, “Restrictions and Limitations on Partitioning”

## **E**

[[index top \[3955\]](#)]

## **expr BETWEEN min AND max**

Section 12.3.2, “Comparison Functions and Operators”

## **expr LIKE pat**

Section 12.5.1, “String Comparison Functions”

## **expr NOT BETWEEN min AND max**

Section 12.3.2, “Comparison Functions and Operators”

## **expr NOT LIKE pat**

Section 12.5.1, “String Comparison Functions”

## **expr NOT REGEXP pat**

Section 12.5.2, “Regular Expressions”

## **expr NOT RLIKE pat**

Section 12.5.2, “Regular Expressions”

## **expr REGEXP pat**

Section 12.5.2, “Regular Expressions”

## **expr RLIKE pat**

Section 12.5.2, “Regular Expressions”

## **expr1 SOUNDS LIKE expr2**

Section 12.5, “String Functions”

## **I**

[[index top \[3955\]](#)]

## **IS**

Section 12.3.1, “Operator Precedence”

## **IS boolean\_value**

Section 12.3.2, “Comparison Functions and Operators”

## **IS NOT boolean\_value**

Section 12.3.2, “Comparison Functions and Operators”

## **IS NOT NULL**

Section 12.3.2, “Comparison Functions and Operators”  
Section C.5.5.3, “Problems with `NULL` Values”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 3.3.4.6, “Working with `NULL` Values”

## **IS NULL**

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 8.13.6, “`IS NULL` Optimization”  
Section 12.3.2, “Comparison Functions and Operators”  
Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”  
Section C.5.5.3, “Problems with `NULL` Values”  
Section 5.1.4, “Server System Variables”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 3.3.4.6, “Working with `NULL` Values”

## **L**

[[index top \[3955\]](#)]

---

## LIKE

Section 4.5.1.4, “mysql Server-Side Help”  
Section 13.8.3, “HELP Syntax”  
Section 13.7.5.4, “SHOW CHARACTER SET Syntax”  
Section 13.7.5.5, “SHOW COLLATION Syntax”  
Section 13.7.5.6, “SHOW COLUMNS Syntax”  
Section 13.7.5.15, “SHOW DATABASES Syntax”  
Section 13.7.5.19, “SHOW EVENTS Syntax”  
Section 13.7.5.25, “SHOW OPEN TABLES Syntax”  
Section 13.7.5.29, “SHOW PROCEDURE STATUS Syntax”  
Section 13.7.5.36, “SHOW STATUS Syntax”  
Section 13.7.5.37, “SHOW TABLE STATUS Syntax”  
Section 13.7.5.38, “SHOW TABLES Syntax”  
Section 13.7.5.39, “SHOW TRIGGERS Syntax”  
Section 13.7.5.40, “SHOW VARIABLES Syntax”  
Section 10.1.9.3, “SHOW Statements and INFORMATION\_SCHEMA”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 12.10, “Cast Functions and Operators”  
Section 8.3.8, “Comparison of B-Tree and Hash Indexes”  
Section 19.32, “Extensions to SHOW Statements”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 3.3.4.7, “Pattern Matching”  
Section 20.2.3.2.1.1, “Pre-Filtering by Instrument”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 6.2.3, “Specifying Account Names”  
Section 12.5.1, “String Comparison Functions”  
Section 9.1.1, “String Literals”  
Section 5.1.5.1, “Structured System Variables”  
Section 11.4.1, “The CHAR and VARCHAR Types”  
Section 11.4.5, “The SET Type”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 5.1.5, “Using System Variables”

## LIKE 'pattern'

Section 13.7.5, “SHOW Syntax”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”

## LIKE ... ESCAPE

Section C.5.8, “Known Issues in MySQL”

## N

[index top [3955]]

## N % M

Section 12.6.1, “Arithmetic Operators”

Section 12.6.2, “Mathematical Functions”

## N MOD M

Section 12.6.1, “Arithmetic Operators”  
Section 12.6.2, “Mathematical Functions”

## NOT

Section 12.3.3, “Logical Operators”  
Section 5.1.7, “Server SQL Modes”

## NOT LIKE

Section 3.3.4.7, “Pattern Matching”  
Section 12.5.1, “String Comparison Functions”

## NOT REGEXP

Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 3.3.4.7, “Pattern Matching”  
Section 12.5.1, “String Comparison Functions”

## NOT RLIKE

Section 3.3.4.7, “Pattern Matching”  
Section 12.5.1, “String Comparison Functions”

## O

[index top [3955]]

## OR

Section 13.7.1.4, “GRANT Syntax”  
Section 9.5, “Expression Syntax”  
Section 8.13.2, “Index Merge Optimization”  
Section 12.3.3, “Logical Operators”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”  
Section 8.13.16.4, “Optimizing Subqueries with EXISTS Strategy”  
Section 3.6.7, “Searching on Two Keys”  
Section 3.3.4.2, “Selecting Particular Rows”  
Section 5.1.7, “Server SQL Modes”  
Section 12.5.1, “String Comparison Functions”  
Section 8.13.2.3, “The Index Merge Sort-Union Access Algorithm”  
Section 8.13.2.2, “The Index Merge Union Access Algorithm”  
Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”

## R

[index top [3955]]

## REGEXP

Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 12.3.1, “Operator Precedence”

---

[Section 3.3.4.7, "Pattern Matching"](#)  
[Section 12.5.2, "Regular Expressions"](#)  
[Section E.7, "Restrictions on Character Sets"](#)

## **RLIKE**

[Section 3.3.4.7, "Pattern Matching"](#)  
[Section 12.5.2, "Regular Expressions"](#)  
[Section E.7, "Restrictions on Character Sets"](#)

## **X**

[\[index top \[3955\]\]](#)

## **XOR**

[Section 12.17.1, "GROUP BY \(Aggregate\) Functions"](#)  
[Section 12.3.3, "Logical Operators"](#)





---

# Option Index

[Symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#) | [Y](#)

## Symbols

[[index top \[3961\]](#)]

--

Section 1.8.5.5, "'--' as the Start of a Comment"  
Section 4.8.2, "[replace](#) — A String-Replacement Utility"

-#

Section 4.6.3.1, "[myisamchk](#) General Options"  
Section 4.5.1.1, "[mysql](#) Options"  
Section 4.6.5, "[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables"  
Section 4.6.6, "[mysql\\_config\\_editor](#) — MySQL Configuration Utility"  
Section 4.6.8, "[mysqlbinlog](#) — Utility for Processing Binary Log Files"  
Section 4.5.2, "[mysqladmin](#) — Client for Administering a MySQL Server"  
Section 4.5.3, "[mysqlcheck](#) — A Table Maintenance Program"  
Section 4.5.4, "[mysqldump](#) — A Database Backup Program"  
Section 4.5.5, "[mysqlimport](#) — A Data Import Program"  
Section 4.5.6, "[mysqlshow](#) — Display Database, Table, and Column Information"  
Section 4.5.7, "[mysqlslap](#) — Load Emulation Client"  
Section 4.7.3, "[my\\_print\\_defaults](#) — Display Options from Option Files"  
Section 4.4.1, "[comp\\_err](#) — Compile MySQL Error Message File"  
Section 4.8.2, "[replace](#) — A String-Replacement Utility"  
Section 5.1.3, "Server Command Options"  
Section 22.4.3, "The DBUG Package"

-1

Section 4.5.3, "[mysqlcheck](#) — A Table Maintenance Program"

-?

Section 4.6.3.1, "[myisamchk](#) General Options"  
Section 4.5.1.1, "[mysql](#) Options"  
Section 4.6.2, "[myisam\\_ftdump](#) — Display Full-Text Index information"  
Section 4.6.4, "[myisamlog](#) — Display MyISAM Log File Contents"  
Section 4.6.5, "[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables"

Section 4.6.6, "[mysql\\_config\\_editor](#) — MySQL Configuration Utility"  
Section 4.6.15, "[mysql\\_waitpid](#) — Kill Process and Wait for Its Termination"  
Section 4.6.16, "[mysql\\_zap](#) — Kill Processes That Match a Pattern"  
Section 4.6.7, "[mysqlaccess](#) — Client for Checking Access Privileges"  
Section 4.6.8, "[mysqlbinlog](#) — Utility for Processing Binary Log Files"  
Section 4.6.10, "[mysqlhotcopy](#) — A Database Backup Program"  
Section 4.5.2, "[mysqladmin](#) — Client for Administering a MySQL Server"  
Section 4.5.3, "[mysqlcheck](#) — A Table Maintenance Program"  
Section 4.5.4, "[mysqldump](#) — A Database Backup Program"  
Section 4.5.5, "[mysqlimport](#) — A Data Import Program"  
Section 4.5.6, "[mysqlshow](#) — Display Database, Table, and Column Information"  
Section 4.5.7, "[mysqlslap](#) — Load Emulation Client"  
Section 4.7.3, "[my\\_print\\_defaults](#) — Display Options from Option Files"  
Section 4.4.1, "[comp\\_err](#) — Compile MySQL Error Message File"  
Section 4.4.4, "[mysql\\_plugin](#) — Configure MySQL Server Plugins"  
Section 4.8.1, "[perror](#) — Explain Error Codes"  
Section 4.8.2, "[replace](#) — A String-Replacement Utility"  
Section 4.8.3, "[resolveip](#) — Resolve Host name to IP Address or Vice Versa"  
Section 5.1.3, "Server Command Options"  
Section 4.2.3.1, "Using Options on the Command Line"

## A

[[index top \[3961\]](#)]

-A

Section 4.5.1.1, "[mysql](#) Options"  
Section 4.5.3, "[mysqlcheck](#) — A Table Maintenance Program"  
Section 4.5.4, "[mysqldump](#) — A Database Backup Program"  
Section 4.6.3.4, "Other [myisamchk](#) Options"

-a

Section 7.6.4, "MyISAM Table Optimization"  
Section 4.6.9, "[mysqldumpslow](#) — Summarize Slow Query Log Files"  
Section 4.5.3, "[mysqlcheck](#) — A Table Maintenance Program"  
Section 4.5.7, "[mysqlslap](#) — Load Emulation Client"  
Section 4.6.3.4, "Other [myisamchk](#) Options"

---

## **--abort-slave-event-count**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--add-drop-database**

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 7.4.1, “Dumping Data in SQL Format with mysqldump”

## **--add-drop-table**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--add-drop-trigger**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--add-locks**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--addtodest**

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

## **--all**

Section 4.6.6, “mysql\_config\_editor — MySQL Configuration Utility”

## **--all-databases**

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”

Section 7.4.1, “Dumping Data in SQL Format with mysqldump”

Section 9.2.3, “Mapping of Identifiers to File Names”

Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”

Section 7.4.2, “Reloading SQL-Format Backups”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

Section 4.6.8.3, “Using mysqlbinlog to Back Up Binary Log Files”

## **--all-in-1**

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

## **--all-tablespaces**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--allow-keywords**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--allow-suspicious-udfs**

Section 5.1.3, “Server Command Options”

Section 22.3.2.6, “User-Defined Function Security Precautions”

## **--allowold**

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

## **--analyze**

Section 4.6.3.1, “myisamchk General Options”

Section 7.6.4, “MyISAM Table Optimization”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.6.3.4, “Other myisamchk Options”

## **--ansi**

Section 1.8.3, “Running MySQL in ANSI Mode”

Section 5.1.3, “Server Command Options”

## **antonio**

Section 6.3.7.3.2.3, “Unix Password Authentication with Proxy Users and Group Mapping”

## **--apply-slave-statements**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--audit-log**

Section 6.3.11.4, “Audit Log Plugin Options and Variables”

Section 6.3.11.1, “Installing the Audit Log Plugin”

## **--auto-generate-sql**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--auto-generate-sql-add-autoincrement**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--auto-generate-sql-execute-number**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--auto-generate-sql-guid-primary**

Section 4.5.7, “mysqlslap — Load Emulation Client”

---

## **--auto-generate-sql-load-type**

Section 4.5.7, “[mysqslap — Load Emulation Client](#)”

## **--auto-generate-sql-secondary-indexes**

Section 4.5.7, “[mysqslap — Load Emulation Client](#)”

## **--auto-generate-sql-unique-query-number**

Section 4.5.7, “[mysqslap — Load Emulation Client](#)”

## **--auto-generate-sql-unique-write-number**

Section 4.5.7, “[mysqslap — Load Emulation Client](#)”

## **--auto-generate-sql-write-number**

Section 4.5.7, “[mysqslap — Load Emulation Client](#)”

## **--auto-rehash**

Section 4.5.1.2, “[mysql Commands](#)”

Section 4.5.1.1, “[mysql Options](#)”

## **--auto-repair**

Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”

## **--auto-vertical-output**

Section 4.5.1.1, “[mysql Options](#)”

## **--autocommit**

Section 5.1.4, “[Server System Variables](#)”

## **B**

[[index top \[3961\]](#)]

### **-B**

Section 4.6.3.3, “[myisamchk Repair Options](#)”

Section 4.5.1.1, “[mysql Options](#)”

Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”

Section 4.5.4, “[mysqldump — A Database Backup Program](#)”

### **-b**

Section 4.5.1.1, “[mysql Options](#)”

Section 4.6.5, “[myisampack — Generate Compressed, Read-Only MyISAM Tables](#)”

Section 4.6.7, “[mysqlaccess — Client for Checking Access Privileges](#)”

Section 4.5.2, “[mysqladmin — Client for Administering a MySQL Server](#)”

Section 4.4.4, “[mysql\\_plugin — Configure MySQL Server Plugins](#)”

Section 4.6.3.4, “[Other myisamchk Options](#)”

Section 5.1.3, “[Server Command Options](#)”

## **--back\_log**

Section 2.6, “[Installing MySQL on Solaris and OpenSolaris](#)”

## **--backup**

Section 4.6.3.3, “[myisamchk Repair Options](#)”

Section 4.6.5, “[myisampack — Generate Compressed, Read-Only MyISAM Tables](#)”

## **--base64-output**

Section 4.6.8.2, “[mysqlbinlog Row Event Display](#)”

Section 16.1.2.1, “[Advantages and Disadvantages of Statement-Based and Row-Based Replication](#)”

Section 4.6.8, “[mysqlbinlog — Utility for Processing Binary Log Files](#)”

Section 16.1.3.1, “[GTID Concepts](#)”

Section 16.1.2.2, “[Usage of Row-Based Logging and Replication](#)”

## **--basedir**

Section 4.4.3, “[mysql\\_install\\_db — Initialize MySQL Data Directory](#)”

Section 4.4.4, “[mysql\\_plugin — Configure MySQL Server Plugins](#)”

Section 4.4.7, “[mysql\\_upgrade — Check and Upgrade MySQL Tables](#)”

Section 4.3.3, “[mysql.server — MySQL Server Startup Script](#)”

Section 4.3.2, “[mysqld\\_safe — MySQL Server Startup Script](#)”

Section 2.9.4, “[MySQL Source-Configuration Options](#)”

Section 5.6, “[Running Multiple MySQL Instances on One Machine](#)”

Section 5.1.3, “[Server Command Options](#)”

Section 5.1.4, “[Server System Variables](#)”

Section 2.10.1.3, “[Starting and Troubleshooting the MySQL Server](#)”

Section 2.10.1, “[Unix Postinstallation Procedures](#)”

## **basedir**

Section 2.3.4.2, “[Creating an Option File](#)”

Section 2.10.1.2, “[Starting and Stopping MySQL Automatically](#)”

Section 2.3.5, “[Troubleshooting a Microsoft Windows MySQL Server Installation](#)”

## **--batch**

Section 4.5.1.3, “[mysql Logging](#)”

Section 4.5.1.1, “[mysql Options](#)”

## **--big-tables**

Section 5.1.3, “[Server Command Options](#)”

---

## **--binary-mode**

Section 4.5.1.2, “mysql Commands”  
Section 4.5.1.1, “mysql Options”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

## **--bind-address**

Section 4.5.1.1, “mysql Options”  
Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 5.1.9.2, “Configuring the MySQL Server to Permit IPv6 Connections”  
Section 5.1.9.4, “Connecting Using IPv6 Nonlocal Host Addresses”  
Section 5.1.9.3, “Connecting Using the IPv6 Local Host Address”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 4.5.4, “mysqldump — A Database Backup Program”  
Section 4.5.5, “mysqlimport — A Data Import Program”  
Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”  
Section 4.3.4, “mysqld\_multi — Manage Multiple MySQL Servers”  
Section 5.1.9, “IPv6 Support”  
Section 5.1.9.5, “Obtaining an IPv6 Address from a Broker”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **--binlog-checksum**

Section 16.1.4.4, “Binary Log Options and Variables”

## **--binlog-do-db**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”  
Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”  
Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.2.4, “The Binary Log”

## **--binlog-format**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.2.4.1, “Binary Logging Formats”

Section 5.1.3, “Server Command Options”  
Section 5.2.4.2, “Setting The Binary Log Format”

## **--binlog-ignore-db**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”  
Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.2.4, “The Binary Log”

## **--binlog-row-event-max-size**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”  
Section 5.2.4.2, “Setting The Binary Log Format”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **--binlog-rows-query-log-events**

Section 16.1.4.4, “Binary Log Options and Variables”

## **--binlog\_format**

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

## **--block-search**

Section 4.6.3.4, “Other myisamchk Options”

## **--bootstrap**

Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”  
Section 5.1.3, “Server Command Options”

## **--brief**

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

## **C**

[index top [3961]]

## **-C**

Section 4.6.3.2, “myisamchk Check Options”  
Section 4.5.1.1, “mysql Options”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 4.5.4, “mysqldump — A Database Backup Program”

---

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”  
Section 5.1.3, “Server Command Options”

## **-c**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.1, “`innochecksum` — Offline InnoDB File Checksum Utility”  
Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

## **--cflags**

Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **--character-set-client-handshake**

Section 5.1.3, “Server Command Options”  
Section 10.1.14.7.1, “The `cp932` Character Set”

## **--character-set-filesystem**

Section 5.1.3, “Server Command Options”

## **--character-set-server**

Section 10.5, “Character Set Configuration”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 16.4.1.2, “Replication and Character Sets”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 5.1.3, “Server Command Options”

## **--character-sets-dir**

Section 4.6.3.3, “`myisamchk` Repair Options”

Section 4.5.1.1, “`mysql` Options”  
Section C.5.2.17, “Can't initialize character set”  
Section 10.5, “Character Set Configuration”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 5.1.3, “Server Command Options”

## **--character\_set\_server**

Section 2.9.4, “MySQL Source-Configuration Options”

## **--charset**

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

## **--check**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--check-only-changed**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--check-upgrade**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

## **--checkpoint**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## **--chroot**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 5.1.3, “Server Command Options”

## **CMAKE\_INSTALL\_PREFIX**

Section 22.2.4.3, “Compiling and Installing Plugin Libraries”



---

Section 2.9.3, “Installing MySQL from a Development Source Tree”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.6.3, “Running Multiple MySQL Instances on Unix”

## **CMAKE\_PREFIX\_PATH**

Section 2.9.4, “MySQL Source-Configuration Options”

## **--collation-server**

Section 10.5, “Character Set Configuration”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 16.4.1.2, “Replication and Character Sets”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 5.1.3, “Server Command Options”

## **--collation\_server**

Section 2.9.4, “MySQL Source-Configuration Options”

## **--column-names**

Section 4.5.1.1, “mysql Options”  
Section 4.2.3.2, “Program Option Modifiers”

## **--column-type-info**

Section 4.5.1.1, “mysql Options”

## **--columns**

Section 4.5.5, “mysqlimport — A Data Import Program”

## **--comments**

Section 4.5.1.1, “mysql Options”  
Section 4.5.4, “mysqldump — A Database Backup Program”

## **--commit**

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”  
Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--comp**

Section 4.2.3, “Specifying Program Options”

## **--compact**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--compatible**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **COMPILATION\_COMMENT**

Section 5.1.4, “Server System Variables”

## **--complete-insert**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--compr**

Section 4.2.3, “Specifying Program Options”

## **--compress**

Section 4.5.1.1, “mysql Options”  
Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 6.3.9.1, “Basic SSL Concepts”  
Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 4.5.4, “mysqldump — A Database Backup Program”  
Section 4.5.5, “mysqlimport — A Data Import Program”  
Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”  
Section 4.5.7, “mysqlslap — Load Emulation Client”  
Section 4.2.3, “Specifying Program Options”

## **--concurrency**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--config-file**

Section 4.7.3, “my\_print\_defaults — Display Options from Option Files”  
Section 4.3.4, “mysqld\_multi — Manage Multiple MySQL Servers”

## **--connect-expired-password**

Section 4.5.1.1, “mysql Options”

## **--console**

Section 14.2.4.5, “InnoDB General Troubleshooting”  
Section 14.2.4.4, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”  
Section 14.2.2.1, “Creating the InnoDB Tablespace”  
Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”  
Section 5.1.3, “Server Command Options”  
Section 2.3.4.4, “Starting the Server for the First Time”  
Section 5.2.2, “The Error Log”

## **--copy**

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

## **--core-file**

Section 22.4.1.4, “Debugging mysqld under gdb”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

---

## **--core-file-size**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 5.1.3, “Server Command Options”

## **--correct-checksum**

Section 4.6.3.3, “`myisamchk` Repair Options”

## **--count**

Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

## **--create**

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--create-options**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--create-schema**

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--csv**

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--cxxflags**

Section 2.9.5, “Dealing with Problems Compiling MySQL”  
Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **D**

[index top [3961]]

## **-D**

Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.5.1.1, “`mysql` Options”  
Section C.5.2.17, “`Can't initialize character set`”  
Section 10.3, “Adding a Character Set”  
Section 2.1.2.2, “Choosing a Distribution Format”  
Section 22.4.1.1, “Compiling MySQL for Debugging”  
Section 21.7.1, “Compiling Programs with `libmysqld`”  
Section 22.4.2, “Debugging a MySQL Client”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

Section 2.9.2, “Installing MySQL from a Standard Source Distribution”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 22.1.1, “MySQL Threads”  
Chapter 17, *Partitioning*  
Section 20.2.1, “Performance Schema Build Configuration”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 14.10, “The `EXAMPLE` Storage Engine”  
Section 14.9, “The `FEDERATED` Storage Engine”  
Section 5.7, “Tracing `mysqld` Using DTrace”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## **-d**

Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.6.1, “`innochecksum` — Offline InnoDB File Checksum Utility”  
Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 4.6.3.4, “Other `myisamchk` Options”  
Section 5.1.4, “Server System Variables”

## **--data-file-length**

Section 4.6.3.3, “`myisamchk` Repair Options”

## **--database**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **--databases**

Section 7.4.5.2, “Copy a Database from one Server to Another”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.1, “Dumping Data in SQL Format with `mysqldump`”

---

Section 7.4.5.1, “Making a Copy of a Database”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 7.4.2, “Reloading SQL-Format Backups”

## **--datadir**

Section 2.3.4.2, “Creating an Option File”  
Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.6.3, “Running Multiple MySQL Instances on Unix”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.6.1, “Setting Up Multiple Data Directories”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section 4.2.3.3, “Using Option Files”

## **datadir**

Section 2.3.4.2, “Creating an Option File”  
Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”  
Section 2.10.1.2, “Starting and Stopping MySQL Automatically”  
Section 2.3.5, “Troubleshooting a Microsoft Windows MySQL Server Installation”  
Section E.10.6, “Windows Platform Limitations”

## **--db**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

## **--debug**

Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 22.4.1.1, “Compiling MySQL for Debugging”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 2.3.4.5, “Starting MySQL from the Windows Command Line”  
Section 22.4.3, “The DEBUG Package”

## **--debug-check**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

## **--debug-info**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”



---

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

### **--debug-sync-timeout**

Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

### **--default-auth**

Section 4.5.1.1, “mysql Options”  
Section 21.8.13, “C API Client Plugin Functions”  
Section 22.2.4.2.3, “Client Plugin Descriptors”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 6.3.7, “Pluggable Authentication”  
Section 6.3.7.1, “The Native Authentication Plugins”  
Section 22.2.4.9.3, “Using the Authentication Plugins”

### **--default-authentication-plugin**

Section 13.7.1.2, “`CREATE USER` Syntax”  
Section 5.1.3, “Server Command Options”

### **--default-character-set**

Section 4.5.1.1, “mysql Options”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 10.5, “Character Set Configuration”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.1.5, “Executing SQL Statements from a Text File”  
Section 5.1.4, “Server System Variables”  
Section 4.5.1.6.2, “Unicode Support on Windows”  
Section 6.3.1, “User Names and Passwords”

### **--default-storage-engine**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.1, “Setting the Storage Engine”  
Section 14.2.2.8, “Turning Off InnoDB”

### **default-storage-engine**

Section 14.1, “Setting the Storage Engine”

### **--default-time-zone**

Section 10.6, “MySQL Server Time Zone Support”  
Section 16.4.1.28, “Replication and Time Zones”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

### **--default-tmp-storage-engine**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.3, “Server Command Options”  
Section 14.2.2.8, “Turning Off InnoDB”

### **--default.key\_buffer\_size**

Section 5.1.5.1, “Structured System Variables”

### **DEFAULT\_CHARSET**

Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 10.1.3.1, “Server Character Set and Collation”

### **DEFAULT\_COLLATION**

Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 10.1.3.1, “Server Character Set and Collation”

### **--defaults-extra-file**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”  
Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”  
Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 4.2.3.3, “Using Option Files”

---

## **--defaults-file**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”

Section 14.2.1.2, “Configuring InnoDB”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 6.1.2.1, “End-User Guidelines for Password Security”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 21.7.3, “Options with the Embedded Server”

Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 5.6.3, “Running Multiple MySQL Instances on Unix”

Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”

Section 5.6.2.1, “Starting Multiple MySQL Instances at the Windows Command Line”

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## **--defaults-group-suffix**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 2.12, “Environment Variables”

## **--delay-key-write**

Section 14.3.1, “MyISAM Startup Options”

Section 8.10.5, “External Locking”

Section 5.1.3, “Server Command Options”

Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## **--delay\_key\_write**

Section 5.1.4, “Server System Variables”

Section 5.1.5, “Using System Variables”

## **--delayed-insert**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--delete**

Section 4.5.5, “`mysqlimport` — A Data Import Program”

## **--delete-master-logs**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--delimiter**

Section 4.5.1.1, “mysql Options”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--demangle**

Section 22.4.1.5, “Using a Stack Trace”

## **--des-key-file**

Section 13.7.6.3, “FLUSH Syntax”

Section 12.13, “Encryption and Compression Functions”

Section 5.1.3, “Server Command Options”

## **--description**

Section 4.6.3.4, “Other `myisamchk` Options”

## **--detach**

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--disable**

Section 4.2.3.2, “Program Option Modifiers”

## **--disable-auto-rehash**

Section 4.5.1.1, “mysql Options”

## **--disable-gtid-unsafe-statements**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

## **--disable-keys**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--disable-log-bin**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **--disable-named-commands**

Section 4.5.1.1, “mysql Options”

## **--disable-plugin\_name**

Section 5.1.8.1, “Installing and Uninstalling Plugins”

## **--disconnect-slave-event-count**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--dryrun**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

---

## **--dump**

Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”

## **--dump-date**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--dump-slave**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **E**

[index top [3961]]

### **-E**

Section 4.5.1.1, “`mysql` Options”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **-e**

Section 4.6.3.2, “`myisamchk` Check Options”

Section 4.6.3.1, “`myisamchk` General Options”

Section 4.6.3.3, “`myisamchk` Repair Options”

Section 4.5.1.1, “`mysql` Options”

Section 13.2.7, “LOAD XML Syntax”

Section 4.6.1, “`innochecksum` — Offline InnoDB File Checksum Utility”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 7.6.2, “How to Check `MyISAM` Tables for Errors”

Section 4.6.3.5, “Obtaining Table Information with `myisamchk`”

Section 4.2.3.1, “Using Options on the Command Line”

## **--embedded**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **--enable-cleartext-plugin**

Section 4.5.1.1, “`mysql` Options”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 6.3.7.5, “The Cleartext Client-Side Authentication Plugin”

## **--enable-named-pipe**

Section C.5.2.2, “Can't connect to [local] MySQL server”

Section 4.2.2, “Connecting to the MySQL Server”

Section 2.3.4.3, “Selecting a MySQL Server Type”

Section 5.1.3, “Server Command Options”

## **--enable-plugin\_name**

Section 5.1.8.1, “Installing and Uninstalling Plugins”

## **--enforce-gtid-consistency**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

Section 1.4, “What Is New in MySQL 5.6”

## **--engine**

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **event-scheduler**

Section 18.4.2, “Event Scheduler Configuration”

## **--event-scheduler**

Section 18.4.2, “Event Scheduler Configuration”

Section 5.1.3, “Server Command Options”

## **--events**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.5.3, “Dumping Stored Programs”

Section 7.4.5.4, “Dumping Table Definitions and Content Separately”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

## **--example**

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

## **--exclude-gtids**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **--execute**

Section 4.5.1.3, “`mysql` Logging”

Section 4.5.1.1, “`mysql` Options”

Section 4.2.3.1, “Using Options on the Command Line”

## **--exit-info**

Section 5.1.3, “Server Command Options”

---

## **--extend-check**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.6.3.3, “`myisamchk` Repair Options”

## **--extended**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--extended-insert**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--external-locking**

Section 14.3.1, “`MyISAM` Startup Options”  
Section 8.10.5, “External Locking”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 8.11.1, “System Factors and Startup Parameter Tuning”

## **--extra-file**

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

## **F**

[index top [3961]]

### **-F**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.1.2, “`mysql` Commands”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

### **-f**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.16, “`mysql_zap` — Kill Processes That Match a Pattern”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 22.4.1.5, “Using a Stack Trace”

## **--fast**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--federated**

Section 14.9, “The `FEDERATED` Storage Engine”

## **--fields-enclosed-by**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

## **--fields-escaped-by**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

## **--fields-optionally-enclosed-by**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

## **--fields-terminated-by**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

## **--fields-xxx**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--fix-db-names**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

---

## **--fix-table-names**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--flush**

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

## **--flush-logs**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 5.2, “MySQL Server Logs”

## **--flush-privileges**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--flush\_time**

Section 22.1.1, “MySQL Threads”

## **--flushlog**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## **--force**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 22.1.2, “The MySQL Test Suite”  
Section 3.5, “Using `mysql` in Batch Mode”

## **--force-if-open**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **--force-read**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **G**

[index top [3961]]

### **-G**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

### **-g**

Section 22.4.1.1, “Compiling MySQL for Debugging”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

### **--gdb**

Section 22.4.1.4, “Debugging `mysqld` under `gdb`”  
Section 5.1.3, “Server Command Options”

### **--general-log**

Section 5.1.3, “Server Command Options”

### **--general\_log**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.2.3, “The General Query Log”  
Section 1.4, “What Is New in MySQL 5.6”

### **--general\_log\_file**

Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.2.3, “The General Query Log”  
Section 1.4, “What Is New in MySQL 5.6”

### **--gtid-mode**

Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section 16.1.3.2, “Setting Up Replication Using GTIDs”  
Section 2.11.1, “Upgrading MySQL”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”



---

Section 1.4, “What Is New in MySQL 5.6”

## **--gtid\_mode**

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

## **H**

[index top [3961]]

### **-H**

Section 4.6.3.1, “myisamchk General Options”

Section 4.5.1.1, “mysql Options”

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.4.1, “comp\_err — Compile MySQL Error Message File”

### **-h**

Section 4.5.1.1, “mysql Options”

Section 4.2.2, “Connecting to the MySQL Server”

Section 4.6.2, “myisam\_ftdump — Display Full-Text Index information”

Section 4.6.6, “mysql\_config\_editor — MySQL Configuration Utility”

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.6.9, “mysqldumpslow — Summarize Slow Query Log Files”

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.5.5, “mysqlimport — A Data Import Program”

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 4.7.4, “resolve\_stack\_dump — Resolve Numeric Stack Trace Dump to Symbols”

Section 4.2.1, “Invoking MySQL Programs”

Section 5.1.3, “Server Command Options”

Section 1.2, “Typographical and Syntax Conventions”

Section 4.2.3.1, “Using Options on the Command Line”

## **HAVE\_EMBEDDED\_PRIVILEGE\_CONTROL**

Section 21.7, “libmysqld, the Embedded MySQL Server Library”

## **--header\_file**

Section 4.4.1, “comp\_err — Compile MySQL Error Message File”

## **--HELP**

Section 4.6.3.1, “myisamchk General Options”

## **--help**

Section 4.6.3.1, “myisamchk General Options”

Section 4.5.1.1, “mysql Options”

Section 4.6.2, “myisam\_ftdump — Display Full-Text Index information”

Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”

Section 4.6.6, “mysql\_config\_editor — MySQL Configuration Utility”

Section 4.6.11, “mysql\_convert\_table\_format — Convert Tables to Use a Given Storage Engine”

Section 4.6.12, “mysql\_find\_rows — Extract SQL Statements from Files”

Section 4.6.14, “mysql\_setpermission — Interactively Set Permissions in Grant Tables”

Section 4.6.15, “mysql\_waitpid — Kill Process and Wait for Its Termination”

Section 4.6.16, “mysql\_zap — Kill Processes That Match a Pattern”

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.6.9, “mysqldumpslow — Summarize Slow Query Log Files”

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.5.5, “mysqlimport — A Data Import Program”

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 4.7.3, “my\_print\_defaults — Display Options from Option Files”

Section 4.7.4, “resolve\_stack\_dump — Resolve Numeric Stack Trace Dump to Symbols”

Section 4.4.1, “comp\_err — Compile MySQL Error Message File”

Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”

Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”

Section 4.8.1, “perror — Explain Error Codes”

---

Section 4.8.3, “`resolveip` — Resolve Host name to IP Address or Vice Versa”  
Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 4.1, “Overview of MySQL Programs”  
Section 5.1.3, “Server Command Options”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 8.11.2, “Tuning Server Parameters”  
Chapter 3, *Tutorial*  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section 4.2.3.3, “Using Option Files”  
Section 4.2.3.1, “Using Options on the Command Line”

### **--hex-blob**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--hexdump**

Section 4.6.8.1, “`mysqlbinlog` Hex Dump Format”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

### **--histignore**

Section 4.5.1.3, “`mysql` Logging”  
Section 4.5.1.1, “`mysql` Options”

### **--host**

Section 4.5.1.1, “`mysql` Options”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 5.1.9.2, “Configuring the MySQL Server to Permit IPv6 Connections”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.6.14, “`mysql_setpermission` — Interactively Set Permissions in Grant Tables”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.2.1, “Invoking MySQL Programs”  
Section 4.2.3.6, “Option Defaults, Options Expecting Values, and the = Sign”  
Section 5.1.3, “Server Command Options”  
Section 1.2, “Typographical and Syntax Conventions”  
Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”  
Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”  
Section 4.2.3.3, “Using Option Files”  
Section 4.2.3.1, “Using Options on the Command Line”

### **host**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.2.3.3, “Using Option Files”

### **--howto**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--html**

Section 4.5.1.1, “`mysql` Options”

|

[index top [3961]]

### **-I**

Section 21.8.3.1, “Building C API Client Programs”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.15, “`mysql_waitpid` — Kill Process and Wait for Its Termination”  
Section 4.6.16, “`mysql_zap` — Kill Processes That Match a Pattern”  
Section 4.8.1, “`perror` — Explain Error Codes”  
Section 4.8.2, “`replace` — A String-Replacement Utility”  
Section 4.8.3, “`resolveip` — Resolve Host name to IP Address or Vice Versa”

### **-i**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

---

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 7.6.2, “How to Check `MyISAM` Tables for Errors”  
Section 22.4.1.3, “Using `pdb` to create a Windows crashdump”

### **--i-am-a-dummy**

Section 4.5.1.1, “`mysql` Options”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”

### **--ignore**

Section 4.5.5, “`mysqlimport` — A Data Import Program”

### **--ignore-builtin-innodb**

Section 14.2.6, “`InnoDB` Startup Options and System Variables”

### **--ignore-db-dir**

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

### **--ignore-lines**

Section 4.5.5, “`mysqlimport` — A Data Import Program”

### **--ignore-spaces**

Section 4.5.1.1, “`mysql` Options”

### **--ignore-table**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--in\_file**

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

### **--include**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

### **--include-gtids**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

### **--include-master-host-port**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--info**

Section 4.8.1, “`pererror` — Explain Error Codes”  
Section 4.8.3, “`resolveip` — Resolve Host name to IP Address or Vice Versa”

### **--Information**

Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”

### **--information**

Section 4.6.3.2, “`myisamchk` Check Options”

### **--init-command**

Section 4.5.1.1, “`mysql` Options”

### **--init-file**

Section 20.2.3, “Performance Schema Runtime Configuration”  
Section C.5.4.1.2, “Resetting the Root Password: Unix Systems”  
Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.4, “The `MEMORY` Storage Engine”

### **--init-rpl-role**

Section 1.4, “What Is New in MySQL 5.6”

### **--init\_connect**

Section 10.1.5, “Configuring the Character Set and Collation for Applications”

### **--innodb**

Section 14.2.6, “`InnoDB` Startup Options and System Variables”  
Section 14.2.2.8, “Turning Off `InnoDB`”

### **innodb-status-file**

Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”

### **--innodb-status-file**

Section 14.2.6, “`InnoDB` Startup Options and System Variables”

### **--innodb-xxx**

Section 5.1.3, “Server Command Options”

### **--innodb\_checksums**

Section 14.2.6, “`InnoDB` Startup Options and System Variables”

### **--innodb\_file\_per\_table**

Section 5.4.1.1, “Enabling and Disabling File-Per-Table Mode”



---

Section 5.1.3, “Server Command Options”  
Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”

## **innodb\_file\_per\_table**

Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 5.1.3, “Server Command Options”

## **--innodb\_rollback\_on\_timeout**

Section 14.2.3.14, “InnoDB Error Handling”  
Section 14.2.6, “InnoDB Startup Options and System Variables”

## **--innodb\_support\_xa**

Section 5.2.4, “The Binary Log”

## **--insert-ignore**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--install**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”  
Section 5.1.3, “Server Command Options”  
Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”  
Section 2.3.4.8, “Starting MySQL as a Windows Service”

## **--install-manual**

Section 5.1.3, “Server Command Options”  
Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”  
Section 2.3.4.8, “Starting MySQL as a Windows Service”

## **--iterations**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **J**

[index top [3961]]

## **-j**

Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

## **--join**

Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”

## **K**

[index top [3961]]

## **-K**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **-k**

Section 4.6.3.3, “myisamchk Repair Options”  
Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”  
Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”

## **--keep\_files\_on\_create**

Section 13.1.14, “CREATE TABLE Syntax”

## **--keepold**

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

## **--key\_buffer\_size**

Section 5.1.3, “Server Command Options”

## **--keys**

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

## **--keys-used**

Section 4.6.3.3, “myisamchk Repair Options”

## **L**

[index top [3961]]

## **-L**

Section 4.5.1.1, “mysql Options”  
Section 21.8.3.1, “Building C API Client Programs”  
Section 4.5.5, “mysqlimport — A Data Import Program”  
Section 2.13.3, “Problems Using the Perl DBI/DBD Interface”  
Section 6.1.6, “Security Issues with LOAD DATA LOCAL”

## **-l**

Section 4.6.3.3, “myisamchk Repair Options”  
Section 21.8.6.39, “mysql\_library\_end()”  
Section 21.8.3.1, “Building C API Client Programs”  
Section 21.8.12, “C API Embedded Server Function Descriptions”  
Section 21.8.5, “C API Function Overview”  
Section 4.6.2, “myisam\_ftdump — Display Full-Text Index information”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”  
Section 4.6.9, “mysqldumpslow — Summarize Slow Query Log Files”

---

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 5.1.3, “Server Command Options”

## **--language**

Section 5.1.3, “Server Command Options”

## **--large-pages**

Section 8.11.4.2, “Enabling Large Page Support”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

## **--lc-messages**

Section 5.1.3, “Server Command Options”

## **--lc-messages-dir**

Section 5.1.3, “Server Command Options”

## **--ldata**

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

## **--ledir**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **--length**

Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”

## **--libmysqld-libs**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **--libs**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **--libs\_r**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

## **--line-numbers**

Section 4.5.1.1, “`mysql` Options”

## **--lines-terminated-by**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

## **--local**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”

## **--local-infile**

Section 4.5.1.1, “`mysql` Options”

Section 13.2.7, “`LOAD XML` Syntax”

Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”

## **--local-load**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **--local-service**

Section 2.3.4.8, “Starting MySQL as a Windows Service”

## **--lock-all-tables**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--lock-tables**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

## **--log**

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 5.1.3, “Server Command Options”

Section 1.4, “What Is New in MySQL 5.6”

## **--log-bin**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

Section 13.4.1.1, “`PURGE BINARY LOGS` Syntax”

Section 7.3.3, “Backup Strategy Summary”

Section 16.1.4.4, “Binary Log Options and Variables”

Section 18.7, “Binary Logging of Stored Programs”

Section 7.2, “Database Backup Methods”

Section 7.3.1, “Establishing a Backup Policy”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.4.5, “How to Report Replication Bugs or Problems”

Section C.5.8, “Known Issues in MySQL”

Section 7.5, “Point-in-Time (Incremental) Recovery Using the Binary Log”

Section 16.1.4.3, “Replication Slave Options and Variables”

---

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 16.3.6, “Switching Masters During Failover”

Section 5.2.4, “The Binary Log”

Section 16.4.4, “Troubleshooting Replication”

Section 16.4.3, “Upgrading a Replication Setup”

Section 7.3.2, “Using Backups for Recovery”

Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

Section 1.4, “What Is New in MySQL 5.6”

### **--log-bin-index**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 5.2.4, “The Binary Log”

### **--log-bin-trust-function-creators**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 18.7, “Binary Logging of Stored Programs”

### **--log-bin-use-v1-row-events**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 2.11.2.1, “Downgrading to MySQL 5.5”

### **--log-error**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 4.2.3.6, “Option Defaults, Options Expecting Values, and the = Sign”

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 5.1.3, “Server Command Options”

Section 5.2.6, “Server Log Maintenance”

Section 2.3.4.5, “Starting MySQL from the Windows Command Line”

Section 2.3.4.4, “Starting the Server for the First Time”

Section 5.2.2, “The Error Log”

### **--log-isam**

Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”

Section 5.1.3, “Server Command Options”

### **--log-output**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

Section 5.1.3, “Server Command Options”

Section 5.2.3, “The General Query Log”

Section 5.2.5, “The Slow Query Log”

### **--log-queries-not-using-indexes**

Section 5.1.3, “Server Command Options”

### **--log-raw**

Section 6.1.2.3, “Passwords and Logging”

Section 5.1.3, “Server Command Options”

Section 5.2.3, “The General Query Log”

### **--log-short-format**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 5.1.3, “Server Command Options”

Section 5.2.5, “The Slow Query Log”

### **--log-slave-updates**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.4.5, “How to Report Replication Bugs or Problems”

Section 16.3.5, “Improving Replication Performance”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.3.6, “Switching Masters During Failover”

Section 5.2.4, “The Binary Log”

Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

### **--log-slow-admin-statements**

Section 5.1.3, “Server Command Options”

Section 5.2.5, “The Slow Query Log”

### **--log-slow-queries**

Section 5.1.3, “Server Command Options”

Section 1.4, “What Is New in MySQL 5.6”

### **--log-slow-slave-statements**

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.2.5, “The Slow Query Log”

### **--log-tc**

Section 5.1.3, “Server Command Options”

### **--log-tc-size**

Section 5.1.3, “Server Command Options”

Section 5.1.6, “Server Status Variables”

### **--log-warnings**

Section C.5.2.9, “MySQL server has gone away”

Section C.5.2.11, “Communication Errors and Aborted Connections”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.3, “Server Command Options”

Section 5.2.2, “The Error Log”

### **--login-path**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”

---

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 4.2.3.3, “Using Option Files”

## **--loose**

Section 4.2.3.2, “Program Option Modifiers”

## **--loose-opt\_name**

Section 4.2.3.3, “Using Option Files”

## **--low-priority**

Section 4.5.5, “`mysqlimport` — A Data Import Program”

## **--low-priority-updates**

Section 13.2.5, “`INSERT` Syntax”

Section 8.10.3, “Concurrent Inserts”

Section 5.1.3, “Server Command Options”

Section 8.10.2, “Table Locking Issues”

## **--lower-case-table-names**

Section 9.2.2, “Identifier Case Sensitivity”

## **M**

[[index top \[3961\]](#)]

### **-m**

Section 4.6.3.2, “`myisamchk` Check Options”

Section 4.5.1.1, “`mysql` Options”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--malloc-lib**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **--master-connect-retry**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-data**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.3.1, “Establishing a Backup Policy”

Section 5.2, “MySQL Server Logs”

Section 16.2.2, “Replication Relay and Status Logs”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## **--master-host**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-info-file**

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.2.2.2, “Slave Status Logs”

## **--master-info-repository**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

Section 16.2.2, “Replication Relay and Status Logs”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.2.2.2, “Slave Status Logs”

Section 1.4, “What Is New in MySQL 5.6”

## **--master-password**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-port**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-retry-count**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl-ca**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl-capath**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl-cert**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl-cipher**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--master-ssl-key**

Section 16.1.4.3, “Replication Slave Options and Variables”

---

### **--master-user**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--master-verify-checksum**

Section 16.1.4.4, “Binary Log Options and Variables”

### **--max**

Section 4.2.3.5, “Using Options to Set Program Variables”

### **--max-binlog-dump-events**

Section 16.1.4.4, “Binary Log Options and Variables”

### **--max-binlog-size**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--max-record-length**

Section 4.6.3.3, “myisamchk Repair Options”  
Section 13.7.2.5, “REPAIR TABLE Syntax”

### **--max-relay-log-size**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--max-seeks-for-key**

Section 8.2.1.4, “How to Avoid Full Table Scans”  
Section C.5.6, “Optimizer-Related Issues”

### **--max\_a**

Section 4.2.3.5, “Using Options to Set Program Variables”

### **--max\_join\_size**

Section 4.5.1.6.4, “Using the `--safe-updates` Option”

### **--maximum**

Section 4.2.3.2, “Program Option Modifiers”

### **--maximum-query\_cache\_size**

Section 4.2.3.2, “Program Option Modifiers”  
Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.5, “Using System Variables”

### **--maximum-var\_name**

Section 5.1.3, “Server Command Options”  
Section 5.1.5, “Using System Variables”

### **--medium-check**

Section 4.6.3.2, “myisamchk Check Options”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

### **--memlock**

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.2.2.3, “Using Raw Disk Partitions for the Shared Tablespace”

### **--method**

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

### **--min-examined-row-limit**

Section 5.1.3, “Server Command Options”

### **--my-plugin**

Section 5.1.8.1, “Installing and Uninstalling Plugins”

### **--my-print-defaults**

Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”

### **--my\_plugin**

Section 5.1.8.1, “Installing and Uninstalling Plugins”

### **--myisam-block-size**

Section 8.9.2.5, “Key Cache Block Size”  
Section 5.1.3, “Server Command Options”

### **--myisam-recover**

Section 5.1.3, “Server Command Options”

### **--myisam-recover-options**

Section 14.3.1, “MyISAM Startup Options”  
Section 8.6.1, “Optimizing MyISAM Queries”  
Section 8.2.5, “Other Optimization Tips”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 7.6.5, “Setting Up a MyISAM Table Maintenance Schedule”  
Section C.5.2.19, “Table-Corruption Issues”  
Section 14.3, “The MyISAM Storage Engine”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in mysqld”

### **--myisam\_sort\_buffer\_size**

Section 4.6.3.6, “myisamchk Memory Usage”

### **MYSQL\_MAINTAINER\_MODE**

Section 2.9.5, “Dealing with Problems Compiling MySQL”

### **MYSQL\_TCP\_PORT**

Section 2.9.3, “Installing MySQL from a Development Source Tree”  
Section 2.9.4, “MySQL Source-Configuration Options”

---

## **MYSQL\_UNIX\_ADDR**

Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”

Section 2.9.3, “Installing MySQL from a Development Source Tree”

Section 2.9.4, “MySQL Source-Configuration Options”

## **--mysqladmin**

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

## **--mysqld**

Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **--mysqld-version**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **N**

[index top [3961]]

### **-N**

Section 4.5.1.1, “`mysql` Options”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

### **-n**

Section 4.6.3.3, “`myisamchk` Repair Options”

Section 4.5.1.1, “`mysql` Options”

Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 4.7.4, “`resolve_stack_dump` — Resolve Numeric Stack Trace Dump to Symbols”

## **--name\_file**

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

## **--named-commands**

Section 4.5.1.1, “`mysql` Options”

## **--ndb**

Section 4.8.1, “`pererror` — Explain Error Codes”

## **net\_retry\_count**

Section 16.2.1, “Replication Implementation Details”

## **net\_write\_timeout**

Section 16.2.1, “Replication Implementation Details”

## **--new**

Section 4.2.3.3, “Using Option Files”

## **--nice**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **--no-auto-rehash**

Section 4.5.1.1, “`mysql` Options”

## **--no-autocommit**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--no-beep**

Section 4.5.1.1, “`mysql` Options”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

## **--no-create-db**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **--no-create-info**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.5.4, “Dumping Table Definitions and Content Separately”

## **--no-data**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.5.4, “Dumping Table Definitions and Content Separately”

## **--no-defaults**

Section 6.2.7, “Causes of Access-Denied Errors”

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”

Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”



---

## **--no-drop**

Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”

## **--no-log**

Section 4.3.4, “[mysqld\\_multi — Manage Multiple MySQL Servers](#)”

## **--no-set-names**

Section 4.5.4, “[mysqldump — A Database Backup Program](#)”

## **--no-symlinks**

Section 4.6.3.3, “[myisamchk Repair Options](#)”

## **--no-tablespaces**

Section 4.5.4, “[mysqldump — A Database Backup Program](#)”

## **--noindices**

Section 4.6.10, “[mysqlhotcopy — A Database Backup Program](#)”

## **--number-char-cols**

Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”

## **--number-int-cols**

Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”

## **--number-of-queries**

Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”

## **--numeric-dump-file**

Section 4.7.4, “[resolve\\_stack\\_dump — Resolve Numeric Stack Trace Dump to Symbols](#)”

## **O**

[[index top \[3961\]](#)]

### **-O**

Section 4.4.1, “[comp\\_err — Compile MySQL Error Message File](#)”

### **-o**

Section 4.6.3.3, “[myisamchk Repair Options](#)”

Section 4.5.1.1, “[mysql Options](#)”

Section 21.7.1, “[Compiling Programs with libmysqld](#)”

Section 4.6.4, “[myisamlog — Display MyISAM Log File Contents](#)”

Section 4.6.8, “[mysqlbinlog — Utility for Processing Binary Log Files](#)”

Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”

Section 8.11.3, “[Optimizing Disk I/O](#)”

## **--offset**

Section 4.6.8, “[mysqlbinlog — Utility for Processing Binary Log Files](#)”

## **--old-alter-table**

Section 5.1.3, “[Server Command Options](#)”

## **--old-style-user-limits**

Section 5.1.3, “[Server Command Options](#)”

Section 6.3.4, “[Setting Account Resource Limits](#)”

## **--old\_server**

Section 4.6.7, “[mysqlaccess — Client for Checking Access Privileges](#)”

Section 4.6.10, “[mysqlhotcopy — A Database Backup Program](#)”

## **ON**

Section 3.3.4.9, “[Using More Than one Table](#)”

## **--one-database**

Section 4.5.1.1, “[mysql Options](#)”

## **--one-thread**

Section 5.1.3, “[Server Command Options](#)”

Section 1.4, “[What Is New in MySQL 5.6](#)”

## **--only-print**

Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”

## **open-files-limit**

Section C.5.2.7, “[Too many connections](#)”

## **--open-files-limit**

Section C.5.2.18, “['File' Not Found and Similar Errors](#)”

Section 14.2.6, “[InnoDB Startup Options and System Variables](#)”

Section 4.3.2, “[mysqld\\_safe — MySQL Server Startup Script](#)”

Section 8.4.3.1, “[How MySQL Opens and Closes Tables](#)”

Section 5.1.3, “[Server Command Options](#)”

## **--opt**

Section 8.5.4, “[Bulk Data Loading for InnoDB Tables](#)”

Section 4.5.4, “[mysqldump — A Database Backup Program](#)”

## **--opt\_name**

Section 4.2.3.3, “[Using Option Files](#)”

## **--optimize**

Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”

---

## **--order-by-primary**

Section 4.5.4, “[mysqldump — A Database Backup Program](#)”

## **--out\_dir**

Section 4.4.1, “[comp\\_err — Compile MySQL Error Message File](#)”

## **--out\_file**

Section 4.4.1, “[comp\\_err — Compile MySQL Error Message File](#)”

## **P**

[[index top \[3961\]](#)]

### **-P**

Section 4.5.1.1, “[mysql Options](#)”  
Section 4.2.2, “[Connecting to the MySQL Server](#)”  
Section 4.6.6, “[mysql\\_config\\_editor — MySQL Configuration Utility](#)”  
Section 4.6.7, “[mysqlaccess — Client for Checking Access Privileges](#)”  
Section 4.6.8, “[mysqlbinlog — Utility for Processing Binary Log Files](#)”  
Section 4.6.10, “[mysqlhotcopy — A Database Backup Program](#)”  
Section 4.5.2, “[mysqladmin — Client for Administering a MySQL Server](#)”  
Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”  
Section 4.5.4, “[mysqldump — A Database Backup Program](#)”  
Section 4.5.5, “[mysqlimport — A Data Import Program](#)”  
Section 4.5.6, “[mysqlshow — Display Database, Table, and Column Information](#)”  
Section 4.5.7, “[mysqslap — Load Emulation Client](#)”  
Section 4.4.4, “[mysql\\_plugin — Configure MySQL Server Plugins](#)”  
Section 4.2.1, “[Invoking MySQL Programs](#)”  
Section 5.1.3, “[Server Command Options](#)”

### **-p**

Section 4.6.3.3, “[myisamchk Repair Options](#)”  
Section 4.5.1.1, “[mysql Options](#)”  
Section 6.3.2, “[Adding User Accounts](#)”  
Section 6.2.7, “[Causes of Access-Denied Errors](#)”  
Section 4.2.2, “[Connecting to the MySQL Server](#)”  
Section 4.6.1, “[innochecksum — Offline InnoDB File Checksum Utility](#)”  
Section 4.6.4, “[myisamlog — Display MyISAM Log File Contents](#)”  
Section 4.6.6, “[mysql\\_config\\_editor — MySQL Configuration Utility](#)”  
Section 4.6.7, “[mysqlaccess — Client for Checking Access Privileges](#)”

Section 4.6.8, “[mysqlbinlog — Utility for Processing Binary Log Files](#)”  
Section 4.6.10, “[mysqlhotcopy — A Database Backup Program](#)”  
Section 4.5.2, “[mysqladmin — Client for Administering a MySQL Server](#)”  
Section 4.5.3, “[mysqlcheck — A Table Maintenance Program](#)”  
Section 4.5.4, “[mysqldump — A Database Backup Program](#)”  
Section 4.5.5, “[mysqlimport — A Data Import Program](#)”  
Section 4.5.6, “[mysqlshow — Display Database, Table, and Column Information](#)”  
Section 4.5.7, “[mysqslap — Load Emulation Client](#)”  
Section 4.4.4, “[mysql\\_plugin — Configure MySQL Server Plugins](#)”  
Section 6.1.2.1, “[End-User Guidelines for Password Security](#)”  
Section 4.2.1, “[Invoking MySQL Programs](#)”  
Section C.5.2.5, “[Password Fails When Entered Interactively](#)”  
Section 2.10.2, “[Securing the Initial MySQL Accounts](#)”  
Section 2.3.4.8, “[Starting MySQL as a Windows Service](#)”  
Section 2.3.4.5, “[Starting MySQL from the Windows Command Line](#)”  
Section 2.3.4.9, “[Testing The MySQL Installation](#)”  
Section 2.3.6, “[Upgrading MySQL on Windows](#)”  
Section 6.3.1, “[User Names and Passwords](#)”  
Section 4.2.3.1, “[Using Options on the Command Line](#)”  
Section 2.3.7, “[Windows Postinstallation Procedures](#)”

## **--pager**

Section 4.5.1.2, “[mysql Commands](#)”  
Section 4.5.1.1, “[mysql Options](#)”

## **--parallel-recover**

Section 4.6.3.3, “[myisamchk Repair Options](#)”

## **--partition**

Section 5.1.3, “[Server Command Options](#)”

## **password**

Section 4.2.3.4, “[Command-Line Options that Affect Option-File Handling](#)”  
Section 4.6.6, “[mysql\\_config\\_editor — MySQL Configuration Utility](#)”  
Section 4.2.3.3, “[Using Option Files](#)”

## **--password**

Section 4.5.1.1, “[mysql Options](#)”  
Section 6.3.2, “[Adding User Accounts](#)”  
Section 6.2.7, “[Causes of Access-Denied Errors](#)”  
Section 4.2.2, “[Connecting to the MySQL Server](#)”  
Section 4.6.6, “[mysql\\_config\\_editor — MySQL Configuration Utility](#)”



Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”

Section 4.6.14, “`mysql_setpermission` — Interactively Set Permissions in Grant Tables”

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 6.1.2.1, “End-User Guidelines for Password Security”

Section 7.3, “Example Backup and Recovery Strategy”

Section 4.2.1, “Invoking MySQL Programs”

Section C.5.2.5, “Password Fails When Entered Interactively”

Section 6.3.7, “Pluggable Authentication”

Section 6.3.7.7, “The Test Authentication Plugin”

Section 6.3.1, “User Names and Passwords”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

Section 4.2.3.1, “Using Options on the Command Line”

**--performance-schema-consumer-consumer\_name**

Section 20.11, “Performance Schema Command Options”

**--performance-schema-instrument**

Section 20.11, “Performance Schema Command Options”

Section 20.2.2, “Performance Schema Startup Configuration”

**--performance-schema-xxx**

Section 5.1.3, “Server Command Options”

--

**performance\_schema\_max\_mutex\_classes**

Section 20.5, “Performance Schema Status Monitoring”

--

**performance\_schema\_max\_mutex\_inst**

Section 20.5, “Performance Schema Status Monitoring”

**pid-file**

Section 2.10.1.2, “Starting and Stopping MySQL Automatically”

**--pid-file**

Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

**--pipe**

Section 4.5.1.1, “`mysql` Options”

Section 4.2.2, “Connecting to the MySQL Server”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 2.3.4.9, “Testing The MySQL Installation”

**--plan**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

**--plugin**

Section 5.1.3, “Server Command Options”

**--plugin-dir**

Section 4.5.1.1, “`mysql` Options”

Section 21.8.13.3, “`mysql_load_plugin()`”

Section 21.8.13, “C API Client Plugin Functions”

Section 22.2.4.2.3, “Client Plugin Descriptors”

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

---

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 6.3.7, “Pluggable Authentication”  
Section E.9, “Restrictions on Pluggable Authentication”  
Section 22.2.4.9.3, “Using the Authentication Plugins”

### **--plugin-ini**

Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”

### **--plugin-**

#### **innodb\_file\_per\_table**

Section 5.1.3, “Server Command Options”

### **--plugin-load**

Section 13.7.3.3, “`INSTALL PLUGIN` Syntax”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 6.3.11.1, “Installing the Audit Log Plugin”  
Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”  
Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 6.1.2.6.1, “Password Validation Plugin Installation”  
Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”  
Section 6.3.7, “Pluggable Authentication”  
Section 22.2.2, “Plugin API Components”  
Section 22.2.4.2, “Plugin Data Structures”  
Section 5.1.3, “Server Command Options”  
Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”  
Section 22.2, “The MySQL Plugin API”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 22.2.4.9.3, “Using the Authentication Plugins”

### **--plugin-load-add**

Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 5.1.3, “Server Command Options”

### **--plugin-sql-mode**

Section 5.1.3, “Server Command Options”

### **--plugin-xxx**

Section 5.1.3, “Server Command Options”

### **--plugin\_dir**

Section 2.9.4, “MySQL Source-Configuration Options”  
Section 22.2.2, “Plugin API Components”

### **--plugin\_name**

Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”

### **--plugindir**

Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”

### **--port**

Section 4.5.1.1, “`mysql` Options”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.6.14, “`mysql_setpermission` — Interactively Set Permissions in Grant Tables”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 4.2.1, “Invoking MySQL Programs”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.6.3, “Running Multiple MySQL Instances on Unix”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”

---

## **--port-open-timeout**

Section 5.1.3, “Server Command Options”

## **--post-query**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--post-system**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--pre-query**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--pre-system**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--preview**

Section 4.6.7, “mysqlaccess — Client for Checking Access Privileges”

## **--print-defaults**

Section 4.2.3.4, “Command-Line Options that Affect Option-File Handling”

Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”

Section 2.11.1, “Upgrading MySQL”

## **--prompt**

Section 4.5.1.2, “mysql Commands”

Section 4.5.1.1, “mysql Options”

## **--protocol**

Section 4.5.1.1, “mysql Options”

Section 4.2.2, “Connecting to the MySQL Server”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.5.5, “mysqlimport — A Data Import Program”

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 5.6.3, “Running Multiple MySQL Instances on Unix”

Section 2.3.4.4, “Starting the Server for the First Time”

Section 2.3.4.9, “Testing The MySQL Installation”

Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”

## **Q**

[index top [3961]]

## **-Q**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **-q**

Section 4.6.3.3, “myisamchk Repair Options”

Section 4.5.1.1, “mysql Options”

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--query**

Section 4.5.7, “mysqlslap — Load Emulation Client”

## **--query-cache-size**

Section 8.10.5, “External Locking”

## **--quick**

Section 4.6.3.6, “myisamchk Memory Usage”

Section 4.6.3.3, “myisamchk Repair Options”

Section 4.5.1.1, “mysql Options”

Section C.5.2.8, “Out of memory”

Section 4.5.1, “mysql — The MySQL Command-Line Tool”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 7.6.1, “Using myisamchk for Crash Recovery”

Section 4.2.3.3, “Using Option Files”

## **--quiet**

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

## **--quote-names**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **R**

[index top [3961]]

## **-R**

Section 7.6.4, “MyISAM Table Optimization”

Section 4.6.4, “myisamlog — Display MyISAM Log File Contents”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.5.4, “mysqldump — A Database Backup Program”

---

Section 4.6.3.4, “Other `myisamchk` Options”  
Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

## **-r**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 22.3.2, “Adding a New User-Defined Function”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 2.2, “Installing MySQL from Generic Binaries on Unix/Linux”  
Section 5.1.3, “Server Command Options”

## **--random-passwords**

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”  
Section 2.5.1, “Installing MySQL from RPM Packages on Linux”  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section 1.4, “What Is New in MySQL 5.6”

## **--raw**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 1.4, “What Is New in MySQL 5.6”

## **--read-from-remote-master**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## **--read-from-remote-server**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 1.4, “What Is New in MySQL 5.6”

## **--read-only**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 16.1.4.3, “Replication Slave Options and Variables”

## **--reconnect**

Section 4.5.1.1, “`mysql` Options”

## **--record\_log\_pos**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## **--recover**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.6.3.6, “`myisamchk` Memory Usage”  
Section 4.6.3.3, “`myisamchk` Repair Options”

## **--regexp**

Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## **--relative**

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

## **--relay-log**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 16.3.5, “Improving Replication Performance”  
Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”  
Section 4.2.3.6, “Option Defaults, Options Expecting Values, and the = Sign”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.2.2.1, “The Slave Relay Log”

## **--relay-log-index**

Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.2.2.1, “The Slave Relay Log”

## **--relay-log-info-file**

Section 16.1.4.3, “Replication Slave Options and Variables”

---

Section 16.2.2.2, “Slave Status Logs”

### **--relay-log-info-repository**

Section 16.2.2, “Replication Relay and Status Logs”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.2.2.2, “Slave Status Logs”

Section 1.4, “What Is New in MySQL 5.6”

### **--relay-log-purge**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--relay-log-recovery**

Section 16.2.2, “Replication Relay and Status Logs”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.12, “The Shutdown Process”

Section 1.4, “What Is New in MySQL 5.6”

### **--relay-log-space-limit**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--relnotes**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--remove**

Section 5.1.3, “Server Command Options”

Section 5.6.2.2, “Starting Multiple MySQL Instances as Windows Services”

Section 2.3.4.8, “Starting MySQL as a Windows Service”

### **--repair**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

### **--replace**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”

### **--replicate-\***

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

Section 16.2.3.3, “Replication Rule Application”

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--replicate-\*-db**

Section 16.2.3.3, “Replication Rule Application”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section E.1, “Restrictions on Stored Programs”

### **--replicate-\*-table**

Section 16.2.3.3, “Replication Rule Application”

### **--replicate-do-db**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 13.3.1, “`START TRANSACTION, COMMIT, and ROLLBACK` Syntax”

Section 16.1.4.4, “Binary Log Options and Variables”

Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

Section 16.3.4, “Replicating Different Databases to Different Slaves”

Section 16.4.1.24, “Replication and Reserved Words”

Section 16.4.1.21, “Replication and Temporary Tables”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.2.4, “The Binary Log”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **--replicate-do-table**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.2.3.2, “Evaluation of Table-Level Replication Options”

Section 16.4.1.24, “Replication and Reserved Words”

Section 16.4.1.21, “Replication and Temporary Tables”

Section 16.2.3.3, “Replication Rule Application”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 14.7, “The `BLACKHOLE` Storage Engine”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **--replicate-ignore-db**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 13.3.1, “`START TRANSACTION, COMMIT, and ROLLBACK` Syntax”

Section 16.1.4.4, “Binary Log Options and Variables”

Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

Section 16.4.1.24, “Replication and Reserved Words”

Section 16.2.3.3, “Replication Rule Application”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.2.4, “The Binary Log”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **--replicate-ignore-table**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”



---

Section 16.2.3.2, “Evaluation of Table-Level Replication Options”

Section 16.4.1.24, “Replication and Reserved Words”

Section 16.4.1.21, “Replication and Temporary Tables”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 14.7, “The `BLACKHOLE` Storage Engine”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **--replicate-rewrite-db**

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **--replicate-same-server-id**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

Section 16.1.4, “Replication and Binary Logging Options and Variables”

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--replicate-wild-do-table**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.2.3.2, “Evaluation of Table-Level Replication Options”

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

Section 16.3.4, “Replicating Different Databases to Different Slaves”

Section 16.4.1.21, “Replication and Temporary Tables”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section E.1, “Restrictions on Stored Programs”

### **--replicate-wild-ignore-table**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.2.3.2, “Evaluation of Table-Level Replication Options”

Section 16.4.1.21, “Replication and Temporary Tables”

Section 16.1.4.3, “Replication Slave Options and Variables”

### **replication-ignore-table**

Section 16.4.1.33, “Replication and Views”

### **--report-host**

Section 13.7.5.34, “`SHOW SLAVE HOSTS` Syntax”

Section 16.1.5.1, “Checking Replication Status”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.4, “Server System Variables”

### **--report-password**

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.4, “Server System Variables”

### **--report-port**

Section 13.7.5.34, “`SHOW SLAVE HOSTS` Syntax”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.4, “Server System Variables”

### **--report-user**

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.4, “Server System Variables”

### **--resetmaster**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

### **--resetslave**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

### **--result-file**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

### **--rhost**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--rollback**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--routines**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.5.3, “Dumping Stored Programs”

Section 7.4.5.4, “Dumping Table Definitions and Content Separately”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

### **--rows**

Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”

### **--rpl-recovery-rank**

Section 1.4, “What Is New in MySQL 5.6”

---

## **--rpm**

Section 4.4.3, “[mysql\\_install\\_db](#) — Initialize MySQL Data Directory”

## **S**

[[index top \[3961\]](#)]

### **-s**

Section 4.5.1.2, “[mysql](#) Commands”

Section 4.5.1.1, “[mysql](#) Options”

Section 7.6.4, “[MyISAM](#) Table Optimization”

Section 4.2.2, “[Connecting to the MySQL Server](#)”

Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”

Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”

Section 4.6.10, “[mysqlhotcopy](#) — A Database Backup Program”

Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”

Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”

Section 4.5.4, “[mysqldump](#) — A Database Backup Program”

Section 4.5.5, “[mysqlimport](#) — A Data Import Program”

Section 4.5.6, “[mysqlshow](#) — Display Database, Table, and Column Information”

Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”

Section 4.4.1, “[comp\\_err](#) — Compile MySQL Error Message File”

Section 4.2.1, “[Invoking MySQL Programs](#)”

Section 4.6.3.4, “[Other myisamchk Options](#)”

### **-s**

Section 4.6.3.1, “[myisamchk](#) General Options”

Section 4.5.1.1, “[mysql](#) Options”

Section 4.6.1, “[innochecksum](#) — Offline InnoDB File Checksum Utility”

Section 4.6.2, “[myisam\\_ftdump](#) — Display Full-Text Index information”

Section 4.6.5, “[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables”

Section 4.6.16, “[mysql\\_zap](#) — Kill Processes That Match a Pattern”

Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”

Section 4.6.9, “[mysqldumpslow](#) — Summarize Slow Query Log Files”

Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”

Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”

Section 4.5.5, “[mysqlimport](#) — A Data Import Program”

Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”

Section 4.7.4, “[resolve\\_stack\\_dump](#) — Resolve Numeric Stack Trace Dump to Symbols”

Section 4.4.7, “[mysql\\_upgrade](#) — Check and Upgrade MySQL Tables”

Section 4.8.1, “[perror](#) — Explain Error Codes”

Section 4.8.2, “[replace](#) — A String-Replacement Utility”

Section 4.8.3, “[resolveip](#) — Resolve Host name to IP Address or Vice Versa”

Section 7.6.2, “[How to Check MyISAM Tables for Errors](#)”

Section 7.6.3, “[How to Repair MyISAM Tables](#)”

Section 5.1.3, “[Server Command Options](#)”

Section 7.6.5, “[Setting Up a MyISAM Table Maintenance Schedule](#)”

## **--safe-mode**

Section 5.1.3, “[Server Command Options](#)”

Section 1.4, “[What Is New in MySQL 5.6](#)”

## **--safe-recover**

Section 4.6.3.1, “[myisamchk](#) General Options”

Section 4.6.3.6, “[myisamchk](#) Memory Usage”

Section 4.6.3.3, “[myisamchk](#) Repair Options”

## **--safe-updates**

Section 4.5.1.2, “[mysql](#) Commands”

Section 4.5.1.1, “[mysql](#) Options”

Section 4.5.1.6.4, “[Using the --safe-updates Option](#)”

## **--safe-user-create**

Section 5.1.3, “[Server Command Options](#)”

## **--secure-auth**

Section 4.5.1.1, “[mysql](#) Options”

Section 6.1.2.4, “[Password Hashing in MySQL](#)”

Section 5.1.3, “[Server Command Options](#)”

## **--secure-file-priv**

Section 6.1.3, “[Making MySQL Secure Against Attackers](#)”

Section 5.1.3, “[Server Command Options](#)”

Section 5.1.4, “[Server System Variables](#)”

Section 2.10.1, “[Unix Postinstallation Procedures](#)”

## **--select\_limit**

Section 4.5.1.6.4, “[Using the --safe-updates Option](#)”

## **server-id**

Section 16.1.1.9, “[Introducing Additional Slaves to an Existing Replication Environment](#)”

Section 16.1, “[Replication Configuration](#)”

Section 16.1.4.2, “[Replication Master Options and Variables](#)”

---

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.1.1.1, “Setting the Replication Master Configuration”

Section 16.1.1.2, “Setting the Replication Slave Configuration”

Section 16.1.1.8, “Setting Up Replication with Existing Data”

### **--server-id**

Section 13.7.5.34, “SHOW SLAVE HOSTS Syntax”

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 16.1.4, “Replication and Binary Logging Options and Variables”

Section 5.1.4, “Server System Variables”

Section 16.4.4, “Troubleshooting Replication”

### **--server-public-key**

Section 4.5.1.1, “mysql Options”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

### **--server-public-key-path**

Section 4.5.1.1, “mysql Options”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

### **--service-startup-timeout**

Section 4.3.3, “mysql.server — MySQL Server Startup Script”

### **--set-auto-increment**

Section 4.6.3.4, “Other myisamchk Options”

### **--set-character-set**

Section 4.6.3.3, “myisamchk Repair Options”

### **--set-charset**

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

Section 4.5.4, “mysqldump — A Database Backup Program”

### **--set-collation**

Section 4.6.3.3, “myisamchk Repair Options”

### **--set-gtid-purged**

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

### **--shared-memory**

Section 4.2.2, “Connecting to the MySQL Server”

Section 5.1.3, “Server Command Options”

Section 5.6.2.1, “Starting Multiple MySQL Instances at the Windows Command Line”

Section 2.3.4.4, “Starting the Server for the First Time”

### **--shared-memory-base-name**

Section 21.8.6.49, “mysql\_options()”

Section 4.2.2, “Connecting to the MySQL Server”

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 5.6, “Running Multiple MySQL Instances on One Machine”

Section 5.1.3, “Server Command Options”

Section 5.6.2.1, “Starting Multiple MySQL Instances at the Windows Command Line”

Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”

### **--short-form**

Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

### **--show-slave-auth-info**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--show-table-type**

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

### **--show-warnings**

Section 4.5.1.1, “mysql Options”

### **--sigint-ignore**

Section 4.5.1.1, “mysql Options”

### **--silent**

Section 4.6.3.1, “myisamchk General Options”

Section 4.5.1.1, “mysql Options”

Section 4.6.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 4.5.5, “mysqlimport — A Data Import Program”

Section 4.5.7, “mysqlslap — Load Emulation Client”

Section 4.8.1, “perror — Explain Error Codes”

Section 4.8.3, “resolveip — Resolve Host name to IP Address or Vice Versa”

Section 4.3.4, “mysqld\_multi — Manage Multiple MySQL Servers”

Section 7.6.5, “Setting Up a MyISAM Table Maintenance Schedule”

### **--single-transaction**

Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”



---

Section 7.2, “Database Backup Methods”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.3.1, “Establishing a Backup Policy”

### **--skip**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.2.3.2, “Program Option Modifiers”  
Section 5.1.3, “Server Command Options”

### **--skip-add-drop-table**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-add-locks**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-auto-rehash**

Section 14.2.4.7, “Troubleshooting InnoDB Data Dictionary Operations”

### **--skip-character-set-client-handshake**

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 10.1.14.7.1, “The `cp932` Character Set”

### **--skip-column-names**

Section 4.5.1.1, “`mysql` Options”

### **--skip-comments**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-concurrent-insert**

Section 5.1.3, “Server Command Options”

### **--skip-database**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

### **--skip-disable-keys**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-dump-date**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-engine\_name**

Section 13.7.5.17, “`SHOW ENGINES` Syntax”

### **--skip-event-scheduler**

Section 5.1.3, “Server Command Options”

### **--skip-events**

Section 7.4.5.3, “Dumping Stored Programs”

### **--skip-extended-insert**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--skip-external-locking**

Section 8.10.5, “External Locking”  
Section 8.12.5.2, “General Thread States”  
Section 5.1.3, “Server Command Options”  
Section 8.11.1, “System Factors and Startup Parameter Tuning”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”

### **--skip-grant-tables**

Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”  
Section 13.7.3.3, “`INSTALL PLUGIN` Syntax”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”  
Section 18.4.2, “Event Scheduler Configuration”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 6.3.7, “Pluggable Authentication”  
Section 2.10.1.1, “Problems Running `mysql_install_db`”  
Section C.5.4.1.3, “Resetting the Root Password: Generic Instructions”  
Section 5.1.3, “Server Command Options”  
Section 4.2.3.1, “Using Options on the Command Line”  
Section 6.2.6, “When Privilege Changes Take Effect”

### **--skip-gtids**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

### **--skip-host-cache**

Section 6.2.7, “Causes of Access-Denied Errors”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

### **--skip-innodb**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

---

Section 14.2.2.8, “Turning Off InnoDB”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”

### **--skip-innodb-checksums**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **--skip-innodb\_adaptive\_hash\_index**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **--skip-innodb\_checksums**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **--skip-innodb\_doublewrite**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **--skip-kill-mysqld**

Section 4.3.2, “mysqld\_safe — MySQL Server Startup Script”

### **--skip-line-numbers**

Section 4.5.1.1, “mysql Options”

### **--skip-lock-tables**

Section 4.5.4, “mysqldump — A Database Backup Program”

### **--skip-log-warnings**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **--skip-name-resolve**

Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”  
Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 2.3.4.9, “Testing The MySQL Installation”

### **--skip-named-commands**

Section 4.5.1.1, “mysql Options”

### **--skip-networking**

Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section C.5.2.9, “MySQL server has gone away”

Section 6.2.7, “Causes of Access-Denied Errors”  
Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”  
Section 6.3.7, “Pluggable Authentication”  
Section C.5.4.1.3, “Resetting the Root Password: Generic Instructions”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 16.4.3, “Upgrading a Replication Setup”

### **skip-networking**

Section 16.1.1.1, “Setting the Replication Master Configuration”  
Section 16.4.4, “Troubleshooting Replication”

### **--skip-new**

Section 13.7.2.4, “OPTIMIZE TABLE Syntax”  
Section 22.4.1, “Debugging a MySQL Server”  
Section 5.1.4, “Server System Variables”

### **--skip-opt**

Section 4.5.4, “mysqldump — A Database Backup Program”

### **--skip-pager**

Section 4.5.1.1, “mysql Options”

### **--skip-partition**

Chapter 17, *Partitioning*  
Section 5.1.3, “Server Command Options”

### **--skip-plugin-innodb\_file\_per\_table**

Section 5.1.3, “Server Command Options”

### **--skip-plugin\_name**

Section 5.1.8.1, “Installing and Uninstalling Plugins”

### **--skip-quick**

Section 4.5.4, “mysqldump — A Database Backup Program”

### **--skip-quote-names**

Section 4.5.4, “mysqldump — A Database Backup Program”

### **--skip-reconnect**

Section 4.5.1.1, “mysql Options”  
Section 21.8.15, “Controlling Automatic Reconnection Behavior”  
Section 4.5.1.6.5, “Disabling mysql Auto-Reconnect”

### **--skip-routines**

Section 7.4.5.3, “Dumping Stored Programs”

---

## **--skip-secure-auth**

Section 4.5.1.1, “mysql Options”  
Section 5.1.3, “Server Command Options”

## **--skip-set-charset**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--skip-show-database**

Section 13.7.5.15, “SHOW DATABASES Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.1.3, “Server Command Options”

## **--skip-slave-start**

Section 13.4.2.1, “CHANGE MASTER TO Syntax”  
Section 13.4.2.5, “START SLAVE Syntax”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.1.3.2, “Setting Up Replication Using GTIDs”  
Section 16.3.7, “Setting Up Replication Using SSL”  
Section 16.1.1.8, “Setting Up Replication with Existing Data”  
Section 16.4.4, “Troubleshooting Replication”  
Section 16.4.3, “Upgrading a Replication Setup”

## **--skip-ssl**

Section 6.3.9.4, “SSL Command Options”

## **--skip-stack-trace**

Section 22.4.1.4, “Debugging mysqld under gdb”  
Section 5.1.3, “Server Command Options”

## **--skip-super-large-pages**

Section 8.11.4.2, “Enabling Large Page Support”  
Section 5.1.3, “Server Command Options”

## **--skip-symbolic-links**

Section 13.1.14, “CREATE TABLE Syntax”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”  
Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”

## **--skip-syslog**

Section 4.3.2, “mysqld\_safe — MySQL Server Startup Script”  
Section 5.2.2, “The Error Log”

## **--skip-thread-priority**

Section 5.1.3, “Server Command Options”

Section 1.4, “What Is New in MySQL 5.6”

## **--skip-triggers**

Section 4.5.4, “mysqldump — A Database Backup Program”  
Section 7.4.5.3, “Dumping Stored Programs”

## **--skip-tz-utc**

Section 4.5.4, “mysqldump — A Database Backup Program”

## **--skip-use-db**

Section 4.6.12, “mysql\_find\_rows — Extract SQL Statements from Files”

## **--skip-version-check**

Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”

## **--skip-warn**

Section 4.6.6, “mysql\_config\_editor — MySQL Configuration Utility”

## **--skip-write-binlog**

Section 11.3.4, “YEAR(2) Limitations and Migrating to YEAR(4)”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 4.4.7, “mysql\_upgrade — Check and Upgrade MySQL Tables”

## **--skip\_grant\_tables**

Section 4.2.3.1, “Using Options on the Command Line”

## **--slave-checkpoint-group**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--slave-checkpoint-period**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--slave-load-tmpdir**

Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 7.2, “Database Backup Methods”  
Section 16.1.4.3, “Replication Slave Options and Variables”

## **slave-max-allowed-packet**

Section 16.1.4.2, “Replication Master Options and Variables”

## **--slave-max-allowed-packet**

Section 16.1.4.2, “Replication Master Options and Variables”

---

## **--slave-net-timeout**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--slave-parallel-workers**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--slave-pending-jobs-size-max**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **slave-rows-search-algorithms**

Section 16.1.4.2, “Replication Master Options and Variables”

## **--slave-skip-errors**

Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.4.1.25, “Slave Errors During Replication”

## **--slave-sql-verify-checksum**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”

## **--slave\_compressed\_protocol**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **--sleep**

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

## **--slow-query-log**

Section 5.1.3, “Server Command Options”

## **--slow-start-timeout**

Section 5.1.3, “Server Command Options”

## **--slow\_query\_log**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.2.5, “The Slow Query Log”  
Section 1.4, “What Is New in MySQL 5.6”

## **--slow\_query\_log\_file**

Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.1.3, “Server Command Options”  
Section 5.2.5, “The Slow Query Log”  
Section 1.4, “What Is New in MySQL 5.6”

## **--socket**

Section 4.5.1.1, “`mysql` Options”  
Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.6.14, “`mysql_setpermission` — Interactively Set Permissions in Grant Tables”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”  
Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section C.5.4.5, “How to Protect or Change the MySQL Unix Socket File”  
Section 4.2.1, “Invoking MySQL Programs”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.6.3, “Running Multiple MySQL Instances on Unix”  
Section 5.1.3, “Server Command Options”  
Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”  
Section 2.3.4.9, “Testing The MySQL Installation”  
Section 5.6.4, “Using Client Programs in a Multiple-Server Environment”

## **socket**

Section 2.4.1, “General Notes on Installing MySQL on Mac OS X”

## **--sort-index**

Section 7.6.4, “`MyISAM` Table Optimization”  
Section 4.6.3.4, “Other `myisamchk` Options”

## **--sort-records**

Section 7.6.4, “`MyISAM` Table Optimization”  
Section 4.6.3.4, “Other `myisamchk` Options”

---

## **--sort-recover**

Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.6.3.6, “`myisamchk` Memory Usage”  
Section 4.6.3.3, “`myisamchk` Repair Options”

## **--spassword**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

## **--sporadic-binlog-dump-fail**

Section 16.1.4.4, “Binary Log Options and Variables”

## **--sql-mode**

Chapter 12, *Functions and Operators*  
Section 1.8.2, “Selecting SQL Modes”  
Section 5.1.3, “Server Command Options”  
Section 5.1.7, “Server SQL Modes”

## **sql-mode**

Section 5.1.7, “Server SQL Modes”

## **--srcdir**

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

## **--ssl**

Section 4.5.1.1, “`mysql` Options”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 5.1.3, “Server Command Options”  
Section 6.3.9.4, “SSL Command Options”

## **--ssl\***

Section 4.5.1.1, “`mysql` Options”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 5.1.3, “Server Command Options”

## **--ssl-ca**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”  
Section 6.3.9.4, “SSL Command Options”  
Section 6.3.9.3, “Using SSL Connections”

## **--ssl-capath**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.3.9.4, “SSL Command Options”

## **--ssl-cert**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”  
Section 6.3.9.4, “SSL Command Options”  
Section 6.3.9.3, “Using SSL Connections”

## **--ssl-cipher**

Section 6.3.9.4, “SSL Command Options”

## **--ssl-crl**

Section 6.3.9.4, “SSL Command Options”

## **--ssl-crlpath**

Section 6.3.9.4, “SSL Command Options”

## **--ssl-key**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.3.9.5, “Setting Up SSL Certificates and Keys for MySQL”  
Section 6.3.9.4, “SSL Command Options”  
Section 6.3.9.3, “Using SSL Connections”

## **--ssl-verify-server-cert**

Section 6.3.9.4, “SSL Command Options”

## **--ssl-xxx**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 6.3.9.2, “Configuring MySQL for SSL”  
Section 5.1.4, “Server System Variables”  
Section 6.3.9.4, “SSL Command Options”

## **--standalone**

Section 22.4.1.2, “Creating Trace Files”  
Section 5.1.3, “Server Command Options”  
Section 2.3.4.5, “Starting MySQL from the Windows Command Line”

## **--start-datetime**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”



---

Section 7.5.1, “Point-in-Time Recovery Using Event Times”

### **--start-position**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 7.5.2, “Point-in-Time Recovery Using Event Positions”

### **--start\_row**

Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”

### **--statefile**

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

### **--stats**

Section 4.6.2, “`myisam_ftdump` — Display Full-Text Index information”

### **--status**

Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”

### **--stop-datetime**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 7.5.1, “Point-in-Time Recovery Using Event Times”

### **--stop-never**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

### **--stop-never-slave-server-id**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”

Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

### **--stop-position**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 7.5.2, “Point-in-Time Recovery Using Event Positions”

### **--suffix**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

### **--super-large-pages**

Section 8.11.4.2, “Enabling Large Page Support”  
Section 5.1.3, “Server Command Options”

### **--superuser**

Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--symbolic-links**

Section 5.1.3, “Server Command Options”

### **--symbols-file**

Section 4.7.4, “`resolve_stack_dump` — Resolve Numeric Stack Trace Dump to Symbols”

## **SYSCONFDIR**

Section 4.2.3.3, “Using Option Files”

### **--sysdate-is-now**

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 12.7, “Date and Time Functions”

Section 16.4.1.14, “Replication and System Functions”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

### **--syslog**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

Section 5.2.2, “The Error Log”

### **--syslog-tag**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”

## **T**

[[index top \[3961\]](#)]

### **-T**

Section 4.6.3.2, “`myisamchk` Check Options”

Section 4.5.1.1, “`mysql` Options”

Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”

Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

Section 5.1.3, “Server Command Options”

### **-t**

Section 4.6.3.3, “`myisamchk` Repair Options”

---

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.5, “`mysampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.16, “`mysqlzap` — Kill Processes That Match a Pattern”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.9, “`mysqldumpslow` — Summarize Slow Query Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 5.1.3, “Server Command Options”

### **--tab**

Section 7.1, “Backup and Recovery Types”  
Section 7.2, “Database Backup Methods”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”  
Section 7.4, “Using `mysqldump` for Backups”

### **--table**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **--table-cache**

Section 1.4, “What Is New in MySQL 5.6”

### **--table\_open\_cache**

Section 8.4.3.1, “How MySQL Opens and Closes Tables”

### **--tables**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--tc-heuristic-recover**

Section 5.1.3, “Server Command Options”

### **--tcp-ip**

Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”

### **--tee**

Section 4.5.1.2, “`mysql` Commands”  
Section 4.5.1.1, “`mysql` Options”

### **--temp-pool**

Section 5.1.3, “Server Command Options”

### **--test**

Section 4.6.5, “`mysampack` — Generate Compressed, Read-Only MyISAM Tables”

### **Text**

Section 1.2, “Typographical and Syntax Conventions”

### **--thread\_cache\_size**

Section 22.4.1.4, “Debugging `mysqld` under `gdb`”

### **--thread\_handling**

Section 5.1.3, “Server Command Options”  
Section 1.4, “What Is New in MySQL 5.6”

### **--thread\_stack**

Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”

### **--timezone**

Section 4.3.2, “`mysqld_safe` — MySQL Server Startup Script”  
Section 10.6, “MySQL Server Time Zone Support”  
Section 16.4.1.28, “Replication and Time Zones”  
Section 5.1.4, “Server System Variables”  
Section C.5.4.6, “Time Zone Problems”

### **tmpdir**

Section 2.3, “Installing MySQL on Microsoft Windows”

### **--tmpdir**

Section 4.6.3.6, “`myisamchk` Memory Usage”  
Section 4.6.3.3, “`myisamchk` Repair Options”  
Section C.5.2.13, “Can't create/write to file”  
Section 4.6.5, “`mysampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 5.1.3, “Server Command Options”  
Section C.5.4.4, “Where MySQL Stores Temporary Files”

### **--to-last-log**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.8.4, “Specifying the `mysqlbinlog` Server ID”  
Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

---

## **--transaction-isolation**

Section 13.3.6, “`SET TRANSACTION` Syntax”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.2.3.1, “The InnoDB Transaction Model and Locking”

## **--transaction-read-only**

Section 13.3.6, “`SET TRANSACTION` Syntax”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **--triggers**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.5.3, “Dumping Stored Programs”

## **--type**

Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”

## **--tz-utc**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## **U**

[index top [3961]]

### **-U**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”

### **-u**

Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 4.2.1, “Invoking MySQL Programs”  
Section 5.1.3, “Server Command Options”  
Section 2.3.4.9, “Testing The MySQL Installation”  
Section 6.3.1, “User Names and Passwords”  
Section 2.3.7, “Windows Postinstallation Procedures”

## **--unbuffered**

Section 4.5.1.1, “`mysql` Options”

## **--unpack**

Section 4.6.3.3, “`myisamchk` Repair Options”  
Section 14.3.3, “MyISAM Table Storage Formats”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”

## **--update-state**

Section 4.6.3.2, “`myisamchk` Check Options”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 14.3, “The MyISAM Storage Engine”

## **--upgrade-system-tables**

Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

## **--use-frm**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

## **--use-mysqld\_safe**

Section 4.3.3, “`mysql.server` — MySQL Server Startup Script”

## **--use-threads**

Section 4.5.5, “`mysqlimport` — A Data Import Program”

## **--user**

Section 4.5.1.1, “`mysql` Options”  
Section C.5.2.18, “`'File' Not Found` and Similar Errors”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.6.11, “`mysql_convert_table_format` — Convert Tables to Use a Given Storage Engine”  
Section 4.6.14, “`mysql_setpermission` — Interactively Set Permissions in Grant Tables”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”



---

Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”  
Section 4.6.10, “[mysqlhotcopy](#) — A Database Backup Program”  
Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”  
Section 4.5.4, “[mysqldump](#) — A Database Backup Program”  
Section 4.5.5, “[mysqlimport](#) — A Data Import Program”  
Section 4.5.6, “[mysqlshow](#) — Display Database, Table, and Column Information”  
Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”  
Section 4.4.3, “[mysql\\_install\\_db](#) — Initialize MySQL Data Directory”  
Section 4.4.7, “[mysql\\_upgrade](#) — Check and Upgrade MySQL Tables”  
Section 4.3.3, “[mysql.server](#) — MySQL Server Startup Script”  
Section 4.3.4, “[mysqld\\_multi](#) — Manage Multiple MySQL Servers”  
Section 4.3.2, “[mysqld\\_safe](#) — MySQL Server Startup Script”  
Section 7.3, “[Example Backup and Recovery Strategy](#)”  
Section 6.1.5, “[How to Run MySQL as a Normal User](#)”  
Section 4.2.1, “[Invoking MySQL Programs](#)”  
Section 6.1.3, “[Making MySQL Secure Against Attackers](#)”  
Section 4.2.3.6, “[Option Defaults, Options Expecting Values, and the = Sign](#)”  
Section 6.3.7, “[Pluggable Authentication](#)”  
Section C.5.4.1.2, “[Resetting the Root Password: Unix Systems](#)”  
Section 5.1.3, “[Server Command Options](#)”  
Section 6.3.7.6, “[The Socket Peer-Credential Authentication Plugin](#)”  
Section 6.3.7.7, “[The Test Authentication Plugin](#)”  
Section 2.10.1, “[Unix Postinstallation Procedures](#)”  
Section 6.3.1, “[User Names and Passwords](#)”  
Section 4.6.8.3, “[Using `mysqlbinlog` to Back Up Binary Log Files](#)”  
Section 4.2.3.3, “[Using Option Files](#)”

## user

Section 4.2.3.4, “[Command-Line Options that Affect Option-File Handling](#)”  
Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”  
Section 4.2.3.3, “[Using Option Files](#)”

## V

[[index top \[3961\]](#)]

## -V

Section 4.6.3.1, “[myisamchk](#) General Options”

Section 4.5.1.1, “[mysql](#) Options”  
Section 4.6.4, “[myisamlog](#) — Display MyISAM Log File Contents”  
Section 4.6.5, “[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”  
Section 4.6.15, “[mysql\\_waitpid](#) — Kill Process and Wait for Its Termination”  
Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”  
Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”  
Section 4.5.4, “[mysqldump](#) — A Database Backup Program”  
Section 4.5.5, “[mysqlimport](#) — A Data Import Program”  
Section 4.5.6, “[mysqlshow](#) — Display Database, Table, and Column Information”  
Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”  
Section 4.7.3, “[my\\_print\\_defaults](#) — Display Options from Option Files”  
Section 4.7.4, “[resolve\\_stack\\_dump](#) — Resolve Numeric Stack Trace Dump to Symbols”  
Section 4.4.1, “[comp\\_err](#) — Compile MySQL Error Message File”  
Section 4.4.4, “[mysql\\_plugin](#) — Configure MySQL Server Plugins”  
Section 4.8.1, “[perror](#) — Explain Error Codes”  
Section 4.8.2, “[replace](#) — A String-Replacement Utility”  
Section 4.8.3, “[resolveip](#) — Resolve Host name to IP Address or Vice Versa”  
Section 5.1.3, “[Server Command Options](#)”  
Section 4.2.3.1, “[Using Options on the Command Line](#)”

## -v

Section 4.6.3.1, “[myisamchk](#) General Options”  
Section 4.5.1.1, “[mysql](#) Options”  
Section 4.6.8.2, “[mysqlbinlog](#) Row Event Display”  
Section 4.6.1, “[innochecksum](#) — Offline InnoDB File Checksum Utility”  
Section 4.6.2, “[myisam\\_ftdump](#) — Display Full-Text Index information”  
Section 4.6.4, “[myisamlog](#) — Display MyISAM Log File Contents”  
Section 4.6.5, “[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”  
Section 4.6.15, “[mysql\\_waitpid](#) — Kill Process and Wait for Its Termination”  
Section 4.6.7, “[mysqlaccess](#) — Client for Checking Access Privileges”  
Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”

---

Section 4.6.9, “[mysqldumpslow](#) — Summarize Slow Query Log Files”  
Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”  
Section 4.5.4, “[mysqldump](#) — A Database Backup Program”  
Section 4.5.5, “[mysqlimport](#) — A Data Import Program”  
Section 4.5.6, “[mysqlshow](#) — Display Database, Table, and Column Information”  
Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”  
Section 4.7.3, “[my\\_print\\_defaults](#) — Display Options from Option Files”  
Section 4.4.4, “[mysql\\_plugin](#) — Configure MySQL Server Plugins”  
Section 4.8.1, “[perror](#) — Explain Error Codes”  
Section 4.8.2, “[replace](#) — A String-Replacement Utility”  
Section 7.6.2, “How to Check [MyISAM](#) Tables for Errors”  
Section 4.6.3.5, “Obtaining Table Information with [myisamchk](#)”  
Section 5.1.3, “Server Command Options”  
Section 4.2.3.1, “Using Options on the Command Line”

## **--validate-password**

Section 6.1.2.6.1, “Password Validation Plugin Installation”  
Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **--var\_name**

Section 4.6.3.1, “[myisamchk](#) General Options”  
Section 4.5.1.1, “[mysql](#) Options”  
Section 14.2.6, “[InnoDB](#) Startup Options and System Variables”  
Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”  
Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 4.5.4, “[mysqldump](#) — A Database Backup Program”  
Section 5.1.3, “Server Command Options”

## **--verbose**

Section 4.6.3.1, “[myisamchk](#) General Options”  
Section 4.5.1.1, “[mysql](#) Options”  
Section 4.6.8.2, “[mysqlbinlog](#) Row Event Display”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 4.6.2, “[myisam\\_ftdump](#) — Display Full-Text Index information”  
Section 4.6.5, “[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables”

Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”  
Section 4.6.11, “[mysql\\_convert\\_table\\_format](#) — Convert Tables to Use a Given Storage Engine”  
Section 4.6.15, “[mysql\\_waitpid](#) — Kill Process and Wait for Its Termination”  
Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”  
Section 4.6.9, “[mysqldumpslow](#) — Summarize Slow Query Log Files”  
Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 4.5.3, “[mysqlcheck](#) — A Table Maintenance Program”  
Section 4.5.4, “[mysqldump](#) — A Database Backup Program”  
Section 4.5.5, “[mysqlimport](#) — A Data Import Program”  
Section 4.5.6, “[mysqlshow](#) — Display Database, Table, and Column Information”  
Section 4.5.7, “[mysqlslap](#) — Load Emulation Client”  
Section 4.7.3, “[my\\_print\\_defaults](#) — Display Options from Option Files”  
Section 4.4.3, “[mysql\\_install\\_db](#) — Initialize MySQL Data Directory”  
Section 4.4.4, “[mysql\\_plugin](#) — Configure MySQL Server Plugins”  
Section 4.4.7, “[mysql\\_upgrade](#) — Check and Upgrade MySQL Tables”  
Section 4.8.1, “[perror](#) — Explain Error Codes”  
Section 4.3.4, “[mysqld\\_multi](#) — Manage Multiple MySQL Servers”  
Section 4.5.1.5, “Executing SQL Statements from a Text File”  
Section 4.6.3.4, “Other [myisamchk](#) Options”  
Section 5.1.3, “Server Command Options”  
Section 2.10.1.3, “Starting and Troubleshooting the MySQL Server”  
Section 8.11.2, “Tuning Server Parameters”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”  
Section 4.2.3.3, “Using Option Files”  
Section 4.2.3.1, “Using Options on the Command Line”

## **--verify-binlog-checksum**

Section 4.6.8, “[mysqlbinlog](#) — Utility for Processing Binary Log Files”

## **--version**

Section 4.6.3.1, “[myisamchk](#) General Options”  
Section 4.5.1.1, “[mysql](#) Options”  
Section 4.6.5, “[myisampack](#) — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.6, “[mysql\\_config\\_editor](#) — MySQL Configuration Utility”  
Section 4.6.11, “[mysql\\_convert\\_table\\_format](#) — Convert Tables to Use a Given Storage Engine”

---

Section 4.6.15, “`mysql_waitpid` — Kill Process and Wait for Its Termination”  
Section 4.6.7, “`mysqlaccess` — Client for Checking Access Privileges”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 4.7.3, “`my_print_defaults` — Display Options from Option Files”  
Section 4.7.2, “`mysql_config` — Display Options for Compiling Clients”  
Section 4.7.4, “`resolve_stack_dump` — Resolve Numeric Stack Trace Dump to Symbols”  
Section 4.4.1, “`comp_err` — Compile MySQL Error Message File”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 4.8.1, “`perror` — Explain Error Codes”  
Section 4.8.3, “`resolveip` — Resolve Host name to IP Address or Vice Versa”  
Section 4.3.4, “`mysqld_multi` — Manage Multiple MySQL Servers”  
Section 5.1.3, “Server Command Options”  
Section 4.2.3.1, “Using Options on the Command Line”

## **--version-check**

Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

## **--vertical**

Section 4.5.1.1, “`mysql` Options”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 1.7, “How to Report Bugs or Problems”

## **W**

[index top [3961]]

### **-W**

Section 4.5.1.1, “`mysql` Options”  
Section 4.2.2, “Connecting to the MySQL Server”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 4.5.6, “`mysqlshow` — Display Database, Table, and Column Information”  
Section 4.5.7, “`mysqlslap` — Load Emulation Client”  
Section 5.1.3, “Server Command Options”

### **-w**

Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.4, “`myisamlog` — Display MyISAM Log File Contents”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--wait**

Section 4.6.3.1, “`myisamchk` General Options”  
Section 4.5.1.1, “`mysql` Options”  
Section 4.6.5, “`myisampack` — Generate Compressed, Read-Only MyISAM Tables”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

### **--warn**

Section 4.6.6, “`mysql_config_editor` — MySQL Configuration Utility”

### **--where**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **--windows**

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

### **--with-mysql-sock**

Section 21.9.4, “MySQL Functions (PDO\_MYSQL) (MySQL (PDO))”

### **--with-pdo-mysql**

Section 21.9.4, “MySQL Functions (PDO\_MYSQL) (MySQL (PDO))”

### **--with-zlib-dir**

Section 21.9.4, “MySQL Functions (PDO\_MYSQL) (MySQL (PDO))”

### **WITH\_DEBUG**

Section 2.9.4, “MySQL Source-Configuration Options”

---

## **WITH\_PERFSCHEMA\_STORAGE\_ENGINE**

Section 20.2.1, “Performance Schema Build Configuration”

## **WITH\_ZLIB**

Section 2.9.4, “MySQL Source-Configuration Options”

## **--write-binlog**

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

## **X**

[index top [3961]]

### **-x**

Section 4.5.1.2, “`mysql` Commands”

Section 4.5.1.1, “`mysql` Options”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **-x**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

## **--xml**

Section 4.5.1.1, “`mysql` Options”

Section 13.2.7, “LOAD XML Syntax”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 12.11, “XML Functions”

## **Y**

[index top [3961]]

### **-Y**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

### **-y**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 4.5.7, “`mysqlslap` — Load Emulation Client”

---

# Privileges Index

[A](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [I](#) | [L](#) | [P](#) | [R](#) | [S](#) | [T](#) | [U](#)

## A

[\[index top \[4005\]\]](#)

### ALL

Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”

### ALL PRIVILEGES

Section 6.2.1, “[Privileges Provided by MySQL](#)”

### ALTER

Section 13.1.1, “[ALTER DATABASE Syntax](#)”  
Section 13.1.6, “[ALTER TABLE Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 13.1.26, “[RENAME TABLE Syntax](#)”  
Section 17.3.3, “[Exchanging Partitions and Subpartitions with Tables](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”  
Section 19.26, “[The INFORMATION\\_SCHEMA TABLE\\_PRIVILEGES Table](#)”

### ALTER ROUTINE

Section 13.1.3, “[ALTER FUNCTION Syntax](#)”  
Section 13.1.4, “[ALTER PROCEDURE Syntax](#)”  
Section 13.1.12, “[CREATE PROCEDURE and CREATE FUNCTION Syntax](#)”  
Section 13.1.21, “[DROP PROCEDURE and DROP FUNCTION Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 18.7, “[Binary Logging of Stored Programs](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”  
Section 5.1.4, “[Server System Variables](#)”  
Section 18.2.2, “[Stored Routines and MySQL Privileges](#)”

## C

[\[index top \[4005\]\]](#)

### CREATE

Section 13.1.6, “[ALTER TABLE Syntax](#)”  
Section 13.1.8, “[CREATE DATABASE Syntax](#)”  
Section 13.1.14, “[CREATE TABLE Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 13.1.26, “[RENAME TABLE Syntax](#)”  
Section 17.3.3, “[Exchanging Partitions and Subpartitions with Tables](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”

### CREATE ROUTINE

Section 13.1.12, “[CREATE PROCEDURE and CREATE FUNCTION Syntax](#)”

Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 18.7, “[Binary Logging of Stored Programs](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”  
Section 5.1.4, “[Server System Variables](#)”  
Section 18.2.2, “[Stored Routines and MySQL Privileges](#)”

### CREATE TABLESPACE

Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”

### CREATE TEMPORARY TABLES

Section 13.1.14, “[CREATE TABLE Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”

### CREATE USER

Section 13.7.1.1, “[ALTER USER Syntax](#)”  
Section 13.7.1.2, “[CREATE USER Syntax](#)”  
Section 13.7.1.3, “[DROP USER Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 13.7.1.5, “[RENAME USER Syntax](#)”  
Section 13.7.1.6, “[REVOKE Syntax](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”

### CREATE VIEW

Section 13.1.7, “[ALTER VIEW Syntax](#)”  
Section 13.1.16, “[CREATE VIEW Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”  
Section E.5, “[Restrictions on Views](#)”  
Section 19.26, “[The INFORMATION\\_SCHEMA TABLE\\_PRIVILEGES Table](#)”

## D

[\[index top \[4005\]\]](#)

### DELETE

Section 13.1.14, “[CREATE TABLE Syntax](#)”  
Section 13.2.2, “[DELETE Syntax](#)”  
Section 13.7.3.2, “[DROP FUNCTION Syntax](#)”  
Section 13.7.1.3, “[DROP USER Syntax](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 13.2.8, “[REPLACE Syntax](#)”  
Section 13.7.3.4, “[UNINSTALL PLUGIN Syntax](#)”  
Section 6.2.5, “[Access Control, Stage 2: Request Verification](#)”  
Section 22.3.2.5, “[Compiling and Installing User-Defined Functions](#)”  
Section 5.1.8.1, “[Installing and Uninstalling Plugins](#)”  
Section 6.2.1, “[Privileges Provided by MySQL](#)”  
Section 14.8, “[The MERGE Storage Engine](#)”  
Section 20.9.1.4, “[The setup\\_objects Table](#)”  
Section 22.3.2.6, “[User-Defined Function Security Precautions](#)”

---

## DROP

Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.1.7, “ALTER VIEW Syntax”  
Section 13.1.16, “CREATE VIEW Syntax”  
Section 13.1.17, “DROP DATABASE Syntax”  
Section 13.1.23, “DROP TABLE Syntax”  
Section 13.1.25, “DROP VIEW Syntax”  
Section 13.7.1.4, “GRANT Syntax”  
Section 13.1.26, “RENAME TABLE Syntax”  
Section 13.1.27, “TRUNCATE TABLE Syntax”  
Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”  
Section 17.3.1, “Management of RANGE and LIST Partitions”  
Section 20.8, “Performance Schema General Table Characteristics”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 20.9.9.1, “The host\_cache Table”  
Section 19.26, “The INFORMATION\_SCHEMA TABLE\_PRIVILEGES Table”  
Section 6.2, “The MySQL Access Privilege System”

## E

[index top [4005]]

## EVENT

Section 13.1.2, “ALTER EVENT Syntax”  
Section 13.1.9, “CREATE EVENT Syntax”  
Section 13.1.18, “DROP EVENT Syntax”  
Section 13.7.1.4, “GRANT Syntax”  
Section 13.7.5.9, “SHOW CREATE EVENT Syntax”  
Section 13.7.5.19, “SHOW EVENTS Syntax”  
Section 18.4.1, “Event Scheduler Overview”  
Section 18.4.3, “Event Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## EXECUTE

Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 13.1.21, “DROP PROCEDURE and DROP FUNCTION Syntax”  
Section 13.7.1.4, “GRANT Syntax”  
Section 18.6, “Access Control for Stored Programs and Views”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.1.4, “Server System Variables”  
Section 18.2.2, “Stored Routines and MySQL Privileges”

## F

[index top [4005]]

## FILE

Section 13.7.1.4, “GRANT Syntax”  
Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 13.2.7, “LOAD XML Syntax”  
Section 13.2.9.1, “SELECT . . . INTO Syntax”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.5.4, “mysqldump — A Database Backup Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with mysqldump”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 12.5, “String Functions”  
Section 11.4.3, “The BLOB and TEXT Types”

## G

[index top [4005]]

## GRANT OPTION

Section 13.7.1.4, “GRANT Syntax”  
Section 13.7.1.6, “REVOKE Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 19.5, “The INFORMATION\_SCHEMA COLUMN\_PRIVILEGES Table”

## I

[index top [4005]]

## INDEX

Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.7.1.4, “GRANT Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 19.26, “The INFORMATION\_SCHEMA TABLE\_PRIVILEGES Table”

## INSERT

Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.7.1.1, “ALTER USER Syntax”  
Section 13.7.2.1, “ANALYZE TABLE Syntax”  
Section 13.7.3.1, “CREATE FUNCTION Syntax for User-Defined Functions”  
Section 13.7.1.2, “CREATE USER Syntax”  
Section 13.1.16, “CREATE VIEW Syntax”  
Section 13.7.1.4, “GRANT Syntax”  
Section 13.2.5, “INSERT Syntax”  
Section 13.7.2.4, “OPTIMIZE TABLE Syntax”  
Section 13.1.26, “RENAME TABLE Syntax”  
Section 13.7.2.5, “REPAIR TABLE Syntax”  
Section 13.2.8, “REPLACE Syntax”  
Section 18.6, “Access Control for Stored Programs and Views”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”



---

Section 6.3.2, “Adding User Accounts”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 14.12.1, “Pluggable Storage Engine Architecture”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.1.3, “Server Command Options”  
Section 19.5, “The `INFORMATION_SCHEMA` `COLUMN_PRIVILEGES` Table”  
Section 19.26, “The `INFORMATION_SCHEMA` `TABLE_PRIVILEGES` Table”  
Section 20.9.1.4, “The `setup_objects` Table”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”  
Section 22.3.2.6, “User-Defined Function Security Precautions”

## L

[index top [4005]]

## LOCK TABLES

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.3.5, “`LOCK TABLES` and `UNLOCK TABLES` Syntax”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 6.2.1, “Privileges Provided by MySQL”

## P

[index top [4005]]

## PROCESS

Section 13.7.1.4, “`GRANT` Syntax”  
Chapter 19, *INFORMATION\_SCHEMA* Tables  
Section 13.7.6.4, “`KILL` Syntax”  
Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”  
Section 13.7.5.30, “`SHOW PROCESSLIST` Syntax”  
Section 6.3.2, “Adding User Accounts”  
Section 18.4.2, “Event Scheduler Configuration”  
Section 8.12.5, “Examining Thread Information”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 19.16, “The `INFORMATION_SCHEMA` `PROCESSLIST` Table”  
Section 20.9.9.3, “The `threads` Table”

## PROXY

Section 13.7.1.4, “`GRANT` Syntax”

Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 6.3.8, “Proxy Users”  
Section 20.9.9.1, “The `host_cache` Table”  
Section 6.3.7.3.2.3, “Unix Password Authentication with Proxy Users and Group Mapping”  
Section 6.3.7.4.2, “Using the Windows Authentication Plugin”

## PROXY ... WITH GRANT OPTION

Section 6.3.8, “Proxy Users”

## R

[index top [4005]]

## REFERENCES

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 19.5, “The `INFORMATION_SCHEMA` `COLUMN_PRIVILEGES` Table”  
Section 19.26, “The `INFORMATION_SCHEMA` `TABLE_PRIVILEGES` Table”

## RELOAD

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 21.8.6.56, “`mysql_refresh()`”  
Section 21.8.6.57, “`mysql_reload()`”  
Section 13.7.6.6, “`RESET` Syntax”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 6.3.2, “Adding User Accounts”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 12.13, “Encryption and Compression Functions”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 20.9.9.1, “The `host_cache` Table”

## REPLICATION CLIENT

Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.5.2, “`SHOW BINARY LOGS` Syntax”  
Section 13.7.5.24, “`SHOW MASTER STATUS` Syntax”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”

## REPLICATION SLAVE

Section 13.7.1.4, “`GRANT` Syntax”  
Section 16.1.1.3, “Creating a User for Replication”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.3.7, “Setting Up Replication Using SSL”

---

## S

[[index top \[4005\]](#)]

### SELECT

[Section 13.7.2.1, “ANALYZE TABLE Syntax”](#)  
[Section 13.7.2.3, “CHECKSUM TABLE Syntax”](#)  
[Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#)  
[Section 13.1.14, “CREATE TABLE Syntax”](#)  
[Section 13.1.15, “CREATE TRIGGER Syntax”](#)  
[Section 13.1.16, “CREATE VIEW Syntax”](#)  
[Section 13.2.2, “DELETE Syntax”](#)  
[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 13.2.5, “INSERT Syntax”](#)  
[Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”](#)  
[Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#)  
[Section 13.7.2.5, “REPAIR TABLE Syntax”](#)  
[Section 13.7.5.14, “SHOW CREATE VIEW Syntax”](#)  
[Section 13.7.5.22, “SHOW GRANTS Syntax”](#)  
[Section 13.2.11, “UPDATE Syntax”](#)  
[Section 18.6, “Access Control for Stored Programs and Views”](#)  
[Section 6.2.5, “Access Control, Stage 2: Request Verification”](#)  
[Section 21.2.6, “Connector/Net Connection String Options Reference”](#)  
[Section 4.6.10, “mysqlhotcopy — A Database Backup Program”](#)  
[Section 4.5.4, “mysqldump — A Database Backup Program”](#)  
[Section 20.8, “Performance Schema General Table Characteristics”](#)  
[Section 6.2.1, “Privileges Provided by MySQL”](#)  
[Section E.5, “Restrictions on Views”](#)  
[Section 19.5, “The INFORMATION\\_SCHEMA COLUMN\\_PRIVILEGES Table”](#)  
[Section 19.26, “The INFORMATION\\_SCHEMA TABLE\\_PRIVILEGES Table”](#)  
[Section 14.8, “The MERGE Storage Engine”](#)  
[Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#)  
[Section 6.2, “The MySQL Access Privilege System”](#)  
[Section 18.3.1, “Trigger Syntax”](#)

### SHOW DATABASES

[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 13.7.5.15, “SHOW DATABASES Syntax”](#)  
[Section 6.2.1, “Privileges Provided by MySQL”](#)  
[Section 5.1.4, “Server System Variables”](#)

### SHOW VIEW

[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 13.7.5.14, “SHOW CREATE VIEW Syntax”](#)  
[Section 4.5.4, “mysqldump — A Database Backup Program”](#)

[Section 6.2.1, “Privileges Provided by MySQL”](#)  
[Section E.5, “Restrictions on Views”](#)  
[Section 19.29, “The INFORMATION\\_SCHEMA VIEWS Table”](#)

### SHUTDOWN

[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 21.8.6.66, “mysql\\_shutdown\(\)”](#)  
[Section 6.2.5, “Access Control, Stage 2: Request Verification”](#)  
[Section 4.3.4, “mysqld\\_multi — Manage Multiple MySQL Servers”](#)  
[Section 6.2.2, “Privilege System Grant Tables”](#)  
[Section 6.2.1, “Privileges Provided by MySQL”](#)  
[Section 16.1.3.2, “Setting Up Replication Using GTIDs”](#)  
[Section 5.1.12, “The Shutdown Process”](#)

### SUPER

[Section 13.1.3, “ALTER FUNCTION Syntax”](#)  
[Section 13.1.5, “ALTER SERVER Syntax”](#)  
[Section 13.1.7, “ALTER VIEW Syntax”](#)  
[Section 13.7.6.1, “BINLOG Syntax”](#)  
[Section 13.1.9, “CREATE EVENT Syntax”](#)  
[Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”](#)  
[Section 13.1.13, “CREATE SERVER Syntax”](#)  
[Section 13.1.15, “CREATE TRIGGER Syntax”](#)  
[Section 13.1.16, “CREATE VIEW Syntax”](#)  
[Section 13.1.22, “DROP SERVER Syntax”](#)  
[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 13.7.6.4, “KILL Syntax”](#)  
[Section 21.8.6.12, “mysql\\_dump\\_debug\\_info\(\)”](#)  
[Section 13.7.1.7, “SET PASSWORD Syntax”](#)  
[Section 13.4.1.3, “SET sql\\_log\\_bin Syntax”](#)  
[Section 13.3.6, “SET TRANSACTION Syntax”](#)  
[Section 13.7.4, “SET Syntax”](#)  
[Section 13.7.5.2, “SHOW BINARY LOGS Syntax”](#)  
[Section 13.7.5.24, “SHOW MASTER STATUS Syntax”](#)  
[Section 13.7.5.30, “SHOW PROCESSLIST Syntax”](#)  
[Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”](#)  
[Section 13.4.2.5, “START SLAVE Syntax”](#)  
[Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#)  
[Section 13.4.2.6, “STOP SLAVE Syntax”](#)  
[Section C.5.2.7, “Too many connections”](#)  
[Section 18.6, “Access Control for Stored Programs and Views”](#)  
[Section 6.3.5, “Assigning Account Passwords”](#)  
[Section 16.1.4.4, “Binary Log Options and Variables”](#)  
[Section 18.7, “Binary Logging of Stored Programs”](#)  
[Section 10.1.5, “Configuring the Character Set and Collation for Applications”](#)  
[Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#)  
[Section 12.13, “Encryption and Compression Functions”](#)  
[Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”](#)



---

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.1.1, “How to Set Up Replication”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 10.7, “MySQL Server Locale Support”

Section 10.6, “MySQL Server Time Zone Support”

Section 6.2.1, “Privileges Provided by MySQL”

Section 16.1.2, “Replication Formats”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 5.1.7, “Server SQL Modes”

Section 5.1.4, “Server System Variables”

Section 5.2.4.2, “Setting The Binary Log Format”

Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 5.2.4, “The Binary Log”

Section 18.4.6, “The Event Scheduler and MySQL Privileges”

Section 5.1.5, “Using System Variables”

Section 20.9.1, “Performance Schema Setup Tables”

Section 6.2.1, “Privileges Provided by MySQL”

Section 19.5, “The `INFORMATION_SCHEMA` `COLUMN_PRIVILEGES` Table”

Section 19.26, “The `INFORMATION_SCHEMA` `TABLE_PRIVILEGES` Table”

Section 14.8, “The `MERGE` Storage Engine”

Section 20.9.1.4, “The `setup_objects` Table”

Section 18.3.1, “Trigger Syntax”

## USAGE

Section 13.7.1.4, “`GRANT` Syntax”

Section 6.2.1, “Privileges Provided by MySQL”

## T

[\[index top \[4005\]\]](#)

### TRIGGER

Section 13.1.15, “`CREATE TRIGGER` Syntax”

Section 13.1.24, “`DROP TRIGGER` Syntax”

Section 13.7.1.4, “`GRANT` Syntax”

Section 13.7.5.39, “`SHOW TRIGGERS` Syntax”

Section 18.6, “Access Control for Stored Programs and Views”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 6.2.1, “Privileges Provided by MySQL”

Section 19.27, “The `INFORMATION_SCHEMA` `TRIGGERS` Table”

## U

[\[index top \[4005\]\]](#)

### UPDATE

Section 13.1.14, “`CREATE TABLE` Syntax”

Section 13.1.15, “`CREATE TRIGGER` Syntax”

Section 13.7.1.4, “`GRANT` Syntax”

Section 13.2.5, “`INSERT` Syntax”

Section 13.7.1.5, “`RENAME USER` Syntax”

Section 13.7.1.6, “`REVOKE` Syntax”

Section 13.2.11, “`UPDATE` Syntax”

Section 18.6, “Access Control for Stored Programs and Views”

Section 20.8, “Performance Schema General Table Characteristics”

Section 20.2.3, “Performance Schema Runtime Configuration”



---

# SQL Modes Index

[A](#) | [D](#) | [E](#) | [H](#) | [I](#) | [M](#) | [N](#) | [O](#) | [P](#) | [R](#) | [S](#) | [T](#)

## A

[\[index top \[4011\]\]](#)

### **ALLOW\_INVALID\_DATES**

Section 12.7, “Date and Time Functions”  
Section 11.3, “Date and Time Types”  
Section C.5.5.2, “Problems Using `DATE` Columns”  
Section 5.1.7, “Server SQL Modes”  
Section 11.3.1, “The `DATE`, `DATETIME`, and `TIMESTAMP` Types”

### **ANSI**

Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”  
Section 9.2.4, “Function Name Parsing and Resolution”  
Section 5.1.7, “Server SQL Modes”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”

### **ANSI\_QUOTES**

Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 9.2, “Schema Object Names”  
Section 5.1.7, “Server SQL Modes”  
Section 9.1.1, “String Literals”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”

## D

[\[index top \[4011\]\]](#)

### **DB2**

Section 5.1.7, “Server SQL Modes”

## E

[\[index top \[4011\]\]](#)

### **ERROR\_FOR\_DIVISION\_BY\_ZERO**

Section 12.19.3, “Expression Handling”  
Section 12.19.5, “Precision Math Examples”  
Section 5.1.7, “Server SQL Modes”

## H

[\[index top \[4011\]\]](#)

### **HIGH\_NOT\_PRECEDENCE**

Section 9.5, “Expression Syntax”  
Section 12.3.1, “Operator Precedence”  
Section 5.1.7, “Server SQL Modes”

## I

[\[index top \[4011\]\]](#)

### **IGNORE\_SPACE**

Section 4.5.1.1, “`mysql` Options”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 9.2.4, “Function Name Parsing and Resolution”  
Section 5.1.7, “Server SQL Modes”

## M

[\[index top \[4011\]\]](#)

### **MAXDB**

Section 11.1.2, “Date and Time Type Overview”  
Section 5.1.7, “Server SQL Modes”  
Section 11.3.1, “The `DATE`, `DATETIME`, and `TIMESTAMP` Types”

### **MSSQL**

Section 5.1.7, “Server SQL Modes”

### **MYSQL323**

Section 5.1.7, “Server SQL Modes”

### **MYSQL40**

Section 5.1.7, “Server SQL Modes”

## N

[\[index top \[4011\]\]](#)

### **NO\_AUTO\_CREATE\_USER**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.3.2, “Adding User Accounts”  
Section 5.1.7, “Server SQL Modes”

### **NO\_AUTO\_VALUE\_ON\_ZERO**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 5.1.7, “Server SQL Modes”  
Section 3.6.9, “Using `AUTO_INCREMENT`”

### **NO\_BACKSLASH\_ESCAPES**

Section 5.1.7, “Server SQL Modes”  
Section 12.5.1, “String Comparison Functions”  
Section 9.1.1, “String Literals”

### **NO\_DIR\_IN\_CREATE**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 16.4.1.9, “Replication and `DIRECTORY` Table Options”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.7, “Server SQL Modes”  
Section 17.2.6, “Subpartitioning”  
Section 5.2.4, “The Binary Log”

---

## **NO\_ENGINE\_SUBSTITUTION**

Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.1.14, “CREATE TABLE Syntax”  
Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 5.1.7, “Server SQL Modes”  
Section 14.1, “Setting the Storage Engine”  
Section 5.1.2.2, “Using a Sample Default Server Configuration File”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”

## **NO\_FIELD\_OPTIONS**

Section 5.1.7, “Server SQL Modes”

## **NO\_KEY\_OPTIONS**

Section 5.1.7, “Server SQL Modes”

## **NO\_TABLE\_OPTIONS**

Section 5.1.7, “Server SQL Modes”

## **NO\_UNSIGNED\_SUBTRACTION**

Section 12.6.1, “Arithmetic Operators”  
Section 12.10, “Cast Functions and Operators”  
Section 11.1.1, “Numeric Type Overview”  
Section 11.2.6, “Out-of-Range and Overflow Handling”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.7, “Server SQL Modes”

## **NO\_ZERO\_DATE**

Section 13.1.14, “CREATE TABLE Syntax”  
Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.10, “Cast Functions and Operators”  
Section 11.3, “Date and Time Types”  
Section C.5.5.2, “Problems Using `DATE` Columns”  
Section 5.1.7, “Server SQL Modes”

## **NO\_ZERO\_IN\_DATE**

Section 13.1.14, “CREATE TABLE Syntax”  
Section 11.3, “Date and Time Types”  
Section C.5.5.2, “Problems Using `DATE` Columns”  
Section 5.1.7, “Server SQL Modes”

## **O**

[index top [4011]]

## **ONLY\_FULL\_GROUP\_BY**

Section 3.3.4.8, “Counting Rows”  
Section 12.17.3, “MySQL Extensions to `GROUP BY`”  
Section 5.1.7, “Server SQL Modes”

## **ORACLE**

Section 5.1.7, “Server SQL Modes”

## **P**

[index top [4011]]

## **PAD\_CHAR\_TO\_FULL\_LENGTH**

Section 5.1.7, “Server SQL Modes”  
Section 11.1.3, “String Type Overview”  
Section 11.4.1, “The `CHAR` and `VARCHAR` Types”

## **PIPES\_AS\_CONCAT**

Section 9.5, “Expression Syntax”  
Section 12.3.1, “Operator Precedence”  
Section 5.1.7, “Server SQL Modes”

## **POSTGRESQL**

Section 5.1.7, “Server SQL Modes”

## **R**

[index top [4011]]

## **REAL\_AS\_FLOAT**

Section 11.1.1, “Numeric Type Overview”  
Section 11.2, “Numeric Types”  
Section 5.1.7, “Server SQL Modes”

## **S**

[index top [4011]]

## **STRICT\_ALL\_TABLES**

Section 6.3.2, “Adding User Accounts”  
Section 1.8.6.3, “Constraints on Invalid Data”  
Section 12.19.3, “Expression Handling”  
Section 5.1.7, “Server SQL Modes”

## **STRICT\_TRANS\_TABLES**

Section 6.3.2, “Adding User Accounts”  
Section 1.8.6.3, “Constraints on Invalid Data”  
Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”  
Section 12.19.3, “Expression Handling”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.2.2, “Using a Sample Default Server Configuration File”

## **T**

[index top [4011]]

## **TRADITIONAL**

Section 11.3.5, “Automatic Initialization and Updating for `TIMESTAMP` and `DATETIME`”  
Section 12.19.3, “Expression Handling”  
Section 5.1.7, “Server SQL Modes”

---

# Statement/Syntax Index

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [K](#) | [L](#) | [O](#) | [P](#) | [R](#) | [S](#) | [T](#) | [U](#) | [W](#) | [X](#)

## A

[[index top \[4013\]](#)]

### ALTER DATABASE

[Section 13.1.1, “ALTER DATABASE Syntax”](#)  
[Section 10.1.3.2, “Database Character Set and Collation”](#)  
[Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”](#)  
[Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”](#)  
[Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”](#)  
[Section 1.8.4, “MySQL Extensions to Standard SQL”](#)  
[Section 16.1.4.3, “Replication Slave Options and Variables”](#)

### ALTER EVENT

[Section 13.1.2, “ALTER EVENT Syntax”](#)  
[Section 13.1.9, “CREATE EVENT Syntax”](#)  
[Section 18.7, “Binary Logging of Stored Programs”](#)  
[Section 18.4.4, “Event Metadata”](#)  
[Section 18.4.1, “Event Scheduler Overview”](#)  
[Section 18.4.3, “Event Syntax”](#)  
[Section 12.14, “Information Functions”](#)  
[Section 16.4.1.6, “Replication of CURRENT\\_USER\( \)”](#)  
[Section 16.4.1.10, “Replication of Invoked Features”](#)  
[Section E.1, “Restrictions on Stored Programs”](#)  
[Section 13.3.3, “Statements That Cause an Implicit Commit”](#)  
[Section 19.7, “The INFORMATION\\_SCHEMA EVENTS Table”](#)  
[Section 18.4.6, “The Event Scheduler and MySQL Privileges”](#)

### ALTER EVENT event\_name ENABLED

[Section 16.4.1.10, “Replication of Invoked Features”](#)

### ALTER FUNCTION

[Section 13.1.3, “ALTER FUNCTION Syntax”](#)  
[Section 18.7, “Binary Logging of Stored Programs”](#)  
[Section 13.3.3, “Statements That Cause an Implicit Commit”](#)  
[Section 18.2.1, “Stored Routine Syntax”](#)

### ALTER PROCEDURE

[Section 13.1.4, “ALTER PROCEDURE Syntax”](#)  
[Section 18.7, “Binary Logging of Stored Programs”](#)

[Section 13.3.3, “Statements That Cause an Implicit Commit”](#)  
[Section 18.2.1, “Stored Routine Syntax”](#)

### ALTER SCHEMA

[Section 13.1.1, “ALTER DATABASE Syntax”](#)

### ALTER SERVER

[Section 13.1.5, “ALTER SERVER Syntax”](#)  
[Section 16.1.4.5, “Global Transaction ID Options and Variables”](#)  
[Section 16.4.1.5, “Replication of CREATE SERVER, ALTER SERVER, and DROP SERVER”](#)  
[Section 13.3.3, “Statements That Cause an Implicit Commit”](#)

### ALTER TABLE

[Section 4.6.3.1, “myisamchk General Options”](#)  
[Section 13.1.6.2, “ALTER TABLE Examples”](#)  
[Section 13.1.6.1, “ALTER TABLE Partition Operations”](#)  
[Section 13.1.6, “ALTER TABLE Syntax”](#)  
[Section 13.7.2.2, “CHECK TABLE Syntax”](#)  
[Section 13.1.11, “CREATE INDEX Syntax”](#)  
[Section 13.1.14, “CREATE TABLE Syntax”](#)  
[Section 13.1.20, “DROP INDEX Syntax”](#)  
[Section 8.8.2, “EXPLAIN Output Format”](#)  
[Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”](#)  
[Section 1.8.6.2, “FOREIGN KEY Constraints”](#)  
[Section 13.7.1.4, “GRANT Syntax”](#)  
[Section 5.4.5, “InnoDB and FOREIGN KEY Constraints”](#)  
[Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”](#)  
[Section 14.2.6, “InnoDB Startup Options and System Variables”](#)  
[Section 14.2.5.7, “InnoDB Strict Mode”](#)  
[Section 13.2.5.2, “INSERT DELAYED Syntax”](#)  
[Section 13.7.6.4, “KILL Syntax”](#)  
[Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”](#)  
[Section 14.8.2, “MERGE Table Problems”](#)  
[Section 14.3.1, “MyISAM Startup Options”](#)  
[Section 14.3.3, “MyISAM Table Storage Formats”](#)  
[Section 21.8.6.35, “mysql\\_info\( \)”](#)  
[Section 13.7.2.4, “OPTIMIZE TABLE Syntax”](#)  
[Section 17.2.3.1, “RANGE COLUMNS partitioning”](#)  
[Section 17.2.1, “RANGE Partitioning”](#)  
[Section 13.1.26, “RENAME TABLE Syntax”](#)  
[Section 13.7.5.41, “SHOW WARNINGS Syntax”](#)  
[Section C.5.7.2, “TEMPORARY Table Problems”](#)  
[Section 11.3.4, “YEAR\( 2 \) Limitations and Migrating to YEAR\( 4 \)”](#)  
[Section 14.2.5.9, “Better Error Handling when Dropping Indexes”](#)  
[Section 16.1.4.4, “Binary Log Options and Variables”](#)  
[Section 10.1.3.4, “Column Character Set and Collation”](#)

---

Section 10.1.13, “Column Character Set Conversion”  
Section 5.5.4, “Combining or Separating DDL Statements”  
Section 14.2.3.3, “Consistent Nonlocking Reads”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 3.3.2, “Creating a Table”  
Section 5.4, “Creating and Using InnoDB Tables and Indexes”  
Section 12.18.4.3, “Creating Spatial Columns”  
Section 12.18.6.1, “Creating Spatial Indexes”  
Section 5.5.8, “DDL for Partitioned Tables”  
Section 5.3.4, “Defragmenting a Table”  
Section 5.4.6.2, “Enabling Compression for a Table”  
Section 5.4.7.1, “Enabling File Formats”  
Section 5.5.5, “Examples of Online DDL”  
Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 12.9, “Full-Text Search Functions”  
Section 8.12.5.2, “General Thread States”  
Section 5.4.6.5, “How Compression Works for InnoDB Tables”  
Section 5.5.7, “How Crash Recovery Works with Online DDL”  
Section C.5.4.3, “How MySQL Handles a Full Disk”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 5.5.6, “Implementation Details of Online DDL”  
Section 12.14, “Information Functions”  
Section C.5.8, “Known Issues in MySQL”  
Section 5.5.9, “Limitations of Online DDL”  
Section E.10.3, “Limits on Table Size”  
Section 17.3.4, “Maintenance of Partitions”  
Section 17.3.2, “Management of HASH and KEY Partitions”  
Section 17.3.1, “Management of RANGE and LIST Partitions”  
Section 14.2.2.5, “Moving or Copying InnoDB Tables to Another Machine”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 11.2.6, “Out-of-Range and Overflow Handling”  
Section 5.4.8.1, “Overview of InnoDB Row Storage”  
Section 5.5.1, “Overview of Online DDL”  
Section 5.4.6.1, “Overview of Table Compression”  
Section 17.3, “Partition Management”  
Section 17.6.4, “Partitioning and Locking”  
Section 17.6.1, “Partitioning Keys, Primary Keys, and Unique Keys”  
Section 5.5.2, “Performance and Concurrency Considerations for Online DDL”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section C.5.7.1, “Problems with ALTER TABLE”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 16.4.1.1, “Replication and AUTO\_INCREMENT”  
Section 16.4.1.24, “Replication and Reserved Words”  
Section 16.2.2, “Replication Relay and Status Logs”  
Section 16.4.1.8.1, “Replication with More Columns on Master or Slave”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section E.5, “Restrictions on Views”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4.2, “Setting The Binary Log Format”  
Section 14.1, “Setting the Storage Engine”  
Section 13.1.14.3, “Silent Column Specification Changes”  
Section 5.4.1.2, “Specifying the Location of a Tablespace”  
Section 5.4.8.2, “Specifying the Row Format for a Table”  
Section 5.4.6.6, “SQL Compression Syntax Warnings and Errors”  
Section 5.5.3, “SQL Syntax for Online DDL”  
Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 11.1.3, “String Type Overview”  
Section 10.1.3.3, “Table Character Set and Collation”  
Section 14.6, “The ARCHIVE Storage Engine”  
Section 19.14, “The INFORMATION\_SCHEMA PARTITIONS Table”  
Section 14.4, “The MEMORY Storage Engine”  
Section 14.3, “The MyISAM Storage Engine”  
Section 5.2.5, “The Slow Query Log”  
Section 5.4.4.1, “Traditional InnoDB Auto-Increment Locking”  
Section 14.2.4.7, “Troubleshooting InnoDB Data Dictionary Operations”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 10.1.11, “Upgrading from Previous to Current Unicode Support”  
Section 3.6.9, “Using AUTO\_INCREMENT”  
Section 13.1.14.2, “Using FOREIGN KEY Constraints”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”  
Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”  
Section 5.4.7.2, “Verifying File Format Compatibility”  
Section 1.4, “What Is New in MySQL 5.6”  
Section C.5.4.2, “What to Do If MySQL Keeps Crashing”  
Section C.5.4.4, “Where MySQL Stores Temporary Files”

---

---

## **ALTER TABLE ... ADD FOREIGN KEY**

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **ALTER TABLE ... DROP FOREIGN KEY**

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **ALTER TABLE ... ENGINE = MEMORY**

Section 16.4.1.20, “Replication and MEMORY Tables”

## **ALTER TABLE ... ENGINE=...**

Section 1.4, “What Is New in MySQL 5.6”

## **ALTER TABLE ...**

### **ENGINE=InnoDB**

Section 16.2.2, “Replication Relay and Status Logs”

## **ALTER TABLE ... EXCHANGE PARTITION**

Section 13.1.6.1, “ALTER TABLE Partition Operations”

Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”

Section 17.6.4, “Partitioning and Locking”

Section 1.4, “What Is New in MySQL 5.6”

## **ALTER TABLE ... IMPORT TABLESPACE**

Section 5.4.1.3, “Copying Tablespaces to Another Server (Transportable Tablespaces)”

## **ALTER TABLE ... OPTIMIZE PARTITION**

Section 17.6.2, “Partitioning Limitations Relating to Storage Engines”

## **ALTER TABLE ... PARTITION BY**

Section 17.6.1, “Partitioning Keys, Primary Keys, and Unique Keys”

## **ALTER TABLE ... PARTITION BY ...**

Section 17.3.1, “Management of RANGE and LIST Partitions”

Section 17.6, “Restrictions and Limitations on Partitioning”

## **ALTER TABLE ... RENAME**

Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”

## **ALTER TABLE ... TRUNCATE PARTITION**

Section 17.3.4, “Maintenance of Partitions”

Section 17.3, “Partition Management”

Section 17.6.4, “Partitioning and Locking”

## **ALTER TABLE ... TRUNCATE PARTITION ALL**

Section 17.3.4, “Maintenance of Partitions”

## **ALTER TABLE EXCHANGE PARTITION**

Section 17.3.3, “Exchanging Partitions and Subpartitions with Tables”

## **ALTER TABLE IGNORE**

Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## **ALTER TABLE t TRUNCATE PARTITION ( )**

Section 13.2.2, “DELETE Syntax”

## **ALTER TABLE tbl\_name FORCE**

Section 13.1.6, “ALTER TABLE Syntax”

## **ALTER TABLESPACE**

Section 19.8, “The INFORMATION\_SCHEMA FILES Table”

## **ALTER USER**

Section 13.7.1.1, “ALTER USER Syntax”

Section 6.3.6, “Password Expiration and Sandbox Mode”

Section 6.2.2, “Privilege System Grant Tables”

Section 1.4, “What Is New in MySQL 5.6”

## **ALTER VIEW**

Section 13.1.7, “ALTER VIEW Syntax”

Section 13.1.16, “CREATE VIEW Syntax”

Section 12.14, “Information Functions”

Section 16.4.1.6, “Replication of CURRENT\_USER( )”

Section E.1, “Restrictions on Stored Programs”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 18.5.2, “View Processing Algorithms”

Section 18.5.1, “View Syntax”

## **ANALYZE TABLE**

Section 4.6.3.1, “myisamchk General Options”

Section 13.1.6, “ALTER TABLE Syntax”

Section 13.7.2.1, “ANALYZE TABLE Syntax”

Section 13.1.11, “CREATE INDEX Syntax”

Section 13.1.14, “CREATE TABLE Syntax”

Section 8.8.2, “EXPLAIN Output Format”



---

Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.8.2, “MERGE Table Problems”  
Section 7.6, “MyISAM Table Maintenance and Crash Recovery”  
Section 13.7.5.23, “SHOW INDEX Syntax”  
Section 14.2.5.8, “Controlling Optimizer Statistics Estimation”  
Section 4.5.3, “mysqlcheck — A Table Maintenance Program”  
Section 8.13.1.3, “Equality Range Optimization of Many-Valued Comparisons”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 8.12.5.2, “General Thread States”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section C.5.8, “Known Issues in MySQL”  
Section 14.2.7, “Limits on InnoDB Tables”  
Section 17.3.4, “Maintenance of Partitions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 8.6.1, “Optimizing MyISAM Queries”  
Section 8.8.1, “Optimizing Queries with EXPLAIN”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.13, “Replication and FLUSH”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 8.2.1.1, “Speed of SELECT Statements”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 5.2.5, “The Slow Query Log”

## B

[index top [4013]]

### BEGIN

Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”  
Section 14.2.3.14, “InnoDB Error Handling”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 16.4.1.29, “Replication and Transactions”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section E.1, “Restrictions on Stored Programs”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.3.1, “The InnoDB Transaction Model and Locking”

### BEGIN ... END

Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”  
Section 13.6.5.1, “CASE Syntax”

Section 13.1.15, “CREATE TRIGGER Syntax”  
Section 13.6.7.2, “DECLARE ... HANDLER Syntax”  
Section 13.6.3, “DECLARE Syntax”  
Section 13.6.5.4, “LEAVE Syntax”  
Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”  
Section 13.6.6.1, “Cursor CLOSE Syntax”  
Section 18.1, “Defining Stored Programs”  
Section 18.4.1, “Event Scheduler Overview”  
Section 13.6.4.1, “Local Variable DECLARE Syntax”  
Section 13.6.4.2, “Local Variable Scope and Resolution”  
Section 13.6, “MySQL Compound-Statement Syntax”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.6.7.6, “Scope Rules for Handlers”  
Section 13.6.2, “Statement Label Syntax”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.3.1, “Trigger Syntax”

## BINLOG

Section 4.6.8.2, “mysqlbinlog Row Event Display”  
Section 13.7.6.1, “BINLOG Syntax”  
Section 4.6.8, “mysqlbinlog — Utility for Processing Binary Log Files”

## C

[index top [4013]]

### CACHE INDEX

Section 13.7.6.2, “CACHE INDEX Syntax”  
Section 13.7.6.5, “LOAD INDEX INTO CACHE Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 8.9.2.4, “Index Preloading”  
Section 8.9.2.2, “Multiple Key Caches”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

### CALL

Section 13.2.1, “CALL Syntax”  
Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.37, “mysql\_insert\_id()”  
Section 21.8.6.45, “mysql\_more\_results()”  
Section 21.8.6.46, “mysql\_next\_result()”  
Section 21.8.6.53, “mysql\_real\_connect()”  
Section 21.8.6.65, “mysql\_set\_server\_option()”  
Section 21.8.10.17, “mysql\_stmt\_next\_result()”  
Section 18.6, “Access Control for Stored Programs and Views”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 21.8.4, “C API Data Structures”



---

Section 21.8.17, “C API Prepared Statement Problems”  
Section 21.8.16, “C API Support for Multiple Statement Execution”  
Section 21.8.19, “C API Support for Prepared `CALL` Statements”  
Section 13.5, “SQL Syntax for Prepared Statements”  
Chapter 18, *Stored Programs and Views*  
Section 18.2.1, “Stored Routine Syntax”  
Section 18.3.1, “Trigger Syntax”

## **CALL p()**

Section 13.6.7.4.3, “`RESIGNAL` with a Condition Value and Optional New Signal Information”

## **CALL stored\_procedure()**

Section 17.6.4, “Partitioning and Locking”

## **CASE**

Section 13.6.5.1, “`CASE` Syntax”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 12.4, “Control Flow Functions”  
Section 13.6.5, “Flow Control Statements”

## **CHANGE MASTER TO**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.4.2.3, “`RESET SLAVE` Syntax”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 16.1.1.5, “Creating a Data Snapshot Using `mysqldump`”  
Section 16.3.9, “Delayed Replication”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.3.1, “GTID Concepts”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 16.4.1.18, “Replication and Master or Slave Shutdowns”  
Section 16.1, “Replication Configuration”  
Section 8.12.5.8, “Replication Slave Connection Thread States”  
Section 8.12.5.6, “Replication Slave I/O Thread States”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 8.12.5.7, “Replication Slave SQL Thread States”  
Section 5.1.6, “Server Status Variables”  
Section 16.1.1.10, “Setting the Master Configuration on the Slave”  
Section 16.1.3.2, “Setting Up Replication Using GTIDs”

Section 16.3.7, “Setting Up Replication Using SSL”  
Section 16.1.1.8, “Setting Up Replication with Existing Data”  
Section 16.1.1.7, “Setting Up Replication with New Master and Slaves”  
Section 16.2.2.2, “Slave Status Logs”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 16.3.6, “Switching Masters During Failover”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 1.4, “What Is New in MySQL 5.6”

## **CHANGE MASTER TO ...**

### **MASTER\_LOG\_FILE = ...**

### **MASTER\_LOG\_POS = ...**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

## **CHECK TABLE**

Section 13.1.6.1, “`ALTER TABLE` Partition Operations”  
Section 13.7.2.2, “`CHECK TABLE` Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 14.2.4.5, “InnoDB General Troubleshooting”  
Section 7.6, “MyISAM Table Maintenance and Crash Recovery”  
Section C.5.2.9, “MySQL server has gone away”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section 11.3.4, “`YEAR(2)` Limitations and Migrating to `YEAR(4)`”  
Section 14.2.2.4, “Backing Up and Recovering an InnoDB Database”  
Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”  
Section 14.3.4.1, “Corrupted MyISAM Tables”  
Section 4.6.1, “`innochecksum` — Offline InnoDB File Checksum Utility”  
Section 4.6.3, “`myisamchk` — MyISAM Table-Maintenance Utility”  
Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”  
Section 4.4.7, “`mysql_upgrade` — Check and Upgrade MySQL Tables”  
Section 8.10.5, “External Locking”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 7.6.3, “How to Repair MyISAM Tables”  
Section 1.7, “How to Report Bugs or Problems”  
Section 17.3.4, “Maintenance of Partitions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”

---

Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section E.3, “Restrictions on Server-Side Cursors”  
Section E.1, “Restrictions on Stored Programs”  
Section E.5, “Restrictions on Views”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 7.6.5, “Setting Up a `MyISAM` Table Maintenance Schedule”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 5.2.5, “The Slow Query Log”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **CHECK TABLE ... EXTENDED**

Section 13.7.2.2, “`CHECK TABLE` Syntax”

## **CHECK TABLE ... FOR UPGRADE**

Section 13.7.2.5, “`REPAIR TABLE` Syntax”  
Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”

## **CHECKSUM TABLE**

Section 13.7.2.3, “`CHECKSUM TABLE` Syntax”  
Section 13.1.14, “`CREATE TABLE` Syntax”

## **CLOSE**

Section 13.6.6.1, “Cursor `CLOSE` Syntax”

## **COMMIT**

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 14.2.3.14, “InnoDB Error Handling”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 13.3.4, “`SAVEPOINT`, `ROLLBACK TO SAVEPOINT`, and `RELEASE SAVEPOINT` Syntax”  
Section 13.3.1, “`START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 8.5.4, “Bulk Data Loading for InnoDB Tables”  
Section 5.4.3, “Converting Tables from `MyISAM` to InnoDB”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 5.4.2, “Grouping DML Operations with Transactions”  
Section 14.2.3.8, “Implicit Transaction Commit and Rollback”  
Section 14.2.7, “Limits on InnoDB Tables”  
Section 14.2.8, “MySQL and the ACID Model”

Section 13.3, “MySQL Transactional and Locking Statements”  
Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”  
Section 16.4.1.29, “Replication and Transactions”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.3.1, “The InnoDB Transaction Model and Locking”  
Section 5.2.4, “The Binary Log”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”  
Section 18.3.1, “Trigger Syntax”  
Section 3.6.9, “Using `AUTO_INCREMENT`”

## **CREATE DATABASE**

Section 13.1.8, “`CREATE DATABASE` Syntax”  
Section 21.8.6.8, “`mysql_create_db()`”  
Section 13.7.5.8, “`SHOW CREATE DATABASE` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 7.1, “Backup and Recovery Types”  
Section 21.8.5, “C API Function Overview”  
Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 7.4.5.2, “Copy a Database from one Server to Another”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.1, “Dumping Data in SQL Format with `mysqldump`”  
Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”  
Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 20.4, “Performance Schema Instrument Naming Conventions”  
Section 7.4.2, “Reloading SQL-Format Backups”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## **CREATE DATABASE dbx**

Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”

---

## CREATE DATABASE IF NOT EXISTS

Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”

## CREATE EVENT

Section 13.1.2, “`ALTER EVENT` Syntax”  
Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.7.5.9, “`SHOW CREATE EVENT` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 18.4.4, “Event Metadata”  
Section 18.4.3, “Event Syntax”  
Section 12.14, “Information Functions”  
Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 19.7, “The `INFORMATION_SCHEMA EVENTS` Table”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## CREATE EVENT IF NOT EXISTS

Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”

## CREATE FUNCTION

Section 13.1.3, “`ALTER FUNCTION` Syntax”  
Section 13.1.10, “`CREATE FUNCTION` Syntax”  
Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.7.3.2, “`DROP FUNCTION` Syntax”  
Section 22.3, “Adding New Functions to MySQL”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 9.2.4, “Function Name Parsing and Resolution”  
Section 12.14, “Information Functions”  
Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.2.1, “Stored Routine Syntax”  
Section 22.3.2.1, “UDF Calling Sequences for Simple Functions”  
Section 2.11.1, “Upgrading MySQL”  
Section 22.3.2.6, “User-Defined Function Security Precautions”

## CREATE INDEX

Section 13.1.11, “`CREATE INDEX` Syntax”

Section 14.2.3.12.3, “`FULLTEXT` Indexes”  
Section 14.2.5.7, “`InnoDB Strict Mode`”  
Section 14.2.5.9, “Better Error Handling when Dropping Indexes”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the `InnoDB memcached` Plugin”  
Section 12.18.6.1, “Creating Spatial Indexes”  
Section 5.4.6.2, “Enabling Compression for a Table”  
Section 5.5.5, “Examples of Online DDL”  
Section 12.9, “Full-Text Search Functions”  
Section 5.4.6.5, “How Compression Works for `InnoDB` Tables”  
Section 5.5.7, “How Crash Recovery Works with Online DDL”  
Section 8.7, “Optimizing for `MEMORY` Tables”  
Section 5.5.1, “Overview of Online DDL”  
Section 5.5.2, “Performance and Concurrency Considerations for Online DDL”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 5.2.5, “The Slow Query Log”

## CREATE LOGFILE GROUP

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## CREATE OR REPLACE VIEW

Section 13.1.7, “`ALTER VIEW` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section E.5, “Restrictions on Views”

## CREATE PROCEDURE

Section 13.1.4, “`ALTER PROCEDURE` Syntax”  
Section 13.2.1, “`CALL` Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 12.14, “Information Functions”  
Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.2.1, “Stored Routine Syntax”

## CREATE SCHEMA

Section 13.1.8, “`CREATE DATABASE` Syntax”

## CREATE SERVER

Section 13.1.5, “`ALTER SERVER` Syntax”  
Section 13.1.13, “`CREATE SERVER` Syntax”  
Section 14.9.3, “`FEDERATED` Storage Engine Notes and Tips”  
Section 13.7.6.3, “`FLUSH` Syntax”

---

Section 14.9.2.2, “Creating a `FEDERATED` Table Using `CREATE SERVER`”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 14.9.2, “How to Create `FEDERATED` Tables”  
Section 16.4.1.5, “Replication of `CREATE SERVER`, `ALTER SERVER`, and `DROP SERVER`”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## **CREATE TABLE**

Section 4.5.1.1, “`mysql` Options”  
Section 13.1.6.1, “`ALTER TABLE` Partition Operations”  
Section 13.1.6, “`ALTER TABLE` Syntax”  
Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.1.11, “`CREATE INDEX` Syntax”  
Section 13.1.13, “`CREATE SERVER` Syntax”  
Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”  
Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 13.8.2, “`EXPLAIN` Syntax”  
Section 1.8.6.2, “`FOREIGN KEY` Constraints”  
Section 17.2.4, “`HASH` Partitioning”  
Section 13.8.3, “`HELP` Syntax”  
Section 5.4.5, “`InnoDB` and `FOREIGN KEY` Constraints”  
Section 14.2.2.6, “`InnoDB` and MySQL Replication”  
Section 14.2.4.5, “`InnoDB` General Troubleshooting”  
Section 14.2.6, “`InnoDB` Startup Options and System Variables”  
Section 14.2.5.7, “`InnoDB` Strict Mode”  
Section 17.2.5, “`KEY` Partitioning”  
Section 17.2.2, “`LIST` Partitioning”  
Section 13.2.7, “`LOAD XML` Syntax”  
Section 14.3.3, “`MyISAM` Table Storage Formats”  
Section 17.2.3.1, “`RANGE COLUMNS` partitioning”  
Section 17.2.1, “`RANGE` Partitioning”  
Section 13.7.5.6, “`SHOW COLUMNS` Syntax”  
Section 13.7.5.12, “`SHOW CREATE TABLE` Syntax”  
Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”  
Section 13.7.5.37, “`SHOW TABLE STATUS` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 13.1.27, “`TRUNCATE TABLE` Syntax”  
Section 13.7.3.4, “`UNINSTALL PLUGIN` Syntax”  
Section 7.1, “Backup and Recovery Types”  
Section 14.2.5.9, “Better Error Handling when Dropping Indexes”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 10.1.3.4, “Column Character Set and Collation”  
Section 8.3.4, “Column Indexes”  
Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”  
Section 14.9.2.1, “Creating a `FEDERATED` Table Using `CONNECTION`”

Section 3.3.2, “Creating a Table”  
Section 5.4, “Creating and Using `InnoDB` Tables and Indexes”  
Section 12.18.4.3, “Creating Spatial Columns”  
Section 12.18.6.1, “Creating Spatial Indexes”  
Section 7.2, “Database Backup Methods”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”  
Section 5.4.6.2, “Enabling Compression for a Table”  
Section 5.4.7.1, “Enabling File Formats”  
Section 5.5.5, “Examples of Online DDL”  
Section 1.8.5.4, “Foreign Key Differences”  
Section 12.9, “Full-Text Search Functions”  
Section 3.4, “Getting Information About Databases and Tables”  
Section 5.4.8, “How `InnoDB` Stores Variable-Length Columns”  
Section 5.4.6.5, “How Compression Works for `InnoDB` Tables”  
Section 17.2.7, “How MySQL Partitioning Handles `NULL`”  
Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 12.14, “Information Functions”  
Section E.10.3, “Limits on Table Size”  
Section 3.3.3, “Loading Data into a Table”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 17.3.1, “Management of `RANGE` and `LIST` Partitions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 8.5.6, “Optimizing `InnoDB` DDL Operations”  
Section 8.4.1, “Optimizing Data Size”  
Section 5.4.8.1, “Overview of `InnoDB` Row Storage”  
Section 5.5.1, “Overview of Online DDL”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section 5.4.6.1, “Overview of Table Compression”  
Section 17.3, “Partition Management”  
Section 17.6.1, “Partitioning Keys, Primary Keys, and Unique Keys”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 17.2, “Partitioning Types”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 7.4.4, “Reloading Delimited-Text Format Backups”  
Section 16.4.1.1, “Replication and `AUTO_INCREMENT`”  
Section 16.4.1.9, “Replication and `DIRECTORY` Table Options”  
Section 16.4.1.2, “Replication and Character Sets”

---

Section 16.4.1.12, “Replication and Fractional Seconds Support”

Section 16.4.1.14, “Replication and System Functions”

Section 16.4.1.4, “Replication of `CREATE TABLE ... SELECT` Statements”

Section 16.4.1.8.1, “Replication with More Columns on Master or Slave”

Section 17.6, “Restrictions and Limitations on Partitioning”

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

Section 5.1.3, “Server Command Options”

Section 5.1.7, “Server SQL Modes”

Section 5.1.4, “Server System Variables”

Section 5.2.4.2, “Setting The Binary Log Format”

Section 14.1, “Setting the Storage Engine”

Section 13.1.14.3, “Silent Column Specification Changes”

Section C.1, “Sources of Error Information”

Section 5.4.1.2, “Specifying the Location of a Tablespace”

Section 5.4.8.2, “Specifying the Row Format for a Table”

Section 5.4.6.6, “SQL Compression Syntax Warnings and Errors”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Chapter 14, *Storage Engines*

Section 11.1.3, “String Type Overview”

Section 17.2.6, “Subpartitioning”

Section 10.1.3.3, “Table Character Set and Collation”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

Section 14.6, “The `ARCHIVE` Storage Engine”

Section 11.4.4, “The `ENUM` Type”

Section 19.14, “The `INFORMATION_SCHEMA PARTITIONS` Table”

Section 14.2, “The `InnoDB` Storage Engine”

Section 14.4, “The `MEMORY` Storage Engine”

Section 14.3, “The `MyISAM` Storage Engine”

Section 13.2.10.1, “The Subquery as Scalar Operand”

Section 5.4.4.1, “Traditional `InnoDB` Auto-Increment Locking”

Section 14.2.4.7, “Troubleshooting `InnoDB` Data Dictionary Operations”

Section 7.4, “Using `mysqldump` for Backups”

Section 3.6.9, “Using `AUTO_INCREMENT`”

Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”

Section 3.3.4.9, “Using More Than one Table”

Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”

Section 8.11.3.1, “Using Symbolic Links”

Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”

Section E.10.6, “Windows Platform Limitations”

## **CREATE TABLE ... LIKE**

Section 13.1.14, “`CREATE TABLE` Syntax”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

## **CREATE TABLE ... SELECT**

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”

Section 1.8.5.1, “`SELECT INTO TABLE` Differences”

Section 13.2.9, “`SELECT` Syntax”

Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”

Section 18.7, “Binary Logging of Stored Programs”

Section 12.10, “Cast Functions and Operators”

Section 14.2.3.3, “Consistent Nonlocking Reads”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section C.5.8, “Known Issues in MySQL”

Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”

Section 16.4.2, “Replication Compatibility Between MySQL Versions”

Section 16.4.1.4, “Replication of `CREATE TABLE ... SELECT` Statements”

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

Section 13.3.3, “Statements That Cause an Implicit Commit”

## **CREATE TABLE ... SELECT ...**

Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”

Section 17.3.1, “Management of `RANGE` and `LIST` Partitions”

## **CREATE TABLE IF NOT EXISTS**

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”

Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”

## **CREATE TABLE IF NOT EXISTS ... LIKE**

Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”

## **CREATE TABLE IF NOT EXISTS ... SELECT**

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”

Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”



---

## **CREATE TABLE `new_table`**

### **SELECT ... FROM**

## **old\_table ...**

Section 13.1.14.1, “CREATE TABLE ... SELECT Syntax”

Section 13.2.9, “SELECT Syntax”

## **CREATE TABLESPACE**

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 19.8, “The `INFORMATION_SCHEMA FILES` Table”

## **CREATE TEMPORARY TABLE**

Section 13.7.1.4, “GRANT Syntax”

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 7.5, “Point-in-Time (Incremental) Recovery Using the Binary Log”

Section 6.2.1, “Privileges Provided by MySQL”

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

Section 5.1.4, “Server System Variables”

Section 14.1, “Setting the Storage Engine”

Section 13.3.3, “Statements That Cause an Implicit Commit”

## **CREATE TRIGGER**

Section 13.1.15, “CREATE TRIGGER Syntax”

Section 13.7.5.13, “SHOW CREATE TRIGGER Syntax”

Section 18.7, “Binary Logging of Stored Programs”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 12.14, “Information Functions”

Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”

Section 16.4.1.6, “Replication of `CURRENT_USER()`”

Section 16.4.1.10, “Replication of Invoked Features”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 18.3.1, “Trigger Syntax”

## **CREATE USER**

Section 13.7.1.2, “CREATE USER Syntax”

Section 13.7.6.3, “FLUSH Syntax”

Section 13.7.1.4, “GRANT Syntax”

Section 6.3.2, “Adding User Accounts”

Section 6.3.5, “Assigning Account Passwords”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 5.1.9.3, “Connecting Using the IPv6 Local Host Address”

Section 16.1.1.3, “Creating a User for Replication”

Section 6.1.2.1, “End-User Guidelines for Password Security”

Section 8.11.4.1, “How MySQL Uses Memory”

Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”

Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”

Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”

Section 5.1.9, “IPv6 Support”

Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”

Section 6.1.2.4, “Password Hashing in MySQL”

Section 22.2.3.8, “Password-Validation Plugins”

Section 6.1.2.3, “Passwords and Logging”

Section 6.3.7, “Pluggable Authentication”

Section 6.2.1, “Privileges Provided by MySQL”

Section 6.3.8, “Proxy Users”

Section 5.1.4, “Server System Variables”

Section 6.2.3, “Specifying Account Names”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 6.2, “The MySQL Access Privilege System”

Section 6.1.2.6, “The Password Validation Plugin”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

Section 6.3.1, “User Names and Passwords”

Section 6.3.7.3.2, “Using the PAM Authentication Plugin”

Section 6.3.7.4.2, “Using the Windows Authentication Plugin”

Section 1.4, “What Is New in MySQL 5.6”

## **CREATE VIEW**

Section 13.1.7, “ALTER VIEW Syntax”

Section 13.1.16, “CREATE VIEW Syntax”

Section 13.7.5.14, “SHOW CREATE VIEW Syntax”

Section 8.12.5.2, “General Thread States”

Section 12.14, “Information Functions”

Section 17.6.4, “Partitioning and Locking”

Section 6.2.1, “Privileges Provided by MySQL”

Section 16.4.1.6, “Replication of `CURRENT_USER()`”

Section E.5, “Restrictions on Views”

Section 9.2, “Schema Object Names”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”

Section 18.5.3, “Updatable and Insertable Views”

Section 18.5.2, “View Processing Algorithms”

Section 18.5.1, “View Syntax”

## **D**

[index top [4013]]

## **DEALLOCATE PREPARE**

Section 13.5.3, “DEALLOCATE PREPARE Syntax”

Section 13.5.1, “PREPARE Syntax”

---

Section E.1, “Restrictions on Stored Programs”  
Section 5.1.6, “Server Status Variables”  
Section 13.5, “SQL Syntax for Prepared Statements”

## DECLARE

Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 13.6.3, “DECLARE Syntax”  
Section 13.6.7.3, “GET DIAGNOSTICS Syntax”  
Section 13.6.7.5, “SIGNAL Syntax”  
Section 13.6.6.2, “Cursor DECLARE Syntax”  
Section 13.6.4.1, “Local Variable DECLARE Syntax”  
Section 13.6.4, “Variables in Stored Programs”

## DECLARE ... CONDITION

Section 13.6.7.1, “DECLARE ... CONDITION Syntax”  
Section 13.6.7.2, “DECLARE ... HANDLER Syntax”  
Section 13.6.7.5, “SIGNAL Syntax”  
Section 13.6.7, “Condition Handling”

## DECLARE ... HANDLER

Section 13.6.7.1, “DECLARE ... CONDITION Syntax”  
Section 13.6.7.2, “DECLARE ... HANDLER Syntax”  
Section 13.6.7, “Condition Handling”  
Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”

## DELETE

Section 4.5.1.1, “mysql Options”  
Section 13.1.6.1, “ALTER TABLE Partition Operations”  
Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 13.1.15, “CREATE TRIGGER Syntax”  
Section 13.1.16, “CREATE VIEW Syntax”  
Section 13.2.2, “DELETE Syntax”  
Section 8.8.3, “EXPLAIN EXTENDED Output Format”  
Section 8.8.2, “EXPLAIN Output Format”  
Section 13.8.2, “EXPLAIN Syntax”  
Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”  
Section 13.7.1.4, “GRANT Syntax”  
Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 14.2.2.6, “InnoDB and MySQL Replication”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 13.2.9.2, “JOIN Syntax”  
Section 13.7.6.4, “KILL Syntax”  
Section 17.2.2, “LIST Partitioning”  
Section 14.8.2, “MERGE Table Problems”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.48, “mysql\_num\_rows()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.13, “mysql\_stmt\_field\_count()”  
Section 21.8.10.18, “mysql\_stmt\_num\_rows()”  
Section 17.2.1, “RANGE Partitioning”  
Section 13.7.1.6, “REVOKE Syntax”  
Section 13.3.6, “SET TRANSACTION Syntax”

Section 13.1.27, “TRUNCATE TABLE Syntax”  
Section 6.3.2, “Adding User Accounts”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 8.6.2, “Bulk Data Loading for MyISAM Tables”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 14.2.3.3, “Consistent Nonlocking Reads”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section C.5.5.6, “Deleting Rows from Related Tables”  
Section 5.5.5, “Examples of Online DDL”  
Section 1.8.5.4, “Foreign Key Differences”  
Chapter 12, *Functions and Operators*  
Section 8.12.5.2, “General Thread States”  
Section 8.2.1.2, “How MySQL Optimizes WHERE Clauses”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section 8.10.1, “Internal Locking Methods”  
Section C.5.8, “Known Issues in MySQL”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”  
Section 5.2.4.4, “Logging Format for Changes to mysql Database Tables”  
Section 17.3.1, “Management of RANGE and LIST Partitions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 8.2.2, “Optimizing DML Statements”  
Section 8.8.1, “Optimizing Queries with EXPLAIN”  
Section 5.5.1, “Overview of Online DDL”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section 17.5, “Partition Selection”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 8.12.5.4, “Query Cache Thread States”  
Section 16.4.1.15, “Replication and LIMIT”  
Section 16.4.1.20, “Replication and MEMORY Tables”  
Section 16.4.1.23, “Replication and the Query Optimizer”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 9.3, “Reserved Words”  
Section E.5, “Restrictions on Views”  
Section 13.2.10.11, “Rewriting Subqueries as Joins”  
Section 3.3.4.1, “Selecting All Data”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 8.2.2.3, “Speed of DELETE Statements”  
Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”  
Section 5.7.1.12, “Statement Probes”  
Section 13.2.10.9, “Subquery Errors”

---

Section 13.2.10, “Subquery Syntax”  
Section 8.10.2, “Table Locking Issues”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 14.4, “The `MEMORY` Storage Engine”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 5.2.4, “The Binary Log”  
Section 6.2, “The MySQL Access Privilege System”  
Section 18.3.1, “Trigger Syntax”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”  
Section 3.6.9, “Using `AUTO_INCREMENT`”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”  
Section 18.3, “Using Triggers”  
Section 1.4, “What Is New in MySQL 5.6”  
Section 21.8.14.2, “What Results You Can Get from a Query”  
Section 6.2.6, “When Privilege Changes Take Effect”  
Section 21.8.14.1, “Why `mysql_store_result()` Sometimes Returns `NULL` After `mysql_query()` Returns Success”  
Section 22.2.4.8, “Writing Audit Plugins”

## **DELETE FROM ... WHERE ...**

Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”

## **DELETE FROM t1,t2**

Section 5.7.1.12, “Statement Probes”

## **DESCRIBE**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 13.8.1, “`DESCRIBE` Syntax”  
Section 13.8.2, “`EXPLAIN` Syntax”  
Section 21.8.10.28, “`mysql_stmt_store_result()`”  
Section 21.8.6.70, “`mysql_store_result()`”  
Section 21.8.6.72, “`mysql_use_result()`”  
Section 13.7.5.6, “`SHOW COLUMNS` Syntax”  
Section 21.8.4, “C API Data Structures”  
Section 21.8.5, “C API Function Overview”  
Section 3.3.2, “Creating a Table”  
Section 19.32, “Extensions to `SHOW` Statements”  
Section 3.4, “Getting Information About Databases and Tables”  
Section 8.2.1.3, “Optimizing `LIMIT` Queries”  
Section 8.8.1, “Optimizing Queries with `EXPLAIN`”  
Section 13.1.14.3, “Silent Column Specification Changes”  
Section 3.6.6, “Using Foreign Keys”  
Section 10.1.12, “UTF-8 for Metadata”

## **DO**

Section 13.1.2, “`ALTER EVENT` Syntax”  
Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.2.3, “`DO` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 12.16, “Miscellaneous Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.6.4, “Partitioning and Locking”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.2.10, “Subquery Syntax”  
Section 19.7, “The `INFORMATION_SCHEMA EVENTS` Table”

## **DROP DATABASE**

Section 13.1.17, “`DROP DATABASE` Syntax”  
Section 21.8.6.11, “`mysql_drop_db()`”  
Section 21.8.5, “C API Function Overview”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.4.1, “Dumping Data in SQL Format with `mysqldump`”  
Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”  
Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 7.5, “Point-in-Time (Incremental) Recovery Using the Binary Log”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section E.10.6, “Windows Platform Limitations”

## **DROP DATABASE IF EXISTS**

Section 16.4.1.7, “Replication of `DROP ... IF EXISTS` Statements”

## **DROP EVENT**

Section 13.1.18, “`DROP EVENT` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 18.4.3, “Event Syntax”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## **DROP FUNCTION**

Section 13.1.3, “`ALTER FUNCTION` Syntax”  
Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”



---

Section 13.1.19, “`DROP FUNCTION` Syntax”  
Section 13.7.3.2, “`DROP FUNCTION` Syntax”  
Section 13.1.21, “`DROP PROCEDURE` and `DROP FUNCTION` Syntax”  
Section 22.3, “Adding New Functions to MySQL”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 9.2.4, “Function Name Parsing and Resolution”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.2.1, “Stored Routine Syntax”  
Section 2.11.1, “Upgrading MySQL”  
Section 22.3.2.6, “User-Defined Function Security Precautions”

## DROP INDEX

Section 13.1.6, “`ALTER TABLE` Syntax”  
Section 13.1.20, “`DROP INDEX` Syntax”  
Section 12.18.6.1, “Creating Spatial Indexes”  
Section 5.5.5, “Examples of Online DDL”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 5.5.1, “Overview of Online DDL”  
Section 5.1.3, “Server Command Options”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 5.2.5, “The Slow Query Log”

## DROP PROCEDURE

Section 13.1.4, “`ALTER PROCEDURE` Syntax”  
Section 13.1.21, “`DROP PROCEDURE` and `DROP FUNCTION` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.2.1, “Stored Routine Syntax”

## DROP SCHEMA

Section 13.1.17, “`DROP DATABASE` Syntax”  
Section 5.1.4, “Server System Variables”

## DROP SERVER

Section 13.1.22, “`DROP SERVER` Syntax”  
Section 13.7.6.3, “`FLUSH` Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 16.4.1.5, “Replication of `CREATE SERVER`, `ALTER SERVER`, and `DROP SERVER`”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## DROP TABLE

Section 4.5.1.1, “`mysql` Options”

Section 13.1.6, “`ALTER TABLE` Syntax”  
Section 13.1.15, “`CREATE TRIGGER` Syntax”  
Section 13.1.23, “`DROP TABLE` Syntax”  
Section 14.9.3, “`FEDERATED` Storage Engine Notes and Tips”  
Section 14.2.6, “`InnoDB` Startup Options and System Variables”  
Section 13.3.5, “`LOCK TABLES` and `UNLOCK TABLES` Syntax”  
Section 14.8.2, “`MERGE` Table Problems”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 13.6.7.5, “`SIGNAL` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 13.1.27, “`TRUNCATE TABLE` Syntax”  
Section 13.7.3.4, “`UNINSTALL PLUGIN` Syntax”  
Section 14.2.3.3, “Consistent Nonlocking Reads”  
Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 22.2.3.2, “Full-Text Parser Plugins”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 8.5.6, “Optimizing `InnoDB` DDL Operations”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section E.5, “Restrictions on Views”  
Section 13.6.7.6, “Scope Rules for Handlers”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4.2, “Setting The Binary Log Format”  
Section 5.4.1.2, “Specifying the Location of a Tablespace”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.4, “The `MEMORY` Storage Engine”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 14.2.4.7, “Troubleshooting `InnoDB` Data Dictionary Operations”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”

## DROP TABLE IF EXISTS

Section 16.4.1.7, “Replication of `DROP ... IF EXISTS` Statements”

## DROP TEMPORARY TABLE

Section 16.1.3.4, “Restrictions on Replication with GTIDs”

## DROP TEMPORARY TABLE IF EXISTS

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

---

## DROP TRIGGER

- Section 13.1.24, “`DROP TRIGGER` Syntax”
- Section 4.5.4, “`mysqldump` — A Database Backup Program”
- Section 16.4.1.10, “Replication of Invoked Features”
- Section 13.3.3, “Statements That Cause an Implicit Commit”
- Section 18.3.1, “Trigger Syntax”

## DROP USER

- Section 13.7.1.3, “`DROP USER` Syntax”
- Section 13.7.6.3, “`FLUSH` Syntax”
- Section 13.7.1.4, “`GRANT` Syntax”
- Section 13.7.1.6, “`REVOKE` Syntax”
- Section 8.11.4.1, “How MySQL Uses Memory”
- Section 12.14, “Information Functions”
- Section 6.2.1, “Privileges Provided by MySQL”
- Section 6.3.3, “Removing User Accounts”
- Section 16.4.1.6, “Replication of `CURRENT_USER()`”
- Section 13.3.3, “Statements That Cause an Implicit Commit”
- Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## DROP USER 'x'@'localhost'

- Section 22.2.4.9.3, “Using the Authentication Plugins”

## DROP VIEW

- Section 13.1.25, “`DROP VIEW` Syntax”
- Section E.5, “Restrictions on Views”
- Section 13.3.3, “Statements That Cause an Implicit Commit”
- Section 13.3.5.3, “Table-Locking Restrictions and Conditions”
- Section 18.5.1, “View Syntax”

## DROP VIEW IF EXISTS

- Section 16.4.1.7, “Replication of `DROP ... IF EXISTS` Statements”

## E

[index top [4013]]

## EXECUTE

- Section 13.2.1, “`CALL` Syntax”
- Section 13.5.2, “`EXECUTE` Syntax”
- Section 13.5.1, “`PREPARE` Syntax”
- Section 21.8.19, “C API Support for Prepared `CALL` Statements”
- Section E.1, “Restrictions on Stored Programs”
- Section E.5, “Restrictions on Views”
- Section 5.1.6, “Server Status Variables”
- Section 13.5, “SQL Syntax for Prepared Statements”

## EXPLAIN

- Section 13.1.6, “`ALTER TABLE` Syntax”

- Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”
- Section 13.8.1, “`DESCRIBE` Syntax”
- Section 8.13.15, “`DISTINCT` Optimization”
- Section 8.8.3, “`EXPLAIN EXTENDED` Output Format”
- Section 8.8.2, “`EXPLAIN` Output Format”
- Section 13.8.2, “`EXPLAIN` Syntax”
- Chapter 19, *INFORMATION\_SCHEMA* Tables
- Section 8.13.6, “`IS NULL` Optimization”
- Section 21.8.10.28, “`mysql_stmt_store_result()`”
- Section 21.8.6.70, “`mysql_store_result()`”
- Section 21.8.6.72, “`mysql_use_result()`”
- Section 8.13.13, “`ORDER BY` Optimization”
- Section 13.2.9, “`SELECT` Syntax”
- Section 13.7.5.41, “`SHOW WARNINGS` Syntax”
- Section 8.13.12.3, “Batched Key Access Joins”
- Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”
- Section 21.8.4, “C API Data Structures”
- Section 21.8.5, “C API Function Overview”
- Section 22.4.1, “Debugging a MySQL Server”
- Section 8.13.3, “Engine Condition Pushdown Optimization”
- Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”
- Section 8.2.1.4, “How to Avoid Full Table Scans”
- Section 8.13.4, “Index Condition Pushdown Optimization”
- Section 13.2.9.3, “Index Hint Syntax”
- Section 8.13.2, “Index Merge Optimization”
- Section 8.13, “Internal Details of MySQL Optimizations”
- Section 8.13.14.1, “Loose Index Scan”
- Section 8.13.11, “Multi-Range Read Optimization”
- Chapter 20, *MySQL Performance Schema*
- Section 17.3.5, “Obtaining Information About Partitions”
- Section C.5.6, “Optimizer-Related Issues”
- Section 8.2.4, “Optimizing `INFORMATION_SCHEMA` Queries”
- Section 8.8.1, “Optimizing Queries with `EXPLAIN`”
- Section 13.2.10.10, “Optimizing Subqueries”
- Section 8.13.16.3, “Optimizing Subqueries in the `FROM` Clause (Derived Tables)”
- Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”
- Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”
- Section 8.13.16.2, “Optimizing Subqueries with Subquery Materialization”
- Section 8.2.5, “Other Optimization Tips”
- Section E.1, “Restrictions on Stored Programs”
- Section C.5.5.7, “Solving Problems with No Matching Rows”
- Section 8.2.1.1, “Speed of `SELECT` Statements”
- Section 13.2.10.8, “Subqueries in the `FROM` Clause”
- Section 8.13.2.1, “The Index Merge Intersection Access Algorithm”
- Section 8.13.1.2, “The Range Access Method for Multiple-Part Indexes”

---

Section 8.8, “Understanding the Query Execution Plan”  
Section 8.13.5, “Use of Index Extensions”  
Section 12.18.6.2, “Using a Spatial Index”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in `mysqld`”  
Section 8.3.6, “Verifying Index Usage”  
Section 1.4, “What Is New in MySQL 5.6”

## **EXPLAIN ... SELECT**

Section 17.3.5, “Obtaining Information About Partitions”

## **EXPLAIN EXTENDED**

Section 8.8.3, “`EXPLAIN EXTENDED` Output Format”  
Section 8.8.2, “`EXPLAIN` Output Format”  
Section 13.8.2, “`EXPLAIN` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 8.13.3, “Engine Condition Pushdown Optimization”  
Section 8.8.1, “Optimizing Queries with `EXPLAIN`”  
Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”  
Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”  
Section 8.13.16.2, “Optimizing Subqueries with Subquery Materialization”

## **EXPLAIN PARTITIONS**

Section 13.8.2, “`EXPLAIN` Syntax”  
Section 17.3.5, “Obtaining Information About Partitions”  
Section 8.8.1, “Optimizing Queries with `EXPLAIN`”

## **EXPLAIN PARTITIONS SELECT**

Section 17.3.5, “Obtaining Information About Partitions”

## **EXPLAIN PARTITIONS SELECT**

### **COUNT( )**

Section 17.2.1, “`RANGE` Partitioning”

## **EXPLAIN SELECT**

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 1.7, “How to Report Bugs or Problems”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.3.5, “Obtaining Information About Partitions”  
Section 8.13.16.3, “Optimizing Subqueries in the `FROM` Clause (Derived Tables)”  
Section 13.2.10.8, “Subqueries in the `FROM` Clause”

## **EXPLAIN SELECT ... ORDER BY**

Section 8.13.13, “`ORDER BY` Optimization”

## **EXPLAIN tbl\_name**

Section 8.8.1, “Optimizing Queries with `EXPLAIN`”

## **F**

[index top [4013]]

## **FETCH**

Section 13.6.6.2, “Cursor `DECLARE` Syntax”  
Section 13.6.6.3, “Cursor `FETCH` Syntax”  
Section E.1, “Restrictions on Stored Programs”

## **FETCH ... INTO var\_list**

Section 13.6.4, “Variables in Stored Programs”

## **FLUSH**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.6.6, “`RESET` Syntax”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section C.5.4.1.3, “Resetting the Root Password: Generic Instructions”  
Section C.5.4.1.2, “Resetting the Root Password: Unix Systems”  
Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”  
Section E.1, “Restrictions on Stored Programs”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## **FLUSH BINARY LOGS**

Section 5.2.6, “Server Log Maintenance”

## **FLUSH DES\_KEY\_FILE**

Section 12.13, “Encryption and Compression Functions”

## **FLUSH HOSTS**

Section C.5.2.6, “Host ‘`host_name`’ is blocked”  
Section 21.8.6.56, “`mysql_refresh()`”  
Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”  
Section 5.1.4, “Server System Variables”  
Section 20.9.9.1, “The `host_cache` Table”

## **FLUSH LOGS**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 21.8.6.56, “`mysql_refresh()`”  
Section 7.3.3, “Backup Strategy Summary”  
Section 7.2, “Database Backup Methods”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 5.2, “MySQL Server Logs”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

---

Section 5.2.6, “Server Log Maintenance”  
Section 5.1.11, “Server Response to Signals”  
Section 5.1.6, “Server Status Variables”  
Section 5.2.2, “The Error Log”  
Section 16.2.2.1, “The Slave Relay Log”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## **FLUSH MASTER**

Section 16.4.1.13, “Replication and `FLUSH`”

## **FLUSH PRIVILEGES**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 21.8.6.56, “`mysql_refresh()`”  
Section 21.8.6.57, “`mysql_reload()`”  
Section 6.3.2, “Adding User Accounts”  
Section 6.3.5, “Assigning Account Passwords”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section 5.1.3, “Server Command Options”  
Section 5.1.11, “Server Response to Signals”  
Section 6.3.4, “Setting Account Resource Limits”  
Section 1.2, “Typographical and Syntax Conventions”  
Section 6.2.6, “When Privilege Changes Take Effect”

## **FLUSH QUERY CACHE**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 8.9.3.4, “Query Cache Status and Maintenance”

## **FLUSH SLAVE**

Section 16.4.1.13, “Replication and `FLUSH`”

## **FLUSH STATUS**

Section 5.1.6, “Server Status Variables”  
Section 8.13.5, “Use of Index Extensions”

## **FLUSH TABLE**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”  
Section 8.13.5, “Use of Index Extensions”

## **FLUSH TABLES**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.2.4, “`HANDLER` Syntax”  
Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 14.8.2, “`MERGE` Table Problems”  
Section 21.8.6.56, “`mysql_refresh()`”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 5.4.1.3, “Copying Tablespaces to Another Server (Transportable Tablespaces)”

Section 4.6.3, “`myisamchk` — MyISAM Table-Maintenance Utility”  
Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”  
Section 8.12.5.2, “General Thread States”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”  
Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”  
Section 8.9.3.4, “Query Cache Status and Maintenance”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

## **FLUSH TABLES ... FOR EXPORT**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 5.4.1.2, “Specifying the Location of a Tablespace”

## **FLUSH TABLES `tbl_list` WITH READ LOCK**

Section 4.6.10, “`mysqlhotcopy` — A Database Backup Program”

## **FLUSH TABLES `tbl_name` ... FOR EXPORT**

Section 13.7.6.3, “`FLUSH` Syntax”

## **FLUSH TABLES `tbl_name` ... WITH READ LOCK**

Section 13.7.6.3, “`FLUSH` Syntax”

## **FLUSH TABLES `tbl_name` WITH READ LOCK**

Section 13.2.4, “`HANDLER` Syntax”

## **FLUSH TABLES WITH READ LOCK**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.3.5, “`LOCK TABLES` and `UNLOCK TABLES` Syntax”  
Section 13.3.1, “`START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax”  
Section 16.1.1.5, “Creating a Data Snapshot Using `mysqldump`”  
Section 7.2, “Database Backup Methods”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 8.12.5.2, “General Thread States”

---

Section 13.3.5.1, “Interaction of Table Locking and Transactions”  
Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## **FLUSH USER RESOURCES**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 6.3.4, “Setting Account Resource Limits”

## **G**

[index top [4013]]

## **GET DIAGNOSTICS**

Section 13.6.7.3, “`GET DIAGNOSTICS` Syntax”  
Section 13.6.7.4, “`RESIGNAL` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 13.6.7.5, “`SIGNAL` Syntax”  
Section 13.6.7, “Condition Handling”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section E.2, “Restrictions on Condition Handling”  
Section E.1, “Restrictions on Stored Programs”  
Section 5.1.4, “Server System Variables”  
Section 13.6.7.5.1, “Signal Condition Information Items”  
Section 1.4, “What Is New in MySQL 5.6”

## **GRANT**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.1.15, “`CREATE TRIGGER` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.1.6, “`REVOKE` Syntax”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 13.7.5.22, “`SHOW GRANTS` Syntax”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 6.3.2, “Adding User Accounts”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 6.3.5, “Assigning Account Passwords”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 5.1.9.3, “Connecting Using the IPv6 Local Host Address”  
Section 16.1.1.3, “Creating a User for Replication”  
Section 6.1.2.1, “End-User Guidelines for Password Security”  
Section 8.11.4.1, “How MySQL Uses Memory”

Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”  
Section 12.14, “Information Functions”  
Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”  
Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”  
Section 5.1.9, “IPv6 Support”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 8.2.3, “Optimizing Database Privileges”  
Section 6.1.2.4, “Password Hashing in MySQL”  
Section 22.2.3.8, “Password-Validation Plugins”  
Section 6.1.2.3, “Passwords and Logging”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 2.10.1.1, “Problems Running `mysql_install_db`”  
Section 6.3.8, “Proxy Users”  
Section 16.4.1.13, “Replication and `FLUSH`”  
Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 16.4.1.22, “Replication of the `mysql` System Database”  
Section 6.1.1, “Security Guidelines”  
Section 5.1.3, “Server Command Options”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 6.3.4, “Setting Account Resource Limits”  
Section 6.2.3, “Specifying Account Names”  
Section 6.3.9.4, “SSL Command Options”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.5.1, “Support for Read-Only Media”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”  
Section 6.2, “The MySQL Access Privilege System”  
Section 6.1.2.6, “The Password Validation Plugin”  
Section 6.3.1, “User Names and Passwords”  
Section 6.3.9.3, “Using SSL Connections”  
Section 6.3.9, “Using SSL for Secure Connections”  
Section 6.3.7.3.2, “Using the PAM Authentication Plugin”  
Section 6.3.7.4.2, “Using the Windows Authentication Plugin”  
Section 1.4, “What Is New in MySQL 5.6”  
Section 6.2.6, “When Privilege Changes Take Effect”

## **GRANT ALL**

Section 13.7.1.4, “`GRANT` Syntax”

## **GRANT EVENT**

Section 18.4.6, “The Event Scheduler and MySQL Privileges”



---

## GRANT USAGE

- Section 13.7.1.4, “GRANT Syntax”
- Section 6.3.5, “Assigning Account Passwords”
- Section 6.3.4, “Setting Account Resource Limits”

## H

[index top [4013]]

### HANDLER

- Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”
- Section 13.7.6.3, “FLUSH Syntax”
- Section 13.2.4, “HANDLER Syntax”
- Section 21.8.6.3, “mysql\_change\_user()”
- Section 21.8.15, “Controlling Automatic Reconnection Behavior”
- Section 1.8, “MySQL Standards Compliance”
- Section 5.1.4, “Server System Variables”

### HANDLER ... CLOSE

- Section 13.7.5.25, “SHOW OPEN TABLES Syntax”

### HANDLER ... OPEN

- Section 13.7.5.25, “SHOW OPEN TABLES Syntax”

### HANDLER ... READ

- Section E.1, “Restrictions on Stored Programs”

### HANDLER OPEN

- Section 13.2.4, “HANDLER Syntax”
- Section 13.1.27, “TRUNCATE TABLE Syntax”

### HELP

- Section 13.8.3, “HELP Syntax”
- Section 5.1.10, “Server-Side Help”
- Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

## I

[index top [4013]]

### IF

- Section 13.6.7.2, “DECLARE ... HANDLER Syntax”
- Section 13.6.5.2, “IF Syntax”
- Section 8.9.4, “Caching of Prepared Statements and Stored Programs”
- Section 12.4, “Control Flow Functions”
- Section 13.6.5, “Flow Control Statements”

### INSERT

- Section 4.5.1.1, “mysql Options”
- Section 13.1.11, “CREATE INDEX Syntax”
- Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”

- Section 13.1.14.1, “CREATE TABLE ... SELECT Syntax”
- Section 13.1.15, “CREATE TRIGGER Syntax”
- Section 13.1.16, “CREATE VIEW Syntax”
- Section 13.6.7.2, “DECLARE ... HANDLER Syntax”
- Section 13.2.2, “DELETE Syntax”
- Section 8.8.3, “EXPLAIN EXTENDED Output Format”
- Section 8.8.2, “EXPLAIN Output Format”
- Section 13.8.2, “EXPLAIN Syntax”
- Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”
- Section 13.7.1.4, “GRANT Syntax”
- Chapter 19, *INFORMATION\_SCHEMA Tables*
- Section 14.2.3.14.1, “InnoDB Error Codes”
- Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”
- Section 14.2.6, “InnoDB Startup Options and System Variables”
- Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”
- Section 13.2.5.1, “INSERT ... SELECT Syntax”
- Section 13.2.5.2, “INSERT DELAYED Syntax”
- Section 13.2.5, “INSERT Syntax”
- Section 17.2.2, “LIST Partitioning”
- Section 13.2.6, “LOAD DATA INFILE Syntax”
- Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”
- Section 14.8.2, “MERGE Table Problems”
- Section C.5.2.9, “MySQL server has gone away”
- Section 21.8.6.1, “mysql\_affected\_rows()”
- Section 21.8.6.37, “mysql\_insert\_id()”
- Section 21.8.6.48, “mysql\_num\_rows()”
- Section 21.8.10.10, “mysql\_stmt\_execute()”
- Section 21.8.10.13, “mysql\_stmt\_field\_count()”
- Section 21.8.10.16, “mysql\_stmt\_insert\_id()”
- Section 21.8.10.18, “mysql\_stmt\_num\_rows()”
- Section 21.8.10.21, “mysql\_stmt\_prepare()”
- Section 21.8.6.70, “mysql\_store\_result()”
- Section 13.7.2.4, “OPTIMIZE TABLE Syntax”
- Section 1.8.6.1, “PRIMARY KEY and UNIQUE Index Constraints”
- Section 17.2.1, “RANGE Partitioning”
- Section 13.2.8, “REPLACE Syntax”
- Section 13.7.5.28, “SHOW PROCEDURE CODE Syntax”
- Section 13.7.5.41, “SHOW WARNINGS Syntax”
- Section 13.2.11, “UPDATE Syntax”
- Section 6.2.5, “Access Control, Stage 2: Request Verification”
- Section 6.3.2, “Adding User Accounts”
- Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”
- Section 7.1, “Backup and Recovery Types”
- Section 18.7, “Binary Logging of Stored Programs”
- Section 8.5.4, “Bulk Data Loading for InnoDB Tables”
- Section 8.6.2, “Bulk Data Loading for MyISAM Tables”
- Section 21.8.5, “C API Function Overview”
- Section 21.8.9, “C API Prepared Statement Function Overview”

---

[Section 21.8.16, “C API Support for Multiple Statement Execution”](#)  
[Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#)  
[Section 6.2.7, “Causes of Access-Denied Errors”](#)  
[Section 10.1.13, “Column Character Set Conversion”](#)  
[Section 8.10.3, “Concurrent Inserts”](#)  
[Section 5.4.4.2, “Configurable `InnoDB` Auto-Increment Locking”](#)  
[Section 1.8.6.3, “Constraints on Invalid Data”](#)  
[Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”](#)  
[Section 14.9.2.1, “Creating a `FEDERATED` Table Using `CONNECTION`”](#)  
[Section 11.5, “Data Type Default Values”](#)  
[Section 11.1.2, “Date and Time Type Overview”](#)  
[Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”](#)  
[Section 4.5.4, “`mysqldump` — A Database Backup Program”](#)  
[Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”](#)  
[Section 7.3.1, “Establishing a Backup Policy”](#)  
[Section 12.19.3, “Expression Handling”](#)  
[Section 1.8.5.4, “Foreign Key Differences”](#)  
[Section 8.12.5.2, “General Thread States”](#)  
[Section 5.4.2, “Grouping DML Operations with Transactions”](#)  
[Section 8.9.3.1, “How the Query Cache Operates”](#)  
[Section 21.8.14.3, “How to Get the Unique ID for the Last Inserted Row”](#)  
[Section 12.14, “Information Functions”](#)  
[Section 8.10.1, “Internal Locking Methods”](#)  
[Section 3.3.3, “Loading Data into a Table”](#)  
[Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”](#)  
[Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”](#)  
[Section 17.3.1, “Management of `RANGE` and `LIST` Partitions”](#)  
[Section 12.16, “Miscellaneous Functions”](#)  
[Section 1.8.4, “MySQL Extensions to Standard SQL”](#)  
[Section 8.6.1, “Optimizing `MyISAM` Queries”](#)  
[Section 8.2.2, “Optimizing DML Statements”](#)  
[Section 8.8.1, “Optimizing Queries with `EXPLAIN`”](#)  
[Section 8.2.5, “Other Optimization Tips”](#)  
[Section 11.2.6, “Out-of-Range and Overflow Handling”](#)  
[Section 5.5.1, “Overview of Online DDL”](#)  
[Section 17.1, “Overview of Partitioning in MySQL”](#)  
[Section 17.5, “Partition Selection”](#)  
[Section 17.6.4, “Partitioning and Locking”](#)  
[Section 12.18.4.4, “Populating Spatial Columns”](#)  
[Section 6.2.1, “Privileges Provided by MySQL”](#)  
[Section 8.12.5.4, “Query Cache Thread States”](#)  
[Section 16.4.1.1, “Replication and `AUTO\_INCREMENT`”](#)  
[Section 16.4.1.26, “Replication and Server SQL Mode”](#)  
[Section 16.4.1.14, “Replication and System Functions”](#)  
[Section 16.4.1.32, “Replication and Variables”](#)  
[Section 16.1.4.2, “Replication Master Options and Variables”](#)  
[Section 16.2.3.3, “Replication Rule Application”](#)  
[Section 17.6, “Restrictions and Limitations on Partitioning”](#)  
[Section E.1, “Restrictions on Stored Programs”](#)  
[Section 5.7.1.7, “Row-Level Probes”](#)  
[Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”](#)  
[Section 5.1.3, “Server Command Options”](#)  
[Section 5.1.7, “Server SQL Modes”](#)  
[Section 5.1.4, “Server System Variables”](#)  
[Section 16.4.1.25, “Slave Errors During Replication”](#)  
[Section 8.2.2.1, “Speed of `INSERT` Statements”](#)  
[Section 14.2.4.6, “Starting `InnoDB` on a Corrupted Database”](#)  
[Section 5.7.1.12, “Statement Probes”](#)  
[Section 13.2.10, “Subquery Syntax”](#)  
[Section 8.10.2, “Table Locking Issues”](#)  
[Section 10.1.7.6, “The `\_bin` and `binary` Collations”](#)  
[Section 14.6, “The `ARCHIVE` Storage Engine”](#)  
[Section 14.7, “The `BLACKHOLE` Storage Engine”](#)  
[Section 19.29, “The `INFORMATION\_SCHEMA VIEWS` Table”](#)  
[Section 14.8, “The `MERGE` Storage Engine”](#)  
[Section 14.3, “The `MyISAM` Storage Engine”](#)  
[Section 5.2.4, “The Binary Log”](#)  
[Section 6.2, “The MySQL Access Privilege System”](#)  
[Section 8.9.3, “The MySQL Query Cache”](#)  
[Section 5.1.12, “The Shutdown Process”](#)  
[Section 5.4.4.1, “Traditional `InnoDB` Auto-Increment Locking”](#)  
[Section 18.3.1, “Trigger Syntax”](#)  
[Section 18.5.3, “Updatable and Insertable Views”](#)  
[Section 3.6.9, “Using `AUTO\_INCREMENT`”](#)  
[Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”](#)  
[Section 18.3, “Using Triggers”](#)  
[Section 1.4, “What Is New in MySQL 5.6”](#)  
[Section 21.8.14.2, “What Results You Can Get from a Query”](#)  
[Section 6.2.6, “When Privilege Changes Take Effect”](#)  
[Section 21.8.14.1, “Why `mysql\_store\_result\(\)` Sometimes Returns `NULL` After `mysql\_query\(\)` Returns Success”](#)  
  
**`INSERT ... ()`**  
[Section 5.7.1.12, “Statement Probes”](#)  
  
**`INSERT ... ON DUPLICATE KEY UPDATE`**  
[Section 14.9.3, “`FEDERATED` Storage Engine Notes and Tips”](#)  
[Section 13.2.5.3, “`INSERT ... ON DUPLICATE KEY UPDATE` Syntax”](#)  
[Section 13.2.5.2, “`INSERT DELAYED` Syntax”](#)  
[Section 13.2.5, “`INSERT` Syntax”](#)  
[Section 14.8.2, “`MERGE` Table Problems”](#)

---

---

Section 21.8.6.1, “`mysql_affected_rows()`”  
Section 21.8.6.37, “`mysql_insert_id()`”  
Section 5.4.4.2, “Configurable `InnoDB` Auto-Increment Locking”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.14, “Information Functions”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”  
Section 12.16, “Miscellaneous Functions”  
Section 17.6.4, “Partitioning and Locking”

## **INSERT ... SELECT**

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”  
Section 13.2.5.3, “`INSERT ... ON DUPLICATE KEY UPDATE` Syntax”  
Section 13.2.5.1, “`INSERT ... SELECT` Syntax”  
Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 13.2.5, “`INSERT` Syntax”  
Section 21.8.6.37, “`mysql_insert_id()`”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 8.10.3, “Concurrent Inserts”  
Section 5.4.4.2, “Configurable `InnoDB` Auto-Increment Locking”  
Section C.5.8, “Known Issues in MySQL”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”  
Section 17.5, “Partition Selection”  
Section 17.6.4, “Partitioning and Locking”  
Section 16.4.1.15, “Replication and `LIMIT`”  
Section 5.1.4, “Server System Variables”  
Section 5.7.1.12, “Statement Probes”  
Section 5.2.4, “The Binary Log”

## **INSERT ... SELECT ON DUPLICATE KEY UPDATE**

Section 13.2.5.3, “`INSERT ... ON DUPLICATE KEY UPDATE` Syntax”  
Section 13.2.5.1, “`INSERT ... SELECT` Syntax”

## **INSERT ... SET**

Section 13.2.5, “`INSERT` Syntax”

## **INSERT ... VALUES**

Section 13.2.5, “`INSERT` Syntax”  
Section 21.8.6.35, “`mysql_info()`”

## **INSERT DELAYED**

Section 13.1.6, “`ALTER TABLE` Syntax”  
Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 13.2.5, “`INSERT` Syntax”  
Section 13.7.6.4, “`KILL` Syntax”  
Section 14.8.2, “`MERGE` Table Problems”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”

Section 8.12.5.3, “Delayed-Insert Thread States”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section C.5.8, “Known Issues in MySQL”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 22.1.1, “MySQL Threads”  
Section 8.6.1, “Optimizing `MyISAM` Queries”  
Section 17.6.4, “Partitioning and Locking”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section E.1, “Restrictions on Stored Programs”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 8.10.2, “Table Locking Issues”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.4, “The `MEMORY` Storage Engine”  
Section 20.9.9.3, “The `threads` Table”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”  
Section 18.5.3, “Updatable and Insertable Views”

## **INSERT IGNORE**

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”  
Section 1.8.6.4, “`ENUM` and `SET` Constraints”  
Section 13.2.5, “`INSERT` Syntax”  
Section 1.8.6.3, “Constraints on Invalid Data”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 12.14, “Information Functions”  
Section 5.1.7, “Server SQL Modes”

## **INSERT IGNORE ... SELECT**

Section 13.2.5.1, “`INSERT ... SELECT` Syntax”

## **INSERT INTO ... SELECT**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.2.5, “`INSERT` Syntax”  
Section 1.8.5.1, “`SELECT INTO TABLE` Differences”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 14.2.3.3, “Consistent Nonlocking Reads”  
Section 1.8.6.3, “Constraints on Invalid Data”  
Section 16.4.1.4, “Replication of `CREATE TABLE ... SELECT` Statements”  
Section 14.4, “The `MEMORY` Storage Engine”

## **INSERT INTO ... SELECT ...**

Section 14.9.3, “`FEDERATED` Storage Engine Notes and Tips”  
Section 21.8.6.35, “`mysql_info()`”  
Section 21.8.14.2, “What Results You Can Get from a Query”



---

## **INSERT INTO ... SELECT FROM memory\_table**

Section 16.4.1.20, “Replication and MEMORY Tables”

## **INSERT LOW\_PRIORITY**

Section 8.2.5, “Other Optimization Tips”

## **INSTALL PLUGIN**

Section 13.7.6.3, “FLUSH Syntax”

Section 13.7.3.3, “INSTALL PLUGIN Syntax”

Section 13.7.5.26, “SHOW PLUGINS Syntax”

Section 6.3.11.4, “Audit Log Plugin Options and Variables”

Section 4.4.4, “mysql\_plugin — Configure MySQL Server Plugins”

Section 8.11.4.1, “How MySQL Uses Memory”

Section 5.1.8.1, “Installing and Uninstalling Plugins”

Section 6.3.11.1, “Installing the Audit Log Plugin”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 5.1.8.2, “Obtaining Server Plugin Information”

Section 6.1.2.6.1, “Password Validation Plugin Installation”

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

Section 6.3.7, “Pluggable Authentication”

Section 14.12.1, “Pluggable Storage Engine Architecture”

Section 22.2.2, “Plugin API Components”

Section 22.2.4.2, “Plugin Data Structures”

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 5.1.3, “Server Command Options”

Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”

Section 19.15, “The INFORMATION\_SCHEMA PLUGINS Table”

Section 22.2, “The MySQL Plugin API”

Section 22.2.4.6, “Writing INFORMATION\_SCHEMA Plugins”

Section 22.2.4.8, “Writing Audit Plugins”

Section 22.2.4.5, “Writing Daemon Plugins”

Section 22.2.4.4, “Writing Full-Text Parser Plugins”

Section 22.2.4.10, “Writing Password-Validation Plugins”

Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

## **ITERATE**

Section 13.6.7.2, “DECLARE ... HANDLER Syntax”

Section 13.6.5.3, “ITERATE Syntax”

Section 13.6.5, “Flow Control Statements”

Section 13.6.2, “Statement Label Syntax”

## **K**

[index top [4013]]

## **KILL**

Section 13.7.1.4, “GRANT Syntax”

Section 13.7.6.4, “KILL Syntax”

Section C.5.2.9, “MySQL server has gone away”

Section 21.8.6.38, “mysql\_kill()”

Section 13.7.5.30, “SHOW PROCESSLIST Syntax”

Section 4.6.16, “mysql\_zap — Kill Processes That Match a Pattern”

Section 8.12.5.2, “General Thread States”

Section 6.2.1, “Privileges Provided by MySQL”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

## **KILL CONNECTION**

Section 13.7.6.4, “KILL Syntax”

Section 13.4.2.6, “STOP SLAVE Syntax”

Section 5.1.12, “The Shutdown Process”

## **KILL QUERY**

Section 13.7.6.4, “KILL Syntax”

Section 13.4.2.6, “STOP SLAVE Syntax”

Section 5.1.12, “The Shutdown Process”

## **L**

[index top [4013]]

## **LEAVE**

Section 13.6.7.2, “DECLARE ... HANDLER Syntax”

Section 13.6.5.4, “LEAVE Syntax”

Section 13.6.5.5, “LOOP Syntax”

Section 13.6.5.7, “RETURN Syntax”

Section 13.6.5, “Flow Control Statements”

Section E.1, “Restrictions on Stored Programs”

Section 13.6.2, “Statement Label Syntax”

## **LOAD DATA**

Section 13.1.15, “CREATE TRIGGER Syntax”

Section 13.2.6, “LOAD DATA INFILE Syntax”

Section 13.2.7, “LOAD XML Syntax”

Section 8.10.3, “Concurrent Inserts”

Section 5.4.4.2, “Configurable InnoDB Auto-Increment Locking”

Section C.5.8, “Known Issues in MySQL”

Section 3.3.3, “Loading Data into a Table”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 17.1, “Overview of Partitioning in MySQL”

Section 17.5, “Partition Selection”

Section 17.6.4, “Partitioning and Locking”

Section 17.6, “Restrictions and Limitations on Partitioning”

Section E.1, “Restrictions on Stored Programs”

---

Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”  
Section 3.3.4.1, “Selecting All Data”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 11.4.4, “The `ENUM` Type”  
Section 9.4, “User-Defined Variables”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”  
Section 1.4, “What Is New in MySQL 5.6”

## **LOAD DATA INFILE**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8.1, “`mysqlbinlog` Hex Dump Format”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 14.3.1, “`MyISAM` Startup Options”  
Section 9.1.7, “`NULL` Values”  
Section 13.2.9.1, “`SELECT ... INTO` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.5, “Audit Log Plugin Restrictions”  
Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 7.1, “Backup and Recovery Types”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 8.10.3, “Concurrent Inserts”  
Section 7.2, “Database Backup Methods”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 4.5.5, “`mysqlimport` — A Data Import Program”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section C.5.4.3, “How MySQL Handles a Full Disk”  
Section 12.14, “Information Functions”  
Section C.5.8, “Known Issues in MySQL”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 8.2.5, “Other Optimization Tips”  
Section 11.2.6, “Out-of-Range and Overflow Handling”  
Section 4.1, “Overview of MySQL Programs”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section C.5.5.3, “Problems with `NULL` Values”  
Section 7.4.4, “Reloading Delimited-Text Format Backups”  
Section 16.4.1.16, “Replication and `LOAD DATA INFILE`”  
Section 16.4.2, “Replication Compatibility Between MySQL Versions”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 8.12.5.7, “Replication Slave SQL Thread States”

Section E.7, “Restrictions on Character Sets”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 8.2.2.1, “Speed of `INSERT` Statements”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 13.2.10, “Subquery Syntax”  
Section 14.4, “The `MEMORY` Storage Engine”  
Section 6.2, “The MySQL Access Privilege System”  
Section 13.2.10.1, “The Subquery as Scalar Operand”  
Section C.5.4.4, “Where MySQL Stores Temporary Files”  
Section E.10.6, “Windows Platform Limitations”

## **LOAD DATA INFILE ...**

Section 21.8.6.35, “`mysql_info()`”  
Section 21.8.14.2, “What Results You Can Get from a Query”

## **LOAD DATA LOCAL**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 21.8.6.49, “`mysql_options()`”  
Section 21.8.6.53, “`mysql_real_connect()`”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 6.1.6, “Security Issues with `LOAD DATA LOCAL`”

## **LOAD DATA LOCAL INFILE**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 21.8.6.64, “`mysql_set_local_infile_handler()`”  
Section 21.8.5, “C API Function Overview”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **LOAD INDEX INTO CACHE**

Section 13.7.6.2, “`CACHE INDEX` Syntax”  
Section 13.7.6.5, “`LOAD INDEX INTO CACHE` Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 8.9.2.4, “Index Preloading”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 13.3.3, “Statements That Cause an Implicit Commit”

## **LOAD INDEX INTO CACHE ...**

## **IGNORE LEAVES**

Section 13.7.6.5, “`LOAD INDEX INTO CACHE` Syntax”

## **LOAD XML**

Section 13.2.7, “`LOAD XML` Syntax”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section 17.5, “Partition Selection”  
Section 1.4, “What Is New in MySQL 5.6”

---

## LOAD XML INFILE

Section 13.2.7, “LOAD XML Syntax”

## LOAD XML LOCAL

Section 13.2.7, “LOAD XML Syntax”

## LOAD XML LOCAL INFILE

Section 13.2.7, “LOAD XML Syntax”

## LOCK TABLE

Section 8.10.3, “Concurrent Inserts”

Section 8.12.5.2, “General Thread States”

Section C.5.7.1, “Problems with ALTER TABLE”

## LOCK TABLES

Section 13.1.8, “CREATE DATABASE Syntax”

Section 13.1.14, “CREATE TABLE Syntax”

Section 13.1.15, “CREATE TRIGGER Syntax”

Section 13.7.6.3, “FLUSH Syntax”

Section 13.7.1.4, “GRANT Syntax”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”

Section 13.3.5.2, “LOCK TABLES and Triggers”

Section 14.8.2, “MERGE Table Problems”

Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”

Section 8.6.2, “Bulk Data Loading for MyISAM Tables”

Section 14.2.3.9, “Deadlock Detection and Rollback”

Section 4.6.10, “mysqlhotcopy — A Database Backup Program”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 14.2.3.10, “How to Cope with Deadlocks”

Section 13.3.5.1, “Interaction of Table Locking and Transactions”

Section 8.10.1, “Internal Locking Methods”

Section 14.2.7, “Limits on InnoDB Tables”

Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

Section 17.6.4, “Partitioning and Locking”

Section 6.2.1, “Privileges Provided by MySQL”

Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”

Section E.1, “Restrictions on Stored Programs”

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

Section 5.1.4, “Server System Variables”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 8.11.1, “System Factors and Startup Parameter Tuning”

Section 8.10.2, “Table Locking Issues”

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”

Section 1.8.5.3, “Transaction and Atomic Operation Differences”

## LOCK TABLES ... READ

Section 13.7.6.3, “FLUSH Syntax”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

## LOCK TABLES ... WRITE

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

## LOOP

Section 13.6.5.3, “ITERATE Syntax”

Section 13.6.5.4, “LEAVE Syntax”

Section 13.6.5.5, “LOOP Syntax”

Section 13.6.5, “Flow Control Statements”

Section 13.6.2, “Statement Label Syntax”

## O

[index top [4013]]

## OPEN

Section 13.6.6.4, “Cursor OPEN Syntax”

## OPTIMIZE TABLE

Section 4.6.3.1, “myisamchk General Options”

Section 13.2.2, “DELETE Syntax”

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 13.7.6.4, “KILL Syntax”

Section 14.8.2, “MERGE Table Problems”

Section 7.6, “MyISAM Table Maintenance and Crash Recovery”

Section 7.6.4, “MyISAM Table Optimization”

Section 13.7.2.4, “OPTIMIZE TABLE Syntax”

Section 22.4.1, “Debugging a MySQL Server”

Section 4.5.3, “mysqlcheck — A Table Maintenance Program”

Section 14.3.3.2, “Dynamic Table Characteristics”

Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

Section 8.12.5.2, “General Thread States”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section C.5.4.3, “How MySQL Handles a Full Disk”

Section C.5.8, “Known Issues in MySQL”

Section 17.3.4, “Maintenance of Partitions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 8.6.1, “Optimizing MyISAM Queries”

Section 8.2.5, “Other Optimization Tips”

Section 6.2.1, “Privileges Provided by MySQL”

Section 16.4.1.13, “Replication and FLUSH”

---

Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 7.6.5, “Setting Up a MyISAM Table Maintenance Schedule”  
Section 8.2.2.2, “Speed of UPDATE Statements”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.3.3.1, “Static (Fixed-Length) Table Characteristics”  
Section 14.6, “The ARCHIVE Storage Engine”  
Section 19.30.23, “The INFORMATION\_SCHEMA INNODB\_FT\_BEING\_DELETED Table”  
Section 19.30.22, “The INFORMATION\_SCHEMA INNODB\_FT\_DELETED Table”  
Section 19.30.21, “The INFORMATION\_SCHEMA INNODB\_FT\_INDEX\_CACHE Table”  
Section 5.1.12, “The Shutdown Process”  
Section 5.2.5, “The Slow Query Log”  
Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”

## P

[index top [4013]]

### PREPARE

Section 13.2.1, “CALL Syntax”  
Section 13.5.3, “DEALLOCATE PREPARE Syntax”  
Section 13.5.2, “EXECUTE Syntax”  
Section 13.5.1, “PREPARE Syntax”  
Section 21.8.19, “C API Support for Prepared CALL Statements”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 8.10.4, “Metadata Locking”  
Section E.1, “Restrictions on Stored Programs”  
Section E.5, “Restrictions on Views”  
Section 5.1.6, “Server Status Variables”  
Section 13.5, “SQL Syntax for Prepared Statements”

### PURGE BINARY LOGS

Section 13.7.1.4, “GRANT Syntax”  
Section 13.4.1.1, “PURGE BINARY LOGS Syntax”  
Section 13.4.1.2, “RESET MASTER Syntax”  
Section 4.5.4, “mysqldump — A Database Backup Program”  
Section 7.3.1, “Establishing a Backup Policy”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.2.6, “Server Log Maintenance”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4, “The Binary Log”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## R

[index top [4013]]

### RELEASE SAVEPOINT

Section 13.3.4, “SAVEPOINT, ROLLBACK TO SAVEPOINT, and RELEASE SAVEPOINT Syntax”

### RENAME TABLE

Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.2.2, “DELETE Syntax”  
Section 13.1.26, “RENAME TABLE Syntax”  
Section 8.12.5.2, “General Thread States”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 14.2.2.5, “Moving or Copying InnoDB Tables to Another Machine”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 8.11.3.1.2, “Using Symbolic Links for MyISAM Tables on Unix”

### RENAME USER

Section 13.7.1.4, “GRANT Syntax”  
Section 13.7.1.5, “RENAME USER Syntax”  
Section 12.14, “Information Functions”  
Section 5.2.4.4, “Logging Format for Changes to mysql Database Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.6, “Replication of CURRENT\_USER( )”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”  
Section 6.2.6, “When Privilege Changes Take Effect”

### REPAIR TABLE

Section 4.6.3.1, “myisamchk General Options”  
Section 13.1.6.1, “ALTER TABLE Partition Operations”  
Section 13.1.6, “ALTER TABLE Syntax”  
Section 13.7.2.2, “CHECK TABLE Syntax”  
Section 13.7.6.4, “KILL Syntax”  
Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 14.8.2, “MERGE Table Problems”  
Section 14.3.1, “MyISAM Startup Options”  
Section 7.6, “MyISAM Table Maintenance and Crash Recovery”  
Section 13.7.2.5, “REPAIR TABLE Syntax”  
Section 11.3.4, “YEAR( 2 ) Limitations and Migrating to YEAR( 4 )”  
Section 14.3.4.1, “Corrupted MyISAM Tables”  
Section 7.2, “Database Backup Methods”  
Section 4.6.3, “myisamchk — MyISAM Table-Maintenance Utility”

---

Section 4.5.3, “`mysqlcheck` — A Table Maintenance Program”

Section 8.10.5, “External Locking”

Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

Section 8.12.5.2, “General Thread States”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section C.5.4.3, “How MySQL Handles a Full Disk”

Section 7.6.3, “How to Repair `MyISAM` Tables”

Section 1.7, “How to Report Bugs or Problems”

Section C.5.8, “Known Issues in MySQL”

Section 17.3.4, “Maintenance of Partitions”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 6.2.1, “Privileges Provided by MySQL”

Section 14.3.4.2, “Problems from Tables Not Being Closed Properly”

Section C.5.7.1, “Problems with `ALTER TABLE`”

Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”

Section 16.4.1.13, “Replication and `FLUSH`”

Section 16.4.1.17, “Replication and `REPAIR TABLE`”

Section 17.6, “Restrictions and Limitations on Partitioning”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 7.6.5, “Setting Up a `MyISAM` Table Maintenance Schedule”

Section 8.6.3, “Speed of `REPAIR TABLE` Statements”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 14.6, “The `ARCHIVE` Storage Engine”

Section 5.1.12, “The Shutdown Process”

Section 5.2.5, “The Slow Query Log”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”

## REPEAT

Section 13.6.7.2, “`DECLARE ... HANDLER` Syntax”

Section 13.6.5.3, “`ITERATE` Syntax”

Section 13.6.5.4, “`LEAVE` Syntax”

Section 13.6.5.6, “`REPEAT` Syntax”

Section 18.1, “Defining Stored Programs”

Section 13.6.5, “Flow Control Statements”

Section 13.6.2, “Statement Label Syntax”

## REPLACE

Section 13.1.14.1, “`CREATE TABLE ... SELECT` Syntax”

Section 13.1.15, “`CREATE TRIGGER` Syntax”

Section 8.8.3, “`EXPLAIN EXTENDED` Output Format”

Section 8.8.2, “`EXPLAIN` Output Format”

Section 13.8.2, “`EXPLAIN` Syntax”

Section 13.2.5.2, “`INSERT DELAYED` Syntax”

Section 13.2.5, “`INSERT` Syntax”

Section 14.8.2, “`MERGE` Table Problems”

Section C.5.2.9, “MySQL server has gone away”

Section 21.8.6.1, “`mysql_affected_rows()`”

Section 13.2.8, “`REPLACE` Syntax”

Section 13.2.11, “`UPDATE` Syntax”

Section 5.4.4.2, “Configurable `InnoDB` Auto-Increment Locking”

Section 11.5, “Data Type Default Values”

Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 12.14, “Information Functions”

Section C.5.8, “Known Issues in MySQL”

Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”

Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 8.8.1, “Optimizing Queries with `EXPLAIN`”

Section 17.1, “Overview of Partitioning in MySQL”

Section 17.5, “Partition Selection”

Section 17.6.4, “Partitioning and Locking”

Section 5.1.3, “Server Command Options”

Section 13.2.10, “Subquery Syntax”

Section 14.6, “The `ARCHIVE` Storage Engine”

Section 1.4, “What Is New in MySQL 5.6”

## REPLACE ... SELECT

Section 5.4.4.2, “Configurable `InnoDB` Auto-Increment Locking”

Section C.5.8, “Known Issues in MySQL”

## RESET

Section 13.7.6.3, “`FLUSH` Syntax”

Section 13.7.6.6, “`RESET` Syntax”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 13.3.3, “Statements That Cause an Implicit Commit”

## RESET MASTER

Section 21.8.6.56, “`mysql_refresh()`”

Section 13.4.1.2, “`RESET MASTER` Syntax”

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.3.6, “Switching Masters During Failover”

Section 5.2.4, “The Binary Log”

Section 1.4, “What Is New in MySQL 5.6”

## RESET QUERY CACHE

Section 8.12.5.4, “Query Cache Thread States”

## RESET SLAVE

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

Section 21.8.6.56, “`mysql_refresh()`”

Section 13.4.1.2, “`RESET MASTER` Syntax”

Section 13.4.2.3, “`RESET SLAVE` Syntax”

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”



---

Section 16.3.9, “Delayed Replication”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 1.4, “What Is New in MySQL 5.6”

## RESIGNAL

Section 13.6.7.1, “`DECLARE ... CONDITION` Syntax”  
Section 13.6.7.2, “`DECLARE ... HANDLER` Syntax”  
Section 13.6.7.4.1, “`RESIGNAL` Alone”  
Section 13.6.7.4.4, “`RESIGNAL` Requires an Active Handler”  
Section 13.6.7.4, “`RESIGNAL` Syntax”  
Section 13.6.7.4.3, “`RESIGNAL` with a Condition Value and Optional New Signal Information”  
Section 13.6.7.4.2, “`RESIGNAL` with New Signal Information”  
Section 13.6.7, “Condition Handling”  
Section 13.6.7.7.2, “Diagnostics Area Information Items”  
Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section E.2, “Restrictions on Condition Handling”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.6.7.6, “Scope Rules for Handlers”  
Section 13.6.7.5.1, “Signal Condition Information Items”

## RETURN

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.6.5.5, “`LOOP` Syntax”  
Section 13.6.5.7, “`RETURN` Syntax”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 13.6.5, “Flow Control Statements”  
Section E.1, “Restrictions on Stored Programs”

## REVOKE

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.1.6, “`REVOKE` Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 12.14, “Information Functions”  
Section 5.1.9, “IPv6 Support”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 1.8.5, “MySQL Differences from Standard SQL”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 2.10.1.1, “Problems Running `mysql_install_db`”  
Section 6.3.8, “Proxy Users”

Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 16.4.1.22, “Replication of the `mysql` System Database”  
Section 6.1.1, “Security Guidelines”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.5.1, “Support for Read-Only Media”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”  
Section 6.2, “The MySQL Access Privilege System”  
Section 6.3.1, “User Names and Passwords”  
Section 6.2.6, “When Privilege Changes Take Effect”

## REVOKE ALL PRIVILEGES

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”

## ROLLBACK

Section 14.2.3.14, “InnoDB Error Handling”  
Section 21.8.6.3, “`mysql_change_user()`”  
Section 13.3.4, “`SAVEPOINT`, `ROLLBACK TO SAVEPOINT`, and `RELEASE SAVEPOINT` Syntax”  
Section 13.3.1, “`START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”  
Section 14.2.3.9, “Deadlock Detection and Rollback”  
Section 5.4.2, “Grouping DML Operations with Transactions”  
Section 12.14, “Information Functions”  
Section 13.3.5.1, “Interaction of Table Locking and Transactions”  
Section 14.2.8, “MySQL and the ACID Model”  
Section 13.3, “MySQL Transactional and Locking Statements”  
Section 16.4.1.29, “Replication and Transactions”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section C.5.5.5, “Rollback Failure for Nontransactional Tables”  
Section 5.1.4, “Server System Variables”  
Section 13.3.2, “Statements That Cannot Be Rolled Back”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.3.1, “The `InnoDB` Transaction Model and Locking”  
Section 5.2.4, “The Binary Log”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”  
Section 18.3.1, “Trigger Syntax”

## ROLLBACK TO SAVEPOINT

Section 13.3.4, “`SAVEPOINT`, `ROLLBACK TO SAVEPOINT`, and `RELEASE SAVEPOINT` Syntax”

---

## S

[\[index top \[4013\]\]](#)

### SAVEPOINT

Section 13.3.4, “[SAVEPOINT, ROLLBACK TO SAVEPOINT, and RELEASE SAVEPOINT Syntax](#)”

### SELECT

Section 4.5.1.1, “[mysql Options](#)”  
Section 13.1.6, “[ALTER TABLE Syntax](#)”  
Section 13.1.7, “[ALTER VIEW Syntax](#)”  
Section 13.2.1, “[CALL Syntax](#)”  
Section 13.1.9, “[CREATE EVENT Syntax](#)”  
Section 13.1.12, “[CREATE PROCEDURE and CREATE FUNCTION Syntax](#)”  
Section 13.1.14.1, “[CREATE TABLE ... SELECT Syntax](#)”  
Section 13.1.14, “[CREATE TABLE Syntax](#)”  
Section 13.1.16, “[CREATE VIEW Syntax](#)”  
Section 13.2.2, “[DELETE Syntax](#)”  
Section 8.8.3, “[EXPLAIN EXTENDED Output Format](#)”  
Section 8.8.2, “[EXPLAIN Output Format](#)”  
Section 13.8.2, “[EXPLAIN Syntax](#)”  
Section 14.9.3, “[FEDERATED Storage Engine Notes and Tips](#)”  
Section 14.2.3.12.3, “[FULLTEXT Indexes](#)”  
Section 13.7.1.4, “[GRANT Syntax](#)”  
Section 12.17.1, “[GROUP BY \(Aggregate\) Functions](#)”  
Section 13.2.4, “[HANDLER Syntax](#)”  
Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 14.2.6, “[InnoDB Startup Options and System Variables](#)”  
Section 13.2.5.3, “[INSERT ... ON DUPLICATE KEY UPDATE Syntax](#)”  
Section 13.2.5.1, “[INSERT ... SELECT Syntax](#)”  
Section 13.2.5.2, “[INSERT DELAYED Syntax](#)”  
Section 13.2.5, “[INSERT Syntax](#)”  
Section 13.2.9.2, “[JOIN Syntax](#)”  
Section 13.7.6.4, “[KILL Syntax](#)”  
Section 13.2.7, “[LOAD XML Syntax](#)”  
Section 14.8.2, “[MERGE Table Problems](#)”  
Section 7.6.4, “[MyISAM Table Optimization](#)”  
Section 21.8.6.1, “[mysql\\_affected\\_rows\(\)](#)”  
Section 21.8.6.17, “[mysql\\_fetch\\_field\(\)](#)”  
Section 21.8.6.22, “[mysql\\_field\\_count\(\)](#)”  
Section 21.8.6.47, “[mysql\\_num\\_fields\(\)](#)”  
Section 21.8.6.48, “[mysql\\_num\\_rows\(\)](#)”  
Section 21.8.10.10, “[mysql\\_stmt\\_execute\(\)](#)”  
Section 21.8.10.11, “[mysql\\_stmt\\_fetch\(\)](#)”  
Section 21.8.10.18, “[mysql\\_stmt\\_num\\_rows\(\)](#)”  
Section 21.8.10.28, “[mysql\\_stmt\\_store\\_result\(\)](#)”  
Section 21.8.6.70, “[mysql\\_store\\_result\(\)](#)”  
Section 21.8.6.72, “[mysql\\_use\\_result\(\)](#)”  
Section 17.2.3.1, “[RANGE COLUMNS partitioning](#)”  
Section 13.2.8, “[REPLACE Syntax](#)”  
Section 13.2.9.1, “[SELECT ... INTO Syntax](#)”

Section 13.2.9, “[SELECT Syntax](#)”  
Section 13.3.6, “[SET TRANSACTION Syntax](#)”  
Section 13.7.4, “[SET Syntax](#)”  
Section 13.7.5.3, “[SHOW BINLOG EVENTS Syntax](#)”  
Section 13.7.5.11, “[SHOW CREATE PROCEDURE Syntax](#)”  
Section 13.7.5.14, “[SHOW CREATE VIEW Syntax](#)”  
Section 13.7.5.18, “[SHOW ERRORS Syntax](#)”  
Section 13.7.5.28, “[SHOW PROCEDURE CODE Syntax](#)”  
Section 13.7.5.30, “[SHOW PROCESSLIST Syntax](#)”  
Section 13.7.5.33, “[SHOW RELAYLOG EVENTS Syntax](#)”  
Section 13.7.5.40, “[SHOW VARIABLES Syntax](#)”  
Section 13.7.5.41, “[SHOW WARNINGS Syntax](#)”  
Section 13.7.5, “[SHOW Syntax](#)”  
Section 13.3.1, “[START TRANSACTION, COMMIT, and ROLLBACK Syntax](#)”  
Section 13.2.9.4, “[UNION Syntax](#)”  
Section 13.2.11, “[UPDATE Syntax](#)”  
Section 12.3.4, “[Assignment Operators](#)”  
Section 14.2.3.6, “[Avoiding the Phantom Problem Using Next-Key Locking](#)”  
Section 18.7, “[Binary Logging of Stored Programs](#)”  
Section 8.6.2, “[Bulk Data Loading for MyISAM Tables](#)”  
Section 21.8.4, “[C API Data Structures](#)”  
Section 21.8.5, “[C API Function Overview](#)”  
Section 21.8.9, “[C API Prepared Statement Function Overview](#)”  
Section 21.8.16, “[C API Support for Multiple Statement Execution](#)”  
Section 21.8.19, “[C API Support for Prepared CALL Statements](#)”  
Section 8.9.4, “[Caching of Prepared Statements and Stored Programs](#)”  
Section 12.3.2, “[Comparison Functions and Operators](#)”  
Section 8.3.8, “[Comparison of B-Tree and Hash Indexes](#)”  
Section 8.10.3, “[Concurrent Inserts](#)”  
Section 5.4.4.2, “[Configurable InnoDB Auto-Increment Locking](#)”  
Section 10.1.4, “[Connection Character Sets and Collations](#)”  
Section 14.2.3.3, “[Consistent Nonlocking Reads](#)”  
Section 5.4.3, “[Converting Tables from MyISAM to InnoDB](#)”  
Section 14.9.2.1, “[Creating a FEDERATED Table Using CONNECTION](#)”  
Section 3.3.1, “[Creating and Selecting a Database](#)”  
Section 13.6.6.2, “[Cursor DECLARE Syntax](#)”  
Section 13.6.6.3, “[Cursor FETCH Syntax](#)”  
Section 14.2.3.9, “[Deadlock Detection and Rollback](#)”  
Section 4.5.7, “[mysqlslap — Load Emulation Client](#)”  
Section 8.4.3.2, “[Disadvantages of Creating Many Tables in the Same Database](#)”  
Section 8.11.5.2, “[DNS Lookup Optimization and the Host Cache](#)”  
Section 5.1.5.2, “[Dynamic System Variables](#)”  
Section 3.2, “[Entering Queries](#)”  
Section 18.4.2, “[Event Scheduler Configuration](#)”

---

Section 10.1.7.8, “Examples of the Effect of Collation”  
Section 1.8.5.4, “Foreign Key Differences”  
Chapter 12, *Functions and Operators*  
Section 8.12.5.2, “General Thread States”  
Section 5.4.2, “Grouping DML Operations with Transactions”  
Section 8.2.1.2, “How MySQL Optimizes `WHERE` Clauses”  
Section 17.2.7, “How MySQL Partitioning Handles `NULL`”  
Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 1.7, “How to Report Bugs or Problems”  
Section 9.2.1, “Identifier Qualifiers”  
Section 13.2.9.3, “Index Hint Syntax”  
Section 12.14, “Information Functions”  
Section 8.10.1, “Internal Locking Methods”  
Section C.5.8, “Known Issues in MySQL”  
Section 13.6.4.2, “Local Variable Scope and Resolution”  
Section 14.2.3.4, “Locking Reads (`SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE`)”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 17.3.1, “Management of `RANGE` and `LIST` Partitions”  
Section 8.3.5, “Multiple-Column Indexes”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Chapter 20, *MySQL Performance Schema*  
Section 12.9.1, “Natural Language Full-Text Searches”  
Section 17.3.5, “Obtaining Information About Partitions”  
Section 8.3, “Optimization and Indexes”  
Section C.5.6, “Optimizer-Related Issues”  
Section 8.5.2, “Optimizing `InnoDB` Transaction Management”  
Section 8.6.1, “Optimizing `MyISAM` Queries”  
Section 8.2.1, “Optimizing `SELECT` Statements”  
Section 8.8.1, “Optimizing Queries with `EXPLAIN`”  
Section 8.13.16.3, “Optimizing Subqueries in the `FROM` Clause (Derived Tables)”  
Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”  
Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”  
Section 4.6.3.4, “Other `myisamchk` Options”  
Section 8.2.5, “Other Optimization Tips”  
Section 5.5.1, “Overview of Online DDL”  
Section 17.5, “Partition Selection”  
Section 17.6.4, “Partitioning and Locking”  
Section 20.4, “Performance Schema Instrument Naming Conventions”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section C.5.5.2, “Problems Using `DATE` Columns”  
Section C.5.5.8, “Problems with Floating-Point Values”  
Section 8.9.3.2, “Query Cache `SELECT` Options”  
Section 8.9.3.4, “Query Cache Status and Maintenance”  
Section 8.12.5.4, “Query Cache Thread States”  
Section 14.5.1, “Repairing and Checking CSV Tables”  
Section 16.2, “Replication Implementation”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 16.4.1.3, “Replication of `CREATE ... IF NOT EXISTS` Statements”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 9.3, “Reserved Words”  
Section E.1, “Restrictions on Stored Programs”  
Section E.5, “Restrictions on Views”  
Section 3.3.4, “Retrieving Information from a Table”  
Section 3.6.7, “Searching on Two Keys”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 3.3.4.1, “Selecting All Data”  
Section 3.3.4.2, “Selecting Particular Rows”  
Section 5.1.4, “Server System Variables”  
Section C.5.5.7, “Solving Problems with No Matching Rows”  
Section 8.2.1.1, “Speed of `SELECT` Statements”  
Section 8.2.2.2, “Speed of `UPDATE` Statements”  
Section 14.2.4.6, “Starting `InnoDB` on a Corrupted Database”  
Section 5.7.1.12, “Statement Probes”  
Section 18.2.1, “Stored Routine Syntax”  
Section 9.1.1, “String Literals”  
Section 13.2.10.8, “Subqueries in the `FROM` Clause”  
Section 13.2.10.6, “Subqueries with `EXISTS` or `NOT EXISTS`”  
Section 13.2.10.9, “Subquery Errors”  
Section 13.2.10, “Subquery Syntax”  
Section 8.10.2, “Table Locking Issues”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 11.4.4, “The `ENUM` Type”  
Section 20.9.9.1, “The `host_cache` Table”  
Section 19.4, “The `INFORMATION_SCHEMA COLUMNS` Table”  
Section 19.7, “The `INFORMATION_SCHEMA EVENTS` Table”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 20.9.9.3, “The `threads` Table”  
Section 5.2.4, “The Binary Log”  
Section 6.2, “The MySQL Access Privilege System”  
Section 8.9.3, “The MySQL Query Cache”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 13.2.10.1, “The Subquery as Scalar Operand”  
Section 18.3.1, “Trigger Syntax”  
Section 1.2, “Typographical and Syntax Conventions”

---



---

Section 2.10.1, “Unix Postinstallation Procedures”  
Section 9.4, “User-Defined Variables”  
Section 8.4.2.4, “Using `PROCEDURE ANALYSE`”  
Section 12.18.6.2, “Using a Spatial Index”  
Section 22.4.1.6, “Using Server Logs to Find Causes of Errors in `mysqld`”  
Section 5.1.5, “Using System Variables”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”  
Section 10.1.12, “UTF-8 for Metadata”  
Section 18.5.1, “View Syntax”  
Section 1.4, “What Is New in MySQL 5.6”  
Section C.5.4.4, “Where MySQL Stores Temporary Files”  
Section 22.2.4.8, “Writing Audit Plugins”

### **SELECT \***

Section 11.4.3, “The `BLOB` and `TEXT` Types”

### **SELECT \* FROM t PARTITION ( )**

Section 17.1, “Overview of Partitioning in MySQL”

### **SELECT \* INTO OUTFILE 'file\_name' FROM tbl\_name**

Section 7.2, “Database Backup Methods”

### **SELECT ... FOR UPDATE**

Section 14.2.3.2, “InnoDB Lock Modes”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 14.2.3.4, “Locking Reads (`SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE`)”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

### **SELECT ... FROM**

Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

### **SELECT ... FROM ... FOR UPDATE**

Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

### **SELECT ... FROM ... LOCK IN SHARE MODE**

Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

### **SELECT ... INTO**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 13.2.9.1, “`SELECT ... INTO` Syntax”  
Section 1.8.5.1, “`SELECT INTO TABLE` Differences”  
Section 13.2.9, “`SELECT` Syntax”  
Section 13.6.4.2, “Local Variable Scope and Resolution”

Section 16.4.1.14, “Replication and System Functions”

### **SELECT ... INTO DUMPFILE**

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 5.1.4, “Server System Variables”

Section 2.10.1, “Unix Postinstallation Procedures”

### **SELECT ... INTO OUTFILE**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”

Section 9.1.7, “`NULL` Values”

Section 13.2.9.1, “`SELECT ... INTO` Syntax”

Section 1.8.5.1, “`SELECT INTO TABLE` Differences”

Section 7.1, “Backup and Recovery Types”

Section 6.2.7, “Causes of Access-Denied Errors”

Section 7.4.3, “Dumping Data in Delimited-Text Format with `mysqldump`”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 6.2.1, “Privileges Provided by MySQL”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”

Section 1.2, “Typographical and Syntax Conventions”

Section E.10.6, “Windows Platform Limitations”

### **SELECT ... INTO OUTFILE 'file\_name'**

Section 13.2.9.1, “`SELECT ... INTO` Syntax”

### **SELECT ... INTO var\_list**

Section E.1, “Restrictions on Stored Programs”

Section 13.6.4, “Variables in Stored Programs”

### **SELECT ... LOCK IN SHARE MODE**

Section 14.2.3.2, “InnoDB Lock Modes”

Section 13.3.6, “`SET TRANSACTION` Syntax”

Section 14.2.3.4, “Locking Reads (`SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE`)”

Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”

### **SELECT ... PARTITION**

Section 1.4, “What Is New in MySQL 5.6”

### **SELECT DISTINCT**

Section 8.12.5.2, “General Thread States”

Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”

Section E.4, “Restrictions on Subqueries”

### **SELECT HIGH\_PRIORITY**

Section 8.2.5, “Other Optimization Tips”

---

## SET

Section 13.7.4, “`SET` Syntax”  
Section 13.7.5.40, “`SHOW VARIABLES` Syntax”  
Section 12.3.4, “Assignment Operators”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 18.7, “Binary Logging of Stored Programs”  
Section 18.1, “Defining Stored Programs”  
Section 4.6.12, “`mysql_find_rows` — Extract SQL Statements from Files”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 5.1.5.2, “Dynamic System Variables”  
Section 18.4.2, “Event Scheduler Configuration”  
Chapter 12, *Functions and Operators*  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section 12.14, “Information Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 17.6.4, “Partitioning and Locking”  
Section 6.3.6, “Password Expiration and Sandbox Mode”  
Section 8.9.3.3, “Query Cache Configuration”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 13.2.10, “Subquery Syntax”  
Section 18.3.1, “Trigger Syntax”  
Section 9.4, “User-Defined Variables”  
Section 4.2.3.5, “Using Options to Set Program Variables”  
Section 5.1.5, “Using System Variables”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”  
Section 13.6.4, “Variables in Stored Programs”

## SET @@global.gtid\_purged

Section 4.5.4, “`mysqldump` — A Database Backup Program”

## SET autocommit

Section 8.5.4, “Bulk Data Loading for `InnoDB` Tables”  
Section 13.3, “MySQL Transactional and Locking Statements”

## SET GLOBAL

Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.4, “`SET` Syntax”  
Section 5.1.5.2, “Dynamic System Variables”  
Section 8.9.2.2, “Multiple Key Caches”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”  
Section 5.1.5, “Using System Variables”

## SET GLOBAL

### sql\_slave\_skip\_counter

Section 13.4.2.4, “`SET GLOBAL sql_slave_skip_counter` Syntax”

## SET GLOBAL TRANSACTION

Section 13.3.6, “`SET TRANSACTION` Syntax”

## SET NAMES

Section 10.1.6, “Character Set for Error Messages”  
Section 12.2, “Type Conversion in Expression Evaluation”

## SET PASSWORD

Section 13.7.1.1, “`ALTER USER` Syntax”  
Section C.5.2.4, “Client does not support authentication protocol”  
Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 13.7.4, “`SET` Syntax”  
Section 6.3.5, “Assigning Account Passwords”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”  
Section 6.1.2.1, “End-User Guidelines for Password Security”  
Section 12.14, “Information Functions”  
Section 2.5.1, “Installing MySQL from RPM Packages on Linux”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 6.3.6, “Password Expiration and Sandbox Mode”  
Section 6.1.2.4, “Password Hashing in MySQL”  
Section 22.2.3.8, “Password-Validation Plugins”  
Section 6.1.2.3, “Passwords and Logging”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 16.4.1.32, “Replication and Variables”  
Section 16.4.1.6, “Replication of `CURRENT_USER()`”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 5.1.4, “Server System Variables”  
Section 6.2.3, “Specifying Account Names”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 6.1.2.6, “The Password Validation Plugin”  
Section 6.3.7.2, “The SHA-256 Authentication Plugin”  
Section 1.4, “What Is New in MySQL 5.6”  
Section 6.2.6, “When Privilege Changes Take Effect”

## SET SESSION

Section 13.7.4, “`SET` Syntax”  
Section 5.1.5.2, “Dynamic System Variables”  
Section 5.1.5, “Using System Variables”

## SET SESSION TRANSACTION

Section 13.3.6, “`SET TRANSACTION` Syntax”

---

## **SET sql\_mode='mode\_value'**

Section 1.8.2, “Selecting SQL Modes”

## **SET TIMESTAMP = value**

Section 8.12.5, “Examining Thread Information”

## **SET TRANSACTION**

Section 13.3.6, “SET TRANSACTION Syntax”

Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 14.2.3.1, “The InnoDB Transaction Model and Locking”

## **SET TRANSACTION ISOLATION**

### **LEVEL**

Section 13.7.4, “SET Syntax”

Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”

## **SHOW**

Section 13.1.9, “CREATE EVENT Syntax”

Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”

Chapter 19, *INFORMATION\_SCHEMA Tables*

Section 21.8.10.28, “mysql\_stmt\_store\_result()”

Section 21.8.6.70, “mysql\_store\_result()”

Section 21.8.6.72, “mysql\_use\_result()”

Section 13.7.5.6, “SHOW COLUMNS Syntax”

Section 13.7.5.23, “SHOW INDEX Syntax”

Section 13.7.5.25, “SHOW OPEN TABLES Syntax”

Section 13.7.5.38, “SHOW TABLES Syntax”

Section 10.1.9.3, “SHOW Statements and INFORMATION\_SCHEMA”

Section 13.7.5, “SHOW Syntax”

Section 21.8.4, “C API Data Structures”

Section 21.8.5, “C API Function Overview”

Section 3.3, “Creating and Using a Database”

Section 13.6.6.2, “Cursor DECLARE Syntax”

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

Section 19.32, “Extensions to SHOW Statements”

Section 9.2.3, “Mapping of Identifiers to File Names”

Section 1.8.4, “MySQL Extensions to Standard SQL”

Section 20.1, “Performance Schema Quick Start”

Section E.1, “Restrictions on Stored Programs”

Section 13.4.1, “SQL Statements for Controlling Master Servers”

Section 19.1, “The INFORMATION\_SCHEMA CHARACTER\_SETS Table”

Section 19.3, “The INFORMATION\_SCHEMA COLLATION\_CHARACTER\_SET\_APPLICABILITY Table”

Section 19.2, “The INFORMATION\_SCHEMA COLLATIONS Table”

Section 19.5, “The INFORMATION\_SCHEMA COLUMN\_PRIVILEGES Table”

Section 19.4, “The INFORMATION\_SCHEMA COLUMNS Table”

Section 19.6, “The INFORMATION\_SCHEMA ENGINES Table”

Section 19.7, “The INFORMATION\_SCHEMA EVENTS Table”

Section 19.8, “The INFORMATION\_SCHEMA FILES Table”

Section 19.9, “The INFORMATION\_SCHEMA GLOBAL\_STATUS and SESSION\_STATUS Tables”

Section 19.10, “The INFORMATION\_SCHEMA GLOBAL\_VARIABLES and SESSION\_VARIABLES Tables”

Section 19.11, “The INFORMATION\_SCHEMA KEY\_COLUMN\_USAGE Table”

Section 19.14, “The INFORMATION\_SCHEMA PARTITIONS Table”

Section 19.15, “The INFORMATION\_SCHEMA PLUGINS Table”

Section 19.16, “The INFORMATION\_SCHEMA PROCESSLIST Table”

Section 19.17, “The INFORMATION\_SCHEMA PROFILING Table”

Section 19.18, “The INFORMATION\_SCHEMA REFERENTIAL\_CONSTRAINTS Table”

Section 19.21, “The INFORMATION\_SCHEMA SCHEMA\_PRIVILEGES Table”

Section 19.20, “The INFORMATION\_SCHEMA SCHEMATA Table”

Section 19.22, “The INFORMATION\_SCHEMA STATISTICS Table”

Section 19.25, “The INFORMATION\_SCHEMA TABLE\_CONSTRAINTS Table”

Section 19.26, “The INFORMATION\_SCHEMA TABLE\_PRIVILEGES Table”

Section 19.23, “The INFORMATION\_SCHEMA TABLES Table”

Section 19.24, “The INFORMATION\_SCHEMA TABLESPACES Table”

Section 19.27, “The INFORMATION\_SCHEMA TRIGGERS Table”

Section 19.28, “The INFORMATION\_SCHEMA USER\_PRIVILEGES Table”

Section 19.29, “The INFORMATION\_SCHEMA VIEWS Table”

Section 5.2.4, “The Binary Log”

Section 10.1.12, “UTF-8 for Metadata”

## **SHOW AUTHORS**

Section 13.7.5.1, “SHOW AUTHORS Syntax”

## **SHOW BINARY LOGS**

Section 13.4.1.1, “PURGE BINARY LOGS Syntax”

Section 13.7.5.2, “SHOW BINARY LOGS Syntax”

Section 6.2.1, “Privileges Provided by MySQL”

---

Section 13.4.1, “SQL Statements for Controlling Master Servers”  
Section 4.6.8.3, “Using `mysqlbinlog` to Back Up Binary Log Files”

## SHOW BINLOG EVENTS

Section 13.7.5.3, “`SHOW BINLOG EVENTS` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 16.1.3.1, “GTID Concepts”  
Section E.3, “Restrictions on Server-Side Cursors”  
Section 13.4.1, “SQL Statements for Controlling Master Servers”

## SHOW CHARACTER SET

Section 13.1.1, “`ALTER DATABASE` Syntax”  
Section 13.7.5.4, “`SHOW CHARACTER SET` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 10.1.2, “Character Sets and Collations in MySQL”  
Section 10.1.14, “Character Sets and Collations That MySQL Supports”  
Section 19.32, “Extensions to `SHOW` Statements”

## SHOW COLLATION

Section 13.1.1, “`ALTER DATABASE` Syntax”  
Section 13.7.5.5, “`SHOW COLLATION` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 21.8.4, “C API Data Structures”  
Section 10.5, “Character Set Configuration”  
Section 10.1.2, “Character Sets and Collations in MySQL”  
Section 10.1.3.5, “Character String Literal Character Set and Collation”  
Section 10.4.2, “Choosing a Collation ID”  
Section 10.1.3.4, “Column Character Set and Collation”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 10.1.3.3, “Table Character Set and Collation”  
Section 19.3, “The `INFORMATION_SCHEMA` `COLLATION_CHARACTER_SET_APPLICABILITY` Table”  
Section 19.2, “The `INFORMATION_SCHEMA` `COLLATIONS` Table”

## SHOW COLUMNS

Section 13.8.2, “`EXPLAIN` Syntax”  
Section 13.7.5.6, “`SHOW COLUMNS` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 19.32, “Extensions to `SHOW` Statements”  
Section 8.2.1.3, “Optimizing `LIMIT` Queries”  
Section 8.8.1, “Optimizing Queries with `EXPLAIN`”

`SHOW COLUMNS FROM tbl_name`  
`LIKE 'enum_col'`  
Section 11.4.4, “The `ENUM` Type”

## SHOW CONTRIBUTORS

Section 13.7.5.7, “`SHOW CONTRIBUTORS` Syntax”

## SHOW COUNT( )

Section 13.7.5.18, “`SHOW ERRORS` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”

## SHOW CREATE DATABASE

Section 13.7.5.8, “`SHOW CREATE DATABASE` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 5.1.4, “Server System Variables”

## SHOW CREATE EVENT

Section 13.7.5.9, “`SHOW CREATE EVENT` Syntax”  
Section 13.7.5.19, “`SHOW EVENTS` Syntax”  
Section 18.4.4, “Event Metadata”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## SHOW CREATE FUNCTION

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.7.5.10, “`SHOW CREATE FUNCTION` Syntax”  
Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”  
Section 1.7, “How to Report Bugs or Problems”  
Section 18.2.3, “Stored Routine Metadata”

## SHOW CREATE PROCEDURE

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.7.5.10, “`SHOW CREATE FUNCTION` Syntax”  
Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”  
Section 1.7, “How to Report Bugs or Problems”  
Section 18.2.3, “Stored Routine Metadata”

## SHOW CREATE SCHEMA

Section 13.7.5.8, “`SHOW CREATE DATABASE` Syntax”

## SHOW CREATE TABLE

Section 13.1.6.1, “`ALTER TABLE` Partition Operations”  
Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 13.8.2, “`EXPLAIN` Syntax”  
Section 17.2.5, “`KEY` Partitioning”  
Section 13.7.5.6, “`SHOW COLUMNS` Syntax”  
Section 13.7.5.12, “`SHOW CREATE TABLE` Syntax”  
Section 10.1.9.3, “`SHOW` Statements and `INFORMATION_SCHEMA`”  
Section 11.5, “Data Type Default Values”

---

Section 3.4, “Getting Information About Databases and Tables”  
Section 14.9.2, “How to Create `FEDERATED` Tables”  
Section 7.6.3, “How to Repair `MyISAM` Tables”  
Section 17.3.1, “Management of `RANGE` and `LIST` Partitions”  
Section 17.3.5, “Obtaining Information About Partitions”  
Section 2.11.4, “Rebuilding or Repairing Tables or Indexes”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 13.1.14.3, “Silent Column Specification Changes”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”  
Section 3.6.6, “Using Foreign Keys”

## **SHOW CREATE TRIGGER**

Section 13.7.5.13, “`SHOW CREATE TRIGGER` Syntax”

## **SHOW CREATE VIEW**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section E.5, “Restrictions on Views”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 18.5.4, “View Metadata”

## **SHOW DATABASES**

Section 13.1.8, “`CREATE DATABASE` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 13.7.5.15, “`SHOW DATABASES` Syntax”  
Section 3.3, “Creating and Using a Database”  
Section 19.32, “Extensions to `SHOW` Statements”  
Section 3.4, “Getting Information About Databases and Tables”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 20.2.1, “Performance Schema Build Configuration”  
Section 6.2.2, “Privilege System Grant Tables”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **SHOW ENGINE**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”

## **SHOW ENGINE INNODB MUTEX**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 5.1.4, “Server System Variables”

## **SHOW ENGINE INNODB STATUS**

Section 5.4.5, “`InnoDB` and `FOREIGN KEY` Constraints”  
Section 14.2.4.4.1, “`InnoDB` Standard Monitor and Lock Monitor Output”

Section 14.2.6, “`InnoDB` Startup Options and System Variables”

Section 14.2.4.4, “`SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors”  
Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 14.2.3.12.6, “Adaptive Hash Indexes”  
Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 14.2.2.5, “Moving or Copying `InnoDB` Tables to Another Machine”  
Section 8.5.7, “Optimizing `InnoDB` Disk I/O”  
Section C.1, “Sources of Error Information”  
Section 14.2.5.1, “Support for Read-Only Media”  
Section 19.30.16, “The `INFORMATION_SCHEMA INNODB_BUFFER_POOL_STATS` Table”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”

## **SHOW ENGINE**

## **PERFORMANCE\_SCHEMA STATUS**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 20.5, “Performance Schema Status Monitoring”

## **SHOW ENGINES**

Section 13.7.5.17, “`SHOW ENGINES` Syntax”  
Section 14.2.2.7, “Checking `InnoDB` Availability”  
Section 20.2.1, “Performance Schema Build Configuration”  
Section 20.1, “Performance Schema Quick Start”  
Section 2.3.4.3, “Selecting a MySQL Server Type”  
Section 5.1.4, “Server System Variables”  
Chapter 14, *Storage Engines*  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 1.4, “What Is New in MySQL 5.6”

## **SHOW ERRORS**

Section 13.6.7.3, “`GET DIAGNOSTICS` Syntax”  
Section 5.4.5, “`InnoDB` and `FOREIGN KEY` Constraints”  
Section 13.6.7.4.3, “`RESIGNAL` with a Condition Value and Optional New Signal Information”  
Section 13.7.5.18, “`SHOW ERRORS` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section 5.1.4, “Server System Variables”  
Section 13.6.7.5.1, “Signal Condition Information Items”  
Section C.1, “Sources of Error Information”

## **SHOW EVENTS**

Section 13.7.5.19, “`SHOW EVENTS` Syntax”  
Section 18.4.4, “Event Metadata”  
Section 16.4.1.10, “Replication of Invoked Features”  
Section 19.7, “The `INFORMATION_SCHEMA EVENTS` Table”



---

Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## SHOW FULL COLUMNS

Section 13.1.14, “CREATE TABLE Syntax”

Section 10.1.9.3, “SHOW Statements and INFORMATION\_SCHEMA”

Section 19.5, “The INFORMATION\_SCHEMA COLUMN\_PRIVILEGES Table”

## SHOW FULL PROCESSLIST

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”

Section 8.12.5, “Examining Thread Information”

## SHOW FULL TABLES

Section 13.7.5.38, “SHOW TABLES Syntax”

Section 4.5.6, “mysqlshow — Display Database, Table, and Column Information”

## SHOW FUNCTION CODE

Section 13.7.5.20, “SHOW FUNCTION CODE Syntax”

Section 13.7.5.28, “SHOW PROCEDURE CODE Syntax”

## SHOW FUNCTION STATUS

Section 13.7.5.21, “SHOW FUNCTION STATUS Syntax”

Section 13.7.5.29, “SHOW PROCEDURE STATUS Syntax”

Section 18.2.3, “Stored Routine Metadata”

## SHOW GLOBAL STATUS

Section 5.1.4, “Server System Variables”

Section 19.9, “The INFORMATION\_SCHEMA GLOBAL\_STATUS and SESSION\_STATUS Tables”

## SHOW GLOBAL VARIABLES

Section 19.10, “The INFORMATION\_SCHEMA GLOBAL\_VARIABLES and SESSION\_VARIABLES Tables”

## SHOW GRANTS

Section 13.7.1.4, “GRANT Syntax”

Section 13.7.1.6, “REVOKE Syntax”

Section 13.7.5.22, “SHOW GRANTS Syntax”

Section 13.7.5.27, “SHOW PRIVILEGES Syntax”

Section 6.3.2, “Adding User Accounts”

Section 6.2.2, “Privilege System Grant Tables”

Section 6.1.1, “Security Guidelines”

Section 6.2, “The MySQL Access Privilege System”

## SHOW INDEX

Section 13.7.2.1, “ANALYZE TABLE Syntax”

Section 8.8.2, “EXPLAIN Output Format”

Section 13.8.2, “EXPLAIN Syntax”

Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 13.7.5.6, “SHOW COLUMNS Syntax”

Section 13.7.5.23, “SHOW INDEX Syntax”

Section 13.2.9.3, “Index Hint Syntax”

Section 14.2.7, “Limits on InnoDB Tables”

Section 4.6.3.4, “Other myisamchk Options”

Section 19.22, “The INFORMATION\_SCHEMA STATISTICS Table”

Section 19.25, “The INFORMATION\_SCHEMA TABLE\_CONSTRAINTS Table”

## SHOW MASTER LOGS

Section 13.7.5.2, “SHOW BINARY LOGS Syntax”

## SHOW MASTER STATUS

Section 13.7.5.24, “SHOW MASTER STATUS Syntax”

Section 16.1.1.5, “Creating a Data Snapshot Using mysqldump”

Section 16.1.3.1, “GTID Concepts”

Section 16.4.5, “How to Report Replication Bugs or Problems”

Section 16.1.1.4, “Obtaining the Replication Master Binary Log Coordinates”

Section 6.2.1, “Privileges Provided by MySQL”

Section 13.4.1, “SQL Statements for Controlling Master Servers”

Section 16.4.4, “Troubleshooting Replication”

## SHOW OPEN TABLES

Section 13.7.5.25, “SHOW OPEN TABLES Syntax”

## SHOW PLUGINS

Section 13.7.3.3, “INSTALL PLUGIN Syntax”

Section 13.7.5.26, “SHOW PLUGINS Syntax”

Section 5.1.8.1, “Installing and Uninstalling Plugins”

Section 6.3.11.1, “Installing the Audit Log Plugin”

Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”

Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”

Section 5.1.8.2, “Obtaining Server Plugin Information”

Chapter 17, *Partitioning*

Section 6.1.2.6.1, “Password Validation Plugin Installation”

Section 6.3.7, “Pluggable Authentication”

Section 22.2.1, “Plugin API Characteristics”

Section 22.2.2, “Plugin API Components”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”

Section 19.15, “The INFORMATION\_SCHEMA PLUGINS Table”

Section 8.11.6.1, “Thread Pool Components and Installation”

---

Section 22.2.4.6, “Writing `INFORMATION_SCHEMA` Plugins”  
Section 22.2.4.8, “Writing Audit Plugins”  
Section 22.2.4.5, “Writing Daemon Plugins”  
Section 22.2.4.4, “Writing Full-Text Parser Plugins”  
Section 22.2.4.10, “Writing Password-Validation Plugins”  
Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

## SHOW PRIVILEGES

Section 13.7.5.27, “`SHOW PRIVILEGES` Syntax”

## SHOW PROCEDURE CODE

Section 13.7.5.20, “`SHOW FUNCTION CODE` Syntax”  
Section 13.7.5.28, “`SHOW PROCEDURE CODE` Syntax”

## SHOW PROCEDURE STATUS

Section 13.7.5.21, “`SHOW FUNCTION STATUS` Syntax”  
Section 13.7.5.29, “`SHOW PROCEDURE STATUS` Syntax”  
Section 18.2.3, “Stored Routine Metadata”

## SHOW PROCESSLIST

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 13.7.1.4, “`GRANT` Syntax”  
Section 14.2.3.14, “InnoDB Error Handling”  
Section 13.7.6.4, “`KILL` Syntax”  
Section 21.8.6.43, “`mysql_list_processes()`”  
Section 13.7.5.30, “`SHOW PROCESSLIST` Syntax”  
Section 13.7.5.31, “`SHOW PROFILE` Syntax”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 13.3.1, “`START TRANSACTION, COMMIT, and ROLLBACK` Syntax”  
Section C.5.2.7, “Too many connections”  
Section 16.1.5.1, “Checking Replication Status”  
Section 5.7.1.2, “Command Probes”  
Section 5.7.1.1, “Connection Probes”  
Section 16.3.9, “Delayed Replication”  
Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”  
Section 18.4.2, “Event Scheduler Configuration”  
Section 8.12.5, “Examining Thread Information”  
Section 8.12.5.2, “General Thread States”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 20.4, “Performance Schema Instrument Naming Conventions”  
Section 20.9.4, “Performance Schema Stage Event Tables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.7.1.6, “Query Execution Probes”  
Section 5.7.1.3, “Query Probes”  
Section 16.2.1, “Replication Implementation Details”  
Section 16.3.6, “Switching Masters During Failover”

Section 19.16, “The `INFORMATION_SCHEMA` `PROCESSLIST` Table”  
Section 20.9.9.3, “The `threads` Table”  
Section 16.4.4, “Troubleshooting Replication”

## SHOW PROFILE

Section 13.7.5.31, “`SHOW PROFILE` Syntax”  
Section 13.7.5.32, “`SHOW PROFILES` Syntax”  
Section 8.12.5, “Examining Thread Information”  
Section 8.12.5.2, “General Thread States”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.1.4, “Server System Variables”  
Section 19.17, “The `INFORMATION_SCHEMA` `PROFILING` Table”

## SHOW PROFILES

Section 13.7.5.31, “`SHOW PROFILE` Syntax”  
Section 13.7.5.32, “`SHOW PROFILES` Syntax”  
Section 2.9.4, “MySQL Source-Configuration Options”  
Section 5.1.4, “Server System Variables”  
Section 19.17, “The `INFORMATION_SCHEMA` `PROFILING` Table”

## SHOW RELAYLOG EVENTS

Section 13.7.5.3, “`SHOW BINLOG EVENTS` Syntax”  
Section 13.7.5.33, “`SHOW RELAYLOG EVENTS` Syntax”  
Section 13.4.2, “SQL Statements for Controlling Slave Servers”

## SHOW SCHEMAS

Section 13.7.5.15, “`SHOW DATABASES` Syntax”

## SHOW SESSION STATUS

Section 19.9, “The `INFORMATION_SCHEMA` `GLOBAL_STATUS` and `SESSION_STATUS` Tables”

## SHOW SESSION VARIABLES

Section 19.10, “The `INFORMATION_SCHEMA` `GLOBAL_VARIABLES` and `SESSION_VARIABLES` Tables”

## SHOW SLAVE HOSTS

Section 13.7.5.34, “`SHOW SLAVE HOSTS` Syntax”  
Section 16.1.5.1, “Checking Replication Status”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 13.4.1, “SQL Statements for Controlling Master Servers”

## SHOW SLAVE STATUS

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 13.4.1.1, “`PURGE BINARY LOGS` Syntax”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”

---

Section 16.1.5.1, “Checking Replication Status”  
Section 16.3.9, “Delayed Replication”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 16.1.3.1, “GTID Concepts”  
Section 16.4.5, “How to Report Replication Bugs or Problems”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 16.2.1, “Replication Implementation Details”  
Section 8.12.5.6, “Replication Slave I/O Thread States”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.3.7, “Setting Up Replication Using SSL”  
Section 16.4.1.25, “Slave Errors During Replication”  
Section 16.2.2.2, “Slave Status Logs”  
Section C.1, “Sources of Error Information”  
Section 13.4.2, “SQL Statements for Controlling Slave Servers”  
Section 16.4.4, “Troubleshooting Replication”

## SHOW STATUS

Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 13.7.5.36, “`SHOW STATUS` Syntax”  
Section 22.2.1, “Plugin API Characteristics”  
Section 8.9.3.4, “Query Cache Status and Maintenance”  
Section 16.4.1.21, “Replication and Temporary Tables”  
Section 16.2.1, “Replication Implementation Details”  
Section 16.4.1.27, “Replication Retries and Timeouts”  
Section E.1, “Restrictions on Stored Programs”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”  
Section 22.2.4.2.2, “Server Plugin Status and System Variables”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 8.13.5, “Use of Index Extensions”  
Section 22.2.4.8, “Writing Audit Plugins”  
Section 22.2.4.4, “Writing Full-Text Parser Plugins”  
Section 22.2.4, “Writing Plugins”

## SHOW STATUS LIKE 'perf%'

Section 20.5, “Performance Schema Status Monitoring”

## SHOW TABLE STATUS

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 13.8.2, “`EXPLAIN` Syntax”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 13.7.5.6, “`SHOW COLUMNS` Syntax”  
Section 13.7.5.37, “`SHOW TABLE STATUS` Syntax”  
Section 14.2.5.8, “Controlling Optimizer Statistics Estimation”

Section 5.4, “Creating and Using InnoDB Tables and Indexes”  
Section 5.3.2, “File Space Management”  
Section 14.2.7, “Limits on InnoDB Tables”  
Section 17.3.5, “Obtaining Information About Partitions”  
Section 14.2.3.12.7, “Physical Row Structure”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 5.4.4.1, “Traditional InnoDB Auto-Increment Locking”

## SHOW TABLES

Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 13.7.5.37, “`SHOW TABLE STATUS` Syntax”  
Section 13.7.5.38, “`SHOW TABLES` Syntax”  
Section C.5.2.16, “Table '`tbl_name`' doesn't exist”  
Section C.5.7.2, “`TEMPORARY` Table Problems”  
Section 3.3.2, “Creating a Table”  
Section 19.32, “Extensions to `SHOW` Statements”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 9.2.3, “Mapping of Identifiers to File Names”  
Section 5.1.3, “Server Command Options”

## SHOW TRIGGERS

Section 13.7.5.39, “`SHOW TRIGGERS` Syntax”  
Section 19.27, “The `INFORMATION_SCHEMA TRIGGERS` Table”  
Section 18.3.2, “Trigger Metadata”

## SHOW VARIABLES

Section 13.7.4, “`SET` Syntax”  
Section 13.7.5.40, “`SHOW VARIABLES` Syntax”  
Section 22.2.4.3, “Compiling and Installing Plugin Libraries”  
Section 18.4.2, “Event Scheduler Configuration”  
Section 20.2.2, “Performance Schema Startup Configuration”  
Section 22.2.1, “Plugin API Characteristics”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 5.6, “Running Multiple MySQL Instances on One Machine”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 5.1.4, “Server System Variables”  
Section 5.1.5, “Using System Variables”  
Section 22.2.4.10, “Writing Password-Validation Plugins”  
Section 22.2.4, “Writing Plugins”

## SHOW WARNINGS

Section 13.1.6, “`ALTER TABLE` Syntax”  
Section 13.1.21, “`DROP PROCEDURE` and `DROP FUNCTION` Syntax”  
Section 8.8.3, “`EXPLAIN EXTENDED` Output Format”  
Section 8.8.2, “`EXPLAIN` Output Format”  
Section 13.6.7.3, “`GET DIAGNOSTICS` Syntax”



---

Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 1.8.6.1, “PRIMARY KEY and UNIQUE Index Constraints”  
Section 13.7.5.18, “SHOW ERRORS Syntax”  
Section 13.7.5.41, “SHOW WARNINGS Syntax”  
Section 13.6.7.5, “SIGNAL Syntax”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 10.4.4.3, “Diagnostics During `Index.xml` Parsing”  
Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”  
Section 9.2.4, “Function Name Parsing and Resolution”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 8.13.16.4, “Optimizing Subqueries with `EXISTS` Strategy”  
Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”  
Section 8.13.16.2, “Optimizing Subqueries with Subquery Materialization”  
Section 5.1.4, “Server System Variables”  
Section 13.6.7.5.1, “Signal Condition Information Items”  
Section C.1, “Sources of Error Information”

## SIGNAL

Section 13.6.7.1, “DECLARE ... CONDITION Syntax”  
Section 13.6.7.2, “DECLARE ... HANDLER Syntax”  
Section 13.6.7.4, “RESIGNAL Syntax”  
Section 13.6.7.5, “SIGNAL Syntax”  
Section 13.6.7, “Condition Handling”  
Section 13.6.7.7.2, “Diagnostics Area Information Items”  
Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”  
Section 13.6.7.7.3, “How the Diagnostics Area is Populated”  
Section 12.14, “Information Functions”  
Section E.2, “Restrictions on Condition Handling”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.6.7.6, “Scope Rules for Handlers”  
Section 13.6.7.5.1, “Signal Condition Information Items”

## START SLAVE

Section 13.4.2.1, “CHANGE MASTER TO Syntax”  
Section 13.4.2.3, “RESET SLAVE Syntax”  
Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”  
Section 13.4.2.5, “START SLAVE Syntax”  
Section 13.4.2.6, “STOP SLAVE Syntax”  
Section 16.3.9, “Delayed Replication”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.3.1, “GTID Concepts”  
Section 6.1.2.3, “Passwords and Logging”

Section 16.1.5.2, “Pausing Replication on the Slave”  
Section 16.3.4, “Replicating Different Databases to Different Slaves”  
Section 16.1.4, “Replication and Binary Logging Options and Variables”  
Section 16.2.1, “Replication Implementation Details”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”  
Section 16.4.1.25, “Slave Errors During Replication”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 16.3.6, “Switching Masters During Failover”  
Section 16.4.4, “Troubleshooting Replication”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”  
Section 1.4, “What Is New in MySQL 5.6”

## START SLAVE UNTIL

Section 16.1.4.3, “Replication Slave Options and Variables”

## START SLAVE UNTIL

### SQL\_AFTER\_MTS\_GAPS

Section 16.1.4.3, “Replication Slave Options and Variables”

## START TRANSACTION

Section 13.6.1, “BEGIN ... END Compound-Statement Syntax”  
Section 14.2.3.14, “InnoDB Error Handling”  
Section 13.3.5, “LOCK TABLES and UNLOCK TABLES Syntax”  
Section 13.3.6, “SET TRANSACTION Syntax”  
Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 5.4.2, “Grouping DML Operations with Transactions”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 13.3.5.1, “Interaction of Table Locking and Transactions”  
Section 14.2.3.4, “Locking Reads (`SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE`)”  
Section 13.3, “MySQL Transactional and Locking Statements”  
Section E.1, “Restrictions on Stored Programs”  
Section 16.3.8, “Semisynchronous Replication”  
Section 5.1.4, “Server System Variables”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 14.2.3.1, “The InnoDB Transaction Model and Locking”  
Section 18.3.1, “Trigger Syntax”  
Section 13.3.7.2, “XA Transaction States”

---

## **START TRANSACTION READ ONLY**

Section 1.4, “What Is New in MySQL 5.6”

## **START TRANSACTION WITH CONSISTENT SNAPSHOT**

Section 14.2.3.3, “Consistent Nonlocking Reads”

## **STOP SLAVE**

Section 13.4.2.1, “CHANGE MASTER TO Syntax”

Section 13.4.1.2, “RESET MASTER Syntax”

Section 13.4.2.3, “RESET SLAVE Syntax”

Section 13.4.2.5, “START SLAVE Syntax”

Section 13.4.2.6, “STOP SLAVE Syntax”

Section 16.1.5.1, “Checking Replication Status”

Section 16.3.9, “Delayed Replication”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 16.1.4.5, “Global Transaction ID Options and Variables”

Section 16.1.1.9, “Introducing Additional Slaves to an Existing Replication Environment”

Section 16.1.5.2, “Pausing Replication on the Slave”

Section 16.1.4, “Replication and Binary Logging Options and Variables”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 16.3.6, “Switching Masters During Failover”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

Section 1.4, “What Is New in MySQL 5.6”

## **STOP SLAVE SQL\_THREAD**

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

## **T**

[index top [4013]]

## **TRUNCATE TABLE**

Section 13.1.15, “CREATE TRIGGER Syntax”

Section 13.2.2, “DELETE Syntax”

Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”

Section 13.2.4, “HANDLER Syntax”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.8.2, “MERGE Table Problems”

Section 13.1.27, “TRUNCATE TABLE Syntax”

Section 14.3.3.3, “Compressed Table Characteristics”

Section 20.9.8.7, “Connection Summary Tables”

Section 5.4.3, “Converting Tables from MyISAM to InnoDB”

Section 20.2.3.2.1, “Event Pre-Filtering”

Section 20.9.8.1, “Event Wait Summary Tables”

Section 20.9.8.5, “File I/O Summary Tables”

Section 8.9.3.1, “How the Query Cache Operates”

Section 5.2.4.4, “Logging Format for Changes to mysql Database Tables”

Section 17.3.4, “Maintenance of Partitions”

Section 17.3.1, “Management of RANGE and LIST Partitions”

Section 20.9.8.4, “Object Wait Summary Table”

Section 8.5.6, “Optimizing InnoDB DDL Operations”

Section 20.9.6, “Performance Schema Connection Tables”

Section 20.2.3.1, “Performance Schema Event Timing”

Section 20.8, “Performance Schema General Table Characteristics”

Section 20.9.8, “Performance Schema Summary Tables”

Section 6.2.1, “Privileges Provided by MySQL”

Section 16.4.1.20, “Replication and MEMORY Tables”

Section 16.4.1.31, “Replication and TRUNCATE TABLE”

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

Section 5.1.4, “Server System Variables”

Section 20.9.8.8, “Socket Summary Tables”

Section 20.9.8.2, “Stage Summary Tables”

Section 20.9.8.3, “Statement Summary Tables”

Section 13.3.3, “Statements That Cause an Implicit Commit”

Section 20.9.6.1, “The accounts Table”

Section 20.9.4.1, “The events\_stages\_current Table”

Section 20.9.4.2, “The events\_stages\_history Table”

Section 20.9.4.3, “The events\_stages\_history\_long Table”

Section 20.9.5.1, “The events\_statements\_current Table”

Section 20.9.5.2, “The events\_statements\_history Table”

Section 20.9.5.3, “The events\_statements\_history\_long Table”

Section 20.9.3.1, “The events\_waits\_current Table”

Section 20.9.3.2, “The events\_waits\_history Table”

Section 20.9.3.3, “The events\_waits\_history\_long Table”

Section 20.9.9.1, “The host\_cache Table”

Section 20.9.6.2, “The hosts Table”

Section 19.30.8, “The INFORMATION\_SCHEMA INNODB\_SYS\_INDEXES Table”

Section 19.30.7, “The INFORMATION\_SCHEMA INNODB\_SYS\_TABLES Table”

Section 14.4, “The MEMORY Storage Engine”

Section 20.9.1.5, “The setup\_timers Table”

Section 20.9.8.6.2, “The table\_io\_waits\_summary\_by\_index\_usage Table”

---

Section 20.9.8.6.1, “The `table_io_waits_summary_by_table` Table”  
Section 20.9.8.6.3, “The `table_lock_waits_summary_by_table` Table”  
Section 20.9.6.3, “The `users` Table”  
Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”  
Section 21.8.14.2, “What Results You Can Get from a Query”

## **TRUNCATE TABLE `host_cache`**

Section 20.9.9.1, “The `host_cache` Table”

## **U**

[index top [4013]]

### **UNINSTALL PLUGIN**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.3.3, “`INSTALL PLUGIN` Syntax”  
Section 13.7.5.26, “`SHOW PLUGINS` Syntax”  
Section 13.7.3.4, “`UNINSTALL PLUGIN` Syntax”  
Section 4.4.4, “`mysql_plugin` — Configure MySQL Server Plugins”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 20.14, “Performance Schema and Plugins”  
Section 14.12.1, “Pluggable Storage Engine Architecture”  
Section 22.2.2, “Plugin API Components”  
Section 22.2.4.2.1, “Server Plugin Library and Plugin Descriptors”  
Section 19.15, “The `INFORMATION_SCHEMA PLUGINS` Table”  
Section 22.2.4.6, “Writing `INFORMATION_SCHEMA` Plugins”  
Section 22.2.4.8, “Writing Audit Plugins”  
Section 22.2.4.5, “Writing Daemon Plugins”  
Section 22.2.4.4, “Writing Full-Text Parser Plugins”  
Section 22.2.4.10, “Writing Password-Validation Plugins”  
Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

### **UNION**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 8.8.2, “`EXPLAIN` Output Format”  
Section 13.2.9, “`SELECT` Syntax”  
Section 13.2.9.4, “`UNION` Syntax”  
Section 21.8.4, “C API Data Structures”  
Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section 12.14, “Information Functions”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”  
Section 11.2.5, “Numeric Type Attributes”

Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”  
Section E.5, “Restrictions on Views”  
Section 10.1.9.1, “Result Strings”  
Section 3.6.7, “Searching on Two Keys”  
Section 5.1.6, “Server Status Variables”  
Section 13.2.10, “Subquery Syntax”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 8.13.1.1, “The Range Access Method for Single-Part Indexes”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 18.5.2, “View Processing Algorithms”  
Section 18.5.1, “View Syntax”  
Section 12.11, “XML Functions”

### **UNION ALL**

Section 13.2.9.4, “`UNION` Syntax”  
Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section 12.14, “Information Functions”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 18.5.2, “View Processing Algorithms”

### **UNION DISTINCT**

Section 13.2.9.4, “`UNION` Syntax”

### **UNLOCK TABLES**

Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.3.5, “`LOCK TABLES` and `UNLOCK TABLES` Syntax”  
Section 13.3.1, “`START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax”  
Section 8.6.2, “Bulk Data Loading for `MyISAM` Tables”  
Section 7.2, “Database Backup Methods”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”  
Section 14.2.3.10, “How to Cope with Deadlocks”  
Section 13.3.5.1, “Interaction of Table Locking and Transactions”  
Section 14.2.7, “Limits on InnoDB Tables”  
Section E.1, “Restrictions on Stored Programs”  
Section 13.3.3, “Statements That Cause an Implicit Commit”  
Section 8.11.1, “System Factors and Startup Parameter Tuning”  
Section 13.3.5.3, “Table-Locking Restrictions and Conditions”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”

### **UPDATE**

Section 4.5.1.1, “`mysql` Options”  
Section 4.6.8.2, “`mysqlbinlog` Row Event Display”  
Section 13.7.2.2, “`CHECK TABLE` Syntax”  
Section 13.1.15, “`CREATE TRIGGER` Syntax”  
Section 13.1.16, “`CREATE VIEW` Syntax”  
Section 8.8.3, “`EXPLAIN EXTENDED` Output Format”

---

Section 8.8.2, “EXPLAIN Output Format”  
Section 13.8.2, “EXPLAIN Syntax”  
Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”  
Section 13.7.1.4, “GRANT Syntax”  
Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 13.2.5.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”  
Section 13.2.5, “INSERT Syntax”  
Section 13.2.9.2, “JOIN Syntax”  
Section 13.7.6.4, “KILL Syntax”  
Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 21.8.6.1, “mysql\_affected\_rows()”  
Section 21.8.6.35, “mysql\_info()”  
Section 21.8.6.37, “mysql\_insert\_id()”  
Section 21.8.6.48, “mysql\_num\_rows()”  
Section 21.8.6.49, “mysql\_options()”  
Section 21.8.10.10, “mysql\_stmt\_execute()”  
Section 21.8.10.16, “mysql\_stmt\_insert\_id()”  
Section 21.8.10.18, “mysql\_stmt\_num\_rows()”  
Section 1.8.6.1, “PRIMARY KEY and UNIQUE Index Constraints”  
Section 13.7.1.7, “SET PASSWORD Syntax”  
Section 13.3.6, “SET TRANSACTION Syntax”  
Section 13.7.5.41, “SHOW WARNINGS Syntax”  
Section 1.8.5.2, “UPDATE Differences”  
Section 13.2.11, “UPDATE Syntax”  
Section 6.2.5, “Access Control, Stage 2: Request Verification”  
Section 6.3.2, “Adding User Accounts”  
Section 16.1.2.1, “Advantages and Disadvantages of Statement-Based and Row-Based Replication”  
Section 6.3.5, “Assigning Account Passwords”  
Section 12.3.4, “Assignment Operators”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 8.6.2, “Bulk Data Loading for MyISAM Tables”  
Section 21.8.5, “C API Function Overview”  
Section 21.8.9, “C API Prepared Statement Function Overview”  
Section 21.8.16, “C API Support for Multiple Statement Execution”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 6.2.7, “Causes of Access-Denied Errors”  
Section 10.1.13, “Column Character Set Conversion”  
Section 5.4.4.2, “Configurable InnoDB Auto-Increment Locking”  
Section 14.2.3.3, “Consistent Nonlocking Reads”  
Section 1.8.6.3, “Constraints on Invalid Data”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 14.9.2.1, “Creating a FEDERATED Table Using CONNECTION”  
Section 11.5, “Data Type Default Values”  
Section 11.1.2, “Date and Time Type Overview”  
Chapter 12, *Functions and Operators*  
Section 8.12.5.2, “General Thread States”  
Section 8.2.1.2, “How MySQL Optimizes WHERE Clauses”  
Section 8.9.3.1, “How the Query Cache Operates”  
Section 12.14, “Information Functions”  
Section 8.10.1, “Internal Locking Methods”  
Section C.5.8, “Known Issues in MySQL”  
Section 14.2.3.4, “Locking Reads (SELECT ... FOR UPDATE and SELECT ... LOCK IN SHARE MODE)”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”  
Section 5.2.4.4, “Logging Format for Changes to mysql Database Tables”  
Section 12.16, “Miscellaneous Functions”  
Section 1.8.4, “MySQL Extensions to Standard SQL”  
Section 8.2.2, “Optimizing DML Statements”  
Section 8.8.1, “Optimizing Queries with EXPLAIN”  
Section 11.2.6, “Out-of-Range and Overflow Handling”  
Section 5.5.1, “Overview of Online DDL”  
Section 17.1, “Overview of Partitioning in MySQL”  
Section 17.5, “Partition Selection”  
Section 17.6.4, “Partitioning and Locking”  
Section 20.2.2, “Performance Schema Startup Configuration”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section C.5.5.2, “Problems Using DATE Columns”  
Section 16.4.1.15, “Replication and LIMIT”  
Section 16.4.1.23, “Replication and the Query Optimizer”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section C.5.4.1.2, “Resetting the Root Password: Unix Systems”  
Section C.5.4.1.1, “Resetting the Root Password: Windows Systems”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section E.5, “Restrictions on Views”  
Section 13.2.10.11, “Rewriting Subqueries as Joins”  
Section 2.10.2, “Securing the Initial MySQL Accounts”  
Section 3.3.4.1, “Selecting All Data”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 16.4.1.25, “Slave Errors During Replication”  
Section 8.2.2.2, “Speed of UPDATE Statements”  
Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”  
Section 5.7.1.12, “Statement Probes”  
Section 13.2.10.9, “Subquery Errors”  
Section 13.2.10, “Subquery Syntax”  
Section 8.10.2, “Table Locking Issues”

---

---

Section 13.3.5.3, “Table-Locking Restrictions and Conditions”  
Section 10.1.7.6, “The `_bin` and `binary` Collations”  
Section 14.6, “The `ARCHIVE` Storage Engine”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 14.8, “The `MERGE` Storage Engine”  
Section 14.3, “The `MyISAM` Storage Engine”  
Section 5.2.4, “The Binary Log”  
Section 6.2, “The MySQL Access Privilege System”  
Section 5.1.12, “The Shutdown Process”  
Section 18.3.1, “Trigger Syntax”  
Section 18.5.3, “Updatable and Insertable Views”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”  
Section 3.6.9, “Using `AUTO_INCREMENT`”  
Section 13.1.14.2, “Using `FOREIGN KEY` Constraints”  
Section 4.5.1.6.4, “Using the `--safe-updates` Option”  
Section 18.3, “Using Triggers”  
Section 1.4, “What Is New in MySQL 5.6”  
Section 21.8.14.2, “What Results You Can Get from a Query”  
Section 6.2.6, “When Privilege Changes Take Effect”  
Section 21.8.14.1, “Why `mysql_store_result()` Sometimes Returns `NULL` After `mysql_query()` Returns Success”

## **UPDATE ... ()**

Section 14.2.3.3, “Consistent Nonlocking Reads”

## **UPDATE ... WHERE ...**

Section 14.2.3.7, “Locks Set by Different SQL Statements in `InnoDB`”

## **UPDATE IGNORE**

Section 13.2.11, “`UPDATE` Syntax”

## **UPDATE t1,t2 ...**

Section 5.7.1.12, “Statement Probes”

## **USE**

Section 4.5.1.1, “`mysql` Options”  
Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Chapter 19, *INFORMATION\_SCHEMA Tables*  
Section 13.8.4, “`USE` Syntax”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 7.4.5.2, “Copy a Database from one Server to Another”  
Section 3.3.1, “Creating and Selecting a Database”  
Section 3.3, “Creating and Using a Database”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 4.5.4, “`mysqldump` — A Database Backup Program”

Section 7.4.1, “Dumping Data in SQL Format with `mysqldump`”

Section 16.2.3.1, “Evaluation of Database-Level Replication and Binary Logging Options”

Section 13.2.9.3, “Index Hint Syntax”

Section 7.4.2, “Reloading SQL-Format Backups”

Section 16.2.3.3, “Replication Rule Application”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 18.2.1, “Stored Routine Syntax”

## **USE db2**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **USE db\_name**

Section 4.5.1.1, “`mysql` Options”

## **USE test**

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

## **W**

[index top [4013]]

## **WHILE**

Section 13.6.5.3, “`ITERATE` Syntax”

Section 13.6.5.4, “`LEAVE` Syntax”

Section 13.6.5.8, “`WHILE` Syntax”

Section 13.6.5, “Flow Control Statements”

Section 13.6.2, “Statement Label Syntax”

## **X**

[index top [4013]]

## **XA COMMIT**

Section 5.1.4, “Server System Variables”

Section 13.3.7.2, “XA Transaction States”

## **XA END**

Section E.6, “Restrictions on XA Transactions”

Section 13.3.7.1, “XA Transaction SQL Syntax”

Section 13.3.7.2, “XA Transaction States”

## **XA PREPARE**

Section 13.3.7.2, “XA Transaction States”

## **XA RECOVER**

Section 13.3.7.1, “XA Transaction SQL Syntax”

Section 13.3.7.2, “XA Transaction States”

## **XA ROLLBACK**

Section 5.1.4, “Server System Variables”

---

Section 13.3.7.2, "XA Transaction States"

**XA START**

Section E.6, "Restrictions on XA Transactions"  
Section 13.3.7.1, "XA Transaction SQL Syntax"  
Section 13.3.7.2, "XA Transaction States"



---

# System Variable Index

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#)

## A

[[index top](#) [4055]]

### [audit\\_log\\_buffer\\_size](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”

### [audit\\_log\\_file](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”  
Section 6.3.11.1, “Installing the Audit Log Plugin”  
Section 6.3.11, “MySQL Enterprise Audit Log Plugin”

### [audit\\_log\\_flush](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”

### [audit\\_log\\_policy](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”  
Section 6.3.11, “MySQL Enterprise Audit Log Plugin”

### [audit\\_log\\_rotate\\_on\\_size](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”

### [audit\\_log\\_strategy](#)

Section 6.3.11.3, “Audit Log Plugin Logging Control”  
Section 6.3.11.4, “Audit Log Plugin Options and Variables”

### [authentication\\_windows\\_log\\_level](#)

Section 5.1.4, “Server System Variables”  
Section 6.3.7.4.2, “Using the Windows Authentication Plugin”

### [authentication\\_windows\\_use\\_principal\\_name](#)

Section 5.1.4, “Server System Variables”  
Section 6.3.7.4.2, “Using the Windows Authentication Plugin”

### [auto\\_increment\\_increment](#)

Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”

Section 16.1.4.2, “Replication Master Options and Variables”  
Section 5.4.4.1, “Traditional InnoDB Auto-Increment Locking”  
Section 3.6.9, “Using AUTO\_INCREMENT”

### [auto\\_increment\\_offset](#)

Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 5.4.4.1, “Traditional InnoDB Auto-Increment Locking”  
Section 3.6.9, “Using AUTO\_INCREMENT”

## AUTOCOMMIT

Section 16.4.1.29, “Replication and Transactions”

### [autocommit](#)

Section 13.2.2, “DELETE Syntax”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 14.2.3.9, “Deadlock Detection and Rollback”  
Section 5.5.5, “Examples of Online DDL”  
Section 13.3.5.1, “Interaction of Table Locking and Transactions”  
Section 14.2.7, “Limits on InnoDB Tables”  
Section 14.2.3.4, “Locking Reads (SELECT ... FOR UPDATE and SELECT ... LOCK IN SHARE MODE)”  
Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”  
Section 16.4.1.29, “Replication and Transactions”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section 5.1.4, “Server System Variables”  
Section 8.11.6.2, “Thread Pool Operation”  
Section 1.8.5.3, “Transaction and Atomic Operation Differences”

### [automatic\\_sp\\_privileges](#)

Section 13.1.4, “ALTER PROCEDURE Syntax”  
Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 5.1.4, “Server System Variables”  
Section 18.2.2, “Stored Routines and MySQL Privileges”

## B

[[index top](#) [4055]]

### [back\\_log](#)

Section 5.1.2.1, “Changes to Server Defaults”

---

Section 5.1.4, “Server System Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **basedir**

Section 13.7.3.3, “`INSTALL PLUGIN` Syntax”  
Section 5.1.4, “Server System Variables”

## **big\_tables**

Section 5.1.4, “Server System Variables”  
Section 1.4, “What Is New in MySQL 5.6”

## **bind\_address**

Section 5.1.4, “Server System Variables”

## **binlog\_cache\_size**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.1.6, “Server Status Variables”  
Section 5.2.4, “The Binary Log”

## **binlog\_checksum**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 5.2.4, “The Binary Log”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 1.4, “What Is New in MySQL 5.6”

## **binlog\_direct\_non\_transactional\_updates**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 16.4.1.29, “Replication and Transactions”

## **binlog\_format**

Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 12.7, “Date and Time Functions”  
Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”  
Section 16.1.2.3, “Determination of Safe and Unsafe Statements in Binary Logging”  
Section 12.14, “Information Functions”  
Section 5.2.4.4, “Logging Format for Changes to `mysql` Database Tables”  
Section 12.6.2, “Mathematical Functions”  
Section 12.16, “Miscellaneous Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.20, “Replication and `MEMORY` Tables”  
Section 16.4.1.29, “Replication and Transactions”  
Section 16.1.2, “Replication Formats”  
Section 16.4.1.22, “Replication of the `mysql` System Database”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section 5.1.3, “Server Command Options”  
Section 5.2.4.2, “Setting The Binary Log Format”  
Section 14.7, “The `BLACKHOLE` Storage Engine”  
Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”

## **binlog\_max\_flush\_queue\_time**

Section 16.1.4.4, “Binary Log Options and Variables”

## **binlog\_order\_commits**

Section 16.1.4.4, “Binary Log Options and Variables”

## **binlog\_row\_image**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 1.4, “What Is New in MySQL 5.6”

## **binlog\_rows\_query\_log\_events**

Section 16.1.4.4, “Binary Log Options and Variables”

## **binlog\_stmt\_cache\_size**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.1.6, “Server Status Variables”

## **bulk\_insert\_buffer\_size**

Section 14.3.1, “`MyISAM` Startup Options”  
Section 5.1.4, “Server System Variables”  
Section 8.2.2.1, “Speed of `INSERT` Statements”

## **C**

## **binlog\_direct\_non\_transactional\_updates**

[index-top (4055)]

## **character\_set\_client**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 13.7.4, “`SET` Syntax”  
Section 13.7.5.9, “`SHOW CREATE EVENT` Syntax”  
Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”  
Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”  
Section 13.7.5.19, “`SHOW EVENTS` Syntax”  
Section 13.7.5.29, “`SHOW PROCEDURE STATUS` Syntax”  
Section 13.7.5.39, “`SHOW TRIGGERS` Syntax”  
Section 21.8.8.1, “C API Prepared Statement Type Codes”  
Section 10.5, “Character Set Configuration”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 19.7, “The `INFORMATION_SCHEMA EVENTS` Table”  
Section 19.19, “The `INFORMATION_SCHEMA ROUTINES` Table”  
Section 19.27, “The `INFORMATION_SCHEMA TRIGGERS` Table”  
Section 19.29, “The `INFORMATION_SCHEMA VIEWS` Table”  
Section 5.2.4, “The Binary Log”



---

## **character\_set\_connection**

Section 10.1.9.2, “`CONVERT( )` and `CAST( )`”  
Section 13.7.4, “`SET` Syntax”  
Section 10.1.3.5, “Character String Literal Character Set and Collation”  
Section 10.1.7.5, “Collation of Expressions”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 12.7, “Date and Time Functions”  
Section 12.13, “Encryption and Compression Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 10.7, “MySQL Server Locale Support”  
Section 16.4.1.32, “Replication and Variables”  
Section 10.1.9.1, “Result Strings”  
Section 5.1.4, “Server System Variables”  
Section 9.1.1, “String Literals”  
Section 10.1.8, “String Repertoire”  
Section 12.2, “Type Conversion in Expression Evaluation”

## **character\_set\_database**

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 13.7.4, “`SET` Syntax”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”

## **character\_set\_filesystem**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 13.2.9.1, “`SELECT ... INTO` Syntax”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 12.5, “String Functions”

## **character\_set\_results**

Section 13.7.4, “`SET` Syntax”  
Section 10.1.6, “Character Set for Error Messages”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 5.1.4, “Server System Variables”  
Section 10.1.12, “UTF-8 for Metadata”

## **character\_set\_server**

Section 10.5, “Character Set Configuration”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 10.1.3.2, “Database Character Set and Collation”  
Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.4.1.2, “Replication and Character Sets”  
Section 16.4.1.32, “Replication and Variables”  
Section 10.1.3.1, “Server Character Set and Collation”  
Section 5.1.4, “Server System Variables”

## **character\_set\_system**

Section 10.5, “Character Set Configuration”  
Section 5.1.4, “Server System Variables”  
Section 10.1.12, “UTF-8 for Metadata”

## **character\_sets\_dir**

Section 10.4.3, “Adding a Simple Collation to an 8-Bit Character Set”  
Section 10.4.4.1, “Defining a UCA Collation using LDML Syntax”  
Section 5.1.4, “Server System Variables”

## **collation\_connection**

Section 10.1.9.2, “`CONVERT( )` and `CAST( )`”  
Section 13.7.4, “`SET` Syntax”  
Section 13.7.5.9, “`SHOW CREATE EVENT` Syntax”  
Section 13.7.5.11, “`SHOW CREATE PROCEDURE` Syntax”  
Section 13.7.5.14, “`SHOW CREATE VIEW` Syntax”  
Section 13.7.5.19, “`SHOW EVENTS` Syntax”  
Section 13.7.5.29, “`SHOW PROCEDURE STATUS` Syntax”  
Section 13.7.5.39, “`SHOW TRIGGERS` Syntax”  
Section 10.1.3.5, “Character String Literal Character Set and Collation”  
Section 10.1.7.5, “Collation of Expressions”  
Section 10.1.4, “Connection Character Sets and Collations”  
Section 12.7, “Date and Time Functions”  
Section 12.13, “Encryption and Compression Functions”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 10.1.9.1, “Result Strings”  
Section 5.1.4, “Server System Variables”  
Section 19.7, “The `INFORMATION_SCHEMA` `EVENTS` Table”  
Section 19.19, “The `INFORMATION_SCHEMA` `ROUTINES` Table”  
Section 19.27, “The `INFORMATION_SCHEMA` `TRIGGERS` Table”  
Section 19.29, “The `INFORMATION_SCHEMA` `VIEWS` Table”  
Section 5.2.4, “The Binary Log”  
Section 12.2, “Type Conversion in Expression Evaluation”

## **collation\_database**

Section 13.1.12, “`CREATE PROCEDURE` and `CREATE FUNCTION` Syntax”  
Section 10.1.4, “Connection Character Sets and Collations”

---

Section 10.1.3.2, “Database Character Set and Collation”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.4.1.32, “Replication and Variables”

Section 5.1.4, “Server System Variables”

Section 5.2.4, “The Binary Log”

## **collation\_server**

Section 10.1.4, “Connection Character Sets and Collations”

Section 10.1.3.2, “Database Character Set and Collation”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.4.1.2, “Replication and Character Sets”

Section 16.4.1.32, “Replication and Variables”

Section 10.1.3.1, “Server Character Set and Collation”

Section 5.1.4, “Server System Variables”

Section 5.2.4, “The Binary Log”

## **completion\_type**

Section 21.8.6.6, “mysql\_commit()”

Section 21.8.6.58, “mysql\_rollback()”

Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”

Section 5.1.4, “Server System Variables”

## **concurrent\_insert**

Section 8.10.3, “Concurrent Inserts”

Section 8.10.1, “Internal Locking Methods”

Section 8.6.1, “Optimizing MyISAM Queries”

Section 8.2.5, “Other Optimization Tips”

Section 5.1.4, “Server System Variables”

## **connect\_timeout**

Section C.5.2.3, “Lost connection to MySQL server”

Section 21.8.6.53, “mysql\_real\_connect()”

Section C.5.2.11, “Communication Errors and Aborted Connections”

Section 5.1.4, “Server System Variables”

## **core\_file**

Section 5.1.4, “Server System Variables”

## **D**

[index top [4055]]

## **daemon\_memcached\_engine\_lib\_name**

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

## **daemon\_memcached\_engine\_lib\_path**

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

## **daemon\_memcached\_option**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

Section 14.2.9.4.1, “Password-Protecting the memcached Interface through SASL”

Section 14.2.9.8, “Troubleshooting the InnoDB memcached Plugin”

## **daemon\_memcached\_r\_batch\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

Section 14.2.9.5.6, “Performing DML and DDL Statements on the Underlying InnoDB Table”

Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”

## **daemon\_memcached\_w\_batch\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”

Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”

Section 14.2.9.3.2, “Installing and Configuring the InnoDB memcached Plugin”

Section 14.2.9.5.6, “Performing DML and DDL Statements on the Underlying InnoDB Table”

Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”

## **DATADIR**

Section 5.4.1.2, “Specifying the Location of a Tablespace”

## **datadir**

Section 2.3, “Installing MySQL on Microsoft Windows”

Section 5.1.4, “Server System Variables”

## **date\_format**

Section 5.1.4, “Server System Variables”

## **datetime\_format**

Section 5.1.4, “Server System Variables”

## **debug**

Section 5.1.4, “Server System Variables”

---

## **debug\_sync**

Section 5.1.4, “Server System Variables”

## **default\_storage\_engine**

Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 14.1, “Setting the Storage Engine”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”

## **default\_tmp\_storage\_engine**

Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 14.1, “Setting the Storage Engine”

## **default\_week\_format**

Section 12.7, “Date and Time Functions”  
Section 17.6.3, “Partitioning Limitations Relating to Functions”  
Section 5.1.4, “Server System Variables”

## **delay\_key\_write**

Section 13.1.14, “CREATE TABLE Syntax”  
Section 5.1.4, “Server System Variables”

## **delayed\_insert\_limit**

Section 13.2.5.2, “INSERT DELAYED Syntax”  
Section 5.1.4, “Server System Variables”

## **delayed\_insert\_timeout**

Section 13.2.5.2, “INSERT DELAYED Syntax”  
Section 5.1.4, “Server System Variables”

## **delayed\_queue\_size**

Section 13.2.5.2, “INSERT DELAYED Syntax”  
Section 5.1.4, “Server System Variables”

## **disable\_gtid\_unsafe\_statements**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

## **disconnect\_on\_expired\_passwords**

Section 6.3.6, “Password Expiration and Sandbox Mode”  
Section 5.1.4, “Server System Variables”

## **div\_precision\_increment**

Section 12.6.1, “Arithmetic Operators”  
Section 5.1.4, “Server System Variables”

## **E**

[index top [4055]]

## **end\_markers\_in\_json**

Section 5.1.4, “Server System Variables”

## **enforce\_gtid\_consistency**

Section 16.1.4.5, “Global Transaction ID Options and Variables”

## **engine\_condition\_pushdown**

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **eq\_range\_index\_dive\_limit**

Section 8.13.1.3, “Equality Range Optimization of Many-Valued Comparisons”  
Section 5.1.4, “Server System Variables”

## **error\_count**

Section 13.7.5.18, “SHOW ERRORS Syntax”  
Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”  
Section 5.1.4, “Server System Variables”  
Section C.1, “Sources of Error Information”

## **event\_scheduler**

Section 18.4.2, “Event Scheduler Configuration”  
Section 21.7.2, “Restrictions When Using the Embedded MySQL Server”  
Section 5.1.4, “Server System Variables”  
Section 18.4.6, “The Event Scheduler and MySQL Privileges”

## **expire\_logs\_days**

Section 13.4.1.1, “PURGE BINARY LOGS Syntax”  
Section 5.2.6, “Server Log Maintenance”  
Section 5.1.4, “Server System Variables”

## **explicit\_defaults\_for\_timestamp**

Section 11.3.5, “Automatic Initialization and Updating for TIMESTAMP and DATETIME”  
Section 11.5, “Data Type Default Values”  
Section 11.1.2, “Date and Time Type Overview”  
Section 5.1.4, “Server System Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 1.4, “What Is New in MySQL 5.6”

## **external\_user**

Section 22.2.4.9.4, “Implementing Proxy User Support in Authentication Plugins”  
Section 6.3.8, “Proxy Users”  
Section 5.1.4, “Server System Variables”  
Section 22.2.4.9.1, “Writing the Server-Side Authentication Plugin”

## **F**

[index top [4055]]

---

## flush

Section 5.1.4, “Server System Variables”

## flush\_time

Section 5.1.2.1, “Changes to Server Defaults”  
Section 5.1.4, “Server System Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## foreign\_key\_checks

Section 5.5.6, “Implementation Details of Online DDL”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 5.5.1, “Overview of Online DDL”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4, “The Binary Log”  
Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## ft\_boolean\_syntax

Section 12.9.2, “Boolean Full-Text Searches”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 5.1.4, “Server System Variables”

## ft\_max\_word\_len

Section 12.9.2, “Boolean Full-Text Searches”  
Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 22.2.3.2, “Full-Text Parser Plugins”  
Section 5.1.4, “Server System Variables”

## ft\_min\_word\_len

Section 12.9.2, “Boolean Full-Text Searches”  
Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 22.2.3.2, “Full-Text Parser Plugins”  
Section 12.9.1, “Natural Language Full-Text Searches”  
Section 5.1.4, “Server System Variables”

## ft\_query\_expansion\_limit

Section 5.1.4, “Server System Variables”

## ft\_stopword\_file

Section 12.9.2, “Boolean Full-Text Searches”  
Section 16.1.1.6, “Creating a Data Snapshot Using Raw Data Files”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 12.9, “Full-Text Search Functions”  
Section 12.9.1, “Natural Language Full-Text Searches”  
Section 5.1.4, “Server System Variables”

## G

[index top [4055]]

## general\_log

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 5.2.3, “The General Query Log”

## general\_log\_file

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 5.2.3, “The General Query Log”

## group\_concat\_max\_len

Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section 5.1.4, “Server System Variables”

## gtid\_done

Section 13.4.1.2, “RESET MASTER Syntax”  
Section 13.7.5.24, “SHOW MASTER STATUS Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”

## gtid\_executed

Section 13.4.1.2, “RESET MASTER Syntax”  
Section 13.7.5.24, “SHOW MASTER STATUS Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.3.1, “GTID Concepts”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## gtid\_lost

Section 13.4.1.2, “RESET MASTER Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”

## gtid\_mode

Section 13.4.2.1, “CHANGE MASTER TO Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 12.15, “GTID Functions”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## gtid\_next

Section 13.1.5, “ALTER SERVER Syntax”  
Section 13.7.2.1, “ANALYZE TABLE Syntax”  
Section 13.7.6.2, “CACHE INDEX Syntax”  
Section 13.4.2.1, “CHANGE MASTER TO Syntax”  
Section 13.7.2.2, “CHECK TABLE Syntax”

---

Section 13.1.13, “`CREATE SERVER` Syntax”  
Section 13.1.22, “`DROP SERVER` Syntax”  
Section 13.7.6.3, “`FLUSH` Syntax”  
Section 13.7.6.5, “`LOAD INDEX INTO CACHE` Syntax”  
Section 13.7.2.4, “`OPTIMIZE TABLE` Syntax”  
Section 13.7.2.5, “`REPAIR TABLE` Syntax”  
Section 13.7.6.6, “`RESET` Syntax”  
Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 13.4.2.6, “`STOP SLAVE` Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.3.1, “GTID Concepts”

## `gtid_owned`

Section 16.1.4.5, “Global Transaction ID Options and Variables”

## `gtid_purged`

Section 13.4.1.2, “`RESET MASTER` Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.3.1, “GTID Concepts”  
Section 16.1.3.3, “Using GTIDs for Failover and Scaleout”

## H

[index top [4055]]

## `have_compress`

Section 5.1.4, “Server System Variables”

## `have_crypt`

Section 5.1.4, “Server System Variables”

## `have_csv`

Section 5.1.4, “Server System Variables”

## `have_dynamic_loading`

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”  
Section 5.1.4, “Server System Variables”

## `have_geometry`

Section 5.1.4, “Server System Variables”

## `have_innodb`

Section 5.1.4, “Server System Variables”

## `have_openssl`

Section 5.1.4, “Server System Variables”

## `have_partitioning`

Chapter 17, *Partitioning*

Section 5.1.4, “Server System Variables”

## `have_profiling`

Section 5.1.4, “Server System Variables”

## `have_query_cache`

Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.4, “Server System Variables”

## `have_rtree_keys`

Section 5.1.4, “Server System Variables”

## `have_ssl`

Section 6.3.9.2, “Configuring MySQL for SSL”  
Section 5.1.4, “Server System Variables”

## `have_symlink`

Section 5.1.4, “Server System Variables”  
Section 8.11.3.1.2, “Using Symbolic Links for `MyISAM` Tables on Unix”  
Section 8.11.3.1.3, “Using Symbolic Links for Databases on Windows”

## `host_cache_size`

Section 5.1.2.1, “Changes to Server Defaults”  
Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 1.4, “What Is New in MySQL 5.6”

## `hostname`

Section 5.1.4, “Server System Variables”

## I

[index top [4055]]

## `identity`

Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”

## `ignore_db_dirs`

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## `init_connect`

Section 10.1.5, “Configuring the Character Set and Collation for Applications”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.4, “Server System Variables”  
Section 20.9.9.1, “The `host_cache` Table”



---

## **init\_file**

Section 5.1.4, “Server System Variables”

## **init\_slave**

Section 16.1.4.3, “Replication Slave Options and Variables”

## **innodb**

Section 14.2.5.1, “Support for Read-Only Media”  
Section 1.4, “What Is New in MySQL 5.6”

## **innodb\_adaptive\_flushing**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_adaptive\_hash\_index**

Section 14.2.3.12.6, “Adaptive Hash Indexes”  
Section 14.2.5.5.4, “Dynamically Changing innodb\_adaptive\_hash\_index”

## **innodb\_adaptive\_max\_sleep\_delay**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_api\_bk\_commit\_interval**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”

## **innodb\_api\_enable\_binlog**

Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”

## **innodb\_api\_enable\_md1**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”

## **innodb\_api\_trx\_level**

Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”

## **innodb\_autoextend\_increment**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 14.2.1.2, “Configuring InnoDB”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_autoinc\_lock\_mode**

Section 8.5.4, “Bulk Data Loading for InnoDB Tables”

Section 5.4.4.2, “Configurable InnoDB Auto-Increment Locking”

## **innodb\_buffer\_pool\_dump\_at\_shutdown**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_buffer\_pool\_dump\_now**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_buffer\_pool\_instances**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 8.9.1, “The InnoDB Buffer Pool”  
Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_buffer\_pool\_load\_abort**

Section 5.1.6, “Server Status Variables”

## **innodb\_buffer\_pool\_load\_at\_startup**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.6, “Server Status Variables”

## **innodb\_buffer\_pool\_load\_now**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.6, “Server Status Variables”

## **innodb\_buffer\_pool\_size**

Section 14.2.3.14.1, “InnoDB Error Codes”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 8.5.7, “Optimizing InnoDB Disk I/O”  
Section 5.1.6, “Server Status Variables”  
Section 8.9.1, “The InnoDB Buffer Pool”  
Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

## **innodb\_change\_buffer\_max\_size**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_change\_buffering**

Section 8.5.7, “Optimizing InnoDB Disk I/O”  
Section 8.5.2, “Optimizing InnoDB Transaction Management”

---

Section 14.2.5.1, “Support for Read-Only Media”

### **innodb\_checksum\_algorithm**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.1.2.1, “Changes to Server Defaults”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

### **innodb\_checksums**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **innodb\_cmp\_per\_index\_enabled**

Section 5.4.6.4, “Monitoring Compression at Runtime”

Section 19.30.2, “The `INFORMATION_SCHEMA`

`INNODB_CMP_PER_INDEX` and

`INNODB_CMP_PER_INDEX_RESET` Tables”

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

### **innodb\_compression\_failure\_threshold\_pct**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.4.6.5, “How Compression Works for InnoDB Tables”

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

### **innodb\_compression\_level**

Section 5.4.6.5, “How Compression Works for InnoDB Tables”

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

### **innodb\_compression\_pad\_pct\_max**

Section 5.4.6.5, “How Compression Works for InnoDB Tables”

Section 5.4.6.3, “Tuning Compression for InnoDB Tables”

### **innodb\_concurrency\_tickets**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.1.2.1, “Changes to Server Defaults”

Section 8.5.8, “Optimizing InnoDB Configuration Variables”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

### **innodb\_data\_file\_path**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.2.2, “Adding, Removing, or Resizing InnoDB Data and Log Files”

Section 5.1.2.1, “Changes to Server Defaults”

Section 14.2.1.2, “Configuring InnoDB”

Section 4.4.3, “`mysql_install_db` — Initialize MySQL Data Directory”

Section 14.2.4.5.1, “Troubleshooting InnoDB I/O Problems”

Section 14.2.2.3, “Using Raw Disk Partitions for the Shared Tablespace”

### **innodb\_data\_home\_dir**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.1.2, “Configuring InnoDB”

Section 14.2.4.5.1, “Troubleshooting InnoDB I/O Problems”

### **innodb\_doublewrite**

Section 14.2.8, “MySQL and the ACID Model”

Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

### **innodb\_fast\_shutdown**

Section 14.2.2.2, “Adding, Removing, or Resizing

InnoDB Data and Log Files”

Section 5.4.7.2.1, “Compatibility Check When InnoDB Is Started”

Section 5.1.12, “The Shutdown Process”

### **innodb\_file\_format**

Section 13.1.14, “`CREATE TABLE` Syntax”

Section 5.4.8.3, “`DYNAMIC` and `COMPRESSED` Row Formats”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.4.7.2.1, “Compatibility Check When InnoDB Is Started”

Section 5.4.7.2.2, “Compatibility Check When a Table Is Opened”

Section 5.4.3, “Converting Tables from `MyISAM` to InnoDB”

Section 14.2.5.5.1, “Dynamically Changing `innodb_file_per_table`”

Section 5.4.6.2, “Enabling Compression for a Table”

Section 5.4.7.1, “Enabling File Formats”

Section 5.4.7.3, “Identifying the File Format in Use”

Section 5.4.8.1, “Overview of InnoDB Row Storage”

Section 5.4.6.6, “SQL Compression Syntax Warnings and Errors”

Section 14.2.5.4, “The Barracuda File Format”

### **innodb\_file\_format\_check**

Section 5.4.7.2.1, “Compatibility Check When InnoDB Is Started”

Section 5.4.7.2.2, “Compatibility Check When a Table Is Opened”

### **innodb\_file\_format\_max**

Section 14.2.6, “InnoDB Startup Options and System Variables”

---

## **innodb\_file\_per\_table**

Section 13.1.14, “CREATE TABLE Syntax”  
Section 13.7.6.3, “FLUSH Syntax”  
Section 14.2.2.6, “InnoDB and MySQL Replication”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.2.4.4.2, “InnoDB Tablespace Monitor Output”  
Section 13.7.2.4, “OPTIMIZE TABLE Syntax”  
Section 14.2.5.6, “TRUNCATE TABLE Reclaims Space”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 5.4.3, “Converting Tables from MyISAM to InnoDB”  
Section 14.2.2.1, “Creating the InnoDB Tablespace”  
Section 14.2.5.5.1, “Dynamically Changing innodb\_file\_per\_table”  
Section 5.4.1.1, “Enabling and Disabling File-Per-Table Mode”  
Section 5.4.6.2, “Enabling Compression for a Table”  
Section 5.5.6, “Implementation Details of Online DDL”  
Section 5.4.1, “Managing InnoDB Tablespaces”  
Section 14.2.8, “MySQL and the ACID Model”  
Section 5.4.8.1, “Overview of InnoDB Row Storage”  
Section 14.2.3.12.7, “Physical Row Structure”  
Section 16.3.4, “Replicating Different Databases to Different Slaves”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.4.6.6, “SQL Compression Syntax Warnings and Errors”  
Section 14.2, “The InnoDB Storage Engine”  
Section 14.2.4.7, “Troubleshooting InnoDB Data Dictionary Operations”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_flush\_log\_at\_timeout**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_flush\_log\_at\_trx\_commit**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 14.2.8, “MySQL and the ACID Model”  
Section 8.5.2, “Optimizing InnoDB Transaction Management”  
Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

## **innodb\_flush\_method**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 8.5.7, “Optimizing InnoDB Disk I/O”  
Section 5.1.6, “Server Status Variables”  
Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

## **innodb\_flush\_neighbors**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_force\_load\_corrupted**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_force\_recovery**

Section 1.7, “How to Report Bugs or Problems”  
Section 8.5.2, “Optimizing InnoDB Transaction Management”  
Section 14.2.4.6, “Starting InnoDB on a Corrupted Database”

## **innodb\_ft\_aux\_table**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 19.30.23, “The INFORMATION\_SCHEMA INNODB\_FT\_BEING\_DELETED Table”  
Section 19.30.22, “The INFORMATION\_SCHEMA INNODB\_FT\_DELETED Table”  
Section 19.30.21, “The INFORMATION\_SCHEMA INNODB\_FT\_INDEX\_CACHE Table”  
Section 19.30.20, “The INFORMATION\_SCHEMA INNODB\_FT\_INDEX\_TABLE Table”

## **innodb\_ft\_cache\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_ft\_enable\_stopword**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 12.9.2, “Boolean Full-Text Searches”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 12.9, “Full-Text Search Functions”  
Section 12.9.1, “Natural Language Full-Text Searches”

## **innodb\_ft\_max\_token\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 12.9.2, “Boolean Full-Text Searches”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

## **innodb\_ft\_min\_token\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 12.9.2, “Boolean Full-Text Searches”  
Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”  
Section 12.9.1, “Natural Language Full-Text Searches”

## **innodb\_ft\_num\_word\_optimize**

Section 14.2.3.12.3, “FULLTEXT Indexes”  
Section 14.2.6, “InnoDB Startup Options and System Variables”



---

Section 13.7.2.4, “OPTIMIZE TABLE Syntax”

### **`innodb_ft_server_stopword_table`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 12.9.2, “Boolean Full-Text Searches”

Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

Section 12.9, “Full-Text Search Functions”

Section 12.9.4, “Full-Text Stopwords”

Section 12.9.1, “Natural Language Full-Text Searches”

Section 19.30.19, “The `INFORMATION_SCHEMA`

`INNODB_FT_DEFAULT_STOPWORD` Table”

### **`innodb_ft_sort_pll_degree`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **`innodb_ft_user_stopword_table`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 12.9.2, “Boolean Full-Text Searches”

Section 12.9.6, “Fine-Tuning MySQL Full-Text Search”

Section 12.9, “Full-Text Search Functions”

Section 12.9.4, “Full-Text Stopwords”

Section 12.9.1, “Natural Language Full-Text Searches”

Section 19.30.19, “The `INFORMATION_SCHEMA`

`INNODB_FT_DEFAULT_STOPWORD` Table”

### **`innodb_io_capacity`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **`innodb_io_capacity_max`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **`innodb_large_prefix`**

Section 2.11.3, “Checking Whether Tables or Indexes Must Be Rebuilt”

Section 14.2.7, “Limits on InnoDB Tables”

### **`innodb_lock_wait_timeout`**

Section 14.2.3.14.1, “InnoDB Error Codes”

Section 14.2.3.9, “Deadlock Detection and Rollback”

Section 14.2.5.5.3, “Dynamically Changing

`innodb_lock_wait_timeout`”

Section 16.4.1.27, “Replication Retries and Timeouts”

Section 16.1.4.3, “Replication Slave Options and

Variables”

### **`innodb_locks_unsafe_for_binlog`**

Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 13.3.6, “SET TRANSACTION Syntax”

Section 14.2.3.3, “Consistent Nonlocking Reads”

Section 14.2.3.7, “Locks Set by Different SQL

Statements in InnoDB”

### **`innodb_log_buffer_size`**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **`innodb_log_file_size`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.1.2.1, “Changes to Server Defaults”

Section 14.2.1.2, “Configuring InnoDB”

Section 4.4.3, “`mysql_install_db` — Initialize

MySQL Data Directory”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

Section 14.2.5.1, “Support for Read-Only Media”

Section 1.4, “What Is New in MySQL 5.6”

### **`innodb_log_files_in_group`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 1.4, “What Is New in MySQL 5.6”

### **`innodb_log_group_home_dir`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **`innodb_lru_scan_depth`**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **`innodb_max_dirty_pages_pct`**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **`innodb_max_purge_lag`**

Section 14.2.3.11, “InnoDB Multi-Versioning”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **`innodb_max_purge_lag_delay`**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **`innodb_monitor_disable`**

Section 19.30.17, “The `INFORMATION_SCHEMA`

`INNODB_METRICS` Table”

### **`innodb_monitor_enable`**

Section 19.30.17, “The `INFORMATION_SCHEMA`

`INNODB_METRICS` Table”

### **`innodb_monitor_reset`**

Section 19.30.17, “The `INFORMATION_SCHEMA`

`INNODB_METRICS` Table”

---

## **innodb\_monitor\_reset\_all**

Section 19.30.17, “The `INFORMATION_SCHEMA` `INNODB_METRICS` Table”

## **innodb\_old\_blocks\_pct**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 8.9.1, “The InnoDB Buffer Pool”

## **innodb\_old\_blocks\_time**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.1.2.1, “Changes to Server Defaults”

Section 8.9.1, “The InnoDB Buffer Pool”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_online\_alter\_log\_max\_size**

Section 5.5.6, “Implementation Details of Online DDL”

## **innodb\_open\_files**

Section 5.1.2.1, “Changes to Server Defaults”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_optimize\_fulltext\_only**

Section 14.2.3.12.3, “FULLTEXT Indexes”

Section 13.7.2.4, “OPTIMIZE TABLE Syntax”

## **innodb\_page\_size**

Section 5.4.3, “Converting Tables from MyISAM to InnoDB”

Section 2.11.2.1, “Downgrading to MySQL 5.5”

Section 5.3.2, “File Space Management”

Section 5.4.6.5, “How Compression Works for InnoDB Tables”

Section 14.2.7, “Limits on InnoDB Tables”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

Section 14.2.3.12.4, “Physical Structure of an InnoDB Index”

Section 14.2.9.8, “Troubleshooting the InnoDB memcached Plugin”

Section 1.4, “What Is New in MySQL 5.6”

## **innodb\_print\_all\_deadlocks**

Section 14.2.4.5, “InnoDB General Troubleshooting”

Section 5.4.3, “Converting Tables from MyISAM to InnoDB”

Section 14.2.3.10, “How to Cope with Deadlocks”

## **innodb\_random\_read\_ahead**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_read\_ahead\_threshold**

Section 14.2.5.11, “More Read-Ahead Statistics”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_read\_io\_threads**

Section 14.2.4.4.1, “InnoDB Standard Monitor and Lock Monitor Output”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 22.1.1, “MySQL Threads”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_rollback\_segments**

Section 8.5.7, “Optimizing InnoDB Disk I/O”

## **innodb\_sort\_buffer\_size**

Section 14.2.6, “InnoDB Startup Options and System Variables”

## **innodb\_stats\_auto\_recalc**

Section 13.1.14, “CREATE TABLE Syntax”

## **innodb\_stats\_method**

Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”

## **innodb\_stats\_on\_metadata**

Section 5.1.2.1, “Changes to Server Defaults”

Section 14.2.5.5.2, “Dynamically Changing `innodb_stats_on_metadata`”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **innodb\_stats\_persistent**

Section 13.7.2.1, “ANALYZE TABLE Syntax”

Section 13.1.11, “CREATE INDEX Syntax”

Section 13.1.14, “CREATE TABLE Syntax”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

## **innodb\_stats\_persistent\_sample\_pages**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

## **innodb\_stats\_sample\_pages**

Section 14.2.5.8, “Controlling Optimizer Statistics Estimation”

## **innodb\_stats\_transient\_sample\_pages**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

## **innodb\_strict\_mode**

Section 13.1.11, “CREATE INDEX Syntax”

Section 13.1.14, “CREATE TABLE Syntax”

Section 14.2.5.7, “InnoDB Strict Mode”

---

Section 5.4.6.5, “How Compression Works for InnoDB Tables”

Section 5.1.7, “Server SQL Modes”

Section 5.4.6.6, “SQL Compression Syntax Warnings and Errors”

### **innodb\_support\_xa**

Section 8.5.2, “Optimizing InnoDB Transaction Management”

Section 14.2.9.5.3, “Tuning Performance of the InnoDB memcached Plugin”

### **innodb\_table\_locks**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.7, “Limits on InnoDB Tables”

### **innodb\_thread\_concurrency**

Section 14.2.4.4.1, “InnoDB Standard Monitor and Lock Monitor Output”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 8.5.8, “Optimizing InnoDB Configuration Variables”

### **innodb\_thread\_sleep\_delay**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **innodb\_undo\_directory**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.5.1, “Support for Read-Only Media”

### **innodb\_undo\_logs**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 5.1.6, “Server Status Variables”

### **innodb\_undo\_tablespaces**

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 14.2.5.1, “Support for Read-Only Media”

### **innodb\_use\_native\_aio**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **innodb\_use\_sys\_malloc**

Section 14.2.6, “InnoDB Startup Options and System Variables”

### **innodb\_write\_io\_threads**

Section 14.2.4.4.1, “InnoDB Standard Monitor and Lock Monitor Output”

Section 14.2.6, “InnoDB Startup Options and System Variables”

Section 22.1.1, “MySQL Threads”

Section 8.5.7, “Optimizing InnoDB Disk I/O”

### **insert\_id**

Section 14.9.3, “FEDERATED Storage Engine Notes and Tips”

Section 5.1.4, “Server System Variables”

### **interactive\_timeout**

Section 21.8.6.53, “mysql\_real\_connect()”

Section C.5.2.11, “Communication Errors and Aborted Connections”

Section 5.1.4, “Server System Variables”

## **J**

[index top [4055]]

### **join\_buffer\_size**

Section 5.1.2.1, “Changes to Server Defaults”

Section 8.13.8, “Nested-Loop Join Algorithms”

Section 5.1.4, “Server System Variables”

Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **K**

[index top [4055]]

### **keep\_files\_on\_create**

Section 5.1.4, “Server System Variables”

### **key\_buffer\_size**

Section 8.6.2, “Bulk Data Loading for MyISAM Tables”

Section 5.4.3, “Converting Tables from MyISAM to InnoDB”

Section 8.8.4, “Estimating Query Performance”

Section 8.11.4.1, “How MySQL Uses Memory”

Section 7.6.3, “How to Repair MyISAM Tables”

Section C.5.8, “Known Issues in MySQL”

Section 8.9.2.2, “Multiple Key Caches”

Section 8.9.2.6, “Restructuring a Key Cache”

Section 5.1.3, “Server Command Options”

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

Section 8.2.2.3, “Speed of DELETE Statements”

Section 8.6.3, “Speed of REPAIR TABLE Statements”

Section 5.1.5.1, “Structured System Variables”

Section 8.9.2, “The MyISAM Key Cache”

Section 8.11.2, “Tuning Server Parameters”

Section 4.2.3.3, “Using Option Files”

### **key\_cache\_age\_threshold**

Section 8.9.2.3, “Midpoint Insertion Strategy”

Section 5.1.4, “Server System Variables”

---

Section 5.1.5.1, “Structured System Variables”

## **key\_cache\_block\_size**

Section 8.9.2.5, “Key Cache Block Size”

Section 8.9.2.6, “Restructuring a Key Cache”

Section 5.1.4, “Server System Variables”

Section 5.1.5.1, “Structured System Variables”

## **key\_cache\_division\_limit**

Section 8.9.2.3, “Midpoint Insertion Strategy”

Section 5.1.4, “Server System Variables”

Section 5.1.5.1, “Structured System Variables”

## **L**

[index top [4055]]

## **large\_files\_support**

Section 17.6, “Restrictions and Limitations on Partitioning”

Section 5.1.4, “Server System Variables”

## **large\_page\_size**

Section 5.1.4, “Server System Variables”

## **large\_pages**

Section 5.1.4, “Server System Variables”

## **last\_insert\_id**

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 16.4.1.32, “Replication and Variables”

Section 5.1.4, “Server System Variables”

## **lc\_messages**

Section 5.1.4, “Server System Variables”

Section 10.2, “Setting the Error Message Language”

## **lc\_messages\_dir**

Section 5.1.4, “Server System Variables”

Section 10.2, “Setting the Error Message Language”

## **lc\_time\_names**

Section 12.7, “Date and Time Functions”

Section 5.2.4.3, “Mixed Binary Logging Format”

Section 10.7, “MySQL Server Locale Support”

Section 16.4.1.32, “Replication and Variables”

Section 5.1.4, “Server System Variables”

Section 12.5, “String Functions”

## **license**

Section 5.1.4, “Server System Variables”

## **local**

Section 13.2.6, “LOAD DATA INFILE Syntax”

Section 13.2.7, “LOAD XML Syntax”

Section 2.9.4, “MySQL Source-Configuration Options”

Section 6.1.6, “Security Issues with LOAD DATA LOCAL”

## **local\_infile**

Section 5.1.4, “Server System Variables”

## **lock\_wait\_timeout**

Section 5.1.4, “Server System Variables”

## **locked\_in\_memory**

Section 5.1.4, “Server System Variables”

## **log**

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 14.2.9.6, “Using the InnoDB memcached Plugin with Replication”

## **log\_bin**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 1.4, “What Is New in MySQL 5.6”

## **log\_bin\_basename**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 1.4, “What Is New in MySQL 5.6”

## **log\_bin\_trust\_function\_creators**

Section 16.1.4.4, “Binary Log Options and Variables”

Section 18.7, “Binary Logging of Stored Programs”

Section 5.1.4, “Server System Variables”

## **log\_bin\_use\_v**

Section 16.1.4.4, “Binary Log Options and Variables”

## **log\_error**

Section 5.1.4, “Server System Variables”

## **log\_output**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”

Section 5.1.4, “Server System Variables”

Section 5.2.3, “The General Query Log”

Section 5.2.5, “The Slow Query Log”

## **log\_queries\_not\_using\_indexes**

Section 5.1.4, “Server System Variables”

Section 5.2.5, “The Slow Query Log”

## **log\_slave\_updates**

Section 16.1.4.4, “Binary Log Options and Variables”

## **log\_slow\_admin\_statements**

Section 5.1.4, “Server System Variables”

---

## log\_slow\_queries

Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## log\_slow\_slave\_statements

Section 16.1.4.3, “Replication Slave Options and Variables”

## log\_throttle\_queries\_not\_using\_indexes

Section 5.1.4, “Server System Variables”  
Section 5.2.5, “The Slow Query Log”

## log\_warnings

Section 5.1.4, “Server System Variables”  
Section 5.2.2, “The Error Log”

## long\_query\_time

Section 4.5.2, “mysqladmin — Client for Administering a MySQL Server”  
Section 5.2, “MySQL Server Logs”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 5.2.5, “The Slow Query Log”

## low\_priority\_updates

Section 5.1.4, “Server System Variables”  
Section 8.10.2, “Table Locking Issues”  
Section 1.4, “What Is New in MySQL 5.6”

## lower\_case\_file\_system

Section 5.1.4, “Server System Variables”

## lower\_case\_table\_names

Section 13.7.1.4, “GRANT Syntax”  
Section 13.7.1.6, “REVOKE Syntax”  
Section 13.7.5.38, “SHOW TABLES Syntax”  
Section 10.1.7.9, “Collation and INFORMATION\_SCHEMA Searches”  
Section 16.2.3, “How Servers Evaluate Replication Filtering Rules”  
Section 1.7, “How to Report Bugs or Problems”  
Section 9.2.2, “Identifier Case Sensitivity”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 19.30.9, “The INFORMATION\_SCHEMA INNODB\_SYS\_COLUMNS Table”  
Section 19.30.7, “The INFORMATION\_SCHEMA INNODB\_SYS\_TABLES Table”  
Section 13.1.14.2, “Using FOREIGN KEY Constraints”

## M

[index top [4055]]

## master\_info\_repository

Section 16.1.4.3, “Replication Slave Options and Variables”

## master\_verify\_checksum

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.2.4, “The Binary Log”

## max\_allowed\_packet

Section 12.17.1, “GROUP BY (Aggregate) Functions”  
Section C.5.2.3, “Lost connection to MySQL server”  
Section C.5.2.9, “MySQL server has gone away”  
Section 21.8.10.26, “mysql\_stmt\_send\_long\_data()”  
Section 21.8.6.72, “mysql\_use\_result()”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section C.5.2.11, “Communication Errors and Aborted Connections”  
Section 12.3.2, “Comparison Functions and Operators”  
Section C.5.5.6, “Deleting Rows from Related Tables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 21.8, “MySQL C API”  
Section C.5.2.10, “Packet Too Large”  
Section 16.4.1.19, “Replication and max\_allowed\_packet”  
Section 16.1.4.2, “Replication Master Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.4, “Server System Variables”  
Section 12.5, “String Functions”  
Section 11.4.3, “The BLOB and TEXT Types”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 4.2.3.3, “Using Option Files”  
Section 1.4, “What Is New in MySQL 5.6”

## max\_binlog\_cache\_size

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.2.4, “The Binary Log”

## max\_binlog\_size

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 5.2, “MySQL Server Logs”  
Section 5.2.6, “Server Log Maintenance”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4, “The Binary Log”  
Section 16.2.2.1, “The Slave Relay Log”

## max\_binlog\_stmt\_cache\_size

Section 16.1.4.4, “Binary Log Options and Variables”

## max\_connect\_errors

Section 13.7.6.3, “FLUSH Syntax”  
Section C.5.2.6, “Host 'host\_name' is blocked”



---

Section 5.1.2.1, “Changes to Server Defaults”  
Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”  
Section 5.1.4, “Server System Variables”  
Section 20.9.9.1, “The `host_cache` Table”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 1.4, “What Is New in MySQL 5.6”

## **max\_connections**

Section C.5.2.18, “`'File' Not Found` and Similar Errors”  
Section C.5.2.7, “Too many connections”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 22.4.1.4, “Debugging `mysqld` under `gdb`”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”  
Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”  
Section 20.12, “Performance Schema System Variables”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 5.1.3, “Server Command Options”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 8.11.6.2, “Thread Pool Operation”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **max\_delayed\_threads**

Section 20.12, “Performance Schema System Variables”  
Section 5.1.4, “Server System Variables”

## **max\_error\_count**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 13.6.7.4, “`RESIGNAL` Syntax”  
Section 13.7.5.18, “`SHOW ERRORS` Syntax”  
Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”  
Section 5.1.4, “Server System Variables”

## **max\_heap\_table\_size**

Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section E.10.3, “Limits on Table Size”  
Section 16.4.1.20, “Replication and `MEMORY` Tables”  
Section 16.4.1.32, “Replication and Variables”  
Section E.3, “Restrictions on Server-Side Cursors”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 14.4, “The `MEMORY` Storage Engine”

## **max\_insert\_delayed\_threads**

Section 5.1.4, “Server System Variables”

## **max\_join\_size**

Section 4.5.1.1, “`mysql` Options”

Section 8.8.2, “`EXPLAIN` Output Format”  
Section 13.7.4, “`SET` Syntax”  
Section 5.1.4, “Server System Variables”  
Section 5.1.5, “Using System Variables”  
Section 1.4, “What Is New in MySQL 5.6”

## **max\_length\_for\_sort\_data**

Section 8.13.13, “`ORDER BY` Optimization”  
Section 5.1.4, “Server System Variables”

## **max\_prepared\_stmt\_count**

Section 13.5.3, “`DEALLOCATE PREPARE` Syntax”  
Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 13.5, “SQL Syntax for Prepared Statements”

## **max\_relay\_log\_size**

Section 16.1.4.4, “Binary Log Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.4, “Server System Variables”  
Section 16.2.2.1, “The Slave Relay Log”

## **max\_seeks\_for\_key**

Section 14.2.7, “Limits on `InnoDB` Tables”  
Section 5.1.4, “Server System Variables”

## **max\_sort\_length**

Section 13.1.14, “`CREATE TABLE` Syntax”  
Section C.5.8, “Known Issues in MySQL”  
Section 5.1.4, “Server System Variables”  
Section 11.4.3, “The `BLOB` and `TEXT` Types”

## **max\_sp\_recursion\_depth**

Section 5.1.4, “Server System Variables”  
Section 18.2.1, “Stored Routine Syntax”

## **max\_tmp\_tables**

Section 5.1.4, “Server System Variables”

## **max\_user\_connections**

Section 13.7.1.4, “`GRANT` Syntax”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 5.1.4, “Server System Variables”  
Section 6.3.4, “Setting Account Resource Limits”

## **max\_write\_lock\_count**

Section 5.1.4, “Server System Variables”  
Section 8.10.2, “Table Locking Issues”

## **metadata\_locks\_cache\_size**

Section 5.1.4, “Server System Variables”

---

## **metadata\_locks\_hash\_instances**

Section 5.1.4, “Server System Variables”

## **min\_examined\_row\_limit**

Section 5.1.4, “Server System Variables”

Section 5.2.5, “The Slow Query Log”

## **myisam\_data\_pointer\_size**

Section 13.1.14, “CREATE TABLE Syntax”

Section E.10.3, “Limits on Table Size”

Section 5.1.4, “Server System Variables”

## **myisam\_max\_sort\_file\_size**

Section 14.3.1, “MyISAM Startup Options”

Section 17.6, “Restrictions and Limitations on Partitioning”

Section 5.1.4, “Server System Variables”

Section 8.6.3, “Speed of REPAIR TABLE Statements”

## **myisam\_mmap\_size**

Section 5.1.4, “Server System Variables”

## **myisam\_recover\_options**

Section 5.1.4, “Server System Variables”

## **myisam\_repair\_threads**

Section 5.1.4, “Server System Variables”

## **myisam\_sort\_buffer\_size**

Section 4.6.3.1, “myisamchk General Options”

Section 13.1.6, “ALTER TABLE Syntax”

Section 14.3.1, “MyISAM Startup Options”

Section 5.1.4, “Server System Variables”

Section 8.6.3, “Speed of REPAIR TABLE Statements”

## **myisam\_stats\_method**

Section 8.3.7, “InnoDB and MyISAM Index Statistics Collection”

Section 5.1.4, “Server System Variables”

## **myisam\_use\_mmap**

Section 8.11.4.1, “How MySQL Uses Memory”

Section 5.1.4, “Server System Variables”

## **N**

[index top [4055]]

## **named\_pipe**

Section 5.1.4, “Server System Variables”

## **net\_buffer\_length**

Section 4.5.1.1, “mysql Options”

Section 4.5.4, “mysqldump — A Database Backup Program”

Section 8.11.4.1, “How MySQL Uses Memory”

Section 21.8, “MySQL C API”

Section 5.1.4, “Server System Variables”

## **net\_read\_timeout**

Section C.5.2.3, “Lost connection to MySQL server”

Section 5.1.4, “Server System Variables”

## **net\_retry\_count**

Section 21.8.6.49, “mysql\_options()”

Section 5.1.4, “Server System Variables”

## **net\_write\_timeout**

Section 5.1.4, “Server System Variables”

## **new**

Section 5.1.4, “Server System Variables”

## **O**

[index top [4055]]

## **old**

Section 13.2.9.3, “Index Hint Syntax”

Section 5.1.4, “Server System Variables”

## **old\_alter\_table**

Section 13.1.6, “ALTER TABLE Syntax”

Section 5.5.1, “Overview of Online DDL”

Section 5.5.2, “Performance and Concurrency Considerations for Online DDL”

Section 5.1.4, “Server System Variables”

Section 5.5.3, “SQL Syntax for Online DDL”

## **old\_passwords**

Section 13.7.1.1, “ALTER USER Syntax”

Section C.5.2.4, “Client does not support authentication protocol”

Section 13.7.1.7, “SET PASSWORD Syntax”

Section 6.3.5, “Assigning Account Passwords”

Section 12.13, “Encryption and Compression Functions”

Section 6.1.2.5, “Implications of Password Hashing Changes in MySQL 4.1 for Application Programs”

Section 21.9, “MySQL PHP API”

Section 6.3.6, “Password Expiration and Sandbox Mode”

Section 6.1.2.4, “Password Hashing in MySQL”

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

## **open\_files\_limit**

Section C.5.2.18, “'File' Not Found and Similar Errors”

---

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”  
Section 17.6, “Restrictions and Limitations on Partitioning”  
Section 5.1.4, “Server System Variables”

### **optimizer\_join\_cache\_level**

Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”  
Section 8.13.12.1, “Join Buffer Management for Block Nested-Loop and Batched Key Access Algorithms”  
Section 5.1.4, “Server System Variables”

### **optimizer\_prune\_level**

Section 8.8.5.1, “Controlling Query Plan Evaluation”  
Section 5.1.4, “Server System Variables”

### **optimizer\_search\_depth**

Section 8.8.5.1, “Controlling Query Plan Evaluation”  
Section 5.1.4, “Server System Variables”

### **optimizer\_switch**

Section 8.13.12.3, “Batched Key Access Joins”  
Section 8.13.12.2, “Block Nested-Loop Algorithm for Outer Joins and Semi-Joins”  
Section 8.8.5.2, “Controlling Switchable Optimizations”  
Section 8.13.3, “Engine Condition Pushdown Optimization”  
Section 8.13.4, “Index Condition Pushdown Optimization”  
Section 8.13.12.1, “Join Buffer Management for Block Nested-Loop and Batched Key Access Algorithms”  
Section 8.13.11, “Multi-Range Read Optimization”  
Section 8.13.16.1, “Optimizing Subqueries with Semi-Join Transformations”  
Section 8.13.16.2, “Optimizing Subqueries with Subquery Materialization”  
Section 5.1.4, “Server System Variables”  
Section 8.13.5, “Use of Index Extensions”  
Section 1.4, “What Is New in MySQL 5.6”

### **optimizer\_trace**

Section 5.1.4, “Server System Variables”  
Section 19.12, “The `INFORMATION_SCHEMA.OPTIMIZER_TRACE` Table”

### **optimizer\_trace\_features**

Section 5.1.4, “Server System Variables”

### **optimizer\_trace\_limit**

Section 5.1.4, “Server System Variables”

### **optimizer\_trace\_max\_mem\_size**

Section 5.1.4, “Server System Variables”

### **optimizer\_trace\_offset**

Section 5.1.4, “Server System Variables”

## **P**

[index top [4055]]

### **performance\_schema**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.1, “Performance Schema Quick Start”  
Section 20.2.2, “Performance Schema Startup Configuration”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_accounts\_size**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.6.1, “The `accounts` Table”

### **performance\_schema\_digests\_size**

Section 20.7, “Performance Schema Statement Digests”  
Section 20.12, “Performance Schema System Variables”  
Section 20.9.8.3, “Statement Summary Tables”

### **performance\_schema\_events\_stages\_history**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.4.3, “The `events_stages_history_long` Table”

### **performance\_schema\_events\_stages\_history**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.4.2, “The `events_stages_history` Table”

### **performance\_schema\_events\_statement**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.5.3, “The `events_statements_history_long` Table”

### **performance\_schema\_events\_statement**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.5.2, “The `events_statements_history` Table”

### **performance\_schema\_events\_waits\_history**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”



---

Section 20.9, “Performance Schema Table Descriptions”  
Section 20.9.3.3, “The `events_waits_history_long` Table”

### **performance\_schema\_events\_waits\_history\_size**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”  
Section 20.9, “Performance Schema Table Descriptions”  
Section 20.9.3.2, “The `events_waits_history` Table”

### **performance\_schema\_hosts\_size**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.6.2, “The `hosts` Table”

### **performance\_schema\_max\_cond\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_cond\_instances**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_file\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_file\_handles**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_file\_instances**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_mutex\_classes**

Section 20.5, “Performance Schema Status Monitoring”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_mutex\_instances**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_rwlock\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_rwlock\_instances**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_socket\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_socket\_instances**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_stage\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_statement\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_table\_handles**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_table\_instances**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_thread\_classes**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_max\_thread\_instances**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_session\_connect\_attrs**

Section 20.12, “Performance Schema System Variables”

### **performance\_schema\_setup\_actors\_size**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.1.1, “The `setup_actors` Table”

### **performance\_schema\_setup\_objects\_size**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.1.4, “The `setup_objects` Table”

---

## **performance\_schema\_users\_size**

Section 20.12, “Performance Schema System Variables”  
Section 20.9.6.3, “The `users` Table”

## **pid\_file**

Section 5.1.4, “Server System Variables”

## **plugin\_dir**

Section 13.7.3.1, “`CREATE FUNCTION` Syntax for User-Defined Functions”  
Section 13.7.3.3, “`INSTALL PLUGIN` Syntax”  
Section 13.7.5.26, “`SHOW PLUGINS` Syntax”  
Section 6.1.2.2, “Administrator Guidelines for Password Security”  
Section 22.2.4.3, “Compiling and Installing Plugin Libraries”  
Section 22.3.2.5, “Compiling and Installing User-Defined Functions”  
Section 5.1.8.1, “Installing and Uninstalling Plugins”  
Section 6.3.11.1, “Installing the Audit Log Plugin”  
Section 6.3.7.3.1, “Installing the PAM Authentication Plugin”  
Section 6.3.7.4.1, “Installing the Windows Authentication Plugin”  
Section 6.1.3, “Making MySQL Secure Against Attackers”  
Section 6.1.2.6.1, “Password Validation Plugin Installation”  
Section 6.3.7, “Pluggable Authentication”  
Section 14.12.1, “Pluggable Storage Engine Architecture”  
Section 22.2.2, “Plugin API Components”  
Section E.9, “Restrictions on Pluggable Authentication”  
Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 19.15, “The `INFORMATION_SCHEMA PLUGINS` Table”  
Section 6.3.7.3, “The PAM Authentication Plugin”  
Section 6.3.7.6, “The Socket Peer-Credential Authentication Plugin”  
Section 6.3.7.7, “The Test Authentication Plugin”  
Section 6.3.7.4, “The Windows Native Authentication Plugin”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 2.10.1, “Unix Postinstallation Procedures”  
Section 22.3.2.6, “User-Defined Function Security Precautions”  
Section 22.2.4.9.3, “Using the Authentication Plugins”  
Section 22.2.4.6, “Writing `INFORMATION_SCHEMA` Plugins”  
Section 22.2.4.8, “Writing Audit Plugins”  
Section 22.2.4.5, “Writing Daemon Plugins”  
Section 22.2.4.4, “Writing Full-Text Parser Plugins”

Section 22.2.4.10, “Writing Password-Validation Plugins”  
Section 22.2.4.7, “Writing Semisynchronous Replication Plugins”

## **port**

Section C.5.2.2, “Can't connect to [local] MySQL server”  
Section 5.1.4, “Server System Variables”

## **preload\_buffer\_size**

Section 5.1.4, “Server System Variables”

## **profiling**

Section 13.7.5.31, “`SHOW PROFILE` Syntax”  
Section 5.1.4, “Server System Variables”  
Section 19.17, “The `INFORMATION_SCHEMA PROFILING` Table”

## **profiling\_history\_size**

Section 13.7.5.31, “`SHOW PROFILE` Syntax”  
Section 5.1.4, “Server System Variables”

## **protocol\_version**

Section 5.1.4, “Server System Variables”

## **proxy\_user**

Section 6.3.8, “Proxy Users”  
Section 5.1.4, “Server System Variables”

## **pseudo\_slave\_mode**

Section 5.1.4, “Server System Variables”

## **pseudo\_thread\_id**

Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”

## **Q**

[index top [4055]]

## **query\_alloc\_block\_size**

Section 5.1.4, “Server System Variables”

## **query\_cache\_limit**

Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.4, “Server System Variables”

## **query\_cache\_min\_res\_unit**

Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.4, “Server System Variables”

## **query\_cache\_size**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 8.9.3.3, “Query Cache Configuration”

---

Section 5.1.4, “Server System Variables”  
Section 8.9.3, “The MySQL Query Cache”  
Section 5.1.5, “Using System Variables”

### **query\_cache\_type**

Section 13.2.9, “`SELECT` Syntax”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 8.9.3.2, “Query Cache `SELECT` Options”  
Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

### **query\_cache\_wlock\_invalidate**

Section 5.1.4, “Server System Variables”

### **query\_prealloc\_size**

Section 5.1.4, “Server System Variables”

## **R**

[index top [4055]]

### **rand\_seed**

Section 5.1.4, “Server System Variables”

### **range\_alloc\_block\_size**

Section 5.1.4, “Server System Variables”

### **read\_buffer\_size**

Section 8.11.4.1, “How MySQL Uses Memory”  
Section 5.1.4, “Server System Variables”  
Section 8.6.3, “Speed of `REPAIR TABLE` Statements”

### **read\_only**

Section 13.7.1.7, “`SET PASSWORD` Syntax”  
Section 13.3.1, “`START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax”  
Section 6.3.5, “Assigning Account Passwords”  
Section 16.3.1.3, “Backing Up a Master or Slave by Making It Read Only”  
Section 6.2.1, “Privileges Provided by MySQL”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 16.1.3.2, “Setting Up Replication Using GTIDs”

### **read\_rnd\_buffer\_size**

Section 8.13.13, “`ORDER BY` Optimization”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 8.13.11, “Multi-Range Read Optimization”  
Section 5.1.4, “Server System Variables”  
Section 8.11.2, “Tuning Server Parameters”

### **relay\_log**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **relay\_log\_basename**

Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 1.4, “What Is New in MySQL 5.6”

### **relay\_log\_index**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **relay\_log\_info\_file**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **relay\_log\_info\_repository**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **relay\_log\_purge**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”  
Section 5.1.4, “Server System Variables”

### **relay\_log\_recovery**

Section 13.4.2.5, “`START SLAVE` Syntax”  
Section 16.1.4.3, “Replication Slave Options and Variables”

### **relay\_log\_space\_limit**

Section 8.12.5.6, “Replication Slave I/O Thread States”  
Section 5.1.4, “Server System Variables”

### **report\_host**

Section 5.1.4, “Server System Variables”

### **report\_password**

Section 5.1.4, “Server System Variables”

### **report\_port**

Section 5.1.4, “Server System Variables”

### **report\_user**

Section 5.1.4, “Server System Variables”

### **rpl\_semi\_sync\_master\_enabled**

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”  
Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 5.1.4, “Server System Variables”

### **rpl\_semi\_sync\_master\_timeout**

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”  
Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

---

Section 5.1.4, “Server System Variables”

## **rpl\_semi\_sync\_master\_trace\_level**

Section 5.1.4, “Server System Variables”

## **rpl\_semi\_sync\_master\_wait\_no\_slave**

Section 5.1.4, “Server System Variables”

## **rpl\_semi\_sync\_slave\_enabled**

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”

Section 16.3.8.2, “Semisynchronous Replication Installation and Configuration”

Section 5.1.4, “Server System Variables”

## **rpl\_semi\_sync\_slave\_trace\_level**

Section 5.1.4, “Server System Variables”

## **S**

[[index top \[4055\]](#)]

### **secure\_auth**

Section 5.1.2.1, “Changes to Server Defaults”

Section 6.1.2.4, “Password Hashing in MySQL”

Section 5.1.4, “Server System Variables”

### **secure\_file\_priv**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”

Section 13.2.9.1, “`SELECT ... INTO` Syntax”

Section 6.1.3, “Making MySQL Secure Against Attackers”

Section 6.2.1, “Privileges Provided by MySQL”

Section 5.1.4, “Server System Variables”

Section 12.5, “String Functions”

### **server\_id**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 4.6.8, “`mysqlbinlog` — Utility for Processing Binary Log Files”

Section 12.16, “Miscellaneous Functions”

Section 5.1.4, “Server System Variables”

Section 6.3.11.2, “The Audit Log File”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **server\_uuid**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 13.4.2.5, “`START SLAVE` Syntax”

Section 16.1.3.1, “GTID Concepts”

Section 16.1.4, “Replication and Binary Logging Options and Variables”

### **sha**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

### **shared\_memory**

Section 5.1.4, “Server System Variables”

### **shared\_memory\_base\_name**

Section 5.1.4, “Server System Variables”

### **skip\_external\_locking**

Section 8.10.5, “External Locking”

Section 5.1.4, “Server System Variables”

### **skip\_name\_resolve**

Section 5.1.4, “Server System Variables”

### **skip\_networking**

Section 5.1.4, “Server System Variables”

### **skip\_show\_database**

Section 5.1.3, “Server Command Options”

Section 5.1.4, “Server System Variables”

### **slave\_allow\_batching**

Section 16.1.4.2, “Replication Master Options and Variables”

### **slave\_checkpoint\_group**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_checkpoint\_period**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_compressed\_protocol**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_exec\_mode**

Section 16.4.1.20, “Replication and `MEMORY` Tables”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 16.1.2.2, “Usage of Row-Based Logging and Replication”

### **slave\_load\_tmpdir**

Section 13.2.6, “`LOAD DATA INFILE` Syntax”

Section 16.1.4.3, “Replication Slave Options and Variables”

Section 5.1.4, “Server System Variables”

### **slave\_max\_allowed\_packet**

Section 16.1.4.2, “Replication Master Options and Variables”

### **slave\_net\_timeout**

Section 13.4.2.1, “`CHANGE MASTER TO` Syntax”

---

Section 16.1.5.1, “Checking Replication Status”  
Section 16.4.1.18, “Replication and Master or Slave Shutdowns”  
Section 8.12.5.6, “Replication Slave I/O Thread States”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.4, “Server System Variables”

### **slave\_parallel\_workers**

Section 13.7.5.35, “SHOW SLAVE STATUS Syntax”  
Section 13.4.2.5, “START SLAVE Syntax”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 1.4, “What Is New in MySQL 5.6”

### **slave\_pending\_jobs\_size\_max**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_rows\_search\_algorithms**

Section 16.1.4.2, “Replication Master Options and Variables”

### **slave\_skip\_errors**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_sql\_verify\_checksum**

Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.2.4, “The Binary Log”

### **slave\_transaction\_retries**

Section 16.4.1.27, “Replication Retries and Timeouts”  
Section 16.1.4.3, “Replication Slave Options and Variables”

### **slave\_type\_conversions**

Section 16.1.4.3, “Replication Slave Options and Variables”

### **slow\_launch\_time**

Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

### **slow\_query\_log**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 5.2.5, “The Slow Query Log”

### **slow\_query\_log\_file**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 5.2.5, “The Slow Query Log”

### **socket**

Section 5.1.4, “Server System Variables”

### **sort\_buffer\_size**

Section 8.13.13, “ORDER BY Optimization”  
Section 7.6.3, “How to Repair MYISAM Tables”  
Section 8.2.1.3, “Optimizing LIMIT Queries”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 1.4, “What Is New in MySQL 5.6”

### **sql\_auto\_is\_null**

Section 12.3.2, “Comparison Functions and Operators”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4, “The Binary Log”

### **sql\_big\_selects**

Section 5.1.4, “Server System Variables”

### **sql\_buffer\_result**

Section 5.1.4, “Server System Variables”

### **sql\_log\_bin**

Section 13.4.1.3, “SET sql\_log\_bin Syntax”  
Section 5.1.4, “Server System Variables”  
Section 16.4.3, “Upgrading a Replication Setup”

### **sql\_log\_off**

Section 5.2.1, “Selecting General Query and Slow Query Log Output Destinations”  
Section 5.1.4, “Server System Variables”  
Section 5.2.3, “The General Query Log”

### **SQL\_MODE**

Section 5.5.5, “Examples of Online DDL”  
Section 5.5.1, “Overview of Online DDL”

### **sql\_mode**

Section 13.1.9, “CREATE EVENT Syntax”  
Section 13.1.12, “CREATE PROCEDURE and CREATE FUNCTION Syntax”  
Section 13.1.15, “CREATE TRIGGER Syntax”  
Section 13.2.6, “LOAD DATA INFILE Syntax”  
Section 13.7.5.14, “SHOW CREATE VIEW Syntax”  
Section 13.7.5.39, “SHOW TRIGGERS Syntax”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 4.4.3, “mysql\_install\_db — Initialize MySQL Data Directory”  
Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”  
Section 12.19.3, “Expression Handling”  
Section 1.7, “How to Report Bugs or Problems”  
Section 5.2.4.3, “Mixed Binary Logging Format”



---

Section C.5.5.2, “Problems Using `DATE` Columns”  
Section 16.4.1.32, “Replication and Variables”  
Section 1.8.3, “Running MySQL in ANSI Mode”  
Section 1.8.2, “Selecting SQL Modes”  
Section 5.1.7, “Server SQL Modes”  
Section 5.1.4, “Server System Variables”  
Section 19.19, “The `INFORMATION_SCHEMA` `ROUTINES` Table”  
Section 19.27, “The `INFORMATION_SCHEMA` `TRIGGERS` Table”  
Section 19.29, “The `INFORMATION_SCHEMA` `VIEWS` Table”  
Section 5.2.4, “The Binary Log”  
Section 5.1.2.2, “Using a Sample Default Server Configuration File”  
Section 5.1.5, “Using System Variables”

## **sql\_notes**

Section 13.7.5.41, “`SHOW WARNINGS` Syntax”  
Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”  
Section 5.1.4, “Server System Variables”

## **sql\_quote\_show\_create**

Section 13.7.5.8, “`SHOW CREATE DATABASE` Syntax”  
Section 13.7.5.12, “`SHOW CREATE TABLE` Syntax”  
Section 5.1.4, “Server System Variables”

## **sql\_safe\_updates**

Section 5.1.4, “Server System Variables”

## **sql\_select\_limit**

Section 5.1.4, “Server System Variables”

## **sql\_slave\_skip\_counter**

Section 13.4.2.4, “`SET GLOBAL sql_slave_skip_counter` Syntax”  
Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”  
Section 16.1.4.5, “Global Transaction ID Options and Variables”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 16.1.3.4, “Restrictions on Replication with GTIDs”

## **sql\_warnings**

Section 5.1.4, “Server System Variables”

## **ssl\_ca**

Section 5.1.4, “Server System Variables”

## **ssl\_capath**

Section 5.1.4, “Server System Variables”

## **ssl\_cert**

Section 5.1.4, “Server System Variables”

## **ssl\_cipher**

Section 5.1.4, “Server System Variables”

## **ssl\_crl**

Section 5.1.4, “Server System Variables”

## **ssl\_crlpath**

Section 5.1.4, “Server System Variables”

## **ssl\_key**

Section 5.1.4, “Server System Variables”

## **storage\_engine**

Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 16.3.2, “Using Replication with Different Master and Slave Storage Engines”

## **stored\_program\_cache**

Section 8.9.4, “Caching of Prepared Statements and Stored Programs”  
Section 5.1.4, “Server System Variables”

## **sync\_binlog**

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 16.1.4.4, “Binary Log Options and Variables”  
Section 14.2.8, “MySQL and the ACID Model”  
Section 8.5.7, “Optimizing InnoDB Disk I/O”  
Section 16.4.1.18, “Replication and Master or Slave Shutdowns”  
Section 5.2.4, “The Binary Log”

## **sync\_frm**

Section 5.1.4, “Server System Variables”

## **sync\_master\_info**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **sync\_relay\_log**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

## **sync\_relay\_log\_info**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 16.4.1.18, “Replication and Master or Slave Shutdowns”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”

---

## **system\_time\_zone**

Section 10.6, “MySQL Server Time Zone Support”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **T**

[index top [4055]]

## **table\_definition\_cache**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 5.1.4, “Server System Variables”

## **table\_open\_cache**

Section C.5.2.18, “*'File' Not Found* and Similar Errors”  
Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.2.1, “Changes to Server Defaults”  
Section 8.12.5.2, “General Thread States”  
Section 8.4.3.1, “How MySQL Opens and Closes Tables”  
Section 8.11.4.1, “How MySQL Uses Memory”  
Section 5.1.3, “Server Command Options”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”  
Section 8.11.2, “Tuning Server Parameters”  
Section 2.11.1.1, “Upgrading from MySQL 5.5 to 5.6”  
Section 1.4, “What Is New in MySQL 5.6”

## **table\_open\_cache\_instances**

Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

## **thread\_cache\_size**

Section 5.1.2.1, “Changes to Server Defaults”  
Section 22.4.1.4, “Debugging `mysqld` under `gdb`”  
Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”  
Section 8.2.5, “Other Optimization Tips”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

## **thread\_concurrency**

Section 5.1.4, “Server System Variables”

## **thread\_handling**

Section 5.1.4, “Server System Variables”  
Section 8.11.6.1, “Thread Pool Components and Installation”

## **thread\_pool\_algorithm**

Section 5.1.4, “Server System Variables”  
Section 19.31.2, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATE` Table”

Section 8.11.6.1, “Thread Pool Components and Installation”

## **thread\_pool\_high\_priority\_connect**

Section 5.1.4, “Server System Variables”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 8.11.6.2, “Thread Pool Operation”

## **thread\_pool\_max\_unused\_threads**

Section 5.1.4, “Server System Variables”  
Section 8.11.6.1, “Thread Pool Components and Installation”

## **thread\_pool\_prio\_kickup\_timer**

Section 5.1.4, “Server System Variables”  
Section 19.31.2, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATE` Table”  
Section 19.31.3, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATS` Table”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 8.11.6.2, “Thread Pool Operation”  
Section 8.11.6.3, “Thread Pool Tuning”

## **thread\_pool\_size**

Section 5.1.4, “Server System Variables”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 8.11.6.2, “Thread Pool Operation”  
Section 8.11.6.3, “Thread Pool Tuning”

## **thread\_pool\_stall\_limit**

Section 5.1.4, “Server System Variables”  
Section 19.31.2, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATE` Table”  
Section 19.31.3, “The `INFORMATION_SCHEMA` `TP_THREAD_GROUP_STATS` Table”  
Section 8.11.6.1, “Thread Pool Components and Installation”  
Section 8.11.6.2, “Thread Pool Operation”  
Section 8.11.6.3, “Thread Pool Tuning”

## **thread\_stack**

Section 8.11.4.1, “How MySQL Uses Memory”  
Section 5.1.4, “Server System Variables”  
Section 18.2.1, “Stored Routine Syntax”

## **time\_format**

Section 5.1.4, “Server System Variables”

## **time\_zone**

Section 13.1.9, “`CREATE EVENT` Syntax”  
Section 12.7, “Date and Time Functions”  
Section 18.4.4, “Event Metadata”  
Section 5.2.4.3, “Mixed Binary Logging Format”

---

Section 10.6, “MySQL Server Time Zone Support”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”  
Section 11.3.1, “The `DATE`, `DATETIME`, and `TIMESTAMP` Types”

## **timed\_mutexes**

Section 13.7.5.16, “`SHOW ENGINE` Syntax”  
Section 5.1.4, “Server System Variables”

## **timestamp**

Section 14.9.3, “`FEDERATED` Storage Engine Notes and Tips”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”

## **tmp\_table\_size**

Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”  
Section E.3, “Restrictions on Server-Side Cursors”  
Section 5.1.6, “Server Status Variables”  
Section 5.1.4, “Server System Variables”

## **tmpdir**

Section C.5.2.13, “`Can't create/write to file`”  
Section 13.2.6, “`LOAD DATA INFILE` Syntax”  
Section 8.13.13, “`ORDER BY` Optimization”  
Section 16.3.1.2, “Backing Up Raw Data from a Slave”  
Section 7.2, “Database Backup Methods”  
Section 5.5.6, “Implementation Details of Online DDL”  
Section 5.5.9, “Limitations of Online DDL”  
Section 16.1.4.3, “Replication Slave Options and Variables”  
Section 5.1.4, “Server System Variables”

## **transaction\_alloc\_block\_size**

Section 5.1.4, “Server System Variables”

## **transaction\_prealloc\_size**

Section 5.1.4, “Server System Variables”

## **tx\_isolation**

Section 13.3.6, “`SET TRANSACTION` Syntax”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **tx\_read\_only**

Section 13.3.6, “`SET TRANSACTION` Syntax”  
Section 5.1.3, “Server Command Options”  
Section 5.1.4, “Server System Variables”

## **U**

[index top [4055]]

## **unique\_checks**

Section 5.4.3, “Converting Tables from `MyISAM` to `InnoDB`”  
Section 5.2.4.3, “Mixed Binary Logging Format”  
Section 16.4.1.32, “Replication and Variables”  
Section 5.1.4, “Server System Variables”  
Section 5.2.4, “The Binary Log”

## **updatable\_views\_with\_limit**

Section 5.1.4, “Server System Variables”  
Section 18.5.3, “Updatable and Insertable Views”

## **V**

[index top [4055]]

## **validate\_password\_dictionary\_file**

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **validate\_password\_length**

Section 12.13, “Encryption and Compression Functions”  
Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **validate\_password\_mixed\_case\_count**

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **validate\_password\_number\_count**

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **validate\_password\_policy**

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”  
Section 6.1.2.6, “The Password Validation Plugin”

## **validate\_password\_special\_char\_count**

Section 6.1.2.6.2, “Password Validation Plugin Options and Variables”

## **validate\_user\_plugins**

Section 5.1.4, “Server System Variables”

## **version**

Section 14.2.6, “`InnoDB` Startup Options and System Variables”  
Section 12.14, “Information Functions”  
Section 5.1.4, “Server System Variables”  
Section 6.3.11.2, “The Audit Log File”

## **version\_comment**

Section 5.1.4, “Server System Variables”



---

## **version\_compile\_machine**

Section 5.1.4, “Server System Variables”

## **version\_compile\_os**

Section 5.1.4, “Server System Variables”

## **W**

[index top [4055]]

## **wait\_timeout**

Section C.5.2.9, “MySQL server has gone away”

Section 21.8.6.53, “mysql\_real\_connect()”

Section C.5.2.11, “Communication Errors and Aborted Connections”

Section 5.1.4, “Server System Variables”

## **warning\_count**

Section 13.7.5.18, “SHOW ERRORS Syntax”

Section 13.7.5.41, “SHOW WARNINGS Syntax”

Section 13.6.7.7.4, “Diagnostics Area-Related System Variables”

Section 13.6.7.5.2, “Effect of Signals on Handlers, Cursors, and Statements”

Section 5.1.4, “Server System Variables”

Section C.1, “Sources of Error Information”



---

# Status Variable Index

[A](#) | [B](#) | [C](#) | [D](#) | [F](#) | [H](#) | [I](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#)  
[| U](#)

## A

[\[index top \[4083\]\]](#)

### [Aborted\\_clients](#)

[Section C.5.2.11, “Communication Errors and Aborted Connections”](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Aborted\\_connects](#)

[Section C.5.2.11, “Communication Errors and Aborted Connections”](#)

[Section 5.1.6, “Server Status Variables”](#)

## B

[\[index top \[4083\]\]](#)

### [Binlog\\_cache\\_disk\\_use](#)

[Section 16.1.4.4, “Binary Log Options and Variables”](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.2.4, “The Binary Log”](#)

### [Binlog\\_cache\\_use](#)

[Section 16.1.4.4, “Binary Log Options and Variables”](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.2.4, “The Binary Log”](#)

### [Binlog\\_stmt\\_cache\\_disk\\_use](#)

[Section 16.1.4.4, “Binary Log Options and Variables”](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Binlog\\_stmt\\_cache\\_use](#)

[Section 16.1.4.4, “Binary Log Options and Variables”](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Bytes\\_received](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Bytes\\_sent](#)

[Section 5.1.6, “Server Status Variables”](#)

## C

[\[index top \[4083\]\]](#)

### [Com\\_flush](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Com\\_stmt\\_reprepare](#)

[Section 8.9.4, “Caching of Prepared Statements and Stored Programs”](#)

### [Compression](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_accept](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_internal](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_max\\_connections](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.1.4, “Server System Variables”](#)

### [Connection\\_errors\\_peer\\_addr](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_select](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_tcpwrap](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Connection\\_errors\\_xxx](#)

[Section 8.11.5.2, “DNS Lookup Optimization and the Host Cache”](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 1.4, “What Is New in MySQL 5.6”](#)

### [Connections](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.1.4, “Server System Variables”](#)

### [Created\\_tmp\\_disk\\_tables](#)

[Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.1.4, “Server System Variables”](#)

[Section 20.9.5.1, “The](#)

[events\\_statements\\_current Table”](#)

### [Created\\_tmp\\_files](#)

[Section 5.1.6, “Server Status Variables”](#)

### [Created\\_tmp\\_tables](#)

[Section 13.7.5.36, “SHOW STATUS Syntax”](#)

[Section 8.4.3.3, “How MySQL Uses Internal Temporary Tables”](#)

[Section 5.1.6, “Server Status Variables”](#)

[Section 5.1.4, “Server System Variables”](#)

[Section 20.9.5.1, “The](#)

[events\\_statements\\_current Table”](#)

---

## D

[\[index top \[4083\]\]](#)

### [Delayed\\_errors](#)

Section 5.1.6, “Server Status Variables”

### [Delayed\\_insert\\_threads](#)

Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 5.1.6, “Server Status Variables”

### [Delayed\\_writes](#)

Section 13.2.5.2, “`INSERT DELAYED` Syntax”  
Section 5.1.6, “Server Status Variables”

## F

[\[index top \[4083\]\]](#)

### [Flush\\_commands](#)

Section 5.1.6, “Server Status Variables”

## H

[\[index top \[4083\]\]](#)

### [Handler\\_commit](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_delete](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_external\\_lock](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_mrr\\_init](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_prepare](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_first](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_key](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_last](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_next](#)

Section 5.1.6, “Server Status Variables”  
Section 8.13.5, “Use of Index Extensions”

### [Handler\\_read\\_prev](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_rnd](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_read\\_rnd\\_next](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_rollback](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_savepoint](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_savepoint\\_rollback](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_update](#)

Section 5.1.6, “Server Status Variables”

### [Handler\\_write](#)

Section 5.1.6, “Server Status Variables”

## I

[\[index top \[4083\]\]](#)

### [Innodb\\_available\\_undo\\_logs](#)

Section 14.2.6, “InnoDB Startup Options and System Variables”  
Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_bytes\\_data](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_bytes\\_dirty](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_dump\\_status](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_load\\_status](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_pages\\_data](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_pages\\_dirty](#)

Section 5.1.6, “Server Status Variables”

### [Innodb\\_buffer\\_pool\\_pages\\_flushed](#)

Section 5.1.6, “Server Status Variables”

---

<b>Innodb_buffer_pool_pages_free</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_data_written</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_pages_latched</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_dblwr_pages_written</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_pages_misc</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_dblwr_writes</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_pages_total</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_have_atomic_builtins</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_read_ahead</b> Section 14.2.5.11, "More Read-Ahead Statistics" Section 5.1.6, "Server Status Variables"	<b>Innodb_log_waits</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_read_ahead_evicted</b> Section 14.2.5.11, "More Read-Ahead Statistics" Section 5.1.6, "Server Status Variables"	<b>Innodb_log_write_requests</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_read_requests</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_log_writes</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_reads</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_num_open_files</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_wait_free</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_os_log_fsyncls</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_buffer_pool_write_requests</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_os_log_pending_fsyncls</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_fsyncls</b> Section 14.2.6, "InnoDB Startup Options and System Variables" Section 5.1.6, "Server Status Variables"	<b>Innodb_os_log_pending_writes</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_pending_fsyncls</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_os_log_written</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_pending_reads</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_page_size</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_pending_writes</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_pages_created</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_read</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_pages_read</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_reads</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_pages_written</b> Section 5.1.6, "Server Status Variables"
<b>Innodb_data_writes</b> Section 5.1.6, "Server Status Variables"	<b>Innodb_row_lock_current_waits</b> Section 5.1.6, "Server Status Variables"
	<b>Innodb_row_lock_time</b> Section 5.1.6, "Server Status Variables"

---

---

## **Innodb\_row\_lock\_time\_avg**

Section 5.1.6, “Server Status Variables”

## **Innodb\_row\_lock\_time\_max**

Section 5.1.6, “Server Status Variables”

## **Innodb\_row\_lock\_waits**

Section 5.1.6, “Server Status Variables”

## **Innodb\_rows\_deleted**

Section 5.1.6, “Server Status Variables”

## **Innodb\_rows\_inserted**

Section 5.1.6, “Server Status Variables”

## **Innodb\_rows\_read**

Section 5.1.6, “Server Status Variables”

## **Innodb\_rows\_updated**

Section 5.1.6, “Server Status Variables”

## **Innodb\_truncated\_status\_writes**

Section 5.1.6, “Server Status Variables”

## **K**

[index top [4083]]

## **Key\_blocks\_not\_flushed**

Section 5.1.6, “Server Status Variables”

## **Key\_blocks\_unused**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Key\_blocks\_used**

Section 5.1.6, “Server Status Variables”

## **Key\_read\_requests**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Key\_reads**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Key\_write\_requests**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Key\_writes**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **L**

[index top [4083]]

## **Last\_query\_cost**

Section 5.1.6, “Server Status Variables”

## **Last\_query\_partial\_plans**

Section 5.1.6, “Server Status Variables”

## **M**

[index top [4083]]

## **Max\_used\_connections**

Section 13.7.6.3, “FLUSH Syntax”

Section 5.1.6, “Server Status Variables”

## **N**

[index top [4083]]

## **Not\_flushed\_delayed\_rows**

Section 13.2.5.2, “INSERT DELAYED Syntax”

Section 5.1.6, “Server Status Variables”

## **O**

[index top [4083]]

## **Open\_files**

Section 5.1.6, “Server Status Variables”

## **Open\_streams**

Section 5.1.6, “Server Status Variables”

## **Open\_table\_definitions**

Section 5.1.6, “Server Status Variables”

## **Open\_tables**

Section 5.1.6, “Server Status Variables”

## **Opened\_files**

Section 5.1.6, “Server Status Variables”

## **Opened\_table\_definitions**

Section 5.1.6, “Server Status Variables”

## **Opened\_tables**

Section 8.4.3.1, “How MySQL Opens and Closes Tables”

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

---

## P

[index top [4083]]

### [Performance\\_schema\\_mutex\\_classes\\_lost](#)

Section 20.5, “Performance Schema Status Monitoring”

### [Performance\\_schema\\_mutex\\_instances\\_lost](#)

Section 20.5, “Performance Schema Status Monitoring”

### [Performance\\_schema\\_session\\_connect\\_attrs\\_lost](#)

Section 20.12, “Performance Schema System Variables”

### [Prepared\\_stmt\\_count](#)

Section 5.1.6, “Server Status Variables”

## Q

[index top [4083]]

### [Qcache\\_free\\_blocks](#)

Section 8.9.3.3, “Query Cache Configuration”  
Section 8.9.3.4, “Query Cache Status and Maintenance”  
Section 5.1.6, “Server Status Variables”

### [Qcache\\_free\\_memory](#)

Section 5.1.6, “Server Status Variables”

### [Qcache\\_hits](#)

Section 8.9.3.1, “How the Query Cache Operates”  
Section 5.1.6, “Server Status Variables”

### [Qcache\\_inserts](#)

Section 5.1.6, “Server Status Variables”

### [Qcache\\_lowmem\\_prunes](#)

Section 8.9.3.3, “Query Cache Configuration”  
Section 8.9.3.4, “Query Cache Status and Maintenance”  
Section 5.1.6, “Server Status Variables”

### [Qcache\\_not\\_cached](#)

Section 5.1.6, “Server Status Variables”

### [Qcache\\_queries\\_in\\_cache](#)

Section 8.9.3.3, “Query Cache Configuration”  
Section 5.1.6, “Server Status Variables”

### [Qcache\\_total\\_blocks](#)

Section 8.9.3.3, “Query Cache Configuration”  
Section 8.9.3.4, “Query Cache Status and Maintenance”

Section 5.1.6, “Server Status Variables”

## Queries

Section 5.1.6, “Server Status Variables”

## Questions

Section 4.5.2, “[mysqladmin](#) — Client for Administering a MySQL Server”  
Section 5.1.6, “Server Status Variables”

## R

[index top [4083]]

### [Rpl\\_semi\\_sync\\_master\\_clients](#)

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_net\\_avg\\_wait](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_net\\_wait\\_time](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_net\\_waits](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_no\\_times](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_no\\_tx](#)

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_status](#)

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”  
Section 16.3.8.3, “Semisynchronous Replication Monitoring”  
Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_timefunc\\_failures](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_tx\\_avg\\_wait](#)

Section 5.1.6, “Server Status Variables”

### [Rpl\\_semi\\_sync\\_master\\_tx\\_wait\\_time](#)

Section 5.1.6, “Server Status Variables”

---

## **Rpl\_semi\_sync\_master\_tx\_waits**

Section 5.1.6, “Server Status Variables”

## **Rpl\_semi\_sync\_master\_wait\_pos\_backlog**

Section 5.1.6, “Server Status Variables”

## **Rpl\_semi\_sync\_master\_wait\_sessions**

Section 5.1.6, “Server Status Variables”

## **Rpl\_semi\_sync\_master\_yes\_tx**

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”

Section 16.3.8.3, “Semisynchronous Replication Monitoring”

Section 5.1.6, “Server Status Variables”

## **Rpl\_semi\_sync\_slave\_status**

Section 16.3.8.1, “Semisynchronous Replication Administrative Interface”

Section 16.3.8.3, “Semisynchronous Replication Monitoring”

Section 5.1.6, “Server Status Variables”

## **Rsa\_public\_key**

Section 5.1.6, “Server Status Variables”

Section 6.3.7.2, “The SHA-256 Authentication Plugin”

## **S**

[index top [4083]]

## **Select\_full\_join**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Select\_full\_range\_join**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Select\_range**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Select\_range\_check**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Select\_scan**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Slave\_heartbeat\_period**

Section 5.1.6, “Server Status Variables”

## **Slave\_heartbeat**

Section 5.1.6, “Server Status Variables”

## **Slave\_open\_temp\_tables**

Section 16.4.1.21, “Replication and Temporary Tables”

Section 5.1.6, “Server Status Variables”

## **Slave\_received\_heartbeats**

Section 5.1.6, “Server Status Variables”

## **Slave\_retried\_transactions**

Section 5.1.6, “Server Status Variables”

## **Slave\_running**

Section 13.7.5.35, “`SHOW SLAVE STATUS` Syntax”

Section 16.2.1, “Replication Implementation Details”

Section 5.1.6, “Server Status Variables”

## **Slow\_launch\_threads**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Slow\_queries**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Sort\_merge\_passes**

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Sort\_range**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Sort\_rows**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Sort\_scan**

Section 5.1.6, “Server Status Variables”

Section 20.9.5.1, “The `events_statements_current` Table”

## **Ssl\_accept\_renegotiates**

Section 5.1.6, “Server Status Variables”

## **Ssl\_accepts**

Section 5.1.6, “Server Status Variables”



---

## **Ssl\_callback\_cache\_hits**

Section 5.1.6, “Server Status Variables”

## **Ssl\_cipher**

Section 5.1.6, “Server Status Variables”  
Section 6.3.9.3, “Using SSL Connections”

## **Ssl\_cipher\_list**

Section 5.1.6, “Server Status Variables”

## **Ssl\_client\_connects**

Section 5.1.6, “Server Status Variables”

## **Ssl\_connect\_renegotiates**

Section 5.1.6, “Server Status Variables”

## **Ssl\_ctx\_verify\_depth**

Section 5.1.6, “Server Status Variables”

## **Ssl\_ctx\_verify\_mode**

Section 5.1.6, “Server Status Variables”

## **Ssl\_default\_timeout**

Section 5.1.6, “Server Status Variables”

## **Ssl\_finished\_accepts**

Section 5.1.6, “Server Status Variables”

## **Ssl\_finished\_connects**

Section 5.1.6, “Server Status Variables”

## **Ssl\_server\_not\_after**

Section 5.1.6, “Server Status Variables”

## **Ssl\_server\_not\_before**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_hits**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_misses**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_mode**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_overflows**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_size**

Section 5.1.6, “Server Status Variables”

## **Ssl\_session\_cache\_timeouts**

Section 5.1.6, “Server Status Variables”

## **Ssl\_sessions\_reused**

Section 5.1.6, “Server Status Variables”

## **Ssl\_used\_session\_cache\_entries**

Section 5.1.6, “Server Status Variables”

## **Ssl\_verify\_depth**

Section 5.1.6, “Server Status Variables”

## **Ssl\_verify\_mode**

Section 5.1.6, “Server Status Variables”

## **Ssl\_version**

Section 5.1.6, “Server Status Variables”

## **T**

[index top [4083]]

## **Table\_locks\_immediate**

Section 8.10.1, “Internal Locking Methods”  
Section 5.1.6, “Server Status Variables”

## **Table\_locks\_waited**

Section 8.10.1, “Internal Locking Methods”  
Section 5.1.6, “Server Status Variables”

## **Table\_open\_cache\_hits**

Section 5.1.6, “Server Status Variables”

## **Table\_open\_cache\_misses**

Section 5.1.6, “Server Status Variables”

## **Table\_open\_cache\_overflows**

Section 5.1.6, “Server Status Variables”

## **Tc\_log\_max\_pages\_used**

Section 5.1.6, “Server Status Variables”

## **Tc\_log\_page\_size**

Section 5.1.6, “Server Status Variables”

## **Tc\_log\_page\_waits**

Section 5.1.6, “Server Status Variables”

## **Threads\_cached**

Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”  
Section 5.1.6, “Server Status Variables”

## **Threads\_connected**

Section 5.1.6, “Server Status Variables”

---

## **Threads\_created**

Section 8.11.5.1, “How MySQL Uses Threads for Client Connections”

Section 5.1.6, “Server Status Variables”

Section 5.1.4, “Server System Variables”

## **Threads\_running**

Section 5.1.6, “Server Status Variables”

## **U**

[[index top \[4083\]](#)]

## **Uptime**

Section 4.5.2, “`mysqladmin` — Client for Administering a MySQL Server”

Section 5.1.6, “Server Status Variables”

## **Uptime\_since\_flush\_status**

Section 5.1.6, “Server Status Variables”

---

# Transaction Isolation Level

## Index

R | S

### R

[[index top \[4091\]](#)]

#### **READ COMMITTED**

[Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#)

[Section 14.2.6, “InnoDB Startup Options and System Variables”](#)

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 14.2.3.3, “Consistent Nonlocking Reads”](#)

[Section 14.2.3.10, “How to Cope with Deadlocks”](#)

[Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”](#)

[Section 8.5.2, “Optimizing InnoDB Transaction Management”](#)

[Section 5.2.4.2, “Setting The Binary Log Format”](#)

[Section 14.2.3.1, “The InnoDB Transaction Model and Locking”](#)

#### **READ UNCOMMITTED**

[Section 14.2.6, “InnoDB Startup Options and System Variables”](#)

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 14.2.9.2, “Architecture of InnoDB and memcached Integration”](#)

[Section 14.2.3.3, “Consistent Nonlocking Reads”](#)

[Section 14.2.9.5.6, “Performing DML and DDL Statements on the Underlying InnoDB Table”](#)

[Section 5.2.4.2, “Setting The Binary Log Format”](#)

[Section 14.2.3.1, “The InnoDB Transaction Model and Locking”](#)

#### **READ-COMMITTED**

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 5.1.3, “Server Command Options”](#)

#### **READ-UNCOMMITTED**

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 5.1.3, “Server Command Options”](#)

#### **REPEATABLE READ**

[Section 14.2.3.5, “InnoDB Record, Gap, and Next-Key Locks”](#)

[Section 14.2.6, “InnoDB Startup Options and System Variables”](#)

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#)

[Section 14.2.3.3, “Consistent Nonlocking Reads”](#)

[Section 14.2.9.5.4, “Controlling Transactional Behavior of the InnoDB memcached Plugin”](#)

[Section 5.2.4.3, “Mixed Binary Logging Format”](#)

[Section 8.5.2, “Optimizing InnoDB Transaction Management”](#)

[Section 14.2.3.1, “The InnoDB Transaction Model and Locking”](#)

[Section 13.3.7, “XA Transactions”](#)

#### **REPEATABLE-READ**

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 5.1.3, “Server Command Options”](#)

[Section 5.1.4, “Server System Variables”](#)

### S

[[index top \[4091\]](#)]

#### **SERIALIZABLE**

[Section 14.2.6, “InnoDB Startup Options and System Variables”](#)

[Section 13.3.6, “SET TRANSACTION Syntax”](#)

[Section 13.3.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#)

[Section 14.2.3.3, “Consistent Nonlocking Reads”](#)

[Section 8.9.3.1, “How the Query Cache Operates”](#)

[Section 14.2.3.7, “Locks Set by Different SQL Statements in InnoDB”](#)

[Section 5.2.4.3, “Mixed Binary Logging Format”](#)

[Section 5.1.3, “Server Command Options”](#)

[Section 14.2.3.1, “The InnoDB Transaction Model and Locking”](#)

[Section 13.3.7, “XA Transactions”](#)

